

Approximate Nearest Neighbors Search Without False Negatives For l_2 For $c > \sqrt{\log \log n}^*$

Piotr Sankowski¹ and Piotr Wygocki²

- 1 University of Warsaw, Poland
sank@mimuw.edu.pl
- 2 University of Warsaw, Poland
wygos@mimuw.edu.pl

Abstract

In this paper, we report progress on answering the open problem presented by Pagh [11], who considered the nearest neighbor search without false negatives for the Hamming distance. We show new data structures for solving the c -approximate nearest neighbors problem without false negatives for Euclidean high dimensional space \mathcal{R}^d . These data structures work for any $c = \omega(\sqrt{\log \log n})$, where n is the number of points in the input set, with poly-logarithmic query time and polynomial pre-processing time. This improves over the known algorithms, which require c to be $\Omega(\sqrt{d})$.

This improvement is obtained by applying a sequence of reductions, which are interesting on their own. First, we reduce the problem to d instances of dimension logarithmic in n . Next, these instances are reduced to a number of c -approximate nearest neighbor search without false negatives instances in $(\mathbb{R}^k)^L$ space equipped with metric $m(x, y) = \max_{1 \leq i \leq L} (\|x_i - y_i\|_2)$.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.3 Probability and Statistics

Keywords and phrases locality sensitive hashing, approximate nearest neighbor search, high-dimensional, similarity search

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2017.63

1 Introduction

The nearest neighbor search has numerous applications in image processing, search engines, recommendation engines, predictions and machine learning. We define the nearest neighbor problem as follows: for a given input set, a query point and a distance R , return a point (optionally all points) from the input set, which is closer to the query point than R in the given metric (typically l_p for $p \in [1, \infty]$), or report that such a point does not exist. The input set and the distance R are known in advance. Hence, the input set may be preprocessed what may result in reducing the query time. The problem in which the distance R is not known during the preprocessing and our task is to find the nearest neighbor can be efficiently reduced to the problem defined as above [7].¹ Unfortunately, the nearest neighbors search, defined as above, appears to be intractable for high dimensional spaces such as l_p^d for large d . The existence of an algorithm with a sub-linear in the data size and not exponential in

* This work was partially supported by grant NCN2014/13/B/ST6/00770 of Polish National Science Center and ERC StG grant TOTAL no. 677651.

¹ Authors used to distinguish between these two problems. The problem in which the radius is known in pre-processing is sometimes called Point Location in Equal Balls (PLEB) [7].



d query time and with not exponential in d pre-processing, would contradict the strong exponential time hypothesis [12]. In order to overcome this obstacle, the c -approximate nearest neighbors problem with $c > 1$, was introduced. In this problem, the query result is allowed to contain points which are within the distance cR from the query point. In other words, the points within the distance R from the query point are classified as neighbors, the points farther than cR are classified as non-neighbors, while the rest may be classified into any of these two categories. This assumption makes the problem easier, for many metric spaces such as l_p when $p \in [1, 2]$ or the Hamming space [7]. On one hand, sub linear in the input size queries are possible. On the other hand, the queries and pre-processing times are polynomial in the dimension of the space.

Locality sensitive hashing (LSH) is one of the major techniques for solving the c -approximate nearest neighbor search. Many LSH functions are mappings which roughly preserve distances. A random LSH function maps two 'close' points to two 'close' hashes with 'large' probability. Analogously, two 'distant' points are mapped to two 'distant' hashes with 'large' probability. Roughly speaking, the LSH is used to reduce the dimension of the input space, which allows to solve the problem in the lower dimensional space. Thus, the efficiency of the algorithm strongly depends on the quality of LSH functions used. The crucial properties of the LSH functions are the probability of false positives and the probability of false negatives. A false negative is a point which is 'close' to the query point, but its hash is 'far away' from the hash of the query point. Analogously, the false positive is a point whose distance to the query point is 'large', but it is mapped to a 'close' hash.

The previously known algorithms for the c -approximate nearest neighbors (see e.g., [2, 4]) give Monte Carlo guaranties for returned points, i.e., an input point close to the query point is returned with some probability. In other words, there might be false negatives. For example, a common choice of the hash functions is $f(x) = \langle x, v \rangle$ or $f(x) = \lfloor \langle x, v \rangle \rfloor$, where v is a vector of numbers drawn independently from some probability distribution [2, 7, 10]. For a Gaussian distribution, $\langle x, v \rangle$ is also Gaussian with zero mean and standard deviation equal to $\|x\|_2$. It is easy to see that these are LSH functions for l_2 , but as mentioned above, they only give probabilistic guaranties. In this paper, we aim to enhance this by focusing on the c -approximate nearest neighbor search without false negatives for l_2 . In other words, we consider algorithms, where a point 'close' to the query point is guaranteed to be returned.

Throughout this paper, we assume that $n \gg d$ and $\exp(d) \gg n$. This represents a situation where the exhaustive scan through all the input points, as well as the usage of data structures exponentially dependant on d , become intractable. The typical values to consider could be $n = 10^9$ and $d = 100$. If not explicitly specified, all statements assume the usage of the l_2 norm.

2 Related Work

There exists an efficient, Monte Carlo c -nearest neighbor algorithm for l_1 [7] with the query and the pre-processing complexity equal to $\mathcal{O}(n^{1/c})$ and $\mathcal{O}(n^{1+1/c})$, respectively. For l_2 in turn, there exists a near to optimal [9] algorithm [2] with the query and the pre-processing complexity equal to $n^{1/c^2+o(1)}$ and $n^{1+1/c^2+o(1)}$, respectively. Moreover, the algorithms presented in [7] work for l_p for any $p \in [1, 2]$. There are also data dependent algorithms, which take into account the actual distribution of the input set [3], which achieve query time $dn^{\rho+o(1)}$ and space $\mathcal{O}(n^{1+\rho+o(1)} + dn)$, where $\rho = 1/(2c^2 - 1)$.

Pagh [11] considered the c -approximate nearest neighbor search without false negatives (NN_{wfn}) for the Hamming space, obtaining the results close to those of [7]. He showed that the bounds of his algorithm for $cR = \log(n/k)$ differ by at most a factor of $\ln 4$ in

the exponent in comparison to the bounds of [7]. Recently, Ahle [1] showed an optimal [9] algorithm for the nearest neighbor without false negatives for Hamming space and Braun-Blanquet metric. Indyk [5] provided a deterministic algorithm for l_∞ for $c = \Theta(\log_{1+\rho} \log d)$ with the storage $\mathcal{O}(n^{1+\rho} \log^{O(1)} n)$ and the query time $\log^{O(1)} n$ for some tunable parameter ρ . He proved that the nearest neighbor without false negatives for l_∞ for $c < 3$ is as hard as the subset query problem, a long-standing combinatorial problem. This indicates that the nearest neighbor without false negatives for l_∞ might be hard to solve for any $c > 1$. Also, Indyk [6] considered deterministic mappings $l_2^n \rightarrow l_1^m$, for $m = n^{1+o(1)}$, which might be useful for constructing efficient algorithms for the nearest neighbor without false negatives [11].

Pacuk et al. [10] presented a schema for solving the nearest neighbor without false negatives for any $p \in [1, \infty]$ for $c = \Omega(d^{\max\{1/2, 1-1/p\}})$. Using the enhanced hash functions, Wygocki [13] presented algorithms with improved complexities. He considered two hashing families, giving different trade-offs between the execution times and the conditions on c . In particular, (Theorem 3, case 2, for $p = 2$ in [13]):

► **Theorem 1** ([13]). *For any $c > \bar{\tau} = 2\sqrt{d}$, there are data structures for the nearest neighbor without false negatives with*

- $\mathcal{O}(n^{1+\frac{\ln 3}{\ln(c/\bar{\tau})}})$ pre-processing time and $\mathcal{O}(d|P| + d \log n + d^2)$ query time for the 'fast query' algorithm,
- $\mathcal{O}(nd \log n)$ pre-processing time and $\mathcal{O}(d(|P| + n^{\frac{\ln 3}{\ln(3c/\bar{\tau})}}))$ query time for the 'fast pre-processing' algorithm,

where $|P|$ is the size of the result.

The dimension reduction with means of random linear mappings was considered previously in a more general context. In particular, Johnson-Lindenstrauss Lemma [8] is the most well known reference here. The concentration bounds used to prove this classic result will be very useful in our reductions:

► **Lemma 2** (Johnson-Lindenstrauss). *Let $Y \in \mathbb{R}^d$ be chosen uniformly from the surface of the d -dimensional sphere. Let $Z = (Y_1, Y_2, \dots, Y_k)$ be the projection onto the first k coordinates, where $k < d$. Then for any $\alpha < 1$:*

$$\mathbb{P} \left[\frac{d}{k} \|Z\|_2^2 \leq \alpha \right] \leq \exp\left(\frac{k}{2}(1 - \alpha + \log \alpha)\right), \quad (1)$$

3 Our contribution

Recently, efficient algorithms were proposed for solving the approximate nearest neighbor search without false negatives for $c = \Omega(\max\{\sqrt{d}, d^{1-1/p}\})$ in l_p for any $p \in [1, \infty]$ [10, 13]. The main problem with these algorithms is the constraint on c . For l_2 , the previous result require c to be of order of $\Omega(\sqrt{d})$, thus the nearest neighbor algorithm were allowed to return points within $\mathcal{O}(\sqrt{d}R)$ radius from query point. We relax this to any c , which makes the presented algorithms usable in practical cases. The contribution of this paper is relaxing this condition and improving the complexity of the algorithms for l_2 :

- We show that the NN_{wfn} can be reduced to d instances of NN_{wfn} in $\mathbb{R}^{\log n}$. For our typical settings of parameters, the factor of d is negligible. As a result, reducing the dimension leads to reducing the complexity of the problem. Moreover, it leads to relaxing the conditions on c to $c = \Omega(\sqrt{\log n})$.
- Further reductions lead to algorithms for any $c = \omega(\sqrt{\log \log n})$. We introduce an algorithm with the $n^{o(1)}$ query time and polynomial pre-processing time, which for large c tents to $n^{1+o(1)}$.

The first reduction is interesting on its own since further work on the problem can be done under the assumption that the dimension of the problem is logarithmic in n . This simplifies the problem at a cost of multiplying the complexities by a factor of d . Also, the authors of [10] proved that their construction is tight for $d = \omega(\log n)$, leaving the case of $d = \Theta(\log n)$ inconclusive.

3.1 Used Methods

In order to relax the conditions on c , we apply a sequence of dimension reductions. In Section 4, we show how to reduce the c -approximate nearest neighbors in l_2^d ($\text{NN}_{\text{wfn}}(c, d)$) to $d/\log(n)$ instances of $\text{NN}_{\text{wfn}}(\mathcal{O}(c), \mathcal{O}(\log n))$. Applying the algorithm of [13] as a black box gives the first improvement over [13]: an efficient algorithm for $c = \Omega(\sqrt{\log \log n})$. The reduction is based on the well-known Johnson-Lindenstrauss Lemma [8]. We introduce $d/\log(n)$ linear mappings, each reduces the dimension of the original problem. Each mapping roughly preserves the length of the vector and additionally at least one of them does not increase the length of the input vector. The property of not increasing the length of the vector is crucial. For two 'close' vectors $x, y \in \mathbb{R}^d$: $\|x - y\|_2 < 1$ and a linear mapping A , Ax and Ay are 'close' if and only if $\|Ax - Ay\|_2 = \|A(x - y)\|_2 < 1$, so A maps a 'small' vector $x - y$, to a 'small' vector $A(x - y)$.

In Section 5, we show further reductions, which enable us to relax the constraint to $c = \omega(\sqrt{\log \log n})$. We extend the reduction from Section 4 by using a number of mapping families. This leads to an interesting sub-problem of solving the approximate nearest neighbors in $(\mathbb{R}^k)^L$, for norm $\text{max-}l_2(x) := \max_{1 \leq i \leq L} \|x_i\|_2$ and the induced metric. This norm is present in literature and was denoted as max-product or l_∞ -product. Apparently, the c -approximate nearest neighbor search in $\text{max-}l_2$ might be solved using the LSH functions family introduced in [13].

This series of reductions leads to our final results. First we reduce the problem to a number of $\text{NN}_{\text{wfn}}(\mathcal{O}(c), \mathcal{O}(\log n))$ instances, each of which is further reduced to a number of problems in $\text{max-}l_2$, which in turn are solved using the LSH functions presented in [13].

The c -approximate nearest neighbors search without false negatives with parameter $c > 1$ and the dimension of the space equal to d , will be denoted as $\text{NN}_{\text{wfn}}(c, d)$. The expected query and pre-processing time complexities of $\text{NN}_{\text{wfn}}(c, d)$ will be denoted as $\text{query}(c, d)$ and $\text{preproc}(c, d)$ respectively. The input set will be denoted as X and it will always be assumed to contain n points. W.l.o.g, throughout this work we will assume, that R – a given radius equals 1 (otherwise, all vectors' lengths might be rescaled by $1/R$). The $\tilde{\mathcal{O}}()$ denotes the complexity up to the poly logarithmic factors i.e., $\tilde{\mathcal{O}}(f(n)) = \mathcal{O}(f(n)\text{poly}(\log(n)))$. $\|\cdot\|_2$ denotes the standard norm in l_2 , i.e., $\|x\|_2 = (\sum_i |x_i|^2)^{1/2}$. The $f(n) = \omega(g(n))$ means that f dominates g asymptotically, i.e., $g(n) = o(f(n))$.

4 Algorithm for $c = \Omega(\sqrt{\log \log n})$

The basic idea is the following: we will introduce a number of linear mappings to transform the d -dimensional problem to a number of problems with dimension reduced to $\mathcal{O}(\log n)$. Then we use the algorithm introduced in [13], to solve these problems in the space with the reduced dimension.

We will introduce d/k^2 linear mappings $A^{(1)}, A^{(2)}, \dots, A^{(d/k)} : \mathbb{R}^d \rightarrow \mathbb{R}^k$, where $k < d$ and show the following properties:

1. for each point $x \in \mathbb{R}^d$, such that $\|x\|_2 \leq 1$, there exists $1 \leq i \leq d/k$, such that $\|A^{(i)}x\|_2 \leq 1$,
2. for each point $x \in \mathbb{R}^d$, such that $\|x\|_2 \geq c$, where $c > 1$, the probability that there exists $1 \leq i \leq d/k$, such that $\|A^{(i)}x\|_2 \leq 1$ is bounded.

The property 1. states, that for a given 'short' vector (with a length smaller than 1), there is always at least one mapping, which transforms this vector to a vector of length smaller than 1. Moreover, we will show, that there exists at least one mapping $A^{(i)}$, which does not increase the length of the vector, i.e., such that $\|A^{(i)}x\|_2 \leq \|x\|_2$. The property 2. states, that we can bound the probability of a 'long' vector ($\|x\|_2 > c$), being mapped to a 'short' one ($\|A^{(i)}x\|_2 \leq 1$). Using the standard concentration measure arguments, we will prove that this probability decays exponentially in k .

4.1 Linear mappings

In this section, we will introduce linear mappings satisfying properties 1. and 2. Our technique will depend on the concentration bound used to prove the classic Johnson-Lindenstrauss Lemma. In Lemma 2, we take a random vector and project it to the first k vectors of the standard basis of \mathbb{R}^d . In our settings, we will project the given vector to a random orthonormal basis which gives the same guaranties. The mapping $A^{(i)}$ consists of k consecutive vectors from the random basis of the \mathbb{R}^d space scaled by $\sqrt{\frac{d}{k}}$. The following reduction describes the basic properties of our construction:

► **Lemma 3 (Reduction Lemma).** *For any parameter $\alpha \geq 1$ and $k < d$, there exist d/k linear mappings $A^{(1)}, A^{(2)}, \dots, A^{(d/k)}$, from \mathbb{R}^d to \mathbb{R}^k , such that:*

1. for each point $x \in \mathbb{R}^d$ such that $\|x\|_2 \leq 1$, there exists $1 \leq i \leq d/k$, such that $\|A^{(i)}x\|_2 \leq 1$,
2. for each point $x \in \mathbb{R}^d$ such that $\|x\|_2 \geq c$, where $c > 1$, for each $i: 1 \leq i \leq d/k$, we have

$$\mathbb{P} \left[\|A^{(i)}x\|_2 \leq \alpha \right] < e^{-k \left(\frac{c-\alpha}{2c} \right)^2}.$$

Proof. Let a_1, a_2, \dots, a_d be a random basis of \mathbb{R}^d . Each of the $A^{(i)}$ mappings is represented by a $k \times d$ dimensional matrix. We will use $A^{(i)}$ for denoting both the mapping and the corresponding matrix. The j th row of the matrix $A^{(i)}$ equals $A_j^{(i)} = \sqrt{\frac{d}{k}} a_{(i-1)k+j}$. In other words, the rows of $A^{(i)}$ consist of k consecutive vectors from the random basis of the \mathbb{R}^d space scaled by $\sqrt{\frac{d}{k}}$.

To prove the first property, observe that $A = \sum_{i=1}^d \langle a_i, x \rangle^2 \leq 1$, since the distance is independent of the basis. Assume on the contrary, that for each i , $\|A^{(i)}x\|_2 > 1$. It follows that $d \geq dA = k \sum_{i=1}^d \|A^{(i)}x\|_2^2 > d$. This contradiction ends the proof of the first property.

For any $x \in \mathbb{R}^d$, such that $\|x\|_2 > c$, the probability:

$$\mathbb{P} \left[\|A^{(i)}x\|_2 \leq \alpha \right] = \mathbb{P} \left[\frac{\|A^{(i)}x\|_2^2}{c^2} \leq \left(\frac{\alpha}{c} \right)^2 \right] \leq \mathbb{P} \left[\frac{\|A^{(i)}x\|_2^2}{\|x\|_2^2} \leq \left(\frac{\alpha}{c} \right)^2 \right].$$

² For simplicity, let us assume that k divides d , this can be achieved by padding extra dimensions with 0's.

Using the fact that $\log x < x - 1 - (x - 1)^2/2$ for $x < 1$ and Lemma 2, the above is bounded as follows:

$$\mathbb{P} \left[\frac{\|A^{(i)}x\|_2^2}{\|x\|_2^2} \leq \left(\frac{\alpha}{c}\right)^2 \right] \leq \exp \left(-\frac{k}{4} \left(1 - \left(\frac{\alpha}{c}\right)^2\right) \right) \leq e^{-k \left(\frac{c-\alpha}{2c}\right)^2},$$

which completes the proof. \blacktriangleleft

4.2 Algorithm

The algorithm works as follows: for each i , we project \mathbb{R}^d to \mathbb{R}^k using A_i and solve the corresponding problem in the smaller space. For each query point, we need to merge the solutions obtained for each subproblem. This results in reducing the $\text{NN}_{\text{wfn}}(c, d)$ to d/k instances of $\text{NN}_{\text{wfn}}(\alpha, k)$.

► **Lemma 4.** *For $1 < \alpha < c$ and $k < d$, the $\text{NN}_{\text{wfn}}(c, d)$ can be reduced to d/k instances of the $\text{NN}_{\text{wfn}}(\alpha, k)$. The expected pre-processing time equals $\mathcal{O}(d^2n + d/k \text{ preproc}(\alpha, k))$ and the expected query time equals $\mathcal{O}(d^2 + d/k e^{-k \left(\frac{c-\alpha}{2c}\right)^2} n + d/k \text{ query}(k, \alpha))$.*

Proof. We use the assumption that $k < d < n$ to simplify the complexities. The pre-processing time consists of:

- d^3 : the time of computing a random orthonormal basis of \mathbb{R}^d .
- d^2n : the time of changing the basis to a_1, a_2, \dots, a_d .
- dnk : the time of computing $A^{(i)}x$ for all $1 \leq i \leq d$ and for all n points.
- $d/k \text{ preproc}(\alpha, k)$: the expected pre-processing time of all subproblems.

The query time consists of:

- d^2 : the time of changing the basis to a_1, a_2, \dots, a_d .
- $d/k e^{-k \left(\frac{c-\alpha}{2c}\right)^2} n$: the expected number of false positives (by Lemma 3).
- $d/k \text{ query}(k, \alpha)$: the expected query time of all subproblems. \blacktriangleleft

The following corollary simplifies the formulas used in Lemma 4 and shows that if $\frac{c}{c-\alpha}$ is bounded, the $\text{NN}_{\text{wfn}}(c, d)$ can be reduced to a number of problems of dimension $\log n$ in an efficient way. Namely, setting $k = \left(\frac{2c}{c-\alpha}\right)^2 \log n$ we get:


► **Corollary 5.** *For any $1 \leq \alpha < c$ and $\gamma \log n < d$, the $\text{NN}_{\text{wfn}}(c, d)$ can be reduced to $d/\log n$ instances of the $\text{NN}_{\text{wfn}}(\alpha, \gamma \log n)$, where $\gamma = \left(\frac{2c}{c-\alpha}\right)^2$ and:*

$$\text{query}(c, d) = \mathcal{O}(d^2 + d/\log(n) \text{ query}(\alpha, \gamma \log n)),$$

$$\text{preproc}(c, d) = \mathcal{O}(d^2n + d/\log(n) \text{ preproc}(\alpha, \gamma \log n)).$$

Combining the above corollary with the results introduced in [13], we can achieve the algorithm with the polynomial pre-processing time and the sub-linear query time. Theorem 1 states, that for any $c > 2\sqrt{d}$, the $\text{NN}_{\text{wfn}}(c, d)$ can be solved in the $\mathcal{O}(n^{1 + \frac{\log 3}{\log(c/\tilde{\tau})}})$ pre-processing time and the query time equal to $\mathcal{O}(d|P| + d \log n + d^2)$, where P is the size of the result set and $\tilde{\tau} = 2\sqrt{d}$. Altogether, setting $\alpha = c/2$ in Corollary 5, we get³:

³ The author of [13] presented multiple algorithms giving different trade-offs between the pre-processing time and the query time. Particularly, the algorithm with the $\mathcal{O}(n \log n)$ processing time and the sub-linear query time was presented. The same can be done for Theorem 6. We omit this to avoid the unnecessary complexity.



■ **Figure 1** Each column describes one family of linear mappings, constructed based on one random, orthonormal basis. The blue path describes one combination of mappings.

► **Theorem 6.** *The $NN_{\text{wfn}}(c, d)$ can be solved for any $c > \tilde{\kappa} = 16\sqrt{\log n}$ with:*

$$\text{query}(c, d) = \mathcal{O}(d|P| + d \log n + d^2)$$

$$\text{preproc}(c, d) = \mathcal{O}(dn^{1 + \frac{\ln 3}{\log(c/\tilde{\kappa})}} / \log(n)).$$

The time complexity of the algorithm is the same as for $c = \Omega(\sqrt{d})$, the pre-processing time is larger by a factor of $d/\log(n)$.

5 The algorithm for $c = \omega(\sqrt{\log(\log(n))})$

In this section we give another algorithm which works for $c = \omega(\sqrt{\log(\log(n))})$. Lemma 2 implies that the $NN_{\text{wfn}}(c, d)$ problem can be reduced to $d/\log(n)$ problems of dimension logarithmic in n . In order to reduce the dimension even more, we will employ L independent families of linear mappings introduced in Section 4. In each of the families, there is at least one mapping, which does not increase the length of the input vector. As a result, there exists a combination of L mappings (each mapping taken from a distinct family) which do not increase the input vector length. Also, for any combination of L mappings, the probability that all the mappings transform a 'long' vectors to a 'short'; one can be bounded. The structure of the mappings is presented in Figure 1.

To formalize the above line of thinking, we introduce the following lemma:

► **Lemma 7.** *For any natural number $L > 0$, there exist d/k L linear mappings $A^{(i,j)} : \mathbb{R}^d \rightarrow \mathbb{R}^k$, where $k < d$, $1 \leq i \leq d/k$ and $1 \leq j \leq L$, such that*

1. *for each point $x \in \mathbb{R}^d$ which satisfies $\|x\|_2 \leq 1$, there exist $1 \leq i_1, i_2, \dots, i_L \leq d$ such that $\|A^{(i_j, j)}x\|_2 \leq 1$, for each $1 \leq j \leq L$.*
2. *for each point $x \in \mathbb{R}^d$ which satisfies $\|x\|_2 \geq c$, where $c > 1$, for each $i_1, i_2, \dots, i_L : 1 \leq i_1, i_2, \dots, i_L \leq d/k$, we have*

$$\mathbb{P} \left[\forall_j : \|A^{(i_j, j)}x\|_2 \leq \alpha \right] < \exp \left(-\frac{kL}{4} \left(\frac{c - \alpha}{c} \right)^2 \right).$$

Proof. For each $j: 1 \leq j \leq L$ we independently sample the orthonormal basis of $\mathbb{R}^d: a_1, a_2, \dots, a_d$. The $A^{(i,j)}$ will be created in the same way as in Lemma 3, namely, the t -th row of $A^{(i,j)}$ equals $A_t^{(i,j)} = \sqrt{\frac{d}{k}} a_{(i-1)k+t}$. The properties (1) and (2) follow directly from Lemma 3. ◀

In order to employ Lemma 7 for a given query q , we need to be able to find all points in X such that a given combination of mappings transforms these points and the query point to 'close' vectors. In other words, we need to find all c -approximate nearest neighbors for the transformed input set $\tilde{X} \subset (\mathbb{R}^k)^L$ in the space equipped with metric: $\max_{1 \leq i \leq L} (\|x_i - y_i\|)$, which is formally defined as follows:

► **Definition 8** (the c -approximate nearest neighbor search in $max\text{-}l_2$). The $max\text{-}l_2\text{-}NN(c, L, k)$ is defined as follows: given a query point $q \in (\mathbb{R}^k)^L$ and a set $\tilde{X} \subset (\mathbb{R}^k)^L$ of n input points, find all input points, such that for each $1 \leq i \leq L$: $\|q_i - \tilde{x}_i\|_2 \leq 1$. Moreover, each \tilde{x} satisfying $\forall_i \|q_i - \tilde{x}_i\|_2 \leq c$, might be returned as well. Finally, each x such that $\exists_i \|q_i - \tilde{x}_i\| > c$, must not be returned.

Using the construction from Lemma 7, the $NN_{\text{wfn}}(c, d)$ problem can be reduced to d^L instances of the $max\text{-}l_2\text{-}NN(\alpha, L, k)$. Each of the instances is represented by indices: $\{i_1, i_2, \dots, i_L\}$ and the corresponding mappings $A^{(i_j, j)}$ for $1 \leq j \leq L$. Each input point $\tilde{x} \in \tilde{X}$ comes from the point $x \in X$ by applying the mappings: $\tilde{x} = (A^{(i_1, 1)}x, \dots, A^{(i_L, L)}x)$. Similarly, the query point q , in $max\text{-}l_2$, is created from the query point q in $NN_{\text{wfn}}(c, d)$, as $(A^{(i_1, 1)}q, \dots, A^{(i_L, L)}q)$.

► **Lemma 9.** The $NN_{\text{wfn}}(c, d)$ can be reduced to $(d/k)^L$ instances of $max\text{-}l_2\text{-}NN(\alpha, L, k)$. The expected pre-processing time equals:

$$\mathcal{O}(Ld^2n + (d/k)^L \text{preproc}_{max\text{-}l_2}(\alpha, L, k))$$

and the expected query time equals:

$$\mathcal{O}(Ld^2 + (d/k)^L e^{-kL(\frac{c-\alpha}{2c})^2} n + (d/k)^L \text{query}_{max\text{-}l_2}(\alpha, L, k)).$$

The proof of the Lemma is analogical to the proof of Lemma 4. The following corollary presents the simplified version of Lemma 9. Setting $k = \lceil L^{-1}(\frac{2c}{c-\alpha})^2 \log n \rceil$ we get:

► **Corollary 10.** For any $1 \leq \alpha < c$, the $NN_{\text{wfn}}(c, d)$ can be reduced to d^L instances of $max\text{-}l_2\text{-}NN(\alpha, L, L^{-1}\gamma \log n)$, where $\gamma = (\frac{2c}{c-\alpha})^2$ and:

$$\text{query}(c, d) = \mathcal{O}(Ld^2 + (d/\log(n))^L \text{query}_{max\text{-}l_2}(\alpha, L, \lceil L^{-1}\gamma \log n \rceil)),$$

$$\text{preproc}(c, d) = \mathcal{O}(Ld^2n + (d/\log(n))^L \text{preproc}_{max\text{-}l_2}(\alpha, L, \lceil L^{-1}\gamma \log n \rceil)).$$

The $max\text{-}l_2\text{-}NN(\alpha, L, k)$ can be trivially solved by dealing with each of the L -dimensional $NN_{\text{wfn}}(\alpha, k)$ problems separately. Unfortunately, this gives unacceptable complexities. In order to improve complexity of algorithm for the $NN_{\text{wfn}}(c, d)$ problem, we need to be able to solve the $max\text{-}l_2\text{-}NN$ more efficiently.

5.1 Solving the c -approximate nearest neighbors in $max\text{-}l_2$

In order to solve this problem, we use the standard LSH technique based on the hash functions \tilde{h} introduced in [13] defined as follows:

$$\tilde{h}(x) = \lfloor \langle w, x \rangle \rfloor, \text{ where } w \text{ is a random vector from the unit sphere } \mathbb{S}^{(d-1)}.$$

We consider two hashes to be 'close' if $|\tilde{h}(x) - \tilde{h}(x')| \leq 1$. Based on \tilde{h} , we introduce a new hash function g . Each of the input points is hashed by g and the reference to this point is kept in a single hash map. For a given query point, we examine all input points which are hashed to the same value as the query point.

Namely, each $\tilde{x} \in (\mathbb{R}^k)^L$ will be hashed by $g(\tilde{x}) := (g_1(\tilde{x}_1), \dots, g_L(\tilde{x}_L))$, where $g_i(x) := (\tilde{h}_1(x), \tilde{h}_2(x), \dots, \tilde{h}_w(x))$ is a hash function defined as a concatenation of w random LSH functions \tilde{h} . The function g can be also seen as a concatenation of wL random hash functions \tilde{h} . If two points are 'close' in the considered $max\text{-}l_2$ metric, then g transforms these points to hashes $p^{(1)}, p^{(2)} \in \mathbb{Z}^{wL}$, such that $|p_i^{(1)} - p_i^{(2)}| \leq 1$ for all $i \in wL$. The pre-processing

algorithm is summarized in the following pseudocode:

Algorithm 1: The pre-processing algorithm

Data: $X \subset (\mathbb{R}^k)^L$ - the set of n input points
Result: $H : \mathbb{Z}^{wL} \rightarrow 2^X$ - the hash map storing for each hash $\alpha \in \mathbb{Z}^{wL}$ the subset of input points with hashes close to α

```

 $H = \emptyset;$ 
for  $x \in X$  do
   $\alpha = g(x);$ 
  for  $\alpha'$  such that  $\|\alpha - \alpha'\|_\infty \leq 1$  do
     $H(\alpha').push(x);$ 
  end
end

```

The query algorithm consists of examining the bucket for $g(\text{query_point})$:

Algorithm 2: The query algorithm

Data: $q \in (\mathbb{R}^k)^L$ - the query point
Result: $P \subset X$ - the set of neighbors of q

```

 $P = \emptyset;$ 
for  $x \in H(g(q))$  do
  if  $x$  is a neighbor of  $q$  then
     $P.push(x);$ 
  end
end

```

The following theorem describes the above algorithm:

► **Theorem 11.** For $L = o(\log n)$ and $c > 2\sqrt{k}$, the $\text{max-}l_2\text{-NN}(c, L, k)$ can be solved in the $\mathcal{O}(kL|P| + k \ln n + Lk^2)$ query time and $\mathcal{O}(n^{1 + \frac{\ln 3}{\ln(c/\tilde{k})}})$ pre-processing time complexity for $\tilde{k} = 2\sqrt{k}$.⁴

Proof. Let us start with the key properties of the LSH family.

► **Observation 12** ('Close' points have 'close' hashes for \tilde{h} (Observation 5 in [13])). For $x, y \in \mathbb{R}^d$, if $\|x - y\| < 1$ then $\forall_{\tilde{h}} |\tilde{h}(x) - \tilde{h}(y)| \leq 1$.

► **Lemma 13** (The probability of false positives for \tilde{h} (Lemma 2 in [13])). For $x, y \in \mathbb{R}^d$ and $c > \tilde{\tau} = 2d^{1/2}$ such that $\|x - y\| > c$, it holds:

$$\tilde{p}_{fp} = \mathbb{P} [|\tilde{h}(x) - \tilde{h}(y)| \leq 1] < \tilde{\tau}/c.$$

Since we consider two hashes to be 'close', when they differ at most by one (see Observation 12), for each hash $\alpha \in \mathbb{Z}^{wL}$ we need to store the reference to every point, that satisfies $\|\alpha - g(x)\|_\infty \leq 1$. Thus, the hash map size is $\mathcal{O}(n3^{wL})$. Computing a single \tilde{h} function in \mathbb{R}^k takes $\mathcal{O}(k)$, so evaluating the $g(x)$ for $x \in (\mathbb{R}^k)^L$ takes $\mathcal{O}(wkL)$. The pre-processing consists of computing the 3^{wL} hashes for each point in the input set. The query consists of computing the hash of the query point, looking up all the points with colliding hashes, filtering out the false positives and returning the neighbors.

⁴ Theorem 11 might be generalized, to any $p \in [1, \infty]$. The generalization is done by applying hash functions suited for l_p . Such hash functions were introduced in [13].

It is easy to derive the following complexities:

For any $c > 2k^{1/2}$ and the number of iterations $w \geq 1$, there exists a $max\text{-}l_2\text{-}NN(c, L, k)$ algorithm with the following properties:

- the pre-processing time: $\mathcal{O}(n(wkL + 3^{wL}))$, where wkL is the time needed to compute the $g(\text{input_point})$ and the $\mathcal{O}(3^{wL})$ is the number of the updated hashes for one input point,
- the expected query time: $\mathcal{O}(kL(|P| + w + n\tilde{p}_{fp}^{wL}))$, where wkL is the time needed to compute the $g(\text{query_point})$, $n\tilde{p}_{fp}^{wL}$ is the number of false positives which need to be ignored, $|P|$ denotes the size of the result set. For each of the candidates, we need to perform a check of complexity $\mathcal{O}(kL)$ to classify the point as a true positive or a false positive.

Above $\tilde{p}_{fp} = \tilde{\tau}/c$ (see Lemma 13).

The number of iterations w can be chosen arbitrarily, so we will choose the optimal value. Denote $a = -\ln \tilde{p}_{fp}$ and $b = \ln 3$, then set w to be:

$$w = \left\lceil \frac{\ln \frac{na}{k}}{a} L^{-1} \right\rceil.$$

Let us assume that n is large enough so that $w \geq 1$. Then, using the fact that $x^{1/x}$ is bounded for $x > 0$ we have:

$$3^{wL} \leq 3 \cdot (3^{\ln \frac{na}{k}})^{1/a} = 3 \cdot \left(\frac{na}{k}\right)^{b/a} = 3 \cdot \left(\frac{n}{k}\right)^{b/a},$$

$$n\tilde{p}_{fp}^{wL} = ne^{-awL} \leq ne^{-a \frac{\ln \frac{na}{k}}{a}} = \frac{k}{a}.$$

Hence, for constant c the expected query time is $\mathcal{O}(kL|P| + k \ln n + Lk^2)$. Subsequently, the pre-processing time is: $\mathcal{O}(n3^{wL}) = \mathcal{O}(n^{1+b/a})$. Substituting a , b and p_{fp} values gives the needed complexity guaranties. ◀

5.2 Putting it All Together

In order to achieve an efficient algorithm for $c = \omega(\log(\log(n)))$, we will make a series of reductions. First, using Corollary 5, we reduce our problem to a number of $NN_{\text{wfn}}(\mathcal{O}(c), \mathcal{O}(\log n))$ problems. Next, these problems are reduced to a number of $max\text{-}l_2\text{-}NN$ problems with dimension k of $\mathcal{O}(\log \log n)$. In the end, we use Theorem 11 to solve the $max\text{-}l_2\text{-}NN$.

► **Theorem 14.** *The $NN_{\text{wfn}}(c, d)$ can be solved with:*

- the pre-processing time $\tilde{\mathcal{O}}(d^2n + dn^{1+\frac{\ln 3}{\ln(c/\mu)}+1/f(n)})$,
- the query time $\tilde{\mathcal{O}}(d^2 + dn^{1/f(n)}|P|)$,

for any $c > \mu = D\sqrt{f(n) \log \log n}$, where $f(n)$ is any function, which satisfies $1/f(n) = o(1)$ and D is some constant.

Proof. There are two consecutive reductions:

1. By Corollary 5, the $NN_{\text{wfn}}(c, d)$ can be reduced to d instances of the $NN_{\text{wfn}}(\alpha_1, k_1)$.
2. By Corollary 10, the $NN_{\text{wfn}}(\alpha_1, k_1)$ can be reduced to k_1^L instances of the $max\text{-}l_2\text{-}NN(\alpha_2, k_2, L)$.

Accordingly, we set:

1. $\alpha_1 = c/2$ and $k_1 = \lceil D_1 \log n \rceil$ in the first reduction
2. $\alpha_2 = c/4$, $k_2 = \lceil D_2 L^{-1} \log n \rceil \leq \lceil D_2 f(n) \log \log n \rceil$ and $L = \lceil \frac{\log n}{f(n) \log \log n} \rceil$ in the second reduction.

The constants D_1 and D_2 are chosen to satisfy Corollaries 5 and 10. k_1^L can be bounded in the following way:

$$k_1^L = \lceil D_1 \log n \rceil^L = \tilde{O}(n^{1/f(n)}) = \tilde{O}(n^{o(1)}).$$

Substituting the complexities for subproblems gives the final complexities. ◀

The function $f(n)$ may be chosen arbitrarily. Slowly increasing $f(n)$ will be chosen for small c close to $\Theta(\log \log n)$. For larger c , one should choose the maximal possible $f(n)$, to optimize the query time complexity.

6 Conclusion and Future Work

We have presented the c -approximate nearest neighbor algorithm without false negatives in l_2 for any $c = \omega(\sqrt{\log \log n})$. Such an algorithm might work very well for high entropy datasets, where the distances tend to be relatively large (see [11] for more details). Also, we showed that the c -approximate nearest neighbor search in l_2^d may be reduced to d instances of the problem in $l_2^{\log n}$. Hence, further research might focus on the instances with dimension logarithmic in n .

Another open problems are to reduce the time complexity of the algorithm and relax the restrictions on the approximation factor c or proving that these restrictions are essential. We wish to match the time complexities given in [7] or show that the achieved bounds are optimal.

Acknowledgments. We would like to thank Andrzej Pacuk and Kamila Wygocka for meaningful discussion. Also, I would like to thank the anonymous reviewers for many comments which greatly increased the quality of this work.

References

- 1 Thomas Dybdahl Ahle. Optimal las vegas locality sensitive data structures. *CoRR*, abs/1704.02054, 2017. URL: <http://arxiv.org/abs/1704.02054>.
- 2 Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM*, 51(1):117–122, 2008. doi:10.1145/1327452.1327494.
- 3 Alexandr Andoni and Ilya Razenshteyn. Optimal data-dependent hashing for approximate near neighbors. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 793–801. ACM, 2015. doi:10.1145/2746539.2746553.
- 4 Bernard Chazelle, Ding Liu, and Avner Magen. Approximate range searching in higher dimension. *Computational Geometry*, 39(1):24–29, 2008. doi:10.1016/j.comgeo.2007.05.008.
- 5 Piotr Indyk. On approximate nearest neighbors in non-euclidean spaces. In *39th Annual Symposium on Foundations of Computer Science, FOCS'98, November 8-11, 1998, Palo Alto, California, USA*, pages 148–155, 1998. doi:10.1109/SFCS.1998.743438.
- 6 Piotr Indyk. Uncertainty principles, extractors, and explicit embeddings of l_2 into l_1 . In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing, STOC'07*, pages 615–620, New York, NY, USA, 2007. ACM. doi:10.1145/1250790.1250881.
- 7 Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on*

- Theory of Computing*, STOC'98, pages 604–613, New York, NY, USA, 1998. ACM. doi:10.1145/276698.276876.
- 8 William Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, volume 26 of *Contemporary Mathematics*, pages 189–206. American Mathematical Society, 1984.
 - 9 Ryan O'Donnell, Yi Wu, and Yuan Zhou. Optimal lower bounds for locality-sensitive hashing (except when q is tiny). *ACM Trans. Comput. Theory*, 6(1):5:1–5:13, March 2014. doi:10.1145/2578221.
 - 10 Andrzej Pacuk, Piotr Sankowski, Karol Wegrzycki, and Piotr Wygocki. Locality-sensitive hashing without false negatives for l_p . In *Computing and Combinatorics - 22nd International Conference, COCOON 2016, Ho Chi Minh City, Vietnam, August 2-4, 2016, Proceedings*, pages 105–118, 2016. doi:10.1007/978-3-319-42634-1_9.
 - 11 Rasmus Pagh. Locality-sensitive hashing without false negatives. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'16, pages 1–9, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=2884435.2884436>.
 - 12 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2):357–365, December 2005. doi:10.1016/j.tcs.2005.09.023.
 - 13 Piotr Wygocki. On fast bounded locality sensitive hashing. *ArXiv e-prints*, 2017. URL: <http://arxiv.org/abs/1704.05902>.