

Tight Lower Bounds for Data-Dependent Locality-Sensitive Hashing

Alexandr Andoni*
Columbia University

Ilya Razenshteyn
CSAIL MIT

July 8, 2015

Abstract

We prove a tight lower bound for the exponent ρ for data-dependent Locality-Sensitive Hashing schemes, recently used to design efficient solutions for the c -approximate nearest neighbor search. In particular, our lower bound matches the bound of $\rho \leq \frac{1}{2c-1} + o(1)$ for the ℓ_1 space, obtained via the recent algorithm from [Andoni–Razenshteyn, STOC’15].

In recent years it emerged that data-dependent hashing is strictly superior to the classical Locality-Sensitive Hashing, when the hash function is data-*independent*. In the latter setting, the best exponent has been already known: for the ℓ_1 space, the tight bound is $\rho = 1/c$, with the upper bound from [Indyk–Motwani, STOC’98] and the matching lower bound from [O’Donnell–Wu–Zhou, ITCS’11].

We prove that, even if the hashing is data-dependent, it must hold that $\rho \geq \frac{1}{2c-1} - o(1)$. To prove the result, we need to formalize the exact notion of data-dependent hashing that also captures the complexity of the hash functions (in addition to their collision properties). Without restricting such complexity, we would allow for obviously infeasible solutions such as the Voronoi diagram of a dataset. To preclude such solutions, we assume that collision probabilities are sufficiently high, as well as that the hash functions have a low description complexity. Both conditions are satisfied by all the known algorithmic results.

*Work done in part while the author was at the Simons Institute for the Theory of Computing, Berkeley University.

1 Introduction

We study lower bounds for the high-dimensional nearest neighbor search problem, which is a problem of major importance in several areas, such as databases, data mining, information retrieval, computer vision, computational geometry, signal processing, etc. This problem suffers from the “curse of dimensionality” phenomenon: either space or query time are exponential in the dimension d . To escape this curse, researchers proposed *approximation* algorithms for the problem. In the (c, r) -approximate near neighbor problem, the data structure may return any data point whose distance from the query is at most cr , for an approximation factor $c > 1$ (provided that there exists a data point within distance r from the query). Many approximation algorithms are known for this problem: e.g., see surveys [Sam06, AI08, And09, WSSJ14].

An influential algorithmic technique for the approximate near neighbor search (ANN) is the *Locality Sensitive Hashing* (LSH) [IM98, HPIM12]. The main idea is to hash the points so that the probability of collision is much higher for points that are close to each other (at distance $\leq r$) than for those which are far apart (at distance $> cr$). Given such hash functions, one can retrieve near neighbors by hashing the query point and retrieving elements stored in buckets containing that point. If the probability of collision is at least p_1 for the close points and at most p_2 for the far points, the algorithm solves the (c, r) -ANN using essentially $O(n^{1+\rho}/p_1)$ extra space and $O(dn^\rho/p_1)$ query time, where $\rho = \frac{\log(1/p_1)}{\log(1/p_2)}$ [HPIM12]. The value of the exponent ρ thus determines the “quality” of the LSH families used.

Consequently, a lot of work focused on understanding the best possible value ρ for LSH, including the sequence of upper bounds [IM98, DIIM04, AI06] and lower bounds [MNP07, OWZ11]. Overall, they established the precise bounds for the best value of ρ : for ℓ_1 the tight bound is $\rho = \frac{1}{c} \pm o(1)$. In general, for ℓ_p , where $1 \leq p \leq 2$, the tight bound is $\rho = \frac{1}{c^p} \pm o(1)$.

Surprisingly, it turns out there exist *more efficient* ANN data structures, which step outside the LSH framework. Specifically, [AINR14, AR15] design algorithms using the concept of *data-dependent hashing*, which is a randomized hash family that itself adapts to the actual given dataset. In particular, the result of [AR15] obtains an exponent $\rho = \frac{1}{2c^p-1} + o(1)$ for the ℓ_p space, thus improving upon the best possible LSH exponent essentially by a factor of 2 for both ℓ_1 and ℓ_2 spaces.

Our result. Here we prove that the exponent $\rho = \frac{1}{2c^p-1}$ from [AR15] is essentially optimal even for data-dependent hashing, and cannot be improved upon. Stating the precise theorem requires introducing the precise model for the lower bound, which we accomplish below. For now, we state our main theorem informally:

Theorem 1.1 (Main, informal). *Any data-dependent hashing scheme for ℓ_1 that achieves probabilities p_1 and p_2 must satisfy*

$$\rho \geq \frac{1}{2c-1} - o(1),$$

as long as p_1, p_2 are sufficiently high, and the description complexity of the hash functions is sufficiently small.

An immediate consequence is that $\rho \geq \frac{1}{2c^p-1} - o(1)$ for all ℓ_p with $1 \leq p \leq 2$, using the embedding from [LLR95].

1.1 Lower Bound Model

To state the precise theorem, we need to describe formally what is data-dependent hashing. First, we state the definition of (data-independent) LSH, as well as define LSH sensitive for a fixed dataset P .

Definition 1.2 (Locality-Sensitive Hashing). *We say that a hash family \mathcal{H} over $\{0, 1\}^d$ is (r_1, r_2, p_1, p_2) -sensitive, if for every $u, v \in \{0, 1\}^d$ one has:*

- if $\|u - v\|_1 \leq r_1$, then $\Pr_{h \sim \mathcal{H}}[h(u) = h(v)] \geq p_1$;
- if $\|u - v\|_1 > r_2$, then $\Pr_{h \sim \mathcal{H}}[h(u) = h(v)] \leq p_2$.

We now refine the notion of data-independent LSH, where we require the distribution to work only for a particular dataset P .

Definition 1.3 (LSH for a dataset P). *A hash family \mathcal{H} over $\{0, 1\}^d$ is said to be (r_1, r_2, p_1, p_2) -sensitive for a dataset $P \subseteq \{0, 1\}^d$, if:*

- for every $v \in \{0, 1\}^d$ and every $u \in P$ with $\|u - v\|_1 \leq r_1$ one has

$$\Pr_{h \sim \mathcal{H}}[h(u) = h(v)] \geq p_1;$$

- $\Pr_{\substack{h \sim \mathcal{H} \\ u, v \sim P}}[h(u) = h(v) \text{ and } \|u - v\|_1 > r_2] \leq p_2$.

Note that the second definition is less stringent than the first one: in fact, if all the points in a dataset are at distance more than r_2 from each other, then an LSH family \mathcal{H} is also LSH for P , but not necessarily vice versa! Furthermore, in the second definition, we require the p_2 property (2nd property) to be only *on average* (in contrast to *every* point as in the first definition). This aspect means that, while Definition 1.3 is certainly necessary for an ANN data structure, it is not obviously *sufficient*. Indeed, the algorithm from [AR15] requires proving additional properties of their partitioning scheme, and in particular analyzes *triples* of points. Since here we focus on lower bound, this aspect will not be important.

We are now ready to introduce data-dependent hashing.

Data-dependent hashing. Intuitively, a data-dependent hashing scheme is one where we can pick the family \mathcal{H} as a function of P , and thus be sensitive for P . An obvious solution would hence be to choose \mathcal{H} to consist of a single hash function h which is just the Voronoi diagram of the dataset P : it will be $(r, cr, 1, 0)$ -sensitive for P and hence $\rho = 0$. However this does not yield a good ANN solution since *evaluating* such a hash function on a query point q is as hard as the original ANN question!

Hence, ideally, our lower bound model would require that the hash function is *computationally efficient* to evaluate. We do not know how to formulate such a condition which

would not make the question as hard as circuit lower bounds or high cell-probe lower bounds, which would be well beyond the scope of this paper.

Instead, we introduce two conditions on the hash functions, which serve as a proxy for the computational complexity of the hash function. The formal statement of the main theorem (Theorem 1.1) will require these two technical conditions. We note that both of these conditions are satisfied by *all* data-dependent (or data-independent) hashing schemes that are currently known; we will discuss these conditions further below.

Theorem 1.4 (Main theorem, full). *Fix the approximation $c > 1$ to be a constant. Suppose the dataset P lives in the Hamming space $\{0, 1\}^d$, where the dataset size $n = |P|$ is such that $n \leq 2^{d/3}$ and $n = \omega(1)$. There exist distance thresholds r and $(c - o(1))r$ with the following property.*

Let $p_1, p_2 > 0$ be collision probabilities from the definition of LSH. Also, let $T \geq 1$ be the number of distinct hash functions h . Assume the following two conditions hold:

- $n^{-o(1)} \leq p_1, p_2 \leq 1/2$,
- $|T| \leq 2^{n^{1-\Omega(1)}}$.

A scheme is data-dependent if there exist T hash functions $\{h_i\}_{1 \leq i \leq T}$ over $\{0, 1\}^d$ such that for every n -point dataset P there exists a distribution \mathcal{H}_P over h_i 's such that the corresponding hash family is $(r, (c - o(1))r, p_1, p_2)$ -sensitive for P . For any such data-dependent scheme it must hold that

$$\rho = \frac{\log 1/p_1}{\log 1/p_2} \geq \frac{1}{2c - 1} - o(1).$$

We describe the two conditions next. We note that without either of the two conditions, Theorem 1.4 is simply not true.

Bounds on probability p_1 and p_2 . The query time of an ANN solution based on a hashing scheme achieving probabilities p_1, p_2 is $O(n^\rho/p_1)$. Hence it is common to assume that p_1 is not too small; e.g., the LSH lower bounds in [MNP07, OWZ11]. Indeed, if we were to allow p_1 and p_2 to be arbitrary small, then we can again construct (data-independent) LSH that achieves $\rho = 0$, though it would be unusable for ANN algorithms. In particular, one can partition the space by carving out balls of diameter $cr - o(1)$, which ensures $p_1 > 0$ and $p_2 = 0$. However p_1 becomes exponentially small in d , and so does the query time as well then.¹

We leave unanswered the following question: what happens if we allow p_1 and p_2 to be as small as $n^{-\alpha}$ for a fixed constant $\alpha > 0$? (See [OWZ11] for a further discussion about this regime.)

Description complexity of hash functions. As mentioned before, without constraining the hash function complexity, there's a perfect data-dependent hash function for a dataset P —the Voronoi diagram of P —but we cannot evaluate such a hash function efficiently. To

¹We note that, in the case when $d = O(\log n)$, such solution may actually become useful for ANN as exponential in d is merely polynomial in n .

preclude such cases, we assume that the hash function has small description complexity: the number of bits to describe the hash function is $n^{1-\Omega(1)}$. Note that such a description complexity would allow a “sampled Voronoi diagram”, where we sample, say, $n^{0.9}$ points and build a Voronoi diagram of these points. Our theorem proves that, in this sampled case, our lower bound on ρ already applies.

There are two further reasons this condition is a reasonable one. First, all the ANN algorithms [IM98, DIIM04, AI06, AINR14, AR15] use hash functions with description complexity only $n^{o(1)}$. For example, to see how [AR15] fits our lower bound model, see Appendix A. Second, in all the data-dependent hashing schemes ([AINR14, AR15]), as well as the optimal data-independent ℓ_2 scheme [AI06], the *evaluation time* of the hash function is in fact comparable with the description complexity.

In fact, a similar assumption was already used in a related setting: namely, to prove a lower bound for ANN under the ℓ_∞ norm [ACP08]. This lower bound applies to decision trees, where it was assumed that each node of the tree has $n^{o(1)}$ description complexity.

1.2 Techniques and Related Work

There are two components to our lower bound, i.e., Theorem 1.4.

The first component is a lower bound for *data-independent* LSH for a *random dataset*. We show that in this case, we must have $\rho \geq \frac{1}{2c-1} - o(1)$. This is in contrast to the lower bound of [OWZ11], who achieve a higher lower bound but for the case when the (far) points are correlated. Our lower bound is closer in spirit to [MNP07], who also consider the case when the far points are random uncorrelated. In fact, this component is a strengthening of the lower bound from [MNP07], and is based crucially on a Fourier-analytic inequality proved there.

We mention that, in [Dub10], Dubiner has also considered the setting of a random dataset for a related problem—finding the closest pair in a given dataset P . Dubiner sets up a certain related “bucketing” model, in which he conjectures the lower bound, which would imply a $\rho \geq \frac{1}{2c-1}$ lower bound for data-independent LSH for a random set. Dubiner verifies the conjecture computationally and claims it is proved in a different manuscript.²

We also point out that, for the ℓ_2 case, the optimal data-independent lower bound $\rho \geq \frac{1}{2c^2-1} - o(1)$ follows from a recent work [AIK⁺15]. In fact, it shows almost exact trade-off between p_1 and p_2 (not only the lower bound on $\rho = \frac{\log(1/p_1)}{\log(1/p_2)}$). Unfortunately, the techniques there are really tailored to the Euclidean case (in particular, a powerful isoperimetric inequality of Feige and Schechtman is used [FS02]) and it is unclear how to extend it to ℓ_1 , as well as, more generally, to ℓ_p for $1 \leq p < 2$.

Our second component focuses on the data-dependent aspect of the lower bound. In particular, we prove that if there exists a data-dependent hashing scheme for a random dataset with a better ρ , then in fact there is also a data-independent such hashing scheme. To accomplish this, we consider the “empirical” average p_1 and p_2 for a random dataset, and prove it is close to the average p_1 and p_2 , for which we can deduce a lower bound from the first component.

²This manuscript does not appear to be available at the moment of writing of the present paper.

In terms of related work, we also must mention the related work of [PTW08, PTW10], who prove cell-probe lower bounds for the same problem of ANN. In particular, their work utilizes the lower bound of [MNP07] as well. Their results are however incomparable to our results: while their results are unconditional, our model allows us to prove a much higher lower bound, in particular matching the best algorithms.

2 Data-independent Lower Bound

In this section we prove the first component of Theorem 1.4. Overall we show a lower bound of $\rho \geq \frac{1}{2c-1} + o(1)$ for data-independent hash families for random datasets. Our proof is a strengthening of [MNP07]. The final statement appears as Corollary 2.7. In the second component, we will use a somewhat stronger statement, Lemma 2.6).

For $u \in \{0, 1\}^d$ and non-negative integer k define a random variable $W_k(u)$ distributed over $\{0, 1\}^d$ to be the resulting point of the standard random walk of length k that starts in u (at each step we flip a random coordinate).

We build on the following inequality from [MNP07] that is proved using Fourier analysis on $\{0, 1\}^d$.

Lemma 2.1 ([MNP07]). *For every hash function $h: \{0, 1\}^d \rightarrow \mathbb{Z}$ and every odd positive integer k one has:*

$$\Pr_{\substack{u \sim \{0, 1\}^d \\ v \sim W_k(u)}}[h(u) = h(v)] \leq \Pr_{u, v \sim \{0, 1\}^d}[h(u) = h(v)]^{\frac{\exp(2k/d) - 1}{\exp(2k/d) + 1}}.$$

The above inequality can be thought of as a lower bound on ρ already. In particular, the left-hand-side quantity is probability of collision of a point and a point generated via a random walk of length k from it. The right-hand side corresponds to collision of random independent points, which are usually at distance $d/2$. One already obtain a lower bound on ρ by considering $k = d/2c$.

To strengthen the lower bound of [MNP07], we analyze carefully the distance between the endpoints of a random walk in the hypercube. In particular the next (somewhat folklore) lemmas show that this distance is somewhat smaller than the (trivial) upper bound of k . Denote X_k the distance $\|W_k(u) - u\|_1$. Note that the distribution of X_k does not depend on a particular starting point u .

Lemma 2.2. *For every k one has*

$$\mathbb{E}[X_k] = \frac{d}{2} \cdot \left(1 - \left(1 - \frac{2}{d}\right)^k\right).$$

Proof. We have that

$$\begin{aligned} \mathbb{E}[X_k \mid X_{k-1} = t] &= \Pr[X_k = t - 1] \cdot (t - 1) + \Pr[X_k = t + 1] \cdot (t + 1) \\ &= \frac{t}{d} \cdot (t - 1) + \left(1 - \frac{t}{d}\right) \cdot (t + 1) = \left(1 - \frac{2}{d}\right) \cdot t + 1. \end{aligned}$$

Thus,

$$\mathbb{E}[X_k] = \left(1 - \frac{2}{d}\right) \cdot \mathbb{E}[X_{k-1}] + 1.$$

Since $\mathbb{E}[X_0] = 0$, we obtain that

$$\mathbb{E}[X_k] = \sum_{i=0}^{k-1} \left(1 - \frac{2}{d}\right)^i = \frac{d}{2} \cdot \left(1 - \left(1 - \frac{2}{d}\right)^k\right).$$

□

We can now prove that the value of X_k indeed concentrates well around the expectation, using concentration inequalities.

Lemma 2.3. *For every k and every $t > 0$ one has*

$$\Pr[X_k \geq \mathbb{E}[X_k] + t \cdot \sqrt{k}] \leq e^{-t^2/4}.$$

Proof. For $t \geq 1$ define Y_t to be the index of the coordinate that got flipped at time t . Obviously, we have that X_k is a (deterministic) function of Y_1, \dots, Y_k . Furthermore, changing one Y_t changes X_k by at most 2. Hence we can apply the McDiarmid's inequality to $X_k = f(Y_1, \dots, Y_k)$:

Theorem 2.4 (McDiarmid's inequality). *Let Y_1, \dots, Y_k be independent random variables and $X = f(Y_1, \dots, Y_k)$ such that changing variable Y_t only changes the value by at most c_t . Then we have that*

$$\Pr[X \geq \mathbb{E}[X] + \varepsilon] \leq \exp\left(-\frac{\varepsilon^2}{\sum_{i=1}^k c_i^2}\right).$$

Hence we obtain that

$$\Pr[X_k \geq \mathbb{E}[X_k] + \varepsilon] \leq \exp\left(-\frac{\varepsilon^2}{4k}\right).$$

Substituting $\varepsilon = t \cdot \sqrt{k}$, we get the result. □

Note that choosing $k \approx \frac{d}{2} \cdot \ln \frac{c}{c-1}$ will mean that distance $X_k = \|W_k(u) - u\|_1$ is now around $d/2c$, i.e., we can actually use random walks longer than the ones considered in [MNP07]. Indeed, from Lemma 2.2 and Lemma 2.3 one can immediately conclude the following corollary.

Corollary 2.5. *Let $c > 1$ be a fixed constant. Suppose that $\gamma = \gamma(d) > 0$ is such that $\gamma = o(1)$ as $d \rightarrow \infty$. Then, there exists $\alpha = \alpha(d)$ such that*

$$\alpha(d) = \frac{1}{2} \cdot \ln \frac{c}{c-1} - o_{c,\gamma}(1)$$

that satisfies the following:

$$\Pr\left[X_{\alpha \cdot d} > \frac{d}{2c}\right] \leq 2^{-\gamma \cdot d}.$$

We are now ready to prove the main lemma of this section, which will be used in the later section on data-dependent hashing. We need to introduce two more definitions. We define “average p_1 ” as:

$$\zeta(c, d, \alpha, h) := \Pr_{\substack{u \sim \{0,1\}^d \\ v \sim W_{\alpha \cdot d}(u)}} \left[h(u) = h(v) \mid \|u - v\|_1 \leq \frac{d}{2c} \right]$$

for $c > 1$, positive integer d , $\alpha > 0$ and a hash function $h: \{0, 1\}^d \rightarrow \mathbb{Z}$. Similarly, we define the “average p_2 ” as

$$\eta(d, \beta, h) := \Pr_{u, v \sim \{0,1\}^d} \left[h(u) = h(v), \|u - v\|_1 > \left(\frac{1}{2} - \beta \right) \cdot d \right]$$

for positive integer d , $\beta > 0$ and a hash function $h: \{0, 1\}^d \rightarrow \mathbb{Z}$.

Lemma 2.6. *Let $c > 1$ be a fixed constant. Suppose that $\gamma = \gamma(d) > 0$ is such that $\gamma = o(1)$ as $d \rightarrow \infty$. Then, there exist $\alpha = \alpha(d)$, $\beta = \beta(d)$ and $\rho = \rho(d)$ such that:*

$$\begin{aligned} \alpha &= \frac{1}{2} \cdot \ln \frac{c}{c-1} - o_{c,\gamma}(1), \\ \beta &= o_\gamma(1), \\ \rho &= \frac{1}{2c-1} - o_{c,\gamma}(1) \end{aligned}$$

as $d \rightarrow \infty$ such that, for every d and every hash function $h: \{0, 1\}^d \rightarrow \mathbb{Z}$, one has

$$\zeta(c, d, \alpha, h) \leq \eta(d, \beta, h)^\rho + 2^{-\gamma \cdot d}.$$

Proof. We use Lemma 2.5 to choose $\alpha = \frac{1}{2} \cdot \ln \frac{c}{c-1} + o_{c,\gamma}(1)$ such that $\alpha \cdot d$ is an odd integer and

$$\Pr \left[X_{\alpha \cdot d} \geq \frac{d}{2c} \right] < \frac{2^{-\gamma \cdot d}}{2}. \quad (1)$$

We can choose $\beta = o_\gamma(1)$ so that

$$\Pr_{u, v \sim \{0,1\}^d} \left[\|u - v\|_1 < \left(\frac{1}{2} - \beta \right) \cdot d \right] < \frac{2^{-\gamma \cdot d}}{2}. \quad (2)$$

(This follows from the standard Chernoff-type bounds.)

Now we apply Lemma 2.1 and get

$$\Pr_{\substack{u \sim \{0,1\}^d \\ v \sim W_{\alpha \cdot d}(u)}} [h(u) = h(v)] \leq \Pr_{u, v \sim \{0,1\}^d} [h(u) = h(v)]^\rho, \quad (3)$$

where

$$\rho = \frac{\exp(2\alpha) - 1}{\exp(2\alpha) + 1} = \frac{1}{2c-1} - o_{c,\gamma}(1).$$

Finally, we combine (3), (1) and (2), and get the desired inequality. \square

The following corollary shows how the above lemma implies a lower bound on data-independent LSH.

Corollary 2.7. *For every $c > 1$ and $\gamma = \gamma(d) > 0$ such that $\gamma = o(1)$ there exists $\beta = \beta(d) > 0$ with $\beta = o_\gamma(1)$ such that if \mathcal{H} is a data-independent $(d/(2c), (1/2 - \beta)d, p_1, p_2)$ -sensitive family, then*

$$p_1 \leq p_2^{\frac{1}{2c-1} - o_{c,\gamma}(1)} + 2^{-\gamma \cdot d}.$$

Proof. We observe that for every $\alpha, \beta > 0$:

- $p_1 \leq \mathbb{E}_{h \sim \mathcal{H}} [\zeta(c, d, \alpha, h)];$
- $p_2 \geq \mathbb{E}_{h \sim \mathcal{H}} [\eta(d, \beta, h)].$

Now we apply Lemma 2.6 together with the following application of Jensen's inequality:

$$\mathbb{E}_{h \sim \mathcal{H}} [\eta(d, \beta, h)^\rho] \leq \mathbb{E}_{h \sim \mathcal{H}} [\eta(d, \beta, h)]^\rho,$$

since $0 < \rho \leq 1$. □

3 Data-Dependent Hashing

We now prove the second component of the main Theorem 1.4 proof. In particular we show that a very good data-dependent hashing scheme would refute Lemma 2.6 from the previous section.

3.1 Empirical Probabilities of Collision

For a particular dataset P , we will be interested in *empirical* probabilities p_1, p_2 — i.e., the equivalents of ζ, μ for a given set P — defined as follows. Let $0 < \delta(d) < 1/3$ be some function. Let P be a random set of points from $\{0, 1\}^d$ of size $2^{\delta(d) \cdot d}$. The *empirical* versions of ζ and η with respect to P are:

$$\begin{aligned} \hat{\zeta}(c, d, \alpha, h, P) &:= \Pr_{\substack{u \sim P \\ v \sim W_{\alpha \cdot d}(u)}} \left[h(u) = h(v) \mid \|u - v\|_1 \leq \frac{d}{2c} \right] \\ \hat{\eta}(d, \beta, h, P) &:= \Pr_{u, v \sim P} \left[h(u) = h(v), \|u - v\|_1 > \left(\frac{1}{2} - \beta \right) \cdot d \right]. \end{aligned}$$

We now want to prove that, for a random dataset P , the empirical ζ, μ are close to the true averages. For this we will need the following auxiliary lemma.

Lemma 3.1. *Let M be an $n \times n$ symmetric matrix with entries from $[0; 1]$ and average ε . Let M' be a principal $n^\delta \times n^\delta$ submatrix of M sampled uniformly with replacement. Then, for every $\theta > 0$, the probability that the maximum of the average over M' and θ does not lie in $[1/2; 2] \cdot \max\{\varepsilon, \theta\}$ is at most $n^\delta \cdot 2^{-\Omega(\theta n^\delta)}$.*

Proof. We need the following version of Bernstein's inequality.

Lemma 3.2. *Suppose that X_1, \dots, X_n are i.i.d. random variables that are distributed over $[0; 1]$. Suppose that $\mathbb{E}[X_i] = \varepsilon$. Then, for every $0 < \theta < 1$, one has*

$$\Pr \left[\max \left\{ \frac{1}{n} \sum_{i=1}^n X_i, \theta \right\} \in \left[\frac{1}{2}; 2 \right] \cdot \max\{\varepsilon, \theta\} \right] \geq 1 - 2^{-\Omega(\theta n)}.$$

We just apply Lemma 3.2 and take union bound over the rows of M' . \square

The following two lemmas are immediate corollaries of Lemma 3.2 and Lemma 3.1, respectively.

Lemma 3.3. *For every $c > 1$, $\alpha > 0$, positive integer d , $\theta > 0$ and a hash function $h: \{0, 1\}^d \rightarrow \mathbb{Z}$, one has*

$$\Pr_P \left[\max\{\hat{\zeta}(c, d, \alpha, h, P), \theta\} \in [1/2; 2] \cdot \max\{\zeta(c, d, \alpha, h), \theta\} \right] \geq 1 - 2^{-\Omega(\theta \cdot 2^{\delta(d) \cdot d})}.$$

Lemma 3.4. *For every $\beta > 0$, positive integer d , $\theta > 0$ and a hash function $h: \{0, 1\}^d \rightarrow \mathbb{Z}$, one has*

$$\Pr_P \left[\max\{\hat{\eta}(d, \beta, h, P), \theta\} \in [1/2; 2] \cdot \max\{\eta(d, \beta, h), \theta\} \right] \geq 1 - 2^{\delta(d) \cdot d} \cdot 2^{-\Omega(\theta \cdot 2^{\delta(d) \cdot d})}.$$

3.2 Proof of Theorem 1.4

We are finally ready to complete the prove of the main result, Theorem 1.4.

Let us first assume that $p_1, p_2 = o(1)$ and then show how to handle the general case. Suppose that $|P| = 2^{\delta \cdot d}$.

Let $\{h_1, h_2, \dots, h_T\}$ be a set of hash functions. Fix $\gamma = \gamma(d) > 0$ such that $\gamma = o(1)$ and $2^{-\gamma d} \ll p_1, p_2$. Then, by Lemma 2.6, there is $\alpha = \alpha(d)$ and $\beta = \beta(d) = o_\gamma(1)$ such that for every $1 \leq i \leq T$

$$\zeta(c, d, \alpha, h_i) \leq \eta(d, \beta, h_i)^{\frac{1}{2c-1} - o_{c,\gamma}(1)} + 2^{-\gamma d}. \quad (4)$$

Let us choose $\theta > 0$ such that

$$T \cdot 2^{\delta \cdot d} \ll 2^{\theta \cdot 2^{\delta \cdot d}}$$

and $\theta \ll p_1, p_2$. If we can not do it, then $T \geq 2^{(1-o(1))\delta d} = 2^{n^{1-o(1)}}$, contradicting theorem hypothesis.

Then, from Lemma 3.3 and Lemma 3.4 we get that, with high probability over the choice of P , one has for every $1 \leq i \leq T$:

- $\max\{\hat{\zeta}(c, d, \alpha, h_i), \theta\} \in [1/2; 2] \cdot \max\{\zeta(c, d, \alpha, h_i), \theta\};$
- $\max\{\hat{\eta}(d, \beta, h_i), \theta\} \in [1/2; 2] \cdot \max\{\eta(d, \beta, h_i), \theta\}.$

Suppose these conditions hold and assume there exist a distribution \mathcal{D} over $[T]$ such that the corresponding hash family is $(d/2c, (1/2 - \beta(d))d, p_1, p_2)$ -sensitive for P . Then,

$$\begin{aligned} p_1 &\leq \mathbb{E}_{i \sim \mathcal{D}} [\hat{\zeta}(c, d, \alpha, h_i)] \leq \mathbb{E}_{i \sim \mathcal{D}} [\max\{\hat{\zeta}(c, d, \alpha, h_i), \theta\}] \leq 2 \cdot \mathbb{E}_{i \sim \mathcal{D}} [\max\{\zeta(c, d, \alpha, h_i), \theta\}] \\ &\leq 2 \cdot \mathbb{E}_{i \sim \mathcal{D}} [\zeta(c, d, \alpha, h_i)] + \theta. \end{aligned} \quad (5)$$

Similarly,

$$\begin{aligned} p_2 &\geq \mathbb{E}_{i \sim \mathcal{D}} [\hat{\eta}(d, \beta, h_i)] \geq \mathbb{E}_{i \sim \mathcal{D}} [\max\{\hat{\eta}(d, \beta, h_i), \theta\}] - \theta \geq \frac{1}{2} \cdot \mathbb{E}_{i \sim \mathcal{D}} [\max\{\eta(d, \beta, h_i), \theta\}] - \theta \\ &\geq \frac{1}{2} \cdot \mathbb{E}_{i \sim \mathcal{D}} [\eta(d, \beta, h_i)] - \theta. \end{aligned} \quad (6)$$

Averaging (4) and applying Jensen's inequality, we have

$$\mathbb{E}_{i \sim \mathcal{D}} [\zeta(c, d, \alpha, h_i)] \leq \mathbb{E}_{i \sim \mathcal{D}} [\eta(d, \beta, h_i)]^{\frac{1}{2c-1}-o(1)} + 2^{-\gamma(d) \cdot d}. \quad (7)$$

Thus, substituting (5) and (6) into (7)

$$\frac{p_1 - \theta}{2} \leq (2(p_2 + \theta))^{\frac{1}{2c-1}-o(1)} + 2^{-\gamma d},$$

which proves the theorem, since $\theta \ll p_1, p_2$, and $2^{-\gamma d} \ll p_1, p_2$, and $p_1, p_2 = o(1)$.

Now let us deal with the case $p_1, p_2 = \Omega(1)$. This can be reduced to the case $p_1, p_2 = o(1)$ by choosing a slowly-growing super constant k and replacing the set of T functions with the set of T^k tuples of length k . This replaces p_1 and p_2 with p_1^k and p_2^k , respectively. In the same time, we choose k so that T^k is still $2^{n^{1-\Omega(1)}}$. Then, we just apply the above proof.

References

- [ACP08] Alexandr Andoni, Dorian Croitoru, and Mihai Pătraşcu. Hardness of nearest neighbor under L-infinity. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 424–433, 2008.
- [AI06] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2006)*, pages 459–468, 2006.
- [AI08] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM*, 51(1):117–122, 2008.
- [AIK⁺15] Alexandr Andoni, Piotr Indyk, Michael Kapralov, Thijs Laarhoven, Ilya Razenshteyn, and Ludwig Schmidt. Practical and optimal LSH for angular distance. Unpublished manuscript, 2015.

- [AINR14] Alexandr Andoni, Piotr Indyk, Huy L. Nguyen, and Ilya Razenshteyn. Beyond locality-sensitive hashing. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '2014)*, pages 1018–1028, 2014.
- [And09] Alexandr Andoni. *Nearest Neighbor Search: the Old, the New, and the Impossible*. PhD thesis, Massachusetts Institute of Technology, 2009.
- [AR15] Alexandr Andoni and Ilya Razenshteyn. Optimal data-dependent hashing for approximate near neighbors. In *Proceedings of the ACM Symposium on the Theory of Computing (STOC '15)*, 2015.
- [DIIM04] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the 20th ACM Symposium on Computational Geometry (SoCG '2004)*, pages 253–262, 2004.
- [Dub10] Moshe Dubiner. Bucketing coding and information theory for the statistical highdimensional nearest-neighbor problem. *IEEE Transactions on Information Theory*, 56(8):4166–4179, 2010.
- [FS02] Uriel Feige and Gideon Schechtman. On the optimality of the random hyperplane rounding technique for MAX CUT. *Random Structures and Algorithms*, 20(3):403–440, 2002.
- [HPIM12] Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate nearest neighbor: towards removing the curse of dimensionality. *Theory of Computing*, 8(1):321–350, 2012.
- [IM98] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the 30th ACM Symposium on the Theory of Computing (STOC '1998)*, pages 604–613, 1998.
- [LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [MNP07] Rajeev Motwani, Assaf Naor, and Rina Panigrahy. Lower bounds on locality sensitive hashing. *SIAM Journal on Discrete Mathematics*, 21(4):930–935, 2007.
- [OWZ11] Ryan O’Donnell, Yi Wu, and Yuan Zhou. Optimal lower bounds for locality sensitive hashing (except when q is tiny). In *Proceedings of Innovations in Computer Science (ICS '2011)*, pages 275–283, 2011.
- [PTW08] Rina Panigrahy, Kunal Talwar, and Udi Wieder. A geometric approach to lower bounds for approximate near-neighbor search and partial match. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2008)*, pages 414–423, 2008.

- [PTW10] Rina Panigrahy, Kunal Talwar, and Udi Wieder. Lower bounds on near neighbor search via metric expansion. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS '2010)*, pages 805–814, 2010.
- [Sam06] Hannan Samet. *Foundations of Multidimensional and Metric Data Structures*. Elsevier, 2006.
- [WSSJ14] Jingdong Wang, Heng Tao Shen, Jingkuan Song, and Jianqiu Ji. Hashing for similarity search: A survey. *CoRR*, abs/1408.2927, 2014.

A Upper bounds are in the model

In this section, we show how the data-dependent hash family from [AR15] fits into the model of the lower bound from Section 1.1.

Let us briefly recall the hash family construction from [AR15]. For simplicity, assume that all points and queries lie on a sphere of radius $R \gg cr$. First, consider a *data-independent* hash family from [AINR14, AR15], called *Spherical LSH*. It gives a good exponent $\rho \leq \frac{1}{2c^2-1} + o(1)$ for distance thresholds r vs $\sqrt{2}R$ (the latter corresponds to a typical distance between a pair of points from the sphere). The main challenge that arises is how to handle distance thresholds r vs cr , where the latter may be much smaller than $\sqrt{2}R$. Here comes the main insight of [AR15].

We would like to process the dataset so that the distance between a typical pair of *data points* is around $\sqrt{2}R$, so to apply Spherical LSH. To accomplish this, we remove all the clusters of radius $(\sqrt{2} - \varepsilon)R$ that contain lots of points (think of $\varepsilon > 0$ being barely sub-constant). We will treat these clusters separately, and will focus on the remainder of the pointset for now. So for the remainder of the pointset, we just apply the Spherical LSH to it. This scheme turns out to satisfy the definition of the data-dependent hash family for the remaining points and for distance thresholds r vs. cr ; in particular, the hash function is sensitive for the remaining set only! The intuition is that, for any potential query point, there is only a small number of data points within distance $(\sqrt{2} - \varepsilon)R$ — otherwise, they would have formed yet another cluster, which we would have removed — and the larger distances are handled well by the Spherical LSH. Thus, *for a typical pair of data points*, Spherical LSH is sensitive for P (see Definition 1.3).

How does [AR15] deal with the clusters? The main observation is that one can enclose such a cluster in a ball of radius $(1 - \Omega(\varepsilon^2))R$, which intuitively makes our problem a little bit simpler (after we reduce the radius enough times, the problem becomes trivial). To answer a query, we query *every cluster*, as well as *one part* of the remainder (partitioned by the Spherical LSH). This can be shown to work overall, in particular, we can control the overall branching and depth of the recursion (see [AR15] for the details).

For both the clusters, as well as the parts obtained from the Spherical LSH, the algorithm recurses on the obtained point subsets. The overall partitioning scheme from [AR15] can be seen as a tree, where the root corresponds to the whole dataset, and every node either corresponds to a cluster or to an application of the Spherical LSH. One nuisance is that, in different parts of the tree the Spherical LSH partitioning obtains different p_1, p_2 (depending on R). Nonetheless, each time it holds that $p_1 \geq p_2^\rho$ for $\rho \leq \frac{1}{2c^2-1} + o(1)$. Hence, a node

terminates as a leaf once its accumulated p_2 (product of p_2 's of "Spherical LSH" nodes along the path from the root) drops below $1/n$.

We now want to argue that the above algorithm can be recast in the framework of data-dependent hashing as per Definition 1.3. We consider a subtree of the overall tree that contains the root and is defined as follows. Fix a parameter $l = n^{-o(1)}$ (it is essentially the target p_2 of the partition). We perform DFS of the tree and cut the tree at any "Spherical LSH" node where the cumulative p_2 drops below l . This subtree gives a partial partition of the dataset P as follows: for "Spherical LSH" nodes we just apply the corresponding partition, and for cluster nodes we "carve" them in the random order. It turns out that if we choose $l = n^{-o(1)}$ carefully, the partition will satisfy Definition 1.3 and the preconditions of Theorem 1.4. In particular, the description complexity of the resulting hash function is $n^{o(1)}$.

Let us emphasize that, while Definition 1.3 is certainly necessary for an ANN data structure based on data-dependent hashing, it is not *sufficient*. In fact, [AR15] prove additional properties of the above partitioning scheme, essentially because the " p_2 property" is on average one (thus, we end up having to understand this partitioning scheme treats *triples* of points).