Optimal Las Vegas Locality Sensitive Data Structures

Full Version

Thomas Dybdahl Ahle
IT University of Copenhagen

June 27 2018

Abstract

We show that approximate similarity (near neighbour) search can be solved in high dimensions with performance matching state of the art (data independent) Locality Sensitive Hashing, but with a guarantee of no false negatives. Specifically, we give two data structures for common problems.

For c-approximate near neighbour in Hamming space we get query time $dn^{1/c+o(1)}$ and space $dn^{1+1/c+o(1)}$ matching that of [Indyk and Motwani, 1998] and answering a long standing open question from [Indyk, 2000a] and [Pagh, 2016] in the affirmative. By means of a new deterministic reduction from ℓ_1 to Hamming we also solve ℓ_1 and ℓ_2 with query time $d^2n^{1/c+o(1)}$ and space $d^2n^{1+1/c+o(1)}$.

For (s_1, s_2)-approximate Jaccard similarity we get query time $dn^{\rho+o(1)}$ and space $dn^{1+\rho+o(1)}$, $\rho = \log \frac{1+s_1}{2s_1} \Big/ \log \frac{1+s_2}{2s_2}$, when sets have equal size, matching the performance of [Pagh and Christiani, 2017].

The algorithms are based on space partitions, as with classic LSH, but we construct these using a combination of brute force, tensoring, perfect hashing and splitter functions à la [Naor et al., 1995]. We also show a new dimensionality reduction lemma with 1-sided error.

1 Introduction

Locality Sensitive Hashing has been a leading approach to high dimensional similarity search (nearest neighbour search) data structures for the last twenty years. Intense research [Indyk and Motwani, 1998, Gionis et al., 1999, Kushilevitz et al., 2000, Indyk, 2000b, Indyk, 2001, Charikar, 2002, Datar et al., 2004, Lv et al., 2007, Panigrahy, 2006, Andoni and Indyk, 2006, Andoni et al., 2014, Andoni et al., 2017a, Becker et al., 2016, Ahle et al., 2017, Aumüller et al., 2017] has applied the concept of space partitioning to many different problems and similarity spaces. These data structures are popular in particular because of their ability to overcome the 'curse of dimensionality' and conditional lower bounds by [Williams, 2005], and give sub-linear query time on worst case instances. They achieve this by being approximate and

Monte Carlo, meaning they may return a point that is slightly further away than the nearest, and with a small probability they may completely fail to return any nearby point.

Definition 1 ((c,r)-Approximate Near Neighbour)**.** *Given a set P of n data points in a metric space (X, dist), build a data structure, such that given any $q \in X$, for which there is an $x \in P$ with $\mathrm{dist}(q,x) \leq r$, we return a $x' \in P$ with $\mathrm{dist}(q,x') \leq cr$.*

A classic problem in high dimensional geometry has been whether data structures existed for (c,r)-Approximate Near Neighbour with Las Vegas guarantees, and performance matching that of Locality Sensitive Hashing. That is, whether we could guarantee that a query will always return an approximate near neighbour, if a near neighbour exists; or simply, if we could rule out false negatives? The problem has seen practical importance as well as theoretical. There is in general no way of verifying that an LSH algorithm is correct when it says 'no near neighbours' - other than iterating over every point in the set, in which case the data structure is entirely pointless. This means LSH algorithms can't be used for many critical applications, such as finger print data bases. Even more applied, it has been observed that tuning the error probability parameter is hard to do well, when implementing LSH [Gionis et al., 1999, Arya et al., 1998]. A Las Vegas data structure entirely removes this problem. Different authors have described the problem with different names, such as 'Las Vegas' [Indyk, 2000a], 'Have no false negatives' [Goswami et al., 2017, Pagh, 2016], 'Have total recall' [Pham and Pagh, 2016], 'Are exact' [Arasu et al., 2006] and 'Are explicit' [Karppa et al., 2016].

Recent years have shown serious progress towards finally solving the problem. In particular [Pagh, 2016] showed that the problem in Hamming space admits a Las Vegas algorithm with query time $dn^{1.38/c+o(1)}$, matching the $dn^{1/c}$ data structure of [Indyk and Motwani, 1998] up to a constant factor in the exponent. In this paper we give an algorithm in the Locality Sensitive Filter framework [Becker et al., 2016, Christiani, 2017], which not only removes the factor 1.38, but improves to $dn^{1/(2c-1)+o(1)}$ in the case $cr \approx d/2$, matching the algorithms of [Andoni et al., 2015] for Hamming space.

We would like to find an approach to Las Vegas LSH that generalizes to the many different situations where LSH is useful. Towards that goal, we present as second algorithm for the approximate similarity search problem under Braun-Blanquet similarity, which is defined for sets $x, y \subseteq [d]$ as $\mathrm{sim}(x,y) = |x \cap y| / \max(|x|,|y|)$. We refer to the following problem definition:

Definition 2 (Approximate similarity search)**.** *Let $P \subseteq \mathcal{P}([d])$ be a set of $|P| = n$ subsets of $[d]$; (here $\mathcal{P}(X)$ denotes the powerset of X.) let $\mathrm{sim} : \mathcal{P}([d]) \times \mathcal{P}([d]) \to [0,1]$ be a similarity measure. For given $s_1, s_2 \in [0,1]$, $s_1 > s_2$, a solution to the "(s_1,s_2)-similarity search problem under sim" is a data structure that supports the following query operation: on input $q \subseteq [d]$, for which there exists a set $x \in P$ with $\mathrm{sim}(x,q) \geq s_1$, return $x' \in P$ with $\mathrm{sim}(x',q) > s_2$.*

The problem has traditionally been solved using the Min-Hash LSH [Broder et al., 1997, Broder, 1997], which combined with the results of Indyk and Motwani [Indyk and Motwani, 1998] gives a data structure with query time dn^ρ and space $dn^{1+\rho}$ for $\rho = \log s_1 / \log s_2$. Recently it was shown by [Pagh and Christiani, 2017] that this could be improved for vectors of equal weight to $\rho = \log \frac{2s_1}{1+s_1} \big/ \log \frac{2s_2}{1+s_2}$. We show that it is possible to achieve this recent result with a data structure that has no false negatives.

1.1 Summary of Contributions

We present the first Las Vegas algorithm for approximate near neighbour search, which gives sub-linear query time for any approximation factor $c > 1$. This solves a long standing open question from [Indyk, 2000a] and [Pagh, 2016]. In particular we get the following two theorems:

Theorem 1. *Let $X = \{0,1\}^d$ be the Hamming space with metric $\operatorname{dist}(x,y) = \|x \oplus y\| \in [0,d]$ where $\oplus$ is "xor" or addition in $\mathbb{Z}_2$. For every choice of $0 < r$, $1 < c$ and $cr \le d/2$, we can solve the (c,r)-approximate near neighbour problem in Hamming space with query time dn^ρ and space usage $dn + n^{1+\rho}$ where $\rho = 1/c + \hat{O}((\log n)^{-1/4})$.*

Note: $\hat{O}$ hides $\log \log n$ factors.

Corollary 1. *When $r/d = \Omega((\log n)^{-1/6})$, we get the improved exponent $\rho = \frac{1-cr/d}{c(1-r/d)} + \hat{O}((\log n)^{-1/3} d/r)$.*

This improves upon theorem 1 when r/d is constant (or slightly sub-constant), including in the important "random case", when $r/d = 1/(2c)$ where we get $\rho = 1/(2c-1) + o(1)$.

Theorem 2. *Let sim be the Braun-Blanquet similarity $\operatorname{sim}(x,y) = |x \cap y| / \max(|x|,|y|)$. For every choice of constants $0 < s_2 < s_1 < 1$, we can solve the (s_1, s_2)-similarity problem over sim with query time dn^ρ and space usage $dn + n^{1+\rho}$ where $\rho = \log s_1 / \log s_2 + \hat{O}((\log n)^{-1/2})$.*

For sets of fixed size w, the dn terms above can be improved to wn. It is also possible to let s_1 and s_2 depend on n with some more work.

The first result matches the lower bounds by [O'Donnell et al., 2014] for "data independent" LSH data structures for Hamming distance and improves upon [Pagh, 2016] by a factor of $\log 4 > 1.38$ in the exponent. By deterministic reductions from ℓ_2 to ℓ_1 [Indyk, 2007] and ℓ_1 to hamming (appendix 6.1), this also gives the best currently known Las Vegas data structures for ℓ_1 and ℓ_2 in $\mathbb{R}^d$. The second result matches the corresponding lower bounds by [Pagh and Christiani, 2017] for Braun-Blanquet similarity and, by reduction, Jaccard similarity. See table 1 for more comparisons.

Detaching the data structures from our constructions, we give the first explicit constructions of large Turán Systems [Sidorenko, 1995], which are families $\mathcal{T}$ of k-subsets of $[n]$, such that any r-subset of $[n]$ is contained in a set in $\mathcal{T}$. Lemma 5 constructs (n, k, r)-Turán Systems using $(n/k)^r e^{\chi}$ sets, where $\chi = O(\sqrt{r}\log r + \log k + \log\log n)$. For small values of k this is sharp with the lower bound of $\binom{n}{r}/\binom{k}{r}$, and our systems can be efficiently decoded, which is likely to have other algorithmic applications.

1.2 Background and Related Work

The arguably most successful technique for similarity search in high dimensions is Locality-Sensitive Hashing (LSH), introduced in 1998 by [Indyk and Motwani, 1998, Har-Peled et al., 2012]. The idea is to make a random space partition in which similar points are likely to be stored in the same region, thus allowing the search space to be pruned substantially. The granularity of the space partition (the size/number of regions) is chosen to balance the expected number of points searched against keeping a (reasonably) small probability of pruning away the actual nearest point. To ensure a high probability of success (good recall) one repeats the above construction, independently at random, a small polynomial (in n) number of times.

In [Pagh, 2016, Arasu et al., 2006] it was shown that one could change the above algorithm to not do the repetitions independently. (Eliminating the error probability of an algorithm by independent repetitions, of course, takes an infinite number of repetitions.) By making correlated repetitions, it was shown possible to reach zero false negatives much faster, after only polynomially many repetitions. This means, for example, that they needed more repetitions than LSH does to get 0.99 success rate, but fewer than LSH needs for success rate $1 - 2^{-n}$.

An alternative to LSH was introduced by [Becker et al., 2016, Dubiner, 2010]. It is referred to as Locality Sensitive Filters, or LSF. While it achieves the same bounds as LSH, LSF has the advantage of giving more control to the algorithm designer for balancing different performance metrics. For example, it typically allows better results for low dimensional data, $d = O(\log n)$, and space/time trade-offs [Andoni et al., 2017a]. The idea is to sample a large number of random sections of the space. In contrast to LSH these sections are not necessarily partitions and may overlap heavily. For example, for points on the sphere S^{d-1} the sections may be defined by balls around the points of a spherical code. One issue compared to LSH is that the number of sections in LSF is very large. This means we need to impose some structure so we can efficiently find all sections containing a particular point. With LSH the space partitioning automatically provided such an algorithm, but for LSF it is common to use a kind of random product code. (An interesting alternative is [Pagh and Christiani, 2017], which uses a random branching processes.) LSF is similar to LSH in that it only approaches 100% success rate as the number of sections goes to infinity.

The work in this paper can be viewed as way of constructing correlated, efficiently

decodable filters for Hamming space and Braun-Blanquet similarity. That is, our filters guarantee that any two close points are contained in a shared section, without having an infinite number of sections. Indeed the number of sections needed is equal to that needed by random constructions for achieving constant success probability, up to $n^{o(1)}$ factors. It is not crucial that our algorithms are in the LSF framework rather than LSH. Our techniques can make correlated LSH space partitions of optimal size as well as filters. However the more general LSF framework allows for us to better show of the strength of the techniques.

One very important line of LSH/LSF research, that we don't touch upon in this paper, is that of data dependency. In the seminal papers [Andoni et al., 2014, Andoni and Razenshteyn, 2015, Andoni et al., 2017a] it was shown that the performance of space partition based data structures can be improved, even in the worst case, by considering the layout of the points in the data base. Using clustering, certain bad cases for LSH/LSF can be removed, leaving only the case of "near random" points to be considered, on which LSH works very well. It seems possible to make Las Vegas versions of these algorithms as well, since our approach gives the optimal performance in these near random cases. However one would need to find a way to derandomize the randomized clustering step used in their approach.

There is of course also a literature of deterministic and Las Vegas data structures not using LSH. As a baseline, we note that the "brute force" algorithm that stores every data point in a hash table, and given a query, $q \in \{0,1\}^d$, looks up every $\sum_{k=1}^{r} \binom{d}{k}$ point of Hamming distance most r. This requires $r \log(d/r) < \log n$ to be sub-linear, so for a typical example of $d = (\log n)^2$ and $r = d/10$ it won't be practical. In [Cole et al., 2004] this was somewhat improved to yield $n(\log n)^r$ time, but it still requires $r = O(\frac{\log n}{\log \log n})$ for queries to be sub-linear. We can also imagine storing the nearest neighbour for every point in $\{0,1\}^d$. Such an approach would give fast (constant time) queries, but the space required would be exponential in r.

In Euclidean space (ℓ_2 metric) the classical K-d tree algorithm [Bentley, 1975] is of course deterministic, but it has query time $n^{1-1/d}$, so we need $d = O(1)$ for it to be strongly sub-linear. Allowing approximation, but still deterministic, [Arya et al., 1998] found a $(\frac{d}{c-1})^d$ algorithm for $c > 1$ approximation. They thus get sublinear queries for $d = O(\frac{\log n}{\log \log n})$.

For large approximation factors [Har-Peled et al., 2012] gave a deterministic data structure with query time $O(d \log n)$, but space and preprocessing more than $n \cdot O(1/(c-1))^d$. In a different line of work, [Indyk, 2000a] gave a deterministic $(d\epsilon^{-1} \log n)^{O(1)}$ query time, fully deterministic algorithm with space usage $n^{O(1/\epsilon^6)}$ for a $3+\epsilon$ approximation.

See Table 1 for an easier comparison of the different results and spaces.

Reference	**Space**	**Exponent, search time**	**Comments**
[Bentley, 1975]	ℓ_2	$1 - 1/d$	Exact algorithm, Fully deterministic.
[Cole et al., 2004]	Hamming	$r\frac{\log\log n}{\log n}$	Sub-linear for $r < \frac{\log n}{\log\log n}$. Exact.
[Arya et al., 1998]	ℓ_2	$d\frac{\log(d/(c-1))}{\log n}$	Sub-linear for $d < \frac{\log n}{\log\log n}$.
[Har-Peled et al., 2012]	Hamming	$o(1)$	c-approximation, Fully deterministic, $(1/(c-1))^d$ space.
[Indyk, 2000a]	Hamming	$o(1)$	$(3+\epsilon)$-approximation, Fully deterministic, $n^{\Omega(1/\epsilon^6)}$ space.
[Arasu et al., 2006]	Hamming	$\approx 3/c$	The paper makes no theoretical claims on the exponent.
[Pagh, 2016]	Hamming	$1.38/c$	Exponent $1/c$ when $r = o(\log n)$ or $(\log n)/(cr) \in \mathbb{N}$.
[Pacuk et al., 2016]	ℓ_p	$O(d^{1-1/p}/c)$	Sub-linear for ℓ_2 when $c = \omega(\sqrt{d})$.
This paper	Hamming, ℓ_1, ℓ_2	$1/c$	Actual exponent is $\frac{1-cr/d}{c(1-r/d)}$ which improves to $1/(2c-1)$ for $cr \approx d/2$.
[Pagh, 2016]	Braun-Blanquet	$1.38\,\frac{1-b_1}{1-b_2}$	Via reduction to Hamming. Requires sets of equal weight.
This paper	Braun-Blanquet	$\frac{\log 1/b_1}{\log 1/b_2}$	See [Pagh and Christiani, 2017] figure 2 for a comparison with [Pagh, 2016].

Table 1: Comparison of Las Vegas algorithms for high dimensional near neighbour problems. The exponent is the value ρ, such that the data structure has query time $n^{\rho+o(1)}$. All listed algorithms, except for [Indyk, 2000a] use less than n^2 space. All algorithms give c-approximations, except for the first two, and for [Indyk, 2000a], which is a $(3+\epsilon)$-approximation.

1.3 Techniques

Our main new technique is a combination of 'splitters' as defined by [Naor et al., 1995, Alon et al., 2006], and 'tensoring' which is a common technique in the LSH literature.

Tensoring means constructing a large space partition $P \subseteq \mathcal{P}(X)$ by taking multiple smaller random partitions $P_1, P_2, \ldots$ and taking all the intersections $P = \{p_1 \cap p_2, \ldots \mid p_1 \in P_1, p_2 \in P_2, \ldots\}$. Often the implicit partition P is nearly as good as a fully random partition of equal size, while it is cheaper to store in memory and allows much faster lookups of which section covers a given point. In this paper we are particularly interested in P_i's that partition different small sub-spaces, such that P is used to increase the dimension of a small, explicit, good partition.

Unfortunately tensoring doesn't seem to be directly applicable for deterministic constructions, since deterministic space partitions tend to have some overhead that gets amplified by the product construction. This is the reason why [Pagh, 2016] constructs hash functions directly using algebraic methods, rather than starting with a small hash function and 'amplifying' as is common for LSH. Algebraic methods are great when they exist, but they tend to be hard to find, and it would be a tough order to find them for every similarity measure we would like to make a data structure for.

It turns out we can use splitters to help make tensoring work deterministically. Roughly, these are generalizations of perfect hash functions. However, where a (d, m, k)-perfect hash family guarantees that for any set $S \subseteq [d]$ of size k, there is a function $\pi : [d] \to [m]$ such that $|\pi(S)| = k$, a (d, m)-splitter instead guarantees that the is some π such that $|S \cap \pi^{-1}(i)| = d/m$ for each $i = 1, \ldots, m$; or as close as possible if m does not divide d. That is, for any S there is some π that 'splits' S evenly between m buckets.

Using splitters with tensoring, we greatly limit the number of combinations of smaller space partitions that are needed to guarantee covering. We use this to amplify partitions found probabilistically and verified deterministically. The random aspect is however only for convenience, since the greedy set cover algorithm would suffice as well, as is done in [Alon et al., 2006]. We don't quite get a general reduction from Monte Carlo to Las Vegas LSH data structures, but we show how two state of the art algorithms may be converted at a negligible overhead.

A final technique to make everything come together is the use of dimensionality reductions. We can't quite use the standard bit-sampling and Johnson–Lindenstrauss lemmas, since those may (though unlikely) increase the distance between originally near points. Instead we use two dimensionality reduction lemmas based on partitioning. Similarly to [Pagh, 2016] and others, we fix a random permutation. Then given a vector $x \in \{0,1\}^d$ we permute the coordinates and partition into blocks $x_1, \ldots, x_{d/B}$ of size B. For some linear distance function, $\text{dist}(x, y) = \text{dist}(x_1, y_1) + \cdots + \text{dist}(x_{d/B}, y_{d/B})$, which implies that for some i we must have $\text{dist}(x_i, y_i) \leq \text{dist}(x, y)B/d$. Running the algorithm separately for each set of blocks

guarantee that we no pair gets mapped too far away from each other, while the randomness of the permutation lets us apply standard Chernoff bounds on how close the remaining points get.

Partitioning, however, doesn't work well if distances are very small, $cr << d$. This is because we need $B = \frac{d}{cr}\epsilon^{-2}\log n$ to get the said Chernoff bounds on distances for points at distance cr. We solve this problem by hashing coordinates into buckets of $\approx cr/\epsilon$ and taking the xor of each bucket. This has the effect of increasing distances and thereby allowing us to partition into blocks of size $\approx \epsilon^{-3}\log n$. A similar technique was used for dimensionality reduction in [Kushilevitz et al., 2000], but without deterministic guarantees. The problem is tackled fully deterministically in [Indyk, 2000a] using codes, but with the slightly worse bound of $\epsilon^{-4}\log n$.

For the second problem of Braun-Blanquet similarity we also need a way to reduce the dimension to a manageble size. Using randomized reductions (for example partitioning), we can reduce to $|x \cap y| \sim \log n$ without introducing too many false positives. However we could easily have e.g. universe size $d = (\log n)^{100}$ and $|x| = |y| = (\log n)^2$, which is much too high a dimension for our splitter technique to work. There is probably no hope of actually reducing d, since increasing $|x|/d$ and $|y|/d$ makes the problem we are trying to solve easier, and such a reduction would thus break LSH lower bounds.

Instead we introduce tensoring technique based on perfect hash functions, which allows us to create Turán Systems with very large universe sizes for very little overhead.

In the process of showing our results, we show a useful bound on the ratio between two binomial coefficients, which may be of separate interest.

1.4 Notation

We use $[d] = \{1, \ldots, d\}$ as convenient notation sets of a given size. Somewhat overloading notation, for a predicate P, we also use the Iversonian notation $[P]$ for a value that is 1 if P is true and 0 otherwise.

For a set $x \subseteq [d]$, we will sometimes think of it as a subset of the universe $[d]$, and at other times as a vector $x \in \{0,1\}^d$, where $x_i = 1$ indicates that $i \in x$. This correspondence goes further, and we may refer to the set size $|x|$ or the vector norm $\|x\|$, which is always the Hamming norm, $\|x\| = \sum_{i=1}^{d} x_i$. Similarly for two sets or points $x, y \in \{0,1\}^d$, we may refer to the inner product $\langle x, y\rangle = \sum_{i=1}^{d} x_i y_i$ or to the size of their intersection $|x \cap y|$.

We use $S \times T = \{(s,t) : s \in S, t \in T\}$ for the cross product, and $x \oplus y$ for symmetric difference (or 'xor'). $\mathcal{P}(X)$ is the power set of X, such that $x \subseteq X \equiv x \in \mathcal{P}(X)$. $\binom{X}{k}$ denotes all subsets of X of size k.

For at set $S \subseteq [d]$ and a vector $x \in \{0,1\}^d$, we let x_S be the projection of x onto S. This is an $|S|$-dimensional vector, consisting of the coordinates $x_S = \langle x_i : i \in S\rangle$ in the natural order of i. For a function $f : [a] \to [b]$ we let $f^{-1} : \mathcal{P}([b]) \to \mathcal{P}([a])$ be the

'pullback' of f, such that $f^{-1}(S) = \{i \in [a] \mid f(i) \in S\}$. For example, for $x \in \{0,1\}^a$, we may write $x_{f^{-1}(1)}$ to be the vector x projected onto the coordinates of $f^{-1}(\{1\})$.

Sometimes when a variable is $\omega(1)$ we may assume it is integral, when this is achievable easily by rounding that only perturbs the result by an insignificant $o(1)$ amount.

The functional $\text{poly}(a, b, \dots)$ means any polynomial combination of the arguments, essentially the same set as $(a \cdot b \dots)^{\pm O(1)}$.

1.5 Organization

We start by laying out the general framework shared between our algorithms. We use a relatively common approach to modern near neighbour data structures, but the overview also helps establish some notation used in the later sections.

The second part of section 2 describes the main ideas and intuition on how we achieve our results. In particular it defines the concept of 'splitters' and how they may be used to create list-decodable codes for various measures. The section finally touches upon the issues we encounter on dimensionality reduction, which we can use to an extent, but which is restricted by our requirement of '1-sided' errors.

In sections 3 and 4 we prove the main theorems from the introduction. The sections follow a similar pattern: First we introduce a filter family and prove its existence, then we show a dimensionality reduction lemma and analyze the resulting algorithm.

2 Overview

Both algorithms in this paper follow the structure of the Locality Sensitive Filter framework, which is as follows: For a given universe U, we define a family $\mathcal{F}$ of 'filters' equipped with a (possibly random) function $F : U \to \mathcal{P}(\mathcal{F})$, which assigns every point a set of filters.

Typically, $\mathcal{F}$ will be a generous covering of U, and $F(x)$ will be the sets that cover the point x. Critically, any pair x, y that is close/similar enough in U must share a filter, such that $F(X) \cap F(Y) \neq \varnothing$. Further we will want that pairs x, y that are sufficiently far/dissimilar only rarely share a filter, such that $E[|F(x) \cap F(Y)|]$ is tiny.

To construct the data structure, we are given a set of data points $P \subseteq U$. We compute $F(x)$ for every $x \in P$ and store the points in a (hash) map $T : \mathcal{F} \to \mathcal{P}(P)$. For any point $x \in P$ and filter $f \in F(x)$, we store $x \in T[f]$. Note that the same x may be stored in multiple different buckets.

To query the data structure with a point $x \in U$, we compute the distance/similarity between x and every point $y \in \bigcup_{f \in F(x)} T[f]$, returning the first suitable candidate, if any.

There are many possible variations of the scheme, such as sampling $\mathcal{F}$ from a distribution of filter families. In case we want a data structure with space/time trade-offs, we can use different $\mathcal{F}$ functions for data points and query points. However in this article we will not include these extensions.

We note that while it is easy to delete and insert new points in the data structure after creation, we are going to choose $\mathcal{F}$ parametrized on the total number of points, $|P|$. This makes our data structure essentially static, but luckily [Overmars and van Leeuwen, 1981] have found general, deterministic reductions from dynamic to static data structures.

2.1 Intuition

The main challenge in this paper will be the construction of filter families $\mathcal{F}$ which are: (i) not too large; (ii) have a $F(\cdot)$ function that is efficient to evaluate; and most importantly, (iii) guarantee that all sufficiently close/similar points always share a filter. The last requirement is what makes our algorithm different from previous results, which only had this guarantee probabilistically.

For concreteness, let us consider the Hamming space problem. Observe that for very low dimensional spaces, $d = (1 + o(1)) \log n$, we can afford to spend exponential time designing a filter family. In particular we can formulate a set cover problem, in which we wish to cover each pair of points at distance $\leq r$ with Hamming balls of radius s. This gives a family that is not much larger than what can be achieved probabilistically, and which is guaranteed to work. Furthermore, this family has sublinear size $(n^{o(1)})$, making $F(x)$ efficient to evaluate, since we can simply enumerate all of the Hamming balls and check if x is contained.

The challenge is to scale this approach up to general d.

Using a standard approach of randomly partitioning the coordinates, we can reduce the dimension to $(\log n)^{1+\epsilon}$. This is basically dimensionality reduction by bit sampling, but it produces $d/\log n$ different subspaces, such that for any pair x, y there is at least one subspace in which their distance is not increased. We are left with a gap from $(\log n)^{1+\epsilon}$ down to $\log n$. Bridging this gap turns out to require a lot more work. Intuitively we cannot hope to simply use a stronger dimensionality reduction, since $\log n$ dimensions only just fit n points in Hamming space and would surely make too many non-similar points collide to be effective.

A natural idea is to construct higher-dimensional filter families by combining multiple smaller families. This is a common technique from the first list decodable error correcting codes, for example [Elias, 1957]: Given a code $\mathcal{C} \subseteq \{0,1\}^d$ with covering radius r, we can create a new code $\mathcal{C}_2 \subseteq \{0,1\}^{2d}$ of size $|\mathcal{C}|^2$ with covering radius $2r$ by taking every pair of code words and concatenating them. Then for a given point $x \in \{0,1\}^{2d}$ we can decode the first and last d coordinates of $x = x_1x_2$ separately in $\mathcal{C}$. This returns two code words c_1, c_2 such that $\text{dist}(x_1, c_1) \leq r$ and $\text{dist}(x_2, c_2) \leq r$. By construction c_1c_2 is in $\mathcal{C}_2$ and $\text{dist}(x_1x_2, c_1c_2) \leq 2r$.

This combination idea gives is nice when it applies. When used with high quality inner codes, the combined code is close to optimal as well. In most cases the properties of $\mathcal{C}$ that we are interested in won't decompose as nicely. With the example of our Hamming ball filter family, consider $x, y \in \{0,1\}^{2d}$ with distance $\text{dist}(x,y) = r$. If we split $x = x_1x_2$ an $y = y_1y_2$ we could decode the smaller vectors individually in a smaller family, however we don't have any guarantee on $\text{dist}(x_1, y_1)$ and $\text{dist}(x_2, y_2)$ individually, so the inner code might fail to return anything at all.

To solve this problem, we use a classic tool for creating combinatorial objects, such as our filter families, called 'splitters'. Originally introduced by [Mairson, 1983, Naor et al., 1995] they are defined as follows:

Definition 3 (Splitter)**.** *A (B,l)-splitter H is a family of functions from $\{1,\ldots,B\}$ to $\{1,\ldots,l\}$ such that for all $S \subseteq \{1,\ldots,B\}$, there is a $h \in H$ that splits S perfectly, i.e., into equal-sized parts $h^{-1}(j) \cap S$, $j = 1,2,\ldots,l$. (or as equal as possible, if l does not divide $|S|$).*

The size of H is at most B^l, and using either a generalization by [Alon et al., 2006] or a simple combinatorial argument, it is possible to ensure that the size of each part $|h^{-1}(j)|$ equals B/l (or as close as possible).

We now explain how splitters help us combine filter families. Let H be a splitter from $\{1,\ldots,2d\}$ to $\{1,2\}$. For any $x, y \in \{0,1\}^{2d}$ we can let S be the set of coordinates on which x and y differ. Then there is a function $h \in H$ such that $|h^{-1}(1) \cap S| = |h^{-1}(2) \cap S| = |S|/2$. (Or as close as possible if $|S|$ is odd.) If we repeat the failed product combination from above for every $h \in H$ we get a way to scale our family from d to $2d$ dimensions, taking the size from $|\mathcal{F}|$ to $(2d)^2|\mathcal{F}|^2$. That is, we only suffer a small polynomial loss. In the end it turns out that the loss suffered from creating filter families using this divide and conquer approach can be contained, thus solving our problem.

An issue that comes up, is that the 'property' we are splitting (such as distance) can often be a lot smaller than the dimensionality d of the points. In particular this original dimensionality reduction may suffer an overhead factor $d/|S|$, which could make it nearly useless if $|S|$ is close to 1. To solve this problem, both of our algorithms employ special half-deterministic dimensionality reductions, which ensures that the interesting properties get 'boosted' and end up taking a reasonable amount of 'space'. These reductions are themselves not complicated, but they are somewhat non-standard, since they can only have a one sided error. For example for Hamming distance we need that the mapped distance is never larger than its expected value, since otherwise we could get false negatives.

For Hamming distance our dimension reduction works by hashing the coordinates randomly from [d] to [m] taking the xor of the coordinates in each bucket. This is related to the β-test in [Kushilevitz et al., 2000]. The idea is that if x and y are different in only a few coordinates, then taking a small random group of coordinates, it is likely to contain at most one where they differ. If no coordinates differ, then

after taking the xor the result will still be the same, but if exactly one (or an odd number) of coordinates differ, the resulting coordinate will be different.

For set similarity things are a bit more hairy. There is no data independent dimensionality reduction that can reduce the size of the domain. In fact this would break the lower bounds of e.g. [Pagh and Christiani, 2017]. Instead we come up with a new construction based on perfect hash functions, which greatly increases the number of filters needed, but only as much as we can afford given the easier sub-problems.

The idea can be motivated as follows: Suppose you have to make a family of sets $\mathcal{T} \subseteq \mathcal{P}([n])$ of size r, such that for each each set $K \subseteq [n]$ of size $|K| = k$ there is an $R \in \mathcal{T}$ such that $R \subseteq K$. Then you might try to extend this to the domain $[2n]$ as follows: For each $R \in \mathcal{T}$ and each $b \in \{0,1\}^r$, make a new set $R' = \{i + nb_i : i \in R\}$ (where b_i is padded appropriately). This creates $2^r|\mathcal{T}|$ new sets, which can be shown to have the property, that for any set $K \subseteq [2n]$ of size $|K| = k$, there is an R' such that $R' \subseteq K$. That is as long as $K \cap (K - n) = \varnothing$, since then we can consider $R \in \mathcal{T}$ such that $R \subseteq (K \bmod n)$. That is R is a subset of K 'folded' into the first $[n]$ elements, and one of the R' will be a subset of K.

Because of the requirement that $|K \bmod n| = k$ we need to use perfect hashing as a part of the construction. However for non-Las Vegas algorithms, a similar approach may be useful, simply using random hashing.

3 Hamming Space Data Structure

We will give an efficient filter family for LSF in Hamming space. Afterwards we will analyze it and choose the most optimal parameters, including dimensionality reduction.

Lemma 1. *For every choice of parameters $B, b \in \mathbb{N}$, $b \leq B$, $0 < r < B/2$ and $s^2 = O(B/\sqrt{b})$, there exists a code $\mathcal{C} \subseteq \{0,1\}^B$ of size $|\mathcal{C}| = \text{poly}(B^{B/b}) \exp(\frac{s^2}{2(1-r/d)})$ with the following properties:*

1. *Given $x \in \{0,1\}^B$ we can find a subset $C(x) \subseteq \{c \in \mathcal{C} : \text{dist}(x,c) \leq B/2 - s\sqrt{B}/2\}$ in time $|C(x)| + \text{poly}(B^{B/b}, e^{s^2 b/B})$.*
2. *For all pairs $x, y \in \{0,1\}^B$ with $\text{dist}(x,y) \leq r$ there is some common nearby code word $c \in C(x) \cap C(y)$.*
3. *The code requires $4^b \,\text{poly}(B^{B/b}, e^{s^2 b/B})$ time for preprocessing and $\text{poly}(B^{B/b}, e^{s^2 b/B})$ space.*

Note that we don't actually guarantee that our 'list-decoding' function $C(x)$ returns *all* nearby code words, just that it returns enough for property (2) which is what we need for the data structure. This is however not intrinsic to the methods

and using a decoding algorithm similar to [Becker et al., 2016] would make it a 'true' list-decoding.

Proof. We first show how to construct a family for $\{0,1\}^b$, then how to enlarge it for $\{0,1\}^B$. We then show that it has the covering property and finally the decoding properties. In order for our probabilistic arguments to go through, we need the following lemma, which follows from Stirling's Approximation:

Lemma 2. *For $t = \frac{d}{2} - \frac{s\sqrt{d}}{2}$, $1 \le s \le d^{1/4}/2$ and $r < d/2$, Let $x, y \in \{0,1\}^d$ be two points at distance $\mathrm{dist}(x,y) = r$, and let $I = |\{z \in \{0,1\}^d : \mathrm{dist}(z,x) \le t, \mathrm{dist}(z,y) \le t\}|$ be the size of the intersection of two hamming balls around x and y of radius t, then*

$$\frac{7}{8d}\exp\left(\tfrac{-s^2}{2(1-r/d)}\right) \le I\,2^{-d} \le \exp\left(\tfrac{-s^2}{2(1-r/d)}\right)$$

Proof is in the appendix.

Let $s' = s\sqrt{b/B}$. Consider any two points $x, y \in \{0,1\}^b$ with distance $\le (r/d)b$. If we choose $a \in \{0,1\}^b$ uniformly at random, by lemma 2 we have probability $p = \mathrm{poly}(b)\exp(\frac{-s'^2}{2(1-r/d)})$ that both x and y have distance at most $t = b/2 - s'\sqrt{b/4}$ with c. By the union bound over pairs in $\{0,1\}^b$, if we sample $p^{-1}b\log 2$ independent as, we get constant probability that some a works for every pair. We can verify that a set A of such filters indeed works for every pair in time $4^b|A|$. By repeatedly sampling sets A and verifying them, we get a working A in expected $O(1)$ tries.[1]

Next we define $\mathcal{C}$. We build a splitter, that is a set Π of functions $\pi : [B] \to [B/b]$, such that for every set $I \subseteq [B]$ there is a $\pi \in \Pi$ such that $\lfloor |I|b/B \rfloor \le |\pi^{-1}(j) \cap I| \le \lceil |I|b/B \rceil$ for $j = 1, \dots, B/b$. By the discussion after definition 3, such a set of size $\mathrm{poly}(B^{B/b})$ exists and can be constructed in time and space proportional to its size. Implicitly we can then define $\mathcal{C}$ by making one code word $c \in \{0,1\}^B$ for every $\pi \in \Pi$ and $1 \le j_1, \dots, j_{B/b} \le |A|$, satisfying the requirement that $c_{\pi^{-1}(j_k)} = A_{j_k}$ for $k = 1 \dots B/b$. That is, for a given set of rows of A and a split of $[B]$, we combine the rows into one row c such that each row occupies a separate part of the split. Note that this is possible, since splitter has all partitions of equal size, b. The created family then has size $|\mathcal{C}| = |\Pi||A|^{B/b} = \mathrm{poly}(B^{B/b})\exp(\frac{-s^2}{2(1-r/d)})$ as promised. Since the only explicit work we had to do was finding A, we have property (3) of the lemma.

We define the decoding function $C(x) \in \mathcal{C}$ for $x \in \{0,1\}^B$ with the following algorithm: For each $\pi \in \Pi$ compute the inner decodings $A_j = \{a \in A : \mathrm{dist}(x_{\pi^{-1}(j)}, a) \le b/2 - s\sqrt{b}/2\}$ for $j = 1, \dots, B/b$. Return the set of all concatenations in the product of the A_j's: $C(x) = \{a_1 \| a_2 \| \dots \| a_{B/b} : a_1 \in A_1, \dots\}$. Computing the A_j's take time $(B/b)|A|$, while computing and concatenating the product takes linear time in the size of the output. This shows property (1).

[1] The randomness is not essential, and we could as well formulate a set cover instance and solve it using the greedy algorithm, which matches the probabilistic construction up to a log factor in size and time.

Finally for property (2), consider a pair $x, y \in \{0,1\}^B$ of distance $\leq r$. Let I be the set of coordinates on which x and y differ. Then there is a function $\pi \in \Pi$ such that x and y differ in at most $|I|b/B = rb/B$ coordinates in each subset $\pi^{-1}(1), \ldots, \pi^{-1}(B/b) \subseteq [B]$. Now for each pair of projected vectors $x_{\pi^{-1}(1)}, y_{\pi^{-1}(1)}, \ldots$ (let's call them $x_1, y_1, \ldots$) there is an $a_j \in A$ such that $\mathrm{dist}(a_j, x_j) \leq b/2 - s'\sqrt{b}/2$ and $\mathrm{dist}(a_j, y_j) \leq b/2 - s'\sqrt{b}/2$. This means that x and y must both have distance at most $(b/2 - s'/2)B/b = B/2 - s\sqrt{B}/2$ to that $c \in \mathcal{C}$ which has $c_{\pi^{-1}(j)} = a_j$ for $j = 1 \ldots B/b$. By the same reasoning, this c will be present in both $C(x)$ and $C(y)$, which proves the lemma. □

Returning to the problem of near neighbour search in $\{0,1\}^d$, it is clear from the $4^b \,\mathrm{poly}(B^{B/b})$ construction time of the above family, that it will not be efficient for dimension $B = (\log n)^{\omega(1)}$. For this reason we will apply the following dimensionality reduction lemma:

Lemma 3. *Given $d \geq cr > r \geq 1$ and $\epsilon, \delta > 0$, define $B = 27\epsilon^{-3} \log 1/\delta$ and $m = 3cr/\epsilon$ and assume $\delta^{-1} \geq m$, then there is a random set F of at most $S = m/B$ functions $f : \{0,1\}^d \to \{0,1\}^B$ with the following properties for every $x, y \in \{0,1\}^d$:*

1. *With probability 1, there is at least one $f \in F$ st.:*

$$\mathrm{dist}(f(x), f(y)) \leq \mathrm{dist}(x, y)/S.$$

2. *If $\mathrm{dist}(x, y) \geq cr$ then for every $f \in F$ with probability at least $1 - \delta$:*

$$\mathrm{dist}(f(x), f(y)) \geq (1 - \epsilon)cr/S.$$

The idea is to randomly throw the d coordinates in $m = 3cr/\epsilon$ buckets, (xor-ing the value if more than two coordinates land in the same group.) For two points with $\leq cr$ differences, this has the effect of rarely colliding two coordinates at which the points disagree, thus preserving distances. It further has the effect of changing the relative distances from arbitarily low r/d to something around ϵ, which allows partitioning the coordinates into groups of around $\epsilon^{-3} \log 1/\delta$ coordinates using Chernoff bounds.

Proof. To prove lemma 3 first notice that if $B \geq d$ we can use the identity function and we are done. If $B \geq m$, then we can duplicate the vector $\lceil m/B \rceil = O(\epsilon^{-2} \log 1/\delta)$ times. Also, by adjusting ϵ by a constant factor, we can assume that B divides m.

For the construction, pick a random function $h : [d] \to [m]$. Define $g : \{0,1\}^d \to \{0,1\}^m$ by 'xor'ing the contents of each bucket, $g(x)_i = \bigoplus_{j \in h^{-1}(i)} x_j$, and let $f_i(x) = g(x)_{(iB,(i+1)B]}$ for $i = 0 \ldots m/B$ be the set of functions in the lemma. We first show that this set has property (1) and then property (2).

Observe that g never increases distances, since for any $x, y \in \{0,1\}^d$ the distance

$$\mathrm{dist}(g(x), g(y)) = \sum_{i=1}^{m} \left[\bigoplus_{j \in h^{-1}(i)} x_j \neq \bigoplus_{j \in h^{-1}(i)} y_j \right]$$

is just $\sum_{i=1}^{m}\left(\sum_{j\in h^{-1}(i)}[x_j \neq y_j] \bmod 2\right)$ which is upper bounded by the number of coordinates at which x and y differ. By averaging, there must be one f_i such that $\mathrm{dist}(f_i(x), f_i(x)) \leq \mathrm{dist}(g(x), g(y))B/m \leq \mathrm{dist}(x, y)/S$.

For the second property, let $R = \mathrm{dist}(x, y) \geq cr$ and let $X_1, \ldots, X_m$ be the random number of differences between x and y in each bucket under h. Let $Y_1, \ldots, Y_m$ be iid. Poisson distributed variables with mean $\lambda = \mathrm{E}\, X_1 = R/m \geq \epsilon/3$. We use the the Poisson trick from [Mitzenmacher and Upfal, 2005] theorem 5.7:

$$\Pr[\sum_{i=1}^{B}(X_i \bmod 2) < x] \leq e\sqrt{m}\Pr[\sum_{i=1}^{B}(Y_i \bmod 2) < x].$$

The probability $\Pr[Y \bmod 2 \equiv 1]$ that a Poisson random variable is odd is $(G(1) - G(-1))/2$ where $G(z) = \sum_i \Pr[Y = i]z^i = e^{\lambda(z-1)}$. This gives us the bound $\Pr[Y_i \bmod 2 \equiv 1] = (1-e^{-2\lambda})/2 \geq (1-e^{-2\epsilon/3})/2 \geq (1-\epsilon/3)\epsilon/3$. We can then bound the probability of an f_i decreasing distances too much, using a Chernoff bound ($\Pr[Z \leq x] \leq \exp(-D[x/B \mid p]B)$):

$$\begin{aligned}
&\Pr[\mathrm{dist}(f_i(x), f_i(y)) \leq (1-\epsilon)cr/S] \\
&\quad \leq e\sqrt{m}\exp(-D[(1-\epsilon)\epsilon/3 \mid (1-\epsilon/3)\epsilon/3]B) \\
&\quad \leq e\sqrt{m}\exp(-2\epsilon^3 B/27).
\end{aligned}$$

Since $cr/S = crB/m = B\epsilon/3$. Here $D[\alpha \mid \beta] = \alpha\log\frac{\alpha}{\beta} + (1-\alpha)\log\frac{1-\alpha}{1-\beta} \geq (\alpha-\beta)^2/(2\beta)$ is the Kullback–Leibler divergence. For our choice of B the error probability is then $e\sqrt{m}\delta^2$ which is less than δ by our assumptions. This proves the lemma. □

Using lemma 3 we can make at most $3cr/\epsilon = O(d/\epsilon)$ data structures, as described below, and be guaranteed that in one of them, we will find a near neighbour at distance $r' = r/S = \epsilon/(3c)B$. In each data structure we will have to reduce the distance cr', at which we expect far points to appear, to $cr'(1-\epsilon)$. This ensures we see at most a constant number of false positives in each data structure, which we can easily filter out. For $\epsilon = o(1)$ this change be swallowed by the approximation factor c, and won't significantly impair our performance.

When using the filter family of lemma 1 for LSF, the time usage for queries and inserting points is dominated by two parts: 1) The complexity of evaluating $C(x)$, and 2) The expected number of points at distance larger than $cr'(1-\epsilon)$ that falls in the same filter as x.

By randomly permuting and flipping the coordinates, we can assume that x is a random point in $\{0,1\}^d$. The expected time to decode $C(x)$ is then

$$\begin{aligned}
&\mathrm{E}\,|C(x)| + \mathrm{poly}(B^{B/b}, e^{s^2 b/B}) \\
&= |\mathcal{C}|\Pr_x[0 \in C(x)] + \mathrm{poly}(B^{B/b}, e^{s^2 b/B}) \\
&\leq \mathrm{poly}(B^{B/b}, e^{s^2 b/B})\exp\left(\tfrac{s^2}{2(1-r'/B)} - \tfrac{s^2}{2}\right).
\end{aligned}$$

For estimating collisions with far points, we can similarly assume that x and y are random points in $\{0,1\}^d$ with fixed distance $cr'(1-\epsilon)$:

$$\begin{aligned}
&\mathrm{E}\,|\{y \in P : C(x) \cap C(y) \neq \varnothing\}| \\
&\leq n\,|\mathcal{C}| \Pr_{x,y}[0 \in C(x), 0 \in C(y)] \\
&\leq B^{O(B/b)} \exp\left(\tfrac{s^2}{2(1-r'/B)} - \tfrac{s^2}{2(1-c(1-\epsilon)r'/B)} + \log n\right) \\
&= B^{O(B/b)} \exp\left(\tfrac{s^2}{2}(\tfrac{1}{1-r'/B} - \tfrac{1}{1-cr'/B} + O(\epsilon)) + \log n\right).
\end{aligned}$$

Finally we should recall that constructing the data structures takes time $4^b\,\mathrm{poly}(e^{s^2 b/B})$ plus n inserts.

We now choose the parameters:

$$\begin{aligned}
s^2/2 &= \tfrac{1-cr'/B}{cr'/B}\log n, & B &= 27\epsilon^{-3}\log n, \\
b &= \log_4 n, & \epsilon &= (\log n)^{-1/4}.
\end{aligned}$$

This makes the code construction time $n^{1+o(1)}$ while evaluating $C(x)$ and looking at far points takes expected time at most $n^{1/c+\tilde{O}(\log n)^{-1/4}}$. To use lemma 1 we have to check that $s^2 = O(B/\sqrt{b}) = O((\log n)^{5/4})$, but $s^2/2 = \frac{1-\epsilon/3}{\epsilon/3}\log n = (\log n)^{5/4}(1-o(1))$ so everything works out. This shows theorem 1.

To get the result of corollary 1, we just need to substitute the dimensionality reduction lemma 3 for a simple partitioning approach. (Lemma 4 below.) The idea is that of [Pagh, 2016] which is to randomly partition the d coordinates in B parts and run the algorithm on those. The important point is that in this case $r'/B = r/d$, so the relative distance is not decreased. We choose parameters

$$\begin{aligned}
s^2/2 &= \tfrac{1-cr/d}{cr/d}\log n & B &= O(\epsilon^{-2}(cr/d)^{-1}\log n), \\
b &= \log_4 n, & \epsilon &= (\log n)^{-1/3}.
\end{aligned}$$

This again makes this makes the code construction time $n^{1+o(1)}$ while evaluating $C(x)$ and looking at far points takes time $n^{\frac{1-c\delta}{c(1-\delta)}+\tilde{O}(\log n)^{-1/3}d/r}$ as in the corollary. Again we need need to check that $s^2 = O(B/\sqrt{b}) = O((\log n)^{7/6})$. This works as long as $r/d = \Omega((\log n)^{-1/6})$, which is the assumption of the corollary.

Lemma 4. *For any $d \geq r \geq 1$ and $\epsilon > 0$ there is a set F of d/B functions, $f : \{0,1\}^d \to \{0,1\}^B$, where $B = 2\epsilon^{-2}d/(cr)\log n$, such that:*

1. *With probability 1, there is at least one $f \in F$ st.:*

$$\mathrm{dist}(f(x), f(y)) \leq \mathrm{dist}(x,y)\,B/d.$$

2. *If $\mathrm{dist}(x,y) \geq cr$ then for every $f \in F$ with probability at least $1 - 1/n$:*

$$\mathrm{dist}(f(x), f(y)) \geq (1-\epsilon)cr\,B/d.$$

Proof. Fix a random permutation. Given $x \in \{0,1\}^d$, shuffle the coordinates using the permutation. Let $f_1(x)$ be the first B coordinates of x, $f_2(x)$ the next B coordinates and so forth. For any $y \in \{0,1\}^d$, after shuffling, the expected number of differences in some block of B coordinates is $\text{dist}(x,y)B/d$. Then the first property follows because $\sum_i \text{dist}(f_i(x), f_i(y)) = \text{dist}(x,y)$ so not all distances can be below the expectation. The second property follows from the Chernoff/Hoeffding bound [Hoeffding, 1963]. □

4 Set Similarity Data Structure

We explore the generality of our methods, by making a Las Vegas version of another very common LSH data structure. Recall the theorem we are trying to prove, from the introduction:

Theorem. *Given a set P of n subsets of $[d]$, define the Braun-Blanquet similarity $\text{sim}(x,y) = |x \cap y| / \max(|x|,|y|)$ on the elements of P. For every choice of $0 < b_2 < b_1 < 1$ there is a data structure on P that supports the following query operation:*

On input $q \subseteq [d]$, for which there exists a set $x \in P$ with $\text{sim}(x,q) \geq b_1$, return an $x' \in P$ with $\text{sim}(x',q) > b_2$. The data structure uses expected time dn^ρ per query, can be constructed in expected time $dn^{1+\rho}$, and takes expected space $n^{1+\rho} + dn$ where $\rho = \frac{\log 1/b_1}{\log 1/b_2} + \hat{O}(1/\sqrt{\log n})$.

By known reductions [Pagh and Christiani, 2017] we can focus on the case where all sets have the same weight, w, which is known to the algorithm. This works by grouping sets in data structures with sets of similar weight and uses no randomization. The price is only a $(\log n)^{O(1)}$ factor in time and space, which is dominated by the $n^{\hat{O}(1/\sqrt{\log n})}$ factor in the theorem.

When two sets have equal weight w, being b-close in Braun-Blanquet similarity coresponds exactly to having an intersection of size bw. Hence for the data structure, when trying to solve the (b_1, b_2)-approximate similarity search problem, we may assume that the intersections between the query and the 'close' sets we are interested in are at least $b_1 w$, while the intersections between the query and the 'far' sets we are not interested in are at most $b_2 w$.

The structure of the algorithm follows the LSF framework as in the previous section. A good filter family for set similarity turns out to be the r-element blocks of a Turán system. This choice is inspired by [Pagh and Christiani, 2017] who sampled subsets with replacement.

Definition 4 ([Turán, 1961, Colbourn and Dinitz, 2006])**.** *A Turán (n,k,r)-system is a collection of r-element subsets, called 'blocks', of an n element set X such that every k element subset of X contains at least one of the blocks.*

Turán systems are well studied on their own, however all known general constructions are either only existential or of suboptimal size. The following lemma

provides the first construction to tick both boxes, and with the added benefit of being efficiently decodable.

An interesting difference between this section and the last, is that we don't know how to do a dimensionality reduction like we did for hamming distance. Instead we are (luckily) able to make an efficiently decodable filter family even for very large dimensional data points.

Lemma 5. *For every n, k, r, where $n > k > r^{3/2}$, there exists a Turán (n, k, r)-system, $\mathcal{T}$, of size $|\mathcal{T}| \leq (n/k)^r e^{\chi}$ where $\chi = O(\sqrt{r}\log r + \log k + \log\log n)$ with the following properties:*

1. *Correctness: For every set $K \subseteq [n]$ of size at least k, there is at least one block $R \in \mathcal{T}$ such that $R \subseteq K$.*
2. *Efficient decoding: Given a set $S \subseteq [n]$, we can find all the blocks it contains $T(S) = \{R \in \mathcal{T} : R \subseteq S\}$ in time $|S||T(S)| + e^{\chi}$. Furthermore, $T(S)$ has expected size $\leq (|S|/k)^r e^{\chi}$.*
3. *Efficient construction: The system is constructible, implicitly, in $e^{r(1+o(1))}$ time and space.*

Notes: A simple volume lower bound shows that an (n, k, r)-system must have at least $\binom{n}{r}/\binom{k}{r} \geq (n/k)^r$ blocks, making our construction optimal up the factor e^{χ}. Using the sharper bound $\binom{n}{r}/\binom{k}{r} \approx (n/k)^r \exp(\frac{r^2}{2k})$ from lemma 6, we get that the factor $\exp(\Omega(\sqrt{r}))$ is indeed tight for $k = O(r^{3/2})$.

The size of the system is in 'expectation', which is sufficient for our purposes, but is in fact fairly easy to fix. On the other hand, the 'expectation' in the size of sets $T(S)$ seems harder to get rid of, which is the reason why the data strcuture is Las Vegas and not deterministic.

4.1 The algorithm

We continue to describe the algorithm and proving theorem 2 using the lemma. The proof of the lemma is at the end and will be the main difficulty of the section.

As discussed, by the reduction of [Pagh and Christiani, 2017] we can assume that all sets have weight w, intersections with close sets have size $\geq b_1 w$ and intersections with far sets have size $\leq b_2 w$. The data structure follows the LSF scheme as in the previous section. For filters we use a Turán $(d, b_1 w, \frac{\log n}{\log 1/b_2})$ design, constructed by lemma 5. Note that if $b_1 w < (\frac{\log n}{\log 1/b_2})^{3/2}$ ($k < r^{3/2}$ in the terms of the lemma), we can simply concatenate the vectors with themselves $O(\log n)$ times. If $b_1 w \leq \frac{\log n}{\log 1/b_2}$ we can simply use all the $\binom{d}{b_1 w}$ sets of size $b_1 w$ as a Turán $(d, b_1 w, b_1 w)$ system and get a fast deterministic data structure.

As in the previous section, given a dataset P of n subsets of $[d]$, we build the data structure by decoding each set in our filter system $\mathcal{T}$. We store pointers from

each set $R \in \mathcal{T}$ to the elements of P in which they are a contained. By the lemma, this takes time $n(w(w/k)^r + 1)e^\chi = wn(1/b_1)^{\frac{\log n}{\log 1/b_2}} e^\chi \leq dn^\rho$, while expected space usage is $n(w/k)^r e^\chi + e^{r(1+o(1))} + dn = n^\rho + dn$ as in the theorem. Building $\mathcal{T}$ takes time $e^{r(1+o(1))} = n^{(1+o(1))/\log 1/b_2} = n^{1+o(1)}$.

Queries are likewise done by decoding the query set x in $\mathcal{T}$ and inspecting each point $y \in P$ for which there exists $R \in \mathcal{T}$ with $R \subseteq y$, until a point y' with $\mathrm{sim}(x, y') > b_2$ is found. Let's call this the candidate set of x. The expected number of false positive points in the candidates is thus

$$\sum_{y \in P} \mathrm{E}[|\{R \in \mathcal{T} : R \subseteq x \cap y\}|] = \sum_{y \in P} \mathrm{E}[|T(x \cap y)|] \leq n(b_2 w/(b_1 w))^{\frac{\log n}{\log 1/b_2}} e^\chi = n^\rho.$$

Computing the actual similarity takes time w, so this step takes time at most $wn^\rho \leq dn^\rho$. We also have to pay for actually decoding $T(x)$, but that takes time $w(w/(b_1 w))^{\frac{\log n}{\log 1/b_2}} e^\chi + e^\chi \leq dn^\rho$ as well.

Finally, to see that the algorithm is correct, if $\mathrm{sim}(x, y) \geq b_1$ we have $|x \cap y| \geq b_1 w$, and so by the Turán property of $\mathcal{T}$ there is $R \in \mathcal{T}$ such that $R \subseteq x \cap y$ which implies $R \subseteq x$ and $R \subseteq y$. This shows that there will always be at least one true high-similarity set in the candidates, which proves theorem 2.

4.2 The proof of lemma 5

Proof. We first prove the lemma in four parts, starting with a small design and making it larger step by step. To more easily describe the small designs, define $a = kr^{-3/2} \log(r^{3/2})$ and $b = \sqrt{r}$. The steps are then

1. Making a $(k^2/(a^2 b), k/(ab), r/b)$ using brute force methods.
2. Use splitters to scale it to $((k/a)^2, k/a, r)$.
3. Use perfect hashing to make it an $(n/a, k/a, r)$ design.
4. Use partitioning to finally make an (n, k, r) design.

We first prove the lemma without worrying about the above values being intergers. Then we'll show that each value is close enough to some integer that we can hide any excess due to approximation in the loss term.

The four steps can also be seen as proofs of the four inequalities:

$$\begin{aligned}
T(n, k, r) &\leq \tbinom{n}{r} / \tbinom{k}{r} (1 + \log \tbinom{n}{k}), \\
T(cn, ck, cr) &\leq \tbinom{cn}{c} T(n, k, r)^c, \\
T(ck^2, k, r) &\leq (k^4 \log ck^2) c^r T(k^2, k, r), \\
T(cn, ck, r) &\leq c T(n, k, r).
\end{aligned}$$

where the c are arbitrary integer constants > 0.

1. For convenience, define $n' = k^2/(a^2 b)$, $k' = k/(ab)$, $r' = r/b$ and assume they are all intergers. Probabilistically we can build a Turán (n', k', r')-system, $\mathcal{T}^{(n')}$, by sampling

$$\ell = \binom{n'}{r'}/\binom{k'}{r'}(1 + \log\binom{n'}{k'}) \leq (n'/k')^{r'} e^{r'^2/k'}(1 + k' \log(en'/k')) = (n'/k')^{r'} r^{5/2}(1 + o(1))$$

independent size r'-sets. (Here we used the bound on $\binom{n'}{r'}/\binom{k'}{r'}$ from lemma 6 in the appendix.) For any size k' subset, K, the probability that it contains none of the r'-sets is $\left(1 - \binom{k'}{r'}/\binom{n'}{r'}\right)^{\ell} \leq e^{-1}/\binom{n'}{k'}$. Hence by the union bound over all $\binom{n'}{k'}$ K-sets, there is constant probability that they all contain an r'-set, making our $\mathcal{T}^{(n')}$ a valid system.

We can verify the correctness of a sampled system, naiively, by trying iterating over all $\binom{n'}{k'}$ K-sets, and for each one check if any of the R-sets is contained. This takes time bounded by

$$\begin{aligned}\binom{n'}{k'}\ell &\leq (en'/k')^{k'}(n'/k')^{r'} r^{5/2}(1 + o(1)) \\ &= \left(\frac{er^{3/2}}{\log(r^{3/2})}\right)^{\frac{r}{\log(r^{3/2})}} \left(\frac{r^2}{\log(r^{3/2})}\right)^{\sqrt{r}} r^{O(1)} \\ &= e^{r + O(r/\log r)}\end{aligned}$$

as in the preprocessing time promised by the lemma. Since we had constant success probability, we can repeat the above steps an expected constant number of times to get a correct system.

Notice that the system has a simple decoding algorithm of brute-force iterating through all ℓ sets in $\mathcal{T}^{(n')}$.

2. To scale up the system, we proceed as in the previous section, by taking a splitter, Π, that is a set of functions $\pi : [bn'] \to [b]$ such that for any set $I \subseteq [bn']$ there is a $\pi \in \Pi$ such that

$$\lfloor |I|/b \rfloor \leq |\pi^{-1}(j) \cap I| \leq \lceil |I|/b \rceil \quad \text{for } j = 1, \ldots, b.$$

In other words, each π partitions $[bn']$ in b sets $[bn'] = \pi^{-1}(1) \cup \ldots \pi^{-1}(b)$ and for any subset $I \subseteq [bn']$ there is a partition which splits it as close to evenly as possible. We discuss the constructions of such sets of functions in the appendix.

For each $\pi \in \Pi$, and distinct $i_1, \ldots, i_b \in [|\mathcal{T}^{(n')}|]$, we make a br'-set, $R \subseteq [bn']$, which we think of as an indicator vector $\in \{0,1\}^{bn'}$, such that $R_{\pi^{-1}(j)} = \mathcal{T}^{(n')}_{i_j}$ for $j = 1 \ldots b$. That is, the new block, restricted to $\pi^{-1}(1), \pi^{-1}(2), \ldots$, will be equal to the i_1th, i_2th, ... blocks of $\mathcal{T}^{(n')}$. Another way to think of this is that we take the i_1th, i_2th, ... blocks of $\mathcal{T}^{(n')}$ considered as binary vectors in $\{0,1\}^{n'}$ and combine them to a bn' block 'spreading' them using π.

The new blocks taken together forms a family $\mathcal{T}^{(bn')}$ of size

$$|\mathcal{T}^{(bn')}| = |\Pi||\mathcal{T}^{(n')}|^b = \binom{bn'}{b}[(n'/k')^{r'} r^{O(1)}]^b \leq (en')^b (n'/k')^{br'} r^{O(b)} = (n'/k')^{br'} r^{O(b)},$$

where we used only the trivial bound $\binom{n}{k} \leq (en/k)^k$ and the fact that $n' = r^{O(1)}$.

We now show correctness of $\mathcal{T}^{(bn')}$. For this, consider any bk'-set $K \subseteq [bn']$. We need to show that there is some br'-set $R \in \mathcal{T}^{(bn')}$ such that $R \subseteq K$. By construction of Π, there is some $\pi \in \Pi$ such that $|\pi^{-1}(j) \cap K| = k'$ for all $j = 1, \ldots, b$. Considering K as an indicator vector, we look up $K_{\pi^{-1}(1)}, \ldots, K_{\pi^{-1}(b)}$ in $\mathcal{T}^{(n')}$, which gives us b disjoint r'-sets, $R'_1, \ldots, R'_b$. By construction of $\mathcal{T}^{(bn')}$ there is a single $R \in \mathcal{T}^{(bn')}$ such that $R_{\pi^{-1}(j)} = R'_j$ for all j. Now, since $R'_j \subseteq K_{\pi^{-1}(j)}$ for all j, we get $R \subseteq K$ and so we have proven $\mathcal{T}^{(bn')}$ to be a correct (bn', bk', br')-system.

Decoding $\mathcal{T}$ is fast, since we only have to do $|\Pi| \cdot b$ lookups in (enumerations of) $\mathcal{T}^{(n')}$. When evaluating $T^{(bn')}(K)$ we make sure we return every br'-set in K. Hence we return the entire "product" of unions:

$$T^{(bn')}(K) = \bigcup_{\pi \in \Pi} \{R_1 \cup \cdots \cup R_b : R_1 \in T^{(n')}(K_{\pi^{-1}(1)}), R_2 \in \ldots\}.$$

In total this takes time $|T^{(bn')}(K)|$ for the union product plus an overhead of $|\Pi| b |\mathcal{T}^{(n')}| \leq (en')^b r^{O(r'+b)} = r^{O(\sqrt{r})}$ for the individual decodings.

3. Next we convert our $((k/a)^2, k/a, r)$ design, $\mathcal{T}^{(bn')}$ (note that $bn' = (k/a)^2$), to a $(n/a, k/a, r)$ design, call it $\mathcal{T}^{(n/a)}$.

Let $\mathcal{H}$ be a perfect hash family of functions $h : [n/a] \to [(k/a)^2]$, such that for every k/a-set, $S \subseteq [n/a]$, there is an $h \in \mathcal{H}$ such that $|h(S)| = k/a$. That is, no element of S gets mapped to the same value. By lemma 3 in [Alon et al., 2006], we can efficiently construct such a family of $(k/a)^4 \log(n/a)$ functions.

We will first describe the decoding function $T^{(n/a)} : \mathcal{P}([n/a]) \to \binom{[n/a]}{k/a}$, and then let $\mathcal{T}^{(n/a)} = T^{(n/a)}([n/a])$. For any set $S \subseteq [n/a]$ to be decoded, for each $h \in \mathcal{H}$, we evaluate $T^{(bn')}(h(S))$ to get all r-sets $R \in \mathcal{T}^{(bn')}$ where $R \subseteq h(S)$. For each such set, we return each set in

$$(h^{-1}(R_1) \cap S) \times (h^{-1}(R_2) \cap S) \times \cdots \times (h^{-1}(R_r) \cap S),$$

where R_i is the ith element of R when considered as a $[bn']^r$ vector.

This takes time equal to the size of the above product (the number of results, $|T(S)|$) plus an overhead of $|\mathcal{H}|$ times the time to decode in $\mathcal{T}^{(bn')}$ which is $|\mathcal{H}| r^{O(\sqrt{r})} = e^{\chi}$ by the previous part. The other part of the decoding time, the actual size $|T^{bn'}(h(S))|$, is dominated by the size of the product. To prove the the 'efficient decoding' part of the lemma we thus have to show that the expected size of $T(S)$ is $\leq (|S|a/k)^r e^{\chi}$ for any $S \subseteq [n/a]$. (Note: this is for a set $S \subseteq [n/a]$, part four of the proof will extend to sets $S \subseteq [n]$ and that factor a in the bound will disappear.)

At this point we will add a random permutation, $\pi : [(k/a)^2] \to [(k/a)^2]$, to the preprocessing part of the lemma. This bit of randomness will allow us to consider the perfect hash-family as 'black box' without worrying that it might conspire in a worst case way with our fixed family $\mathcal{T}^{(bn')}$. We apply this permutation to each function of $\mathcal{H}$, so we are actually returning

$$T^{(n/a)}(S) = \bigcup \left\{ \begin{array}{r} (h^{-1}(\pi^{-1}R_1) \cap S) \times (h^{-1}(\pi^{-1}R_2) \cap S) \times \cdots \times (h^{-1}(\pi^{-1}R_r) \cap S) \\ : \text{for all } R \in T^{(bn')}(\pi h(S)) \text{ and } h \in \mathcal{H}. \end{array} \right\}$$

We can then show for any $S \subseteq [n/a]$:

$$\begin{aligned} E_\pi[|T^{(n/a)}(S)|] &= E_\pi\left[\sum_{h\in\mathcal{H},\, R\in T^{(bn')}(\pi h(S))} \left|(h^{-1}(\pi^{-1}R_1) \cap S) \times \cdots \times (h^{-1}(\pi^{-1}R_r) \cap S)\right|\right] \\ &= \sum_{h\in\mathcal{H},\, R\in\mathcal{T}^{(bn')}} E_\pi\left[|(h^{-1}(\pi^{-1}R_1) \cap S)| \cdots |(h^{-1}(\pi^{-1}R_r) \cap S)| \cdot [R \subseteq \pi h(S)]\right] \\ &= |\mathcal{T}^{(bn')}| \sum_{h\in\mathcal{H}} E_\pi\left[|(h^{-1}(\pi^{-1}R_1) \cap S)| \cdots |(h^{-1}(\pi^{-1}R_r) \cap S)|\right] \qquad (1) \\ &\le |\mathcal{T}^{(bn')}| \sum_{h\in\mathcal{H}} E_\pi\left[|h^{-1}(\pi_1) \cap S|\right]^r \qquad (2) \\ &= |\mathcal{T}^{(bn')}|\,|\mathcal{H}|\,(|S|/(k/a)^2)^r \\ &= (|S|a/k)^r e^{\chi}. \end{aligned}$$

For (1) we used that

$$[R \subseteq h(S)] = [\forall_{r\in R} r \in h(S)] = [\forall_{r\in R} h^{-1}(r) \cap S \neq \varnothing] \qquad (3)$$

and so the value in the expectation was already 0 exactly when the Iversonian bracket was zero.

For (2) we used the Maclaurin's Inequality [Ben-Ari and Conrad, 2014] which says that $E(X_1X_2\ldots X_r) \le (EX_1)^r$ when the X_is are values sampled identically, uniformely at random without replacement from som set of non-negative values. In our case those values were sizes of partitions $h^{-1}(1) \cap S, \ldots, h^{-1}(bn') \cap S$, which allowed us to bound the expectation as if h had been a random function.

We need to show that $T^{(n/a)}$ is a correct decoding function, that is $T^{(n/a)}(S) = \{R \in \mathcal{T}^{(n/a)} : R \subseteq S\}$, and the correctness of $\mathcal{T}^{(n/a)}$, that is $|S| \ge k/a$ implies $T^{(n/a)}(S) \neq \varnothing$.

For this, first notice that T is monotone, that is if $S \subseteq U$ then $T(S) \subseteq T(U)$. This follows because $R \subseteq \pi h(S) \implies R \subseteq \pi h(U)$ and that each term $h^{-1}(R_i) \cap S$ is monotone in the size of S. This means we just need to show that $T(K)$ returns something for every $|K| = k$, since then $\mathcal{T} = T([n/a]) = T(\bigcup_K K) \supseteq \bigcup T(K)$ will return all these things.

Hence, consider a k-set, $K \subseteq [n/a]$. By the property of $\mathcal{H}$, there must be some $h \in \pi H$ such that $|h(K)| = k$, and by correctness of $\mathcal{T}^{(bn')}$ we know there is some r-set,

$R \in T^{(bn')}(h(K))$. Now, for these h and R, since $R \subseteq h(K)$ and using (3) we have that $(h^{-1}(R_1) \cap K) \times \ldots$ is non-empty, which is what we wanted. Conversely, consider some $R \in \mathcal{T}^{(n/a)} = T^{(n/a)}([n/a])$ such that $R \subseteq K$, then $R \in h^{-1}(R'_1) \times h^{-1}(R'_2) \ldots$ for some $R' \in \mathcal{T}^{(bn')}$ and $h(R) \subseteq h(K)$. However $h(R)$ is exactly R', since $R_i \in h^{-1}(R'_i) \implies h(R_i) = R'_i$, which shows that $T^{(n/a)}(K)$ returns all the set we want.

4. Finally we convert our $(n/a, k/a, r)$ design, $\mathcal{T}^{(n/a)}$ to an (n, k, r) design, call it $\mathcal{T}$. We do this by choosing a random permutation $\pi : [n] \to [n]$ and given any set $S \subseteq [n]$ we decode it as

$$T(S) = T^{(n/a)}(\pi S \cap \{1, \ldots, n/a\}) \cup \cdots \cup T^{(n/a)}(\pi S \cap \{n - n/a + 1, \ldots, n\}).$$

To see that this is indded an (n, k, r) design, consider any set $K \subseteq [n]$ of size $|K| = k$, there must be one partition $K \cap \{1 \ldots, n/a\}, \ldots$ that has at least the average size k/a. Since $\mathcal{T}^{(n/a)}$ is a $(n/a, k/a, r)$ design, it will contain a set $R \subseteq K \cap \{in - n/a + 1, \ldots, in\} \subseteq K$ which we return.

It remains to analyze the size of $T(S)$, which may of course get large if we are so unlucky that π places much more than the expected number of elements in one of the partitions. In expectation this turns out to not be a problem, as we can see as follows:

$$\begin{aligned}
E_\pi |T(S)| &= \sum_i E_\pi \left[\left| T^{(n/a)}(\pi S \cap p_i) \right| \right] \\
&= a \sum_s E \left[\left| T^{(n/a)}(\pi S \cap p_1) \right| \, \middle| \, |\pi S \cap p_1| = s \right] \Pr[|\pi S \cap p_1| = s] \\
&= a \sum_s (sa/k)^r e^\chi \binom{|S|}{s} \binom{n - |S|}{n/a - s} \Big/ \binom{n}{n/a} \\
&\leq e^\chi \sum_s \frac{\binom{s}{r}}{\binom{k/a}{r}} \frac{\binom{|S|}{s}\binom{n-|S|}{n/a-s}}{\binom{n}{n/a}} \\
&= e^\chi \frac{\binom{|S|}{r}}{\binom{k/a}{r}\binom{n}{n/a}} \sum_s \binom{|S| - r}{s - r} \binom{n - |S|}{n/a - s} \\
&= e^\chi \frac{\binom{|S|}{r}\binom{n-r}{n/a-r}}{\binom{k/a}{r}\binom{n}{n/a}} = e^\chi \frac{\binom{|S|}{r}\binom{n/a}{r}}{\binom{k/a}{r}\binom{n}{r}} \\
&\leq e^\chi (|S|a/k)^r e^{r^2/(k/a)} a^{-r} = (|S|/k)^r e^\chi.
\end{aligned}$$

Here we used Vandermonde convolution to complete the sum over s, and then eventually lemma 6 to bound the binomial ratios. This completes the proof of lemma 5. □

4.3 Integrality considerations

In the proof, we needed the following quantities to be integral: $b = r/b = \sqrt{r}$, $a = kr^{-3/2}\log(r^{3/2})$, $k^2/(a^2b) = k/a = r^{3/2}/\log(r^{3/2})$, $k/(ab) = r/\log(r^{3/2})$, n/a.

It suffices to have $\sqrt{r}$ and $r/\log(r^{3/2})$ integral, and that the later divides k.

It is obviously ridiculous to require that r is a square number. Or is it? You can actually make a number square by just changing it by a factor $1 + 2/\sqrt{r}$. That would only end up giving us an $e^{O(\sqrt{r})}$, so maybe not so bad?

To make $r/\log(r^{3/2})$ integral, we can multiply with a constant. Since it didn't matter that we divided by a log, surely it doesn't matter that we multiply with a constant.

To make $r/\log(r^{3/2})$ divide k, we need k to have some divisors. We can't just round k to say, a power of two, since that could potentially change it by a constant factor, which would come out of $(n/k)^r$. We can change k with at most $1 + 1/\sqrt{r}$. So $1 + 1/\sqrt{k}$ would be just fine. Of course we can change it by an additive $r/\log(r^{3/2})$. That corresponds to a factor about $1 + r/k$. Since $k > r^{3/2}$ that is fine! Or maybe we'll subtract that, because then it is still a valid (n, k, r) design. In the same way, if we round r up to the nearest square root, we don't have to make the changes in the later calculations, but they can be kept intirely inside the lemma.

5 Conclusion and Open Problems

We have seen that, perhaps surprisingly, there exists a relatively general way of creating efficient Las Vegas versions of state-of-the art high-dimensional search data structures.

As bi-products we found efficient, explicit constructions of large Turán systems and covering codes for pairs. We also showed an efficient way to do dimensionality reduction in hamming space without false negatives.

The work leaves a number of open questions for further research:

1) Can we make a completely deterministic high-dimensional data structure for the proposed problems? Cutting the number of random bits used for Las Vegas guarantees would likewise be interesting. The presented algorithms both use $O(d \log d)$ bits, but perhaps limited independence could be used to show that $O(\log d)$ suffice?

2) Does there exist Las Vegas data structures with performance matching that of data-dependent LSH data structures? This might connect to the previous question, since a completely deterministic data structure would likely have to be data-dependent. However the most direct approach would be to repeat [Andoni et al., 2017b], but find Las Vegas constructions for the remaining uses of Monte Carlo randomness, such as clustering.

3) By reductions, our results extend to ℓ_2 and ℓ_1 with exponent $n^{1/c}$. This is optimal for ℓ_1, but for ℓ_2 we would hope to get n^{1/c^2}. Can our techniques be applied to yield such a data structure? Are there other interesting LSH algorithms that can be made Las Vegas using our techniques? The author conjectures that a space/time trade-off version of the presented algorithm should follow easily following the approach of [Andoni et al., 2017b, Laarhoven, 2015, Christiani, 2017].

4) In the most general version, we we get the overhead term $(\log n)^{-1/4}$ in our ρ value. Some previous known LSH data structures also had large terms, such as [Andoni and Indyk, 2006], which had a $(\log n)^{-1/3}$ term and [Andoni et al., 2017b], which has $(\log \log n)^{-\Theta(1)}$, but in general most algorithms have at most a $(\log n)^{-1/2}$ term.

   Can we improve the overhead of the approach given in this paper? Alternatively, is there a completely different approach, that has a smaller overhead?

5.1 Acknowledgements

The author would like to thank Rasmus Pagh, Tobias Christiani and the members of the Scalable Similarity Search project for many rewarding discussions on derandomization and set similarity data structures. Further thanks to Omri Weinstein, Rasmus Pagh, Martin Aumüller, Johan Sivertsen and the anonymous reviewers for useful comments on the manuscript; and to the people at University of Texas, Austin, for hosting me while doing much of this work. An extra thanks to Piotr Wygocki for pointing out the need for a deterministic reduction from ℓ_1 to Hamming space.

The research leading to these results has received funding from the European Research Council under the European Union's 7th Framework Programme (FP7/2007-2013) / ERC grant agreement no. 614331.

References

[Ahle et al., 2017] Ahle, T. D., Aumüller, M., and Pagh, R. (2017). Parameter-free locality sensitive hashing for spherical range reporting. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 239–256. SIAM.

[Alon et al., 2006] Alon, N., Moshkovitz, D., and Safra, S. (2006). Algorithmic construction of sets for k-restrictions. *ACM Transactions on Algorithms (TALG)*, 2(2):153–177.

[Andoni and Indyk, 2006] Andoni, A. and Indyk, P. (2006). Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *Foundations*

of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on, pages 459–468. IEEE.

[Andoni et al., 2015] Andoni, A., Indyk, P., Laarhoven, T., Razenshteyn, I., and Schmidt, L. (2015). Practical and optimal lsh for angular distance. In *Advances in Neural Information Processing Systems*, pages 1225–1233.

[Andoni et al., 2014] Andoni, A., Indyk, P., Nguyen, H. L., and Razenshteyn, I. (2014). Beyond locality-sensitive hashing. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1018–1028. Society for Industrial and Applied Mathematics.

[Andoni et al., 2017a] Andoni, A., Laarhoven, T., Razenshteyn, I., and Waingarten, E. (2017a). Optimal hashing-based time-space trade-offs for approximate near neighbors. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 47–66. SIAM.

[Andoni et al., 2017b] Andoni, A., Laarhoven, T., Razenshteyn, I., and Waingarten, E. (2017b). Optimal hashing-based time-space trade-offs for approximate near neighbors. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 47–66. Society for Industrial and Applied Mathematics.

[Andoni and Razenshteyn, 2015] Andoni, A. and Razenshteyn, I. (2015). Optimal data-dependent hashing for approximate near neighbors. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, pages 793–801. ACM.

[Arasu et al., 2006] Arasu, A., Ganti, V., and Kaushik, R. (2006). Efficient exact set-similarity joins. In *Proceedings of the 32nd international conference on Very large data bases*, pages 918–929. VLDB Endowment.

[Arya et al., 1998] Arya, S., Mount, D. M., Netanyahu, N. S., Silverman, R., and Wu, A. Y. (1998). An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM (JACM)*, 45(6):891–923.

[Aumüller et al., 2017] Aumüller, M., Christiani, T., Pagh, R., and Silvestri, F. (2017). Distance-sensitive hashing. *arXiv preprint arXiv:1703.07867*.

[Becker et al., 2016] Becker, A., Ducas, L., Gama, N., and Laarhoven, T. (2016). New directions in nearest neighbor searching with applications to lattice sieving. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 10–24. SIAM.

[Ben-Ari and Conrad, 2014] Ben-Ari, I. and Conrad, K. (2014). Maclaurin's inequality and a generalized bernoulli inequality. *Mathematics Magazine*, 87(1):14–24.

[Bentkus, 2005] Bentkus, V. (2005). A lyapunov-type bound in rd. *Theory of Probability & Its Applications*, 49(2):311–323.

[Bentley, 1975] Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517.

[Broder, 1997] Broder, A. Z. (1997). On the resemblance and containment of documents. In *Compression and Complexity of Sequences 1997. Proceedings*, pages 21–29. IEEE.

[Broder et al., 1997] Broder, A. Z., Glassman, S. C., Manasse, M. S., and Zweig, G. (1997). Syntactic clustering of the web. *Computer Networks and ISDN Systems*, 29(8-13):1157–1166.

[Charikar, 2002] Charikar, M. S. (2002). Similarity estimation techniques from rounding algorithms. In *Proceedings of the thiry-fourth annual ACM Symposium on Theory of Computing*, pages 380–388. ACM.

[Christiani, 2017] Christiani, T. (2017). A framework for similarity search with space-time tradeoffs using locality-sensitive filtering. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 31–46. SIAM.

[Colbourn and Dinitz, 2006] Colbourn, C. J. and Dinitz, J. H. (2006). *Handbook of combinatorial designs*. CRC press.

[Cole et al., 2004] Cole, R., Gottlieb, L.-A., and Lewenstein, M. (2004). Dictionary matching and indexing with errors and don't cares. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 91–100. ACM.

[Datar et al., 2004] Datar, M., Immorlica, N., Indyk, P., and Mirrokni, V. S. (2004). Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*, pages 253–262. ACM.

[Dubiner, 2010] Dubiner, M. (2010). Bucketing coding and information theory for the statistical high-dimensional nearest-neighbor problem. *IEEE Transactions on Information Theory*, 56(8):4166–4179.

[Elias, 1957] Elias, P. (1957). List decoding for noisy channels. In *1957-IRE WESCON Convention Record, Pt.* Citeseer.

[Gionis et al., 1999] Gionis, A., Indyk, P., and Motwani, R. (1999). Similarity search in high dimensions via hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases*, pages 518–529. Morgan Kaufmann Publishers Inc.

[Goswami et al., 2017] Goswami, M., Pagh, R., Silvestri, F., and Sivertsen, J. (2017). Distance sensitive bloom filters without false negatives. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 257–269. SIAM.

[Har-Peled et al., 2012] Har-Peled, S., Indyk, P., and Motwani, R. (2012). Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory of computing*, 8(1):321–350.

[Hoeffding, 1963] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30.

[Indyk, 2000a] Indyk, P. (2000a). Dimensionality reduction techniques for proximity problems. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 371–378. Society for Industrial and Applied Mathematics.

[Indyk, 2000b] Indyk, P. (2000b). *High-dimensional computational geometry*. PhD thesis, Stanford University.

[Indyk, 2001] Indyk, P. (2001). On approximate nearest neighbors under l_∞ norm. *Journal of Computer and System Sciences*, 63(4):627–638.

[Indyk, 2007] Indyk, P. (2007). Uncertainty principles, extractors, and explicit embeddings of l2 into l1. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 615–620. ACM.

[Indyk and Motwani, 1998] Indyk, P. and Motwani, R. (1998). Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613. ACM.

[Karppa et al., 2016] Karppa, M., Kaski, P., Kohonen, J., and Catháin, P. Ó. (2016). Explicit correlation amplifiers for finding outlier correlations in deterministic subquadratic time. *Proceedings of the 24th European Symposium Of Algorithms*.

[Kushilevitz et al., 2000] Kushilevitz, E., Ostrovsky, R., and Rabani, Y. (2000). Efficient search for approximate nearest neighbor in high dimensional spaces. *SIAM Journal on Computing*, 30(2):457–474.

[Laarhoven, 2015] Laarhoven, T. (2015). Tradeoffs for nearest neighbors on the sphere. *arXiv preprint arXiv:1511.07527*.

[Lv et al., 2007] Lv, Q., Josephson, W., Wang, Z., Charikar, M., and Li, K. (2007). Multi-probe lsh: efficient indexing for high-dimensional similarity search. In *Proceedings of the 33rd International Conference on Very Large Data Bases*, pages 950–961. VLDB Endowment.

[Mairson, 1983] Mairson, H. G. (1983). The program complexity of searching a table. In *Foundations of Computer Science, 1983., 24th Annual Symposium on*, pages 40–47. IEEE.

[Mitzenmacher and Upfal, 2005] Mitzenmacher, M. and Upfal, E. (2005). *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge university press.

[Naor and Naor, 1993] Naor, J. and Naor, M. (1993). Small-bias probability spaces: Efficient constructions and applications. *SIAM journal on computing*, 22(4):838–856.

[Naor et al., 1995] Naor, M., Schulman, L. J., and Srinivasan, A. (1995). Splitters and near-optimal derandomization. In *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*, pages 182–191. IEEE.

[Overmars and van Leeuwen, 1981] Overmars, M. H. and van Leeuwen, J. (1981). Worst-case optimal insertion and deletion methods for decomposable searching problems. *Information Processing Letters*, 12(4):168–173.

[O'Donnell et al., 2014] O'Donnell, R., Wu, Y., and Zhou, Y. (2014). Optimal lower bounds for locality-sensitive hashing (except when q is tiny). *ACM Transactions on Computation Theory (TOCT)*, 6(1):5.

[Pacuk et al., 2016] Pacuk, A., Sankowski, P., Wegrzycki, K., and Wygocki, P. (2016). Locality-sensitive hashing without false negatives for lp. In *International Computing and Combinatorics Conference*, pages 105–118. Springer.

[Pagh, 2016] Pagh, R. (2016). Locality-sensitive hashing without false negatives. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1–9. SIAM.

[Pagh and Christiani, 2017] Pagh, R. and Christiani, T. (2017). Beyond minhash for similarity search. Proceedings of the forty-ninth annual ACM symposium on Theory of computing.

[Panigrahy, 2006] Panigrahy, R. (2006). Entropy based nearest neighbor search in high dimensions. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1186–1195. Society for Industrial and Applied Mathematics.

[Pham and Pagh, 2016] Pham, N. and Pagh, R. (2016). Scalability and total recall with fast coveringlsh. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 1109–1118. ACM.

[Sidorenko, 1995] Sidorenko, A. (1995). What we know and what we do not know about turán numbers. *Graphs and Combinatorics*, 11(2):179–199.

[Topsøe, 2006] Topsøe, F. (2006). Some bounds for the logarithmic function. *Inequality theory and applications*, 4:137.

[Turán, 1961] Turán, P. (1961). Research problems. *Közl MTA Mat. Kutató Int*, 6:417–423.

[Williams, 2005] Williams, R. (2005). A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2):357–365.

6 Appendix

6.1 Explicit reduction from ℓ_1 to Hamming

Theorem 3. *For $d, r, c \geq 1$ and a set of points $P \subseteq \mathbb{R}^d$ of size $|P| = n$, there is a function $f : \mathbb{R}^d \to \{0,1\}^b$ where $b = 2d^2\epsilon^{-3}\log n$, such that for any two points $x \in \mathbb{R}^d, y \in P$,*

1. *if $\|x - y\|_1 \leq r$ then $\|f(x) - f(y)\|_1 \leq (1+\epsilon)S$,*
2. *if $\|x - y\|_1 \geq cr$ then $\|f(x) - f(y)\|_1 \geq (1-\epsilon)cS$,*

and the scale factor $S = b/(2dcr) = (d\log n)/(cr\epsilon^3)$.

Proof. First notice that we can assume all coordinates are at most rn. This can be assured by imposing a grid of side length $2rn$ over the points of P, such that no point is closer than distance r from a cell boundary. Since points $x, y \in \mathbb{R}^d$ in different cells must now be more than distance r from each other, we can set the embedded distance to cS by prefixing points from different cells with a different code word. The grid can be easily found in time $O(dn)$ by sweeping over each dimension seperately.

Notice that for actual data structure purposes, we can even just process each cell seperately and don't have to worry about separating them.

By splitting up each coordinate into positive and negative parts, we can further assume each coordinate of each vector is positive. This costs a factor of 2 in d.

Next, if we have an $2\epsilon r/d$ grid, then there is always a grid point within ℓ_1-distance ϵr. That means if we multiply each coordinate by $d/(2\epsilon r)$ and round the coordinates to nearest integer, we get distances are changed by at most ϵr.

We are now ready for the main trick of the reduction. Let M be the largest coordinate, which we can assume is at most dn/ϵ, and $R = dc/(2\epsilon)$ be the value of cr after scaling and rounding. For each coordinate we map $[M] \to [\lfloor M/R \rfloor]^R$ by $h(c) := \langle \lfloor \frac{x}{R} \rfloor, \lfloor \frac{x+1}{R} \rfloor, \ldots, \lfloor \frac{x+R-1}{R} \rfloor \rangle$. The point of this mapping is that for every $c_1, c_2 \in [M]$, $\mathrm{dist}(h(c_1), h(c_2)) = \min(|c_1 - c_2|, R)$, where dist is hamming distance.

100	12	12	12	12	13	13	13	13
105	13	13	13	13	13	13	13	14
	*	*	*	*				*

Figure 1: Mapping 100 and 105 to $[\lfloor 100/8 \rfloor]^8$ preserving ℓ_1 distance in Hamming distance.

All that's left is to use a code with good minimum and maximum distance to map down into $\{0,1\}$. A random code with bit length $k = 4\epsilon^{-2}(\log 4n)$ suffices. To see this, let X be a binomial random variable, $X \sim B(k, 1/2)$. Then

$$\Pr[(1-\epsilon)k/2 \le C \le (1+\epsilon)k/2] \le 2e^{-\epsilon^2 k/2} \le 1/(8n^2)$$

so by union bound over all $\binom{M/R}{2} \le 2n^2$ pairs of values, we have constant probability that the code works. For a given code, we can check this property deterministically in time kn^2, so we can use rejection sampling and generate the code in time $\approx O(n^2)$. Of course, n^2 time may be too much. Luckily there are also explicit codes with the property, such as those by Naor and Naor [Naor and Naor, 1993].

The complete construction follows by concatenating the result of h on all coordinates. □

See [Indyk, 2007] for an explicit reduction from ℓ_2 to ℓ_1.

6.2 The Ratio of Two Binomial Coefficients

Classical bounds for the binomial coefficient: $(n/k)^k \le \binom{n}{k} \le (en/k)^k$ give us simple bounds for binomial ratios, when $n \ge m$: $(n/em)^k \le \binom{n}{k}\Big/\binom{m}{k} \le (en/m)^k$. The factor e on both sides can often be a nuisance.

Luckily tighter analysis show, that they can nearly always be either removed or reduced. Using the fact that $\frac{n-i}{m-i}$ is increasing in i for $n \ge m$, we can show $\binom{n}{k}\Big/\binom{m}{k} = \prod_{i=0}^{k-1} \frac{n-i}{m-i} \ge \prod_{i=0}^{k-1} \frac{n}{m} = \left(\frac{n}{m}\right)^k$. This is often sharp enough, but on the upper bound side, we need to work harder to get results.

Let $\mathrm{H}(x) = x \log 1/x + (1-x) \log 1/(1-x)$ be the binary entropy function,

Lemma 6. *For $n \ge m \ge k \ge 0$ we have the following bounds:*

$$\left(\frac{n}{m}\right)^k \le \left(\frac{n}{m}\right)^k \exp\left(\frac{n-m}{nm}\frac{k(k-1)}{2}\right) \le \binom{n}{k}\Big/\binom{m}{k} \le \frac{\exp\left(n\,\mathrm{H}(k/n)\right)}{\exp\left(m\,\mathrm{H}(k/m)\right)} \le \left(\frac{n}{m}\right)^k e^{k^2/m}$$

If $m \geq n$ we can simply flip the inequalities and swap n for m. Note that $(n/em)^k \leq (n/m)^k$ and $e^{k^2/m} \leq e^k$, so the bounds strictly tighten the simple bounds states above.

Especially the entropy bound is quite sharp, since we can also show: $\binom{n}{k}/\binom{m}{k} \geq \frac{\exp((n+1)\,\mathrm{H}(k/(n+1)))}{\exp((m+1)\,\mathrm{H}(k/(m+1)))}$, though for very small values of k, the lower bound in the theorem is actually even better. We can also get a feeling for the sharpness of the bounds, by considering the series expansion of the entropy bound at $k/m \to 0$: $\frac{\exp(n\,\mathrm{H}(k/n))}{\exp(m\,\mathrm{H}(k/m))} = \left(\frac{n}{m}\right)^k \exp(\frac{n-m}{nm}\frac{k^2}{2} + O(k^3/m^2))$.

For the proofs, we'll use some inequalities on the logarithmic function from [Topsøe, 2006]:

$$\log(1+x) \geq x/(1+x) \tag{4}$$
$$\log(1+x) \geq 2x/(2+x) \text{ for } x \geq 0 \tag{5}$$
$$\log(1+x) \leq x(2+x)/(2+2x) \text{ for } x \geq 0. \tag{6}$$

In particular (5) and (6) imply the following bounds for the entropy function:

$$\mathrm{H}(x) \leq x \log 1/x + x(2-x)/2 \tag{7}$$
$$\mathrm{H}(x) \geq x \log 1/x + 2x(1-x)/(2-x), \tag{8}$$

which are quite good for small x.

We'll prove theorem 6 one inequality at a time, starting from the left most:

Proof. The first inequality follows simply from $\frac{n-m}{nm}\frac{k(k-1)}{2} \geq 0$, which is clear from the conditions on $n \geq m \geq k$.

The second inequality we prove by using (4), which implies $1+x \geq \exp(x/(1+x))$, to turn the product into a sum:

$$\begin{aligned}
\binom{n}{k}\Big/\binom{m}{k} &= \prod_{i=0}^{k-1} \frac{n-i}{m-i} \\
&= \left(\frac{n}{m}\right)^k \prod_{i=0}^{k-1} \frac{1-i/n}{1-i/m} \\
&= \left(\frac{n}{m}\right)^k \prod_{i=0}^{k-1} \left(1 + \frac{i/m - i/n}{1-i/m}\right) \\
&\geq \left(\frac{n}{m}\right)^k \exp\left(\sum_{i=0}^{k-1} \frac{i(n-m)}{(n-i)m}\right) \\
&\geq \left(\frac{n}{m}\right)^k \exp\left(\sum_{i=0}^{k-1} i\frac{n-m}{nm}\right) \\
&= \left(\frac{n}{m}\right)^k \exp\left(\frac{k(k-1)}{2}\frac{n-m}{nm}\right).
\end{aligned}$$

For the entropy upper bound we will use an integration bound, integrating $\log(n-i)/(m-i)$ by parts:

$$\begin{aligned}
\binom{n}{k}\Big/\binom{m}{k} &= \prod_{i=0}^{k-1}\frac{n-i}{m-i} \\
&= \exp\left(\sum_{i=0}^{k-1}\log\frac{n-i}{m-i}\right) \\
&\le \exp\left(\int_0^k 1\log\frac{n-x}{m-x}dx\right) \\
&= \exp\left(x\log\frac{n-x}{m-x}\Big|_0^k - \int_0^k x\left(\frac{1}{m-x}-\frac{1}{n-x}\right)dx\right) \\
&= \exp\left(k\log\frac{n-k}{m-k} + \int_0^k\left(\frac{m}{m-x}-\frac{n}{n-x}\right)dx\right) \\
&= \exp\left(k\log\frac{n-k}{m-k} - \left|m\log\frac{1}{m-x} - n\log\frac{1}{n-x}\right|_0^k\right) \\
&= \exp\left(n\,\mathrm{H}(k/n) - m\,\mathrm{H}(k/m)\right).
\end{aligned}$$

The integral bound holds because $\log\frac{n-i}{m-i}$ is increasing in i, and so $\log\frac{n-i}{m-i} \le \int_i^{i+1}\log\frac{n-x}{m-x}dx$. We see that $\frac{n-i}{m-i}$ is increasing by observing $\frac{n-i}{m-i} = \frac{n}{m} + \frac{in/m-i}{m-i}$ where the numerator and denominator of the last fraction are both positive. The entropy lower bound, mentioned in the discussion after the theorem, follows similarly from integration, using $\log\frac{n-i}{m-i} \ge \int_{i-1}^{i}\log\frac{n-x}{m-x}dx$.

For the final upper bound, we use the bounds (7) and (8) on $\mathrm{H}(k/n)$ and $\mathrm{H}(k/m)$ respectively:

$$\frac{\exp\left(n\,\mathrm{H}(k/n)\right)}{\exp\left(m\,\mathrm{H}(k/m)\right)} \le \left(\frac{n}{m}\right)^k \exp\left(\frac{k^2}{2}\left(\frac{1}{m-k/2}-\frac{1}{n}\right)\right) \le \left(\frac{n}{m}\right)^k \exp\left(\frac{k^2}{m}\right).$$

□

6.3 Proof of lemma 2

Recall the lemma:

Lemma 2. *For $t = \frac{d}{2} - \frac{s\sqrt{d}}{2}$, $1 \le s \le d^{1/4}/2$ and $r < d/2$, Let $x, y \in \{0,1\}^d$ be two points at distance $\mathrm{dist}(x,y) = r$, and let $I = |\{z \in \{0,1\}^d : \mathrm{dist}(z,x) \le t, \mathrm{dist}(z,y) \le t\}|$ be the size of the intersection of two hamming balls around x and y of radius t, then*

$$\frac{7}{8d}\exp\left(\tfrac{-s^2}{2(1-r/d)}\right) \le I\,2^{-d} \le \exp\left(\tfrac{-s^2}{2(1-r/d)}\right)$$

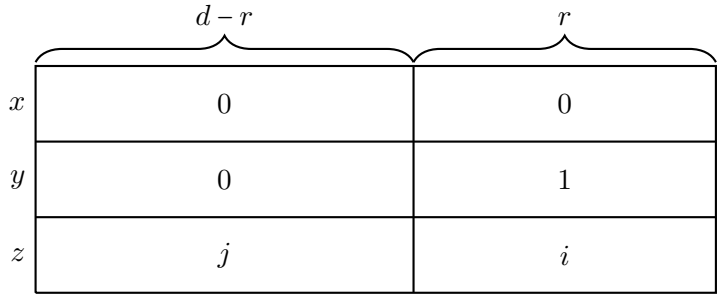

Figure 2: To calculate how many points are within distance t from two points x and y, we consider without loss of generality $x = 0 \dots 0$. For a point, z, lying in the desired region, we let i specify the number of 1's where x and y differ, and j the number of 1's where they are equal. With this notation we get $d(x,z) = i + j$ and $d(y,z) = j + r - i$.

Proof. From figure 2 we have that $I = \sum_{\substack{i+j\le t \\ j-i\le t-r}} \binom{r}{i}\binom{d-r}{j}$, and from monotonicity (and figure 3) it is clear that $\binom{r}{r/2}\binom{d-r}{t-r/2} \le I \le \sum_{\substack{0\le i\le r \\ 0\le j\le t-r/2}} \binom{r}{i}\binom{d-r}{j}$.

We expand the binomials using Stirling's approximation: $\frac{\exp(nH(k/n))}{\sqrt{8(1-k/n)k}} \le \binom{n}{k} \le \sum_{i\le k}\binom{n}{i} \le \exp(nH(k/n))$ where $H(x) = x\log\frac{1}{x} + (1-x)\log\frac{1}{1-x}$ is the binary entropy function, which we bound as $\log 2 - 2(\frac{1}{2}-x)^2 - 4(\frac{1}{2}-x)^4 \le H(x) \le \log 2 - 2(\frac{1}{2}-x)^2$. We then have for the upper bound:

$$I2^{-d} \le 2^{r-d}\exp\Big[(d-r)H\Big(\tfrac{t-r/2}{d-r}\Big)\Big] \le \exp\Big[-\tfrac{s^2}{2(1-r/d)}\Big]$$

And for the lower bound:

$$\begin{aligned} I2^{-d} \ge 2^{-d}\binom{r}{r/2}\binom{d-r}{t-r/2} &\ge \frac{2^{r-d}}{\sqrt{2r}}\frac{\exp\Big[(d-r)H\Big(\frac{t-r/2}{d-r} - \log 2\Big)\Big]}{\sqrt{8(1-\frac{t-r/2}{d-r})(t-r/2)}} \\ &\ge \exp[-\tfrac{s^2}{2(1-r/d)}]\frac{\exp[-\frac{s^4}{4(1-r/d)^3 d}]}{\sqrt{4r(d-r)(1-\frac{ds^2}{(d-r)^2})}} \\ &\ge \exp[-\tfrac{s^2}{2(1-r/d)}]\frac{1}{d}\frac{1-2s^4/d}{\sqrt{1-4s^2/d}}, \end{aligned}$$

where for the last inequality we used the bound $e^x \ge 1 + x$. The last factor is monotone in s and we see that for $s \le d^{1/4}/2$ it is $\ge \frac{7}{8}\left(1 - 1/\sqrt{d}\right)^{-1/2} \ge \frac{7}{8}$, which gives the theorem. □

The factor of $1/d$ can be sharpened a bit, e.g. by using the two dimensional Berry-Essen theorem from [Bentkus, 2005].

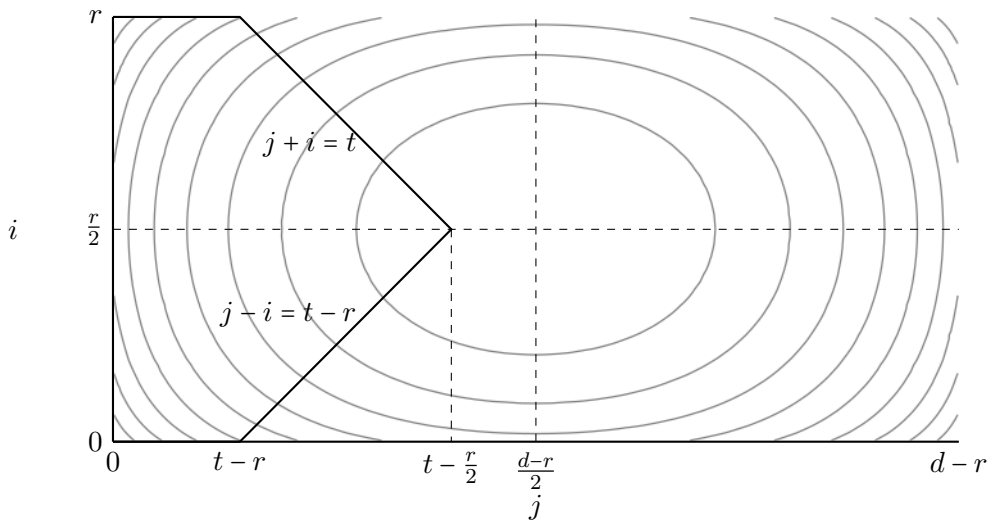

Figure 3: A contour plot over the two dimensional binomial. The pentagon on the left marks the region over which we want to sum. For the upper bound we sum i from 0 to r and j from 0 to $t - r/2$.