Paired compressed cover trees guarantee a near linear parametrized complexity for all k-nearest neighbors search in an arbitrary metric space

Yury Elkin
Materials Innovation Factory and Computer Science department, University of Liverpool, UK
yura.elkin@gmail.com

Vitaliy Kurlin
Materials Innovation Factory and Computer Science department, University of Liverpool, UK
vitaliy.kurlin@gmail.com

─── **Abstract** ───

This paper studies the important problem of finding all k-nearest neighbors to points of a query set Q in another reference set R within any metric space. Our previous work defined compressed cover trees and corrected the key arguments in several past papers for challenging datasets. In 2009 Ram, Lee, March, and Gray attempted to improve the time complexity by using pairs of cover trees on the query and reference sets. In 2015 Curtin with the above co-authors used extra parameters to finally prove a time complexity for $k = 1$. The current work fills all previous gaps and improves the nearest neighbor search based on pairs of new compressed cover trees. The novel imbalance parameter of paired trees allowed us to prove a better time complexity for any number of neighbors $k \geq 1$.

1 All k-nearest neighbors problem and overview of new results

The nearest neighbor search was one of the first data-driven problems [7]. Briefly, the problem is to find all $k \geq 1$ nearest neighbors in a reference set R for all points from a query set Q. Both sets live in an ambient space X with a distance d satisfying all metric axioms. For example, $X = \mathbb{R}^m$ with the Euclidean metric, a query set Q can be a single point.

The naive approach to find first nearest neighbors of points from Q within R is proportional to the product $|Q| \cdot |R|$ of the sizes of Q, R. The aim in nearest neighbor search is to reduce the brute-force complexity $O(|Q| \cdot |R|)$ to a near linear time in the maximum size of Q, R.

The *exact* k-nearest neighbor problem asks for exact (true) k-nearest neighbors of every query point q. The *probabilistic* neighbor problem [14] aims to find exact k-nearest neighbors with a given probability. The *approximate* neighbor problem [2], [13],[1],[17] aims to find for every query point $q \in Q$ its approximate neighbor $r \in R$ satisfying $d(q, r) \leq (1+\epsilon)d(q, \text{NN}(q))$, where $\epsilon > 0$ is fixed and $\text{NN}(q)$ is the exact first nearest neighbor of q.

Arguably the first approach to the neighbor search was developed in 1974 by using a *quadtree* [11], which hierarchically indexed a reference set $R \subset \mathbb{R}^2$. In higher dimensions, a KD-tree [3] smarter subdivided a subset of $R \subset \mathbb{R}^m$ at every level into two subsets instead of 2^m subsets. [10, Section 1] described further developments since 1970s in detail. Section 2 will technically compare the most recent results after formalizing the key concepts below.

We formally define the k-*nearest neighbor set* $\text{NN}_k(q)$, which might contain several points in a singular case when these points are at the exactly same distance from a query point q.
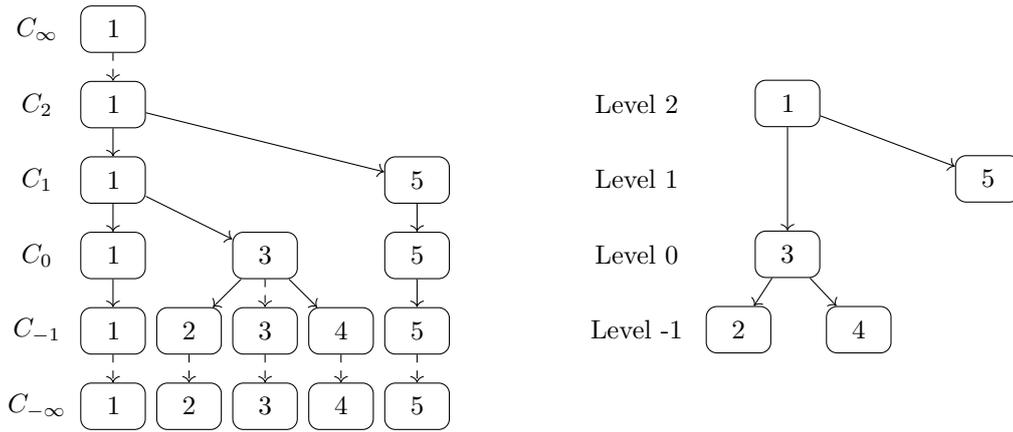
■ **Figure 1 Left:** an implicit cover tree for the finite set of reference points $R = \{1, 2, 3, 4, 5\}$ with the Euclidean distance $d(x, y) = |x - y|$, initially introduced in [5, Section 2]. **Right:** a new compressed cover tree in Definition 3.1 corrects the past complexity results for neighbor search.

▶ **Definition 1.1** (k-nearest neighbor set NN_k)**.** Let Q, R be finite subsets of a metric space. For any points $q \in Q$ and $r \in R$, the *neighborhood* $N(q; u) = \{p \in R \mid d(q, p) \leq d(q, u)\}$ consists of all points that are non-strictly closer to q than u is close to q. For any integer $k \geq 1$, the *k-nearest neighbor set* $\mathrm{NN}_k(q)$ consists of all points $u \in R$ such that $|N(q; u)| \geq k$ and any other $v \in R$ with $d(v, q) > d(u, q)$ has a larger neighborhood of size $|N(q; v)| > k$. ■

For $Q = R = \{0, 1, 2, 3\}$, the nearest neighbor sets of 0 are $\mathrm{NN}_1(0) = \{1\}$, $\mathrm{NN}_2(0) = \{2\}$, $\mathrm{NN}_3(0) = \{3\}$. Due to the neighborhoods $N(1; 0) = \{0, 1, 2\} = N(1; 2)$, both sets $\mathrm{NN}_1(1) = \{0, 2\} = \mathrm{NN}_2(1)$ consist of two points $0, 2$ at equal distance to 1. The 3-nearest neighbor set $\mathrm{NN}_3(1) = \{3\}$ is a single point. Because of a potential ambiguity of an exact k-nearest neighbor, Problem 1.2 below allows any neighbors within a set $\mathrm{NN}_k(q)$. In the above example, 0 can be chosen as a 1st neighbor of 1, then 2 as a 2nd neighbor of 1, or vice versa.

▶ **Problem 1.2** (all k-nearest neighbors search)**.** For any finite sets Q, R in a metric space and for any integer $k \geq 1$, design an algorithm to exactly find a distinct points $p_i \in \mathrm{NN}_i(q) \subseteq R$ for all $i = 1, ..., k$ and all points $q \in Q$ such that the total complexity is near linear in $n = \max\{|Q|, |R|\}$ with hidden constants depending on structures of the sets Q, R. ■

We solve Problem 1.2 by using new compressed cover trees on both sets Q, R. [5, Section 2] introduced a first version of a cover tree, which implicitly repeats every data point at infinitely many levels, see a comparison of two trees on the same R in Fig. 1. The new tree $\mathcal{T}(R)$ in the right hand side of Fig. 1 has nodes at levels $-1, 0, 1, 2$, so its height is $|H(\mathcal{T}(R))| = 4$.

The height parameter $|H(\mathcal{T}(R))|$ will play an important role to correctly estimate the time complexity, see Definition 3.10. The new concept of the imbalance $I(\mathcal{T}(R), \mathcal{T}(Q))$ in Definition 5.1 is the final parameter needed to express the search time complexities.

Section 2 discusses challenging cases that were overlooked in the past. Sections 3 and 4 introduce key concepts and prove auxiliary results. Section 5 estimates the time complexity for a traversal of paired trees in Theorem 5.7. Section 6 describes Algorithm 6.3 whose correctness and complexity are proved in Theorems 6.5, 6.8 solving Problem 1.2 for $k \geq 1$.

2   Review of past work and challenges in the k-nearest neighbor search

First we review the past searches by using a single tree on a reference set R. Then we consider the faster searchers by using a paired trees (a dual tree approach) on both Q, R.

Search on a single tree. The highly-cited paper [5] by Beygelzimer, Kakade, Langford in 2006 introduced the first concept of a cover tree to prove a worst-case time complexity for a 1st nearest neighbor search ($k = 1$). This tree was motivated by the earlier navigating nets [13] and is called an *implicit cover tree* to distinguish other versions in Fig. 1 and 4.

[5, Theorem 6] implied that an implicit cover tree T on a reference set R can be built in time $O(c^6|R|\log|R|)$, where $|R|$ is the size (number of points) of R, c is an expansion constant depending a structure of R, see Definition 3.4. Assuming that T is already built, [5, Theorem 5] expressed the time complexity to find a 1st nearest neighbor in R for all points $q \in Q$ as $O(c^{12}|Q|\log|R|)$. However, both statements relied on [5, Lemma 4.3] showing that a depth of a single point in a tree is $O(c^2\log|R|)$. In 2015 [9, Section 5.3] pointed out [5, Lemma 4.3] cannot be used to estimate the number of internal recursions in [5, Theorem 5]. In [10, section 3] it was shown that the original time complexity theorem for the cover tree construction algorithm [5, Theorem 6] had similar issues. Otherwise, the past argument guarantees the total complexity $O(c^{12}|R|^2)$ for $Q = R$ as for a brute-force search for $k = 1$. In 2021 [10, section 3] described typical sets R that indeed require more careful estimates.

The above facts motivated a new concept of a compressed cover tree to clarify the worst-case complexity for all reference sets R, see Definition 3.1, which [10, Theorem 6.5] showed that, after constructing a compressed cover tree $\mathcal{T}(R)$ on R, we can find all k-nearest neighbors of q in R in time $O(c^{10}k\log(k) \cdot |H(\mathcal{T}(R))|)$ for any $k \geq 1$, where $|H(\mathcal{T}(R))|$ is the height of $\mathcal{T}(R)$, see Definition 3.10. The new height parameter $|H(\mathcal{T}(R))|$ has the upper bound $\log_2(\Delta)$, where the aspect ratio Δ is the diameter of R divided by the minimum distance between points in R. Another approach uses paired trees on both sets Q, R.

Search on paired trees. In 2001 Gray and Moore [12] considered two KD-trees on the reference and query sets Q, R. The approach was often called a dual-tree search. Since the trees on Q, R are independent, there is no real duality between them, so we call such trees *paired*. In 2006 [4] proposed to consider paired cover trees. [4, Theorem 3.7] claimed that [4, Algorithm 5] finds all 1st nearest-neighbors for a query set Q in time $O(c^{16}\max\{|Q|, |R|\})$. In the case $Q = R$, [4, Algorithm 5] seems to return only the self-neighbor q for any point $q \in R$. A sketched proof of [4, Theorem 3.7] claimed similarities with the argument in [5, Theorem 5] without details that could clarify the concerns in the single tree case.

In 2009 [16, Theorem 3.1] revisited the time complexity for all 1st nearest neighbors and claimed the upper bound $O(c(R)^{12}c(Q)^{4\kappa}\max\{|Q|, |R|\})$, where $c(Q), c(R)$ are expansion constants of the query set Q and reference set R. The degree of bichromaticity κ is a parameter of both sets Q, R, see [16, Definition 3.1]. We have found the following issues.

First, Counterexample A.5 shows that [16, Algorithm 1] for $Q = R$ returns for any query point $q \in Q$ the same point q as its first neighbor. Second, Remark A.7 explains several possible interpretations of [16, Definition 3.1] for the parameter κ. Third, [16, Theorem 3.1] similarly to [5, Theorem 5] relied on the same estimate of recursions in the proof of [5, Lemma 4.3]. Counterexample A.8 explains step-by-step why the time complexity of [16, Algorithm 1] requires a clearer definition of κ and more detailed justifications.

In 2015 Curtin with the authors above [8] introduced other parameters: the imbalance I_t in [8, Definition 3] and θ in [8, Definition 4]. These parameters measured extra recursions that occurred due to possible imbalances in trees built on Q, R, which was missed in the past. [8, Theorem 2] shows that, for constructed cover trees on a query set Q and a reference set R, Problem 1.2 for $k = 1$ (only 1st nearest neighbors) can be solved in time

$$O\Big(c^{O(1)}\big(|R| + |Q| + I_t + \theta\big)\Big). \tag{$*$}$$

The main novelty of this work is Theorem 6.8 proving a better time complexity in Problem 1.2 for any $k \geq 1$ (all k-nearest neighbors), which is a bit simplified here as

$$O\Big(c^{O(1)}k\log(k) \cdot \big(|H(\mathcal{T}(R))| + |Q| + I(\mathcal{T}(R), \mathcal{T}(Q))\big)\Big). \tag{$**$}$$

The new single imbalance I of both compressed cover trees on Q, R, roughly combines the above parameters I_t and θ from [8], see Definition 5.1. The height $|H(\mathcal{T}(R))|$ is bounded above by $\log_2(\Delta)$, where Δ is the aspect ratio of the reference set R, see Definition 3.11.

In practice, we often need to quickly find nearest neighbors for a small number of data objects within a much larger dataset. So the reference set R is usually larger than a query set Q, which might be a single point. Hence $(**)$ improves the previous time complexity in $(*)$ from [8, Theorem 2] by replacing the large term $|R|$ by the height $|H(\mathcal{T}(R))|$.

Applying the upper bound $I(\mathcal{T}(R), \mathcal{T}(Q)) \leq |Q| \cdot |H(\mathcal{T}(R))|$ from Lemma 5.2, the complexity in $(**)$ reduces to [10, Theorem 6.5]. Hence the new imbalance $I(\mathcal{T}(R), \mathcal{T}(Q))$ for paired trees improves the recent approach [10] to Problem 1.2 based on a single tree.

Finally, in the case of a full batch search for $Q = R$, Lemma 5.4 shows that if $\mathcal{T}(R)$ is a balanced tree, then $I(\mathcal{T}(R), \mathcal{T}(R)) = O(|R|) = O(|Q|)$. Hence in all cases the new parameterized complexity is near linear in the maximum size $\max\{|R|, |Q|\}$ and any $k \geq 1$.
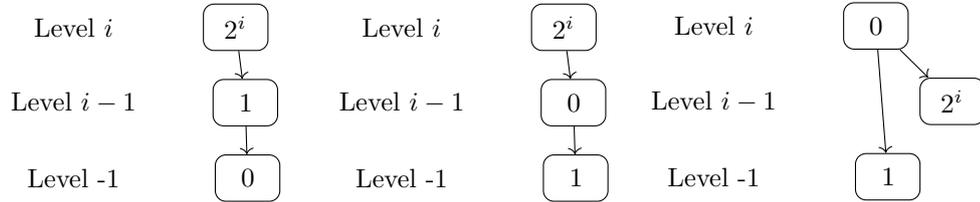
Figure 2 For any integer $i \geq 2$, the reference set $R = \{0, 1, 2^i\}$ with Euclidean metric has at least three compressed cover trees $\mathcal{T}(R)$ satisfying all conditions of Definition 3.1.

3   A compressed cover tree and other tools for k-nearest neighbors

▶ **Definition 3.1** (a compressed cover tree $\mathcal{T}(R)$, [10, Definition 2.1]). Let R be a finite set in an ambient space X with a metric d. A *compressed cover tree* $\mathcal{T}(R)$ has the vertex set R with a root $r \in R$ and a *level* function $l : R \to \mathbb{Z}$ satisfying the conditions below.

(3.1a) *Root condition* : the level of a root node r is $l(r) \geq 1 + \max_{p \in R - r} l(p)$.

(3.1b) *Covering condition* : every non-root node $q \in R$ has a unique *parent* p and a level $l(q) < l(p)$ so that $d(q, p) \leq 2^{l(q)+1}$; then p has a single link to its *child* node q in $\mathcal{T}(R)$.

(3.1c) *Separation condition* : for any $i \in \mathbb{Z}$, the *cover set* $C_i = \{p \in R \mid l(p) \geq i\}$ has a minimum separation distance $d_{\min}(C_i) = \min\limits_{p \in C_i} \min\limits_{q \in C_i - p} d(p, q) > 2^i$.

Since there is a 1-1 correspondence between all reference points in R and all nodes of the tree $\mathcal{T}(R)$, the same notation p can refer to a point of the reference set R or to a node of a tree $\mathcal{T}(R)$. Set $l_{\max} = 1 + \max_{p \in R-r} l(p)$ and $l_{\min} = \min_{p \in R} l(p)$. For a node $p \in \mathcal{T}(R)$, let $\text{Children}(p)$ denote the set of all children of p, including p itself, which will be convenient later.

For any node $p \in \mathcal{T}(R)$, define the *node-to-root* path as a unique sequence of nodes $w_0, \ldots, w_m$ such that $w_0 = p$, w_m is the root and w_{j+1} is the parent of w_j for all $j = 0, \ldots, m-1$. A node $q \in \mathcal{T}(R)$ is called a *descendant* of another node p if p belongs to the node-to-root path of q. A node p is an *ancestor* of q if q belongs to the node-to-root path of p. The set of all descendants of a node p is denoted by $\text{Descendants}(p)$ and includes the node p itself. ∎
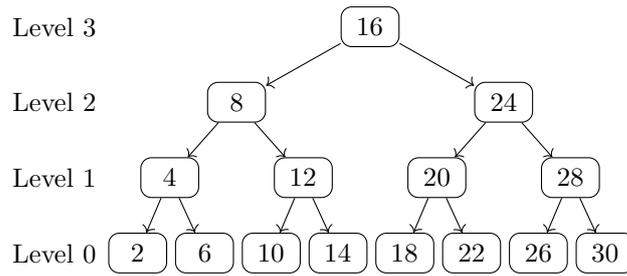
Figure 3 Compressed cover tree $\mathcal{T}(R)$ built on set R defined in Example 3.2 with root 16.

▶ **Example 3.2** (tree $\mathcal{T}(R)$ in Fig. 3). Let $\mathbb{R}$ be the real line with Euclidean metric $d = |x-y|$. Set $R = \{2, 4, 6, \ldots, 30\} \subset \mathbb{R}$. Figure 3 shows a compressed cover tree on the set R with root 16. The cover sets of $\mathcal{T}(R)$ are $C_0 = \{2, 4, 6, \ldots, 30\}$, $C_1 = \{4, 8, 12, 16, 20, 24, 28\}$, $C_2 = \{8, 16, 24\}$ and $C_3 = \{16\}$. We check the conditions of Definition 3.1 below.

- Root condition (3.1a): $\max_{p \in R \setminus \{16\}} d(p, 16) = 14$ and $\lceil \log_2(14) \rceil - 1 = 3$, so $l(16) = 2$.
- Covering condition (3.1b) : for any $i \in 0, 1, 2, 3$, let p_i be any point with level $l(p_i) = i$. Then $d(p_0, p_1) = 2 \le 2^1$. $d(p_1, p_2) = 4 \le 2^2$ and $d(p_2, p_3) = 8 \le 2^3$.
- Condition (3.1c) : $d_{\min}(C_0) = 2 > 1 = 2^0$, $d_{\min}(C_1) = 4 > 2 = 2^1$, $d_{\min}(C_2) = 8 > 2^2$. ∎

A cover tree was defined in [5, Section 2] as a tree version of a navigating net [13, Section 2.1]. All nodes at a fixed level $i \in \mathbb{Z} \cup \{\pm\infty\}$ form the cover set C_i from Definition 3.1, which can have nodes at different levels in a compressed cover tree. Any point $p \in C_i$ had a single parent in the set C_{i+1}, which satisfied conditions (3.1b,c). The original tree was called an implicit representation of a cover tree in [5, Section 2]. This tree in Figure 4 (left) contains infinite repetitions of every point $p \in R$ and will be called an *implicit cover tree*.

For practical implementations, the authors of [5] had to use another version of a cover tree that they named as an explicit representation of a cover tree and we call an *explicit cover tree*. Here is the full defining quote at the end of [5, Section 2]: "The explicit representation of the tree coalesces all nodes in which the only child is a self-child". In an explicit cover tree, if a subpath of every node-to-root path consists of all identical nodes without other children, all these identical nodes collapse to a single node, see Figure 4 (middle).

Since an explicit cover tree still repeats points, Definition 3.1 is well-motivated by including every point exactly once to save memory and simplify later algorithms, see Figure 4 (right).

▶ **Example 3.3** (a short train line tree). Let G be an unoriented metric graph consisting of two vertices r, q connected by three different edges e, h, g of lengths $|e| = 2^6$, $|h| = 2^3$,
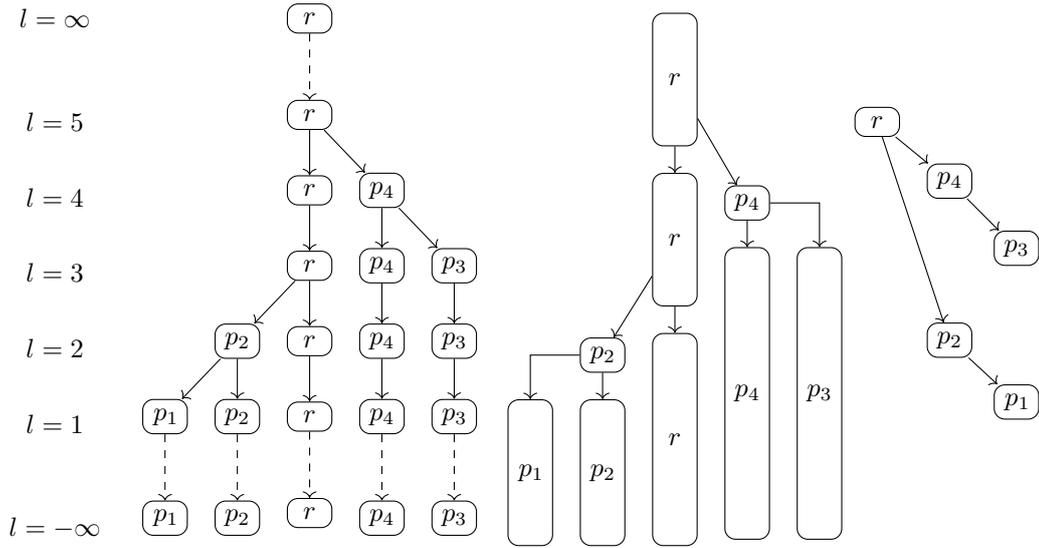
Figure 4 A comparison of past trees and a new compressed cover tree in Example 3.3. **Left:** an implicit cover tree contains infinite repetitions of points. **Middle:** an explicit cover tree also with repeated points. **Right:** a compressed cover tree from Definition 3.1 includes every point exactly once.

$|g| = 1$. Let p_4 be the middle point of the edge e. Let p_3 be the middle point of the subedge (p_4, q). Let p_2 be the middle point of the edge h. Let p_1 be the middle point of subedge (p_2, q). Set $R = \{p_1, p_2, p_3, p_4, r\}$. To construct a compressed cover tree $\mathcal{T}(R)$, set the level $l(p_i) = i$, choose the root r to be the parent of both p_2 and p_4, p_4 to be the parent of p_3 and p_2 to be the parent of p_1. Then $\mathcal{T}(R)$ satisfies Definition 3.1, see three trees in Fig. 4. ∎

In any metric space X, let $\bar{B}(p, t) \subseteq X$ be the closed ball with a center p and a radius t. If this metric space is finite, $|\bar{B}(p, t)|$ denotes the number of points in $\bar{B}(p, t)$.

▶ **Definition 3.4** (expansion constants c [5] and c_m [10, Definition 2.4]). Let R be a finite subset of a space X with a metric d. The *expansion constant $c(R)$* is the smallest real number $c(R) \geq 2$ such that $|\bar{B}(p, 2t)| \leq c(R) \cdot |\bar{B}(p, t)|$ for any $p \in R$ and radius $t \geq 0$. The *minimized expansion constant* $c_m(R) = \inf\limits_{R \subseteq A \subseteq X} c(A)$ is minimized over all finite sets A that cover R. ∎

▶ **Lemma 3.5** (properties of c_m [10, Lemma 2.5]). *For any finite sets $R \subseteq U$ in a metric space, we have $c_m(R) \leq c_m(U)$, $c_m(R) \leq c(R)$.*

[10, Example 2.6] showed that the expansion constant $c(R)$ can be $O(|R|)$. Example 3.6 shows that minimized expansion constant $c_m(R)$ can be much smaller than $c(R)$.

▶ **Example 3.6** (expansion constants). Let $(\mathbb{R}, d)$ be the Euclidean line. For any integer $n > 10$, set $R = \{4^i \mid i \in [1, n]\}$ and let $Q = \{i \mid i \in [1, 4^n]\}$. If $0 < \epsilon < 10^{-9}$, then $\bar{B}(4^n, 2 \cdot 4^{n-1} - \epsilon) = \{2^n\}$ and $\bar{B}(2^n, 2(2 \cdot 4^{n-1} - \epsilon)) = R$. Therefore $c(R) = n$. On the other hand, for any $q \in Q$ and any $t \in \mathbb{R}$, we have the balls $\bar{B}(q, t) = \mathbb{Z} \cap [q - t, q + t]$ and $\bar{B}(q, t) = \mathbb{Z} \cap [q - 2t, q + 2t]$, hence $c(Q) \leq 4$. Lemma 3.5 implies that $c_m(R) \leq c_m(Q) \leq c(Q) \leq 4$. ∎

▶ **Lemma 3.7** (a distance bound on descendants [10, Lemma 2.8]). *Let R be a finite set with a metric d. In a compressed cover tree $\mathcal{T}(R)$, let q be any descendant of a node p. Let the node-to-root path of q have a node $u \in \mathrm{Children}(p) \setminus \{p\}$. Then $d(p, q) \leq 2^{l(u)+2} \leq 2^{l(p)+1}$.* ∎
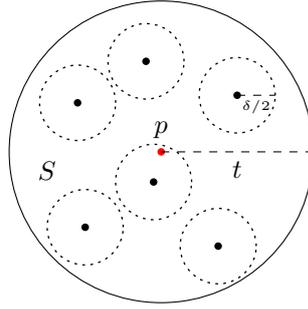
Figure 5 This volume argument proves Lemma 3.8. By using an expansion constant, we can find an upper bound for the number of smaller balls of radius $\frac{\delta}{2}$ that can fit inside a larger ball $\bar{B}(p,t)$.

Lemma 3.8 uses the idea of [8, Lemma 1] to show that if S is a δ-sparse subset of a metric space X, then S has at most $(c_m(S))^\mu$ points in the ball $\bar{B}(p,r)$, where $c_m(S)$ is the minimized expansion constant of S, while μ depends on δ, r.

▶ **Lemma 3.8** (packing [10, Lemma 2.9]). *Let S be a finite δ-sparse set in a metric space (X,d), so $d(a,b) > \delta$ for all $a,b \in S$. Then, for any point $p \in X$ and any radius $t > \delta$, we have $|\bar{B}(p,t) \cap S| \le (c_m(S))^\mu$, where $\mu = \lceil \log_2(\frac{4t}{\delta} + 1) \rceil$.* ∎

[13, Section 1.1] defined $\dim(X)$ of a metric space (X,d) as the minimum number m such that every set $U \subseteq X$ can be covered by 2^m sets whose diameter is a half of the diameter of U. If U is finite, Lemma 3.8 for $\delta = \frac{r}{2}$ shows that $\dim(X) \le \sup_{A \subseteq X}(c_m(A))^4 \le \sup_{A \subseteq X} \inf_{A \subseteq B \subseteq X}(c(B))^4$, where A and B are finite subsets of X.

Let $T(R)$ be an implicit cover tree [5] on a finite set R. [5, Lemma 4.1] showed that the number of children of any node $p \in T(R)$ has the upper bound $(c(R))^4$. Lemma 3.9 generalizes [5, Lemma 4.1] for a compressed cover tree.

▶ **Lemma 3.9** (width bound [10, Lemma 2.10]). *Let R be a finite subset of a metric space (X,d). For any compressed cover tree $\mathcal{T}(R)$, any node p has at most $(c_m(R))^4$ children at every level i, where $c_m(R)$ is the minimized expansion constant of the set R.* ∎

The following concept of a height was introduced and motivated in [10, Definition 2.11].

▶ **Definition 3.10** (the height of a compressed cover tree, [10, Definition 2.11]). For a compressed cover tree $\mathcal{T}(R)$ on a finite set R, the *height set* is $H(\mathcal{T}(R)) = \{l_{\max}, l_{\min}\} \cup \{i \mid C_{i-1} \setminus C_i \ne \emptyset\}$, whose size $|H(\mathcal{T}(R))|$ is called the *height* of $\mathcal{T}(R)$. ∎

By condition (3.1b), the height $|H(\mathcal{T}(R))|$ counts the number of levels i whose cover sets C_i include new points that were absent on higher levels. Since any point can appear alone at its own level, $|H(\mathcal{T})| \le |R|$ is the worst case upper bound of the height. The following parameters help prove an upper bound for the height $|H(\mathcal{T}(R))|$ in Lemma 3.12.

▶ **Definition 3.11** (diameter and aspect ratio). For any finite set R with a metric d, the *diameter* is $\text{diam}(R) = \max_{p \in R} \max_{q \in R} d(p,q)$. The *aspect ratio* [13] is $\Delta(R) = \dfrac{\text{diam}(R)}{d_{\min}(R)}$. ∎

▶ **Lemma 3.12** (an upper bound for the height $|H(\mathcal{T}(R))|$, see [10, Lemma 2.13]). *Any finite set R has the upper bound $|H(\mathcal{T}(R))| \le 1 + \log_2(\Delta(R))$.* ∎

If the aspect ratio $\Delta(R) = O(\text{Poly}(|R|))$ polynomially depends on the size $|R|$, then $|H(\mathcal{T}(R))| \le O(\log(|R|))$. All auxiliary lemmas are proved in appendix B.
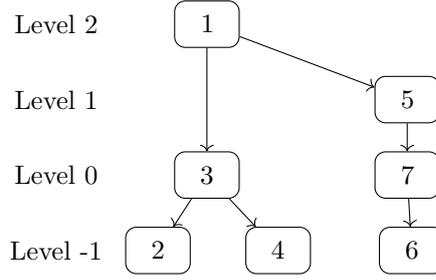
Figure 6 A compressed cover tree $\mathcal{T}(R)$ on the set $R = \{1, 2, 3, 4, 5, 6, 7\}$. Let $\mathcal{S}_i(p, \mathcal{T}(R))$ be a distinctive descendant set from Definition 4.1. Then $V_2(1) = \emptyset, V_1(1) = \{5\}, V_0(1) = \{3, 5, 7\}$, and $\mathcal{S}_2(1, \mathcal{T}(R)) = \{1, 2, 3, 4, 5, 6, 7\}$, $\mathcal{S}_1(1, \mathcal{T}(R)) = \{1, 2, 3, 4\}$, $\mathcal{S}_0(1, \mathcal{T}(R)) = \{1\}$.

4 Distinctive descendant sets and other tools for neighbor candidates

This section we introduce a few more concepts and results from [10] for the new theorems in sections 5 and 6. The first concept is a distinctive descendant set at a level i of a node p in a compressed cover tree. Briefly, it is the set of descendants of a copy of p at the level i in the implicit cover tree on R. Another important concept is a λ-point in Definition 4.8, which will be used in Algorithm 6.2 as an approximation for a k-nearest neighbor. Lemma 4.16 about a β-point will play a major role in the proof of main Theorem 6.8.

▶ **Definition 4.1** (Distinctive descendant sets). Let $R \subseteq X$ be a finite reference set with a cover tree $\mathcal{T}(R)$. For any node $p \in \mathcal{T}(R)$ in a compressed cove tree on a finite set R and $i \leq l(p) - 1$, set $V_i(p) = \{u \in \text{Descendants}(p) \mid i \leq l(u) \leq l(p) - 1\}$. If $i \geq l(p)$, then set $V_i(p) = \emptyset$. For any level $i \leq l(p)$, the *distinctive descendant set* is $\mathcal{S}_i(p, \mathcal{T}(R)) = \text{Descendants}(p) \setminus \bigcup_{u \in V_i(p)} \text{Descendants}(u)$ and has the size $|\mathcal{S}_i(p, \mathcal{T}(R))|$. ■

▶ **Lemma 4.2** (Distinctive descendant set inclusion property). *In conditions of Definition 4.1 let $p \in R$ and let i, j be integers satisfying $l_{\min}(\mathcal{T}(R)) \leq i \leq j \leq l(p) - 1$. Then $\mathcal{S}_i(p, \mathcal{T}(R)) \subseteq \mathcal{S}_j(p, \mathcal{T}(R))$.*

▶ **Definition 4.3** (Children(p, i) and Next$(p, i, \mathcal{T}(R))$ for a compressed cover tree). In a compressed cover tree $\mathcal{T}(R)$ on a set R, for any level i and a node $p \in R$, set Children$(p, i) = \{a \in \text{Children}(p) \mid l(a) = i\}$. Let Next$(p, i, \mathcal{T}(R))$ be the maximal level j satisfying $j < i$ and Children$(p, i) \neq \emptyset$. For every node p, we store its set of children in a linked hash map so that

(1) any key i gives access to Children(p, i),
(2) every Children(p, i) has access to Children$(p, \text{Next}(p, i, \mathcal{T}(R)))$,
(3) we can directly access $\max\{j \mid \text{Children}(p, j) \neq \emptyset\}$. ■

Essential levels of a node $p \in \mathcal{T}(R)$ has 1-1 correspondence to the set consisting of all nodes containing p in the explicit representation of cover tree in [5], see Figure 4 middle.

▶ **Definition 4.4** (Essential levels of a node). Let $R \subseteq X$ be a finite reference set with a cover tree $\mathcal{T}(R)$. Let $q \in \mathcal{T}(R)$ be a node. Let (t_i) for $i \in \{0, 1, ..., n\}$ be a sequence of $H(\mathcal{T}(R))$ in such a way that $t_0 = l(q)$, $t_n = l_{\min}(\mathcal{T}(R))$ and for all i we have $t_{i+1} = \text{Next}(q, t_i, \mathcal{T}(R))$. Define the set of essential indices $\mathcal{E}(q, \mathcal{T}(R)) = \{t_i + 1 \mid i \in \{0, ..., n\}\}$. ■

▶ **Lemma 4.5** (Number of essential levels). *Let $R \subseteq X$ be a finite reference set with a cover tree $\mathcal{T}(R)$. Then $\sum_{p \in R} |\mathcal{E}(p, \mathcal{T}(R))| \leq 2 \cdot |R|$, where $\mathcal{E}(p, \mathcal{T}(R))$ appears in Definition 4.4.* ■

Proof. Let us prove this claim by induction on size $|R|$. In basecase $R = \{r\}$ and therefore $|\mathcal{E}(r, \mathcal{T}(R))| = 1$. Assume now that the claim holds for any tree $\mathcal{T}(R)$, where $|R| = m$ and let us prove that if we add any node $v \in X \setminus R$ to tree $\mathcal{T}(R)$, then $\sum_{p \in R} |\mathcal{E}(p, \mathcal{T}(R \cup \{v\}))| \leq 2 \cdot |R| + 2$. Assume that we have added u to $\mathcal{T}(R)$, in such a way that v is its new parent. Then $|\mathcal{E}(p, \mathcal{T}(R \cup \{v\}))| = |\mathcal{E}(p, \mathcal{T}(R))| + 1$ and $|\mathcal{E}(v, \mathcal{T}(R \cup \{v\}))| = 1$. We have:

$$\sum_{p \in R \cup \{u\}} |\mathcal{E}(p, \mathcal{T}(R))| = \sum_{p \in R} |\mathcal{E}(p, \mathcal{T}(R))| + 1 + |\mathcal{E}(v, \mathcal{T}(R \cup \{v\}))| \leq 2 \cdot |R| + 2 \leq 2(|R \cup \{v\}|)$$

which completes the induction step.   ◄

1: **Function** : CountDistinctiveDescendants(Node p, a level i of $\mathcal{T}(R)$)
2: **Output** : Number of distinctive descendants of node p.
3: **if** $i > l_{\min}(\mathcal{T}(Q))$ **then**
4:     **for** $q \in \text{Children}(p)$ having $l(p) = i - 1$ or $q = p$ **do**
5:         Set $s = 0$
6:         $j \leftarrow 1 + \text{Next}(q, i - 1, \mathcal{T}(R))$
7:         $s \leftarrow s + \text{CountDistinctiveDescendants}(q, j)$
8:     **end for**
9: **else**
10:     Set $s = 1$
11: **end if**
12: Set $|\mathcal{S}_i(p)| = s$ and **return** s

■ **Algorithm 4.6** This algorithm returns sizes of distinctive descendant set $\mathcal{S}_i(p, \mathcal{T}(R))$ for all essential levels $i \in \mathcal{E}(p, \mathcal{T}(R))$

▶ **Lemma 4.7** ([10, Lemma 5.3]). *Let R be a finite subset of a metric space. Let $\mathcal{T}(R)$ be a compressed cover tree on R. Then, Algorithm 4.6 computes the sizes $|\mathcal{S}_i(p, \mathcal{T}(R))|$ for all $p \in R$ and essential levels $i \in \mathcal{E}(p, \mathcal{T}(R))$ in time $O(|R|)$.* ■

Let i be an arbitrary level. Set $j = \text{Next}(p, i, \mathcal{T}(R))$. Definition 4.3 of Next implies that $l(q) \notin [j + 1, i - 1]$ for all $q \in \text{Descendants}(p)$. Therefore we have $V_i(p) = V_j(p)$ and $\mathcal{S}_i(p, \mathcal{T}(R)) = \mathcal{S}_j(p, \mathcal{T}(R))$. It follows that $|\mathcal{S}_i(p, \mathcal{T}(R))|$ can change only for the indices $i \in H(\mathcal{T}(R))$. Lemma 4.7 shows that all the essential distinctive descendants sets of compressed cover tree $\mathcal{T}(R)$ can be precomputed in $O((c_m(R))^4 \cdot |R| \cdot |H(\mathcal{T}(R))|)$ time.

Recall the neighborhood $N(q; r) = \{p \in C \mid d(q, p) \leq d(q, r)\}$ from Definition 1.1.

▶ **Definition 4.8** (λ-point [10, Definition 5.6]). Fix a query point q in a metric space (X, d) and fix any level $i \in \mathbb{Z}$. Let $\mathcal{T}(R)$ be its compressed cover tree on a finite reference set $R \subseteq X$. Let C be a subset of a cover set C_i from Definition 3.1 satisfying $\sum_{p \in C} |\mathcal{S}_i(p, \mathcal{T}(R))| \geq k$, where $\mathcal{S}_i(p, \mathcal{T}(R))$ is the distinctive descendant set from Definition 4.1. For any $k \geq 1$, define $\lambda_k(q, C)$ as a point $\lambda \in C$ that minimizes $d(q, \lambda)$ subject to $\sum_{p \in N(q; \lambda)} |\mathcal{S}_i(p, \mathcal{T}(R))| \geq k$. ■

▶ **Lemma 4.10** (time complexity of a λ-point [10, Lemma 5.7]). *In the notations of Definition 4.8, the time complexity of Algorithm 4.9 is $O(|C| \log(k) \max_i |R_i|)$.* ■

▶ **Lemma 4.11** (separation lemma [10, Lemma 5.8]). *In the conditions of Definition 4.1, let $p \neq q$ be nodes of $\mathcal{T}(R)$ with $l(p) \geq i$, $l(q) \geq i$. Then $\mathcal{S}_i(p, \mathcal{T}(R)) \cap \mathcal{S}_i(q, \mathcal{T}(R)) = \emptyset$.* ■

1: **Input:** A point $q \in X$, a subset C of level set C_i of a compressed cover tree $\mathcal{T}(R)$, an integer $k \in \mathbb{Z}$
2: Initialize an empty max-binary heap B and an empty array D.
3: **for** $p \in C$ **do**
4:     add p to B with priority $d(q, p)$
5:     **if** $|H| \geq k$ **then**
6:         remove the point with a maximal value from B
7:     **end if**
8: **end for**
9: Transfer points from the binary heap B to the array D in reverse order.
10: Find the smallest index j such that $\sum_{t=0}^{j} \mathcal{S}_i(D[t], \mathcal{T}(R)) \geq k$.
11: **return** $\lambda = D[j]$.

■ **Algorithm 4.9** Computing a λ-point of Definition 4.8 in line 3 of Algorithm 6.2, [10, Lemma 5.6]

▶ **Lemma 4.12** (sum lemma [10, Lemma 5.9]). *In the notations of Definition 4.8 for any subset $V \subseteq C$ of a set C, we have $|\bigcup_{p \in V} \mathcal{S}_i(p, \mathcal{T}(R))| = \sum_{p \in V} |\mathcal{S}_i(p, \mathcal{T}(R))|$.* ∎

By Lemma 4.12 we can assume in Definition 4.8 that $|\bigcup_{p \in C} \mathcal{S}_i(p, \mathcal{T}(R))| \geq k$.

▶ **Lemma 4.13** ([10, Lemma 5.11]). *In the notations of Definition 4.1, let $p \in \mathcal{T}(R)$ be any node. If $w \in \mathcal{S}_i(p, \mathcal{T}(R))$ then either $w = p$ or there exists $a \in \mathrm{Children}(p) \setminus \{p\}$ such that $l(a) < i$ and $w \in \mathrm{Descendants}(a)$.* ∎

▶ **Lemma 4.14.** *In the notations of Definition 4.1, let $p \in \mathcal{T}(R)$ be any node. If $w \in \mathcal{S}_i(p, \mathcal{T}(R))$ then $d(w, p) \leq 2^{i+1}$.* ∎

Proof. By Lemma 4.13 either $w = \gamma$ or $w \in \mathrm{Descendants}(a)$ for some $a \in \mathrm{Children}(\gamma) \setminus \{\gamma\}$ for which $l(a) < i$. If $w = \gamma$, then trivially $d(\gamma, w) \leq 2^i$. Else w is a descendant of a, which is a child of node γ on level $i - 1$ or below, therefore by Lemma 3.7 we have $d(\gamma, w) \leq 2^i$ anyway. ◀

▶ **Lemma 4.15** ([10, Lemma 5.12]). *Let R be a finite subset of a metric pace. Let $\mathcal{T}(R)$ be a compressed cover tree on R. Let $R_i \subseteq C_i$, where C_i is the ith cover set of $\mathcal{T}(R)$. Set $\mathcal{C}(R_i) = \{a \in \mathrm{Children}(p) \text{ for some } p \in R_i \mid l(a) \geq i - 1\}$. Then*

$$\bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R)) = \bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R)).$$

▶ **Lemma 4.16** (β-point [10, Lemma 5.13]). *In the notations of Definition 4.8, let $C \subseteq C_i$ so that $\cup_{p \in C} \mathcal{S}_i(p, \mathcal{T}(R))$ contains all k-nearest neighbors of q. Set $\lambda = \lambda_k(q, C)$. Then R has a point β among the first k nearest neighbors of q such that $d(q, \lambda) \leq d(q, \beta) + 2^{i+1}$.* ∎

5   A parametrized complexity of a traversal algorithm for paired trees

To fix the problem occurring in Counterexample A.8, Definition 5.1 will introduce below a new concept of the imbalance inspired by [8, Definitions 2-4]. Any balanced tree from Example 3.2 will have a near linear imbalance $I(\mathcal{T}(R), \mathcal{T}(R)) = O(|R|)$.

Main Theorem 5.7 of this section will express the time complexity of neighbor search based on paired trees in terms of $|Q|$, the height $|H(\mathcal{T}(R))|$, the new imbalance $I(\mathcal{T}(Q), \mathcal{T}(R))$, and time complexities of auxiliary functions UpdateCandidates() and FinalCandidates().

▶ **Definition 5.1** (the imbalance I of paired compressed cover trees). Let R, Q be finite subsets of an ambient space. Let $\mathcal{T}(Q), \mathcal{T}(R)$ be compressed cover trees on Q, R, respectively. Fro any $i \in \mathbb{Z}$, set $H_i(\mathcal{T}(R)) = H(\mathcal{T}(R)) \cap [l_{\min}(\mathcal{T}(R)), i]$, where the height set $H(\mathcal{T}(R))$ was introduced in Definition 3.10. Then the *imbalance* of the paired trees $(\mathcal{T}(Q), \mathcal{T}(R))$ is

$$I(\mathcal{T}(Q), \mathcal{T}(R)) = \sum_{q \in \mathcal{T}(Q)} |H_{l(q)}(\mathcal{T}(R))|. \tag{5.1}$$

Definition 5.1 combines the parameter θ from [8, Definition 4] and the original imbalance I_t from [8, Definition 3] into the single new imbalance $I(\mathcal{T}(Q), \mathcal{T}(R))$.

Lemma 5.2 below gives a rough upper bound for the imbalance of paired cover trees. This rough bound weakens the new time complexity from main Theorem 6.8 based on paired trees to the previous time complexity from [10, Theorem 6.4] based on a single tree.

▶ **Lemma 5.2** (imbalance bounds). *Let R, Q be finite subsets of some ambient metric space. Let $\mathcal{T}(Q), \mathcal{T}(R)$ be compressed cover trees on Q, R, respectively. Then*

$$I(\mathcal{T}(Q), \mathcal{T}(R)) \leq |Q| \cdot |H(\mathcal{T}(R))|. \tag{5.2}$$

If $\mathcal{T}(Q) = \mathcal{T}(R)$, then $I(\mathcal{T}(R), \mathcal{T}(R)) \geq |H(\mathcal{T}(R))|$ and $I(\mathcal{T}(R), \mathcal{T}(R)) \geq |R|$.

Proof. The bound in (5.2) is obtained as follows:

$$I(\mathcal{T}(Q), \mathcal{T}(R)) = \sum_{q \in \mathcal{T}(Q)} |H_{l(q)}(\mathcal{T}(R))| \leq \sum_{q \in \mathcal{T}(Q)} |H(\mathcal{T}(R))| \leq |Q| \cdot |H(\mathcal{T}(R))|.$$

Let r be the root of the tree $\mathcal{T}(R)$. The second bound is obtained as follows:

$$I(\mathcal{T}(R), \mathcal{T}(R)) = \sum_{q \in \mathcal{T}(Q)} |H_{l(q)}(\mathcal{T}(R))| \geq |H_{l(r)}(\mathcal{T}(R))| \geq |H(\mathcal{T}(R))|.$$

Finally, $I(\mathcal{T}(R), \mathcal{T}(R)) = \sum_{q \in \mathcal{T}(Q)} |H_{l(q)}(\mathcal{T}(R))| \geq \sum_{q \in \mathcal{T}(Q)} 1 \geq |Q| = |R|.$ ◀

▶ **Example 5.3** (imbalance computation). The sets $Q = R = \{1, 2, 3, 4, 5\}$ have compressed cover trees $\mathcal{T}(R) = \mathcal{T}(Q)$ in Figure 1. Since $\mathcal{T}(R)$ has nodes on levels $-1, 0, 1, 2$, the height set is $H(\mathcal{T}(R)) = \{-1, 0, 1, 2\}$. The levels of all nodes are $l(1) = 2$, $l(5) = 1$, $l(3) = 0$, $l(2) = -1$ and $l(4) = -1$. Then $H_2(\mathcal{T}(R)) = H(\mathcal{T}(R))$, $H_1(\mathcal{T}(R)) = \{-1, 0, 1\}$, $H_0(\mathcal{T}(R)) = \{-1, 0\}$, $H_{-1}(\mathcal{T}(R)) = \{-1\}$. By Definition 5.1 the imbalance is $I(\mathcal{T}(Q), \mathcal{T}(R)) = |H_2(\mathcal{T}(R))| + |H_1(\mathcal{T}(R))| + |H_0(\mathcal{T}(R))| + 2 \cdot |H_{-1}(\mathcal{T}(R))| = 4 + 3 + 2 + 2 \cdot 1 = 11$, so $|Q| \cdot |H(\mathcal{T}(R))| = 5 \cdot 4 = 20$.

▶ **Lemma 5.4** (imbalance of a balanced tree pair). *Let R be a finite set with a metric and size $|R| = \frac{t^{m+1} - 1}{t - 1}$ for some integers m, t. Let a compressed cover tree $\mathcal{T}(R)$ have a structure of a balanced tree so that $H(\mathcal{T}(R)) = \{l_0, ..., l_m\}$, $|H(\mathcal{T}(R))| = m + 1$, and $l_i < l_{i+1}$ for all i. Then for $0 \leq i \leq m$, the tree $\mathcal{T}(R)$ has exactly t^i nodes at the level l_i. The imbalance is*

$$I(\mathcal{T}(R), \mathcal{T}(R)) = \left(1 + \frac{1}{t-1}\right)|R| - \frac{|H(\mathcal{T}(R))|}{t - 1} = O(|R|).$$

Proof. The computations below follow directly from Definition 5.1: $I(\mathcal{T}(R), \mathcal{T}(R)) = \sum_{i=0}^{m} |H_{l_i}(\mathcal{T}(R))| \cdot t^i = \sum_{i=0}^{m} (m - i + 1) \cdot t^i = \frac{t^{m+2} - mt + m - 2t + 1}{(t - 1)^2} = |R| + \frac{t^{m+1} - mt + m - t}{(t-1)^2} = |R| + \frac{|R|}{t - 1} + \frac{-mt + m - t + 1}{(t - 1)^2} = (1 + \frac{1}{t - 1})|R| - \frac{|H(\mathcal{T}(R))|}{t - 1} = O(|R|).$ ◀

Algorithm 5.5 adapts [8, Algorithm 1] to compressed cover trees by replacing the functions BaseCase() and Score() by the new ones. In Algorithm 5.5 the function FindCandidates() replaces the function BaseCase() to compute a final output for a given query point when we reach the lowest level of $\mathcal{T}(R)$. The function UpdateCandidates() replaces the function Score() to move neighbor candidates from R_i to R_{i-1}. In Section 6 Algorithm 5.5 will be updated to Algorithm 6.3 to solve Problem 1.2 finding all k-nearest neighbors of Q in R.

We usually call Algorithm 5.5 with the arguments $(l_{\max}(\mathcal{T}(R)), l_{\max}(\mathcal{T}(Q)), q, \{r\})$, where q, r are root nodes of compressed cover trees on Q, R, respectively.

1: Function **PairedTreeTraversal**(a level $i \in \mathbb{Z}$, a level $j \in \mathbb{Z}$, a query node q of $\mathcal{T}(Q)$, a subset R_i of the cover set C_i of $\mathcal{T}(R)$)
2: **if** $i = l_{\min}(\mathcal{T}(R))$ **then**
3:     FinalCandidates(i, j, q, R_i)
4: **end if**
5: **if** $\max(l_{\min}(\mathcal{T}(R)), j) < i$ **then**
6:     {Reference expansion lines 5-11}
7:     Set $\mathcal{C}(R_i) = \{a \in \text{Children}(p) \text{ for some } p \in R_i \mid l(a) \geq i - 1\}$
        {recall that Children(p) contains the node p }
8:     $R_{i-1} \leftarrow \text{UpdateCandidates}(i, j, q, \mathcal{C}(R_i))$
9:     Set $t \leftarrow 1 + \max_{a \in R_{i-1}} \text{Next}(a, i - 1, \mathcal{T}(R))$
10:    Set $R_t \leftarrow R_{i-1}$
11:    **call PairedTreeTraversal**(q, j, R_t) {recall that R_t is a subset of C_t}
12: **else**
13:    {query expansion lines 12-18}
14:    **for** $q' \in \text{Children}(q)$ having $l(q') = j - 1$ or $q' = q$ **do**
15:       $j' \leftarrow 1 + \text{Next}(q', j, \mathcal{T}(Q))$
16:       **call PairedTreeTraversal**(q', j', R_i)
17:    **end for**
18: **end if**

■ **Algorithm 5.5** We adapt [8, Algorithm 1] for compressed cover trees, see Complexity Theorem 5.7.

▶ **Lemma 5.6** (reference expansions in Algorithm 5.5). *Let Q, R be finite subsets of a metric space. Let $\mathcal{T}(Q)$ and $\mathcal{T}(R)$ be compressed cover trees on Q and R, respectively. Let us call Algorithm 5.5 with the input consisting of root nodes of $\mathcal{T}(Q)$, $\mathcal{T}(R)$ and their levels. The algorithm runs at most $I(\mathcal{T}(Q), \mathcal{T}(R)) + |H(\mathcal{T}(R))|$ the reference expansions (lines 5-11).*

Proof. Let ξ be the total number of reference expansions. Let q be any non-root query node and let p be its parent. We note that any call for PairedTreeTraversal(i, j, q, R_i) comes from the query expansion (lines 10-15) in PairedTreeTraversal$(i', l(q), p, R_{i'})$, see Algorithm 5.5.

Then $l(q) \geq i' \geq i$. Since q is the root node of $\mathcal{T}(Q)$, we can visit all levels from $H(\mathcal{T}(R))$, so $\xi \leq \sum\limits_{q \in Q \setminus \text{root}(\mathcal{T}(Q))} |H_{l(q)}(\mathcal{T}(R))| + |H(\mathcal{T}(R))| \leq I(\mathcal{T}(Q), \mathcal{T}(R)) + |H(\mathcal{T}(R))|$.   ◄

Theorem 5.7 adapts [8, Theorem 1] to compressed cover trees and has the following advantages. [8, Theorem 1] used the single size $N = \max |Q|, |R|$ to estimate the time complexity. We replace this maximum size N by $|Q| + |H(\mathcal{T}(R))|$, which is usually smaller than N for online search in a large reference set R, especially if Q is a single point.

▶ **Theorem 5.7** (time complexity of paired trees traversal). *Let* Q, R *be finite subsets of a metric space. Let* $\mathcal{T}(Q)$ *and* $\mathcal{T}(R)$ *be compressed cover trees on the sets* Q *and* R, *respectively. Let* B, S *be the time complexities of the functions FinalCandidates() and UpdateCandidates() operation, respectively. Let* $I(\mathcal{T}(Q), \mathcal{T}(R))$ *be the imbalance from Definition 5.1. Then Algorithm 5.5 has time complexity* $O\Big((B + S) \cdot \big(|H(\mathcal{T}(R))| + |Q| + I(\mathcal{T}(Q), \mathcal{T}(R))\big)\Big)$.

Proof. First, split the algorithm into two parts: reference expansions (lines 5 - 11) and query expansions (lines 12–18). The time complexity of the algorithm is bounded by the total number of reference expansions times the maximum time of a reference expansion plus the total number of all query expansions times maximal time of a query expansion.

The time of a reference expansion (lines 5 - 11) dominated by the sum $O(B + S)$ of the time complexities of the functions FinalCandidates() and UpdateCandidates(). Query expansions (lines 12–18) are launched for every combiantion $q \in Q$ and $j \in \mathcal{E}(q, \mathcal{T}(Q))$ of Definition 4.4. By Lemma 4.5 their total number is $2 \cdot |Q|$ Hence the total time complexity of all query expansions (lines 5 -11) is $O(|Q|)$. By Lemma 5.6 the number of reference expansions is at most $|H(\mathcal{T}(R))| + I(\mathcal{T}(R), \mathcal{T}(Q))$. The final time complexity is the sum $O(|Q|)+$
$O((B+S)(|H(\mathcal{T}(R))|+I(\mathcal{T}(Q), \mathcal{T}(R)))) = O\big((B+S)(|Q|+|H(\mathcal{T}(R))|+I(\mathcal{T}(Q), \mathcal{T}(R)))\big)$. ◀

Theorem 5.7 holds even if a tree $\mathcal{T}(Q)$ fails the separation condition (3.1c). This fact allows us to find a more balanced tree $\mathcal{T}(Q)$ to minimize the imbalance $I(\mathcal{T}(Q), \mathcal{T}(R))$ and substantially decrease the time complexity of the paired tree traversal in Algorithm 5.5.

▶ **Corollary 5.8** (case $Q = R$). *Let* R *have a compressed cover tree* $\mathcal{T}(R)$. *Let* B, S *be the time complexities of the functions FinalCandidates() and UpdateCandidates(), respectively. Algorithm 5.5 for* $Q = R$ *has time complexity* $O\Big((B + S) \cdot I(\mathcal{T}(R), \mathcal{T}(R))\Big)$.

Proof. It follows from Theorem 5.7 and Lemma 5.2 as $I(\mathcal{T}(R), \mathcal{T}(R)) \geq |R| \geq H(\mathcal{T}(R))$. ◀

6 A new algorithm for all k-nearest neighbors based on paired trees

In this section we assume that a query set Q and a reference set R are finite subsets of an ambient metric space X with a metric d. Algorithm 6.3 will use paired trees $\mathcal{T}(R)$ and $\mathcal{T}(Q)$ to solve Problem 1.2 for any $k \geq 1$. Main Theorem 6.8 proves a near linear time complexity for Algorithm 6.3, which runs Algorithm 5.5 by using Algorithm 6.1 for the function FinalCandidates() and Algorithm 6.2 for the function UpdateCandidates().

Both functions FinalCandidates() and UpdateCandidates() need current levels i, j of compressed cover trees $\mathcal{T}(R), \mathcal{T}(Q)$, respectively. Other input arguments are a query point q and a subset R_i of a reference set R. In general, the function FinalCandidates() can use a level j of a tree $\mathcal{T}(Q)$. However, our implementation of Algorithm 6.1 does not need j.

During the tree traversal for any $q \in Q$, we maintain the array N_q (of candidate neighbors) consisting of pairs (a, t), where $a \in R$ and $t = d(q, a)$. To minimize the run time of line 2 in Algorithm 6.1, we implement the array N_q as a max-heap [6, section 6.5], which has time complexity $O(\log(|N_q|))$ for adding elements, finding and removing a maximal element.

▶ **Lemma 6.4** (true k-nearest neighbors are in a candidate set for all levels i). *Let* Q, R *be finite subsets of a space* X *with a metric* d. *Let let* $k \leq |R|$ *be an integer. Assume that we have already constructed cover trees* $\mathcal{T}(Q)$ *and* $\mathcal{T}(R)$ *on* Q *and* R, *respectively. Consider any recursion* (i, j, q, R_i) *of line 3 in Algorithm 6.3. Then the union* $\bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R))$ *contains all* k-*nearest neighbors of any point* $q' \in \text{Descendants}(q)$.

1: **Input** : a level i of $\mathcal{T}(R)$, a level j of $\mathcal{T}(Q)$, a query point q, a subset $R_i \subset R$
2: Select the k-nearest neighbors of q in the set R_i and them into the array N_q

■ **Algorithm 6.1** This function FinalCandidates() is used in Algorithm 5.5 to get final Algorithm 6.3.

1: **Input** : a level i of $\mathcal{T}(R)$, a level j of $\mathcal{T}(Q)$, a query point $q \in Q$,
   a subset $\mathcal{C}(R_i)$ of C_{i-1} of a compressed cover tree $\mathcal{T}(R)$
2: **Output** : a subset $R_{i-1} \subset C_{i-1}$ consisting of all nodes that were not pruned.
3: Compute a point $\lambda = \lambda_k(q, \mathcal{C}(R_i))$ from Definition 4.8 by Algorithm 4.9.
4: **return** $R_{i-1} = \{a \in \mathcal{C}(R_i) \mid d(q,a) \le d(q,\lambda) + 2^{i+1} + 2^{j+2}\}$

■ **Algorithm 6.2** This function UpdateCandidates() is used in Algorithm 5.5 to get Algorithm 6.3.

Proof. Assuming the contrary, let a recursion (i, j, q, R_i) fail the required claim for some i, j, q. Define the order on $\mathbb{Z} \times \mathbb{Z}$ so that $(a, b) \le (c, d)$ if $b < d$ or if $b = d$ and $a \le b$.

For every point q, let $(i, j)_q$ be the maximal element for which (i, j, q, R_i) doesn't satisfy the claim. Set $q = \mathrm{argmax}_{q \in Q}(i, j)_q$ and $(i, j) = \max_{q \in Q}(i, j)_q$. Fix an arbitrary descendant $q' \in \mathrm{Descendants}(q)$. Since $j \le l(q) - 1$, Lemma 3.9 implies that $d(q', q) \le 2^{j+1}$. Choose

$$\beta \in \bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R)) \setminus \bigcup_{p \in R_{i-1}} \mathcal{S}_{i-1}(p, \mathcal{T}(R)).$$

By Lemma 4.15 we have

$$\bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R)) = \bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R)). \tag{1}$$

Let λ be a point computed in line 3 of Algorithm 6.2. Equation (1) implies that $|\bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R))| \ge k$. Then the point λ is well-defined by Definition 4.8. Since

$$\beta \in \bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R)),$$

there exists $\alpha \in \mathcal{C}(R_i)$ satisfying $\beta \in \mathcal{S}_{i-1}(\alpha, \mathcal{T}(R))$. The initial assumption means that $\alpha \notin R_{i-1}$. Due to line 4 of Algorithm 6.2, we get

$$d(\alpha, q) > d(p, \lambda) + 2^{i+1} + 2^{j+2}. \tag{2}$$

Let w be an arbitrary point in the union $\bigcup_{p \in N(q;\lambda)} \mathcal{S}_{i-1}(p, \mathcal{T}(R))$. Then $w \in \mathcal{S}_{i-1}(\gamma, \mathcal{T}(R))$ for some $\gamma \in N(q; \lambda)$. By Lemma 4.13 either $w = \gamma$ or $w \in \mathrm{Descendants}(a)$ for a point $a \in \mathrm{Children}(\gamma) \setminus \{\gamma\}$ at a level $l(a) < i$. If $w = \gamma$, then trivially $d(\gamma, w) \le 2^i$. Else w is a descendant of a, which is a child of node γ on level $i - 1$ or below.

1: **Input:** a query tree $\mathcal{T}(Q)$, a reference tree $\mathcal{T}(R)$, a number of neighbors $k \ge 1$
2: **Output:** arrays N_q of k nearest neighbors for all points $q \in Q$
3: run Algorithm 5.5 for the arguments
   $(l_{\max}(\mathcal{T}(R)), l_{\max}(\mathcal{T}(Q)), \mathrm{root}(\mathcal{T}(Q)), \{\mathrm{root}(\mathcal{T}(R))\})$ by using
   - the function UpdateCandidates() is implemented in Algorithm 6.2
   - the function FinalCandidates() is implemented in Algorithm 6.1
4: **return** the array N_q for all points $q \in Q$.

■ **Algorithm 6.3** The new algorithm for all k-nearest neighbors by using paired trees on Q, R

Then by Lemma 3.7 we have $d(\gamma, w) \leq 2^i$ anyway. By Definition 4.8 since $\gamma \in N(q; \lambda)$ we have $d(q, \gamma) \leq d(q, \lambda)$. We apply (2) and the triangle inequality below:

$$d(q', w) \leq d(q', q) + d(q, p) + d(p, w) \leq 2^{j+1} + d(\lambda, p) + 2^i < d(\alpha, q) - 2^i - 2^{j+1}. \quad (3)$$

On the other hand, β is a descendant of α, so we can estimate

$$d(\beta, q') \geq d(\alpha, q) - d(\alpha, \beta) - d(q, q') \geq d(\alpha, q) - 2^i - 2^{j+1}. \quad (4)$$

Inequalities (3), (4) imply that $d(q', w) < d(q', \beta)$. The point w was arbitrarily chosen in the union $\bigcup_{p \in N(q; \lambda)} \mathcal{S}_{i-1}(p, \mathcal{T}(R))$ containing at least k points. Hence the point β cannot be among the first k nearest neighbors of q'. This final contradiction proves the lemma. ◄

▶ **Theorem 6.5** (correctness of Algorithm 6.3 for all k-nearest neighbors). Given any compressed cover trees $\mathcal{T}(Q)$ and $\mathcal{T}(R)$ on a query set Q and a reference set R, respectively, Algorithm 6.3 solves Problem 1.2 by finding all k-nearest neighbors in R for all points $q \in Q$.

Proof. The nodes of $\mathcal{T}(R)$ cannot have children below the level $l_{\min}$, so $\bigcup_{p \in R_{l_{\min}}} \mathcal{S}_i(p, \mathcal{T}(R)) = R_{l_{\min}}$. Then all the k-nearest neighbors of any $q \in Q$ are contained in the set $R_{l_{\min}}$ obtained in the last recursion of Algorithm 5.5 involving q. Lemma 6.4 finishes the proof. ◄

▶ **Lemma 6.6** (time complexity of the function FinalCandidates()). *For $k \leq |R|$, Algorithm 6.1 with the input arguments (i, j, q, R_i) has the time complexity $O(\log(k) \cdot |R_i|)$.*

Proof. The time complexity of Algorithm 6.1 is dominated by line 2, where we place points of R_i into the array N_q one by one until we find k neighbors. At every step we iterate over the remaining points in the subset R_i by adding them to the array N_q one by one and by removing the node having a maximum distance to q. Since adding a new element and removing a maximal element from N_q takes $O(\log(|N_q|))$ time, the time complexity of the function FinalCandidates() is $O(|R_i| \cdot \log(|N_q|)) = O(\log(k) \cdot |R_i|)$ due to $|N_q| \leq k + 1$. ◄

▶ **Lemma 6.7** (complexity of the function UpdateCandidates()). *For $k \leq |R|$, Algorithm 6.2 with the input arguments (i, j, q, R_i) has the time complexity $O(c_m(R)^4 \cdot \log(k) \cdot |R_i|)$.*

Proof. Lemma 3.9 says that $|\mathcal{C}(R_i)| \leq (c_m(R))^4 \cdot |R_i|$. By Lemma 4.10 line 3 of Algorithm 6.2 takes time $O(\log(k) \cdot \mathcal{C}(R_i)) \leq O(c_m(R)^4 \cdot \log(k) \cdot |R_i|)$. Line 4 takes time $c_m(R)^4 \cdot |R_i|$. So the time complexity is dominated by line 3 and the lemma holds. ◄

Theorem 6.8 will substantially extend [8, Theorem 2] solving Problem 1.2 only for $k1$ to compressed cover trees to find all k-nearest neighbors for any $k \geq 1$. We compare the time complexity $c(R)^4 c_{qr}^5 (N + I_t(\mathcal{T}_q)) + \theta)$ from [8, Theorem 2] with Theorem 6.8 below.

The sum of the past parameters $I_t(\mathcal{T}_q)$ and θ is replaced by the new imbalance $I(\mathcal{T}(Q), \mathcal{T}(R))$. The size $N = \max(|Q|, |R|)$ reduces to $O(|Q| + |H(\mathcal{T}(R))|)$. The expansion constant $c(R)$ is replaced by the minimized expansion constant $c_m(R)$ from Definition 3.4. The factor c_{qr}^5 becomes $\max\{c_{qr}^3 k, c_m(R)^7\} \log(k)$. The appearance of k is natural, because Problem 1.2 is solved for any $k \geq 1$. The difference in powers of c_{qr}^5 and $c_m(R)^7$ is explained by the need to include more nodes into the subset R_{i-1} in line 4 of Algorithm 6.2.

▶ **Theorem 6.8** (Complexity for k-nearest neighbors based on paired trees). In the notations of Theorem 6.5, let $c_{qr} = \max_{q \in Q} c(R \cup \{q\})$. Then the time complexity of Algorithm 6.3 which finds k-nearest neighbors for all $q \in Q$ in set R as stated in Problem 1.2 is

$$O\Big(c_m(R)^4 \cdot \max\{c_{qr}^3 k, (c_m(R))^7\} \cdot \log(k) \cdot \big(|H(\mathcal{T}(R))| + |Q| + I(\mathcal{T}(R), \mathcal{T}(Q))\big)\Big), \quad (6.8)$$

where the imbalance $I(\mathcal{T}(R), \mathcal{T}(Q))$ was introduced in Definition 5.1.

Proof. By Theorem 5.7 the complexity of Algorithm 5.5 is bounded by

$$O\Big((B+S)\cdot(I(\mathcal{T}(Q),\mathcal{T}(R))+|H(\mathcal{T}(R))|+|Q|)\Big). \tag{5}$$

Therefore it suffices to bound the complexity B of the function FinalCandidates() and the complexity S of the function UpdateCandidates(). By Lemma 6.6 and Lemma 6.7 we have

$$B+S=O(\log(k)\cdot|R_i|)+O(c_m(R)^4\cdot\log(k)\cdot|R_i|)\leq O(c_m(R)^4\cdot\log(k)\cdot|R_i|).$$

Let us now bound the maximal size of the subset R_i. If $i=l_{\max}(\mathcal{T}(R))$ then $|R_i|=1$. If $i<l_{\max}(\mathcal{T}(R))$, it suffices to estimate the size $|R_{i-1}|$. Consider $R_{i-1}=\{a\in\mathcal{C}(R_i)\mid d(q,a)\leq d+2^{j+2}+2^{i+1}\}$, where $d=d(q,\lambda_k(q,\mathcal{C}(R_i)))$. Since $|R_i|$ can only change in the reference expansion lines 5-11 of Algorithm 5.5, we get $j+1\leq i$. Then $\mathcal{C}(R_i)\subseteq C_{i-1}$ gives

$$R_{i-1}=B(q,d+2^{i+1}+2^{j+2})\cap\mathcal{C}(R_i)\subseteq B(q,d+2^{i+2})\cap C_{i-1}. \tag{6}$$

Case $d\geq 2^{i+2}$. By Lemma 6.4 the set $\cup_{p\in R_i}\mathcal{S}_i(p,\mathcal{T}(R))$ contains all k-nearest neighbors of q. By Lemma 4.15 we have

$$\cup_{p\in\mathcal{C}(R_i)}\mathcal{S}_{i-1}(p,\mathcal{T}(R))=\cup_{p\in R_i}\mathcal{S}_i(p,\mathcal{T}(R)).$$

Therefore $\cup_{p\in\mathcal{C}(R_i)}\mathcal{S}_{i-1}(p,\mathcal{T}(R))$ contains all k-nearest neighbors of q.

Lemma 4.16 applied for the index $i-1$ says that the reference set R has a point β among all k-nearest neighbors of q such that $2^{i+2}\leq d\leq d(q,\beta)+2^i$. Then $3\cdot 2^i\leq d(q,\beta)$ and

$$d+2^{i+2}\leq d(q,\beta)+2^i+2^{i+2}\leq(1+\frac{5}{3})d(q,\beta)\leq 4d(q,\beta).$$

Combining the above estimate with inclusion (6) and Definition 3.4, we get

$$|R_{i-1}|\leq|B(q,d(q,\beta)+2^i+2^{i+2})|\leq|B(q,4d(q,\beta))|\leq c_{qr}^3|B(q,\frac{d(q,\beta)}{2})|.$$

If $|B(q,\frac{d(q,\beta)}{2})|\geq k$, then $B(q,\frac{d(q,\beta)}{2})$ contains k points α with $d(q,\alpha)\leq\frac{d(q,\beta)}{2}$. Then the point β cannot be among the first k nearest neighbors of q, which contradicts a choice of β above. Then $|B(q,\frac{d(q,\beta)}{2})|<k$, so $|R_{i-1}|\leq c_{qr}^3\cdot k$.

Case $d\leq 2^{i+2}$. Set $t=2^{i+3}$. Inclusion (6) implies that

$$R_{i-1}\subseteq B(q,d+2^{i+2})\subseteq B(q,t)\cap C_{i-1}.$$

By separation condition (3.1c) for $\mathcal{T}(R)$, all points in C_{i-1} are separated by $\delta=2^{i-1}$. Then $\frac{4t}{\delta}+1=2^6+1\leq 2^7$. Lemma 3.8 implies that $|R_{i-1}|\leq|B(p,t)\cap C_{i-1}|\leq(c_m(R))^7$.

The two cases above provide the upper bound $|R_{i-1}|\leq\max\{c_{qr}^3 k,(c_m(R))^7\}$ for any i.

Then $O\Big(B+S\Big)=O\Big(c_m(R)^4\cdot\log(k)\cdot|R_i|\Big)=O\Big(c_m(R)^4\cdot\max\{c_{qr}^3 k,(c_m(R))^7\}\cdot\log(k)\Big)$

is obtained from the time complexity in (5) after replacing $|R_{i-1}|$ by its upper bound. ◄

▶ **Corollary 6.9** (simplified time complexity for all k-nearest neighbors based on paired trees). *The time complexity of Algorithm 6.3 from Theorem 6.9 has the simpler upper bound*

$$O\Big(c_{qr}^{11}\cdot k\log(k)\cdot\big(|H(\mathcal{T}(R))|+|Q|+I(\mathcal{T}(R),\mathcal{T}(Q))\big)\Big), \tag{6.9}$$

see the height $|H(\mathcal{T}(R))|$ and the imbalance $I(\mathcal{T}(R),\mathcal{T}(Q))$ in Definitions 3.10 and 5.1.

7 Conclusions, discussion and further applications of the new results

This paper was motivated by the remaining challenges in the neighbor search based on paired trees, see Problem 1.2 in section 1. Section 2 described key steps in the past proofs that require more justifications, see detailed Counterexamples A.5 and A.8 in appendix A.

Main Theorem 6.8 has completely solved Problem 1.2 for any $k \geq 1$ in any metric space and also improved the recent time complexity based on a single tree in [10, Theorem 6.5].

In comparison with the 2015 estimate in [8, Theorem 2], Theorem 6.8 extended the time complexity to any number of neighbors $k \geq 1$. Even for $k = 1$, the complexity was improved due to the new height and imbalance. These parameters capture the hardness of the query and reference sets for neighbor search. If all parameters have constant bounds, the final complexity $O(k \log k \max\{|Q|, |R|\})$ is near linear with respect to the key variables.

As a practical application, the time complexities in [10, Theorem 6.5] and Theorem 6.8 helped prove a near linear time complexity for new invariants of crystals in [18, Theorem 14]. These invariants are formed by distances from atoms to their nearest neighbors in an infinite periodic crystal. A modest desktop over a couple of days completed 400M+ pairwise comparisons of the above neighbor-based invariants for all 660K+ periodic crystals in the world's largest Cambridge Structural Database. The resulting five pairs of impossible duplicates are now investigated by the journals that published the underlying papers.

One can extend the developed methods in several directions. Similarly to [10, Theorem 7.4], we conjecture that the $(1 + \epsilon)$-approximate version of Problem 1.2 can be solved in time $O\Big(c_m(R)^{\log_2(O(1/\epsilon))} \log(k) \cdot (I(\mathcal{T}(R), \mathcal{T}(Q)) + \log_2(\Delta) + |Q|) + k\Big)$, where $c_m(R)$ is the minimized expansion constant from Definition 3.4, Δ is the aspect ratio of R, see Definition 3.11.

[4, Section 2.3.1] hints at a possibility to improve [10, Algorithm 4.3] for a compressed cover tree to reduce the complexity $O(c_m(R)^{O(1)} \cdot |R| \cdot |H(\mathcal{T}(R))|)$ to $O(c_m(R)^{O(1)} \cdot I(\mathcal{T}(R), \mathcal{T}(R))$. The time complexity for a Minimum Spanning Tree [15] estimates the number of recursions with the same issues as in the proof of [5, Theorem 5] discussed in section 2. The challenging datasets from Counterexamples A.5 and A.8 in appendix A imply that the proof of [15, Theorem 5.1] can be improved by using the novel concept of the imbalance of paired trees.

The authors are grateful to all reviewers for their valuable time and comments in advance.

References

1   Alexandr Andoni, Piotr Indyk, and Ilya Razenshteyn. Approximate nearest neighbor search in high dimensions. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 3287–3318. World Scientific, 2018.

2   Sunil Arya and David M Mount. Approximate nearest neighbor queries in fixed dimensions. In *SODA*, volume 93, pages 271–280. Citeseer, 1993.

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

4   Alina Beygelzimer, Sham Kakade, and John Langford. Extended version of "cover trees for nearest neighbor". URL: `https://hunch.net/~jl/projects/cover_tree/paper/paper.pdf`.

5   Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of ICML*, pages 97–104, 2006.

6   Thomas Cormen. *Introduction to algorithms*. MIT Press McGraw-Hill, Cambridge, Mass. New York, 1990.

7   Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.

8   Ryan R Curtin, Dongryeol Lee, William B March, and Parikshit Ram. Plug-and-play dual-tree algorithm runtime analysis. *J. Mach. Learn. Res.*, 16:3269–3297, 2015.

9   Ryan Ross Curtin. *Improving dual-tree algorithms*. PhD thesis, Georgia Institute of Technology, 2015.

10   Yury Elkin and Vitaliy Kurlin. A new compressed cover tree guarantees a near linear parameterized complexity for all k-nearest neighbors search in metric spaces, 2021. `arXiv:arXiv:2111.15478`.

11   Raphael A Finkel and Jon Louis Bentley. Quad trees a data structure for retrieval on composite keys. *Acta informatica*, 4(1):1–9, 1974.

12   Alexander G Gray and Andrew W Moore. N-body'problems in statistical learning. *Advances in neural information processing systems*, pages 521–527, 2001.

13   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. Citeseer, 2004.

14   S Manocha and Mark A Girolami. An empirical analysis of the probabilistic k-nearest neighbour classifier. *Pattern Recognition Letters*, 28(13):1818–1824, 2007.

15   W. March, P. Ram, and A. Gray. Fast euclidean minimum spanning tree: algorithm, analysis, and applications. In *Proceedings of SIGKDD*, pages 603–612, 2010.

16   Parikshit Ram, Dongryeol Lee, William March, and Alexander Gray. Linear-time algorithms for pairwise statistical problems. *Advances in Neural Information Processing Systems*, 22:1527–1535, 2009.

17   Mengzhao Wang, Xiaoliang Xu, Qiang Yue, and Yuxiang Wang. A comprehensive survey and experimental comparison of graph-based approximate nearest neighbor search. *arXiv preprint arXiv:2101.12631*, 2021.

18   Daniel Widdowson, Marco Mosca, Angeles Pulido, Vitaliy Kurlin, and Andrew Cooper. Average minimum distances of periodic point sets. *MATCH Communications in Mathematical and in Computer Chemistry, to appear*. URL: `https://arxiv.org/abs/2009.02488`.

A   **Challenging data for a nearest neighbor search based on paired trees**

This section discusses k-nearest neighbor Algorithm A.1 based on paired trees from [16].

Counterexample A.5 shows how Algorithm A.1 cannot be used for finding non-trivial nearest neighbors in the case $Q = R$. Counterexample A.8 provides a cover tree that should, according to the proof of [16, Theorem 3.1], run at most $O(\max\{|Q|, |R|\} \cdot \sqrt{\max\{|Q|, |R|\}})$ reference expansions (lines 5-8) during the whole execution of Algorithm A.1.

The step-by-step execution of Algorithm A.1 will show that the number of reference expansions has a lower bound $O(\max\{|Q|, |R|\}^2)$. These issues are resolved in sections 5 and 6 by rewriting the algorithm for compressed cover trees and by using the new parameters: the height of a compressed cover tree in Definition 3.10 and the imbalance in Definition 5.1.

Recall that [16, End of Section 1] defined the all-nearest-neighbor problem as follows. "**All Nearest-neighbors:** For all queries $q \in Q$ find $r^*(q) \in R$ such that $r^*(q) = \operatorname{argmin}_{r \in R} d(q, r)$". For $Q = R$, the last formula produces trivial self-neighbors.

1: **Function** FindAllNN(a node $q_j \in T(Q)$, a subset R_i of a cover set C_i of $T(R)$).
2: **if** $i = -\infty$ **then**
3:     for each $q_j \in L(q_j)$ **return** $\operatorname{argmin}_{r \in R_{-\infty}} d(q, r)$
4:     {here $L(q_j)$ is the set of all descendants of the node q_j}
5: **else if** $j < i$ **then**
6:     $\mathcal{C}(R_i) = \{\text{Children}(r) \mid r \in R_i\}$ {in original pseudo-code the notation is $R = \mathcal{C}(R_i)$}
7:     $R_{i-1} = \{r \in R \mid d(q_j, r) \le d(q_j, R) + 2^i + 2^{j+2}\}$
8:     FindAllNN(q_{j-1}, R_i) { q_{j-1} is the same point as q_j on one level below}
9: **else**
10:     for each $p_{j-1} \in \text{Children}(q_j)$ FindAllNN(p_{j-1}, R_i)
11: **end if**

■ **Algorithm A.1**   Original [16, Algorithm 1] is analyzed in Counterexamples A.5 and A.8.

In original Algorithm A.1, the node q_j has a level j, a reference subset $R_i \subset R$ is a subset of C_i for an implicit cover tree $T(R)$. The algorithm is called for a pair $q_j, R_i = \{r\}$, where q_j is the root of the query tree at the maximal level $j = l_{\max}(T(Q))$, and r is the root of the reference tree at the maximal level $i = l_{\max}(T(R))$. Split Algorithm A.1 into these blocks:

lines 2-4 : FinalCandidates,

lines 5-9 : reference expansion,

lines 9-11 : query expansion.

▶ **Example A.2** (tall imbalanced tree [10, Example 3.1]). For any integer $m > 10$, let G be a metric graph pictured in Figure 7 that has two vertices r, q and $m + 1$ edges (e_i) for $i \in \{0, ..., m\}$, and the length of each edge e_i is $|e_i| = 2^{m \cdot i + 2}$ for $i \ge 1$. Finally, set $|e_0| = 1$.

For every $i \in \{1, ..., m^2\}$ if i is divisible by m we set p_i be the middle point of $e_{i/m}$ and for every other i we define p_i to be the middle point of segment (p_{i+1}, q).

Let d be the induced shortest path metric on the continuous graph G. Then $d(q, r) = 1$, $d(r, p_i) = 2^{i+1} + 1$, $d(q, p_i) = 2^i$ and if $i > j$ and $\lceil \frac{i}{m} \rceil = \lceil \frac{j}{m} \rceil$ we have $d(p_j, p_i) = \sum_{t=j+1}^{i} 2^t$. We consider the reference set $R = \{r\} \cup \{p_i \mid i \in \{1, 2, 3, ..., m^2\}\}$ with the metric d.

Let us define a compressed cover tree $\mathcal{T}(R)$ by setting r to be the root node and $l(p_i) = i$ for all i. If i is divisible by m, we set r to be the parent of p_i. If i is not divisible by m, we set p_{i+1} to be the parent of p_i. For every i divisible by m, the point p_i is in the middle of edge $e_{i/m}$, hence $d(p_i, r) \le 2^{i+1}$.

For every i not divisible by m, by the definition p_i is middle point of (p_{i+1}, q) and therefore we have $d(p_i, p_{i+1}) \le 2^{i+1}$. Since for any point p_i distance to its parent is at most 2^{i+1}, it follows that $\mathcal{T}(R)$ satisfies the covering condition of Definition 3.1.

For any integer t, the cover set is $C_t = \{r\} \cup \{p_i \mid i \ge t\}$. We will prove that C_t satisfies the separation condition. Let $p_i \in C_t$. If i is divisible by m, then $d(r, p_i) = 2^{i+1} \ge 2^{t+1} > 2^t$. If i is not divisible by m, then $d(r, p_i) = d(r, q) + d(q, p_i) = 1 + 2^{i+1} > 2^t$.

Then the root r is separated from the other points by the distance 2^t. Consider arbitrary points p_i and p_j with indices $i > j \ge t$ and $\lceil \frac{i}{m} \rceil = \lceil \frac{j}{m} \rceil$. Then

$$d(p_i, p_j) = \sum_{s=j+1}^{i} 2^s \ge 2^{j+1} \ge 2^{t+1} > 2^t.$$

On the other hand, if $i > j \geq t$ and $\lceil \frac{i}{m} \rceil \neq \lceil \frac{j}{m} \rceil$, then

$$d(p_i, p_j) = d(p_i, q) + d(p_j, q) \geq 2^i + 2^j \geq 2^{j+1} \geq 2^{t+1} > 2^t.$$

For any t, we have shown that all pairwise combinations of points of C_t satisfy the separation condition in Definition 3.1. Hence this condition holds for the whole tree $\mathcal{T}(R)$. ∎

Recall that in [5, Section 2] the explicit representation of cover tree was defined as "the explicit representation of the tree coalesces all nodes in which the only child is a self-child". Simplest way to interpret this is to consider cover sets C_i and define $p \in C_i$ to be an explicit node, if p has child at level $i - 1$. By [5, Lemma 4.3] the depth of any node p is "defined as the number of explicit grandparent nodes on the path from the root to p in the lowest level in which p is explicit". The explicit depth of a node p in any compressed tree $\mathcal{T}$ will be defined in Definition A.3 using the simplest interpretation of the quotes above.

▶ **Lemma A.3** (explicit depth for a compressed cover tree [10, Lemma 3.2])**.** *Let R be a finite subset of a metric space with a metric d. Let $\mathcal{T}(R)$ be a compressed cover tree on R. For any $p \in \mathcal{T}(R)$, let $s = (w_0, ..., w_m)$ be a node-to-root path of p. Then the explicit depth $D(p)$ of node p belonging to compressed cover tree can be interpreted as the sum*

$$D(p) = \sum_{i=0}^{m-1} |\{q \in \mathrm{Children}(w_{i+1}) \mid l(q) \in [l(w_i), l(w_{i+1}) - 1]\}|.$$

∎

Proof. Note that the node-to-root path of an implicit cover tree on R has $l(w_{j+1}) - l(w_j) - 1$ extra copies of w_{j+1} between every w_j and w_{j+1} for any index $j \in [0, m - 1]$. Recall that a node is called explicit, if it has non-trivial children. Therefore there will be exactly

$$|\{q \in \mathrm{Children}(w_{i+1}) \mid l(q) \in [l(w_i), l(w_{i+1}) - 1]\}|$$

explicit nodes between w_j and w_{j+1}. It remains to take the total sum. ◀

▶ **Lemma A.4** ([10, Lemma 3.3])**.** *Let $\mathcal{T}(R)$ be a compressed cover tree on the set R from Example A.2 for some $m \in \mathbb{Z}$. For any $p \in R$, the explicit depth $D(p)$ of Definition A.3 has the upper bound $2m + 1$.* ∎

Proof. For any p_i, if i is divisible by m, then r is the parent of p_i. By definition, the explicit depth is $D(p_i) = |\{p \in \mathrm{Children}(r) \mid l(p) \in [l(p_i), m^2]\}|$. Since r contains children on every level j, where j is divisible by m, we have $D(p_i) = m - \frac{i}{m} + 1$.

Let us now consider an index i that is not divisible by m. Note that p_{j+1} is the parent of p_j for all $j \in [i, m \cdot \lceil i/m \rceil - 1]$. Then the path consisting of all ancestors of p_i from p_i to the root node r has the form $(p_i, p_{i+1}, ..., p_{m \cdot \lceil i/m \rceil}, r)$. It follows that

$$D(p_i) = \sum_{j=i}^{m \cdot \lceil i/m \rceil - 1} |\{p \in \mathrm{Children}(p_j) \mid l(p) \in [l(p_j), l(p_{j+1}) - 1]\}| + D(p_{m \cdot \lceil i/m \rceil}).$$

Since $i \geq m \cdot (\lceil i/m \rceil - 1) + 1$ and $\lceil i/m \rceil \geq 1$, we get the required upper bound:

$$D(p_i) = (m \cdot \lceil i/m \rceil - i) + (m - \lceil i/m \rceil + 1) \leq m + (m + 1) = 2m + 1.$$

◀
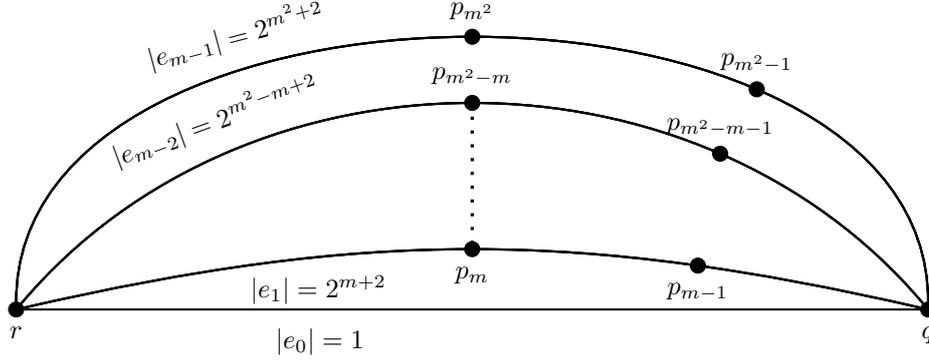
Figure 7 Illustration of a graph G and a point cloud R defined in Example A.2

▶ **Counterexample A.5.** In the notations of ExampleA.2, m is a parameter of R. Build a compressed cover tree $\mathcal{T}(R)$ as in Figure 8. Set $Q = R$. First we show that Algorithm A.1 returns the trivial neighbor for every point for $\mathcal{T}(Q) = \mathcal{T}(R)$.

We start the simulation with the query node r on the level $m^2 + 1$, which has the reference subset $R_{m^2+1} = \{r\}$. The query node and the reference set are at the same levels, so we run the query expansions (lines 9-11). The node r has p_{m^2} and r as its children. Hence the algorithm goes into the branches that have p_{m^2} as the query node and into the branches that have r as the query node.

Let us focus on all recursions having p_{m^2} as the query node. In the first recursion involving the node p_{m^2}, we have $i = m^2 + 1, j = m^2$. Thus $j < i$ and we run reference expansions (lines 5-9). The node r has two children at the level m^2, so $\mathcal{C}(R_i) = \{p_{m^2}, r\}$. Since $d(p_{m^2}, p_{m^2}) = 0$ and $d(p_{m^2}, r) = 2^{m^2+1}$ on line 7, we have:

$$R_{m^2} = \{r \in \mathcal{C}(R_i) \mid d(q_j, r) \leq 2^{m^2+1} + 2^{m^2+2}\} = \{p_{m^2}, r\}.$$

Similarly, for $i = m^2, j = m^2 - 1, q_j = p_{m^2}$, we have $\mathcal{C}(R_i) = \{p_{m^2}, p_{m^2-1}, r\}$ and since $d(p_{m^2-1}, p_{m^2}) = 2^{m^2}$ and $d(r, p_{m^2}) = 2^{m^2+1}$ we have:

$$R_{m^2-1} = \{r \in \mathcal{C}(R_i) \mid d(q_j, r) \leq 2^{m^2} + 2^{m^2+1}\} = \{p_{m^2}, p_{m^2-1}\}.$$

For $i = m^2 - 1, j = m^2 - 2, q_j = p_{m^2}$, we have $\mathcal{C}(R_i) = \{p_{m^2}, p_{m^2-1}, p_{m^2-2}\}$. Since $d(p_{m^2-1}, p_{m^2}) = 2^{m^2}$ and $d(p_{m^2-2}, p_{m^2}) = 2^{m^2} + 2^{m^2-1}$, we have:

$$R_{m^2-2} = \{r \in \mathcal{C}(R_i) \mid d(q_j, r) \leq 2^{m^2-1} + 2^{m^2}\} = \{p_{m^2}, p_{m^2-1}, p_{m^2-2}\}.$$

Finally, for $i = m^2 - 2, j = m^2 - 3, q_j = p_{m^2}$, we have $\mathcal{C}(R_i) = \{p_{m^2}, p_{m^2-1}, p_{m^2-2}, p_{m^3-3}\}$ and $d(p_{m^2}, p_{m^3-3}) = 2^{m^2} + 2^{m^2-1} + 2^{m^2-2}$. The previous inequalities imply that

$$R_{m^2-3} = \{r \in \mathcal{C}(R_i) \mid d(q_j, r) \leq 2^{m^2-2} + 2^{m^2-1}\} = \{p_{m^2}\}.$$

Since $R_{m^2-3} = \{p_{m^2}\}$, the nearest neighbor of p_{m^2} will be chosen to be p_{m^2}. The same argument can be repeated for all $p_t \in R$. It follows that Algorithm A.1 finds trivial nearest neighbor for every point $p_t \in R$. ■

▶ **Example A.6.** To avoid the issue of finding trivial nearest neighbors in Counterexample A.5, we will modify Example A.2. For any integer $m > 100$, let G be a metric graph that has 2

Level $m^2 + 1$ — r

Level m^2 — p_{m^2}

Level $m^2 - 1$ — p_{m^2-1}

Level $m \cdot (m - 1) + 1$ — p_{m^2-m+1}

Level $i \cdot m$ — $p_{i \cdot m}$

Level $i \cdot m - 1$ — $p_{i \cdot m-1}$

Level $m \cdot (i - 1) + 1$ — $p_{i \cdot m-m+1}$

Level m — p_m

Level $m - 1$ — p_{m-1}

Level 1 — p_1
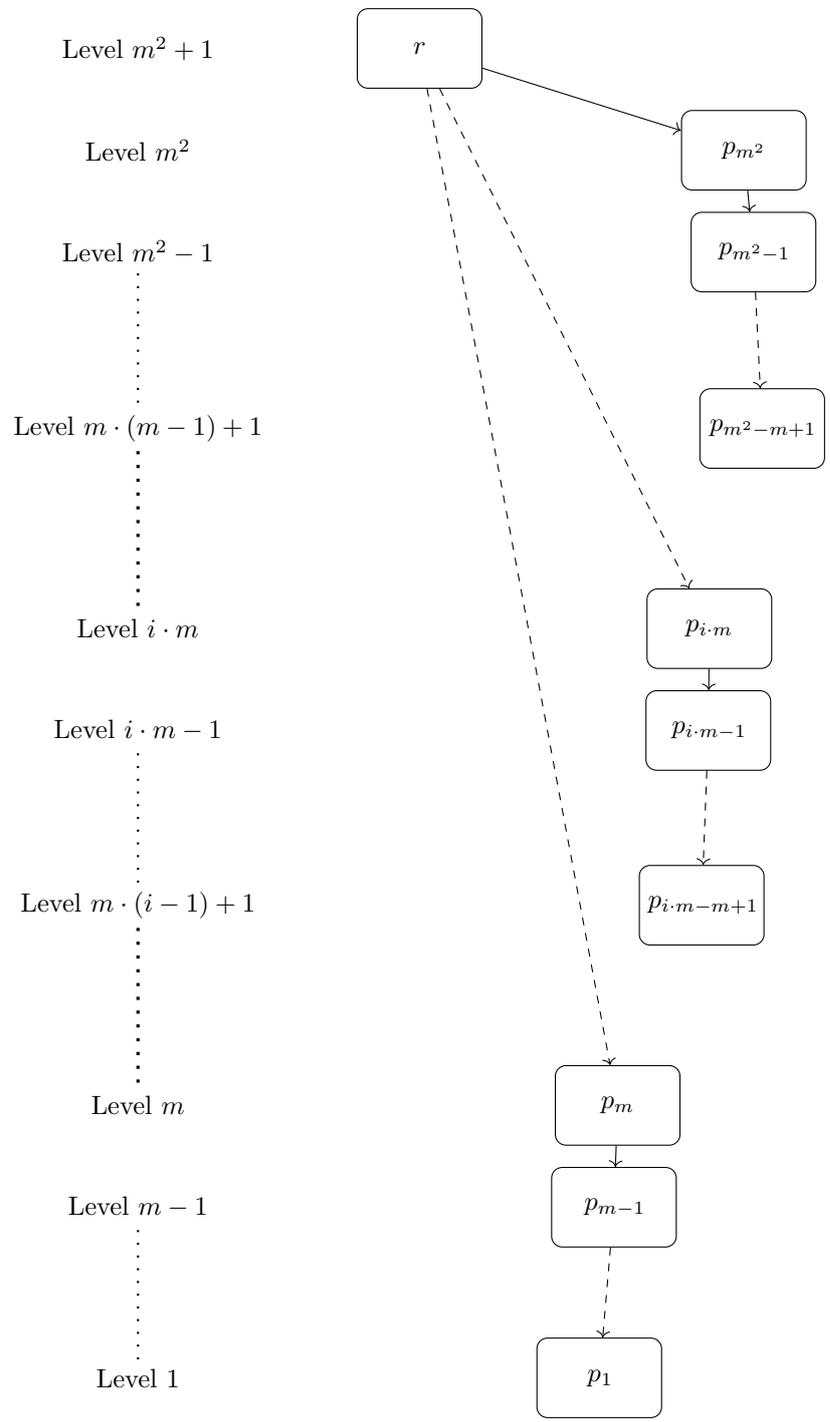
Figure 8 Illustration of the a compressed cover tree $\mathcal{T}(R)$ defined in Example A.2

vertices r and q and $2m - 1$ edges $\{e_0\} \cup \{e_1, ..., e_{m-1}, h_1, ..., h_{m-1}\}$. The edge-lengths are $|e_i| = 2^{i \cdot m + 2}$ and $|h_i| = 2^{i \cdot m + 2}$ for all $i \in [1, m]$, finally $|e_0| = 1$.

For every $i \in \{1, ..., m^2\}$, if i is divisible by m, we set q_i to be the middle point of $e_{i/m}$ and r_i to be the middle point of $h_{i/m}$. For every other i not divisible by m, we define q_i to be the middle point of segment (q_{i+1}, q) and r_i to be the middle point of segment (r_{i+1}, q).

Let d be the shortest path metric on the graph G. Then $d(q_i, r) = d(q_i, q) + 1 = 2^{i+1} + 1$, $d(q_i, r_j) = 2^{i+1} + 2^{j+1}$ and $d(q, r) = 1$. Let $R = \{r, r_{m^2}, r_{m^2-1}, ..., r_1\}$ and let $Q = \{r, q_{m^2}, q_{m^2-1}, ..., q_1\}$. Let compressed cover trees $\mathcal{T}(Q), \mathcal{T}(R)$ have the same structure as the compressed cover tree $\mathcal{T}(R)$ in Example A.2. ∎

▶ **Remark A.7.** [16, Definition 3.1] introduced the degree of bichromaticity κ as follows.

"**Definition 3.1** Let S and T be cover trees built on query set Q and reference set R respectively. Consider a dual-tree algorithm with the property that the scales of S and T are kept as close as possible – i.e. the tree with the larger scale is always descended. Then, the degree of bichromaticity κ of the query-reference pair (Q, R) is the maximum number of descends in S between any two descends in T".

There are at least two different interpretation of this definition. Our best interpretation is that κ the maximal number of levels in T containing at least one node between any two consecutive levels of S. However, if q is a leaf node of S, but there are still many levels between level of q and $l_{\min}(T)$, it is not clear from the definition if κ includes these levels.

[8, page 3284] pointed out that " Our results are similar to that of Ram et al. (2009a), but those results depend on a quantity called the constant of bichromaticity, denoted κ, which has unclear relation to cover tree imbalance. The dependence on κ is given as $c_q^{4\kappa}$, which is not a good bound, especially because κ may be much greater than 1 in the bichromatic case (where $S_q = S_r$)".

To keep track of the indices i, j the function call FindAllNN(q_j, R_i) will be expressed as FindAllNN(i, j, q_j, R_i) in Counterexample A.5.

▶ **Counterexample A.8.** We will now show that in addition to the problems in the pseudocode the proof of [16, Theorem 3.1] is incorrect. Let us consider the following quote from its proof.

"**Theorem 3.1** Given a reference set R of size N and expansion constant c_R, a query set Q of size $O(N)$ and expansion constant c_Q, and bounded degree of bichromaticity κ of the (Q, R) pair, the FindAllNN subroutine of Algorithm 1 computes the nearest neighbor in R of each point in Q in $O(c_R^{12} c_Q^{4\kappa} N)$ time

[*Partial proof:*] Since at any level of recursion, the size of R [Corresponding to $\mathcal{C}(R_i)$ in Algorithm A.1] is bounded by $c_R^4 \max_i R_i$ (width bound), and the maximum depth of any point in the explicit tree is $O(c_R^2 \log(N))$ (depth bound), the number of nodes encountered in Line 6 is $O(c_R^6 \max_i |R_i| \log(N))$. Since the traversal down the query tree causes duplication, and the duplication of any reference node is upper bounded by $c_Q^{4\kappa}$, Line 6 [corresponds to line 8 in Algorithm A.1] takes at most $c_Q^{4\kappa} c_R^6 \max_i |R_i| \log(N)$ in the whole algorithm. "

The above arguments claimed the algorithm runs Line 8 at most this number of times:

$$\#(\text{Line 8}) \leq \max_{p \in R} D(p) \cdot \max_i \mathcal{C}(R_i) \cdot (\text{number of duplications}). \tag{7}$$

Let $X, \mathcal{T}(R), \mathcal{T}(Q), R, Q$ be as in Example A.6 for some parameter m. We will consider the simulation of Algorithm A.1 on pair $(\mathcal{T}(Q), \mathcal{T}(R))$. We note first Lemma A.4 applied on $\mathcal{T}(R)$ provides $\max_{p \in R} D(p) \leq 2m + 1$ Similarly as in [10, Counterexample 3.7], a contradiction will be achieved by showing that R_i and a set of its children $\mathcal{C}(R_i)$ will have constant size bound on any recursion (i, j) of Algorithm A.1.

Since $\mathcal{T}(R)$ contains at most one children on every level i we have $|\mathcal{C}(R_i)| \leq |R_i| + 1$ for any recursion of FindAllNN algorithm. For any $i > m^2$ denote r_i and q_i to be r. Note first that since $l(q_t) = t$ for any $t \in [1, m^2]$, then q_t is recursed into from FindAllNN$(t + 1, t + 1, p, R_i)$, where p is parent node of q_t. Therefore it follows that $t \geq i + 1$ in any stage of the recursion. Let us prove that for any $i \in [1, m^2 + 1]$ following two claims hold: (1) Function FindAllNN$(i, j = i - 1, q_t, R_i)$ is called for all $t \geq i + 1$ and (2) We have $R_i = \{r_{i+1}, r_i, r\}$ in this stage of the algorithm. The claim will be proved by induction on i. Let us first prove case $i = 2m + 1$. Note that Algorithm A.1 is originally launched from FindAllNN$(2m + 1, 2m + 1, r, \{r\})$, therefore the first claim holds. Second claim holds trivially since $r_{2m+2} = r$ and $r_{2m+1} = r$.

Assume now that the claim holds for some i, let us show that the claim will always hold for $i - 1$. Assume that FindAllNN$(i, j = i - 1, q_t, R_i)$ was called for some $t \geq i + 1$. Since $j = i - 1$, we perform a reference expansion (lines 5-9). By line 6 and induction assumption we have $\mathcal{C}(R_i) = \{r, r_{i+1}, r_i, r_{i-1}\}$. Assume first that $q_t = r$. Recall that for any $u \in [1, m^2]$ we have $d(r, r_u) \geq 2^{u+1}$. It follows that

$$R_{i-1} = \{r' \in \mathcal{C}(R_i) \mid d(r, r') \leq 2^i + 2^{i+1}\} = \{r, r_i, r_{i-1}\}.$$

Let us now consider case $q_t \neq r$. We have $d(r, q_t) = 2^{t+1}$ and $d(q_t, r_{u+1}) = 2^{t+1} + 2^{u+2}$ for any $u \in [1, m^2 + 1]$. Therefore

$$R_{i-1} = \{r' \in C_{i-1} \mid d(q_t, r') \leq d(q_t, r) + 2^i + 2^{i+1} \leq 2^{t+1} + 2^i + 2^{i+1}\}.$$

It follows that $R_{i-1} = \{r, r_i, r_{i-1}\}$. In both cases we proceed to line 8 where we launch FindAllNN$(i - 1, i - 1, q_t, R_{i-1})$. After proceeding into the recursion we have $j = i$ and therefore query-expansion (lines 9-11) will be performed. Note that q_t was chosen so that $t \geq i + 1$. Since every q_{t-1} is either a child of r or q_t it follows that FindAllNN$(i - 1, i - 2, q_{t'}, R_{i-1})$ will be called for all $t' \geq t - 1 \geq i$. Then condition (2) of the induction claim holds as well.

It remains to show that Algorithm A.1 $(q, R_i = \{r\})$ has $O(m^4)$ low bound on the number of times reference expansions (lines 5-9) are performed. Let ξ be the number of times Algorithm A.1 performs reference expansions. For every $q' \in Q$ denote $\xi(q')$ to be the total number of reference expansions performed for q'. Recall that any query node q' is introduced in the query expansion (lines 9-11) for parameters $(i = u + 1, j = u + 1, p, R_i)$, where p is the parent node of q'. Since R_i is non empty for all levels $[1, u]$ we have $\xi(q_u) \geq u - 1$ for all u. Thus

$$\xi = \sum_{q' \in Q} \xi(q') \geq \sum_{u=2}^{m^2+1} u - 2 = O(m^4).$$

There are different interpretations for the number of duplications. Note that the query tree $\mathcal{T}(Q)$ has exactly one new child on every level and that trees $\mathcal{T}(Q)$ and $\mathcal{T}(R)$ contain exactly the same levels. By using the definitions the number of duplications should be 2. However, since there can be other interpretations for the number of duplications, we make a rough estimate that the number of duplication is upper bounded by the number of nodes in query tree $O(m^2)$. By using Inequality (7), we obtain the following contradiction:

$$O(m^4) = \xi \leq \max_{p \in R} D(p) \cdot (\max_i \mathcal{C}(R_i)) \cdot (\text{number of duplications}) \leq (2m + 1) \cdot 4 \cdot m^2 \leq O(m^3).$$

B   Detailed proofs of all auxiliary lemmas from sections 3 and 4

All lemmas from sections 3 and 4 are re-stated below (with the same numbers) for convenience.

▶ **Lemma 3.5** (properties of c_m [10, Lemma 2.5]). *For any finite sets $R \subseteq U$ in a metric space, we have $c_m(R) \leq c_m(U)$, $c_m(R) \leq c(R)$.*

Proof. The proof follows trivially from Definition 3.4. ◀

▶ **Lemma 3.7** (a distance bound on descendants [10, Lemma 2.8]). *Let R be a finite set with a metric d. In a compressed cover tree $\mathcal{T}(R)$, let q be any descendant of a node p. Let the node-to-root path of q have a node $u \in \text{Children}(p) \setminus \{p\}$. Then $d(p,q) \leq 2^{l(u)+2} \leq 2^{l(p)+1}$.* ■

Proof. Let $(w_0, ..., w_m)$ be a subpath of the node-to-root path for $w_0 = q$, $w_{m-1} = u$, $w_m = p$. Then $d(w_i, w_{i+1}) \leq 2^{l(w_i)+1}$ for any i. The first required inequality follows from the triangle inequality below:

$$d(p,q) \leq \sum_{j=0}^{m-1} d(w_j, w_{j+1}) \leq \sum_{j=0}^{m-1} 2^{l(w_j)+1} \leq \sum_{t=l_{\min}}^{l(u)+1} 2^t \leq 2^{l(u)+2}$$

Finally, $l(u) \leq l(p) - 1$ implies that $d(p,q) \leq 2^{l(p)+1}$. ◀

▶ **Lemma 3.8** (packing [10, Lemma 2.9]). *Let S be a finite δ-sparse set in a metric space (X, d), so $d(a, b) > \delta$ for all $a, b \in S$. Then, for any point $p \in X$ and any radius $t > \delta$, we have $|\bar{B}(p,t) \cap S| \leq (c_m(S))^\mu$, where $\mu = \lceil \log_2(\frac{4t}{\delta} + 1) \rceil$.* ■

Proof. Assume that $d(p, q) > t$ for any point $q \in S$. Then $\bar{B}(p,t) \cap S = \emptyset$ and the lemma holds trivially. Otherwise $\bar{B}(p,t) \cap S$ is non-empty. By Definition 3.4 of a minimized expansion constant, for any $\epsilon > 0$, we can always find a set A such that $S \subseteq A \subseteq X$ and

$$|B(q, 2s) \cap A| \leq (c_m(S) + \epsilon) \cdot |B(q, s) \cap A| \tag{8}$$

for any $q \in A$ and $s \in \mathbb{R}$. Note that for any $u \in \bar{B}(p,t) \cap S$ we have $\bar{B}(u, \frac{\delta}{2}) \subseteq \bar{B}(u, t + \frac{\delta}{2})$. Therefore, for any point $q \in \bar{B}(p,t) \cap S$, we get

$$\bigcup_{u \in \bar{B}(p,t) \cap S} \bar{B}(u, \frac{\delta}{2}) \subseteq \bar{B}(p, t + \frac{\delta}{2}) \subseteq \bar{B}(q, 2t + \frac{\delta}{2})$$

Since all the points of S were separated by δ, we have

$$|\bar{B}(p,t) \cap S| \cdot \min_{u \in \bar{B}(p,t) \cap S} |\bar{B}(u, \frac{\delta}{2}) \cap A| \leq \sum_{u \in \bar{B}(p,t) \cap S} |\bar{B}(u, \frac{\delta}{2}) \cap A| \leq |\bar{B}(q, 2t + \frac{\delta}{2}) \cap A|$$

In particular, by setting $q = \text{argmin}_{a \in S \cap \bar{B}(p,t)} |\bar{B}(a, \frac{\delta}{2})|$, we get:

$$|\bar{B}(p,t) \cap S| \cdot |\bar{B}(q, \frac{\delta}{2}) \cap A| \leq |\bar{B}(q, 2t + \frac{\delta}{2}) \cap A| \tag{9}$$

Inequality (8) applied μ times for the radii $s_i = \dfrac{2t + \frac{\delta}{2}}{2^i}$, $i = 1, ..., \mu$, implies that:

$$|\bar{B}(q, 2t + \frac{\delta}{2}) \cap A| \leq (c_m(S) + \epsilon)^\mu |\bar{B}(q, \frac{2t + \frac{\delta}{2}}{2^\mu}) \cap A| \leq (c_m(S) + \epsilon)^\mu |\bar{B}(q, \frac{\delta}{2}) \cap A|. \tag{10}$$

By combining inequalities (9) and (10), we get

$$|\bar{B}(p,t) \cap S| \leq \frac{|\bar{B}(q, 2t + \frac{\delta}{2}) \cap A|}{|\bar{B}(q, \frac{\delta}{2}) \cap A|} \leq (c_m(S) + \epsilon)^{\mu}.$$

The required inequality is obtained by letting $\epsilon \to 0$. ◄

▶ **Lemma 3.9** (width bound [10, Lemma 2.10]). *Let R be a finite subset of a metric space (X,d). For any compressed cover tree $\mathcal{T}(R)$, any node p has at most $(c_m(R))^4$ children at every level i, where $c_m(R)$ is the minimized expansion constant of the set R.* ■

Proof. By the covering condition in Definition 3.1 for a compressed cover tree $\mathcal{T}(R)$, any child q of a node p located on the level i has $d(q,p) \leq 2^{i+1}$. Thus the number of children of the node p at the level i at most $|\bar{B}(p, 2^{i+1})|$.

The separation condition in Definition 3.1 implies that the cover set C_i is 2^i-sparse in X. We can now apply Lemma 3.8 for $t = 2^{i+1}$ and $\delta = 2^i$. Since $4 \cdot \frac{t}{\delta} + 1 \leq 4 \cdot 2 + 1 \leq 2^4$, we get $|\bar{B}(q, 2^{i+1}) \cap C_i| \leq (c_m(C_i))^4$.

The lemma follows by noting that Lemma 3.5 implies $(c_m(C_i))^4 \leq (c_m(R))^4$. ◄

▶ **Lemma 3.12** (an upper bound for the height $|H(\mathcal{T}(R))|$, see [10, Lemma 2.13]). *Any finite set R has the upper bound $|H(\mathcal{T}(R))| \leq 1 + \log_2(\Delta(R))$.* ■

Proof. $|H(\mathcal{T}(R))| \leq l_{\max} - l_{\min} + 1$ by Definition 3.10. We estimate $l_{\max} - l_{\min}$ as follows.

Let $p \in R$ be a point with $\mathrm{diam}(R) = \max_{q \in R} d(p,q)$. Then R is covered by the closed ball $\bar{B}(p; \mathrm{diam}(R))$. Hence the cover set C_i at the level $i = \log_2(\mathrm{diam}(R))$ consists of a single point p. The separation condition in Definition 3.1 implies that $l_{\max} \leq \log_2(d_{\max}(R))$.

Since any distinct points $p, q \in R$ have a distance $d(p,q) \geq d_{\min}(R)$, the covering condition implies that no new points can enter the cover set C_i at the level $i = [\log_2(d_{\min}(R))]$, so $l_{\min} \geq \log_2(d_{\min}(R))$. So $|H(\mathcal{T}(R))| \leq 1 + l_{\max} - l_{\min} \leq 1 + \log_2 \frac{\mathrm{diam}(R)}{d_{\min}(R)}$. ◄

▶ **Lemma 4.7** ([10, Lemma 5.3]). *Let R be a finite subset of a metric space. Let $\mathcal{T}(R)$ be a compressed cover tree on R. Then, Algorithm 4.6 computes the sizes $|\mathcal{S}_i(p, \mathcal{T}(R))|$ for all $p \in R$ and essential levels $i \in \mathcal{E}(p, \mathcal{T}(R))$ in time $O(|R|)$.* ■

Proof. By Lemma 4.5 we have $\sum_{p \in R} |\mathcal{E}(p, \mathcal{T}(R))| \leq 2 \cdot |R|$. Since CountDistinctiveDescendants is called once for every any combination $p \in R$ and $i \in \mathcal{E}(p, \mathcal{T}(R))$ it follows that the time complexity of Algorithm 4.6 is $O(R)$. ◄

▶ **Lemma 4.10** (time complexity of a λ-point [10, Lemma 5.7]). *In the notations of Definition 4.8, the time complexity of Algorithm 4.9 is $O(|C| \log(k) \max_i |R_i|)$.* ■

Proof. To compute $\lambda = \lambda_k(q, R)$ in Algorithm 4.9, we need to select k elements by using an ordering from the set C. This selection takes takes at most $|C| \cdot \log(k)$ time by using a binary heap data structure [6, section 6.5]. ◄

▶ **Lemma 4.11** (separation lemma [10, Lemma 5.8]). *In the conditions of Definition 4.1, let $p \neq q$ be nodes of $\mathcal{T}(R)$ with $l(p) \geq i$, $l(q) \geq i$. Then $\mathcal{S}_i(p, \mathcal{T}(R)) \cap \mathcal{S}_i(q, \mathcal{T}(R)) = \emptyset$.* ■

Proof. Without loss of generality assume that $l(p) \geq l(q)$. If q is not a descendant of p, the lemma holds trivially due to Descendants$(q) \cap$ Descendants$(p) = \emptyset$.

If q is a descendant of p, then $l(q) \leq l(p) - 1$ and therefore $q \in V_i(p)$. It follows that $\mathcal{S}_i(p, \mathcal{T}(R)) \cap$ Descendants$(q) = \emptyset$ and therefore

$$\mathcal{S}_i(p, \mathcal{T}(R)) \cap \mathcal{S}_i(q, \mathcal{T}(R)) \subseteq \mathcal{S}_i(p, \mathcal{T}(R)) \cap \text{Descendants}(q) = \emptyset.$$

◄

▶ **Lemma 4.12** (sum lemma [10, Lemma 5.9]). *In the notations of Definition 4.8 for any subset $V \subseteq C$ of a set C, we have $|\bigcup\limits_{p \in V} \mathcal{S}_i(p, \mathcal{T}(R))| = \sum\limits_{p \in V} |\mathcal{S}_i(p, \mathcal{T}(R))|$.* ∎

Proof. The proof follows from Lemma 4.11 by noting that any $p \in V \subseteq C$ has $l(p) \geq i$. ◄

▶ **Lemma 4.13** ([10, Lemma 5.11]). *In the notations of Definition 4.1, let $p \in \mathcal{T}(R)$ be any node. If $w \in \mathcal{S}_i(p, \mathcal{T}(R))$ then either $w = p$ or there exists $a \in$ Children$(p) \setminus \{p\}$ such that $l(a) < i$ and $w \in$ Descendants(a).* ∎

Proof. Let $w \in \mathcal{S}_i(p)$ be an arbitrary node satisfying $w \neq p$. Let s be the node-to-root path of w. The inclusion $\mathcal{S}_i(p) \subseteq$ Descendants(p) implies that $w \in$ Descendants(p).

Let $a \in$ Children$(p) \setminus \{p\}$ be a child on the path s. If $l(a) \geq i$ then $a \in V_i(p)$. Note that $w \in$ Descendants(a). Therefore $w \notin \mathcal{S}_i(p)$, which is a contradiction. Hence $l(a) < i$. ◄

▶ **Lemma 4.15** ([10, Lemma 5.12]). *Let R be a finite subset of a metric pace. Let $\mathcal{T}(R)$ be a compressed cover tree on R. Let $R_i \subseteq C_i$, where C_i is the ith cover set of $\mathcal{T}(R)$. Set $\mathcal{C}(R_i) = \{a \in$ Children(p) for some $p \in R_i \mid l(a) \geq i - 1\}$. Then*

$$\bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R)) = \bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R)).$$

Proof. Let $a \in \bigcup_{p \in C} \mathcal{S}_{i-1}(p, \mathcal{T}(R))$ be an arbitrary node. Then there is $v \in C$ having $a \in \mathcal{S}_{i-1}(v, \mathcal{T}(R))$. Let $w \in R_i$ be an ancestor of v that has the lowest level among all ancestors from in R_i. The node v has always an ancestor in R_i. Indeed, by the choice of $v \in C$, either $v \in R_i$ or v has a parent in R_i. Hence $a \in$ Descendants(w). Since w has a minimal level among all ancestors of v, we conclude that $a \notin \bigcup_{u \in V_i(w)}$ Descendants(u). Then

$$a \in \mathcal{S}_i(w, \mathcal{T}(R)) \subseteq \bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R)).$$

Assume now that $a \in \bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R))$. Then $a \in \mathcal{S}_i(v, \mathcal{T}(R))$ for some $w \in R_i$. Let w have no children at the level $i - 1$. Then $V_i(w) = V_{i-1}(w)$ and

$$a \in \mathcal{S}_{i-1}(w, \mathcal{T}(R)) \subseteq \bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R)).$$

Assume now that w has children at the level $i - 1$. If there exists $b \in$ Children(w) for which $a \in$ Descendants(b). Since $V_{i-1}(b) = \emptyset$, we conclude that

$$a \in \mathcal{S}_{i-1}(b, \mathcal{T}(R)) \subseteq \bigcup_{p \in C} \mathcal{S}_{i-1}(p, \mathcal{T}(R)).$$

To prove the converse inclusion assume that $a \notin \mathrm{Descendants}(b)$ for all $b \in \mathrm{Children}(w)$ with $l(b) = i - 1$. Then $a \in \mathrm{Descendants}(w)$ and $a \notin \mathrm{Descendants}(b')$ for any $b' \in V_i(w)$. Then $a \in \mathcal{S}_{i-1}(w, \mathcal{T}(R))$ and the proof finishes:

$$\bigcup_{p \in R_i} \mathcal{S}_i(p, \mathcal{T}(R)) \subseteq \bigcup_{p \in \mathcal{C}(R_i)} \mathcal{S}_{i-1}(p, \mathcal{T}(R)).$$

◀

▶ **Lemma 4.16** (β-point [10, Lemma 5.13])**.** *In the notations of Definition 4.8, let $C \subseteq C_i$ so that $\cup_{p \in C} \mathcal{S}_i(p, \mathcal{T}(R))$ contains all k-nearest neighbors of q. Set $\lambda = \lambda_k(q, C)$. Then R has a point β among the first k nearest neighbors of q such that $d(q, \lambda) \leq d(q, \beta) + 2^{i+1}$.* ■

Proof. We show that R has a point β among the first k nearest neighbors of q such that

$$\beta \in \bigcup_{p \in C} \mathcal{S}_i(p, \mathcal{T}(R)) \setminus \bigcup_{p \in N(q,\lambda) \setminus \{\lambda\}} \mathcal{S}_i(p, \mathcal{T}(R)).$$

Lemma 4.12 and Definition 4.8 imply that

$$\left| \bigcup_{p \in N(q,\lambda) \setminus \{\lambda\}} \mathcal{S}_i(p, \mathcal{T}(R)) \right| = \sum_{p \in N(q,\lambda) \setminus \{\lambda\}} |\mathcal{S}_i(p, \mathcal{T}(R))| < k.$$

Since $\cup_{p \in C} s_i(p, \mathcal{T}(R))$ contains all k-nearest neighbors of q, a required point $\beta \in R$ exists.

Let us now show that β satisfies $d(q, \lambda) \leq d(q, \beta) + 2^{i+1}$. Let $\gamma \in C \setminus N(q, \lambda) \cup \{\lambda\}$ be an ancestor of β. Since $\gamma \notin N(q, \lambda) \setminus \{\lambda\}$, we get $d(\gamma, q) \geq d(q, \lambda)$. The triangle inequality says that $d(q, \gamma) \leq d(q, \beta) + d(\gamma, \beta)$. Finally, Lemma 3.7 implies that $d(\gamma, \beta) \leq 2^{i+1}$. Then

$$d(q, \lambda) \leq d(q, \gamma) \leq d(q, \beta) + d(\gamma, \beta) \leq d(q, \beta) + 2^{i+1}$$

So β is a desired k-nearest neighbor satisfying the condition $d(q, \lambda) \leq d(q, \beta) + 2^{i+1}$. ◀