

APPROXIMATE NEAREST NEIGHBOR SEARCH IN HIGH DIMENSIONS

ALEXANDR ANDONI, PIOTR INDYK AND ILYA RAZENSHTEYN

Abstract

The nearest neighbor problem is defined as follows: Given a set P of n points in some metric space (X, D) , build a data structure that, given any point q , returns a point in P that is closest to q (its “nearest neighbor” in P). The data structure stores additional information about the set P , which is then used to find the nearest neighbor without computing all distances between q and P . The problem has a wide range of applications in machine learning, computer vision, databases and other fields.

To reduce the time needed to find nearest neighbors and the amount of memory used by the data structure, one can formulate the *approximate* nearest neighbor problem, where the goal is to return any point $p' \in P$ such that the distance from q to p' is at most $c \cdot \min_{p \in P} D(q, p)$, for some $c \geq 1$. Over the last two decades many efficient solutions to this problem were developed. In this article we survey these developments, as well as their connections to questions in geometric functional analysis and combinatorial geometry.

1 Introduction

The *nearest neighbor* problem is defined as follows: Given a set P of n points in a metric space defined over a set X with distance function D , build a data structure¹ that, given any “query” point $q \in X$, returns its “nearest neighbor” $\arg \min_{p \in P} D(q, p)$. A particularly interesting and well-studied case is that of nearest neighbor in geometric spaces, where $X = \mathbb{R}^d$ and the metric D is induced by some norm. The problem has a wide range of applications in machine learning, computer vision, databases and other fields, see [Shakhnarovich, Darrell, and Indyk \[2006\]](#) and [Andoni and Indyk \[2008\]](#) for an overview.

This research was supported by NSF and Simons Foundation.

MSC2010: primary 68W20; secondary 52A21, 46B09, 46B85, 68P05.

¹See [Section 1.1](#) for a discussion about the computational model.

A simple solution to this problem would store the set P in memory, and then, given q , compute all distances $\mathbb{D}(q, p)$ for $p \in P$ and select the point p with the minimum distance. Its disadvantage is the computational cost: computing all n distances requires at least n operations. Since in many applications n can be as large as 10^9 (see e.g., [Sundaram, Turmukhametova, Satish, Mostak, Indyk, Madden, and Dubey \[2013\]](#)), it was necessary to develop faster methods that find the nearest neighbors without explicitly computing all distances from q . Those methods compute and store additional information about the set P , which is then used to find nearest neighbors more efficiently. To illustrate this idea, consider another simple solution for the case where $X = \{0, 1\}^d$. In this case, one could precompute and store in memory the answers to all 2^d queries $q \in X$; given q , one could then return its nearest neighbor by performing only a single memory lookup. Unfortunately, this approach requires memory of size 2^d , which again is inefficient (d is at least 10^3 or higher in many applications).

The two aforementioned solutions can be viewed as extreme points in a tradeoff between the time to answer a query (“query time”) and the amount of memory used (“space”).² The study of this tradeoff dates back to the work of [Minsky and Papert \[1969, p. 222\]](#), and has become one of the key topics in the field of *computational geometry* [Preparata and Shamos \[1985\]](#). During the 1970s and 1980s many efficient solutions have been discovered for the case when $(X, \mathbb{D}) = (\mathbb{R}^d, \ell_2)$ and d is a constant independent of n . For example, for $d = 2$, one can construct a data structure using $O(n)$ space with $O(\log n)$ query time [Lipton and Tarjan \[1980\]](#). Unfortunately, as the dimension d increases, those data structures become less and less efficient. Specifically, it is known how construct data structures with $O(d^{O(1)} \log n)$ query time, but using $n^{O(d)}$ space ([Meiser \[1993\]](#), building on [Clarkson \[1988\]](#)).³ Furthermore, there is evidence that data structures with query times of the form $n^{1-\alpha} d^{O(1)}$ for some constant $\alpha > 0$ might be difficult to construct efficiently.⁴

The search for efficient solutions to the nearest neighbor problem has led to the question whether better space/query time bounds could be obtained if the data structure was allowed to report *approximate* answers. In the *c-approximate nearest neighbor* problem, the data structure can report any point $p' \in P$ within distance $c \cdot \min_{p \in P} \mathbb{D}(q, p)$ from q ; the parameter $c \geq 1$ is called “approximation factor”. The work of [Arya and Mount](#)

²There are other important data structure parameters, such as the time needed to construct it. For the sake of simplicity, we will mostly focus on query time and space.

³This exponential dependence on the dimension is due to the fact that those data structures compute and store the *Voronoi decomposition* of P , i.e., the decomposition of \mathbb{R}^d into cells such that all points in each cell have the same nearest neighbor in P . The combinatorial complexity of this decomposition could be as large as $n^{\Omega(d)}$ [Carathéodory \[1911\]](#).

⁴If such a data structure could be constructed in polynomial time $n^{O(1)}$, then the Strong Exponential Time Hypothesis [Vassilevska Williams \[2018\]](#) would be false. This fact essentially follows from [R. Williams \[2005\]](#), see the discussion after Theorem 1 in [Ahle, Pagh, Razenshteyn, and Silvestri \[2016\]](#).

[1993] and Bern [1993] showed that allowing $c > 1$ indeed leads to better data structures, although their solutions still retained exponential dependencies on d in the query time or space bounds Arya and Mount [1993] or required that the approximation factor c be polynomial in the dimension d Bern [1993]. These bounds have been substantially improved over the next few years, see e.g., Clarkson [1994], Chan [1998], Arya, Mount, Netanyahu, Silverman, and A. Y. Wu [1998], and Kleinberg [1997] and the references therein.

In this article we survey the “second wave” of approximate nearest neighbor data structures, whose query time and space bounds are polynomial in the dimension d .⁵ At a high level, these data structures are obtained in two steps. In the first step, the approximate *nearest* neighbor problem is reduced to its “decision version”, termed approximate *near* neighbor (see e.g. Har-Peled, Indyk, and Motwani [2012]). The second step involves constructing a data structure for the latter problem. In this survey we focus mostly on the second step.

The approximate near neighbor problem is parameterized by an approximation factor $c > 1$ as well as a “scale parameter” $r > 0$, and defined as follows.

Definition 1.1 ((c, r) -Approximate Near Neighbor). *Given a set P of n points in a metric space (X, D) , build a data structure \mathcal{S} that, given any query point $q \in X$ such that the metric ball $B_D(q, r) = \{p \in X : D(p, q) \leq r\}$ contains a point in P , \mathcal{S} returns any point in $B_D(q, cr) \cap P$.*

Note that the definition does not specify the behavior of the data structure if the ball $B_D(q, r)$ does not contain any point in P . We omit the index D when it is clear from the context.

The above definition applies to algorithms that are *deterministic*, i.e., do not use random bits. However, most of the approximate near neighbor algorithms in the literature are *randomized*, i.e., generate and use random bits while constructing the data structure. In this case, the data structure \mathcal{S} is a random variable, selected uniformly at random from some distribution. This leads to the following generalization.

Definition 1.2 ((c, r, δ) -Approximate Near Neighbor). *Given a set P of n points in a metric space (X, D) , build a data structure \mathcal{S} that, given any query point $q \in X$ such that $B(q, r) \cap P \neq \emptyset$,*

$$\Pr_{\mathcal{S}}[\mathcal{S} \text{ returns any point in } B(q, cr) \cap P] \geq 1 - \delta$$

The probability of failure δ of the data structure can be reduced by independently repeating the process several times, i.e., creating several data structures. Therefore, in the

⁵Due to the lack of space, we will not cover several important related topics, such as data structures for point-sets with low intrinsic dimension Clarkson [2006], approximate furthest neighbor, approximate nearest line search Mahabadi [2014] and other variants of the problem.

rest of the survey we will set δ to an arbitrary constant, say, $1/3$. We will use (c, r) -ANN to denote $(c, r, 1/3)$ -Approximate Near Neighbor.

1.1 Computational model. For the purpose of this survey, a data structure of size M is an array $A[1 \dots M]$ of numbers (“the memory”), together with an associated algorithm that, given a point q , returns a point in P as specified by the problem. The entries $A[i]$ of A are called “memory cells”. Due to lack of space, we will not formally define other details of the computational model, in particular what an algorithm is, how to measure its running time, what is the range of the array elements $A[i]$, etc. There are several ways of defining these notions, and the material in this survey is relatively robust to the variations in the definitions. We note, however, that one way to formalize these notions is to restrict all numbers, including point coordinates, memory entries, etc, to rational numbers of the form a/b , where $a \in \{-n^{O(1)} \dots n^{O(1)}\}$ and $b = n^{O(1)}$, and to define query time as the maximum number of memory cells accessed to answer any query q .

For an overview of these topics and formal definitions, the reader is referred to [Miltersen \[1999\]](#). For a discussion specifically geared towards mathematical audience, see [Fefferman and Klartag \[2009\]](#).

2 Data-independent approach

The first approach to the approximate near neighbor problem has been via data-independent data structures. These are data structures where the memory cells accessed by the query algorithm do not depend on the data set P , but only on q and (for randomized data structures) the random bits used to construct the data structure. In this section, we describe two methods for constructing such data structures, based on *oblivious dimension-reduction*, and on *randomized space partitions*. These methods give ANN data structures for the ℓ_1 and ℓ_2 spaces in particular.

2.1 ANN via dimension reduction. As described in the introduction, there exist ANN data structures with space and query time at most exponential in the dimension d . Since exponential space/time bounds are unaffordable for large d , a natural approach is to perform a *dimension reduction* beforehand, and then solve the problem in the lower, reduced dimension. The main ingredient of such an approach is a map $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ that preserves distances up to a $c = 1 + \varepsilon$ factor, where $k = O(\log n)$. Then a space bound exponential in k becomes polynomial in n .

Such dimension-reducing maps f indeed exist for the ℓ_2 norm if we allow randomization, as first shown in the influential paper by Johnson and Lindenstrauss:

Lemma 2.1 (Johnson and Lindenstrauss [1984]). Fix dimension $d \geq 1$ and a “target” dimension $k < d$. Let A be the projection of \mathbb{R}^d to its k -dimensional subspace selected uniformly at random (with respect to the Haar measure), and define $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ as $f(x) = \frac{\sqrt{d}}{\sqrt{k}}Ax$. Then, there is a universal constant $C > 0$, such that for any $\varepsilon \in (0, 1/2)$, and any $x, y \in \mathbb{R}^d$, we have that

$$\Pr_A \left[\frac{\|f(x) - f(y)\|}{\|x - y\|} \in (1 - \varepsilon, 1 + \varepsilon) \right] \geq 1 - e^{-C\varepsilon^2k}.$$

We can now apply this lemma, with $k = O\left(\frac{\log n}{\varepsilon^2}\right)$, to a set of points P to show that the map f has a $(1 + \varepsilon)$ distortion on P , with probability at least $2/3$. Most importantly the map f is “oblivious”, i.e., it does not depend on P .

We now show how to use Lemma 2.1 to design a $(1 + O(\varepsilon), r)$ -ANN data structure with the following guarantees.

Theorem 2.2 (Indyk and Motwani [1998] and Har-Peled, Indyk, and Motwani [2012]). Fix $\varepsilon \in (0, 1/2)$ and dimension $d \geq 1$. There is a $(1 + O(\varepsilon), r)$ -ANN data structure over (\mathbb{R}^d, ℓ_2) achieving $Q = O(d \cdot \frac{\log n}{\varepsilon^2})$ query time, and $S = n^{O(\log(1/\varepsilon)/\varepsilon^2)} + O(d(n + k))$ space. The time needed to build the data structure is $O(S + ndk)$.

Proof sketch. First, assume there is a $(1 + \varepsilon, r)$ -ANN data structure \mathcal{Q} for the k -dimensional ℓ_2 space, achieving query time $Q(n, k)$ and space bounded by $S(n, k)$. For $k = O\left(\frac{\log n}{\varepsilon^2}\right)$, we consider the map f from Lemma 2.1. For the dataset P , we compute $f(P)$ and preprocess this set using \mathcal{Q} (with the scale parameter $r(1 + \varepsilon)$). Then, for a query point $q \in \mathbb{R}^d$, we query the data structure \mathcal{Q} on $f(q)$. This algorithm works for a fixed dataset P and query q with $5/6$ probability, by applying Lemma 2.1 to the points in the set $P \cup \{q\}$. The map f preserves all distances between P and q up to a factor of $1 + \varepsilon$.

We now construct \mathcal{Q} with space $S(n, k) = n \cdot (1/\varepsilon)^{O(k)}$ and $Q(n, k) = O(k)$, which yields the stated bound for $k = O\left(\frac{\log n}{\varepsilon^2}\right)$. Given the scale parameter r , we discretize the space \mathbb{R}^k into cubes of sidelength $\varepsilon r / \sqrt{k}$, and consider the set S of cubes that intersect any ball $B(p', r)$ where $p' \in f(P)$. Using standard estimates on the volume of ℓ_2 balls, one can prove that $|S| \leq n \cdot (1/\varepsilon)^{O(k)}$. The data structure then stores the set S in a dictionary data structure.⁶ For a query $f(q)$, we just compute the cube that contains $f(q)$, and check whether it is contained in set S using the dictionary data structure. We note that there is an additional $1 + \varepsilon$ factor loss from discretization since the diameter of a cube is εr . \square

⁶In the dictionary problem, we are given a set S of elements from a discrete universe U , and we need to answer queries of the form “given x , is $x \in S$?”. This is a classic data structure problem and has many solutions. One concrete solution is via hashing Cormen, Leiserson, Rivest, and Stein [2001], which achieves space of $O(|S|)$ words, each of $O(\log |U|)$ bits, and query time of $O(1)$ in expectation.

A similar approach was introduced [Kushilevitz, Ostrovsky, and Rabani \[2000\]](#) in the context of the Hamming space $\{0, 1\}^d$. An important difference is that there is no analog of [Lemma 2.1](#) for the Hamming space [Brinkman and M. Charikar \[2005\]](#).⁷ Therefore, [Kushilevitz, Ostrovsky, and Rabani \[2000\]](#) introduce a weaker notion of randomized dimension reduction, which works only for a *fixed scale* r .

Lemma 2.3 ([Kushilevitz, Ostrovsky, and Rabani \[ibid.\]](#)). *Fix the error parameter $\varepsilon \in (0, 1/2)$, dimension $d \geq 1$, and scale $r \in [1, d]$. For any $k \geq 1$, there exists a randomized map $f : \{0, 1\}^d \rightarrow \{0, 1\}^k$ and an absolute constant $C > 0$, satisfying the following for any fixed $x, y \in \{0, 1\}^d$:*

- if $\|x - y\|_1 \leq r$, then $\Pr_f[\|f(x) - f(y)\|_1 \leq k/2] \geq 1 - e^{-C\varepsilon^2 k}$;
- if $\|x - y\|_1 \geq (1 + \varepsilon)r$, then $\Pr_f[\|f(x) - f(y)\|_1 > (1 + \varepsilon/2) \cdot k/2] \geq 1 - e^{-C\varepsilon^2 k}$.

The map f can be constructed via a random projection over $GF(2)$. That is, take $f(x) = Ax$, where A is a $k \times d$ matrix for $k = O(\log(n)/\varepsilon^2)$, with each entry being 1 with some fixed probability p , and zero otherwise. The probability p depends solely on r . The rest of the algorithm proceeds as before, with the exception that the “base” data structure \mathcal{R} is particularly simple: just store the answer for any dimension-reduced query point $f(q) \in \{0, 1\}^k$. Since there are only $2^k = n^{O(1/\varepsilon^2)}$ such possible queries, and computing $f(q)$ takes $O(dk)$ time, we get the following result.

Theorem 2.4 ([Kushilevitz, Ostrovsky, and Rabani \[ibid.\]](#)). *Fix $\varepsilon \in (0, 1/2)$ and dimension $d \geq 1$. There is a $(1 + O(\varepsilon), r)$ -ANN data structure over $(\{0, 1\}^d, \ell_1)$ using $n^{O(1/\varepsilon^2)} + O(d(n + k))$ space and $O(d \cdot \frac{\log n}{\varepsilon^2})$ query time.*

As a final remark, we note we cannot obtain improved space bounds by improving the dimension reduction lemmas [2.1](#) and [2.3](#). Indeed the above lemma are tight as proven in [Jayram and Woodruff \[2013\]](#). There was however work on improving the *run-time complexity* for computing a dimension reduction map, improving over the naïve bound of $O(dk)$; see [Ailon and Chazelle \[2009\]](#), [Dasgupta, R. Kumar, and Sarlós \[2010\]](#), [Ailon and Liberty \[2013\]](#), [Krahmer and Ward \[2011\]](#), [Nelson, Price, and Wootters \[2014\]](#), and [Kane and Nelson \[2014\]](#).

2.2 ANN via space partitions: Locality-Sensitive Hashing. While dimension reduction yields ANN data structure with polynomial space, this is not enough in applications, where one desires space as close as possible to linear in n . This consideration led to the

⁷In fact, it has been shown that spaces for which analogs of [Lemma 2.1](#) hold are “almost” Hilbert spaces [Johnson and A. Naor \[2009\]](#).

following, alternative approach, which yields smaller space bounds, albeit at the cost of increasing the query time to something of the form n^ρ where $\rho \in (0, 1)$.

The new approach is based on randomized space partitions, and specifically on Locality-Sensitive Hashing, introduced in [Indyk and Motwani \[1998\]](#).

Definition 2.5 (Locality-Sensitive Hashing (LSH)). *Fix a metric space (X, D) , scale $r > 0$, approximation $c > 1$ and a set U . Then a distribution \mathcal{H} over maps $h : X \rightarrow U$ is called (r, cr, p_1, p_2) -sensitive if the following holds for any $x, y \in X$:*

- if $D(x, y) \leq r$, then $\Pr_h[h(x) = h(y)] \geq p_1$;
- if $D(x, y) > cr$, then $\Pr_h[h(x) = h(y)] \leq p_2$.

The distribution \mathcal{H} is called an LSH family, and has quality $\rho = \rho(\mathcal{H}) = \frac{\log 1/p_1}{\log 1/p_2}$.

In what follows we require an LSH family to have $p_1 > p_2$, which implies $\rho < 1$. Note that LSH mappings are also oblivious: the distribution \mathcal{H} does not depend on the point-set P or the query q .

Using LSH, [Indyk and Motwani \[ibid.\]](#) show how to obtain the following ANN data structure.

Theorem 2.6 ([Indyk and Motwani \[ibid.\]](#)). *Fix a metric $\mathfrak{M} = (X, D)$, a scale $r > 0$, and approximation factor $c > 1$. Suppose the metric admits a (r, cr, p_1, p_2) -sensitive LSH family \mathcal{H} , where the map $h(\cdot)$ can be stored in σ space, and, for given x , can be computed in τ time; similarly, assume that computing distance $D(x, y)$ takes $O(\tau)$ time. Let $\rho = \rho(\mathcal{H}) = \frac{\log 1/p_1}{\log 1/p_2}$. Then there exists a (c, r) -ANN data structure over \mathfrak{M} achieving query time $Q = O(n^\rho \cdot \tau \frac{\log_{1/p_2} n}{p_1})$ and space $S = O(n^{1+\rho} \cdot \frac{1}{p_1} + n^\rho \frac{1}{p_1} \cdot \sigma \cdot \log_{1/p_2} n)$ (in addition to storing the original dataset P). The time needed to build this data structure is $O(S \cdot \tau)$.*

While we describe some concrete LSH families later on, for now, one can think of the parameters τ, σ as being proportional to the dimension of the space (although this is not always the case).

The overall idea of the algorithm is to use an LSH family as a pre-filter for the dataset P . In particular, for a random partition h from the family \mathcal{H} , the query point q will likely collide with its near neighbor (with probability at least p_1), but with few points at a distance $\geq cr$, in expectation at most $p_2 \cdot n$ of them. Below we show how an extension of this idea yields [Theorem 2.6](#).

Proofsketch. Given an LSH family \mathcal{H} , we can build a new, derived LSH family via a certain tensoring operation. In particular, for an integer $k \geq 1$, consider a new distribution \mathcal{G}_k over maps $g : X \rightarrow U$, where $g(\cdot)$ is obtained by picking k i.i.d. functions h_1, \dots, h_k

chosen from \mathcal{H} and setting $g(x) = (h_1(x), h_2(x), \dots, h_k(x))$. Then, if \mathcal{H} is (r, cr, p_1, p_2) -sensitive, \mathcal{G}_k is (r, cr, p_1^k, p_2^k) -sensitive. Note that the parameter ρ of the hash family does not change, i.e., $\rho(\mathcal{G}_k) = \rho(\mathcal{H})$.

The entire ANN data structure is now composed of L dictionary data structures (e.g., hash tables discussed in the previous section), where $L, k \geq 1$ are parameters to fix later. The i -th hash table is constructed as follows. Pick a map g_i uniformly at random from \mathcal{G}_k , and store the set $g_i(P)$ in the dictionary structure. At the query time, we iterate over $i = 1 \dots L$. For a given i , we compute $g_i(q)$, and use the dictionary structure to obtain the set of “candidate” points $Q_i = \{p \in P : g_i(p) = g_i(q)\}$. For each candidate point we compute the distance from q to that point. The process is stopped when all Q_i ’s are processed, or when a point within distance cr to q is found, whichever happens first.

To analyze the success probability, we note that the dictionary structure i succeeds if $p^* \in Q_i$, where p^* is the assumed point at distance at most r from q . This happens with probability at least p_1^k . Thus, we can take $L = O(1/p_1^k)$ such dictionary structures, and thus guarantee success with a constant probability.

The expected query time is $O(L(k\tau + Ln \cdot p_2^k \cdot \tau))$, which includes both the computation of the maps $g_1(q), \dots, g_L(q)$ and the of distances to the candidates in sets Q_1, \dots, Q_L . We can now derive the value of k that minimizes the above, obtaining $k = \lceil \log_{1/p_2} n \rceil \leq \log_{1/p_2} n + 1$, and hence $L = O(n^\rho/p_1)$.

Finally, note that the space usage is $O(Ln)$ for the dictionary structures, plus $O(Lk\sigma)$ for the description of the maps. \square

2.3 Space partitions: LSH constructions. [Theorem 2.6](#) assumes the existence of an LSH family \mathcal{H} with a parameter $\rho < 1$. In what follows we show a few examples of such families.

1. *Hamming space* $\{0, 1\}^d$, with $\rho = 1/c$. The distribution \mathcal{H} is simply a projection on a random coordinate i : $\mathcal{H} = \{h_i : h_i(x) = x_i, i = 1, \dots, d\}$. This family is $(r, cr, 1 - r/d, 1 - cr/d)$ -sensitive, and hence $\rho \leq 1/c$ [Indyk and Motwani \[1998\]](#).

This LSH scheme is near-optimal for the Hamming space, as described in [Section 2.4](#). We also note that, since ℓ_2 embeds isometrically into ℓ_1 (see [Section 6](#)), this result extends to ℓ_2 as well.

2. *Euclidean space* (\mathbb{R}^d, ℓ_2) , with $\rho < 1/c$. In [Datar, Immorlica, Indyk, and Mirrokni \[2004\]](#), the authors introduced an LSH family which slightly improves over the above construction. It is based on random projections in ℓ_2 . In particular, define a random map $h(x)$ as $h(x) = \lfloor \frac{\langle x, g \rangle}{wr} + b \rfloor$, where g is a random d -dimensional Gaussian vector, $b \in [0, 1]$, and $w > 0$ is a fixed parameter. It can be shown that, for any fixed $c > 1$, there exists $w > 0$ such that $\rho < 1/c$.

3. *Euclidean space* (\mathbb{R}^d, ℓ_2) , with $\rho \rightarrow 1/c^2$. In [Andoni and Indyk \[2006\]](#), the authors showed an LSH family with a much better ρ , which later turned out to be optimal (see [Section 2.4](#)). At its core, the main idea is to partition the space into Euclidean balls.⁸ It proceeds in two steps: 1) perform a random dimension reduction A to dimension t (a parameter), and 2) partition \mathbb{R}^t into balls. Since it is impossible to partition the space \mathbb{R}^t into balls precisely⁹ when $t \geq 2$, instead one performs “ball carving”. The basic idea is to consider a sequence of randomly-centered balls B_1, B_2, \dots , each of radius wr for some parameter $w > 1$, and define the map $h(x)$, for $x \in \mathbb{R}^d$, to be the index i of the first ball B_i containing the point Ax . Since we want to cover an infinite space with balls of finite volume, the above procedure needs to be modified slightly to terminate in finite time. The modified procedure runs in time $T = t^{O(t)}$.

Overall, optimizing for w , one can obtain $\rho = 1/c^2 + \frac{O(\log t)}{\sqrt{t}}$ which tends to $1/c^2$ as $t \rightarrow \infty$. The time to hash is $\tau = O(Tt + dt)$, where T depends exponentially on the parameter t , i.e., $T = t^{\Theta(t)}$. For the ANN data structure, the optimal choice is $t = O(\log n)^{2/3}$, resulting in $\rho = 1/c^2 + \frac{O(\log \log n)}{(\log n)^{1/3}}$.

The ℓ_2 LSH families can be extended to other ℓ_p 's. For $p < 1$, [Datar, Immorlica, Indyk, and Mirrokni \[2004\]](#) showed one can use method 2 as described above, but using p -stable distributions instead of Gaussians. See [Section 6](#) for other extensions for $p > 1$.

We remark that there is a number of other widely used LSH families, including min-hash [Broder \[1997\]](#) and [Broder, Glassman, Manasse, and Zweig \[1997\]](#) and simhash [M. S. Charikar \[2002\]](#), which apply to different notions of similarity between points. See [Andoni and Indyk \[2008\]](#) for an overview.

2.4 Space partitions: impossibility results. It is natural to explore the limits of LSH families and ask what is the best ρ one can obtain for a given metric space as a function of the approximation $c > 1$. In [Motwani, A. Naor, and Panigrahy \[2007\]](#) and [O'Donnell, Y. Wu, and Zhou \[2014\]](#), it was proven that the LSH families [Indyk and Motwani \[1998\]](#) and [Andoni and Indyk \[2006\]](#) from the previous section are near-optimal: for the Hamming space, we must have $\rho \geq 1/c - o(1)$, and for the Euclidean space, $\rho \geq 1/c^2 - o(1)$. Below is the formal statement from [O'Donnell, Y. Wu, and Zhou \[2014\]](#).

Theorem 2.7. *Fix dimension $d \geq 1$ and approximation $c \geq 1$. Let \mathcal{H} be a (r, cr, p_1, p_2) -sensitive LSH family over the Hamming space, and suppose $p_2 \geq 2^{-o(d)}$. Then $\rho \geq 1/c - o_d(1)$.*

⁸In contrast, the above LSH family can be seen as partitioning the space into cubes: when considering the k -tensored family $\mathcal{G} = \mathcal{H}^k$, the resulting map $g \in \mathcal{G}$ is equivalent to performing a random dimension reduction (by multiplying by a random $k \times d$ Gaussian matrix), followed by discretization of the space into cubes.

⁹This is also termed *tessellation* of the space.

Note that the above theorem also immediately implies $\rho \geq 1/c^2 - o(1)$ for the Euclidean space, by noting that $\|x - y\|_1 = \|x - y\|_2^2$ for binary vectors x and y .

Finally, we remark that some condition on p_2 is necessary, as there exists an LSH family with $p_2 = 0$, $p_1 = 2^{-O(d)}$ and hence $\rho = 0$. To obtain the latter, one can use the “ball carving” family of [Andoni and Indyk \[2006\]](#), where the balls have radius $wr = cr/2$. Note however that such a family results in query time that is at least exponential in d , which LSH algorithms are precisely designed to circumvent.

3 (More) Deterministic algorithms

A drawback of data structures described in the previous section is that they allow “false negatives”: with a controllable but non-zero probability, the data structure can report nothing even if the ball $B(q, r)$ is non-empty. Although most of the data structures described in the literature have this property, it is possible to design algorithms with stronger guarantees, including deterministic ones.

The first step in this direction was an observation (already made in [Kushilevitz, Ostrovsky, and Rabani \[2000\]](#)) that for a finite metric (X, D) supported by (c, r) -ANN data structures, it is possible to construct a data structure that provides accurate answers to *all* queries $q \in X$. This is because one can construct and use $O(\log |X|)$ independent data structures, reducing the probability of failure to $\frac{1}{3^{|X|}}$. By taking a union bound over all $q \in X$, the constructed data structure works, with probability at least $2/3$, for all queries X . Note that the space and query time bounds of the new data structure are $O(\log |X|)$ times larger than the respective bounds for (c, r) -ANN. Unfortunately, the algorithm for constructing such data structures has still a non-zero failure probability, and no deterministic polynomial-time algorithm for this task is known.

The first deterministic polynomial-time algorithm for constructing a data structure that works for all queries $q \in X$ appeared in [Indyk \[2000a\]](#). It was developed for d -dimensional Hamming spaces, and solved a (c, r) -ANN with an approximation factor $c = 3 + \varepsilon$ for any $\varepsilon > 0$. The data structure had $d(1/\varepsilon)^{O(1)}$ query time and used $dn^{(1/\varepsilon)^{O(1)}}$ space. It relied on two components. The first component, “densification”, was a deterministic analog of the mapping in [Lemma 2.3](#), which was shown to hold with $k = (d/\varepsilon)^{O(1)}$. Retrospectively, the mapping can be viewed as being induced by an adjacency matrix of an *unbalanced expander* [Guruswami, Umans, and Vadhan \[2009\]](#).

Definition 3.1 (Expander). An (r, α) -unbalanced expander is a bipartite simple graph $G = (U, V, E)$, $|U| = d$, $|V| = k$, with left degree Δ such that for any $X \subset U$ with $|X| \leq r$, the set of neighbors $N(X)$ of X has size $|N(X)| \geq (1 - \alpha)\Delta|X|$.

Given such a graph G , one can construct a mapping $f = f_G : \{0, 1\}^d \rightarrow \Sigma^k$ for some finite alphabet Σ by letting $f(x)_j$ to be the concatenation of all symbols x_i such

that $(i, j) \in E$. Let $H(x, y)$ be the Hamming distance between x and y , i.e., the number of coordinates on which x and y differ. We have that:

- $H(f(x), f(y)) \leq \Delta H(x, y)$, since each difference between a and b contributes to at most Δ differences between $f(x)$ and $f(y)$. In particular $H(f(x), f(y)) \leq \Delta r(1 - \varepsilon)$ if $H(x, y) \leq r(1 - \varepsilon)$.
- if $H(x, y) \geq r$, then $H(f(x), f(y)) \geq (1 - \alpha)\Delta r$ (from the expansion property).

Thus, setting $\alpha = \varepsilon/2$ yields guarantees analogous to [Lemma 2.3](#), but using a deterministic mapping, and with coordinates of $f(x)$ in Σ , not $\{0, 1\}$. To map into binary vectors, we further replace each symbol $f(x)_j$ by $C(f(x)_j)$, where $C : \Sigma \rightarrow \{0, 1\}^s$ is an *error-correcting code*, i.e., having the property that for any distinct $a, b \in \Sigma$ we have $H(C(a), C(b)) \in [s(1/2 - \varepsilon), s(1/2 + \varepsilon)]$. We then use off-the-shelf constructions of expanders [Guruswami, Umans, and Vadhan \[ibid.\]](#) and codes [Guruswami, Rudra, and Sudan \[2014\]](#) to obtain the desired mapping $g = C \circ f : \{0, 1\}^d \rightarrow \{0, 1\}^{ks}$.

The second component partitions the coordinates of points $g(x)$ into blocks $S_1 \dots S_t$ of size $\log(n)/\varepsilon^{O(1)}$ such that an analog of [Lemma 2.3](#) holds for all projections $g(x)_{S_l}$ and $g(y)_{S_l}$, where $x, y \in P, l = 1 \dots t$. Such a partitioning can be shown to exist using the probabilistic method, and can be computed deterministically in time polynomial in n via the method of conditional probabilities. Unfortunately, this property does not extend to the case where one of the points (say, x) is a query point from $X - P$. Nevertheless, by averaging, there must be at least one block S_l such that $H(g(x)_{S_l}, g(y)_{S_l}) \leq H(g(x), g(y))/t$, where y is the nearest neighbor of x in P . It can be then shown that an approximate near neighbor of $g(x)_{S_l}$ in $\{g(y)_{S_l} : y \in P\}$ is an approximate nearest neighbor of x in P . Finding the nearest neighbor in the space restricted to a single block S_l can be solved via exhaustive storage using $n^{1/\varepsilon^{O(1)}}$ space, as in [Theorem 2.4](#).

Perhaps surprisingly, the above construction is the only known example of a polynomial-size deterministic approximate near neighbor data structure with a constant approximation factor. However, more progress has been shown for an “intermediary” problem, where the data structure avoids false negatives by reporting a special symbol \perp .

Definition 3.2 ((c, r, δ) -Approximate Near Neighbor Without False Negatives (ANNWFN)). *Given a set P of n points in a metric space (X, D) , build a data structure \mathcal{S} that, given any query point $q \in X$ such that $B(q, r) \cap P \neq \emptyset$, \mathcal{S} returns an element of $(B(q, cr) \cap P) \cup \{\perp\}$, and $\Pr_{\mathcal{S}}[\mathcal{S} \text{ returns } \perp] \leq \delta$.*

A $(1 + \varepsilon, r, \delta)$ -ANNWFN data structure with bounds similar to those in [Theorem 2.4](#) was given in [Indyk \[2000a\]](#). It used densification and random block partitioning as described above. However, thanks to randomization, block partitioning could be assumed to hold even for the query point with high probability.

Obtaining “no false negatives” analogs of [Theorem 2.2](#) turned out to be more difficult. The first such data structure was presented in [Pagh \[2016\]](#), for the Hamming space, achieving query time of the form (roughly) $dn^{1.38/c}$. Building on that work, very recently, [Ahle \[2017\]](#) improved the bound to (roughly) $dn^{1/c}$, achieving the optimal runtime exponent.

In addition to variants of densification and random block partitioning, the latter algorithm uses a generalization of the space partitioning method from [Section 2.2](#), called *locality sensitive filtering*. Such objects can be constructed deterministically in time and space roughly exponential in the dimension. Unfortunately, random block partitioning leads to blocks whose length is larger than $\log n$ by at least a (large) constant, which results in large (although polynomial) time and space bounds. To overcome this difficulty, [Ahle \[ibid.\]](#) shows how to combine filters constructed for dimension d to obtain a filter for dimension $2d$. This is achieved by using *splitters* [M. Naor, Schulman, and Srinivasan \[1995\]](#), which can be viewed as families of partitions of $\{1 \dots 2d\}$ into pairs of sets $(S_1, \overline{S_1}), (S_2, \overline{S_2}), \dots$ of size d , such that for any x, y , there is a pair $(S_l, \overline{S_l})$ for which $H(x_{S_l}, y_{S_l}) = H(x_{\overline{S_l}}, y_{\overline{S_l}}) \pm 1$. The construction multiplies the space bound by a factor quadratic in d , which makes it possible to apply it a small but super-constant number of times to construct filters for (slightly) super-logarithmic dimension.

4 Data-dependent approach

In the earlier sections, we considered ANN data structures that are based on random and deterministic space partitions. The unifying feature of all of the above approaches is that the partitions used are independent of the dataset. This “data-independence” leads to certain barriers: for instance, the *best possible* LSH exponent is $\rho \geq 1/c - o(1)$ for the ℓ_1 distance and $\rho \geq 1/c^2 - o(1)$ for ℓ_2 (see [Section 2.4](#)). In this section, we show how to improve upon the above results significantly if one allows partitions to depend on the dataset.

This line of study has been developed in a sequence of recent results [Andoni, Indyk, H. L. Nguyễn, and Razenshteyn \[2014\]](#), [Andoni and Razenshteyn \[2015\]](#), and [Andoni, Laarhoven, Razenshteyn, and Waingarten \[2017\]](#). However, even before these works, the data-dependent approach had been very popular in practice (see, e.g., surveys [Wang, Shen, Song, and Ji \[2014\]](#) and [Wang, Liu, S. Kumar, and Chang \[2016\]](#)). Indeed, real-world datasets often have some implicit or explicit structure, thus it pays off to tailor space partitions to a dataset at hand. However, the theoretical results from [Andoni, Indyk, H. L. Nguyễn, and Razenshteyn \[2014\]](#), [Andoni and Razenshteyn \[2015\]](#), and [Andoni, Laarhoven, Razenshteyn, and Waingarten \[2017\]](#) improve upon data-independent partitions for *arbitrary* datasets. Thus, one must show that *any* set of n points has some structure that makes the ANN problem easier.

4.1 The result. In [Andoni and Razenshteyn \[2015\]](#) (improving upon [Andoni, Indyk, H. L. Nguyen, and Razenshteyn \[2014\]](#)), the following result has been shown.

Theorem 4.1. *For every $c > 1$, there exists a data structure for (c, r) -ANN over (\mathbb{R}^d, ℓ_2) with space $n^{1+\rho} + O(nd)$ and query time $n^\rho + dn^{o(1)}$, where*

$$\rho \leq \frac{1}{2c^2 - 1} + o(1).$$

This is much better than the best LSH-based data structure, which has $\rho = \frac{1}{c^2} + o(1)$. For instance, for $c = 2$, the above theorem improves the query time from $n^{1/4+o(1)}$ to $n^{1/7+o(1)}$, while using less memory.

Next, we describe the new approach at a high level.

4.2 Simplification of the problem. Before describing new techniques, it will be convenient to introduce a few simplifications. First, we can assume that $d = \log^{1+o(1)} n$, by applying [Lemma 2.1](#). Second, we can reduce the general ANN problem over (\mathbb{R}^d, ℓ_2) to the *spherical* case: where dataset and queries lie on the unit sphere $S^{d-1} \subset \mathbb{R}^d$ (see [Razenshteyn \[2017\]](#), pages 55–56). Both the dimension reduction and the reduction to the spherical case incur a negligible loss in the approximation¹⁰. After the reduction to the spherical case, the distance to the near neighbor r can be made to be any function of the number of points n that tends to zero as $n \rightarrow \infty$ (for example, $r = \frac{1}{\log \log n}$).

4.3 Data-independent partitions for a sphere. In light of the above discussion, we need to solve the (c, r) -ANN problem for S^{d-1} , where $d = \log^{1+o(1)} n$ and $r = o(1)$. Even though the final data structure is based on data-dependent partitions, we start with developing a *data-independent* LSH scheme for the unit sphere, which will be later used as a building block.

The LSH scheme is parametrized by a number $\eta > 0$. Consider a sequence of i.i.d. samples from a standard d -dimensional Gaussian distribution $N(0, 1)^d$: $g_1, g_2, \dots, g_t, \dots \in \mathbb{R}^d$. The hash function $h(x)$ of the point $x \in S^{d-1}$ is then defined as $\min_t \{t \geq 1 \mid \langle x, g_t \rangle \geq \eta\}$. This LSH family gives the following exponent ρ for distances r and cr :

$$(1) \quad \rho = \frac{\log 1/p_1}{\log 1/p_2} = \frac{4 - c^2 r^2}{4 - r^2} \cdot \frac{1}{c^2} + \delta(r, c, \eta),$$

where $\delta(r, c, \eta) > 0$ and $\delta(r, c, \eta) \rightarrow 0$ as $\eta \rightarrow \infty$. Thus, the larger the value of the threshold η is, the more efficient the resulting LSH scheme is. At the same time, η affects the efficiency of hash functions. Indeed, one can show that with very high probability

¹⁰Approximation c reduces to approximation $c - o(1)$.

$\max_{x \in S^{d-1}} h(x) \leq e^{(1+o(1))\eta^2/2} \cdot d^{O(1)}$, which bounds the hashing time as well as the number of Gaussian vectors to store.

Consider the expression (1) for the exponent ρ in more detail. If $r = o(1)$, then we obtain $\rho = \frac{1}{c^2} + o(1)$, which matches the guarantee of the best data-independent LSH for ℓ_2 . This is hardly surprising, since, as was mentioned above, the general ANN problem over ℓ_2 can be reduced to the (c, r) -ANN problem over the sphere for $r = o(1)$. If $r \approx 2/c$, then ρ is close to zero, and, indeed, the $(c, 2/c)$ -ANN problem on the sphere is trivial (any point can serve as an answer to any valid query).

Between these two extremes, there is a point $r \approx \frac{\sqrt{2}}{c}$ that is crucial for the subsequent discussion. Since the distance between a pair of random points on S^{d-1} is close to $\sqrt{2}$ with high probability, the problem where r is *slightly* smaller than $\frac{\sqrt{2}}{c}$ has the following interpretation: if one is guaranteed to have a data point within distance r from the query, find a data point that is a bit closer to the query than a typical point on the sphere. For $r \approx \frac{\sqrt{2}}{c}$, the Equation (1) gives exponent $\rho \approx \frac{1}{2c^2-1}$, which is significantly smaller than the bound $\frac{1}{c^2}$ one is getting for $r = o(1)$. Later, using a certain data-dependent partitioning procedure, we will be able to reduce the *general* ANN problem on the sphere to this intermediate case of $r \approx \frac{\sqrt{2}}{c}$, thus obtaining the ANN data structure with the exponent $\rho = \frac{1}{2c^2-1} + o(1)$. This significantly improves upon the best possible LSH for ℓ_2 from Section 2, which yields $\rho = \frac{1}{c^2} + o(1)$.

4.4 Data-dependent partitions. We now describe at a high level how to obtain a data structure with space $n^{1+\rho}$ and query time n^ρ , where $\rho = \frac{1}{2c^2-1} + o(1)$, for the (c, r) -ANN problem on the sphere for general $r > 0$. If $r \geq \frac{\sqrt{2}}{c} - o(1)$, then we can simply use the data-independent LSH described above. Now suppose r is nontrivially smaller than $\frac{\sqrt{2}}{c}$.

We start with finding and removing *dense low-diameter clusters*. More precisely, we repeatedly find a point $u \in S^{d-1}$ such that $|P \cap B(u, \sqrt{2} - \varepsilon)| \geq \tau n$, where $\varepsilon, \tau = o(1)$, and set $P := P \setminus B(u, \sqrt{2} - \varepsilon)$. We stop when there are no more dense clusters remaining. Then we proceed with clusters and the remainder separately. Each cluster is enclosed in a ball of radius $1 - \Omega(\varepsilon^2)$ and processed recursively. For the remainder, we sample one partition from the data-independent LSH family described above, apply it to the dataset, and process each resulting part of the dataset recursively. During the query stage, we (recursively) query the data structure for *every* cluster (note that the number of clusters is at most $1/\tau$), and for the remainder we query (again, recursively) a part of the partition, where the query belongs to. Each step of the aforementioned procedure makes progress as follows. For clusters, we decrease the radius by a factor of $1 - \Omega(\varepsilon^2)$. It means that we come slightly closer to the ideal case of $r \approx \frac{\sqrt{2}}{c}$, and the instance corresponding to the cluster becomes easier. For the remainder, we use the fact that there are at most τn data

points closer than $\sqrt{2} - \varepsilon$ to the query. Thus, when we apply the data-independent LSH, the expected number of data points in the same part as the query is at most $(\tau + p_2)n$, where p_2 is the probability of collision under the LSH family for points at the distance $\sqrt{2} - \varepsilon$. We set $\tau \ll p_2$, thus the number of colliding data points is around p_2n . At the same time, the probability of collision with the near neighbor is at least p_1 , where p_1 corresponds to the distance r . Since $r < \frac{\sqrt{2}}{c}$, we obtain an effective exponent of at most $\frac{1}{2c^2-1} + o(1)$. Note that we need to keep extracting the clusters recursively to be able to apply the above reasoning about the remainder set in each step.

One omission in the above high-level description is that the clusters are contained in smaller *balls* rather than *spheres*. This is handled by partitioning balls into thin annuli and treating them as spheres (introducing negligible distortion).

4.5 Time–space trade-off. In [Andoni, Laarhoven, Razenshteyn, and Waingarten \[2017\]](#), [Theorem 4.1](#) has been extended to provide a smooth time–space trade-off for the ANN problem. Namely, it allows to decrease the query time at a cost of increasing the space and vice versa.

Theorem 4.2. *For every $c > 1$ and every ρ_s, ρ_q such that*

$$(2) \quad c^2 \sqrt{\rho_q} + (c^2 - 1) \sqrt{\rho_s} \geq \sqrt{2c^2 - 1},$$

there exists a data structure for (c, r) -ANN over (\mathbb{R}^d, ℓ_2) with space $n^{1+\rho_s+o(1)} + O(nd)$ and query time $n^{\rho_q+o(1)} + dn^{o(1)}$.

The bound (2) interpolates between:

- The near-linear space regime: $\rho_s = 0, \rho_q = \frac{2}{c^2} - \frac{1}{c^4}$;
- The “balanced” regime: $\rho_s = \rho_q = \frac{1}{2c^2-1}$, where it matches [Theorem 4.1](#);
- The fast queries regime: $\rho_s = \left(\frac{c^2}{c^2-1}\right)^2, \rho_q = 0$.

For example, for $c = 2$, one can obtain any of the following trade-offs: space $n^{1+o(1)}$ and query time $n^{7/16+o(1)}$, space $n^{8/7+o(1)}$ and query time $n^{1/7+o(1)}$, and space $n^{16/9+o(1)}$ and query time $n^{o(1)}$.

[Theorem 4.2](#) significantly improves upon the previous ANN data structures in various regimes [Indyk and Motwani \[1998\]](#), [Kushilevitz, Ostrovsky, and Rabani \[2000\]](#), [Indyk \[2000b\]](#), [Panigrahy \[2006\]](#), and [Kapralov \[2015\]](#). For example, it improves the dependence on ε in [Theorem 2.2](#) from $O(\log(1/\varepsilon)/\varepsilon^2)$ to $O(1/\varepsilon^2)$.

4.6 Impossibility results. Similarly to the data-independent case, it is natural to ask whether exponent $\rho = \frac{1}{2c^2-1} + o(1)$ from [Theorem 4.1](#) is optimal for data-dependent space partitions. In [Andoni and Razenshiteyn \[2016\]](#), it was shown that the above ρ is near-optimal in a properly formalized framework of data-dependent space partitions. This impossibility result can be seen as an extension of the results discussed in [Section 2.4](#).

Specifically, [Andoni and Razenshiteyn \[ibid.\]](#) show that $\rho \geq \frac{1}{2c^2-1}$, where $\rho = \frac{\log 1/p_1}{\log 1/p_2}$ for p_1 and p_2 being certain natural counterparts of the LSH collision probabilities for the data-dependent case, even when we allow the distribution on the partitions to depend on a dataset. This result holds under two further conditions. First, as in [Section 2.4](#), we need to assume that p_2 is not too small.

The second condition is specific to the data-dependent case, necessary to address another necessary aspect of the space partition. For any dataset, where all the points are sufficiently well separated, we can build an “ideal” space partition, with $\rho = 0$, simply by considering its Voronoi diagram. However, this is obviously not a satisfactory space partition: it is algorithmically hard to compute fast where in the partition a fixed query point q falls to — in fact, it is precisely equivalent to the original nearest neighbor problem! Hence, to be able to prove a meaningful lower bound on ρ , we would need to restrict the space partitions to have low run-time complexity (e.g., for a given point q , we can compute the part where q lies in, in time $n^{o(1)}$). This precise restriction is well beyond reach of the current techniques (it would require proving computational lower bounds). Instead, [Andoni and Razenshiteyn \[ibid.\]](#) use a different, proxy restriction: they require that the *description complexity* of partitions is $n^{1-\Omega(1)}$. The latter restriction is equivalent to saying that the distribution of partitions (which may depend on the given dataset) is supported on a fixed (universal) family of partitions of the size $2^{n^{1-\Omega(1)}}$. This restriction, for instance, rules out the Voronoi diagram, since the latter has a description complexity of $\Omega(n)$. Furthermore, the description complexity of a randomized partition is a good proxy for the run-time complexity of a partition because in all the known constructions of random space partitions with a near-optimal ρ , the run-time complexity is at least the description complexity, which makes the requirement meaningful.

Overall, under the above two conditions, [Andoni and Razenshiteyn \[ibid.\]](#) show that $\rho \geq \frac{1}{2c^2-1} - o(1)$ for data-dependent random space partitions, and hence [Theorem 4.1](#) is essentially optimal in this framework.

4.7 ANN for ℓ_∞ . In this subsection we will describe another type of data-dependent data structure, for the ℓ_∞ norm. Historically, this was the first example of a data-dependent partitioning procedure used for ANN over high-dimensional spaces.

Theorem 4.3 (Indyk [2001]). *For every $0 < \varepsilon < 1$, there exists a deterministic data structure for $(c, 1)$ -ANN for $(\mathbb{R}^d, \ell_\infty)$ with approximation $c = O\left(\frac{\log \log d}{\varepsilon}\right)$, space $O(dn^{1+\varepsilon})$ and query time $O(d \log n)$.*

The algorithm relies on the following structural lemma.

Lemma 4.4. *Let $P \subset \mathbb{R}^d$ be a set of n points and $0 < \varepsilon < 1$. Then:*

1. *Either there exists an ℓ_∞ -ball of radius $O\left(\frac{\log \log d}{\varepsilon}\right)$ that contains $\Omega(n)$ points from P , or*
2. *There exists a “good” coordinate $i \in \{1, 2, \dots, d\}$ and a threshold $u \in \mathbb{R}$ such that for the sets $A = \{p \in P \mid p_i < u - 1\}$, $B = \{p \in P \mid u - 1 \leq p_i \leq u + 1\}$ and $C = \{p \in P \mid p_i > u + 1\}$ one has:*

$$(3) \quad \left(\frac{|A| + |B|}{n}\right)^{1+\varepsilon} + \left(\frac{|B| + |C|}{n}\right)^{1+\varepsilon} \leq 1$$

and $|A|/n, |C|/n \geq \Omega(1/d)$.

Using this lemma, we can build the data structure for $(c, 1)$ -ANN for $(\mathbb{R}^d, \ell_\infty)$ recursively. If there exists a ball $B(x, R)$ with $R = O\left(\frac{\log \log d}{\varepsilon}\right)$ such that $|P \cap B(x, R)| \geq \Omega(n)$ (Case 1), then we store x and R and continue partitioning $P \setminus B(x, R)$ recursively. If there exists a good coordinate $i \in \{1, 2, \dots, d\}$ and a threshold $u \in \mathbb{R}$ (Case 2), then we define sets A, B, C as in the above lemma and partition $A \cup B$ and $B \cup C$ recursively. We stop as soon as we reach a set that consists of $O(1)$ points.

The query procedure works as follows. Suppose there is a point in P within distance 1 from q (“the near neighbor”). If we are in Case 1, we check if the query point q lies in $B(x, R + 1)$. If it does, we return any data point from $B(x, R)$; if not, we query the remainder recursively. On the other hand, if we are in Case 2, we query $A \cup B$ if $q_i \leq u$, and $B \cup C$ otherwise. In this case we recurse on the part which is guaranteed to contain a near neighbor.

Overall, we always return a point within distance $O\left(\frac{\log \log d}{\varepsilon}\right)$, and it is straightforward to bound the query time by bounding the depth of the tree. We obtain the space bound of $O(dn^{1+\varepsilon})$ by using the property (3) to bound the number of times points that are replicated in the Case 2 nodes.

Surprisingly, the approximation $O(\log \log d)$ turns out to be *optimal* in certain restricted models of computation [Andoni, Croitoru, and Patrascu \[2008\]](#) and [Kapralov and Panigrahy \[2012\]](#), including for the approach from [Indyk \[2001\]](#).

5 Closest pair

A problem closely related to ANN is the closest pair problem, which can be seen as an “offline” version of ANN. Here, we are given a set P of n points, and we need to find a pair $p, q \in P$ of distinct points that minimize their distance.

A trivial solution is to compute the distance between all possible $\binom{n}{2}$ pairs of points and take the one that minimizes the distance. However this procedure has quadratic running time. As for the nearest neighbor problem, there is evidence that for, say, d -dimensional ℓ_2 space, the closest pair problem cannot be solved in time $n^{2-\alpha}d^{O(1)}$ for any constant $\alpha > 0$.

As with c -ANN, we focus on the approximate version of the problem. Furthermore, we consider the *decision version*, where we need to find a pair of points that are below a certain threshold r . The formal definition (for the randomized variant) follows.

Definition 5.1 ((c, r) -approximate close pair problem, or (c, r) -CP). *Given a set of points $P \subset X$ of size n , if there exist distinct points $p^*, q^* \in X$ with $D(p^*, q^*) \leq r$, find a pair of distinct points $p, q \in P$ such that $D(p, q) \leq cr$, with probability at least $2/3$.*

The (c, r) -CP problem is closely related to the (c, r) -ANN problem because we can solve the former using a data structure for the latter. In particular, one can run the following procedure: partition P into two sets A, B randomly; build (c, r) -ANN on the set A ; query every point $q \in B$. It is easy to see that one such run succeeds in solving a (c, r) -approximate close pair with probability at least $1/2 \cdot 2/3$. Repeating the procedure 3 times is enough to guarantee a success probability of $2/3$. If (c, r) -ANN under the desired metric can be solved with query time $Q(n)$ and preprocessing time $S(n)$, we obtain a solution for (c, r) -CP running in time $O(S(n) + nQ(n))$. For example, applying the reduction from above for (\mathbb{R}^d, ℓ_p) space for $p \in \{1, 2\}$, we immediately obtain an algorithm running in $O(dn^{1+\rho})$ time, where $\rho = \frac{1}{2c^p-1} + o(1)$ (Section 4).

Focusing on the case of ℓ_2 , and approximation $c = 1 + \varepsilon$, the above algorithm has runtime $O(n^{2-4\varepsilon+O(\varepsilon^2)d})$. It turns out that, for the ℓ_2 norm, one can obtain algorithms with a better dependence on ε , for small ε . In particular, the line of work from Valiant [2015], Karppa, Kaski, and Kohonen [2016], and Alman, Chan, and R. Williams [2016] led to the following algorithm:

Theorem 5.2 (Alman, Chan, and R. Williams [2016]). *Fix dimension $d \geq 1$, $r > 0$, and $\varepsilon \in (0, 1/2)$. Then, for any set of n points in \mathbb{R}^d , one can solve the $(1 + \varepsilon, r)$ -CP over ℓ_2 in time $O(n^{2-\Omega(\varepsilon^{1/3}/\log(1/\varepsilon))} + dn)$, with constant probability.*

Note that the running time bound in the above theorem is better than that obtained using LSH data structures, for small enough ε .

The main new technical ingredient is the *fast matrix multiplication* algorithm. In particular, suppose we want to multiply two matrices of size $n \times m$ and $m \times n$. Doing so naïvely takes time $O(n^2m)$. Starting with the work of Strassen [1969], there has been substantial work to improve this run-time; see also V. V. Williams [2012]. Below we state the running time of a fast matrix multiplication algorithm due to Coppersmith [1982], which is most relevant for this section.

Theorem 5.3 (Coppersmith [ibid.]). *Fix $n \geq 1$ and let $m \geq 1$ be such that $m \leq n^{0.172}$. One can compute the product of two matrices of sizes $n \times m$ and $m \times n$ in $O(n^2 \log^2 n)$ time.*

5.1 Closest pair via matrix multiplication. We now sketch the algorithm for the closest pair from Valiant [2015], which obtains $O(n^{2-\Omega(\sqrt{\varepsilon})}d)$ time. The algorithm is best described in terms of *inner products*, as opposed to distances as before. In particular, suppose we have a set of points $P \subset \mathbb{S}^d$ of unit norm, where all pairs of points have inner product in the range $[-\theta, \theta]$, except for one “special” pair that has inner product at least $c\theta$, for some scale $\theta > 0$ and approximation $c = 1 + \varepsilon$. Now the problem is to find this special pair—we term this problem (c, θ) -IP problem. We note that we can reduce $(1 + \varepsilon, r)$ -CP over ℓ_2 to $(1 + \Omega(\varepsilon), 1/2)$ -IP, by using the embedding of Schoenberg [1942], or Lemma 2.3 of Kushilevitz, Ostrovsky, and Rabani [2000].

A natural approach to the the IP problem is to multiply two $n \times d$ matrices: if we consider the matrix M where the rows are the points of P , then MM^t will have a large off-diagonal entry precisely for the special pair of points. This approach however requires at least n^2 computation time, since even the output of MM^t has size n^2 . Nevertheless, an extension of this approach gives a better run-time when c is very large (and hence $\theta < 1/c$ very small, i.e., all points except for the special pair are near-orthogonal). In particular, partition randomly the vectors from P into n/g groups $S_1, \dots, S_{n/g}$, each of size $O(g)$. For each group i , we sum the vectors S_i with random signs, obtaining vectors $v_i = \sum_{p_j \in S_i} \chi_j p_j$, where p_j are the points in P and χ_j are Rademacher random variables. Now the algorithm forms a matrix M with v_i 's as rows, and computes MM^t using fast matrix multiplication (Theorem 5.3). The two special points are separated with probability $1 - g/n$. Conditioning on this event, without loss of generality, we can assume that they are in group 1 and 2 respectively. Then, it is easy to note that $|(MM^t)_{12}| \approx \Theta(c \cdot \theta)$, whereas, for $(i, j) \neq (1, 2)$ and $i \neq j$, we have that $|(MM^t)_{ij}| \approx O(g \cdot \theta)$ with constant probability. Hence, we can identify the special pair in the product MM^t as long as $c \gg g$, and yields runtime $O(n^2/g^2)$, i.e., a $g^2 \ll c^2$ speed-up over the naïve algorithm (note that Theorem 5.3 requires that $d < n^{0.172}$).

The above approach requires c to be very large, and hence the challenge is whether we can reduce the case of $c = 1 + \varepsilon$ to the case of large c . Indeed, one method is

to use tensoring: for a fixed parameter k and any two vectors $x, y \in \mathbb{R}^d$, we consider $x^{\otimes k}, y^{\otimes k} \in \mathbb{R}^{d^k}$, for which $\langle x^{\otimes k}, y^{\otimes k} \rangle = (\langle x, y \rangle)^k$. Thus tensoring reduces the problem of $(1 + \varepsilon, 1/2)$ -IP to $((1 + \varepsilon)^k, 2^{-k})$ -IP, and hence we hope to use the above algorithm for $c = (1 + \varepsilon)^k \approx e^{\varepsilon k}$. If we use $t = \zeta \ln n$, for small constant ζ , we obtain $c = n^{\varepsilon \zeta}$, and hence we obtain a speed-up of $g^2 \approx c^2 = n^{2\varepsilon \zeta}$. One caveat here is that, after tensoring the vectors, we obtain vectors of dimension d^k , which could be much larger than n —then even writing down such vectors would take $\Omega(n^2)$ time. Yet, one can use a dimension reduction method, like [Lemma 2.1](#), to reduce dimension to $O(\frac{\log n}{\theta^k}) = \tilde{O}(n^{\xi \ln 2})$, which is enough to preserve all inner products up to additive, say, $0.1 \cdot \theta^k$. There are further details (e.g., we cannot afford to get high-dimensional vectors in the first place, even if we perform dimension-reduction), see [Valiant \[2015\]](#) and [Karppa, Kaski, and Kohonen \[2016\]](#) for more details.

The above algorithm yields a speed-up of the order of $n^{O(\varepsilon)}$, i.e., comparable to the speed-up via the LSH methods. To obtain a better speed-up, like in the [Theorem 5.2](#), one can replace the tensoring transformation with a more efficient one. Indeed, one can employ an *asymmetric* embedding $f, g : \mathbb{R}^d \rightarrow \mathbb{R}^m$, with the property that for any unit-norm vectors x, y , we have that $\langle f(x), g(y) \rangle = p(\langle x, y \rangle)$, where $p(\cdot)$ is a polynomial of choice. In particular, we require a polynomial $p(\cdot)$ that is small on the interval $[-\theta, \theta]$, as large as possible on $[(1 + \varepsilon)\theta, 1]$, and $p(1)$ is bounded. Note that the tensoring operation implements such an embedding with $p(a) = a^k$ and where $f(x) = g(x) = x^{\otimes k}$. However, there are more efficient polynomials: in fact, the optimal such polynomial is the Chebyshev polynomial. For example, for the degree- k Chebyshev polynomial $T_k(\cdot)$, we have that $T_k(1 + \varepsilon)/T_k(1) \approx e^{\sqrt{\varepsilon}k}$, which is in contrast to the above polynomial $p(a) = a^k$, for which $p(1 + \varepsilon)/p(1) \approx e^{\varepsilon k}$.

Using the Chebyshev polynomials, one can obtain a runtime of $n^{2-\Omega(\sqrt{\varepsilon})}$ for the IP and hence CP problem. To obtain the improved result from [Theorem 5.2](#), [Alman, Chan, and R. Williams \[2016\]](#) employ *randomized* polynomials, i.e., a distribution over polynomials where $p(\cdot)$ is small/large only with a certain probability. Without going into further details, the theorem below states the existence of such polynomials, which are used to obtain $n^{2-\Omega(\varepsilon^{1/3}/\log(1/\varepsilon))}$ run-time for the $(1 + \varepsilon, r)$ -CP problem.

Theorem 5.4 ([Alman, Chan, and R. Williams \[ibid.\]](#)). *Fix $d \geq 1, \theta \geq 1, s \geq 1$, and $\varepsilon > 0$. There exists a distribution over polynomials $P : \{0, 1\}^d \rightarrow \mathbb{R}$ of degree $O(\varepsilon^{-1/3} \log s)$, such that we have the following for any $x \in \{0, 1\}^d$:*

- if $\sum_{i=1}^d x_i \leq \theta$, then $|P(x)| \leq 1$ with probability at least $1 - 1/s$;
- if $\sum_{i=1}^d x_i \in (\theta, (1 + \varepsilon)\theta)$, then $|P(x)| > 1$ with probability at least $1 - 1/s$;
- if $\sum_{i=1}^d x_i > (1 + \varepsilon)\theta$, then $|P(x)| \geq s$ with probability at least $1 - 1/s$.

6 Extensions

In this section, we discuss several techniques that significantly extend the class of spaces which admit efficient ANN data structures.

6.1 Metric embeddings. So far, we have studied the ANN problem over the ℓ_1 , ℓ_2 and ℓ_∞ distances. A useful approach is to *embed* a metric of interest into $\ell_1/\ell_2/\ell_\infty$ and use one of the data structures developed for the latter spaces.

6.1.1 Deterministic embeddings.

Definition 6.1. For metric spaces $\mathfrak{M} = (X, D_X)$, $\mathfrak{N} = (Y, D_Y)$ and for $D \geq 1$, we say that a map $f: X \rightarrow Y$ is a bi-Lipschitz embedding with distortion D if there exists $\lambda > 0$ such that for every $x_1, x_2 \in X$ one has:

$$\lambda D_X(x_1, x_2) \leq D_Y(f(x_1), f(x_2)) \leq D \cdot \lambda D_X(x_1, x_2).$$

A bi-Lipschitz embedding of \mathfrak{M} into \mathfrak{N} with distortion D together with a data structure for (c, r) -ANN over \mathfrak{N} immediately implies a data structure for (cD, r') -ANN over \mathfrak{M} , where $r' = \frac{r}{\lambda D}$. However, space and query time of the resulting data structure depend crucially on the *computational efficiency* of the embedding, since, in particular, the query procedure requires evaluating the embedding on a query point.

As the following classic results show, any finite-dimensional normed or finite metric space can be embedded into finite-dimensional ℓ_∞ with small distortion.

Theorem 6.2 (Fréchet [1906] and Kuratowski [1935]). *If \mathfrak{M} is a finite metric space, which consists of N points, then \mathfrak{M} embeds into $(\mathbb{R}^N, \ell_\infty)$ with distortion $D = 1$ (isometrically).*

Theorem 6.3 (see, e.g., Wojtaszczyk [1996]). *For every $\varepsilon > 0$, every normed space $(\mathbb{R}^d, \|\cdot\|)$ embeds with distortion $1 + \varepsilon$ into $(\mathbb{R}^{d'}, \ell_\infty)$, where $d' = O(1/\varepsilon)^d$, via a linear map.*

However, the utility of Theorems 6.2 and 6.3 in the context of the ANN problem is limited, since the required dimension of the target ℓ_∞ space is very high (in particular, Theorem 6.3 gives a data structure with *exponential* dependence on the dimension). Moreover, even if we allow the distortion D of an embedding to be a large constant, the target dimension can not be improved much. As has been shown in Matoušek [1997], one needs at least $N^{\Omega(1/D)}$ -dimensional ℓ_∞ to “host” all the N -point metrics with distortion D . For d -dimensional norms, even as simple as ℓ_2 , the required dimension is $2^{\Omega D(d)}$ Figiel, Lindenstrauss, and Milman [1977] and Ball [1997].

More generally, (lower-dimensional) ℓ_∞ turns out to be not so useful of a target space, and only a handful of efficient embeddings into ℓ_∞ are known (for instance, such an embedding has been constructed in [Farach-Colton and Indyk \[1999\]](#) for the Hausdorff distance). Luckily, the situation drastically improves, if we allow *randomized* embeddings, see [Section 6.1.2](#) for the examples.

Instead of ℓ_∞ , one can try to embed a metric of interest into ℓ_1 or ℓ_2 . Let us list a few cases, where such embeddings lead to efficient ANN data structures.

- Using the result from [Johnson and Schechtman \[1982\]](#), one can embed (\mathbb{R}^d, ℓ_p) for $1 < p \leq 2$ into $(\mathbb{R}^{d'}, \ell_1)$ with distortion $1 + \varepsilon$, where $d' = O(d/\varepsilon^2)$. Moreover, the corresponding map is linear and hence efficient to store and apply. This reduction shows that the ANN problem over ℓ_p for $1 < p \leq 2$ is no harder than for the ℓ_1 case. However, later in this section we will show how to get a better ANN algorithm for the ℓ_p case using a different embedding.
- For the Wasserstein-1 distance (a.k.a. the Earth-Mover distance in the computer science literature) between probability measures defined on $\{1, 2, \dots, d\}^k$, one can use the results from [M. S. Charikar \[2002\]](#), [Indyk and Thaper \[2003\]](#), and [A. Naor and Schechtman \[2007\]](#), to embed it into $(\mathbb{R}^{d^{O(k)}}, \ell_1)$ with distortion $O(k \log d)$.
- The Levenshtein distance (a.k.a. edit distance) over the binary strings $\{0, 1\}^d$ can be embedded into $(\mathbb{R}^{d^{O(1)}}, \ell_1)$ with distortion $2^{O(\sqrt{\log d \log \log d})}$ [Ostrovsky and Rabani \[2007\]](#).

Let us note that there exist generic results concerned with embeddings into ℓ_1/ℓ_2 similar to [Theorem 6.2](#) and [Theorem 6.3](#).

Theorem 6.4 ([Bourgain \[1985\]](#) and [Linial, London, and Rabinovich \[1995\]](#)). *Any N -point metric embeds into $(\mathbb{R}^{O(\log N)}, \ell_2)$ with distortion $O(\log N)$.*

Theorem 6.5 ([John \[1948\]](#) and [Ball \[1997\]](#)). *Any normed space $(\mathbb{R}^d, \|\cdot\|)$ embeds into (\mathbb{R}^d, ℓ_2) with distortion \sqrt{d} via a linear map.*

[Theorem 6.4](#) does not give an embedding efficient enough for the ANN applications: computing it in one point requires time $\Omega(N)$. At the same time, [Theorem 6.5](#) is efficient and, together with ℓ_2 data structures, gives an ANN data structure for a general d -dimensional norm with approximation $O(\sqrt{d})$.

Since the ANN problem is defined for two specific distance scales (r and cr), we do not need the full power of bi-Lipschitz embeddings and sometimes can get away with weaker notions of embeddability. For example, the following theorem follows from the results of [Schoenberg \[1937\]](#).

In the theorem, $\ell_2(\mathbb{N})$ denotes the space of infinite sequences $(a_i)_{i=1}^\infty$ such that $\sum_i |a_i|^2 < +\infty$ and the norm of the sequence $\|a\|_2$ is equal to $(\sum_i |a_i|^2)^{1/2}$.

Theorem 6.6. *For every $1 \leq p < 2$ and every $d \geq 1$, there exists a map $f : \mathbb{R}^d \rightarrow \ell_2(\mathbb{N})$ such that for every $x, y \in \mathbb{R}^d$, one has:*

$$\|f(x) - f(y)\|_2^2 = \|x - y\|_p^p.$$

This embedding allows to use an ANN data structure for ℓ_2 with approximation c to get an ANN data structure for ℓ_p with approximation $c^{2/p}$. However, for this we need to make the embedding computationally efficient. In particular, the target must be finite-dimensional. This can be done, see H. Nguyen [2014] for details. As a result, for the ℓ_p distance for $1 \leq p < 2$, we are able to get the result similar to the one given by Theorem 4.2, where in (2) c^2 is replaced with c^p everywhere.

6.1.2 Randomized embeddings. It would be highly desirable to utilize the fact that every metric embeds well into ℓ_∞ (Theorems 6.2 and 6.3) together with the ANN data structure for ℓ_∞ from Section 4.7. However, as discussed above, spaces as simple as (\mathbb{R}^d, ℓ_1) or (\mathbb{R}^d, ℓ_2) require the target ℓ_∞ to have $2^{\Omega(d)}$ dimensions to be embedded with small distortion. It turns out, this can be remedied by allowing embeddings to be *randomized*. In what follows, we will consider the case of (\mathbb{R}^d, ℓ_1) , and then generalize the construction to other metrics.

The randomized embedding of (\mathbb{R}^d, ℓ_1) into $(\mathbb{R}^d, \ell_\infty)$ is defined as follows: we generate d i.i.d. samples u_1, u_2, \dots, u_d from the exponential distribution with parameter 1, and then the embedding f maps a vector $x \in \mathbb{R}^d$ into

$$\left(\frac{x_1}{u_1}, \frac{x_2}{u_2}, \dots, \frac{x_d}{u_d} \right).$$

Thus, the resulting embedding is linear. Besides that, it is extremely efficient to store (d numbers) and apply ($O(d)$ time).

Let us now understand how $\|f(x)\|_\infty$ is related to $\|x\|_1$. The analysis uses (implicitly) the *min-stability* property of the exponential distribution. One has for every $t > 0$:

$$\Pr_f[\|f(x)\|_\infty \leq t] = \prod_{i=1}^d \Pr \left[\frac{|x_i|}{u_i} \leq t \right] = \prod_{i=1}^d \Pr \left[u_i \geq \frac{|x_i|}{t} \right] = \prod_{i=1}^d e^{-|x_i|/t} = e^{-\|x\|_1/t}.$$

The random variable $\|f(x)\|_\infty$ does not have a finite first moment, however its mode is in the point $t = \|x\|_1$, which allows us to use $\|f(x)\|_\infty$ to estimate $\|x\|_1$. It is immediate

to show that for every $\delta > 0$, there exist $C_1, C_2 > 1$ with $C_1 = O(\log(1/\delta))$ and $C_2 = O(1/\delta)$ such that for every x , one has:

$$(4) \quad \Pr_f \left[\|f(x)\|_\infty \geq \frac{\|x\|_1}{C_1} \right] \geq 1 - \delta$$

and

$$(5) \quad \Pr_f [\|f(x)\|_\infty \leq C_2 \cdot \|x\|_1] \geq 1 - \delta$$

Thus, the map f has distortion $O\left(\frac{\log(1/\delta)}{\delta}\right)$ with probability $1 - \delta$. However, unlike the deterministic case, the randomized guarantees (4) and (5) are not sufficient for the reduction between ANN data structures (if $\delta \gg 1/n$). This is because the lower bound on $\|f(x)\|_\infty$ must apply simultaneously to all “far” points. In order to obtain a desired reduction, we need to use slightly different parameters. Specifically, for $0 < \varepsilon < 1$ one has:

$$\Pr_f \left[\|f(x)\|_\infty \geq \Omega\left(\frac{\|x\|_1}{\log n}\right) \right] \geq 1 - \frac{1}{10n}$$

and

$$\Pr_f \left[\|f(x)\|_\infty \leq O\left(\frac{\|x\|_1}{\varepsilon \cdot \log n}\right) \right] \geq n^{-\varepsilon}.$$

This allows us to reduce the $(c/\varepsilon, r)$ -ANN problem over (\mathbb{R}^d, ℓ_1) to $n^{O(\varepsilon)}$ instances of the (c, r') -ANN problem over $(\mathbb{R}^d, \ell_\infty)$. Indeed, we sample $n^{O(\varepsilon)}$ i.i.d. maps f_i as described above and solve the ANN problem over ℓ_∞ on the image of f_i . Far points remain being far with probability $1 - 1/10n$ each. Using the linearity of expectation and the Markov inequality, we observe that, with probability at least 0.9, no far point come close enough to the query point. At the same time, with probability at least $n^{-\varepsilon}$, the near neighbor does not move too far away, so, with high probability, at least one of the $n^{O(\varepsilon)}$ data structures succeeds. This reduction is quite similar to the use of Locality-Sensitive Hashing in [Section 2.2](#).

As a result, we get an ANN data structure for (\mathbb{R}^d, ℓ_1) with approximation $O\left(\frac{\log \log d}{\varepsilon^2}\right)$, query time $O(dn^\varepsilon)$ and space $O(dn^{1+\varepsilon})$. This is worse than the best ANN data structure for ℓ_1 based on (data-dependent) space partitions. However, the technique we used is very versatile and generalizes easily to many other distances. The ℓ_1 embedding was first used in [Andoni, Indyk, and Krauthgamer \[2009\]](#). Later, it was generalized [Andoni \[2009\]](#) to ℓ_p spaces for $p \geq 1$. To get such an embedding, one can divide every coordinate by the $(1/p)$ -th power of an exponential random variable. Finally, in [Andoni, H. L. Nguyen, Nikolov, Razenshteyn, and Waingarten \[2017\]](#) the same technique has been shown to work for Orlicz norms and top- k norms, which we define next.

Definition 6.7. Let $\psi : [0; +\infty) \rightarrow [0; +\infty)$ be a non-negative monotone increasing convex function with $\psi(0) = 0$. Then, an Orlicz norm $\|\cdot\|_\psi$ over \mathbb{R}^d is given by its unit ball K_ψ , defined as follows:

$$K_\psi = \left\{ x \in \mathbb{R}^d \mid \sum_{i=1}^d \psi(|x_i|) \leq 1 \right\}.$$

Clearly, ℓ_p norm for $p < \infty$ is Orlicz for $\psi(t) = t^p$.

Definition 6.8. For $1 \leq k \leq d$, we define the top- k norm of a vector from \mathbb{R}^d as the sum of k largest absolute values of the coordinates.

The top-1 norm is simply ℓ_∞ , while top- d corresponds to ℓ_1 .

To embed an Orlicz norm $\|\cdot\|_\psi$ into ℓ_∞ , we divide the coordinates using a random variable X with the c.d.f. $F_X(t) = \Pr[X \leq t] = 1 - e^{-\psi(t)}$. To embed the top- k norm, we use a truncated exponential distribution. All of the above embeddings introduce only a constant distortion.

Let us note that for the ℓ_p norms one can achieve approximation $2^{O(p)}$ [A. Naor and Rabani \[2006\]](#) and [Bartal and Gottlieb \[2015\]](#), which is an improvement upon the above $O(\log \log d)$ bound if p is sufficiently small.

6.2 ANN for direct sums. In this section we describe a vast generalization of the ANN data structure for ℓ_∞ from [Section 4.7](#). Namely, we will be able to handle *direct sums* of metric spaces.

Definition 6.9. Let $M_1 = (X_1, D_1)$, $M_2 = (X_2, D_2)$, ..., $M_k = (X_k, D_k)$ be metric spaces and let $\|\cdot\|$ be a norm over \mathbb{R}^k . Then the $\|\cdot\|$ -direct sum of M_1, M_2, \dots, M_k denoted by $\left(\bigoplus_{i=1}^k M_i\right)_{\|\cdot\|}$ is a metric space defined as follows. The ground set is the Cartesian product $X_1 \times X_2 \times \dots \times X_k$. The distance function D is given by the following formula.

$$D((x_1, x_2, \dots, x_k), (y_1, y_2, \dots, y_k)) = \|(D_1(x_1, y_1), D_2(x_2, y_2), \dots, D_k(x_k, y_k))\|.$$

It turns out that for many interesting norms $\|\cdot\|$ the following holds. If for metrics M_1, M_2, \dots, M_k there exist efficient ANN data structures, then the same holds for $\left(\bigoplus_{i=1}^k M_i\right)_{\|\cdot\|}$ (with a mild loss in the parameters).

The first result of this kind was shown in [Indyk \[2002\]¹¹](#) for the case of ℓ_∞ -direct sums. In what follows we denote by d the “complexity” of each metric M_i . That is, we assume

¹¹In [Indyk \[2002\]](#), a slightly weaker version of [Theorem 6.10](#) has been stated. First, it assumed *deterministic* data structures for the spaces M_i . This is straightforward to address by boosting the probability of success

it takes $O(d)$ time to compute the distance between two points, and that a point requires $O(d)$ space to store.

Theorem 6.10. *Let $c > 1$, $r > 0$ and $0 < \varepsilon < 1$. Suppose that each M_i admits a (c, r) -ANN data structure for n -point sets with space $n^{1+\rho}$ (in addition to storing the dataset) for some $\rho \geq 0$ and query time $Q(n)$. Then, there exists a data structure for (c', r) -ANN over $\left(\bigoplus_{i=1}^k M_i\right)_\infty$, where $c' = O\left(\frac{c \log \log n}{\varepsilon}\right)$, the space is $O(n^{1+\rho+\varepsilon})$ (in addition to storing the dataset), and the query time is $Q(n) \cdot \log^{O(1)} n + O(dk \log n)$.*

Informally speaking, compared to data structures for M_i , the data structure for $\left(\bigoplus_i M_i\right)_\infty$ loses $\frac{\log \log n}{\varepsilon}$ in approximation, n^ε in space, and $\log^{O(1)} n$ in query time.

Later, the result of Indyk [2002] was significantly extended Indyk [2004], Andoni, Indyk, and Krauthgamer [2009], Andoni [2009], and Andoni, H. L. Nguyễn, Nikolov, Razenshteyn, and Waingarten [2017], to support $\|\cdot\|$ -direct sums where $\|\cdot\|$ is an ℓ_p norm, an Orlicz norm, or a top- k norm. The main insight is that we can use the randomized embeddings of various norms into ℓ_∞ developed in Section 6.1.2, to reduce the case of $\|\cdot\|$ -direct sums to the case of ℓ_∞ -direct sums. Indeed, we described how to reduce the ANN problem over several classes of norms to n^ε instances of ANN over the ℓ_∞ distance at a cost of losing $O(1/\varepsilon)$ in the approximation. It is not hard to see that the exact same approach can be used to reduce the ANN problem over $\left(\bigoplus_{i=1}^k M_i\right)_{\|\cdot\|}$ to n^ε instances of ANN over $\left(\bigoplus_{i=1}^k M_i\right)_\infty$ also at a cost of losing $O(1/\varepsilon)$ in approximation.

6.3 Embeddings into direct sums. As Section 6.2 shows, for a large class of norms $\|\cdot\|$, we can get an efficient ANN data structure for any $\|\cdot\|$ -direct sum of metrics that admit efficient ANN data structures. This gives a natural approach to the ANN problem: embed a metric of interest into such a direct sum.

This approach has been successful in several settings. In Indyk [2002], the Fréchet distance between two sequences of points in a metric space is embedded into an ℓ_∞ -direct

for data structures for M_i using repetition. Second, the resulting space bound Indyk [2002] was worse. An improvement to the space bound has been described in Appendix A of the arXiv version of Andoni, H. L. Nguyễn, Nikolov, Razenshteyn, and Waingarten [2017]. Finally, the paper Indyk [2002] assumes ANN data structures for M_i with a slightly stronger guarantee. Namely, each point is assigned a *priority* from 1 to n , and if the near neighbor has priority t , we must return a point with priority at most t . It is not hard to solve the version with priorities using a standard ANN data structure (with $\log^{O(1)} n$ loss in space and query time). A naïve reduction builds an ANN data structure for points with priority at most t for every t . Then, we can run a binary search over the resulting priority. However, this gives a *linear* in n loss in space. To rectify this, we use a standard data structure technique: the decomposition of an interval into $O(\log n)$ *dyadic intervals*, i.e., intervals of the form $[2^k \cdot l + 1; 2^k \cdot (l + 1)]$ for integer k, l . Thus, we build an ANN data structure for every dyadic interval of priorities. This still gives $O(n)$ ANN data structures, however, each data point participates in at most $O(\log n)$ of them.

sums of Fréchet distances between shorter sequences. Together with [Theorem 6.10](#), this was used to obtain an ANN data structure for the Fréchet distance. In [Andoni, Indyk, and Krauthgamer \[2009\]](#), it is shown how to embed the Ulam metric (which is the edit distance between permutations of length d) into $\left(\bigoplus^d \left(\bigoplus^{O(\log d)} (\mathbb{R}^d, \ell_1)\right)_{\ell_\infty}\right)_{\ell_2^3}$ with a constant distortion which gives an ANN data structure with doubly-logarithmic approximation. At the same time, the Ulam distance requires distortion $\Omega\left(\frac{\log d}{\log \log d}\right)$ to embed into ℓ_1 [Andoni and Krauthgamer \[2010\]](#). This shows that (lower-dimensional) direct sums form a strictly more “powerful” class of spaces than ℓ_1 or ℓ_2 . Finally, in [Andoni, H. L. Nguyen, Nikolov, Razenshteyn, and Waingarten \[2017\]](#), it is shown that *any* symmetric norm over \mathbb{R}^d is embeddable into $\left(\bigoplus_{i=1}^{d^{O(1)}} \left(\bigoplus_{j=1}^d X_{ij}\right)_1\right)_\infty$ with constant distortion, where X_{ij} is \mathbb{R}^d equipped with the top- j norm. Together with the results from [Section 6.1.2](#) and [Section 6.2](#), this gives an ANN algorithm with approximation $(\log \log n)^{O(1)}$ for general *symmetric*¹² norms.

6.4 ANN for general norms. For *general* d -dimensional norms, the best known ANN data structure is obtained by combining [Theorem 6.5](#) with an efficient ANN data structure for ℓ_2 (for example, the one given by [Theorem 4.1](#)). This approach gives approximation $O(\sqrt{d}/\varepsilon)$ for space $d^{O(1)} \cdot n^{1+\varepsilon}$ and query time $d^{O(1)} \cdot n^\varepsilon$ for every constant $0 < \varepsilon < 1$. Very recently, the approximation $O(\sqrt{d}/\varepsilon)$ has been improved to $O\left(\frac{\log d}{\varepsilon^2}\right)$ [Andoni, A. Naor, Nikolov, Razenshteyn, and Waingarten \[2017\]](#) for the same space and time bounds if one is willing to relax the model of computation to the *cell-probe model*, where the query procedure is charged for *memory accesses*, but any computation is free. This ANN data structure heavily builds on a recent geometric result from [A. Naor \[2017\]](#): a bi-Lipschitz embedding (see [Definition 6.1](#)) of the shortest-path metric of *any* N -node expander graph [Hoory, Linial, and Wigderson \[2006\]](#) into an *arbitrary* d -dimensional normed space must have distortion at least $\Omega(\log_d N)$. At a very high level, this non-embeddability result is used to claim that any large bounded-degree graph, which *does* embed into a normed space, can not be an expander, and hence it must have a sparse cut. The existence of the sparse cut is then used, via a duality argument, to build a (data-dependent) random space partition family for a general d -dimensional normed space. The latter family is used to obtain the final data structure.

This approach can be further extended for several norms of interest to obtain proper, *time-efficient* ANN data structures, with even better approximations. For instance, [Andoni, A. Naor, Nikolov, Razenshteyn, and Waingarten \[2017\]](#) show how to get ANN with approximation $O(p)$ for the ℓ_p norms, improving upon the bound $2^{O(p)}$ from [A. Naor](#)

¹²Under permutations and negations of the coordinates.

and Rabani [2006] and Bartal and Gottlieb [2015]. Finally, for the Schatten- p norms of matrices, defined as the ℓ_p norm of the vector of singular values, one obtains approximation $O(p)$ as well, while the previous best approximation was polynomial in the matrix size (by relating the Schatten- p norm to the Frobenius norm).

Acknowledgments. The authors would like to thank Assaf Naor, Tal Wagner, Erik Waingarten and Fan Wei for many helpful comments.

References

- Thomas D Ahle, Rasmus Pagh, Ilya Razenshteyn, and Francesco Silvestri (2016). “On the complexity of maximum inner product search”. In: *Proc. 8th 35th ACM Symposium on Principles of Database Systems (PODS)* (cit. on p. 3306).
- Thomas Dybdahl Ahle (2017). “Optimal Las Vegas Locality Sensitive Data Structures”. In: *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*. IEEE (cit. on p. 3316).
- Nir Ailon and Bernard Chazelle (2009). “The Fast Johnson–Lindenstrauss Transform and Approximate Nearest Neighbors”. *SIAM J. Comput.* 39.1, pp. 302–322 (cit. on p. 3310).
- Nir Ailon and Edo Liberty (2013). “An Almost Optimal Unrestricted Fast Johnson–Lindenstrauss Transform”. *ACM Transactions on Algorithms* 9.3, p. 21 (cit. on p. 3310).
- Josh Alman, Timothy M Chan, and Ryan Williams (2016). “Polynomial representations of threshold functions and algorithmic applications”. In: *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 467–476 (cit. on pp. 3322, 3324).
- Alexandr Andoni (2009). “Nearest neighbor search: the old, the new, and the impossible”. PhD thesis. Massachusetts Institute of Technology (cit. on pp. 3328, 3330).
- Alexandr Andoni, Dorian Croitoru, and Mihai Patrascu (2008). “Hardness of nearest neighbor under L-infinity”. In: *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 424–433 (cit. on p. 3321).
- Alexandr Andoni and Piotr Indyk (2006). “Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions”. In: *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 459–468 (cit. on pp. 3313, 3314).
- (2008). “Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions”. *Communications of the ACM* 51.1, p. 117 (cit. on pp. 3305, 3313).
- Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer (2009). “Overcoming the ℓ_1 non-embeddability barrier: algorithms for product metrics”. In: *Proceedings of the twentieth*

- Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 865–874 (cit. on pp. [3328](#), [3330](#), [3331](#)).
- Alexandr Andoni, Piotr Indyk, Huy L Nguyễn, and Ilya Razenshteyn (2014). “Beyond locality-sensitive hashing”. In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 1018–1028 (cit. on pp. [3316](#), [3317](#)).
- Alexandr Andoni and Robert Krauthgamer (2010). “The computational hardness of estimating edit distance”. *SIAM Journal on Computing* 39.6, pp. 2398–2429 (cit. on p. [3331](#)).
- Alexandr Andoni, Thijs Laarhoven, Ilya Razenshteyn, and Erik Waingarten (2017). “Optimal hashing-based time-space trade-offs for approximate near neighbors”. In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 47–66 (cit. on pp. [3316](#), [3319](#)).
- Alexandr Andoni, Assaf Naor, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten (2017). “Data-Dependent Hashing via Nonlinear Spectral Gaps”. *Manuscript* (cit. on p. [3331](#)).
- Alexandr Andoni, Huy L Nguyễn, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten (2017). “Approximate near neighbors for general symmetric norms”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, pp. 902–913 (cit. on pp. [3328](#), [3330](#), [3331](#)).
- Alexandr Andoni and Ilya Razenshteyn (2015). “Optimal data-dependent hashing for approximate near neighbors”. In: *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. ACM, pp. 793–801 (cit. on pp. [3316](#), [3317](#)).
- (2016). “Tight Lower Bounds for Data-Dependent Locality-Sensitive Hashing”. In: *Proceedings of the 32nd International Symposium on Computational Geometry, SoCG 2016, June 14-18, 2016, Boston, MA, USA*, 9:1–9:11 (cit. on p. [3320](#)).
- Sunil Arya and David M Mount (1993). “Approximate Nearest Neighbor Queries in Fixed Dimensions.” In: *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. Vol. 93, pp. 271–280 (cit. on pp. [3306](#), [3307](#)).
- Sunil Arya, David M Mount, Nathan S Netanyahu, Ruth Silverman, and Angela Y Wu (1998). “An optimal algorithm for approximate nearest neighbor searching fixed dimensions”. *Journal of the ACM (JACM)* 45.6, pp. 891–923 (cit. on p. [3307](#)).
- Keith Ball (1997). “An elementary introduction to modern convex geometry”. *Flavors of geometry* 31, pp. 1–58 (cit. on pp. [3325](#), [3326](#)).
- Yair Bartal and Lee-Ad Gottlieb (2015). “Approximate Nearest Neighbor Search for ℓ_p -Spaces ($2 < p < \infty$) via Embeddings”. arXiv: [1512.01775](#) (cit. on pp. [3329](#), [3331](#), [3332](#)).
- Marshall Bern (1993). “Approximate closest-point queries in high dimensions”. *Information Processing Letters* 45.2, pp. 95–99 (cit. on pp. [3306](#), [3307](#)).

- Jean Bourgain (1985). “On Lipschitz embedding of finite metric spaces in Hilbert space”. *Israel Journal of Mathematics* 52.1, pp. 46–52 (cit. on p. 3326).
- Bo Brinkman and Moses Charikar (2005). “On the impossibility of dimension reduction in ℓ_1 ”. *Journal of the ACM (JACM)* 52.5, pp. 766–788 (cit. on p. 3310).
- Andrei Z Broder (1997). “On the resemblance and containment of documents”. In: *Compression and Complexity of Sequences 1997. Proceedings*. IEEE, pp. 21–29 (cit. on p. 3313).
- Andrei Z Broder, Steven C Glassman, Mark S Manasse, and Geoffrey Zweig (1997). “Syntactic clustering of the web”. *Computer Networks and ISDN Systems* 29.8-13, pp. 1157–1166 (cit. on p. 3313).
- Constantin Carathéodory (1911). “Über den Variabilitätsbereich der FourierÖschen Konstanten von positiven harmonischen Funktionen”. *Rendiconti Del Circolo Matematico di Palermo (1884-1940)* 32.1, pp. 193–217 (cit. on p. 3306).
- Timothy M Chan (1998). “Approximate nearest neighbor queries revisited”. *Discrete & Computational Geometry* 20.3, pp. 359–373 (cit. on p. 3307).
- Moses S Charikar (2002). “Similarity estimation techniques from rounding algorithms”. In: *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*. ACM, pp. 380–388 (cit. on pp. 3313, 3326).
- Kenneth L Clarkson (1988). “A randomized algorithm for closest-point queries”. *SIAM Journal on Computing* 17.4, pp. 830–847 (cit. on p. 3306).
- (1994). “An algorithm for approximate closest-point queries”. In: *Proceedings of the tenth annual symposium on Computational geometry*. ACM, pp. 160–164 (cit. on p. 3307).
- (2006). “Nearest-neighbor searching and metric space dimensions”. *Nearest-neighbor methods for learning and vision: theory and practice*, pp. 15–59 (cit. on p. 3307).
- Don Coppersmith (1982). “Rapid multiplication of rectangular matrices”. *SIAM Journal on Computing* 11.3, pp. 467–471 (cit. on p. 3323).
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein (2001). *Introduction to Algorithms*. 2nd. MIT Press (cit. on p. 3309).
- Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós (2010). “A sparse johnson: Lindenstrauss transform”. In: *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, pp. 341–350 (cit. on p. 3310).
- Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni (2004). “Locality-sensitive hashing scheme based on p-stable distributions”. In: *Proceedings of the twentieth annual symposium on Computational geometry*. ACM, pp. 253–262 (cit. on pp. 3312, 3313).
- Martin Farach-Colton and Piotr Indyk (1999). “Approximate nearest neighbor algorithms for Hausdorff metrics via embeddings”. In: *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 171–179 (cit. on p. 3326).

- Charles Fefferman and Bo'az Klartag (2009). “Fitting a C^m -Smooth Function to Data I”. *Annals of Mathematics*, pp. 315–346 (cit. on p. 3308).
- Tadeusz Figiel, Joram Lindenstrauss, and Vitali D Milman (1977). “The dimension of almost spherical sections of convex bodies”. *Acta Mathematica* 139.1, pp. 53–94 (cit. on p. 3325).
- M Maurice Fréchet (1906). “Sur quelques points du calcul fonctionnel”. *Rendiconti del Circolo Matematico di Palermo (1884-1940)* 22.1, pp. 1–72 (cit. on p. 3325).
- Venkatesan Guruswami, Atri Rudra, and Madhu Sudan (2014). “Essential coding theory” (cit. on p. 3315).
- Venkatesan Guruswami, Christopher Umans, and Salil Vadhan (2009). “Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes”. *Journal of the ACM (JACM)* 56.4, p. 20 (cit. on pp. 3314, 3315).
- Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani (2012). “Approximate Nearest Neighbor: Towards Removing the Curse of Dimensionality.” *Theory of computing* 8.1, pp. 321–350 (cit. on pp. 3307, 3309).
- Shlomo Hoory, Nathan Linial, and Avi Wigderson (2006). “Expander graphs and their applications”. *Bulletin of the American Mathematical Society* 43.4, pp. 439–561 (cit. on p. 3331).
- Piotr Indyk (2000a). “Dimensionality reduction techniques for proximity problems”. In: *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, pp. 371–378 (cit. on pp. 3314, 3315).
- (2000b). “High-dimensional computational geometry”. PhD thesis. stanford university (cit. on p. 3319).
- (2001). “On approximate nearest neighbors under ℓ_∞ norm”. *Journal of Computer and System Sciences* 63.4, pp. 627–638 (cit. on p. 3321).
- (2002). “Approximate nearest neighbor algorithms for Fréchet distance via product metrics”. In: *Proceedings of the eighteenth annual symposium on Computational geometry*. ACM, pp. 102–106 (cit. on pp. 3329, 3330).
- (2004). “Approximate nearest neighbor under edit distance via product metrics”. In: *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, pp. 646–650 (cit. on p. 3330).
- Piotr Indyk and Rajeev Motwani (1998). “Approximate nearest neighbors: towards removing the curse of dimensionality”. In: *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. ACM, pp. 604–613 (cit. on pp. 3309, 3311–3313, 3319).
- Piotr Indyk and Nitin Thaper (2003). “Fast image retrieval via embeddings”. *Workshop on Statistical and Computational Theories of Vision* (cit. on p. 3326).
- Thathachar S Jayram and David P Woodruff (2013). “Optimal bounds for Johnson-Lindenstrauss transforms and streaming problems with subconstant error”. *ACM Transactions on Algorithms (TALG)* 9.3, p. 26 (cit. on p. 3310).

- Fritz John (1948). “Extremum Problems with Inequalities as Subsidiary Conditions”. In: *Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948*. Interscience Publishers, Inc., New York, N. Y., pp. 187–204 (cit. on p. 3326).
- William B Johnson and Joram Lindenstrauss (1984). “Extensions of Lipschitz mappings into a Hilbert space”. *Contemporary mathematics* 26.189–206, p. 1 (cit. on p. 3309).
- William B Johnson and Assaf Naor (2009). “The Johnson-Lindenstrauss lemma almost characterizes Hilbert space, but not quite”. In: *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 885–891 (cit. on p. 3310).
- William B Johnson and Gideon Schechtman (1982). “Embedding ℓ_p^m into ℓ_1^n ”. *Acta Mathematica* 149.1, pp. 71–85 (cit. on p. 3326).
- Daniel M. Kane and Jelani Nelson (2014). “Sparsifier Johnson-Lindenstrauss Transforms”. *J. ACM* 61.1, p. 4 (cit. on p. 3310).
- Michael Kapralov (2015). “Smooth tradeoffs between insert and query complexity in nearest neighbor search”. In: *Proceedings of the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. ACM, pp. 329–342 (cit. on p. 3319).
- Michael Kapralov and Rina Panigrahy (2012). “NNS lower bounds via metric expansion for $l(\infty)$ and EMD”. In: *International Colloquium on Automata, Languages, and Programming*. Springer, pp. 545–556 (cit. on p. 3321).
- Matti Karppa, Petteri Kaski, and Jukka Kohonen (2016). “A faster subquadratic algorithm for finding outlier correlations”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 1288–1305 (cit. on pp. 3322, 3324).
- Jon M Kleinberg (1997). “Two algorithms for nearest-neighbor search in high dimensions”. In: *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*. ACM, pp. 599–608 (cit. on p. 3307).
- Felix Kraemer and Rachel Ward (2011). “New and Improved Johnson-Lindenstrauss Embeddings via the Restricted Isometry Property”. *SIAM J. Math. Analysis* 43.3, pp. 1269–1281 (cit. on p. 3310).
- Casimir Kuratowski (1935). “Quelques problèmes concernant les espaces métriques non-séparables”. *Fundamenta Mathematicae* 25.1, pp. 534–545 (cit. on p. 3325).
- Eyal Kushilevitz, Rafail Ostrovsky, and Yuval Rabani (2000). “Efficient search for approximate nearest neighbor in high dimensional spaces”. *SIAM Journal on Computing* 30.2, pp. 457–474 (cit. on pp. 3310, 3314, 3319, 3323).
- Nathan Linial, Eran London, and Yuri Rabinovich (1995). “The geometry of graphs and some of its algorithmic applications”. *Combinatorica* 15.2, pp. 215–245 (cit. on p. 3326).
- Richard J Lipton and Robert E Tarjan (1980). “Applications of a Planar Separator Theorem.” *SIAM Journal on Computing* 9.3, pp. 615–627 (cit. on p. 3306).

- Sepideh Mahabadi (2014). “Approximate nearest line search in high dimensions”. In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 337–354 (cit. on p. 3307).
- Jiří Matoušek (1997). “On embedding expanders into ℓ_p spaces”. *Israel Journal of Mathematics* 102.1, pp. 189–197 (cit. on p. 3325).
- Stefan Meiser (1993). “Point location in arrangements of hyperplanes”. *Information and Computation* 106.2, pp. 286–303 (cit. on p. 3306).
- Peter Bro Miltersen (1999). “Cell probe complexity—a survey”. In: *Proceedings of the 19th Conference on the Foundations of Software Technology and Theoretical Computer Science, Advances in Data Structures Workshop*, p. 2 (cit. on p. 3308).
- Marvin Minsky and Seymour A Papert (1969). *Perceptrons: An introduction to computational geometry*. MIT press (cit. on p. 3306).
- Rajeev Motwani, Assaf Naor, and Rina Panigrahy (2007). “Lower bounds on locality sensitive hashing”. *SIAM Journal on Discrete Mathematics* 21.4, pp. 930–935 (cit. on p. 3313).
- Assaf Naor (2017). “A Spectral Gap Precludes Low-Dimensional Embeddings”. In: *Proceedings of the 33rd International Symposium on Computational Geometry, SoCG 2017* (cit. on p. 3331).
- Assaf Naor and Yuval Rabani (2006). “On Approximate Nearest Neighbor Search in ℓ_p , $p > 2$ ”. Manuscript (cit. on pp. 3329, 3331).
- Assaf Naor and Gideon Schechtman (2007). “Planar Earthmover is not in L_1 ”. *SIAM Journal on Computing* 37.3, pp. 804–826 (cit. on p. 3326).
- Moni Naor, Leonard J Schulman, and Aravind Srinivasan (1995). “Splitters and near-optimal derandomization”. In: *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 182–191 (cit. on p. 3316).
- Jelani Nelson, Eric Price, and Mary Wootters (2014). “[New constructions of RIP matrices with fast multiplication and fewer rows](#)”. In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pp. 1515–1528 (cit. on p. 3310).
- Huy L. Nguyễn (2014). “[Algorithms for High Dimensional Data](#)”. Princeton University Thesis (cit. on p. 3327).
- Ryan O’Donnell, Yi Wu, and Yuan Zhou (2014). “Optimal lower bounds for locality-sensitive hashing (except when q is tiny)”. *ACM Transactions on Computation Theory (TOCT)* 6.1, p. 5 (cit. on p. 3313).
- Rafail Ostrovsky and Yuval Rabani (2007). “Low distortion embeddings for edit distance”. *Journal of the ACM (JACM)* 54.5, p. 23 (cit. on p. 3326).
- Rasmus Pagh (2016). “Locality-sensitive hashing without false negatives”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, pp. 1–9 (cit. on p. 3316).

- Rina Panigrahy (2006). “Entropy based nearest neighbor search in high dimensions”. In: *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. SIAM, pp. 1186–1195 (cit. on p. 3319).
- Franco P Preparata and Michael Ian Shamos (1985). “Introduction”. In: *Computational Geometry*. Springer, pp. 1–35 (cit. on p. 3306).
- Ilya Razenshteyn (2017). “High-Dimensional Similarity Search and Sketching: Algorithms and Hardness”. PhD thesis. Massachusetts Institute of Technology (cit. on p. 3317).
- Isaac J. Schoenberg (1937). “On certain metric spaces arising from Euclidean spaces by a change of metric and their imbedding in Hilbert space”. *Annals of mathematics*, pp. 787–793 (cit. on p. 3326).
- (Mar. 1942). “Positive definite functions on spheres”. *Duke Math. J.* 9.1, pp. 96–108 (cit. on p. 3323).
- Gregory Shakhnarovich, Trevor Darrell, and Piotr Indyk (2006). *Nearest-Neighbor Methods in Learning and Vision: Theory and Practice (Neural Information Processing)*. The MIT press (cit. on p. 3305).
- Volker Strassen (1969). “Gaussian elimination is not optimal”. *Numerische mathematik* 13.4, pp. 354–356 (cit. on p. 3323).
- Narayanan Sundaram, Aizana Turmukhmetova, Nadathur Satish, Todd Mostak, Piotr Indyk, Samuel Madden, and Pradeep Dubey (2013). “Streaming similarity search over one billion tweets using parallel locality-sensitive hashing”. *Proceedings of the VLDB Endowment* 6.14, pp. 1930–1941 (cit. on p. 3306).
- Gregory Valiant (2015). “Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem”. *Journal of the ACM (JACM)* 62.2, p. 13 (cit. on pp. 3322–3324).
- Virginia Vassilevska Williams (2018). “On some fine-grained questions in algorithms and complexity”. In: *Proc. of the International Congress of Mathematicians, Rio de Janeiro 2018*. Ed. by Boyan Sirakov, Paulo Ney de Souza, and Marcelo Viana. In this volume. World Scientific, pp. 3441–3482 (cit. on p. 3306).
- Jingdong Wang, Heng Tao Shen, Jingkuan Song, and Jianqiu Ji (2014). “Hashing for similarity search: A survey”. arXiv: 1408.2927 (cit. on p. 3316).
- Jun Wang, Wei Liu, Sanjiv Kumar, and Shih-Fu Chang (2016). “Learning to hash for indexing big data: a survey”. *Proceedings of the IEEE* 104.1, pp. 34–57 (cit. on p. 3316).
- Ryan Williams (2005). “A new algorithm for optimal 2-constraint satisfaction and its implications”. *Theoretical Computer Science* 348.2-3, pp. 357–365 (cit. on p. 3306).
- Virginia Vassilevska Williams (2012). “Multiplying matrices faster than Coppersmith-Winograd”. In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, pp. 887–898 (cit. on p. 3323).

Przemysław Wojtaszczyk (1996). *Banach spaces for analysts*. Vol. 25. Cambridge University Press (cit. on p. 3325).

Received 2017-12-01.

ALEXANDR ANDONI
andoni@cs.columbia.edu

PIOTR INDYK
indyk@mit.edu

ILYA RAZENSHTEYN
ilyaraz@mit.edu

