

# On the I/O Complexity of the k-Nearest Neighbors Problem

Mayank Goswami\*  
mayank.goswami@qc.cuny.edu  
Queens College, CUNY  
Flushing, New York, USA

Riko Jacob†  
rikj@itu.dk  
IT University of Copenhagen  
København S, Denmark

Rasmus Pagh‡  
pagh@itu.dk  
BARC & IT University of Copenhagen  
København S, Denmark

## ABSTRACT

We consider static, external memory indexes for exact and approximate versions of the  $k$ -nearest neighbor ( $k$ -NN) problem, and show new lower bounds under a standard *indivisibility* assumption:

- Polynomial space indexing schemes for high-dimensional  $k$ -NN in Hamming space cannot take advantage of block transfers:  $\Omega(k)$  block reads are needed to answer a query.
- For the  $\ell_\infty$  metric the lower bound holds even if we allow  $c$ -approximate nearest neighbors to be returned, for  $c \in (1, 3)$ .
- The restriction to  $c < 3$  is necessary: For every metric there exists an indexing scheme in the *indexability model* of Hellerstein et al. using space  $O(kn)$ , where  $n$  is the number of points, that can retrieve  $k$  3-approximate nearest neighbors using optimal  $\lceil k/B \rceil$  I/Os, where  $B$  is the block size.
- For specific metrics, data structures with better approximation factors are possible. For  $k$ -NN in Hamming space and every approximation factor  $c > 1$  there exists a polynomial space data structure that returns  $k$   $c$ -approximate nearest neighbors in  $\lceil k/B \rceil$  I/Os.

To show these lower bounds we develop two new techniques: First, to handle that approximation algorithms have more freedom in deciding which result set to return we develop a *relaxed* version of the  $\lambda$ -set workload technique of Hellerstein et al. This technique allows us to show lower bounds that hold in  $d \geq n$  dimensions. To extend the lower bounds down to  $d = O(k \log(n/k))$  dimensions, we develop a new deterministic dimension reduction technique that may be of independent interest.

## CCS CONCEPTS

- **Theory of computation → Lower bounds and information complexity; Nearest neighbor algorithms; Sorting and searching; Expander graphs and randomness extractors; Random projections**

\*Supported by NSF grants CRII-1755791 and CCF-1910873.

†Part of this work done during the second Hawaiian workshop on parallel algorithms and data structures, University of Hawaii at Manoa, Hawaii, USA.

‡Work done in part while visiting Google Research. Supported by Investigator Grant 16582, Basic Algorithms Research Copenhagen (BARC), from the VILLUM Foundation, and by funding from the European Research Council under the European Union's 7th Framework Programme (FP7/2007-2013) / ERC grant agreement no. 614331.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODS'20, June 14–19, 2020, Portland, OR, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7108-7/20/06...\$15.00

<https://doi.org/10.1145/3375395.3387649>

and metric embeddings; Data structures and algorithms for data management; Computational geometry; • **Mathematics of computing → Approximation algorithms; • Information systems → Indexed file organization; Search engine indexing; Search index compression; • Hardware → External storage.**

## KEYWORDS

Nearest neighbors, I/O complexity, indexability model

### ACM Reference Format:

Mayank Goswami, Riko Jacob, and Rasmus Pagh. 2020. On the I/O Complexity of the  $k$ -Nearest Neighbors Problem. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS'20), June 14–19, 2020, Portland, OR, USA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3375395.3387649>

## 1 INTRODUCTION

The DEEP1B data set [8] is among the largest image data sets that has been examined in the similarity search literature. From each of  $n = 10^9$  images, a 96-dimensional vector has been extracted from an intermediate layer of a pre-trained deep neural network, a state-of-the-art method for semantically meaningful feature vectors [9]. Such feature vectors can be thought of as compressed representations of the images that, for example, can be used to estimate the similarity of two images. In many use cases, though, it is not enough to substitute the images with their feature vectors, but we also need to be able to access the corresponding images. Though the size of the raw image data behind DEEP1B is not stated in [8], an estimate would be 1 MB per image on average, or 1000 TB in total. Clearly, retrieving similar images in a data set of this size is beyond what is possible on a single machine, and even just indexing the set of feature vectors would require an amount of internal memory that is larger than what is present in most servers.

*k*-nearest neighbors. In the  $k$ -nearest neighbors ( $k$ -NN) problem we want to construct a data structure on a set  $P$  of  $n$  points in some metric space that, given an integer  $k > 0$  and a query point  $q$ , finds the  $k$  closest points to  $q$  in  $P$ . We will be focusing on data structures that can be constructed in polynomial time and space. The  $k$ -NN problem is believed to be hard in high dimensions even for  $k = 1$ , and the brute-force algorithm that considers all data points in  $P$  essentially optimal. In particular, Williams (see [4]) proved that for constant  $\varepsilon > 0$ , no  $O(n^{1-\varepsilon})$  query time data structure is possible for  $\omega(\log n)$ -dimensional Hamming space, assuming the Strong Exponential Time Hypothesis.

Because of the hardness of the problem most research has revolved around *approximate* solutions. The  $c$ -approximate  $k$ -NN ( $(c, k)$ -NN) problem asks to return  $k$  points from  $P$  with distance at most  $cr$  to  $q$ , where  $r$  is the distance to the  $k$ th nearest neighbor

of  $q$ . It is known that  $(c, k)$ -NN is equivalent, up to polylogarithmic factors, to the simpler *near neighbor* problem: Given an upper bound  $r_+ \geq r$ , return a point within distance  $cr_+$  [16]. We refer to [5] for more background on recent developments in approximate near neighbor search.

*Models of computation.* Motivated by large-scale similarity search applications we consider models of hardware aimed at massive data sets. The *external memory model* [3] abstracts modern, block-oriented storage where memory consists of blocks each capable of holding  $B$  data items. The cost of an algorithm or data structure is measured in terms of the number of block accesses, referred to as *I/Os*. When considering the  $k$ -NN problem we let  $B$  denote the number of vectors that fit in a block.

Distributing similarity search onto many machines has also been considered as a way of scaling to large data sets [10, 21, 29]. We can interpret a static, external memory data structure as an abstraction of a large, distributed system in which each server holds  $B$  data items (and information associated with them). In this context the parameter  $B$  may be relatively large in comparison to  $n$ . For example, to store the  $n = 10^9$  DEEP1B vectors of dimension  $d = 96$ , and associated raw data, we could imagine  $s = 10^4$  servers each holding  $B = 10^6$  data items (so that on average each item is replicated 10 times to achieve redundancy). The number of I/Os needed to answer a query then equals the number of servers that need to be involved when answering a  $k$ -NN query.

Most lower bounds for I/O efficient algorithms are shown under the assumption that data items are *indivisible* in the sense that they are treated as atomic units, and that a block contains a well-defined set of at most  $B$  data items. The *indexability model* [17, 18], introduced at PODS '97, formalized external memory data structures for queries that return a set of  $k$  data items under the indivisibility assumption. For a given data structure, the complexity of a query family  $Q$  is the smallest number of blocks that must be read to answer a query, in the worst case over all queries  $q \in Q$ . There does not need to be a constructive procedure for *identifying* the correct blocks. In particular, the nearest neighbor problem ( $k = 1$ ) is trivial in the indexability model since the block containing the answer to the query is given for free; the search aspect is completely removed from consideration, and the algorithm would return the block in one I/O. Though the original indexability model does not accommodate notions of approximation, it can be naturally extended to the setting where there is a set of at least  $k$  elements that are valid answers to a query and we are required to return any  $k$  of them.

## 1.1 Our results

The complexity of general  $(c, k)$ -NN queries in the I/O model lies between two extremes:

- There exist  $\lceil k/B \rceil$  blocks (or servers) that contains a set of  $k$  valid answers to the query, and
- No block (or server) contains more than one valid query answer, so  $k$  block reads are needed.

Since we do not care about constant factors we can, for simplicity, assume that  $k < n/2$ , since otherwise a trivial brute-force algorithm that reads all points is optimal within a factor of 2. We give several upper and lower bounds for  $(c, k)$ -NN that suggest a dichotomy for polynomial space data structures. For various choices of metric

Metric	Dimension	Approximation factor	I/O bound
any	any	any	$\geq \lceil k/B \rceil$
$\ell_\infty$	$\Omega(k \log(n/k))$	$3 - \varepsilon$	$\Omega(k)$
Hamming	$\Omega(k \log(n/k))$	$1 + \frac{1}{4k}$	$\Omega(k)$
any	any	3	$\leq \lceil k/B \rceil$
Hamming	any	$1 + \varepsilon$	$\leq \lceil k/B \rceil$

**Table 1: Our I/O lower bounds (first three rows) and upper bounds (last two rows) on the  $(c, k)$ -NN problem for data structures using polynomial space, where  $\varepsilon > 0$  can be any constant. The lower bound for  $\ell_\infty$  holds assuming  $k < n^{1-\varepsilon}$ .**

space, number of dimensions, and approximation factor we see that it is *either* possible to achieve  $O(k/B)$  I/Os, or  $\Omega(k)$  I/Os are provably required.

Our results are summarized in Table 1. For our lower bounds, the dimension column represents the minimum (asymptotic) number of dimensions required for the bounds to hold. Since it is possible to decrease the I/O complexity in lower dimensions [35], a condition on the number of dimensions is needed. Our upper bounds do not depend on the number of dimensions, except indirectly through the definition of  $B$  as the number of vectors in a block. We stress that the I/O upper bounds for an indexing scheme do not imply the existence of a data structure with the same guarantees – for a data structure in the I/O model we would expect an additional search cost. Our main theorems are:

**THEOREM 1.1. (L-infinity metric lower bound)** Consider an indexing scheme for  $(c, k)$ -NN in  $d$ -dimensional  $\ell_\infty$  space with  $c < 3$  and with worst case query time  $\lceil k/\alpha \rceil$  I/Os, where  $1 \leq \alpha \leq B$ . For sufficiently large  $d = O(k \log(n/k))$ , the indexing scheme must use  $\Omega\left(\left(\frac{1}{2e} \sqrt{\frac{n\alpha}{kB}}\right)^\alpha\right)$  blocks of space.

**THEOREM 1.2. (General metric indexing scheme)** Given a set  $P$  of  $n$  points in any metric space, there exists a 3-approximate indexing scheme that uses  $n\lceil k/B \rceil$  blocks of space (where  $B$  is the block size) and returns 3-approximate,  $k$  nearest neighbors in optimal  $\lceil k/B \rceil$  I/Os.

**THEOREM 1.3. (Hamming metric lower bound)** Consider an indexing scheme for  $(c, k)$ -NN in  $d$ -dimensional Hamming space with  $c < 1 + \frac{1}{4k}$  and with worst case query time  $\lceil k/\alpha \rceil$  I/Os, where  $1 \leq \alpha \leq B$ . For sufficiently large  $d = O(k \log(n/k))$ , the indexing scheme must use  $\Omega\left(\left(\frac{1}{2e} \sqrt{\frac{n\alpha}{kB}}\right)^\alpha\right)$  blocks of space.

**THEOREM 1.4. (Hamming indexing scheme)** Given a set  $P$  of  $n$  points in  $d$ -dimensional Hamming space and a constant  $c > 1$ , there exists a  $c$ -approximate indexing scheme that uses  $O(n^{O(1)}d\lceil k/B \rceil)$  blocks of space and returns  $c$ -approximate,  $k$  nearest neighbors in optimal  $\lceil k/B \rceil$  I/Os.

*Discussion.* We are not aware of any studies of  $c$ -approximate,  $k$ -nearest neighbors in the I/O model. However, our lower bounds suggest that such data structures cannot simultaneously have good output sensitivity (I/O complexity in terms of  $k/B$ ) and low space usage. The currently best data structures for high-dimensional,  $c$ -approximate  $r$ -near neighbor (return 1 point within known distance  $r$ ) have query time lower bounded by  $(n/B)^{f(c)}$  I/Os, for some non-increasing function  $f$  with  $f(1) = 1$  that depends on the metric. Depending on the choice of  $c$  and  $k$  this “search cost” may be smaller or larger than the cost of reporting the  $c$ -approximate  $k$ -nearest neighbors in the indexability model. In general, we would expect the reporting cost to dominate only when  $c$  and  $k$  are sufficiently large.

Our lower bounds assume that the number of dimensions is sufficiently large. An attempt to bypass the lower bounds would thus be to utilize some kind of *dimension reduction*, such as the Johnson-Lindenstrauss lemma. Unfortunately, this will not work in the settings of Theorems 1.1 and 1.3. This is because dimension reduction is not possible for  $\ell_\infty$  [28], and because dimension reduction for Hamming space with distortion  $c$  requires  $\Omega(\log(n)/(c-1)^2)$  dimensions [26], which is  $\Omega(k^2)$ .

## 1.2 Related work

*Lower bounds on nearest neighbors in restricted models.* A well-known work of Berchtold et al. [12] analyzes the performance of certain types of nearest neighbor indexes on *random* data sets. More recently, Pestov and Stojmirović [33] and Pestov [32] showed lower bounds for high-dimensional similarity search in a model of data structures that encompasses many tree-based indexing methods. These results do not consider approximation, and their algorithmic models do not encompass modern algorithmic approaches to approximate similarity search such as locality-sensitive hashing.

*Data structure lower bounds based on indivisibility.* There is a rich literature, starting with the seminal paper of Aggarwal and Vitter [3], giving lower bounds on I/O-efficiency under an indivisibility assumption. Such results in the context of data structures are known for dynamic dictionaries [37], planar point enclosure [7], range sampling [20], and many variants of orthogonal range reporting [1, 2, 17, 27, 38]. Below we elaborate on the most closely related works on orthogonal range queries.

To our best knowledge the high-dimensional  $k$ -NN problem has not been *explicitly* studied in the indexability model [17]. However, in  $n$ -dimensional Hamming space it is straightforward to use the  $k$ -set workload technique of [17] to show that even obtaining  $k-1$  I/Os is not possible (for  $k \ll n$ ) unless the indexing scheme uses quadratic space. Our lower bound technique is a generalization of the  $k$ -set workload that allows us to deal with approximation as well as space usage larger than quadratic. It also allows us to show lower bounds all the way down to  $O(k \log n)$  dimensions, as opposed to  $n$  dimensions.

*Orthogonal range queries.* *Orthogonal range reporting* in  $d$  dimensions asks to report all points in  $P$  lying inside a query range  $Q$  that is a cross product of intervals. Note that  $k$ -NN in the  $\ell_\infty$  metric is the special case of orthogonal range reporting where all intervals have the same length. Hellerstein et al. [17] showed that in order

to answer orthogonal range reporting queries in  $O(k/B^{1-\varepsilon})$  I/Os, for some  $\varepsilon \in (0, 1)$ , the data structure needs to use space  $\Omega(n/\varepsilon^{d-1})$ . In particular, in dimension  $d = \omega(\log n)$  a polynomial space data structure needs  $k/B^{o(1)}$  I/Os.

*Approximate*  $d$ -dimensional range reporting has been studied in the I/O model: Streppel and Yi [35] show that for a query rectangle  $q$  and constant  $\varepsilon > 0$ , allowing the data structure to report points at distance up to  $\varepsilon \cdot \text{diam}(q)$  from  $q$  makes it possible to report  $k$  points in  $2^{O(d)} + O(k/B)$  I/Os (plus a logarithmic search cost, which does not apply to the indexability model). To our best knowledge, no lower bounds were known for approximate range queries before our work.

*Lower bounds based on computational assumptions.* Unconditional lower bounds for  $(c, k)$ -NN in the cell probe model [11, 13, 14, 31] only match upper bounds in the regime where space is very large. To better understand the complexity for, say, sub-quadratic space usage, a possibility is to base lower bounds on computational assumptions such as the Strong Exponential Time Hypothesis (SETH), or the weaker Orthogonal Vectors Hypothesis (OVH). Recently, Rubinstein [34] showed that under either of these hypotheses, for each constant  $\delta > 0$ , achieving an approximation factor of  $c = 1 + o(1)$  is not possible for a data structure with polynomial space and construction time unless the query time exceeds  $n^{1-\delta}$ . Already in 2001 Indyk showed that in the  $\ell_\infty$  metric,  $(c, k)$ -NN with approximation factor  $c < 3$  is similarly hard [22]. (Though Indyk links the lower bound to a different problem, it can be easily checked that the same conclusion follows from the more recent SETH and OVH assumptions.)

*Upper bounds.* The I/O complexity of the near neighbor problem was studied by Gionis et al. [15], focusing on Hamming space. (Their approach extends to other spaces that have good locality-sensitive hash functions.) For every approximation factor  $c > 1$ , they show that  $O((n/B)^{1/c})$  I/Os suffices to retrieve *one* near neighbor, using a data structure of size subquadratic in  $n$ . It seems that the same algorithm can be adapted to return  $k$  near neighbors at an additional cost of  $O(k/B)$  I/Os. Tao et al. [36] extended these results to handle nearest neighbor queries, but they do not consider the case of  $k$  nearest neighbors.

For the Euclidean and Hamming metrics with constant approximation factor  $c > 1$  it is known how to get  $n^{o(1)}$  query time for the near neighbor problem ( $k = 1$ ) with a polynomial space data structure (see [6] and its references). For  $c > \sqrt{3}$  Kapralov [24] even showed how to achieve this by a *single* probe to the data structure, returning a pointer to the result.

*Organization:* The rest of this paper is organized as follows. In Section 2 we develop notation and proceed in Section 2.1 to extend the indexability result of [19] to approximate indexing schemes. Section 3 describes the indexing scheme promised in Theorem 1.2, followed by Section 4 that proves lower bounds in the  $L$ -infinity metric (Theorem 1.1). Section 5 contains our results on the Hamming metric (Theorem 1.3 and Theorem 1.4), followed by conclusion and open problems in Section 6.

## 2 PRELIMINARIES

The external memory model of computation (due to Aggarwal and Vitter [3]) has a main memory of size  $M$  and an infinite external memory, both realized as arrays. Data is stored in external memory, and is transferred to/from main memory (where computation happens) in I/Os, or *block transfers*, where a block holds  $B$  data items. Computation in main memory is free — the cost of an algorithm is the number of I/Os it performs.

We will be using the following definitions from Hellerstein et al. [17, 18]. For brevity we will refer to a subset of  $I$  of size  $b$  as a *b-subset*.

**DEFINITION 2.1 (WORKLOAD).** *A workload  $W = (D, I, Q)$  consists of a non-empty set  $D$  (the domain), a nonempty finite set  $I \subseteq D$  (the instance, whose size we denote by  $n$ ), and a set  $Q$  of subsets of  $I$  (the query set).*

**DEFINITION 2.2 (INDEXING SCHEME).** *An indexing scheme  $S = (W, \mathcal{B})$  consists of a workload  $W = (D, I, Q)$  and a set  $\mathcal{B}$  of  $B$ -subsets of  $I$ , where  $B$  is the block size of the indexing scheme.*

**DEFINITION 2.3 (COVER SET).** *A cover set  $C_Q \subseteq \mathcal{B}$  for a query  $Q \in Q$  is a minimum-size subset of the blocks such that  $Q \subseteq \bigcup_{b \in C_Q} b$ .*

We will assume that the blocks are chosen such that  $C_Q$  exists for every query  $Q \in Q$ .

**DEFINITION 2.4 ( $\lambda$ -SET WORKLOAD).** *The  $\lambda$ -set workload is a workload with instance  $I = \{1, \dots, n\}$  whose query set  $Q$  is the set of all  $\lambda$ -subsets of the instance.*

While [17, 18] measure performance in terms of *redundancy* and *access overhead*, we find it more natural to define performance in the same way as the I/O model. The *space usage* of an indexing scheme is the number of blocks,  $|\mathcal{B}|$ . The *query time* of  $Q \in Q$  is the size of a cover set,  $|C_Q|$ . Observe that the time of a query of size  $\lambda$  can range from  $\lceil \lambda/B \rceil$  I/Os to  $\lambda$  I/Os, depending on how many blocks are needed to cover it.

All static data structures we know of in the external memory model with a given space usage and query cost translate directly to an indexing scheme with the same, or better, performance. This is because these data structures store, or can be adapted to store,  $O(B)$  vectors explicitly in each block, such that the set of result vectors is explicitly present in the blocks that are read when answering a query. This means that any lower bounds in the indexability model strongly suggest lower bounds also in the external memory model.

### 2.1 Approximate indexing schemes

**DEFINITION 2.5 ( $c$ -APPROXIMATE  $k$ -NEAREST NEIGHBORS PROBLEM,  $(c, k)$ -NN).** *Let  $M$  be a metric space with distance function  $d(\cdot, \cdot)$ , and  $c > 1$  be a constant. Given  $P \subset M$ ,  $|P| = n$ , and a positive integer  $k \leq n$ , construct a data structure that upon receiving a query  $q$ , returns  $Q \subset P$ ,  $Q = \{q^1, \dots, q^k\}$ , such that for all  $q^i \in Q$ ,  $d(q, q^i) \leq c d(q, p_q^k)$ , where  $p_q^k \in P$  is the  $k$ th nearest neighbor of  $q$  in  $P$ .*

We will consider indexing schemes for  $(c, k)$ -NN that can depend on the parameters  $c$  and  $k$ , i.e., the parameters are known when the data structure is constructed. The set of points to be stored is denoted by  $P$ . Observe that for  $(c, k)$ -NN, there can be at most

$\binom{n}{k}$  distinct query answers ( $\lambda = k$  in the  $\lambda$ -set workload above). Intuitively, the more queries there are in a workload, the higher is the space needed by the data structure. But in the  $(c, k)$ -NN problem, for a query  $q$  there is no one “right” answer, as any set of  $k$  approximate near neighbors form a valid output. Thus we find that the definitions in the original indexability model need to be extended to capture approximate data structures. It turns out that even though approximate near neighbors are allowed, it is possible to ensure that at least half of the  $k$  returned points are among the  $k$  nearest to the query point. This motivates the following definition:

**DEFINITION 2.6 (RELAXED  $\lambda$ -SET WORKLOAD).** *The relaxed  $\lambda$ -set workload is a workload whose instance is  $\{1, \dots, n\}$ . The query set is the set of all  $\binom{n}{\lambda}$  subsets of size  $\lambda$ . Given a query  $q$  corresponding to such a  $\lambda$ -subset  $Q$ , the indexing scheme must report  $\lambda$  elements, at least half of which must come from  $Q$ .*

We next show a space-I/O tradeoff for the relaxed workload.

**LEMMA 2.1 (RELAXED SET WORKLOAD TRADEOFF).** *Any indexing scheme for the relaxed  $\lambda$  set workload with a space usage of  $s$  blocks and a query time of  $t$  block accesses must have*

$$t \geq \lambda \log_s \left( \frac{1}{2e} \sqrt{\frac{n}{tB}} \right)$$

**PROOF.** Recall that by the relaxation, the indexing only needs to find a subset  $\hat{Q} \subset Q$  of size  $|\hat{Q}| = \lambda/2$  in the index. In other words, with each subset  $\hat{Q}$  that can be retrieved by the index, the algorithm can add any  $\lambda/2$  elements and arrive at a valid query  $Q$ .

We upper bound the total number of distinct  $\lambda$ -sets reported by an indexing scheme using  $s$  space and  $t$  query I/Os as follows:

- choose the set of  $t$  blocks to retrieve from the index ( $\binom{s}{t}$  choices),
- choose the  $\lambda/2$  elements to use from these at most  $tB$  distinct elements (at most  $\binom{tB}{\lambda/2}$  choices),
- choose  $\lambda/2$  arbitrary other elements (at most  $\binom{n}{\lambda/2}$  choices).

The total number of such combinations should be at least  $\binom{n}{\lambda}$ , which is the possible set of queries, or  $\lambda$ -subsets. Using the inequalities  $\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k \leq n^k$  this gives us:

$$\begin{aligned} \binom{s}{t} \binom{tB}{\lambda/2} \binom{n}{\lambda/2} &\geq \binom{n}{\lambda} \\ \Rightarrow s^t \left( \frac{tBe}{\lambda/2} \right)^{\lambda/2} \left( \frac{ne}{\lambda/2} \right)^{\lambda/2} &\geq \left( \frac{n}{\lambda} \right)^\lambda \\ \Rightarrow s^{2t/\lambda} \left( \frac{tBe}{\lambda/2} \right) \left( \frac{ne}{\lambda/2} \right) &\geq \left( \frac{n}{\lambda} \right)^2 \\ \Rightarrow s^{2t/\lambda} &\geq \frac{n}{4e^2 t B} \\ \Rightarrow t &\geq \lambda \log_s \left( \frac{1}{2e} \sqrt{\frac{n}{tB}} \right). \end{aligned} \tag{1}$$

□

One may be interested in a query time of  $\lceil \lambda/\alpha \rceil$  I/Os, where  $1 \leq \alpha \leq B$ .

**COROLLARY 2.2.** *Any indexing scheme for the relaxed  $\lambda$ -set workload with worst case query time of  $\lceil \lambda/\alpha \rceil$  I/Os, where  $1 \leq \alpha \leq B$  must use  $\Omega\left(\left(\frac{1}{2e}\sqrt{\frac{n\alpha}{\lambda B}}\right)^\alpha\right)$  blocks of space.*

This lower bound is essentially tight: Achieving a query time of  $\lceil \lambda/\alpha \rceil$  I/Os using  $O\left(\binom{n}{\alpha}\right)$  space is easy by just preprocessing all  $\alpha$  combinations of the  $n$  items. By doing an analogous calculation for the exact workload, we are able to give the following tradeoff for the standard  $\lambda$ -set workload.

**LEMMA 2.3 ( $\lambda$ -SET WORKLOAD TRADEOFF).** *Any indexing scheme for the  $\lambda$ -set workload with worst case query time of  $\lceil \lambda/\alpha \rceil$  I/Os, where  $1 \leq \alpha \leq B$  must use  $\Omega\left(\left(\frac{n\alpha}{e\lambda B}\right)^\alpha\right)$  blocks of space.*

Lemma 2.3 generalizes Theorem 7.2 in [17] which considered the case of  $\alpha = 1$ .

### 3 A 3-APPROXIMATION INDEXING SCHEME FOR GENERAL METRICS

We prove Theorem 1.2 in this section, that asserts that (a relatively simple) indexing scheme provides a 3-approximation for the  $k$ -NN problem in any metric space. Note that we are only presenting an indexing scheme as opposed to a data structure; i.e., we assume that once the query is given, an oracle provides the smallest set of blocks that contain a valid answer to the query.

Let  $P$  be a set of  $n$  points in a metric space. Consider the indexing scheme  $I$  that consists of the set of  $\ell = n\lceil k/B \rceil$  blocks  $\mathcal{B} = \{b_1, b_2, \dots, b_\ell\}$  where we store  $p_i \in P$  and the set  $P_i = \{p_{i,1}, p_{i,2}, \dots, p_{i,k}\}$  of the  $k$ -nearest neighbors of the point  $p_i$  (including itself). This requires  $\lceil k/B \rceil$  blocks per element of  $P$ , as claimed. For a query point  $q$  let  $p_{i^*}$  be the nearest neighbor of  $q$ . The oracle then returns the set  $P_{i^*}$ , i.e., the  $k$  nearest neighbors of  $q$ 's nearest neighbor, using  $\lceil k/B \rceil$  I/Os.

**THEOREM 3.1.** *Let  $\{p_1^q, \dots, p_k^q\}$  be the set of the exact  $k$  nearest neighbors of a query  $q$ , and let  $\{p_{i^*,1}, \dots, p_{i^*,k}\}$  be the set of points returned by the indexing scheme described above. Then for any  $1 \leq j \leq k$ ,  $d(q, p_{i^*,j}) \leq 3d(q, p_k^q)$ . That is, all the returned points are within a factor 3 of the distance of the query to its  $k$ th nearest neighbor.*

**PROOF.** The proof is a case analysis. Let  $D$  be the smallest ball centered at  $p_1^q := p_{i^*}$  that contains  $p_{i^*,k}$  (the “ $k$  nearest neighbor disk” of  $p_1^q$ ). The cases are:

(1) Both  $q$  and  $p_k^q$  are outside  $D$ . For each  $j$ , we have that

$$\begin{aligned} d(q, p_{i^*,j}) &\leq d(p_{i^*,j}, p_1^q) + d(p_1^q, q) \quad (\text{triangle inequality}) \\ &\leq d(p_{i^*,j}, p_1^q) + d(q, p_k^q) \quad (p_k^q \text{ is farther from } q \text{ than } p_1^q) \\ &\leq d(q, p_1^q) + d(q, p_k^q) \quad (q \text{ is outside } D) \\ &\leq 2d(q, p_k^q) . \end{aligned}$$

(2)  $q$  is inside  $D$ , and  $p_k^q$  is outside  $D$ . For each  $j$ , we have that

$$\begin{aligned} d(q, p_{i^*,j}) &\leq d(p_{i^*,j}, p_1^q) + d(p_1^q, q) \quad (\text{triangle inequality}) \\ &\leq d(p_{i^*,j}, p_1^q) + d(q, p_k^q) \quad (p_k^q \text{ is farther from } q \text{ than } p_1^q) \\ &\leq d(p_k^q, p_1^q) + d(q, p_k^q) \quad (p_k^q \text{ is outside } D) \\ &\leq d(p_k^q, q) + d(q, p_1^q) + d(q, p_k^q) \quad (\text{triangle inequality}) \\ &\leq 3d(q, p_k^q) . \end{aligned}$$

(3) Both  $q$  and  $p_k^q$  are inside  $D$ . In this case, we first claim that it is sufficient to consider that for some  $i \notin \{1, k\}$ ,  $p_i^q$  is outside  $D$ . If that was not the case, we have reported all  $\{p_i^q\}$ , i.e., the exact  $k$  nearest neighbors of  $q$ . For each  $j$  we now have that

$$\begin{aligned} d(q, p_{i^*,j}) &\leq d(q, p_1^q) + d(p_1^q, p_{i^*,j}) \quad (\text{triangle inequality}) \\ &\leq d(q, p_k^q) + d(p_1^q, p_{i^*,j}) \quad (p_k^q \text{ is farther from } q \text{ than } p_1^q) \\ &\leq d(q, p_k^q) + d(p_1^q, p_i^q) \quad (p_i^q \text{ is outside } D) \\ &\leq d(q, p_k^q) + d(p_1^q, q) + d(q, p_i^q) \quad (\text{triangle inequality}) \\ &\leq 3d(q, p_k^q) . \end{aligned}$$

(4)  $q$  is outside  $D$ , and  $p_k^q$  is inside  $D$ . In this case we have that

$$\begin{aligned} d(q, p_{i^*,j}) &\leq d(p_{i^*,j}, p_1^q) + d(p_1^q, q) \quad (\text{triangle inequality}) \\ &\leq d(q, p_1^q) + d(p_1^q, q) \quad (q \text{ is outside } D; p_{i^*,j} \text{ is inside}) \\ &\leq 2d(q, p_1^q) \quad (p_1^q \text{ is farther from } q \text{ than } p_k^q) \end{aligned}$$

□

*Tightness of the analysis:* Next, we give an example of  $n = k + 1$  points where the indexing scheme above does not achieve a better approximation factor than  $3(1 - \varepsilon)$ , for arbitrary  $\varepsilon > 0$ . Consider the scenario when all points including the query lie on a line:

$$q = 1/2, \quad p_1^q = 0, \quad p_k^q = 1 + \varepsilon/2, \quad p_{k+1}^q = -1, \quad \text{and} \\ p_i^q = -\varepsilon i/2k \text{ for all } i \notin \{1, k, k+1\} .$$

The disk  $D$  contains all points but  $p_k^q, p_{k+1}^q$  will be reported, and  $d(q, p_{k+1}^q) = 3/2 > 3(1 - \varepsilon)(1/2 + \varepsilon/2) = 3(1 - \varepsilon)d(q, p_k^q)$ .

### 4 LOWER BOUNDS

In this section we prove Theorems 1.1 and 1.3.

#### 4.1 Lower bound for the L-infinity metric

We warm up with a proof of a lower bound in  $n$  dimensions, a slight variant of a reduction by Indyk [22].

**LEMMA 4.1.** *Consider the approximate  $k$ -NN problem in  $\mathbb{R}^n$  with  $\ell_\infty$  metric. There exists a set of  $n$  points  $p_i$  for  $i \in \{1, \dots, n\}$  that has a  $k$ -set workload. Specifically, for every subset  $I \subset [n]$  there exists a query point  $q_I$  such that the  $\ell_\infty$ -distance is*

$$\|p_i - q_I\|_\infty = \begin{cases} 1/2 & \text{for } i \in I \\ 3/2 & \text{for } i \notin I \end{cases} .$$

*Hence, the  $(c, k)$ -NN for  $c < 3$  in  $n$ -dimensional  $\ell_\infty$  space leads to a  $k$ -set workload.*

PROOF. The set consists of the  $n$  unit vectors  $p_i = e_i$  (where only the  $i$ th entry is 1 and all other entries are 0). For a set  $I \subset \{1, \dots, n\}$  the query vector is defined as

$$q_I = \begin{cases} +1/2 & \text{for } i \in I \\ -1/2 & \text{for } i \notin I \end{cases}$$

□

To deterministically reduce the dimensionality of the space we use an expander and switch to relaxed  $k$ -set workloads. Expanders were previously used for deterministic embeddings of Euclidean space into  $\ell_1$  by Indyk [23]. There is a vast literature on expanders and the results we are using are standard by now. For the sake of concreteness, we take the definitions and precise results almost literally from [30]. We define  $(m, \delta, 1/3)$ -expander graphs and state some results concerning these graphs. For the rest of this paper we will assume  $\delta$  to divide  $1/3$ , as this makes statements and proofs simpler. This will be without loss of generality, as the statements we show do not change when rounding  $\delta$  down to the nearest such value. Let  $G = (U, V, E)$  be a bipartite graph with left vertex set  $U$ , right vertex set  $V$ , and edge set  $E$ . We denote the set of neighbors of a set  $S \subseteq U$  by  $\Gamma(S) = \cup_{v \in S} \{v \mid (s, v) \in E\}$ , and use  $\Gamma(x)$  as a shorthand for  $\Gamma(\{x\})$ ,  $x \in U$ .

DEFINITION 4.1 (DEFINITION 3 OF [30]). *A bipartite graph  $G = (U, V, E)$  is  $\delta$ -regular if the degree of all nodes in  $U$  is  $\delta$ . A bipartite  $\delta$ -regular graph  $G = (U, V, E)$  is an  $(m, \delta, 1/3)$ -expander if for each  $S \subseteq U$  with  $|S| \leq m$  it holds that  $|\Gamma(S)| \geq (1 - 1/3)\delta|S|$ .*

LEMMA 4.2 (COROLLARY 5 IN [30]). *For every constant  $\varepsilon > 0$  there exists an  $(m, \delta, \varepsilon)$ -expander  $G = (U, V, E)$  with  $|U| = u$ ,  $\delta = O(\log(2u/m))$  and  $|V| = O(m \log(2u/m))$ .*

Next, we discuss how to give an analogue of the hard query set in Lemma 4.1 with  $O(k \log n)$  dimensions.

LEMMA 4.3. *Let  $n, k(n)$  be arbitrary integer parameters. Consider  $(c, k)$ -NN with  $c < 3$ . There exists a set  $D$  of  $n$  points in dimension  $O(k \log n)$  such that for any  $I \subset [n]$  with  $|I| = k$  there exists a query point  $q_I$  such that the set  $I' = \{p \in D \mid \|p - q_I\|_\infty \leq 1/2\}$  of potential answer points has  $I \subset I'$  and  $|I'| \leq 3k/2$ .*

PROOF. Fix  $n$  and  $k(n)$ . Let  $m = k(n)$  and let  $G = (U, V, E)$  be an  $(m, \delta = O(\log(n/m)), 1/3)$ -expander with  $|U| = n$  and  $|V| = d = O(k\delta)$ . For concreteness we take  $U = [n]$ .

Construct the set of points  $p_1, \dots, p_n$  where

$$(p_i)_j = \begin{cases} 1 & \text{if } j \in \Gamma(i) \\ 0 & \text{otherwise} \end{cases}$$

Define the query point for set  $I$  with  $|I| = k$  where

$$(p_I)_j = \begin{cases} +1/2 & \text{for } j \in \Gamma(I) \\ -1/2 & \text{otherwise} \end{cases}$$

It is easy to see that  $\|p_i - p_I\|_\infty = 1/2$  for all  $i \in I$ , so  $p_I$  has at least  $k$  neighbors at distance  $1/2$ . It remains to show that this leads to a relaxed  $k$ -set workload, i.e., that for any  $q_I$  any set of  $1/2$ -near points in the set has at least  $k/2$  points in common with  $I$ . Fix a subset  $I$  and consider the query point  $q_I$ . Let  $I' = \{i \mid \|p_i - p_I\|_\infty \leq 1/2\}$  be the indices of the points of with distance at most  $1/2$  to  $q_I$ . Observe that  $I \subseteq I'$  by construction. Let  $I^* = I' \setminus I$  be the indices of “unintended near points”. Observe that every point  $p_i$  that does

not fulfill  $\Gamma(i) \subset \Gamma(I)$  has  $\|p_i - p_I\|_\infty = 3/2$  and cannot be reported since  $c < 3$ . Hence

$$I^* = \{i \mid \Gamma(i) \subseteq \Gamma(I)\} \setminus I.$$

Observe  $|\Gamma(I)| \leq \delta|I|$  and  $|\Gamma(I \cup I^*)| \geq (1 - 1/3)\delta|I \cup I^*|$ , by the definition of  $G$  and  $|\Gamma(I \cup I^*)| = |\Gamma(I)|$  by the definition of  $I^*$ . Combining this we get  $\delta|I| \geq \frac{2}{3}\delta|I \cup I^*|$ , or  $|I| \geq \frac{2}{3}(|I| + |I^*|)$ . That is,  $\frac{1/3}{2/3}|I| \geq |I^*|$ , as desired. Hence we have a relaxed  $k$ -set workload. □

This means that we can apply Lemma 2.1 and Corollary 2.2:

COROLLARY 4.4. *(Theorem 1.1) Any indexing scheme for  $(c, k)$ -NN in  $O(k \log(n/k))$ -dimensional  $\ell_\infty$  space with  $c < 3$  with worst case query time of  $\lceil k/\alpha \rceil$  I/Os, where  $1 \leq \alpha \leq B$  must use  $\Omega\left(\left(\frac{1}{2e} \sqrt{\frac{n\alpha}{kB}}\right)^\alpha\right)$  blocks of space.*

## 4.2 Lower Bound for the Hamming metric

We now prove Theorem 1.3, giving a lower bound on indexes on sets of vectors in Hamming space with approximation factor  $c > 1$ . This directly implies a lower bound in the  $\ell_1$  metric, as well as a lower bound for  $\ell_2$  with approximation factor  $\sqrt{c}$ .

LEMMA 4.5. *Consider the approximate  $k$ -NN problem in Hamming space of dimension  $d$ . There exists a set of  $n$  points  $p_i, i \in \{1, \dots, n\}$  of dimension  $d = n$  that is a  $k$ -set workload, i.e., for every subset  $I \subset [n]$  with  $|I| = k$  there exists a query point  $q_I$  such that the Hamming (or  $\ell_1$ ) distance is:*

$$\|p_i - q_I\|_1 = \begin{cases} k-1 & \text{for } i \in I \\ k+1 & \text{for } i \notin I \end{cases}.$$

Hence, the  $(c, k)$ -NN for  $c < 1 + 2/(k-1)$  in  $n$ -dimensional Hamming space leads to a  $k$ -set workload.

PROOF. The set consists of the  $n$  unit vectors  $e_i$ . The query vector for set  $I$  is the characteristic vector of  $I$ . It is easy to verify that the distances from the query to the vectors in  $I$  is  $k-1$ , and to those not in  $I$  is  $k+1$ . Note that  $(k+1)/(k-1) = 1 + 2/(k-1)$ , which gives the bound on  $c$ . □

LEMMA 4.6. *Consider the approximate  $k$ -NN problem in Hamming space of dimension  $d$ . There exists a set of  $n$  points  $p_i, i \in \{1, \dots, n\}$ , and  $d = O(k \log(n/k))$  that is a relaxed  $k$ -set workload, i.e., for every subset  $I \subset [n]$  with  $|I| = k$  there exists a query point  $q_I$  such that the Hamming ( $\ell_1$ )-distance is*

$$\|p_i - q_I\|_1 = \begin{cases} \leq d_I & \text{for } i \in I \\ \geq d_I \left(1 + \frac{1}{4(k-1)}\right) & \text{for } i \notin I' \end{cases}$$

for an  $I' \supset I$  with  $|I'| \leq 1.5|I|$ . Hence,  $c < 1 + \frac{1}{4(k-1)}$  leads to a relaxed  $k$ -set workload for  $(c, k)$ -NN in  $O(k \log n)$ -dimensional Hamming space.

PROOF. Fix  $n$  and  $k(n)$ . Let  $G = (U, V, E)$  be an  $(m = k(n), \delta = O(\log n), 1/4)$ -expander with  $|U| = n$  and  $|V| = d = O(k \log n)$ , which exists by Lemma 4.2. For concreteness we take  $U = [n]$ .

Construct the set of points  $p_1, \dots, p_n$  in the same way as earlier:

$$(p_i)_h = \begin{cases} 1 & \text{for } h \in \Gamma(i) \\ 0 & \text{otherwise} \end{cases}.$$

We define the query point  $q_I$  for each set  $I$  with  $|I| = k$  where

$$(q_I)_h = \begin{cases} 1 & \text{for } h \in \Gamma(I) \\ 0 & \text{otherwise} \end{cases}.$$

In other words, the vectors are the characteristic vectors of the neighbor-sets in  $G$ .

We define  $d_I := \|p_j - q_I\|_1$ . For  $i \in I$  we have  $d_i = |\Gamma(I)| - \delta$  and

$$d_I \geq k\delta(1 - 1/4) - \delta = \delta(k(1 - 1/4) - 1).$$

For any  $j$  (in particular  $j \notin I$ ) we have  $d_j = \|p_j - q_I\|_1 = |\Gamma(I)| - \delta + 2|\Gamma(j) \setminus \Gamma(I)|$ , leading to the definition  $\gamma_j = |\Gamma(j) \setminus \Gamma(I)|$ . We can set the distance threshold ratio such that the number of unintended neighbors is again at most  $k/2$ :

We can set the distance threshold ratio to  $1 + 1/(4(k - 1))$  such that the set of unintended near neighbors are the  $j$  with  $\gamma_j < \delta/4$ :

$$d_j/d_I = 1 + \gamma_j/d_I \leq 1 + 1/(4(k - 1)) \Rightarrow \gamma_j \leq d_I/(4(k - 1)) \leq \delta/4$$

To calculate the set of indices of unintended neighbors we define

$$I^* = \{j \mid \gamma_j \leq \delta/4\} \setminus I.$$

Then

$$\begin{aligned} |\Gamma(I \cup I^*)| &\leq \delta|I| + (\delta/4)|I^*| \\ |\Gamma(I \cup I^*)| &\geq (1 - 1/4)\delta|I \cup I^*| = \frac{3}{4}\delta(|I| + |I^*|) \end{aligned}$$

leading to

$$\begin{aligned} \delta|I| + (\delta/4)|I^*| &\geq \frac{3}{4}\delta(|I| + |I^*|) \\ \Rightarrow (\delta/4)|I^*| &\geq \frac{3}{4}\delta|I^*| - |I|\delta/4 \\ \Rightarrow |I|\delta/4 &\geq \frac{3}{4}\delta|I^*| - (\delta/4)|I^*|. \end{aligned}$$

Hence

$$|I^*| \leq |I| \frac{1/4}{\frac{3}{4} - 1/4} = k/2.$$

This means that the described workload is a relaxed  $k$ -set workload. Applying Lemma 2.1 and Corollary 2.2 now completes the proof of Theorem 1.3.  $\square$

## 5 HAMMING METRIC INDEXING SCHEME

The general 3-approximate indexing scheme described in Section 3 can be improved for specific metrics. In this section we prove Theorem 1.4. For any given approximation factor  $c > 1$  we wish to construct an indexing scheme that answers  $(c, k)$ -NN queries in  $d$ -dimensional Hamming space using polynomial space and with  $\lceil k/B \rceil$  I/Os. Our construction is an application of the dimension reduction technique of Kushilevitz et al. [25].

For each  $r \in \{1, \dots, d\}$  we create a data structure that handles the case where the  $k$ th closest point to  $q$  is at distance  $r$ . The data structure must report  $k$  points that have distance at most  $cr$  from  $q$ . The central idea of [25] is to use a randomized mapping

$$t : \{0, 1\}^d \rightarrow \{0, 1\}^D,$$

where  $D = O(\log n)$ , such that for each  $q \in \{0, 1\}^d$  with high probability for all  $x, y \in P$ :

$$d(q, x) \leq r \wedge d(q, y) > cr \implies d(t(q), t(x)) < d(t(q), t(y)). \quad (2)$$

(We note that the required dimension  $D$  grows as  $c$  approaches 1, hence we need to keep  $c$  fixed.) Consider the mapped multiset  $t(P) = \{t(x) \mid x \in P\}$  and create a data structure that for each  $i \in \{0, 1\}^D$  lists, for the  $k$  nearest neighbors of  $i$  in  $t(P)$ , the corresponding vectors in  $P$  (breaking ties arbitrarily), using  $\lceil k/B \rceil$  blocks. If (2)

holds then list  $i = t(q)$  contains only  $c$ -approximate  $k$ -nearest neighbors of  $q$ . To eliminate the error probability, choose  $O(d)$  such random mappings and construct corresponding data structures: With high probability there will be no query  $q \in \{0, 1\}^d$  that does not have at least one data structure that returns a correct result. If this fails for some  $q$ , start over from the beginning and choose new mappings.

The total space usage is  $O(2^D d \lceil k/B \rceil)$ , which is polynomial in  $n$  and  $d$ , as desired. Queries can be answered in  $\lceil k/B \rceil$  I/Os since we are taking full advantage of the power of the indexability model: To answer a query  $q$  it is necessary to know which mapping  $t$  can be used to answer it correctly, and where in storage the blocks with index  $t(q)$  reside.

Of course, we can also get an algorithm in the standard I/O model with a multiplicative query time overhead of  $O(d)$  by querying *all* repetitions and returning the  $k$  closest points seen.

## 6 CONCLUSION AND OPEN PROBLEMS

We have shown that nontrivial lower bounds can be shown in the indexability model, even under approximation. The main open problem that we leave is whether our hardness result for Hamming distance can be extended to approximation factor  $c = 1 + \Omega(1)$ , where the constant in  $\Omega(1)$  is independent of  $k$ . This would give an unconditional analogue of the recent conditional lower bound of Rubinstein [34].

**Acknowledgement.** We thank the anonymous reviewers for constructive feedback.

## REFERENCES

- [1] Peyman Afshani. 2012. Improved pointer machine and I/O lower bounds for simplex range reporting and related problems. In *Proceedings of 28th Symposium on Computational Geometry (SoCG)*. ACM, 339–346.
- [2] Peyman Afshani, Lars Arge, and Kasper Dalggaard Larsen. 2009. Orthogonal range reporting in three and higher dimensions. In *50th IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE, 149–158.
- [3] Alok Aggarwal and Jeffrey S Vitter. 1988. The input/output complexity of sorting and related problems. *Commun. ACM* 31 (1988), 1116–1127. Issue 9.
- [4] Josh Alman and Ryan Williams. 2015. Probabilistic Polynomials and Hamming Nearest Neighbors. In *Proceedings of 56th symposium on Foundations of Computer Science (FOCS)*. 136–150. <https://doi.org/10.1109/FOCS.2015.18>
- [5] Alexandr Andoni, Piotr Indyk, and Ilya Razenshteyn. 2018. Approximate nearest neighbor search in high dimensions. *arXiv preprint 1806.09823* (2018). Also appears in proceedings of ICM 2018.
- [6] Alexandr Andoni, Thijss Laarhoven, Ilya P. Razenshteyn, and Erik Waingarten. 2017. Optimal Hashing-based Time-Space Trade-offs for Approximate Near Neighbors. In *Proceedings of 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 47–66. <https://doi.org/10.1137/1.9781611974782.4>
- [7] Lars Arge, Vasilis Samoladas, and Ke Yi. 2009. Optimal external memory planar point enclosure. *Algorithmica* 54, 3 (2009), 337–352.
- [8] Artem Babenko and Victor Lempitsky. 2016. Efficient indexing of billion-scale datasets of deep descriptors. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. 2055–2063.
- [9] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. 2014. Neural codes for image retrieval. In *European conference on computer vision*. Springer, 584–599.
- [10] Bahman Bahmani, Ashish Goel, and Rajendra Shinde. 2012. Efficient distributed locality sensitive hashing. In *Proceedings of 21st ACM international conference on information and knowledge management (CIKM)*. ACM, 2174–2178.
- [11] Omer Barkol and Yuval Rabani. 2002. Tighter Lower Bounds for Nearest Neighbor Search and Related Problems in the Cell Probe Model. *J. Comput. Syst. Sci.* 64, 4 (June 2002), 873–896. <https://doi.org/10.1006/jcss.2002.1831>
- [12] Stefan Berchtold, Christian Böhm, Daniel A. Keim, and Hans-Peter Kriegel. 1997. A Cost Model For Nearest Neighbor Search in High-Dimensional Data Space. In *Proceedings of 6th ACM Symposium on Principles of Database Systems (PODS)*.

[13] Allan Borodin, Rafail Ostrovsky, and Yuval Rabani. 1999. Lower bounds for high dimensional nearest neighbor search and related problems. In *Proceedings of 31st Symposium on Theory of Computation (STOC)*. 312–321.

[14] Amit Chakrabarti and Oded Regev. 2004. An Optimal Randomised Cell Probe Lower Bound for Approximate Nearest Neighbour Searching. In *Proceedings of 45th IEEE Symposium on Foundations of Computer Science (FOCS '04)*. 473–482. <https://doi.org/10.1109/FOCS.2004.12>

[15] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *Proceedings of 25th International Conference on Very Large Data Bases (VLDB)*. Morgan Kaufmann, 518–529. <http://www.vldb.org/conf/1999/P49.pdf>

[16] Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. 2012. Approximate Nearest Neighbor: Towards Removing the Curse of Dimensionality. *Theory of Computing* 8, 1 (2012), 321–350. <https://doi.org/10.4086/toc.2012.v008a014>

[17] Joseph Hellerstein, Elias Koutsoupias, Daniel Miranker, Christos Papadimitriou, and Samoladas Vasilis. 2002. On a Model of Indexability and Its Bounds for Range Queries. *J. ACM* 49, 1 (2002), 35–55.

[18] Joseph Hellerstein, Elias Koutsoupias, and Christos Papadimitriou. 1997. On the Analysis of Indexing Schemes. In *Proceedings of 16th ACM Symposium on Principles of Database Systems (PODS)*. ACM, 249–256.

[19] Joseph M. Hellerstein, Jeffrey F. Naughton, and Avi Pfeffer. 1995. Generalized Search Tree for Database Systems. In *Proceedings of 21th International Conference on Very Large Data Bases (VLDB)*. ACM, 562–573.

[20] Xiaocheng Hu, Miao Qiao, and Yufei Tao. 2014. Independent range sampling. In *Proceedings of 33rd ACM Symposium on Principles of database systems (PODS)*. ACM, 246–255.

[21] Xiao Hu, Ke Yi, and Yufei Tao. 2019. Output-Optimal Massively Parallel Algorithms for Similarity Joins. *ACM Transactions on Database Systems* 44, 2 (April 2019), 1–36. <https://doi.org/10.1145/3311967>

[22] Piotr Indyk. 2001. On approximate nearest neighbors under  $l_\infty$  norm. *J. Comput. System Sci.* 63, 4 (2001), 627–638.

[23] Piotr Indyk. 2007. Uncertainty principles, extractors, and explicit embeddings of  $L_2$  into  $L_1$ . In *Proceedings of 39th ACM Symposium on Theory of Computing (STOC)*. 615–620. <https://doi.org/10.1145/1250790.1250881>

[24] Michael Kapralov. 2015. Smooth tradeoffs between insert and query complexity in nearest neighbor search. In *Proceedings of 34th ACM Symposium on Principles of Database Systems (PODS)*. ACM, 329–342.

[25] Eyal Kushilevitz, Rafail Ostrovsky, and Yuval Rabani. 2000. Efficient search for approximate nearest neighbor in high dimensional spaces. *SIAM J. Comput.* 30, 2 (2000), 457–474.

[26] Kasper Green Larsen and Jelani Nelson. 2017. Optimality of the Johnson-Lindenstrauss lemma. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 633–638.

[27] Kasper Green Larsen and Freek Van Walderveen. 2013. Near-optimal range reporting structures for categorical data. In *Proceedings of 24th ACM-SIAM symposium on discrete algorithms (SODA)*. SIAM, 265–276.

[28] Jiří Matoušek. 1996. On the distortion required for embedding finite metric spaces into normed spaces. *Israel Journal of Mathematics* 93, 1 (1996), 333–344.

[29] Marius Muja and David G Lowe. 2014. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence* 36, 11 (2014), 2227–2240.

[30] Anna Östlin and Rasmus Pagh. 2002. One-Probe Search. In *Proceedings of 29th international colloquium on automata, languages and programming (ICALP)*. 439–450.

[31] Rina Panigrahy, Kunal Talwar, and Udi Wieder. 2010. Lower bounds on near neighbor search via metric expansion. In *2010 IEEE 51st Symposium on Foundations of Computer Science*. IEEE, 805–814.

[32] Vladimir Pestov. 2013. Lower bounds on performance of metric tree indexing schemes for exact similarity search in high dimensions. *Algorithmica* 66, 2 (2013), 310–328.

[33] Vladimir Pestov and Aleksandar Stojmirović. 2006. Indexing schemes for similarity search: An illustrated paradigm. *Fundamenta Informaticae* 70, 4 (2006), 367–385.

[34] Aviad Rubinstein. 2018. Hardness of approximate nearest neighbor search. In *Proceedings of 50th ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25–29, 2018*, Ilias Diakonikolas, David Kempe, and Monika Henzinger (Eds.). ACM, 1260–1268. <https://doi.org/10.1145/3188745.3188916>

[35] Micha Streppel and Ke Yi. 2011. Approximate range searching in external memory. *Algorithmica* 59, 2 (2011), 115–128.

[36] Yufei Tao, Ke Yi, Cheng Sheng, and Panos Kalnis. 2010. Efficient and accurate nearest neighbor and closest pair search in high-dimensional space. *ACM Trans. Database Syst.* 35, 3 (2010), 20:1–20:46. <https://doi.org/10.1145/1806907.1806912>

[37] Zhewei Wei, Ke Yi, and Qin Zhang. 2009. Dynamic external hashing: The limit of buffering. In *Proceedings of 21st symposium on Parallelism in algorithms and architectures*. ACM, 253–259.

[38] Ke Yi. 2009. Dynamic indexability and lower bounds for dynamic one-dimensional range query indexes. In *Proceedings of 28th ACM Symposium on Principles of Database Systems (PODS)*. ACM, 187–196.