

---

# Fast $k$ -Nearest Neighbour Search via Dynamic Continuous Indexing

---

Ke Li

Jitendra Malik

University of California, Berkeley, CA 94720, United States

KE.LI@EECS.BERKELEY.EDU

MALIK@EECS.BERKELEY.EDU

## Abstract

Existing methods for retrieving  $k$ -nearest neighbours suffer from the curse of dimensionality. We argue this is caused in part by inherent deficiencies of space partitioning, which is the underlying strategy used by most existing methods. We devise a new strategy that avoids partitioning the vector space and present a novel randomized algorithm that runs in time linear in dimensionality of the space and sub-linear in the intrinsic dimensionality and the size of the dataset and takes space constant in dimensionality of the space and linear in the size of the dataset. The proposed algorithm allows fine-grained control over accuracy and speed on a per-query basis, automatically adapts to variations in data density, supports dynamic updates to the dataset and is easy-to-implement. We show appealing theoretical properties and demonstrate empirically that the proposed algorithm outperforms locality-sensitivity hashing (LSH) in terms of approximation quality, speed and space efficiency.

## 1. Introduction

The  $k$ -nearest neighbour method is commonly used both as a classifier and as subroutines in more complex algorithms in a wide range domains, including machine learning, computer vision, graphics and robotics. Consequently, finding a fast algorithm for retrieving nearest neighbours has been a subject of sustained interest among the artificial intelligence and the theoretical computer science communities alike. Work over the past several decades has produced a rich collection of algorithms; however, they suffer from one key shortcoming: as the ambient or intrinsic dimen-

sionality<sup>1</sup> increases, the running time and/or space complexity grows rapidly; this phenomenon is often known as the curse of dimensionality. Finding a solution to this problem has proven to be elusive and has been conjectured to be fundamentally impossible (Minsky & Seymour, 1969). In this era of rapid growth in both the volume and dimensionality of data, it has become increasingly important to devise a fast algorithm that scales to high dimensional space.

We argue that the difficulty in overcoming the curse of dimensionality stems in part from inherent deficiencies in space partitioning, the strategy that underlies most existing algorithms. Space partitioning is a divide-and-conquer strategy that partitions the vector space into a finite number of cells and keeps track of the data points that each cell contains. At query time, exhaustive search is performed over all data points that lie in cells containing the query point. This strategy forms the basis of most existing algorithms, including  $k$ -d trees (Bentley, 1975) and locality-sensitive hashing (LSH) (Indyk & Motwani, 1998).

While this strategy seems natural and sensible, it suffers from critical deficiencies as the dimensionality increases. Because the volume of any region in the vector space grows exponentially in the dimensionality, either the number or the size of cells must increase exponentially in the number of dimensions, which tends to lead to exponential time or space complexity in the dimensionality. In addition, space partitioning limits the algorithm’s “field of view” to the cell containing the query point; points that lie in adjacent cells have no hope of being retrieved. Consequently, if a query point falls near cell boundaries, the algorithm will fail to retrieve nearby points that lie in adjacent cells. Since the number of such cells is exponential in the dimensionality, it is intractable to search these cells when dimensionality is high. One popular approach used by LSH and spill trees (Liu et al., 2004) to mitigate this effect is to partition the space using overlapping cells and search over all points that lie in any of the cells that contain the query point. Because

---

*Proceedings of the 33<sup>rd</sup> International Conference on Machine Learning*, New York, NY, USA, 2016. JMLR: W&CP volume 48. Copyright 2016 by the author(s).

<sup>1</sup>We use *ambient dimensionality* to refer to the original dimensionality of the space containing the data in order to differentiate it from *intrinsic dimensionality*, which measures density of the dataset and will be defined precisely later.

the ratio of surface area to volume grows in dimensionality, the number of overlapping cells that must be used increases in dimensionality; as a result, the running time or space usage becomes prohibitively expensive as dimensionality increases. Further complications arise from variations in data density across different regions of the space. If the partitioning is too fine, most cells in sparse regions of the space will be empty and so for a query point that lies in a sparse region, no points will be retrieved. If the partitioning is too coarse, each cell in dense regions of the space will contain many points and so for a query point that lies in a dense region, many points that are not the true nearest neighbours must be searched. This phenomenon is notably exhibited by LSH, whose performance is highly sensitive to the choice of the hash function, which essentially defines an implicit partitioning. A good partitioning scheme must therefore depend on the data; however, such data-dependent partitioning schemes would require possibly expensive preprocessing and prohibit online updates to the dataset. These fundamental limitations of space partitioning beg an interesting question: is it possible to devise a strategy that does not partition the space, yet still enables fast retrieval of nearest neighbours?

In this paper, we present a new strategy for retrieving  $k$ -nearest neighbours that avoids discretizing the vector space, which we call dynamic continuous indexing (DCI). Instead of partitioning the space into discrete cells, we construct continuous indices, each of which imposes an ordering on data points such that closeness in position serves as an approximate indicator of proximity in the vector space. The resulting algorithm runs in time linear in ambient dimensionality and sub-linear in intrinsic dimensionality and the size of the dataset, while only requiring space constant in ambient dimensionality and linear in the size of the dataset. Unlike existing methods, the algorithm allows fine-grained control over accuracy and speed at query time and adapts to varying data density on-the-fly while permitting dynamic updates to the dataset. Furthermore, the algorithm is easy-to-implement and does not rely on any complex or specialized data structure.

## 2. Related Work

Extensive work over the past several decades has produced a rich collection of algorithms for fast retrieval of  $k$ -nearest neighbours. Space partitioning forms the basis of the majority of these algorithms. Early approaches store points in deterministic tree-based data structures, such as  $k$ -d trees (Bentley, 1975), R-trees (Guttman, 1984) and X-trees (Berchtold et al., 1996; 1998), which effectively partition the vector space into a hierarchy of half-spaces, hyperrectangles or Voronoi polygons. These methods achieve query times that are logarithmic in the number of data

points and work very well for low-dimensional data. Unfortunately, their query times grow exponentially in ambient dimensionality because the number of leaves in the tree that need to be searched increases exponentially in ambient dimensionality; as a result, for high-dimensional data, these algorithms become slower than exhaustive search. More recent methods like spill trees (Liu et al., 2004), RP trees (Dasgupta & Freund, 2008) and virtual spill trees (Dasgupta & Sinha, 2015) extend these approaches by randomizing the dividing hyperplane at each node. Unfortunately, the number of points in the leaves increases exponentially in the intrinsic dimensionality of the dataset.

In an effort to tame the curse of dimensionality, researchers have considered relaxing the problem to allow  $\epsilon$ -approximate solutions, which can contain any point whose distance to the query point differs from that of the true nearest neighbours by at most a factor of  $(1 + \epsilon)$ . Tree-based methods (Arya et al., 1998) have been proposed for this setting; unfortunately, the running time still exhibits exponential dependence on dimensionality. Another popular method is locality-sensitive hashing (LSH) (Indyk & Motwani, 1998; Datar et al., 2004), which relies on a hash function that implicitly defines a partitioning of the space. Unfortunately, LSH struggles on datasets with varying density, as cells in sparse regions of the space tend to be empty and cells in dense regions tend to contain a large number of points. As a result, it fails to return any point on some queries and requires a long time on some others. This motivated the development of data-dependent hashing schemes based on  $k$ -means (Paulevé et al., 2010) and spectral partitioning (Weiss et al., 2009). Unfortunately, these methods do not support dynamic updates to the dataset or provide correctness guarantees. Furthermore, they incur a significant pre-processing cost, which can be expensive on large datasets. Other data-dependent algorithms outside the LSH framework have also been proposed, such as (Fukunaga & Narendra, 1975; Brin, 1995; Nister & Stewenius, 2006; Wang, 2011), which work by constructing a hierarchy of clusters using  $k$ -means and can be viewed as performing a highly data-dependent form of space partitioning. For these algorithms, no guarantees on approximation quality or running time are known.

One class of methods (Orchard, 1991; Arya & Mount, 1993; Clarkson, 1999; Karger & Ruhl, 2002) that does not rely on space partitioning uses a local search strategy. Starting from a random data point, these methods iteratively find a new data point that is closer to the query than the previous data point. Unfortunately, the performance of these methods deteriorates in the presence of significant variations in data density, since it may take very long to navigate a dense region of the space, even if it is very far from the query. Other methods like navigating nets (Krauthgamer & Lee, 2004), cover trees (Beygelzimer et al., 2006) and rank

cover trees (Houle & Nett, 2015) adopt a coarse-to-fine strategy. These methods work by maintaining coarse subsets of points at varying scales and progressively searching the neighbourhood of the query with decreasing radii at increasingly finer scales. Sadly, the running times of these methods again exhibit exponential dependence on the intrinsic dimensionality.

We direct interested readers to (Clarkson, 2006) for a comprehensive survey of existing methods.

### 3. Algorithm

---

#### Algorithm 1 Algorithm for data structure construction

---

**Require:** A dataset  $D$  of  $n$  points  $p^1, \dots, p^n$ , the number of simple indices  $m$  that constitute a composite index and the number of composite indices  $L$

**function** CONSTRUCT( $D, m, L$ )

$\{u_{jl}\}_{j \in [m], l \in [L]} \leftarrow mL$  random unit vectors in  $\mathbb{R}^d$

$\{T_{jl}\}_{j \in [m], l \in [L]} \leftarrow mL$  empty binary search trees or skip lists

**for**  $j = 1$  **to**  $m$  **do**

**for**  $l = 1$  **to**  $L$  **do**

**for**  $i = 1$  **to**  $n$  **do**

$\bar{p}_{jl}^i \leftarrow \langle p^i, u_{jl} \rangle$

                Insert  $(\bar{p}_{jl}^i, i)$  into  $T_{jl}$  with  $\bar{p}_{jl}^i$  being the key and  $i$  being the value

**end for**

**end for**

**end for**

**return**  $\{(T_{jl}, u_{jl})\}_{j \in [m], l \in [L]}$

**end function**

---

The proposed algorithm relies on the construction of continuous indices of data points that support both fast searching and online updates. To this end, we use one-dimensional random projections as basic building blocks and construct  $mL$  simple indices, each of which orders data points by their projections along a random direction. Such an index has both desired properties: data points can be efficiently retrieved and updated when the index is implemented using standard data structures for storing ordered sequences of scalars, like self-balancing binary search trees or skip lists. This ordered arrangement of data points exploits an important property of the  $k$ -nearest neighbour search problem that has often been overlooked: it suffices to construct an index that approximately preserves the relative *order* between the true  $k$ -nearest neighbours and the other data points in terms of their distances to the query point without necessarily preserving all pairwise distances. This observation enables projection to a much lower dimensional space than the Johnson-Lindenstrauss transform (Johnson & Lindenstrauss, 1984). We show in the following section that with high probability, one-dimensional random projection preserves the relative order of two points whose distances to the query point differ significantly re-

gardless of the ambient dimensionality of the points.

---

#### Algorithm 2 Algorithm for $k$ -nearest neighbour retrieval

---

**Require:** Query point  $q$  in  $\mathbb{R}^d$ , binary search trees/skip lists and their associated projection vectors  $\{(T_{jl}, u_{jl})\}_{j \in [m], l \in [L]}$ , and maximum tolerable failure probability  $\epsilon$

**function** QUERY( $q, \{(T_{jl}, u_{jl})\}_{j \in [m], l \in [L]}, \epsilon$ )

$C_l \leftarrow$  array of size  $n$  with entries initialized to 0  $\forall l \in [L]$

$\bar{q}_{jl} \leftarrow \langle q, u_{jl} \rangle \forall j \in [m], l \in [L]$

**for**  $i = 1$  **to**  $n$  **do**

**for**  $l = 1$  **to**  $L$  **do**

$S_l \leftarrow \emptyset$

**for**  $j = 1$  **to**  $m$  **do**

$(\bar{p}_{jl}^{(i)}, h_{jl}^{(i)}) \leftarrow$  the node in  $T_{jl}$  whose key is the  $i^{\text{th}}$  closest to  $\bar{q}_{jl}$

$C_l[h_{jl}^{(i)}] \leftarrow C_l[h_{jl}^{(i)}] + 1$

**end for**

**for**  $j = 1$  **to**  $m$  **do**

**if**  $C_l[h_{jl}^{(i)}] = m$  **then**

$S_l \leftarrow S_l \cup \{h_{jl}^{(i)}\}$

**end if**

**end for**

**if** *IsStoppingConditionSatisfied*( $i, S_l, \epsilon$ ) **then**

**break**

**end if**

**end for**

**return**  $k$  points in  $\bigcup_{l \in [L]} S_l$  that are the closest in Euclidean distance in  $\mathbb{R}^d$  to  $q$

**end function**

---

We combine each set of  $m$  simple indices to form a composite index in which points are ordered by the maximum difference over all simple indices between the positions of the point and the query in the simple index. The composite index enables fast retrieval of a small number of data points, which will be referred to as candidate points, that are close to the query point along several random directions and therefore are likely to be truly close to the query. The composite indices are not explicitly constructed; instead, each of them simply keeps track of the number of its constituent simple indices that have encountered any particular point and returns a point as soon as all its constituent simple indices have encountered that point.

At query time, we retrieve candidate points from each composite index one by one in the order they are returned until some stopping condition is satisfied, while omitting points that have been previously retrieved from other composite indices. Exhaustive search is then performed over candidate points retrieved from all  $L$  composite indices to identify the subset of  $k$  points closest to the query. Please refer to Algorithms 1 and 2 for a precise statement of the construction and querying procedures.

Because data points are retrieved according to their positions in the composite index rather than the regions of space they lie in, the algorithm is able to automatically adapt to




Figure 1. (a) Examples of order-preserving (shown in green) and order-inverting (shown in red) projection directions. Any projection direction within the shaded region inverts the relative order of the vectors by length under projection, while any projection directions outside the region preserves it. The size of the shaded region depends on the ratio of the lengths of the vectors. (b) Projection vectors whose endpoints lie in the shaded region would be order-inverting. (c) Projection vectors whose endpoints lie in the shaded region would invert the order of both long vectors relative to the short vector. Best viewed in colour.

changes in data density as dynamic updates are made to the dataset without requiring any pre-processing to estimate data density at construction time. Also, unlike existing methods, the number of retrieved candidate points can be controlled on a per-query basis, enabling the user to easily trade off accuracy against speed. We develop two versions of the algorithm, a data-independent and a data-dependent version, which differ in the stopping condition that is used. In the former, the number of candidate points is indirectly preset according to the global data density and the maximum tolerable failure probability; in the latter, the number of candidate points is chosen adaptively at query time based on the local data density in the neighbourhood of the query. We analyze the algorithm below and show that its failure probability is independent of ambient dimensionality, its running time is linear in ambient dimensionality and sub-linear in intrinsic dimensionality and the size of the dataset and its space complexity is independent of ambient dimensionality and linear in the size of the dataset.

### 3.1. Properties of Random 1D Projection

First, we examine the effect of projecting  $d$ -dimensional vectors to one dimension, which motivates its use in the proposed algorithm. We are interested in the probability that a distant point appears closer than a nearby point under projection; if this probability is low, then each simple index approximately preserves the order of points by distance to the query point. If we consider displacement vectors between the query point and data points, this probability is then equivalent to the probability of the lengths of these vectors inverting under projection.

**Lemma 1.** *Let  $v^l, v^s \in \mathbb{R}^d$  such that  $\|v^l\|_2 > \|v^s\|_2$ , and  $u \in \mathbb{R}^d$  be a unit vector drawn uniformly at random. Then the probability of  $v^s$  being at least as long as  $v^l$  under*

*projection  $u$  is at most  $1 - \frac{2}{\pi} \cos^{-1} \left( \frac{\|v^s\|_2}{\|v^l\|_2} \right)$ .*

*Proof.* Assuming that  $v^l$  and  $v^s$  are not collinear, consider the two-dimensional subspace spanned by  $v^l$  and  $v^s$ , which we will denote as  $P$ . (If  $v^l$  and  $v^s$  are collinear, we define  $P$  to be the subspace spanned by  $v^l$  and an arbitrary vector that's linearly independent of  $v^l$ .) For any vector  $w$ , we use  $w^\parallel$  and  $w^\perp$  to denote the components of  $w$  in  $P$  and  $P^\perp$  such that  $w = w^\parallel + w^\perp$ . For  $w \in \{v^s, v^l\}$ , because  $w^\perp = 0$ ,  $\langle w, u \rangle = \langle w^\parallel, u^\parallel \rangle$ . So, we can limit our attention to  $P$  for this analysis. We parameterize  $u^\parallel$  in terms of its angle relative to  $v^l$ , which we denote as  $\theta$ . Also, we denote the angle of  $u^\parallel$  relative to  $v^s$  as  $\psi$ . Then,

$$\begin{aligned}
 & \Pr \left( |\langle v^l, u \rangle| \leq |\langle v^s, u \rangle| \right) \\
 &= \Pr \left( |\langle v^l, u^\parallel \rangle| \leq |\langle v^s, u^\parallel \rangle| \right) \\
 &= \Pr \left( \|v^l\|_2 \|u^\parallel\|_2 |\cos \theta| \leq \|v^s\|_2 \|u^\parallel\|_2 |\cos \psi| \right) \\
 &\leq \Pr \left( |\cos \theta| \leq \frac{\|v^s\|_2}{\|v^l\|_2} \right) \\
 &= 2 \Pr \left( \theta \in \left[ \cos^{-1} \left( \frac{\|v^s\|_2}{\|v^l\|_2} \right), \pi - \cos^{-1} \left( \frac{\|v^s\|_2}{\|v^l\|_2} \right) \right] \right) \\
 &= 1 - \frac{2}{\pi} \cos^{-1} \left( \frac{\|v^s\|_2}{\|v^l\|_2} \right)
 \end{aligned}$$

□

Observe that if  $|\langle v^l, u \rangle| \leq |\langle v^s, u \rangle|$ , the relative order of  $v^l$  and  $v^s$  by their lengths would be inverted when projected along  $u$ . This occurs when  $u^\parallel$  is close to orthogonal to  $v^l$ , which is illustrated in Figure 1a. Also note that the probability of inverting the relative order of  $v^l$  and  $v^s$  is small when  $v^l$  is much longer than  $v^s$ . On the other hand, this probability is high when  $v^l$  and  $v^s$  are similar in length,

which corresponds to the case when two data points are almost equidistant to the query point. So, if we consider a sequence of vectors ordered by length, applying random one-dimensional projection will likely perturb the ordering locally, but will preserve the ordering globally.

Next, we build on this result to analyze the order-inversion probability when there are more than two vectors. Consider the sample space  $B = \{u \in \mathbb{R}^d \mid \|u\|_2 = 1\}$  and the set  $U(v^s, v^l) = \{u \in B \mid |\cos \theta| \leq \|v^s\|_2 / \|v^l\|_2\}$ , which is illustrated in Figure 1b, where  $\theta$  is the angle between  $u$  and  $v^l$ . If we use  $\text{area}(U)$  to denote the area of the region formed by the endpoints of all vectors in the set  $U$ , then we can rewrite the above bound on the order-inversion probability as:

$$\begin{aligned} \Pr\left(\left|\langle v^l, u \rangle\right| \leq \left|\langle v^s, u \rangle\right|\right) &\leq \Pr\left(u \in U(v^s, v^l)\right) \\ &= \frac{\text{area}(U)}{\text{area}(B)} \\ &= 1 - \frac{2}{\pi} \cos^{-1}\left(\frac{\|v^s\|_2}{\|v^l\|_2}\right) \end{aligned}$$

**Lemma 2.** Let  $\{v_i^l\}_{i=1}^N$  be a set of vectors such that  $\|v_i^l\|_2 > \|v^s\|_2 \forall i \in [N]$ . Then the probability that there is a subset of  $k'$  vectors from  $\{v_i^l\}_{i=1}^N$  that are all not longer than  $v^s$  under projection is at most  $\frac{1}{k'} \sum_{i=1}^N \left(1 - \frac{2}{\pi} \cos^{-1}\left(\frac{\|v^s\|_2}{\|v_i^l\|_2}\right)\right)$ . Furthermore, if  $k' = N$ , this probability is at most  $\min_{i \in [N]} \left\{1 - \frac{2}{\pi} \cos^{-1}\left(\frac{\|v^s\|_2}{\|v_i^l\|_2}\right)\right\}$ .

*Proof.* For a given subset  $I \subseteq [N]$  of size  $k'$ , the probability that all vectors indexed by elements in  $I$  are not longer than  $v^s$  under projection  $u$  is at most  $\Pr(u \in \bigcap_{i \in I} U(v^s, v_i^l)) = \text{area}(\bigcap_{i \in I} U(v^s, v_i^l)) / \text{area}(B)$ . So, the probability that this occurs on some subset  $I$  is at most  $\Pr(u \in \bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v^s, v_i^l)) = \text{area}\left(\bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v^s, v_i^l)\right) / \text{area}(B)$ .

Observe that each point in  $\bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v^s, v_i^l)$  must be covered by at least  $k'$   $U(v^s, v_i^l)$ 's. So,

$$k' \cdot \text{area}\left(\bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v^s, v_i^l)\right) \leq \sum_{i=1}^N \text{area}\left(U(v^s, v_i^l)\right)$$

It follows that the probability this event occurs on some subset  $I$  is bounded above by  $\frac{1}{k'} \sum_{i=1}^N \frac{\text{area}(U(v^s, v_i^l))}{\text{area}(B)} = \frac{1}{k'} \sum_{i=1}^N \left(1 - \frac{2}{\pi} \cos^{-1}\left(\frac{\|v^s\|_2}{\|v_i^l\|_2}\right)\right)$ .

If  $k' = N$ , we use the fact that  $\text{area}\left(\bigcap_{i \in [N]} U(v^s, v_i^l)\right) \leq \min_{i \in [N]} \left\{\text{area}(U(v^s, v_i^l))\right\}$  to obtain the desired result.  $\square$

Intuitively, if this event occurs, then there are at least  $k'$  vectors that rank above  $v^s$  when sorted in nondecreasing order by their lengths under projection. This can only occur when the endpoint of  $u$  falls in a region on the unit sphere corresponding to  $\bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v^s, v_i^l)$ . We illustrate this region in Figure 1c for the case of  $d = 3$ .

**Theorem 3.** Let  $\{v_i^l\}_{i=1}^N$  and  $\{v_{i'}^s\}_{i'=1}^{N'}$  be sets of vectors such that  $\|v_i^l\|_2 > \|v_{i'}^s\|_2 \forall i \in [N], i' \in [N']$ . Then the probability that there is a subset of  $k'$  vectors from  $\{v_i^l\}_{i=1}^N$  that are all not longer than some  $v_{i'}^s$  under projection is at most  $\frac{1}{k'} \sum_{i=1}^N \left(1 - \frac{2}{\pi} \cos^{-1}\left(\frac{\|v_{\max}^s\|_2}{\|v_i^l\|_2}\right)\right)$ , where  $\|v_{\max}^s\|_2 \geq \|v_{i'}^s\|_2 \forall i' \in [N']$ .

*Proof.* The probability that this event occurs is at most  $\Pr\left(u \in \bigcup_{i' \in [N']} \bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v_{i'}^s, v_i^l)\right)$ . We observe that for all  $i, i'$ ,  $\{\theta \mid |\cos \theta| \leq \|v_{i'}^s\|_2 / \|v_i^l\|_2\} \subseteq \{\theta \mid |\cos \theta| \leq \|v_{\max}^s\|_2 / \|v_i^l\|_2\}$ , which implies that  $U(v_{i'}^s, v_i^l) \subseteq U(v_{\max}^s, v_i^l)$ .

If we take the intersection followed by union on both sides, we obtain  $\bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v_{i'}^s, v_i^l) \subseteq \bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v_{\max}^s, v_i^l)$ . Because this is true for all  $i'$ ,  $\bigcup_{i' \in [N']} \bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v_{i'}^s, v_i^l) \subseteq \bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v_{\max}^s, v_i^l)$ .

Therefore, this probability is bounded above by  $\Pr\left(u \in \bigcup_{I \subseteq [N]: |I|=k'} \bigcap_{i \in I} U(v_{\max}^s, v_i^l)\right)$ . By Lemma 2, this is at most  $\frac{1}{k'} \sum_{i=1}^N \left(1 - \frac{2}{\pi} \cos^{-1}\left(\frac{\|v_{\max}^s\|_2}{\|v_i^l\|_2}\right)\right)$ .  $\square$

### 3.2. Data Density

We now formally characterize data density by defining the following notion of local relative sparsity:

**Definition 4.** Given a dataset  $D \subseteq \mathbb{R}^d$ , let  $B_p(r)$  be the set of points in  $D$  that are within a ball of radius  $r$  around a point  $p$ . We say  $D$  has local relative sparsity of  $(\tau, \gamma)$  at a point  $p \in \mathbb{R}^d$  if for all  $r$  such that  $|B_p(r)| \geq \tau$ ,  $|B_p(\gamma r)| \leq 2|B_p(r)|$ , where  $\gamma \geq 1$ .

Intuitively,  $\gamma$  represents a lower bound on the increase in radius when the number of points within the ball is doubled. When  $\gamma$  is close to 1, the dataset is dense in the neighbourhood of  $p$ , since there could be many points in  $D$  that are almost equidistant from  $p$ . Retrieving the nearest neighbours of such a  $p$  is considered ‘‘hard’’, since it would be difficult to tell which of these points are the true nearest neighbours without computing the distances to all these points exactly.

We also define a related notion of global relative sparsity, which we will use to derive the number of iterations the outer loop of the querying function should be executed and

a bound on the running time that is independent of the query:

**Definition 5.** A dataset  $D$  has global relative sparsity of  $(\tau, \gamma)$  if for all  $r$  and  $p \in \mathbb{R}^d$  such that  $|B_p(r)| \geq \tau$ ,  $|B_p(\gamma r)| \leq 2|B_p(r)|$ , where  $\gamma \geq 1$ .

Note that a dataset with global relative sparsity of  $(\tau, \gamma)$  has local relative sparsity of  $(\tau, \gamma)$  at every point. Global relative sparsity is closely related to the notion of *expansion rate* introduced by (Karger & Ruhl, 2002). More specifically, a dataset with global relative sparsity of  $(\tau, \gamma)$  has  $(\tau, 2^{(1/\log_2 \gamma)})$ -expansion, where the latter quantity is known as the expansion rate. If we use  $c$  to denote the expansion rate, the quantity  $\log_2 c$  is known as the *intrinsic dimension* (Clarkson, 2006), since when the dataset is uniformly distributed, intrinsic dimensionality would match ambient dimensionality. So, the intrinsic dimension of a dataset with global relative sparsity of  $(\tau, \gamma)$  is  $1/\log_2 \gamma$ .

### 3.3. Data-Independent Version

In the data-independent version of the algorithm, the outer loop in the querying function executes for a preset number of iterations  $\tilde{k}$ . The values of  $L$ ,  $m$  and  $\tilde{k}$  are fixed for all queries and will be chosen later.

We apply the results obtained above to analyze the algorithm. Consider the event that the algorithm fails to return the correct set of  $k$ -nearest neighbours – this can only occur if a true  $k$ -nearest neighbour is not contained in any of the  $S_l$ 's, which entails that for each  $l \in [L]$ , there is a set of  $\tilde{k} - k + 1$  points that are not the true  $k$ -nearest neighbours but are closer to the query than the true  $k$ -nearest neighbours under some of the projections  $u_{1l}, \dots, u_{ml}$ . We analyze the probability that this occurs below and derive the parameter settings that ensure the algorithm succeeds with high probability. Please refer to the supplementary material for proofs of the following results.

**Lemma 6.** For a dataset with global relative sparsity  $(k, \gamma)$ , there is some  $\tilde{k} \in \Omega(\max(k \log(n/k), k(n/k)^{1-\log_2 \gamma}))$  such that the probability that the candidate points retrieved from a given composite index do not include some of the true  $k$ -nearest neighbours is at most some constant  $\alpha < 1$ .

**Theorem 7.** For a dataset with global relative sparsity  $(k, \gamma)$ , for any  $\epsilon > 0$ , there is some  $L$  and  $\tilde{k} \in \Omega(\max(k \log(n/k), k(n/k)^{1-\log_2 \gamma}))$  such that the algorithm returns the correct set of  $k$ -nearest neighbours with probability of at least  $1 - \epsilon$ .

The above result suggests that we should choose  $\tilde{k} \in \Omega(\max(k \log(n/k), k(n/k)^{1-\log_2 \gamma}))$  to ensure the algorithm succeeds with high probability. Next, we analyze the time and space complexity of the algorithm. Proofs of the following results are found in the supplementary material.

**Theorem 8.** The algorithm takes  $O(\max(dk \log(n/k), dk(n/k)^{1-1/d'})$  time to retrieve the  $k$ -nearest neighbours at query time, where  $d'$  denotes the intrinsic dimension of the dataset.

**Theorem 9.** The algorithm takes  $O(dn + n \log n)$  time to preprocess the data points in  $D$  at construction time.

**Theorem 10.** The algorithm requires  $O(d + \log n)$  time to insert a new data point and  $O(\log n)$  time to delete a data point.

**Theorem 11.** The algorithm requires  $O(n)$  space in addition to the space used to store the data.

### 3.4. Data-Dependent Version

Conceptually, performance of the proposed algorithm depends on two factors: how likely the index returns the true nearest neighbours before other points and when the algorithm stops retrieving points from the index. The preceding sections primarily focused on the former; in this section, we take a closer look at the latter.

One strategy, which is used by the data-independent version of the algorithm, is to stop after a preset number of iterations of the outer loop. Although simple, such a strategy leaves much to be desired. First of all, in order to set the number of iterations, it requires knowledge of the global relative sparsity of the dataset, which is rarely known a priori. Computing this is either very expensive in the case of large datasets or infeasible in the case of streaming data, as global relative sparsity may change as new data points arrive. More importantly, it is unable to take advantage of the local relative sparsity in the neighbourhood of the query. A method that is capable of adapting to local relative sparsity could potentially be much faster because query points tend to be close to the manifold on which points in the dataset lie, resulting in the dataset being sparse in the neighbourhood of the query point.

Ideally, the algorithm should stop as soon as it has retrieved the true nearest neighbours. Determining if this is the case amounts to asking if there exists a point that we have not seen lying closer to the query than the points we have seen. At first sight, because nothing is known about unseen points, it seems not possible to do better than exhaustive search, as we can only rule out the existence of such a point after computing distances to all unseen points. Somewhat surprisingly, by exploiting the fact that the projections associated with the index are random, it is possible to make inferences about points that we have never seen. We do so by leveraging ideas from statistical hypothesis testing.

After each iteration of the outer loop, we perform a hypothesis test, with the null hypothesis being that the complete set of the  $k$ -nearest neighbours has not yet been retrieved.

Rejecting the null hypothesis implies accepting the alternative hypothesis that all the true  $k$ -nearest neighbours have been retrieved. At this point, the algorithm can safely terminate while guaranteeing that the probability that the algorithm fails to return the correct results is bounded above by the significance level. The test statistic is an upper bound on the probability of missing a true  $k$ -nearest neighbour. The resulting algorithm does not require any prior knowledge about the dataset and terminates earlier when the dataset is sparse in the neighbourhood of the query; for this reason, we will refer to this version of the algorithm as the data-dependent version.

More concretely, as the algorithm retrieves candidate points, it computes their true distances to the query and maintains a list of  $k$  points that are the closest to the query among the points retrieved from all composite indices so far. Let  $\tilde{p}^{(i)}$  and  $\tilde{p}_i^{\max}$  denote the  $i^{\text{th}}$  closest candidate point to  $q$  retrieved from all composite indices and the farthest candidate point from  $q$  retrieved from the  $i^{\text{th}}$  composite index respectively. When the number of candidate points exceeds  $k$ , the algorithm checks if  $\prod_{l=1}^L \left(1 - \left(\frac{2}{\pi} \cos^{-1} \left(\frac{\|\tilde{p}^{(k)} - q\|_2}{\|\tilde{p}_l^{\max} - q\|_2}\right)\right)^m\right) \leq \epsilon$ , where  $\epsilon$  is the maximum tolerable failure probability, after each iteration of the outer loop. If the condition is satisfied, the algorithm terminates and returns  $\{\tilde{p}^{(i)}\}_{i=1}^k$ .

We show the correctness and running time of this algorithm below. Proofs of the following results are found in the supplementary material.

**Theorem 12.** *For any  $\epsilon > 0$ ,  $m$  and  $L$ , the data-dependent algorithm returns the correct set of  $k$ -nearest neighbours of the query  $q$  with probability of at least  $1 - \epsilon$ .*

**Theorem 13.** *On a dataset with global relative sparsity  $(k, \gamma)$ , given fixed parameters  $m$  and  $L$ , the data-dependent algorithm takes  $O\left(\max\left(dk \log\left(\frac{n}{k}\right), dk \left(\frac{n}{k}\right)^{1-\log_2 \gamma}, \frac{2d}{(1 - \frac{m}{\sqrt{1-\sqrt{\epsilon}}})^{d'}}\right)\right)$  time with high probability to retrieve the  $k$ -nearest neighbours at query time, where  $d'$  denotes the intrinsic dimension of the dataset.*

Note that we can make the denominator of the last argument arbitrarily close to 1 by choosing a large  $L$ .

## 4. Experiments

We compare the performance of the proposed algorithm to that of LSH, which is arguably the most popular method for fast nearest neighbour retrieval. Because LSH is designed for the approximate setting, under which the performance metric of interest is how close the points returned by the algorithm are to the query rather than whether the returned points are the true  $k$ -nearest neighbours, we empirically

evaluate performance in this setting. Because the distance metric of interest is the Euclidean distance, we compare to Exact Euclidean LSH (E<sup>2</sup>LSH) (Datar et al., 2004), which uses hash functions designed for Euclidean space.

We compare the performance of the proposed algorithm and LSH on the CIFAR-100 and MNIST datasets, which consist of  $32 \times 32$  colour images of various real-world objects and  $28 \times 28$  grayscale images of handwritten digits respectively. We reshape the images into vectors, with each dimension representing the pixel intensity at a particular location and colour channel of the image. The resulting vectors have a dimensionality of  $32 \times 32 \times 3 = 3072$  in the case of CIFAR-100 and  $28 \times 28 = 784$  in the case of MNIST, so the dimensionality under consideration is higher than what traditional tree-based algorithms can handle. We combine the training set and the test set of each dataset, and so we have a total of 60,000 instances in CIFAR-100 and 70,000 instances in MNIST.

We randomize the instances that serve as queries using cross-validation. Specifically, we randomly select 100 instances to serve as query points and designate the remaining instances as data points. Each algorithm is then used to retrieve approximate  $k$ -nearest neighbours of each query point among the set of all data points. This procedure is repeated for ten folds, each with a different split of query vs. data points.

We compare the number of candidate points that each algorithm requires to achieve a desired level of approximation quality. We quantify approximation quality using the approximation ratio, which is defined as the ratio of the radius of the ball containing approximate  $k$ -nearest neighbours to the radius of the ball containing true  $k$ -nearest neighbours. So, the smaller the approximation ratio, the better the approximation quality. Because dimensionality is high and exhaustive search must be performed over all candidate points, the time taken for compute distances between the candidate points and the query dominates the overall running time of the querying operation. Therefore, the number of candidate points can be viewed as an implementation-independent proxy for the running time.

Because the hash table constructed by LSH depends on the desired level of approximation quality, we construct a different hash table for each level of approximation quality. On the other hand, the indices constructed by the proposed method are not specific to any particular level of approximation quality; instead, approximation quality can be controlled at query time by varying the number of iterations of the outer loop. So, the same indices are used for all levels of approximation quality. Therefore, our evaluation scheme is biased towards LSH and against the proposed method.

We adopt the recommended guidelines for choosing param-




Figure 2. (a-b) Comparison of the speed of the proposed method (labelled as DCI) and LSH on (a) CIFAR-100 and (b) MNIST. Each curve represents the mean number to candidate points required to achieve varying levels of approximation quality over ten folds and the shaded area represents  $\pm 1$  standard deviation. (c) Comparison of the space efficiency of the proposed method and LSH on CIFAR-100 and MNIST. The height of each bar represents the average amount of memory used by each method to achieve the performance shown in (a) and (b).

eters of LSH and used 24 hashes per table and 100 tables. For the proposed method, we used  $m = 25$  and  $L = 2$  on CIFAR-100 and  $m = 15$  and  $L = 3$  on MNIST, which we found to work well in practice. In Figures 2a and 2b, we plot the performance of the proposed method and LSH on the CIFAR-100 and MNIST datasets for retrieving 25-nearest neighbours.

For the purposes of retrieving nearest neighbours, MNIST is a more challenging dataset than CIFAR-100. This is because the instances in MNIST form dense clusters, whereas the instances in CIFAR-100 are more visually diverse and so are more dispersed in space. Intuitively, if the query falls inside a dense cluster of points, there are many points that are very close to the query and so it is difficult to distinguish true nearest neighbours from points that are only slightly farther away. Viewed differently, because true nearest neighbours tend to be extremely close to the query on MNIST, the denominator for computing the approximation ratio is usually very small. Consequently, returning points that are only slightly farther away than the true nearest neighbours would result in a large approximation ratio. As a result, both the proposed method and LSH require far more candidate points on MNIST than on CIFAR-100 to achieve comparable approximation ratios.

We find the proposed method achieves better performance than LSH at all levels of approximation quality. Notably, the performance of LSH degrades drastically on MNIST, which is not surprising since LSH is known to have difficulties on datasets with large variations in data density. On the other hand, the proposed method requires 61.3% – 78.7% fewer candidate points than LSH to achieve the same approximation quality while using less than 1/20 of the memory.

## 5. Conclusion

In this paper, we delineated the inherent deficiencies of space partitioning and presented a new strategy for fast retrieval of  $k$ -nearest neighbours, which we dub dynamic continuous indexing (DCI). Instead of discretizing the vector space, the proposed algorithm constructs continuous indices, each of which imposes an ordering of data points in which closeby positions approximately reflect proximity in the vector space. Unlike existing methods, the proposed algorithm allows granular control over accuracy and speed on a per-query basis, adapts to variations to data density on-the-fly and supports online updates to the dataset. We analyzed the proposed algorithm and showed it runs in time linear in ambient dimensionality and sub-linear in intrinsic dimensionality and the size of the dataset and takes space constant in ambient dimensionality and linear in the size of the dataset. Furthermore, we demonstrated empirically that the proposed algorithm compares favourably to LSH in terms of approximation quality, speed and space efficiency.

**Acknowledgements.** Ke Li thanks the Berkeley Vision and Learning Center (BVLC) and the Natural Sciences and Engineering Research Council of Canada (NSERC) for financial support. The authors also thank Lap Chi Lau and the anonymous reviewers for feedback.

## References

- Arya, Sunil and Mount, David M. Approximate nearest neighbor queries in fixed dimensions. In *SODA*, volume 93, pp. 271–280, 1993.
- Arya, Sunil, Mount, David M, Netanyahu, Nathan S, Silverman, Ruth, and Wu, Angela Y. An optimal algorithm for approximate nearest neighbor searching fixed



- dimensions. *Journal of the ACM (JACM)*, 45(6):891–923, 1998.
- Bentley, Jon Louis. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- Berchtold, Stefan, Keim, Daniel A., and peter Kriegel, Hans. The X-tree : An Index Structure for High-Dimensional Data. In *Very Large Data Bases*, pp. 28–39, 1996.
- Berchtold, Stefan, Ertl, Bernhard, Keim, Daniel A, Kriegel, H-P, and Seidl, Thomas. Fast nearest neighbor search in high-dimensional space. In *Data Engineering, 1998. Proceedings., 14th International Conference on*, pp. 209–218. IEEE, 1998.
- Beygelzimer, Alina, Kakade, Sham, and Langford, John. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pp. 97–104. ACM, 2006.
- Brin, Sergey. Near neighbor search in large metric spaces. *VLDB*, pp. 574584, 1995.
- Clarkson, Kenneth L. Nearest neighbor queries in metric spaces. *Discrete & Computational Geometry*, 22(1):63–93, 1999.
- Clarkson, Kenneth L. Nearest-neighbor searching and metric space dimensions. *Nearest-neighbor methods for learning and vision: theory and practice*, pp. 15–59, 2006.
- Dasgupta, Sanjoy and Freund, Yoav. Random projection trees and low dimensional manifolds. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pp. 537–546. ACM, 2008.
- Dasgupta, Sanjoy and Sinha, Kaushik. Randomized partition trees for nearest neighbor search. *Algorithmica*, 72(1):237–263, 2015.
- Datar, Mayur, Immorlica, Nicole, Indyk, Piotr, and Mirrokni, Vahab S. Locality-sensitive hashing scheme based on  $p$ -stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*, pp. 253–262. ACM, 2004.
- Fukunaga, Keinosuke and Narendra, Patrenahalli M. A branch and bound algorithm for computing  $k$ -nearest neighbors. *Computers, IEEE Transactions on*, 100(7):750–753, 1975.
- Guttman, Antonin. *R-trees: a dynamic index structure for spatial searching*, volume 14. ACM, 1984.
- Houle, Michael E and Nett, Michael. Rank-based similarity search: Reducing the dimensional dependence. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(1):136–150, 2015.
- Indyk, Piotr and Motwani, Rajeev. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pp. 604–613. ACM, 1998.
- Johnson, William B and Lindenstrauss, Joram. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Karger, David R and Ruhl, Matthias. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pp. 741–750. ACM, 2002.
- Krauthgamer, Robert and Lee, James R. Navigating nets: simple algorithms for proximity search. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 798–807. Society for Industrial and Applied Mathematics, 2004.
- Liu, Ting, Moore, Andrew W, Yang, Ke, and Gray, Alexander G. An investigation of practical approximate nearest neighbor algorithms. In *Advances in Neural Information Processing Systems*, pp. 825–832, 2004.
- Minsky, Marvin and Seymour, Papert. Perceptrons: an introduction to computational geometry. pp. 222, 1969.
- Nister, David and Stewenius, Henrik. Scalable recognition with a vocabulary tree. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pp. 2161–2168. IEEE, 2006.
- Orchard, Michael T. A fast nearest-neighbor search algorithm. In *Acoustics, Speech, and Signal Processing, 1991. ICASSP-91., 1991 International Conference on*, pp. 2297–2300. IEEE, 1991.
- Paulevé, Loïc, Jégou, Hervé, and Amsaleg, Laurent. Locality sensitive hashing: A comparison of hash function types and querying mechanisms. *Pattern Recognition Letters*, 31(11):1348–1358, 2010.
- Wang, Xueyi. A fast exact  $k$ -nearest neighbors algorithm for high dimensional search using  $k$ -means clustering and triangle inequality. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pp. 1293–1299. IEEE, 2011.
- Weiss, Yair, Torralba, Antonio, and Fergus, Rob. Spectral hashing. In *Advances in Neural Information Processing Systems*, pp. 1753–1760, 2009.