Nearest Neighbor Search Under Uncertainty

Blake Mason[1] **Ardhendu Tripathy**[2] **Robert Nowak**[1]

[1]Department of Electrical and Computer Engineering., University of Wisconsin, Madison, Wisconsin, USA
[2]Computer Science Department., Missouri University of Science and Technology, Rolla, Missouri, USA

Abstract

Nearest Neighbor Search (NNS) is a central task in knowledge representation, learning, and reasoning. There is vast literature on efficient algorithms for constructing data structures and performing exact and approximate NNS. This paper studies NNS under Uncertainty (NNSU). Specifically, consider the setting in which an NNS algorithm has access only to a stochastic distance oracle that provides a noisy, unbiased estimate of the distance between any pair of points, rather than the exact distance. This models many situations of practical importance, including NNS based on human similarity judgements, physical measurements, or fast, randomized approximations to exact distances. A naive approach to NNSU could employ any standard NNS algorithm and repeatedly query and average results from the stochastic oracle (to reduce noise) whenever it needs a pairwise distance. The problem is that a sufficient number of repeated queries is unknown in advance; e.g., a point may be distant from all but one other point (crude distance estimates suffice) or it may be close to a large number of other points (accurate estimates are necessary). This paper shows how ideas from cover trees and multi-armed bandits can be leveraged to develop an NNSU algorithm that has optimal dependence on the dataset size and the (unknown) geometry of the dataset.

1 INTRODUCTION

This paper considers Nearest Neighbor Search under Uncertainty (NNSU). To motivate the NNSU problem, consider the following application. Suppose we have a database of N genomic sequences of length L of different species and wish to query the database to find the closest relative of a newly discovered species. Computing the exact distance in $\mathbb{R}^L$ between two sequences requires $O(L)$ operations. Using efficient NNS algorithms, one can construct a data structure in $O(NL \log N)$ time and perform a NN search in $O(L \log N)$. To reduce this complexity, one can randomly subsample the sequences at $\ell \ll L$ locations and compute an unbiased estimate of distance in $O(\ell)$ time and find nearest neighbors in $O(\ell \log(n))$ operations. If an algorithm can manage the added uncertainty of this procedure, this can improve computational complexity greatly.

In other settings, uncertainty is naturally present in the measurements due to the presence of noise. For instance, researchers gathering data to map the topology of the internet, rely on noisy one-way-delay measurements to infer distance between servers Eriksson et al. [2010] and, detecting the closest server to a given host can be challenging due to this noise. In preference learning, researchers gather pairwise comparisons from people and attempt to learn which items are most preferred Jamieson and Nowak [2011], Yue and Joachims [2009]. To simplify such problems, human judgments are frequently modelled as distance comparisons Shepard [1962], Kruskal [1964], Mason et al. [2017] with the smallest distance representing the most preferred item, but the noise inherent to human judgements makes these problems challenging. In general, we define the NNSU problem as follows:

> **NNSU - Problem Statement:** Consider a set of n points $\mathcal{X} = \{\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n\}$ in a metric space $(\mathcal{M}, d)$. The metric is unknown, but for any pair of points we can query a stochastic oracle for a noisy, unbiased estimate of their distance. The problem is to use the oracle to efficiently build a data structure such that for any new query point $\boldsymbol{q}$, it returns its nearest neighbor $\boldsymbol{x}_q := \min_{\boldsymbol{x}_i \in \mathcal{X}} d(\boldsymbol{q}, \boldsymbol{x}_i)$ with probability at least $1 - \delta$ in as few additional oracle queries as possible.

The NNSU problem is very different from the standard Nearest Neighbor Search (NNS) problem[1]. To develop

---

[1]Note that NNSU is distinct from the problem of Approximate

noise-tolerant methods, it is tempting to simply extend an NNS algorithm by repeatedly calling the stochastic oracle a *pre-specified* r times each time a distance measurement is needed and averaging the results. The simplicity of this idea masks the difficulty of doing it correctly. If r is too small, an algorithm will make errors, but if r is too large, the number of distance queries and hence the computational cost will be unnecessarily high. To control the probability of error, r can be no smaller than $(d(\boldsymbol{q}, \boldsymbol{x}_q) - d(\boldsymbol{q}, \boldsymbol{x}_{q'}))^{-2}$ where $\boldsymbol{x}_{q'}$ is the second nearest neighbor to $\boldsymbol{q}$. Since this quantity depends on knowledge of the nearest neighbor distances, it is impossible to specify r in practice. Additionally, even if $(d(\boldsymbol{q}, \boldsymbol{x}_q) - d(\boldsymbol{q}, \boldsymbol{x}_{q'}))^{-2}$ were known, setting r in proportion to it is only necessary to handle the worst case; far fewer calls to the stochastic oracle are sufficient to eliminate points that are far from $\boldsymbol{q}$. Our proposed algorithm is *adaptive* to the level of noise and the geometry of $\mathcal{X}$. It is guaranteed to minimize the number of calls to the distance oracle for a specified probability of error.

1.1 RELATED WORK

One of the first algorithms for the NNS problem is the classical `KD-Tree` algorithm by Bentley [1975]. `KD-Tree`s first build a tree-based data structure in $O(n \log(n))$ computations. By paying this cost up front, when presented with a query point, the tree can be used to return the nearest neighbor in only $O(\log(n))$ computations. The core drawback of the `KD-Tree` is that it lacks a uniform accuracy guarantee and may fail to return the correct nearest neighbor for certain query points Dasgupta and Sinha [2013]. A variety of methods attempt to achieve similar computational complexity for the NNS problem while also providing robust accuracy guarantees Dasgupta and Sinha [2013], Krauthgamer and Lee [2004], Beygelzimer et al. [2006], and we refer the reader to Bhatia et al. [2010] for a survey of techniques and theoretical results. Another line of work attempts to instead return an approximate nearest neighbor that is almost as close as the true nearest neighbor. We refer the reader to Wang and Banerjee [2014], Andoni et al. [2018] for an overview.

A related problem to NNSU is the noisy nearest neighbor graph problem. In this setting, one wishes to learn the graph that connects each node in $\mathcal{X}$ to its nearest neighbor. Mason et al. [2019] provide an adaptive method that utilizes multi-armed bandits to find the nearest neighbor graph from a set of n points in $O(n \log(n)\Delta^2)$ samples in favorable settings and $O(n^2\Delta^2)$ at worst where Δ^2 is a problem dependent parameter quantifying the effect of the noise. Their method is adaptive to noise, but requires additional assumptions to achieve the optimal rate of $O(n \log(n))$ samples for that

---

Nearest Neighbor identification as in this case, the measurements themselves are corrupted by noise, but one wishes to find an exact nearest neighbor. We extend to finding an approximate nearest neighbor under uncertainty as well in this work.

problem. Bagaria et al. [2017] study nearest neighbor search in high-dimensional data. In their setting, the distance function is a known, componentwise function. Their algorithm subsamples to approximate distances. This improves dependence on dimension, though the dependence on n is linear.

1.2 MAIN CONTRIBUTIONS

In this paper, we leverage recent developments in multi-armed bandits to solve the nearest neighbor search problem using noisy measurements. We design a novel extension of the `Cover-Tree` algorithm from Beygelzimer et al. [2006] for the NNSU problem. A main innovation is reducing the problem of learning a cover of a set from noisy data to all-ϵ-good identification for multi-armed bandits studied in Mason et al. [2020]. Additionally, we make use of efficient methods for adaptive hypothesis testing to minimize calls to the stochastic oracle Jamieson and Jain [2018]. We refer to the resulting algorithm as the `Bandit-Cover-Tree`. We show that it requires $O(n \log^2(n)\kappa)$ calls to the distance oracle for construction and $O(\log(n)\kappa)$ calls for querying the nearest neighbor of a new point, where κ captures the effect of noise and the geometry of $\mathcal{X}$. This nearly matches the state of the art for the NNS problem despite the added challenge of uncertainty in the NNSU problem. Furthermore, we show how to extend `Bandit-Cover-Tree` for approximate nearest neighbor search instead. Finally, we demonstrate that `Bandit-Cover-Tree` can be used to learn a nearest neighbor graph in $O(n \log^2(n)\kappa)$ distance measurements, nearly matching the optimal rate given in Mason et al. [2019] but without requiring the additional assumptions from that work.

1.3 NOTATION

Let $(\mathcal{M}, d)$ be a metric space with distance function d satisfying the standard axioms. Given $\boldsymbol{x}_i, \boldsymbol{x}_j \in \mathcal{X}$, let $d(\boldsymbol{x}_i, \boldsymbol{x}_j) = d_{i,j}$. For a query point $\boldsymbol{q}$ define $\boldsymbol{x}_q := \arg\min_{\boldsymbol{x} \in \mathcal{X}} d(x_i, \boldsymbol{q})$. For a query point $\boldsymbol{q}$ and $\boldsymbol{x}_i \in \mathcal{X}$, let $d_{q,i}$ denote $d(\boldsymbol{q}, \boldsymbol{x}_i)$. Additionally, for a set $\mathcal{S}$, define the distance from any point $\boldsymbol{x} \in (\mathcal{M}, d)$ to $\mathcal{S}$ as $d(\boldsymbol{x}, \mathcal{S}) := \inf_{z \in S} d_{\boldsymbol{x}, z}$, the smallest distance a point in $\mathcal{S}$. 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

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

where η is a zero-mean subGaussian random variable assumed to have scale parameter $\sigma = 1$[2]. We let $\hat{d}_{i,j}(s)$ denote the empirical mean of the s queries of the distance oracle,

---

[2] $\sigma = 1$ can be relaxed to any subGuassian distribution. The subGaussian assumption itself can be relaxed by considering other confidence widths in our algorithm.

Q(i,j). The number of Q(i,j) queries made until time t is denoted as $T_{i,j}(t)$. All guarantees will be shown to hold with probability $1 - \delta$ where we refer to $\delta > 0$ as the failure probability.

2   COVER TREES FOR NEAREST NEIGHBOR SEARCH

Before presenting our method and associated results, we the review `Cover Tree` algorithm from Beygelzimer et al. [2006] for the NNS problem. As the name suggests, a cover tree is a tree-based data structure where each level of the tree forms a *cover* of $\mathcal{X}$.

Definition 1. *Given a set $\mathcal{X} \subset (\mathcal{M}, d)$, a set $C \subset \mathcal{X}$ is a cover of $\mathcal{X}$ with resolution ϵ if for all $\boldsymbol{x} \in \mathcal{X}$, there exists a $c \in C$ such that $d(\boldsymbol{x}, c) \leq \epsilon$.*

Each level is indexed by an integer i which decreases as one descends the tree. To avoid additional notation, we will represent the top of the tree as level ∞ and the bottom as $-\infty$ though in practice one would record integers i_{top} and i_{bottom} denoting the top and bottom level of the tree and need only explicitly store the tree between these levels. Each node in the tree corresponds to a point in $\mathcal{X}$, but points in $\mathcal{X}$ may correspond to multiple nodes in the tree. Reviewing Beygelzimer et al. [2006], let C_i denote the set of nodes at level i. The cover tree algorithm is designed so that each level of the tree i obeys three invariants:

1. **nesting:** $C_i \subset C_{i-1}$. Hence, the points corresponding to nodes at level i are also correspond to nodes in all lower levels.

2. **covering tree:** For every $p \in C_{i-1}$, there exists a $q \in C_i$ such that $d(p, q) \leq 2^i$, and child node p is connected to parent node q in the cover tree.

3. **Separation:** For any $p, q \in C_i$ with $p \neq q$, $d(p, q) > 2^i$.

These invariants are originally derived from Krauthgamer and Lee [2004]. Beygelzimer et al. [2006] show that their routines obey these invariants, and we will make use of them in our proofs and to build intuition. The core idea of this method is that the i^{th} level of the tree is a cover of resolution 2^i. Given a query point q, when navigating the tree at level i, one identifies all possible ancestor nodes of $\boldsymbol{x}_q$ at that level. When descending the tree, this set is refined until only a single parent is possible, $\boldsymbol{x}_q$ itself. The *nesting* invariance allows one to easily traverse the tree from parent to child. The *covering tree* invariance connects $\boldsymbol{x}_q$ to its parents and ancestors so that one may traverse the tree with query point $\boldsymbol{q}$ and end at $\boldsymbol{x}_q$. Lastly, the *separation* invariance ensures that there is a single parent to each node which avoids redundancy and improves memory complexity.

In order to quantify the sample complexity of cover trees, we require a notion of the effective dimensionality of the points $\mathcal{X}$. In this paper, we will make use of the *expansion constant* as in [Haghiri et al., 2017, Krauthgamer and Lee, 2004, Beygelzimer et al., 2006]. Let $B(\boldsymbol{x}, r)$ denote the ball of radius $r > 0$ centered at $\boldsymbol{x} \in (\mathcal{M}, d)$ according to the distance measure d.

Definition 2. *The expansion constant of a set of points $\mathcal{S}$ is the smallest $c \geq 2$ such that $|\mathcal{S} \cap B(\boldsymbol{x}, 2r)| \leq c|\mathcal{S} \cap B(\boldsymbol{x}, r)|$ for any $\boldsymbol{x} \in \mathcal{S}$ and any $r > 0$.*

The expansion constant is sensitive to the geometry in $\mathcal{X}$. For example, for points in a low-dimensional subspace, $c = O(1)$ independent of the ambient space. By contrast, if points that are spread out on the surface of a sphere in $\mathbb{R}^d$, $c = O(2^d)$. In general c is smaller if the pairwise distances between points in $\mathcal{X}$ are more varied.

3   THE BANDIT COVER TREE ALGORITHM FOR NNSU

The `Bandit Cover Tree` algorithm is comprised of three methods. Given a cover tree and query point $\boldsymbol{q}$, `Noisy-Find-Nearest` finds $\boldsymbol{q}$'s nearest neighbor. `Noisy-Insert` allows one to insert points into a tree and can be used to construct a new tree. `Noisy-Remove` can remove points from the tree and is deferred to the appendix.

3.1   FINDING NEAREST NEIGHBORS WITH A COVER TREE

We begin by discussing how to identify the nearest neighbor of a query point $\boldsymbol{q}$ in the set $\mathcal{X}$ given a cover tree $\mathcal{T}$ on $\mathcal{X}$. Throughout, we take $\mathcal{T}$ to denote the cover tree. We may identify $\mathcal{T}$ by the covers at each level: $\mathcal{T} := \{C_\infty, \cdots, C_i, C_{i-1}, \cdots, C_{-\infty}\}$. Assume that we are given a fixed query point $\boldsymbol{q}$, and the expansion constant of the set $\mathcal{X} \cup \{\boldsymbol{q}\}$ is bounded by c. Our method to find $\boldsymbol{x}_q$ is `Noisy-Find-Nearest` and is given in Algorithm 1. It is inspired by Beygelzimer et al. [2006], but a crucial difference is how the algorithm handles noise. At a high level, as it proceeds down each level i, it keeps track of a set $Q_i \subset C_i$ of all possible ancestors of $\boldsymbol{x}_q$. Given Q_i, to proceed to level $i - 1$, the algorithm first computes the set of children of nodes in Q_i given by the set $Q \subset C_{i-1}$. $Q_{i-1} \subset Q$ is then computed by forming a *cover* of Q at resolution 2^{i-1}. Ideally, this is the set

$$\tilde{Q}_{i-1} := \left\{ j \in Q : d(q, j) \leq \min_{k \in Q} d(q, k) + 2^{i-1} \right\} \quad (2)$$

However, constructing this set requires calling the stochastic distance oracle and presents a trade-off:

1. By using many repeated queries to the stochastic oracle, we can more confidently declare if $j \in \tilde{Q}_{i-1}$ at the expense of higher sample complexity.

2. By using fewer calls to the oracle, we can more quickly declare if $j \in \tilde{Q}_{i-1}$, at the expense of lower accuracy.

To handle the noise, we make use of anytime confidence widths, $C_\delta(t)$ such that for an empirical estimate of the distance $d_{q,j}$: $\hat{d}_{q,j}(t)$ of t samples, we have $\mathbb{P}(\bigcup_{t=1}^{\infty} |\hat{d}_{q,j}(t) - d_{q,j}| > C_\delta(t)) \leq \delta$. For this work, we take $C_\delta(t) = \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$ akin to Howard et al. [2018]. This allows us to quantify how certain or uncertain an estimate of any distance is. Furthermore, to guard against settings where the number of samples needed to detect the set $\tilde{Q}_{i-1}$ in Eq (2) is large, we introduce some slack and instead detect a set Q_i such that it contains every $j \in Q$ such that $d_{q,j} \leq d(\boldsymbol{q}, Q) + 2^{i-1}$ and no j worse than $d_{q,j} \leq d(\boldsymbol{q}, Q) + 2^{i-1} + 2^{i-2}$. In particular, this contains all points in $\tilde{Q}_{i-1}$ and none that are much worse. This has the advantage of controlling the number of calls to the stochastic oracle. As the algorithm descends the tree, i decreases and the slack of 2^{i-2} goes to zero ensuring that the algorithm returns the correct nearest neighbor. To handle both of these challenges, we reduce the problem of learning $\tilde{Q}_{i-1}$ to the problem of finding all ϵ-good arms in mutli-armed bandits studied in Mason et al. [2020] and refer to this method as `Identify-Cover` given in Alg. 2.

This process of exploring children nodes and computing cover sets terminates when the algorithm reaches level $i = -\infty$ and $Q_{-\infty}$ contains $\boldsymbol{x}_q$. Note that in practice the algorithm could terminate when it reaches the bottom of the cover tree, $i = i_{\text{bottom}}$. At this level, it calls a simple elimination scheme `find-smallest-in-set` given in the appendix that queries the stochastic oracle and returns $\boldsymbol{x}_q \in Q_{-\infty}$. To control the overall probability of error in `Noisy-Find-Nearest`, each call to `Identify-Cover` and `find-smallest-in-set` is run with failure probability δ/α for α chosen to control the family-wise error rate.

3.1.1 Approximate NNSU

In some applications, finding an exact nearest neighbor may be unnecessary and an *approximate* nearest neighbor may be sufficient. In the noiseless case, this has been shown to improve the complexity of nearest neighbor search Andoni et al. [2018]. An ϵ-approximate nearest neighbor is defined as any point in the set

$$\{i : d(q,i) \leq (1+\epsilon) \min_j d(q,j)\}.$$

In this section, we show to how to extend `Noisy-Find-Nearest` to return an ϵ-approximate nearest neighbor instead of the exact nearest neighbor and provide psuedocode for this method in the

---

Algorithm 1 `Noisy-Find-Nearest`

Require: Cover tree $\mathcal{T}$, failure probability δ, expansion constant c if known, query point $\boldsymbol{q}$, callable distance oracle $\mathsf{Q}(\cdot, \cdot)$, subroutine `Identify-Cover`.
1: Let $Q_\infty = C_\infty, \alpha = n + 1$
2: **for** $i = \infty$ down to $i = -\infty$ **do**
3:      Let $Q = \bigcup_{p \in Q_i} \text{children}(p)$
4:      **if** c is known: **then**
5:          Let $\alpha = \min\left\{ \left\lceil \frac{\log(n)}{\log(1+1/c^2)} + 1 \right\rceil, n \right\} + 1$
6:      \\ Identify the set: $\{j \in Q : d(q,j) \leq \min_{k \in Q} d(q,k) + 2^{i-1}\}$
7:      $Q_{i-1} = \texttt{Identify-Cover}(Q, \delta/\alpha, \mathsf{Q}(\cdot,\cdot), \boldsymbol{q}, i)$
8: **return** `find-smallest-in-set`$(Q_{-\infty}, \delta/\alpha)$

Algorithm 2 `Identify-Cover`

Require: Failure probability δ, query point $\boldsymbol{q}$, Oracle $\mathsf{Q}(\cdot, \cdot)$, and set $\boldsymbol{Q}$, Cover resolution 2^i
1: Query oracle once for each point in Q
2: Initialize $T_j \leftarrow 1$, update $\hat{d}_{q,j}$ for each $j \in Q$
3: Empirical cover set: $\widehat{G} = \{j : \hat{d}_{q,j} \leq \max_k \hat{d}_{q,k} + 2^i\}$
4: $U_t = \min_k \hat{d}_{q,k}(T_k) + C_{\delta/|Q|}(T_k) + 2^i$
5: $L_t = \min_k \hat{d}_{q,k}(T_k) - C_{\delta/|Q|}(T_k) + 2^i + 2^{i-1}$
6: Known points: $K = \{j : \hat{d}_{q,j}(T_j) - C_{\delta/|Q|}(T_j) > U_t$ or $\hat{d}_{q,j}(T_j) + C_{\delta/|Q|}(T_j) < L_t\}$
7: **while** $K \neq Q$ **do**
8:      $j_1(t) = \arg\min_{j \in \widehat{G} \setminus K} \hat{d}_{q,j}(T_j) + C_{\delta/|Q|}(T_j)$
9:      $j_2(t) = \arg\max_{j \in \widehat{G}^c \setminus K} \hat{d}_{q,j}(T_j) - C_{\delta/|Q|}(T_j)$
10:      $j^*(t) = \arg\min_j \hat{d}_{q,j}(T_j) - C_{\delta/|Q|}(T_j)$
11:      Call oracle $\mathsf{Q}(q, j_1)$, $\mathsf{Q}(q, j_2)$, and $\mathsf{Q}(q, j^*)$
12:      Update $T_{j_1}, \hat{d}_{q,j_1}, T_{j_2}, \hat{d}_{q,j_2}, T_{j^*}, \hat{d}_{q,j^*}$
13:      Update bounds L_t, U_t, sets $\widehat{G}, K$
14: **return** The set cover set with resolution 2^i: $\{j : \hat{d}_{q,j}(T_j) + C_{\delta/|Q|}(T_j) < L_t\}$

---

appendix. To do so, add an additional line that exits the for loop if $d(q, Q_i) \geq 2^{i+1}(1 + 1/\epsilon)$. Then, return $\boldsymbol{x}_{q_\epsilon} = \arg\min_{j \in Q_i} d(\boldsymbol{q}, \boldsymbol{x}_j)$. To see why this condition suffices, note that the nesting and covering tree invariances jointly imply that $d(\boldsymbol{x}_q, Q_i) \leq 2^{i+1}$. Hence, by the triangle inequality

$$d(\boldsymbol{q}, Q_i) \leq d(\boldsymbol{q}, \boldsymbol{x}_q) + d(\boldsymbol{x}_q, Q_i) \leq d(\boldsymbol{q}, \boldsymbol{x}_q) + 2^{i+1}.$$

Since we exit when $d(q, Q_i) \geq 2^{i+1}(1 + 1/\epsilon)$,

$$2^{i+1}(1 + 1/\epsilon) \leq d(\boldsymbol{q}, \boldsymbol{x}_q) + 2^{i+1} \iff 2^{i+1} \leq \epsilon d(\boldsymbol{q}, \boldsymbol{x}_q).$$

Therefore,

$$d(\boldsymbol{q}, Q_i) \leq d(\boldsymbol{q}, \boldsymbol{x}_q) + 2^{i+1} \leq (1+\epsilon)d(\boldsymbol{q}, \boldsymbol{x}_q).$$

Hence, there exists an ϵ-approximate nearest neighbor in Q_i. If this condition is never satisfied, which occurs when

ϵ is vanishingly small, the algorithm terminates normally and returns the exact nearest neighbor. We give pseudocode for this procedure in the appendix. The algorithm is similar to `Noisy-Find-Nearest` except that it makes use of a simple thresholding bandit, akin to the one we use in `Noisy-Insert` to check if $d(q, Q_i) \geq 2^{i+1}(1 + 1/\epsilon)$ by querying the stochastic oracle.

3.2 BUILDING AND ALTERING A COVER TREE

In this section we demonstrate how to insert points into a cover tree by calling the stochastic oracle. Constructing a cover tree can be achieved by simply beginning with an empty tree and inserting points one at a time. Removing points from a cover tree is similar operationally to insertion with slight complications and we defer it to the appendix.

3.2.1 Insertion

Suppose we have access to a cover tree $\mathcal{T}$ on a set $\mathcal{S}$. We wish to insert a point p into $\mathcal{T}$ such that we now have a cover tree on the set $\mathcal{S} \cup \{p\}$. Intuitively, the insertion algorithm can be thought of as beginning at the highest resolution cover, at level $-\infty$ and climbing back up the tree, inserting p in each cover set C_i for all i until it reaches a level i_p such that $\min_{j \in C_{i_p}} d(p, j) \leq 2^{i_p}$ where a suitable parent node exists. The algorithm then chooses a parent $p' \in C_{i_p}$ for p such that $d(p, p') \leq 2^{i_p}$ and terminates. As trees are traditionally traversed via their roots not their leaves, we state this algorithm recursively beginning at the top.

We provide pseudocode in Algorithm 3. The algorithm draws on ideas from Beygelzimer et al. [2006] but includes additional logic to handle uncertainty. In particular, lines 2-8 implement a simple thresholding bandit based on the techniques of Jamieson and Jain [2018] for adaptive hypothesis testing with family-wise probability of error control. This allows us to identify all points within 2^i of the nearest in the set Q. This must be done for every candidate level i that p might be inserted into until it reaches level i_p.

This requires careful handling to ensure the algorithm succeeds with high probability. Each time the thresholding operation is called, there is a chance that the algorithm makes a mistake. `Noisy-Insert` is recursive and performs a this operation in every recursive call. Hence, if it makes an error on any call, the algorithm may fail. Thus, we must ensure that the probability of `Noisy-Insert` making an error in any recursive call is at most δ. Since the level i_p where p is added is unknown and depends on the query point p, the number of recursive calls before success is unknown to the algorithm. Therefore, it is not a priori obvious how to perform the appropriate Bonferroni correction to ensure the probability of an error in any recursive call is bounded by δ. A seemingly attractive approach is to use a summable sequence of δ_i depending on the level i such that $\sum_i \delta_i = \delta$.

Algorithm 3 `Noisy-Insert`

Require: Cover tree $\mathcal{T}$ on n points, cover set Q_i, failure probability δ, point p to be inserted, callable distance oracle $\mathsf{Q}(\cdot, \cdot)$, level i

Optional: Empirical estimates of $\hat{d}_{p,i}$ and T_i for any $i \in C_i$

1: Let $Q = \bigcup_{j \in Q_i}$ children(j)
2: Query oracle once for each point in $Q \cap \{j : T_j = 0\}$
3: Set $T_i \leftarrow 1$, update $\hat{d}_{p,i}$ for each $j \in Q \cap \{j : T_j = 0\}$
4: \\ compute the set $\{j \in Q : d(p, j) \leq 2^i\}$
5: Known points: $K = \{j : \hat{d}_{p,j}(T_j) + C_{\delta/n}(T_j) \leq 2^i$ or $\hat{d}_{p,j}(T_j) - C_{\delta/n}(T_j) > 2^i\}$
6: **while** $|K| \neq |Q|$ **do**
7: $\quad$ Compute $j^*(t) = \arg\min_{j \notin K} \hat{d}_{p,j}(T_j) - C_{\delta/n}(T_j)$
8: $\quad$ Call oracle $\mathsf{Q}(p, j^*)$ and update $T_{j^*}, \hat{d}_{p,j^*}, K$
9: \\ If $d(p, Q) > 2^i$
10: $Q_{i-1} = \{j \in Q : \hat{d}_{p,j}(T_j) + C_{\delta/n}(T_j) \leq 2^i\}$
11: **if** $Q_{i-1} = \emptyset$ **then**
12: $\quad$ **Return:** "no parent found"
13: **else**
14: $\quad$ lower $=$ `Noisy-Insert`$(p, \mathcal{T}, Q_{i-1}, i - 1, \delta, \mathsf{Q}(\cdot, \cdot), \{j \in Q : \hat{d}_{p,j}(T_j)\}, \{j \in Q : T_j\})$
15: $\quad$ **if** $Q_i \cap Q_{i-1} \neq \emptyset$ and lower $=$ "no parent found" **then**
16: $\quad\quad$ Choose any $p' \in Q_i \cap Q_{i-1}$
17: $\quad\quad$ Insert p in children(p') \\ Assign a parent to p
18: $\quad\quad$ **Return:** "parent found"
19: $\quad$ **else**
20: $\quad\quad$ **Return:** "no parent found"

For instance, $\delta_i = \delta \cdot 2^{-i}$ would be suitable if the root of $\mathcal{T}$ is at level 1. However, due to repeated calls to lines 2-8, this would lead to an additional multiplicative factor of the height of the cover tree affecting the sample complexity.

Instead, `Noisy-Insert` *shares* samples between rounds. By the nesting invariance, we have that $C_i \subset C_{i-1}$. Therefore, when we descend the tree from level i to $i - 1$, we already have samples of the distance of some points in C_{i-1} to p. We simply reuse these samples and share them from round to round. Furthermore, since $\mathcal{T}$ is assumed to be a cover tree on n points, we trivially union bound each confidence width to hold with probability $1 - \delta/n$ such that all bounds for all recursive calls holds with probability at least $1 - \delta$.

4 THEORETICAL GUARANTEES OF BANDIT COVER TREE

We measure performance along several axes: 1) how much memory is necessary to store the data structure, 2) how many calls to the distance oracle are needed at query time, 3) How many calls to the distance oracle are needed to for construction, and 4) the accuracy of the data structure in

returning the correct nearest neighbor and performing other operations. Since we only have access to a stochastic oracle, ensuring the accuracy of the `Bandit Cover Tree` (BCT) is especially delicate.

4.1 MEMORY

We begin by showing that a cover tree can efficiently be stored. Naively, a cover tree $\mathcal{T}$ on $\mathcal{X}$ can be stored using $O(n(i_{\text{top}} - i_{\text{bottom}}))$ memory where $n = |\mathcal{X}|$ and $i_{\text{top}} - i_{\text{bottom}}$ is the height of the tree. This follows from each level having at most n nodes trivially and there being $i_{\text{top}} - i_{\text{bottom}}$ levels. For a well balanced tree, we expect that $i_{\text{top}} - i_{\text{bottom}} = O(\log(n))$ leading to an overall memory complexity of $O(n \log(n))$. In fact, it is possible to do better.

Lemma 3. *A bandit cover tree requires $O(n)$ space to be stored.*

$O(n)$ memory is possible due to the nesting and covering tree invariants. By the nesting invariant, if a point p is present in the i^{th} cover C_i, then it is present in the cover sets of all lower levels. By the covering tree invariant, each point has a unique parent in the tree. Therefore, to store a cover tree, one need only store 1) which level of the tree each point first appears in 2) each point's parent node in the level above where it appears. After a node first appears in the tree, it may be represented implicitly in all lower levels. In particular, when all methods query the children of any node p, we define $p \in \text{children}(p)$. Formally, we define an implicit and explicit node as follows:

Definition 4. *A node $p \in C_i$ in level i of cover tree $\mathcal{T}$ is explicit if $p \notin \{C_\infty, \cdots, C_{i+1}\}$ (i.e. it is not present in all previous levels.) Otherwise, any $p \in C_i$ is referred to as implicit.*

4.2 ACCURACY

Next, we show that BCT is accurate with high probability. This requires three guarantees:

1. Search accuracy: Given a cover tree $\mathcal{T}$ on set $\mathcal{X}$ and query point $\boldsymbol{q}$, `Noisy-Find-Nearest` correctly identifies $\boldsymbol{q}$'s nearest neighbor with high probability.

2. Insertion accuracy: Given a cover tree $\mathcal{T}$ and a point p to be inserted, `Noisy-Insert` returns a valid cover tree that includes p with high probability.

3. Removal accuracy: Given a cover tree $\mathcal{T}$ and a point p to be removed, `Noisy-Remove` returns a valid cover tree that without p (deferred to appendix).

4.2.1 Search Accuracy

We begin by showing that for any $\boldsymbol{q} \in (\mathcal{M}, d)$ `Noisy-Find-Nearest` returns $\boldsymbol{x}_q$ with high probability. The proof is deferred to the appendix, but the argument is sketched as follows. First, by appealing to results from Mason et al. [2020], we can guarantee that `Identify-Cover` succeeds with probability $1 - \delta/\alpha$. Second, we show that for the choice of α in `Noisy-Find-Nearest`, the probability that any call to `Identify-Cover` fails is bounded by $1 - \delta$. Finally, we show that if we correctly identify the necessary cover sets, which happens with probability at least $1 - \delta$, we my appeal to Theorem 2 of Beygelzimer et al. [2006] to ensure that our method returns $\boldsymbol{x}_q$.

Lemma 5. *Fix any $\delta \leq 1/2$ and a query point $\boldsymbol{q}$. Let $\mathcal{T}$ be a cover tree on a set $\mathcal{X}$. `Noisy-Find-Nearest` returns $\boldsymbol{x}_q \in \mathcal{X}$ with probability at least $1 - \delta$.*

4.2.2 Insertion Accuracy

The proof that `Noisy-Insert` succeeds with high probability in adding any point p to a cover tree $\mathcal{T}$ follows the argument of Lemma 5 similarly. In particular, given a set Q, it hinges on the idea that we can use a threshold bandit to identify the set $\{j \in Q : d(p, j) \leq 2^i\}$ correctly with high probability by calling the stochastic oracle. If this procedure succeeds, then `Noisy-Insert` has correctly identified the subset of Q that are within 2^i of the point p to be inserted. If $\{j \in Q : d(p, j) \leq 2^i\} = \emptyset$, we may instead verify that $d(p, Q) > 2^i$ and can compute a new cover set Q_{i-1}. The following lemma ensures that `Noisy-Insert` succeeds with high probability.

Lemma 6. *Fix any $\delta > 0$. Let $\mathcal{T}$ be a cover tree on a set $\mathcal{X}$ and p be a point to insert. `Noisy-Insert` correctly returns a cover tree on $\mathcal{X} \cup \{p\}$ with probability at least $1 - \delta$.*

To construct a cover tree from scratch given a set $\mathcal{X}$, one need only call `Noisy-Insert` n times, once for each point in $\mathcal{X}$ beginning with an empty tree and run each call with failure probability δ/n. The following corollary ensures that this leads to a cover tree satisfying the three invariances with probability $1 - \delta$.

Corollary 7 (Construction of a Cover Tree). *Fix any $\delta > 0$ and a set $\mathcal{X} \in (\mathcal{M}, d)$. Calling `Noisy-Insert` with failure probability δ/n iteratively over the points in $\mathcal{X}$ yields a cover tree that satisfies the nesting, covering tree, and separation invariances with probability $1 - \delta$.*

4.3 QUERY TIME COMPLEXITY

In the previous section, we proved that the algorithm succeeds with probability $1 - \delta$ for search and insertion, with

removal deferred to the appendix. In this section we begin to answer the question of how many calls to the stochastic distance oracle it requires to perform these operations.

We begin by analyzing the query time complexity of `Bandit Cover Tree`: the number of calls to the distance oracle made when answering a nearest neighbor query. To do so, we will make use of the *expansion constant*, a data-dependent measure of dimensionality. In particular, for a query point q. We assume that the set $\mathcal{X} \cup \{q\}$ has an expansion constant of c as defined in Definition 2. Note that this quantity is for analysis purposes only and is not required by the algorithm. `Noisy-Find-Nearest` can take in the expansion constant or a bound on it if it is known, but this is not required by the algorithm. In particular, knowing a bound on c helps control the height of the tree. This allows `Noisy-Find-Nearest` to use a smaller union bound and leads to a tighter dependence on n.

To bound query time, we appeal to the concept of explicit and implicit nodes as discussed in Section 4.1 and Definition 4. Each point in $\mathcal{X}$ may correspond to multiple nodes in the tree. The first time a point appears as a node at the highest level where it is present, we say that it is *explicitly represented* and is *implicitly represented* in all levels thereafter by the nesting invariant of cover trees. Recall that the set of nodes at each level i of the cover tree is denoted C_i. `Noisy-Find-Nearest` proceeds by computing a cover set $Q_i \subset C_i$ at each level of the tree. Extending the concept of explicit and implicit representations of nodes, we say that a cover set Q_i is implicitly represented if it only contains nodes that are implicitly represented. This plays an important role in our computation of query time. We may use the size of the last explicitly represented cover to help bound to complexity of the `Identify-Cover` subroutine in Algorithm 2. This routine is called for every level of the tree. Then, by bounding the height of the tree, we can control the total number of calls to the distance oracle.

Theorem 8. *Fix $\delta < 1/2$, a cover tree $\mathcal{T}$ on set $\mathcal{X}$, and a query point q. Let $|\mathcal{X}| = n$ and assume that the expansion rate of $\mathcal{X} \cup \{q\}$ is c (unknown to the algorithm). If `Noisy-Find-Nearest` succeeds, which occurs with probability $1 - \delta$, then `Noisy-Find-Nearest` returns q's nearest neighbor in at most*

$$O\left(c^{17} \log(n) \log\left(\frac{n}{\delta}\right) \kappa\right)$$

calls to the stochastic distance oracle where parameter κ is defined in the proof and satisfies

$$\kappa \leq B \max\left\{[d(q, x_q) - d(q, x_{q'})]^{-2}, d_{\min}^{-2}\right\}$$

where $x_{q'}$ is q's second nearest neighbor, $d_{\min} := \min_{j,k}(x_j, x_k)$ is the smallest pairwise distance in $\mathcal{X}$, and $B > 0$ is a universal constant.[3]

---

[3] The dependence on c is large theoretically, however Beygelzimer et al. [2006] suffer similar theoretical dependencies on c but show that in practice the dependence is more mild.

Remark 9. *The dependence of c^{17} can be improved to c^7 matching Beygelzimer et al. [2006] by redefining $L_t = \min_j \hat{d}_{q,j}(T_j) - C_{\delta/|Q|}(T_j) + 2^i$ in the pseudocode of `Identify-Cover`. For most instances, this will lead to better performance, though a potentially worse worst-case bound on κ in pathological cases.*

The above result scales as $O\left(c^{O(1)} \log^2(n)\right)$. The following corollary highlights that if c is known, this can be improved to $O\left(c^{O(1)} \log(n) \log(\log(n))\right)$.

Corollary 10. *Under the same conditions as Theorem 8, if any bound $\tilde{c} \geq c$ on the expansion rate c is given to the algorithm, the number of queries to the distance oracle is at most*

$$O\left(\tilde{c}^{17} \log(n) \log\left(\frac{\tilde{c} \log(n)}{\delta}\right) \kappa\right)$$

calls to the stochastic oracle.

The term κ captures the average effect of noise on this problem. It is similar to the term $\overline{\Delta^{-2}}$ in given in Mason et al. [2019] for identifying the Nearest Neighbor Graph of $\mathcal{X}$ from a stochastic distance oracle. As the noise variance goes to 0, this term becomes 1. Furthermore, κ is adaptive to the geometry of $\mathcal{X}$ and in most cases is much smaller than its bound in Theorem 8. Intuitively, the bound on κ states that the number of repeated samples is never worse that what is necessary to 1) distinguish between q's nearest neighbor and second nearest neighbor and 2) distinguish between any two points in $\mathcal{X}$. Crucially, however, `Noisy-Find-Nearest` *adapts* to the level of noise, the geometry of $\mathcal{X}$ and need not know these values a priori.

4.4 INSERTION TIME AND BUILD TIME COMPLEXITY

Next, we bound the number of calls to the distance oracle necessary to insert a new point p into a cover tree $\mathcal{T}$.

Theorem 11. *Fix $\delta > 0$, a cover tree $\mathcal{T}$ on set $\mathcal{X}$, and a point to insert p. Let $|\mathcal{X}| = n$ and assume that the expansion rate of $\mathcal{X} \cup \{p\}$ is c. Run `Noisy-Insert` with failure probability $1 - \delta$ and pass it the root level cover set: $C_{i_{top}}$ and level $i = i_{top}$. If `Noisy-Insert` succeeds, which occurs with probability $1 - \delta$, then it returns a cover tree on $\mathcal{X} \cup \{p\}$ in at most*

$$O\left(c^7 \log(n) \log\left(\frac{n}{\delta}\right) \overline{\kappa_p}\right)$$

calls to the noisy distance oracle where parameter $\overline{\kappa_p}$ is defined in the proof and depends on $\mathcal{X}$ and p.

Remark 12. *As in the statement of Theorem 8, the term $\overline{\kappa_p}$ captures the average effect of noise on this problem, and as the noise variance goes to 0, this term becomes 1.*

As discussed in Section 3.2.1, to construct a cover tree from scratch, one need only call `Noisy-Insert` on each point in $\mathcal{X}$ and add them the tree one at a time. The following theorem bounds the complexity of this process.

Theorem 13. *Fix $\delta > 0$ and set n points $\mathcal{X}$. Assume that the expansion rate of $\mathcal{X}$ is c. Calling `Noisy-Insert` with failure probability δ/n on each point in $\mathcal{X}$ one at a time, returns a cover tree $\mathcal{T}$ on $\mathcal{X}$ correctly with probability at least $1 - \delta$ in at most*

$$O\left(c^7 n \log(n) \log\left(\frac{n}{\delta}\right) \widetilde{\kappa}\right)$$

calls to the noisy distance oracle where $\widetilde{\kappa} := \frac{1}{n} \sum_{i \in \mathcal{X}} \overline{\kappa}_i$ for $\overline{\kappa}_i$ defined in the proof of Theorem 11.

4.5 EXTENSION TO THE NEAREST NEIGHBOR GRAPH PROBLEM

In this section we extend the results of Mason et al. [2019] for learning *nearest neighbor graphs* from a stochastic distance oracle. In this setting, given a set $\mathcal{X} \in (\mathcal{M}, d)$, one wishes to learn the directed graph $G(\mathcal{X})$ that connects each $\boldsymbol{x} \in \mathcal{X}$ to its nearest neighbor in $\mathcal{X} \backslash \{\boldsymbol{x}\}$. While Mason et al. [2019] also leverage a bandit subroutine, the bandit algorithm considered in this work is markedly different from the one considered in Mason et al. [2019]. This problem is well studied in the noiseless regime and at times referred to as the all nearest neighbor problem Clarkson [1983], Sankaranarayanan et al. [2007], Vaidya [1989]. Mason et al. [2019] provide the first algorithm that learns the nearest neighbor graph using only a stochastic oracle and show that in special cases it achieves a rate of $O(n \log(n) \Delta^{-2})$ which is matches the optimal rate for the noiseless problem with only a multiplicative penalty of Δ^{-2} accounting for the effect of noise. Unfortunately, the condition necessary to show that result is stringent. In fact, it can be shown that the separation condition of Mason et al. [2019] is equivalent to the special case that $c = 2$. The authors conjectured that the condition was not necessary and showed empirically that it should be possible to achieve $O(n \log(n))$ complexity in the noisy setting without their separation condition. This motivates the question of if similar performance can be achieved under more general conditions than those needed in Mason et al. [2019]. We answer this question in the affirmative and show how to extend the `Bandit-Cover-Tree` to achieve near optimal performance on the nearest neighbor graph problem without the need for any additional assumptions on the data. This proceeds in two steps:

1. Build a cover tree $\mathcal{T}$ on $\mathcal{X}$ with probability $1 - \delta/2$.
2. For each $\boldsymbol{x} \in \mathcal{X}$, find its nearest neighbor in $\mathcal{X} \backslash \{\boldsymbol{x}\}$ using $\mathcal{T}$ with probability $1 - \delta/2n$.

The above is sufficient to specify each edge of the nearest neighbor graph. Note that $\mathcal{T}$ is a cover tree on $\mathcal{X}$ not $\mathcal{X} \backslash \{\boldsymbol{x}\}$

so we cannot blindly use `Noisy-Find-Nearest` as it will return $\boldsymbol{x}$ itself. To find $\boldsymbol{x}$'s nearest neighbor in $\mathcal{X} \backslash \{\boldsymbol{x}\}$, one may instead modify `Noisy-Find-Nearest` so that `Identify-Cover` is called on the set $Q_i \backslash \{\boldsymbol{x}\}$ instead of Q_i. Then, when the algorithm terminates, the final set $Q_{i_{\text{bottom}}}$ will have 2 points: $\boldsymbol{x}$ and its nearest neighbor in $\mathcal{X} \backslash \{\boldsymbol{x}\}$. A simple union bound ensures that this process succeeds with probability $1 - \delta$. The following Lemma bounds the total number of samples.

Lemma 14. *Via the above procedure, `Bandit-Cover-Tree` returns the nearest neighbor graph of $\mathcal{X}$ with probability $1 - \delta$ in*

$$O\left(c^7 n \log(n) \log\left(\frac{n}{\delta}\right) \widetilde{\kappa} + \sum_{\boldsymbol{x} \in \mathcal{X}} c^7 \log(n) \log\left(\frac{n}{\delta}\right) \kappa_{\boldsymbol{x}}\right)$$

for $\widetilde{\kappa}$ defined in Theorem 13 and κ_x as in Theorem 8.

The proof follows by combining the guarantees of Theorems 8 and 13. In particular, the above bound scales as $O(c^7 n \log^2(n) \kappa)$ which matches the rate of Mason et al. [2019] with an additional $\log$ factor. Importantly, the above result makes *no assumptions* on the set $\mathcal{X}$, and instead the bound scales with the expansion rate, c. Hence, we achieve near optimal performance for the nearest neighbor graph problem with under *far more general* conditions.

5 CONCLUSION

In this paper, we introduced the `Bandit Cover Tree` framework for the Nearest Neighbor Search under Uncertainty (NNSU) problem. `BCT` builds on top of the `Cover Tree` algorithm by Beygelzimer et al. [2006]. We present three methods, `Noisy-Find-Nearest`, `Noisy-Insert`, and `Noisy-Remove`, and bound their accuracy, memory footprint, build complexity, insertion and removal complexities, and query complexities. In particular, we show a query complexity that is $O(\log(n))$, insertion complexity that is $O(\log^2(n))$, removal complexity that is $O(\log^2(n))$, and a construction complexity of $O(n \log^2(n))$. The query complexity matches the state of the art n dependence for the NNS problem of $O(\log(n))$ despite the added uncertainty of the NNSU problem. The additional $\log(n)$ term present in the insertion and removal guarantee stems from a union bound which is necessary when dealing with a stochastic oracle, though it may be possible that by giving the algorithm access to the expansion constant, this can be improved to an additional doubly logarithmic term instead. Hence, the insertion, construction, and removal complexities are also near the state of the art for NNS. Lastly a memory footprint of $O(n)$ and accuracy of $1 - \delta$ are both optimal in this problem. In particular, a tree with n leaves requires $\Omega(n)$ memory to store, and though a dependence on δ is unavoidable when dealing

with a stochastic oracle, BCTs enjoy a probability $1 - \delta$ probability of success for any specified $\delta > 0$.

Note that we focus on controlling the number of calls to the distance oracle in this work and assume that an individual call requires $O(1)$ work. We expect that for practical problems the bulk of the computational effort will be many repeated calls to the oracle. It is an open question for future work to control the computational complexity in the noisy regime considering all operations, not just calls to the stochastic oracle.

Furthermore, the results depend strongly on the expansion rate, c. Some works such as Haghiri et al. [2017], Dasgupta and Sinha [2013] trade accuracy for improved dependence on c or other measures of dimension. Instead, these algorithms guarantee that x_q is correctly returned with probability at least $1 - \delta_{c,n}$ for any q. Though $\delta_{c,n}$ is tunable, it often depends on c or other parameters potentially unknown to the practitioner. These methods often achieve good empirical performance, however. It may be interesting to develop methods that achieve lower theoretical accuracy but enjoy a softer dependence on c.

References

Alexandr Andoni, Piotr Indyk, and Ilya Razenshteyn. Approximate nearest neighbor search in high dimensions. *arXiv preprint arXiv:1806.09823*, 7, 2018.

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.

Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.

Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104. ACM, 2006.

Nitin Bhatia et al. Survey of nearest neighbor techniques. *arXiv preprint arXiv:1007.0085*, 2010.

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.

Sanjoy Dasgupta and Kaushik Sinha. Randomized partition trees for exact nearest neighbor search. In *Conference on Learning Theory*, pages 317–337, 2013.

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.

Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Pac bounds for multi-armed bandit and markov decision processes. In *International Conference on Computational Learning Theory*, pages 255–270. Springer, 2002.

Siavash Haghiri, Debarghya Ghoshdastidar, and Ulrike von Luxburg. Comparison based nearest neighbor search. *arXiv preprint arXiv:1704.01460*, 2017.

Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Uniform, nonparametric, non-asymptotic confidence sequences. *arXiv preprint arXiv:1810.08240*, 2018.

Kevin Jamieson and Lalit Jain. A bandit approach to multiple testing with false discovery control. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 3664–3674, Red Hook, NY, USA, 2018. Curran Associates Inc.

Kevin G Jamieson and Robert Nowak. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems*, pages 2240–2248, 2011.

Robert Krauthgamer and James R Lee. Navigating nets: simple algorithms for proximity search. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 798–807. Society for Industrial and Applied Mathematics, 2004.

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

Blake Mason, Lalit Jain, and Robert Nowak. Learning low-dimensional metrics. In *Advances in neural information processing systems*, pages 4139–4147, 2017.

Blake Mason, Ardhendu Tripathy, and Robert Nowak. Learning nearest neighbor graphs from noisy distance samples. In *Advances in Neural Information Processing Systems*, pages 9586–9596, 2019.

Blake Mason, Lalit Jain, Ardhendu Tripathy, and Robert Nowak. Finding all $\{\epsilon\}$-good arms in stochastic bandits. *Advances in Neural Information Processing Systems*, 2020.

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.

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

Pravin M Vaidya. An o(n logn) algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4(2):101–115, 1989.

Huahua Wang and Arindam Banerjee. Randomized block coordinate descent for online and stochastic optimization. *arXiv preprint arXiv:1407.0107*, 2014.

Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1201–1208, 2009.

CONTENTS

A ADDITIONAL METHODS AND PSUEDOCODE

A.1 FINDING THE CLOSEST POINT IN A SET

In this subsection, we provide the additional method used by `Noisy-Find-Nearest` to find the closest point among a set of alternatives.

Algorithm 4 find-smallest-in-set

Require: Set Q, query point $\boldsymbol{q}$, callable distance oracle $Q(\cdot, \cdot)$
1: **if** $|Q| = 1$ **then return** $i \in Q$
2: Call oracle for each $j \in Q$ and initialize $T_j \to 1$
3: Define $j^* = \arg\min_{j \in Q} \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n}(T_{q,j})$
4: Define set $K = \{j : \hat{d}_{q,j^*}(T_{q,j^*}) + C_{\delta/n}(T_{q,j^*}) < \hat{d}_{q,j}(T_{q,j}) - C_{\delta/n}(T_{q,j})\}$
5: **while** $|Q \backslash K| > 1$ **do**
6:     Call oracle $Q(q, j)$ for every $j \in Q \backslash K$, update T_js and set K
    return Return $i \in Q \backslash K$

A.2 APPROXIMATE NEAREST NEIGHBORS

Here we provide psuedocode for the `Approx-Noisy-Find-Nearest` routine. It is similar to `Noisy-Find-Nearest` except that it uses a thresholding bandit similar to the one used in `Noisy-Insert` and `Noisy-Remove` to check the condition $d(\boldsymbol{q}, Q_i) \geq 2^{i-1}(1 + 1/\epsilon)$.

Algorithm 5 `Approx-Noisy-Find-Nearest`

Require: Cover tree $\mathcal{T}$, failure probability δ, expansion constant c if known, query point $\boldsymbol{q}$, callable distance oracle $Q(\cdot, \cdot)$, subroutines `Identify-Cover` and `find-smallest-in-set`.
1: Let $Q_\infty = C_\infty$, $\alpha = n$
2: **for** $i = \infty$ down to $i = -\infty$ **do**
3:     Let $Q = \bigcup_{p \in Q_i}$ children(p)
4:     **if** c is known: **then**
5:         Let $\alpha = \min\left\{\left\lceil \frac{\log(n)}{\log(1 + 1/c^2)} + 1 \right\rceil, n\right\}$
6:     \\ Identify the set: $\{j \in Q : d(q, j) \leq \min_{k \in Q} d(q, k) + 2^{i-1}\}$
7:     $Q_{i-1} = $ `Identify-Cover` $(Q, \delta/\alpha, Q(\cdot, \cdot), \boldsymbol{q}, i)$
8:     \\ Check if $d(q, Q_{i-1}) \geq 2^i(1 + 1/\epsilon)$
9:     Call oracle for each $j \in Q_{i-1}$ and initialize $T_j \to 1$
10:     Define set $K = \{j : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n}(T_{q,j}) \leq 2^i(1 + 1/\epsilon)$ or $\hat{d}_{q,j}(T_{q,j}) - C_{\delta/n}(T_{q,j}) > 2^i(1 + 1/\epsilon)\}$
11:     **while** $|K| \neq |Q_{i-1}|$ **do**
12:         Call oracle $Q(q, j^*)$ for $j^*(t) = \arg\min_{j \notin K} \hat{d}_{q,j}(T_{q,j}) - C_{\delta/n}(T_{q,j})$
13:         **if** $\exists j \in Q_{i-1} : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n}(T_{q,j}) \leq 2^i(1 + 1/\epsilon)$ **then**
14:             **Continue** to next iteration of For loop
15:     \\ We have satisfied the condition that $d(\boldsymbol{q}, Q_{i-1}) \geq 2^i(1 + 1/\epsilon)$ and can return the closest.
        return `find-smallest-in-set`(Q_{i-1})
16: **return** $Q_{-\infty}$, a singleton set containing $\boldsymbol{x}_q$.

A.3 ALGORITHM TO REMOVE A POINT FROM A COVER TREE

A.3.1 Algorithm

Given a cover tree $\mathcal{T}$ and a point p that exists in the tree, we show how to remove p efficiently such that the resulting tree also forms a cover tree the obeys the nesting, covering tree, and separation invariances. The process of removal is similar to insertion but slightly more complicated as we must find new parents for all of p's children in $\mathcal{T}$. We provide pseudocode

in Algorithm 6. Similar to `Noisy-Insert` it adapts a threshold bandit routine to minimize the number of calls to the stochastic oracle while controlling the probability of error.

Algorithm 6 `Noisy-Remove`

Require: Cover tree $\mathcal{T}$ on n points, past cover sets $\{Q_i, \cdots, Q_\infty\}$, failure probability δ, point p to be removed, callable distance oracle $Q(\cdot, \cdot)$, level i

Require: Empirical estimates $\hat{d}_{p,j}$ and T_j for all $j \in C_i$ \\ let both be 0 if no samples have been collected

1: Let $Q = \bigcup_{p \in Q_i} \text{children}(p)$
2: \\ Remove from lower levels first
3: $\mathcal{T}, \{j, k : \hat{d}_{j,k}(T_{j,k})\}, \{j, k : T_{j,k}\} \leftarrow$ `Noisy-Remove`$(\mathcal{T}, \{Q_{i-1}, \cdots, Q_\infty\}, \delta, p, Q(\cdot, \cdot), i - 1, \{j, k : \hat{d}_{j,k}(T_{j,k})\}, \{j, k : T_{j,k}\})$
4: **if** $p \in Q$ **then**
5:     Remove p from C_{i-1} and children($\text{parent}(p)$)
6:     **for** $q \in \text{children}(p)$ **do**
7:         $i' \leftarrow i - 1$
8:         \\ compute the set $\{j \in Q_{i'} : d(q, j) \leq 2^{i'}\}$
9:         Query oracle once for each point in $Q \cap \{j : T_{q,j} = 0\}$
10:         Initialize $T_{q,j} \leftarrow 1$, update $\hat{d}_{q,j}$ for each $j \in Q \cap \{j : T_{q,j} = 0\}$
11:         Known points: $K = \{j : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n^2}(T_{q,j}) \leq 2^{i'} \text{ or } \hat{d}_{q,j}(T_{q,j}) - C_{\delta/n^2}(T_{q,j}) > 2^{i'}\}$
12:         **while** $|K| \neq |Q_{i'}|$ **do**
13:             Call oracle $Q(q, j^*)$ for $j^*(t) = \arg\min_{j \notin K} \hat{d}_{q,j}(T_{q,j}) - C_{\delta/n^2}(T_{q,j})$
14:             Update $T_{q,j^*}, \hat{d}_{q,j^*}$
15:             Update set K
16:         **while** $\{j \in Q_{i'} : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n^2}(T_{q,j}) \leq 2^{i'}\} = \emptyset$ **do**
17:             Add q to the sets $C_{i'}$ and $Q_{i'}$. Increment i'.
18:             \\ for the incremented i', recompute the set $\{j \in Q_{i'} : d(q, j) \leq 2^{i'}\}$
19:             Known points: $K = \{j : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n^2}(T_{q,j}) \leq 2^{i'} \text{ or } \hat{d}_{q,j}(T_{q,j}) - C_{\delta/n^2}(T_{q,j}) > 2^{i'}\}$
20:             **while** $|K| \neq |Q_{i'}|$ **do**
21:                 Call oracle $Q(q, j^*)$ for $j^*(t) = \arg\min_{j \notin K} \hat{d}_{q,j}(T_j) - C_{\delta/n^2}(T_j)$
22:                 Update $T_{q,j^*}, \hat{d}_{q,j^*}$
23:                 Update set K
24:         Choose any $q' \in \{j \in Q_{i'} : \hat{d}_{q,j}(T_j) + C_{\delta/n}(T_j) \leq 2^{i'}\}$
25:         make $q' = \text{parent}(q)$

Note that when we compute the distance to the set $Q_{i'}$ in lines 13 and 20, due to the nesting invariance of cover trees and the fact that samples are reused between rounds, $T_i > 0$ for all $i \in Q_{i'}$. Hence, it is unnecessary to collect any initial samples. Note that `Noisy-Remove` not only queries distances to the point p to be removed but also to p's children in order to find them new parents. Because of this, we index samples as $T_{i,j}$ denoting the number of calls to the distance oracle $Q(i, j)$. Furthermore, we union bound with δ/n^2 instead of δ/n where the additional factor of n derives from a trivial bound that p has at most $n - 1$ children nodes. If the expansion rate is known, it is possible to union bound as $\delta/c^4 n$ since Lemma 23 bounds the number of children as c^4 for any node p. However, as the other factor of n remains unchanged by this, the improvement from $\log(n^2/\delta)$ to $\log(c^4 n/\delta)$ only impacts constant factors in the sample complexity.

A.3.2 Theoretical Guarantees

The following lemma guarantees that `Noisy-Remove` succeeds with high probability. The argument is similar to that of Lemma 6 in that it relies on correctly identifying the set $\{i \in Q : d(p, i) \leq 2^i\}$.

Lemma 15. *Fix any $\delta > 0$. Let $\mathcal{T}$ be a cover tree on a set $\mathcal{X}$ and $p \in \mathcal{X}$ be a point to remove. `Noisy-Remove` correctly returns a cover tree on $\mathcal{X}\backslash\{p\}$ with probability at least $1 - \delta$.*

Next we bound the number of calls to the distance oracle necessary to remove a point p from a cover tree $\mathcal{T}$.

Theorem 16. *Fix $\delta > 0$, a cover tree $\mathcal{T}$ on set $\mathcal{X}$, and a point $p \in \mathcal{X}$ to remove. Let $|\mathcal{X}| = n$ and assume that the expansion rate of $\mathcal{X}$ is c. Run `Noisy-Remove` with failure probability $1 - \delta$ and pass it the root level cover set: $C_{i_{top}}$. If `Noisy-Remove` succeeds, which occurs with probability $1 - \delta$, then it returns a cover tree on $\mathcal{X}\backslash\{p\}$ in at most*

$$O\left(c^{11} \log^2(n) \log\left(\frac{n^2}{\delta}\right) \widehat{\kappa}_p\right)$$

calls to the noisy distance oracle where parameter $\widehat{\kappa}_p$ is defined in the proof and depends on $\mathcal{X}$ and p.

Remark 17. *As in the statement of Theorem 8, the term $\widehat{\kappa}_p$ captures the average effect of noise on this problem, and as the noise variance goes to 0, this term becomes 1.*

B PROOFS

B.1 MEMORY AND ACCURACY PROOFS

First, we prove Lemma 3 which bounds the memory necessary to build, store, or search a `Bandit Cover Tree`

Proof of Lemma 3. By the nesting invariance, once a point $p \in \mathcal{X}$ enters the cover tree at a level i, it is present in all lower levels. By the covering tree invariance, p has a unique parent node. When computing estimates of distances, a single empirical estimate of $d(\boldsymbol{q}, p)$ and confidence width is necessary to save for each point p. Hence, for each point $p \in \mathcal{X}$, at most 4 numbers are necessary to store. Therefore, the `Bandit-Cover-Tree` can be stored in $O(n)$ memory.

$\square$

Next we prove several Lemmas that guarantee that the `Noisy-Find-Nearest`, `Noisy-Insert`, and `Noisy-Remove` methods all succeed with high probability. We begin with search accuracy.

Proof of Lemma 5. `Noisy-Find-Nearest` is adapted from `Find-Nearest`, but uses `Identify-Cover` to identify individual cover sets at it descends the tree and `find-smallest-in-set` to return the smallest in the final set. By Theorem 2 of Beygelzimer et al. [2006], `Find-Nearest` succeeds with probability 1. Hence, we wish to show that the probability of `Identify-Cover` or `find-smallest-in-set` making an error in any call is at most δ.

By Theorem 18, `Identify-Cover` given a failure probability of δ/α succeeds with probability $1 - \delta/\alpha$. Though `Identify-Cover` may return a set that is slightly larger than $\{j \in Q : d(\boldsymbol{q}, \boldsymbol{x}_j) \le d(\boldsymbol{q}, Q) + 2^i\}$, the returned set always contains this set if `Identify-Cover` succeeds. This implies that no grandparent node of $\boldsymbol{x}_q$ is ever removed. Hence, this does not affect the proof of correctness.

By definition, $\alpha = \min\left\{\left\lceil \frac{\log(n)}{\log(1+1/c^2)} + 1 \right\rceil, n\right\} + 1$. By Lemma 24, $i_{\text{top}} - i_{\text{bottom}} \le \alpha - 1$. Hence, `Identify-Cover` is called at most $\alpha - 1$ times (once per level during traversal). `find-smallest-in-set` is run a single time with failure probability δ/α, and by Theorem 19 makes an error with probability at most δ/α. A union bound implies that the probability of an error in any call is at most δ, completing the proof. $\square$

Next, we show that `Noisy-Insert` succeeds with high probability.

Proof of Lemma 6. To show correctness, we begin by verifying that the thresholding bandit routine in lines $2-9$ does not fail in any recursive call of `Noisy-Insert`. Define the event

$$\mathcal{E} := \bigcap_{k \in [n]} \bigcap_{t=1}^{\infty} \left\{ |\hat{d}_{p,k}(t) - d_{p,k}| \le C_{\delta/n}(t) \right\}$$

By definition of $C_\delta(t)$, we have that $P(\mathcal{E}^c) \le \delta$ where we have used the assumption that $|\mathcal{X}| = n$ so there are only n different points explicitly represented in $\mathcal{T}$. For the remainder of the proof, we assume that $\mathcal{E}$ occurs and show that it leads to correctness in the thresholding bandit routine.

Let $\mu > 0$ denote any threshold. On $\mathcal{E}$, we have that

$$\{k : d_{p,k} < \mu\} \supset \left\{k : \hat{d}_{p,k} + C_{\delta/n}(T_k(t)) < \mu\right\}$$

for any set of values $\{k : T_k(t)\}$ at any time t. Similarly, we have that

$$\{k : d_{p,k} > \mu\} \supset \left\{k : \hat{d}_{p,k} - C_{\delta/n}(T_k(t)) > \mu\right\}.$$

Therefore, if the algorithm stops sampling when either

$$\hat{d}_{p,k} + C_{\delta/n}(T_k(t)) < \mu$$

or

$$\hat{d}_{p,k} - C_{\delta/n}(T_k(t)) > \mu$$

for every point k and any value of μ, then on event $\mathcal{E}$, we have identified the sets $\{k : d_{p,k} < \mu\}$ and $\{k : d_{p,k} > \mu\}$ correctly. In particular, the above superset relation holds with equality.

Applying this to $\mu = 2^i$ for different values of i in the algorithm we see that at all times the thresholding bandit implemented in lines $2 - 9$ succeeds. Since these bounds are shared between recursive calls of the algorithm, the routine succeeds in every recursive call.

We conclude by showing that if these sets have been computed correctly, which occurs when $\mathcal{E}$ occurs, then the algorithm correctly computes all quantities needed for the `Insert` algorithm in Beygelzimer et al. [2006].

First, note that after the thresholding bandit terminates

$$\{\text{Threshold bandit has terminated}\} \cap \{j \in Q : \hat{d}_{p,j}(T_j) + C_{\delta/n}(T_j) \leq 2^i\} = \emptyset \implies d(p, Q) > 2^i.$$

Next, note that when the thresholding bandit routine terminates at line 8, we have that

$$\left\{k \in Q : \hat{d}_{p,k} + C_{\delta/n}(T_k(t)) \leq 2^i\right\} = \left\{k \in Q : d_{p,k} \leq 2^i\right\} = Q_{i-1}$$

where the last equality holds by definition in the `Insert` algorithm.

Finally, by the nesting invariance of cover trees and the definition of children in a tree, we have that $Q_i \subset Q$. Hence $Q_i \cap Q_{i-1} \neq \emptyset$ is equivalent to the condition that $d(p, Q_i) \leq 2^i$. This implies that we have computed (under uncertainty) the same quantities needed in for the `Insert` algorithm of Beygelzimer et al. [2006] for the NNS problem. Applying Theorem 3 therein completes the proof. $\square$

Finally, we show that `Noisy-Remove` succeeds with high probability.

Proof of Lemma 15. Similar to the proof for `Noisy-Insert`, we begin by showing correctness of the thresholding bandit in lines $9 - 16$. Again, we must verify that it does not fail in any recursive call of `Noisy-Remove`. Define the event

$$\mathcal{E} := \bigcap_{j,k \in [n]} \bigcap_{t=1}^{\infty} \left\{|\hat{d}_{j,k}(t) - d_{j,k}| \leq C_{\delta/n^2}(t)\right\}$$

By definition of $C_\delta(t)$, we have that $P(\mathcal{E}^c) \leq \delta$ where we have used the assumption that $|\mathcal{X}| = n$ so there are only n different points explicitly represented in $\mathcal{T}$ and there are at most $\binom{n}{2}$ pairs of distances. Note that especially for higher levels i of the tree it is possible that $|Q_i| < n$. Hence a weaker union bound is possible. As $|Q_i|$ is unknown a priori, we take the naive bound that $|Q_i| \leq n$ for any round i, though it is possible to instead alter the union bound in different recursive calls to `Noisy-Remove`.

For the remainder of the proof, we assume that $\mathcal{E}$ occurs and show that it leads to correctness in the thresholding bandit routine. Let $\mu > 0$ denote any threshold. On $\mathcal{E}$, we have that for any point q (in particular any child of node p)

$$\{k : d_{q,k} < \mu\} \supset \left\{k : \hat{d}_{q,k}(T_{q,k}(t)) + C_{\delta/n^2}(T_{q,k}(t)) < \mu\right\}$$

for any set of values $\{k : T_{q,k}(t)\}$ at any time t. Similarly, we have that

$$\{k : d_{p,k} > \mu\} \supset \left\{k : \hat{d}_{q,k}(T_{q,k}(t)) - C_{\delta/n^2}(T_{q,k}(t)) > \mu\right\}.$$

Therefore, if the algorithm stops sampling when either

$$\hat{d}_{q,k}(T_{q,k}(t)) + C_{\delta/n^2}(T_{q,k}(t)) < \mu$$

or

$$\hat{d}_{q,k}(T_{q,k}(t)) - C_{\delta/n^2}(T_{q,k}(t)) > \mu$$

for every point k and any value of μ, then on event $\mathcal{E}$, we have identified the sets $\{k : d_{p,k} < \mu\}$ and $\{k : d_{p,k} > \mu\}$ correctly. In particular, the above superset relation holds with equality.

Applying this to $\mu = 2^{i'}$ for different values of i' in the algorithm we see that at all times, the thresholding bandit implemented in lines $9 - 16$ succeeds. Since these bounds are shared between recursive calls of the algorithm, the routine succeeds in every recursive call.

We conclude by showing that if these sets have been computed correctly, which occurs when $\mathcal{E}$ occurs, then the algorithm correctly computes all quantities needed for the `Remove` algorithm in Beygelzimer et al. [2006].

First, note that when the thresholding bandit routine terminates

$$\{\text{Thresholding bandit terminated}\} \cap \{j \in Q : \hat{d}_{q,j}(T_j) + C_{\delta/n^2}(T_{q,j}) \le 2^{i'}\} = \emptyset \implies d(q, Q) > 2^{i'}.$$

Second, note that if the set

$$\{j \in Q_{i'} : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n^2}(T_j) \le 2^{i'}\}$$

is nonempty, then for any q' contained in it, we have that $d_{q,q'} \le 2^{i'}$. Therefore, on $\mathcal{E}$, `Noisy-Remove` computes the same quantities as `Remove`. Applying Theorem 4 of Beygelzimer et al. [2006] completes the proof. $\square$

B.2 QUERY TIME COMPLEXITY

Now we turn our attention to the number of calls to the distance oracle needed by `Bandit Cover Tree`. We begin by proving a bound on the number of oracle calls made by `Noisy-Find-Nearest`.

Proof of Theorem 8. Assume that `Noisy-Find-Nearest` succeeds, which occurs with probability $1 - \delta$. Let i_{top} and i_{bottom} represent the top and bottom of $\mathcal{T}$. We begin by bounding the number of oracle calls drawn in an arbitrary round. Calls to the oracle only occur within the `Identify-Cover` routine. Suppose `Identify-Cover` is called on set Q and run with failure probability δ/n since $\delta/\alpha \ge \delta/n$ deterministically. By Theorem 18 the number of calls to the distance oracle made by `Identify-Cover` is bounded by

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{j \in Q} \min\left\{\max\left\{\frac{1}{(d_{\min} + 2^{i-1} - d_{q,j})^2}, \frac{1}{(d_{q,j} + \kappa_i - d_{\min})^2}\right\}, 2^{-2i}\right\} \tag{3}$$

where c_1 includes constants and doubly logarithmic terms, $d_{\min} = \min_{j \in Q} d_{q,j}$, and $\kappa_i = \min|d_{\min} + 2^{i-1} - d_{q,j}|$.

Define κ_i^{avg} to be the arithmetic means of the summands in the above sum in Equation 3. Recall that for the cover set Q_i at level i, Q is defined as $Q = \bigcup_{p \in Q_i} \text{children}(p)$ in the `Noisy-Find-Nearest` algorithm. Hence, we may compactly bound Equation 3 as

$$c_1 \log\left(\frac{n}{\delta}\right) \left| \bigcup_{p \in Q_i} \text{children}(p) \right| \kappa_i^{\text{avg}}$$

since we have replaces a $\min$ by one of the terms inside.

Applying this, we may sum over all levels of the tree $\mathcal{T}$ and bound the total number of oracle calls as

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{i=i_{\text{top}}}^{i_{\text{bottom}}} \left| \bigcup_{p \in Q_i} \text{children}(p) \right| \kappa_i^{\text{avg}}$$

where we have written the outer sum index in order of *descending* i since $i_{\text{top}} > i_{\text{bottom}}$. This is done to reflect to the process of descending tree $\mathcal{T}$ and counting the number of oracle calls taken at each level. We proceed by bounding $\left| \bigcup_{p \in Q_i} \text{children}(p) \right|$.

For the set Q defined in the algorithm, define the set

$$\tilde{Q}_i := \{j \in Q : d(\boldsymbol{q}, \boldsymbol{x}_j) \leq d(\boldsymbol{q}, Q) + 2^i\}.$$

By design, Identify-Cover returns a set Q_i such that

$$\tilde{Q}_i = \{j \in Q : d(\boldsymbol{q}, \boldsymbol{x}_j) \leq d(\boldsymbol{q}, Q) + 2^i\} \subset Q_i \subset \{j \in Q : d(\boldsymbol{q}, \boldsymbol{x}_j) \leq d(\boldsymbol{q}, Q) + 2^{i+1}\} = \tilde{Q}_{i+1}.$$

By the expansion rate, $|\tilde{Q}_{i+1}| \leq c|\tilde{Q}_i|$ which implies that $|Q_i| \leq c|\tilde{Q}_i|$. By Lemma 23, the number of children of any node is bounded by c^4. Therefore, we have that

$$
\left| \bigcup_{p \in Q_i} \text{children}(p) \right| \leq \left| \bigcup_{p \in \tilde{Q}_{i-1}} \text{children}(p) \right|
$$

$$
= \left| \bigcup_{p \in \tilde{Q}_{i+1} \backslash \tilde{Q}_i} \text{children}(p) \cup \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right|
$$

$$
\overset{a}{\leq} c^4 |\tilde{Q}_{i+1} \backslash \tilde{Q}_i| \left| \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right|
$$

$$
\overset{b}{\leq} c^4 (c-1) |\tilde{Q}_i| \left| \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right|
$$

$$
\overset{c}{\leq} c^5 \left| \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right|^2
$$

Inequality a follows by the bound on the number of children of any node. Inequality b is a consequence of the expansion rate on $\mathcal{X} \cup \{\boldsymbol{q}\}$. To see this observe that we have that $|\tilde{Q}_{i+1}| \leq c|\tilde{Q}_i|$ and $|\tilde{Q}_{i+1}| = |\tilde{Q}_i| + |\tilde{Q}_{i+1} \backslash \tilde{Q}_i|$. Hence $|\tilde{Q}_{i+1} \backslash \tilde{Q}_i| \leq (c-1)|\tilde{Q}_i|$. Inequality c is a consequence of the nesting invariance which implies that $p \in \text{children}(p)$ for any node p in the cover tree.

Let Q^* be the final explicit Q. That is $Q^* = \tilde{Q}_{i^*}$ where i^* is the lowest level such that an explicit node exists in the cover set Q_{i^*}. For all levels i such that $i > i^*$ (levels above i^*), explicit nodes have yet to be added. Hence

$$
\left| \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right| \leq \left| \bigcup_{p \in \tilde{Q}_{i^*}} \text{children}(p) \right|.
$$

For all levels below i^*, by definition, no new nodes are added as all remaining nodes are implicit. Therefore, for $i < i^*$ (lower levels of the tree),

$$
\left| \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right| \leq \left| \bigcup_{p \in \tilde{Q}_{i^*}} \text{children}(p) \right|.
$$

In particular, Q_{i^*} maximizes $\left| \bigcup_{p \in \tilde{Q}_i} \text{children}(p) \right|$. By Lemma 22, we have that $\left| \bigcup_{p \in \tilde{Q}_{i^*}} \text{children}(p) \right| \leq c^5$. Next, for clarity, define

$$\overline{\kappa} := \frac{1}{i_{\text{top}} - i_{\text{bottom}}} \sum_{i=i_{\text{top}}}^{i_{\text{bottom}}} \kappa_i^{\text{avg}},$$

the average of the κ_i^{avg} terms that appears in each level. Plugging in both above pieces, we have that

$$
c_1 \log\left(\frac{n}{\delta}\right) \sum_{i=i_{\text{top}}}^{i_{\text{bottom}}} \left| \bigcup_{p \in Q_i} \text{children}(p) \right| \kappa_i^{\text{avg}} \leq c_1 c^5 \log\left(\frac{n}{\delta}\right) \sum_{i=i_{\text{top}}}^{i_{\text{bottom}}} \left| \bigcup_{p \in \tilde{Q}_{i^*}} \text{children}(p) \right|^2 \kappa_i^{\text{avg}} \leq c_1 c^{15} \log\left(\frac{n}{\delta}\right) (i_{\text{top}} - i_{\text{bottom}}) \overline{\kappa}
$$

By Lemma 24, we have that $i_{\text{top}} - i_{\text{bottom}} = O(c^2 \log(n))$. Plugging this in we have that the total number of oracle taken by `Identify-Cover` calls is bounded by

$$O\left(c^{17} \log(n) \log\left(\frac{n}{\delta}\right) \overline{\kappa}\right).$$

It remains to bound the number or oracle calls drawn by `find-smallest-in-set` to find x_q in the set $Q_{i_{\text{bottom}}}$. One the high probability event that all calls to `Identify-Cover` succeed, which occurs with probability at least $1 - \delta$, Lemma 5 implies that $x_q \in Q_{i_{\text{bottom}}}$. Hence, $d(q, Q_{i_{\text{bottom}}} = d(q, x_q)$. By Theorem 19, the total number of samples drawn by `find-smallest-in-set` is bounded by

$$c_2 \log\left(\frac{n}{\delta}\right) \sum_{j \in Q_{i_{\text{bottom}}} \setminus \{x_q\}} (d(q, x_q) - d(q, x_j))^{-2} = c_2 \log\left(\frac{n}{\delta}\right) |Q_{i_{\text{bottom}}}| \kappa' \leq c_2 \log\left(\frac{n}{\delta}\right) c^6 \kappa'$$

where κ' is the empirical average of $(d(q, x_q) - d(q, x_j))^{-2}$ for all $j \in \setminus\{x_q\}$ and we have used the fact that $|Q_{i_{\text{bottom}}}| \leq c|\tilde{Q}_{i^*}| \leq c^6$. Therefore, the total complexity of `Noisy-Find-Nearest` is

$$O\left(c^{17} \log(n) \log\left(\frac{n}{\delta}\right) \overline{\kappa} + \log\left(\frac{n}{\delta}\right) c^6 \kappa'\right). = O\left(c^{17} \log(n) \log\left(\frac{n}{\delta}\right) \kappa\right)$$

where $\kappa = \max(\overline{\kappa}, \kappa')$. To complete the proof, note that $\kappa' \leq ((d(q, x_q) - d(q, x_{q'}))^{-2}$ where $x_{q'}$ is the second closest point to q. To bound $\overline{\kappa}$, note that each summand in Equation (3) may be bounded instead as 2^{-2i} in level i. As κ_i^{avg} is an average of these terms, it too is bounded by 2^{-2i}. $\overline{\kappa}$ is the average of κ_i^{avg} in each round. As 2^{-2i} monotonically decreases in i, we have that $\overline{\kappa} \leq \max_i \kappa_i^{\text{avg}} \leq 2^{-2i_{\text{bottom}}}$. Hence, by Lemma 21, the bottom level is $i = O(\log_2(d_{\min}))$ where $d_{\min} := \min_{j,k \in \mathcal{X}} d(x_j, x_k)$. Hence We have that $2^{-2i_{\text{bottom}}}$ is $O(d_{\min}^{-2})$. Therefore, $\kappa \leq O\left(\max\{[d(q, x_q) - d(q, x_{q'})]^{-2}, d_{\min}^{-2}\}\right)$. $\square$

Proof of Corollary 10. The proof follows identically, except that we make use $\tilde{c}$ as a bound on c throughout and may plug in

$$\alpha = \left\lceil \frac{\log(n)}{\log(1 + 1/\tilde{c}^2)} + 1 \right\rceil + 1 = O(\tilde{c}^2 \log(n)).$$

Hence the term from the union bound becomes $O\left(\log\left(\frac{\tilde{c}^2 \log(n)}{\delta}\right)\right)$ instead of $\log\left(\frac{n}{\delta}\right)$. $\square$

B.3 INSERTION TIME COMPLEXITY

Next, we bound the number of oracle calls made by `Noisy-Insert`.

Proof of Theorem 11. We begin by analyzing the complexity of the thresholding bandit subroutine in lines $2 - 9$ of `Noisy-Insert`. Assume the same event $\mathcal{E}$ from the proof of Lemma 6 that holds with probability $1 - \delta$. We proceed by bounding the number of rounds any point j in the set Q may be $i^*(t)$ before it must enter the set K. Summing up the complexity for all points in Q bounds the complexity of this routine.

Suppose we wish to insert p into level i. Assume for point $j \in C_i$ that $d_{p,j} \leq 2^i$. Assume that

$$T_j \geq \frac{c'}{(d_{p,j} - 2^i)^2} \log\left(\frac{n}{\delta} \log\left(\frac{n}{\delta(d_{p,j} - 2^i)^2}\right)\right)$$

for a sufficiently large constant c'. Then

$$\hat{d}_{p,j}(T_j) + C_{\delta/n}(T_j) \overset{\mathcal{E}}{\leq} d_{p,j} + 2C_{\delta/n}(T_j) \overset{(a)}{<} d_{p,j} + 2^i - d_{p,j} = 2^i,$$

implying that j must be in the set K. Inequality (a) follows from Lemma 20. This argument may be repeated for $j \in C_i$ such that $d_{p,j} > 2^i$ by instead considering the lower confidence bound on $\hat{d}_{p,j}(T_j)$.

Summing over all j in the set Q, no more than

$$\sum_{j \in Q} \frac{c'}{(d_{p,j} - 2^i)^2} \log\left(\frac{n}{\delta} \log\left(\frac{n}{\delta(d_{p,j} - 2^i)^2}\right)\right) \leq c_1 |Q| \log\left(\frac{n}{\delta}\right) \kappa_i^{\text{avg}}$$

calls to the oracle are made between lines 2 and 8 for a problem independent constant c_1. Similar to the proof of Theorem 8, we define κ_i^{avg} to be average of the summands including doubly logarithmic terms for brevity and clarity.

By definition, in level i $Q = \bigcup_{j \in Q_i} \mathrm{children}(j)$. In the worst case, p is added at the leaf level, i_{bottom}. Noisy-Insert descends down the tree via recursive calls. This happens at most $i_{\mathrm{top}} - t_{\mathrm{bottom}}$ times. Summing over all levels, the total number of calls to the oracle is bounded by

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{i=i_{\mathrm{top}}}^{i_{\mathrm{bottom}}} \left| \bigcup_{j \in Q_i} \mathrm{children}(j) \right| \kappa_i^{\mathrm{avg}}$$

where the sum is indexed from the largest i to the smallest to reflect descending the tree. As in the proof of Theorem 8, there exists a level i^* which maximizes $\left| \bigcup_{j \in Q_i} \mathrm{children}(j) \right|$ and we have that $\left| \bigcup_{j \in Q_{i^*}} \mathrm{children}(j) \right| \le c^5$. Define

$$\overline{\kappa}_p = \frac{1}{i_{\mathrm{top}} - i_{\mathrm{bottom}}} \sum_{i=i_{\mathrm{top}}}^{i_{\mathrm{bottom}}} \kappa_i^{\mathrm{avg}}$$

Plugging this in,

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{i=i_{\mathrm{top}}}^{i_{\mathrm{bottom}}} \left| \bigcup_{j \in Q_i} \mathrm{children}(j) \right| \kappa_i^{\mathrm{avg}} \le c_1 c^5 \log\left(\frac{n}{\delta}\right) (i_{\mathrm{top}} - i_{\mathrm{bottom}})\overline{\kappa}_p$$

$$\le O\left(c^7 \log(n) \log\left(\frac{n}{\delta}\right) \overline{\kappa}_p\right).$$

where the final inequality follows from Lemma 24. $\qquad\square$

Proof of Theorem 13. Begin with an empty tree. Noisy-Insert can trivially place x_1 as a root and replace the root as necessary. We run Noisy-Insert with failure probability δ/n. Since placing the first node at the root is trivial, make the inductive hypothesis that we have a correct tree $\mathcal{T}$ built on a strict subset $\mathcal{S} \subsetneq \mathcal{X}$ and wish to insert a point $p \in \mathcal{X}\backslash\mathcal{S}$ to $\mathcal{T}$ using Noisy-Insert. By Lemma 6, this process succeeds with probability $1 - \delta/n$. By Theorem 11, this requires no more than

$$O\left(c^7 \log(n) \log\left(\frac{n}{\delta}\right) \overline{\kappa}_p\right)$$

calls to the noisy distance oracle.

A union bound over inserting the n points in $\mathcal{X}$ implies correctness. Summing the above expression for every point $p \in \mathcal{X}$ bounds the total sample complexity necessary for construction, stated in the Theorem. In particular, $\widetilde{\kappa}$ is the arithmetic mean of the individual $\overline{\kappa}_p$s. $\qquad\square$

B.4 REMOVAL TIME COMPLEXITY

Finally, we bound the number of oracle calls needed by Noisy-Remove

Proof of Theorem 16. We begin by analyzing the complexity of the Thresholding bandit subroutine in lines $9 - 16$ of Noisy-Remove. Assume the same event $\mathcal{E}$ from the proof of Lemma 15 that holds with probability $1 - \delta$. Fix an arbitrary $q \in \mathrm{children}(p)$. We proceed by bounding the number of rounds any point j in the set $Q_{i'}$ may be $i^*(t)$ before it must enter the set K. Summing up the complexity for all points in Q bounds the complexity of this routine.

Assume for point $j \in C_{i'}$ that $d_{q,j} \le 2^{i'}$. Assume that

$$T_{q,j} \ge \frac{c'}{(d_{q,j} - 2^{i'})^2} \log\left(\frac{n^2}{\delta} \log\left(\frac{n^2}{\delta(d_{q,j} - 2^{i'})^2}\right)\right)$$

for a sufficiently large constant c'. Then

$$\hat{d}_{q,j}(T_{q,j}) + C_{\delta/n^2}(T_{q,j}) \overset{\mathcal{E}}{\le} d_{q,j} + 2C_{\delta/n^2}(T_{q,j}) \overset{(a)}{<} d_{q,j} + 2^{i'} - d_{p,j} = 2^{i'},$$

implying that j must be in the set K. Inequality (a) follows from Lemma 20. This argument may be repeated for $j \in C_i$ such that $d_{q,j} > 2^{i'}$ by instead considering the lower confidence bound on $\hat{d}_{q,j}(T_{q,j})$.

Summing over all j in the set $Q_{i'}$, no more than

$$\sum_{j \in Q_{i'}} \frac{c'}{(d_{q,j} - 2^{i'})^2} \log\left(\frac{n^2}{\delta} \log\left(\frac{n^2}{\delta(d_{q,j} - 2^{i'})^2}\right)\right) \leq c_1 |Q_{i'}| \log\left(\frac{n^2}{\delta}\right) \kappa_{q,i'}^{\text{avg}}$$

calls to the oracle are made between lines 9 and 16 for a problem independent constant c_1. Similar to the proof of Theorem 8, we define $\kappa_{q,i'}^{\text{avg}}$ to be average of the summands including doubly logarithmic terms for brevity and clarity.

If

$$\{j \in Q_{i'} : \hat{d}_{q,j}(T_{q,j}) + C_{\delta/n^2}(T_{q,j}) \leq 2^{i'}\} = \emptyset,$$

i' is reset to $i' + 1$ and the algorithm proceeds to the next level up the tree. An identical computation is performed as in lines $9 - 16$ and a similar bound applies as the above. The only difference is that since the Thresholding bandit is now comparing to a threshold of $2^{i'+1}$ (or alternatively $2^{i'}$ for the incremented value of i') we incur a dependence on $\kappa_{q,i'+1}^{\text{avg}}$ instead. In particular, the number of oracle queries drawn between lines 20 and 25 is at most

$$\sum_{j \in Q_{i'+1}} \frac{c'}{(d_{q,j} - 2^{i'+1})^2} \log\left(\frac{n^2}{\delta} \log\left(\frac{n^2}{\delta(d_{q,j} - 2^{i'+1})^2}\right)\right) \leq c_1 |Q_{i'+1}| \log\left(\frac{n^2}{\delta}\right) \kappa_{q,i'+1}^{\text{avg}}.$$

This process repeats until the conditional in the while loop is no longer satisfied. Naively, this happens at most $i_{\text{top}} - i + 1$ times as i' is initialized as $i - 1$ and is incremented until potentially it reaches the top level of $\mathcal{T}$, i_{top}. Summing this quantity over all levels of the tree, we may bound the total number of oracle calls as

$$\sum_{i'=i_{\text{top}}}^{i-1} c_1 |Q_{i'}| \log\left(\frac{n^2}{\delta}\right) \kappa_{q,i'}^{\text{avg}}.$$

As in the proof of Theorem 8, there exists a level i^* which maximizes $\left|\bigcup_{j \in Q_i} \text{children}(j)\right|$ and we have that $\left|\bigcup_{j \in Q_{i^*}} \text{children}(j)\right| \leq c^5$. As $Q_i \subset \bigcup_{j \in Q_i} \text{children}(j)$ by the nesting invariance, c^5 likewise bounds Q_i for all i. Define

$$\kappa_q^{\text{avg}}(i) = \frac{1}{i_{\text{top}} - i + 1} \sum_{i=i_{\text{top}}}^{i-1} \kappa_{q,i}^{\text{avg}}$$

Plugging this in, we may bound the total number of oracle calls (for the distance of any point to q) by

$$\sum_{i'=i_{\text{top}}}^{i-1} c_1 |Q_{i'}| \log\left(\frac{n^2}{\delta}\right) \kappa_{q,i'}^{\text{avg}} \leq c_1 c^5 (i_{\text{top}} - i + 1) \log\left(\frac{n^2}{\delta}\right) \kappa_q^{\text{avg}}(i)$$

$$\leq c_1 c^5 (i_{\text{top}} - i_{\text{bottom}}) \log\left(\frac{n^2}{\delta}\right) \kappa_q^{\text{avg}}(i).$$

Next, by Lemma 24, $i_{\text{top}} - i_{\text{bottom}} = O(c^2 \log(n))$. Therefore,

$$c_1 c^5 (i_{\text{top}} - i_{\text{bottom}}) \log\left(\frac{n^2}{\delta}\right) \kappa_q^{\text{avg}}(i) \leq O\left(c^7 \log(n) \log\left(\frac{n^2}{\delta}\right) \kappa_q^{\text{avg}}(i)\right)$$

The above bounds the number of calls to the distance oracle needed for any child q of p, the point to be removed. Due to the 'For' loop in line 6, this process is repeated for all $q \in \text{children}(p)$. By Lemma 23, the number of children of any node $p \in \mathcal{T}$ is at most c^4. Define

$$\kappa^p(i) := \frac{1}{|\text{children}(p)|} \sum_{q \in \text{children}(p)} \kappa_q^{\text{avg}}(i)$$

where the superscript p and the parenthetical i denote that this quantity depends on all children of p and level i of the tree. Then

$$\sum_{q \in \text{children}(p)} \kappa_q^{\text{avg}}(i) = |\text{children}(p)| \frac{1}{|\text{children}(p)|} \sum_{q \in \text{children}(p)} \kappa_q^{\text{avg}}(i)$$

$$\leq c^4 \frac{1}{|\text{children}(p)|} \sum_{q \in \text{children}(p)} \kappa_q^{\text{avg}}(i)$$

$$= c^4 \kappa^p(i)$$

Therefore, summing over all $q \in \text{children}(p)$, we can bound the total number of calls to the distance oracle drawn in lines 6 to 29 of `Noisy-Remove` by

$$O\left(c^9 \log(n) \log\left(\frac{n^2}{\delta}\right) \kappa^p(i)\right).$$

As `Noisy-Remove` is recursive, it remains to sum the complexity of all recursive calls. The above bound depends on the level i on which `Noisy-Remove` is called only through the term $\kappa^p(i)$. In the worst case, p is present in every level and the 'If' condition in line 4 is true for every recursive call. Hence, we sum the above expression over every level i. Define

$$\widehat{\kappa}_p := \frac{1}{i_{\text{top}} - i_{\text{bottom}}} \sum_{i=i_{\text{top}}}^{i_{\text{bottom}}} \kappa^p(i).$$

The total number of oracle calls is bounded as

$$\sum_{i=i_{\text{top}}}^{\text{bottom}} O\left(c^9 \log(n) \log\left(\frac{n^2}{\delta}\right) \kappa^p(i)\right) = O\left(c^9 (i_{\text{top}} - i_{\text{bottom}}) \log(n) \log\left(\frac{n^2}{\delta}\right) \widehat{\kappa}_p\right)$$

$$= O\left(c^{11} \log^2(n) \log\left(\frac{n^2}{\delta}\right) \widehat{\kappa}_p\right)$$

completing the proof. $\qquad\square$

C EXTERNAL RESULTS USED IN PROOFS

C.1 RESULTS ABOUT BANDIT SUBROUTINES

Theorem 18 (Based on Theorem 3.1 of Mason et al. [2020])**.** *Fix $\delta < 1/2$, a query point q, a set of points Q, and $i \in \mathbb{Z}$. With probability at least $1 - \delta$, `Identify-Cover` returns a set Q_i such that*

$$\{j \in Q : d_{q,j} \leq d(\boldsymbol{q}, Q) + 2^{i-1}\} \subset Q_i \subset \{j \in Q : d_{q,j} \leq d(\boldsymbol{q}, Q) + 2^{i-1} + 2^{i-2}\}$$

in at most calls to the distance oracle made by `Identify-Cover` is bounded by

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{j \in Q} \min\left\{\max\left\{\frac{1}{(d(\boldsymbol{q}, Q) + 2^{i-1} - d_{q,j})^2}, \frac{1}{(d_{q,j} + \kappa_i - d(\boldsymbol{q}, Q))^2}\right\}, 2^{-2i}\right\} \qquad (4)$$

where c_1 includes constants and doubly logarithmic terms, and $\kappa_i = \min|d_{\min} + 2^{i-1} - d_{q,j}|$.

Proof. This result follows by noticing that `Identify-Cover` is equivalent to $(\text{ST})^2$ with parameters $\epsilon = 2^{i-1}$ and $\gamma = 2^{i-2}$ from Mason et al. [2020] if all distance are multiplied by -1. Hence, Theorem 3.1 of Mason et al. [2020] implies the result. $\qquad\square$

Next we give a bound on the complexity of `find-smallest-in-set` which is equivalent to the bandit algorithm `Successive-Elimination` given the negatives of the distances.

Theorem 19 (Even-Dar et al. [2002], Theorem 2)**.** *Fix $\delta > 0$, a query point $\boldsymbol{q}$, and a set Q. Let $\boldsymbol{x}^* \in Q = \arg\min_{\boldsymbol{x} \in Q} d(\boldsymbol{q}, \boldsymbol{x})$. `find-smallest-in-set` returns $\boldsymbol{x}^*$ with probability at least $1 - \delta$ in at most*

$$c_2 \log\left(\frac{n}{\delta}\right) \sum_{j \in Q_{i_{bottom}} \setminus \{\boldsymbol{x}^*\}} (d(\boldsymbol{q}, \boldsymbol{x}_q) - d(\boldsymbol{q}, \boldsymbol{x}_j))^{-2}$$

calls to the distance oracle.

Lemma 20 (Mason et al. [2020], Lemma F.2)**.** *For $\delta < 2e^{-e/2}$, $\Delta \leq 2$,*

$$t \geq \frac{4}{\Delta^2} \log\left(\frac{2}{\delta} \log_2\left(\frac{12}{\delta\Delta^2}\right)\right) \implies C_\delta(t) = \sqrt{\frac{4\log(\log_2(2t)/\delta)}{t}} \leq \Delta.$$

C.2 RESULTS ABOUT COVER TREES

First we state a result from Krauthgamer and Lee [2004]. Though this result pertains to the `Navigating-Nets` data structure, as the cover tree is based on the same set of invarainces, the same result applies to cover trees.

Lemma 21 (Krauthgamer and Lee [2004], Lemma 2.3). *The $i_{top} = O(\log_2(d_{\max})$ and $i_{bottom} = O(\log_2(d_{\min}))$ where $d_{\max} := \max_{i,j \in \mathcal{X}} d_{i,j}$ and $d_{\min} := \min_{i,j \in \mathcal{X}} d_{i,j}$.*

The remained of the statements are from Beygelzimer et al. [2006] for the NNS problem. The core idea of many of the above results is that by managing the uncertainty in the NNSU problem, we can use results from NNS.

Lemma 22 (Beygelzimer et al. [2006], proof of Theorem 5). *For expansion constant c, the set $\tilde{Q}_i := \{j \in Q : d_{\mathbf{q},j} \leq d(\mathbf{q}, Q) + 2^i\}$ as defined in the proof of* `Noisy-Find-Nearest` *obeys $\max_i |\tilde{Q}_i| \leq c^5$.*

Lemma 23 (Beygelzimer et al. [2006], Lemma 4.1). *The number of children of any node of a cover tree is bounded by c^4.*

Lemma 24 (Beygelzimer et al. [2006]). *The maximum explicit depth of any node in a cover tree is $\frac{\log(n)}{\log(1+1/c^2)} = O(c^2 \log(n))$. Hence $i_{top} - i_{bottom} = O(c^2 \log(n))$.*