



POLITECNICO DI TORINO
Repository ISTITUZIONALE

Analysis of SparseHash: An efficient embedding of set-similarity via sparse projections

Original

Analysis of SparseHash: An efficient embedding of set-similarity via sparse projections / Valsesia, Diego; Fosson, Sophie M.; Ravazzi, Chiara; Bianchi, Tiziano; Magli, Enrico. - In: PATTERN RECOGNITION LETTERS. - ISSN 0167-8655. - 128(2019), pp. 93-99. [10.1016/j.patrec.2019.08.014]

Availability:

This version is available at: 11583/2750152 since: 2019-09-06T12:04:00Z

Publisher:

Elsevier

Published

DOI:10.1016/j.patrec.2019.08.014

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2019. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.patrec.2019.08.014>

(Article begins on next page)

Analysis of SparseHash: an efficient embedding of set-similarity via sparse projections

D. Valsesia, S. M. Fosson, C. Ravazzi, T. Bianchi, E. Magli

September 5, 2019

Abstract

Embeddings provide compact representations of signals in order to perform efficient inference in a wide variety of tasks. In particular, random projections are common tools to construct Euclidean distance-preserving embeddings, while hashing techniques are extensively used to embed set-similarity metrics, such as the Jaccard coefficient. In this letter, we theoretically prove that a class of random projections based on sparse matrices, called SparseHash, can preserve the Jaccard coefficient between the supports of sparse signals, which can be used to estimate set similarities. Moreover, besides the analysis, we provide an efficient implementation and we test the performance in several numerical experiments, both on synthetic and real datasets.

1 Introduction

Retrieving meaningful information from large amounts of data is a complex task, often impossible if those data have to be analyzed in their original domain. For this motivation, compact representations have been studied in different frameworks, ranging from information retrieval (see, e.g., [1, 2, 36]) to signal processing (see, e.g., [16, 10]).

In information retrieval, hash functions are widely used to map data into compact representations, see, e.g., [39] and references therein. Traditional hash functions compress arbitrary data to fixed length representations, preserve exact matches, and minimize collisions between different objects. An important purpose of hashing techniques is the evaluation of the similarity between sets of generic objects. This has many applications in information retrieval. An example is the search of near-duplicate documents: this can be performed by finding the number of bag-of-words shared by different documents, which are said to be near-duplicate if this number overcomes a given threshold. A popular technique to measure set similarity is the min-wise hashing (also known as MinHash) proposed by [7], and further analysed by [8, 20, 15, 13, 33]. MinHash approximately preserves the Jaccard coefficient, which is a popular similarity

metric for bag-of-words and similar representations, and is used in a wide range of applications, see, e.g., [26].

In signal processing, compact signal representations are usually referred to as embeddings, see [22, 5, 38]. Formally, an embedding is a transformation that maps a set of signals in a high dimensional space to a set in a lower dimensional space, in such a way that the geometry of the set is approximately preserved. The most famous embedding is probably that proposed by [24], which preserves Euclidean distances using random projections.

Hash functions and embeddings bear many similarities. For example, a class of efficient indexing techniques known as locality sensitive hashing (LSH) can be constructed using both traditional hash functions and embeddings, based on random projections, as studied by [2]. Hence, it is not surprising that results in one field can be exploited to obtain significant advancements in the other field, and vice versa.

A novel embedding for the Jaccard coefficient, called SparseHash, was proposed by [37]. SparseHash builds on the concept of sparsity. A signal $u \in \mathbb{R}^n$ is said to be sparse if it has few non-zero components. The index set of its non-zero components is called *support*. SparseHash can efficiently evaluate the similarity between sparse signals, in terms of support overlap. Its rationale is based on recent results by [3, 31, 32], which show that the sparsity level of a signal, that is, the size of its support, can be efficiently estimated from compressed linear projections obtained through sparse random matrices.

Sparsity is envisaged also in information retrieval, as several compact representations, e.g., bag-of-words, produce sparse features, see, e.g. [19]. Considering the example of near-duplicate documents, typically each document contains only few bag-of-words with respect to the general vocabulary. Given this observation, we can highlight a duality between sparse signals and sets of generic objects: a set $S \subseteq \Omega = \{1, \dots, n\}$ can be represented as a signal $u \in \{0, 1\}^n$, whose entries are $u_i = 1$, if $i \in S$, and $u_i = 0$, otherwise, and generally we expect that the number k of ones is much smaller than n . Similarity between sets can then be interpreted as the overlap between the supports of sparse signals. In this perspective, SparseHash is a natural alternative to MinHash, which can be considered as benchmark for this kind of problems.

This letter extends our preliminary work [37]. We propose two metrics to measure similarity in the embedded domain that depend linearly and nonlinearly, respectively, on the original Jaccard coefficient. Moreover, we present a deeper theoretical analysis of performance in terms of estimation of set similarities. We also introduce a new algorithm that implements SparseHash more efficiently, i.e., with the same asymptotic complexity of the MinHash bottom sketch by [14]. However, compared to the bottom sketch, SparseHash has binary measurements rather than real-valued ones, yielding improved compression efficiency.

The letter is organized as follows. In Section 2, we introduce MinHash and random projections. In Section 3, we illustrate SparseHash and discuss its implementation. In Section 4, we provide the theoretical analysis. Section 5 is devoted to numerical experiments; conclusions are reported in Section 6.

2 Related work

MinHash is a hashing method conceived to preserve the Jaccard coefficient. Specifically, MinHash generates $m \geq 1$ independent hash functions h_1, \dots, h_m which return an integer value for each element of the original set S ; then it selects the m minimum values $\min_{u \in S} h_1(u), \dots, \min_{u \in S} h_m(u)$. Different variants of MinHash have been devised to obtain more compact storage or lower computational complexity of the hashing operation. In particular, we mention (a) b -bit minwise hashing by [26], which quantizes the hashes over b bits instead of using integers, and (b) bottom- m sketch by [14], which selects the m smallest values of a single hash function instead of the minima of m independent hash functions; this conceptually substitutes the “sampling with replacement” operation performed by MinHash with a “sampling without replacement”. Clearly, the bottom- m sketch has a lower complexity in computing the hashes.

We mention that techniques have been proposed that aim at extending MinHash beyond Jaccard similarity, by assigning weights to the elements of the set, see [40, 21]. This was motivated by bag-of-words representations where the count associated to each element may carry additional information. In the context of hashing functions, BitShred was also proposed by [23], however it achieves substantially biased estimates of the Jaccard coefficient when compared to MinHash. Besides, Bloom filters have been proposed by [4], which are space-efficient representations for set membership queries. They essentially hash each element of a set into positions of a bit array, and they could also be used to estimate the size of the intersection and union between sets. Bloom filters present some challenges like false positive errors, i.e., wrongly indicating that a non-member element is a member of the set (see [28]); a comparison with them is postponed to future work.

Concerning random projections, one of the most famous methods for dimensionality reduction has been introduced by [24], which preserves Euclidean distances. Several extensions have been later proposed, that embed the angle between signals (see [12, 22]) or control the maximum distance that is embedded (see [6]). Finally, we notice that sparse random matrices have received some attention for embedding ℓ_2 or ℓ_1 distances in [25].

3 SparseHash

In this section, we illustrate SparseHash. In particular, we describe how to efficiently implement it, without generating the projection matrix, and we propose a novel faster approximated implementation. In the following, we denote the support of $u \in \mathbb{R}^n$ as $\text{supp}(u) = \{i \in \{1, \dots, n\} : u_i \neq 0\}$; k is the sparsity level, that is, the cardinality of $\text{supp}(u)$, and $\Sigma_k := \{u \in \mathbb{R}^n : |\text{supp}(u)| \leq k\}$.

We call random projection an algorithm that projects a vector $u \in \mathbb{R}^n$ onto a lower-dimensional subspace \mathbb{R}^m by multiplying it by a random matrix $A \in \mathbb{R}^{m \times n}$, $m < n$, see [16]. The obtained vector $y = Au$ is referred to as measurement, and \mathbb{R}^m is known as reduced space. The intuitive idea is that

a properly designed random mapping projects data points onto a randomly selected subspace approximately preserving distances. Generally, dense random matrices are considered in the literature.

SparseHash consists of computing binary-quantized sparse random projections $y = |\text{sign}(Au)|$, where $A \in \mathbb{R}^{m \times n}$ is a γ -sparsified random matrix, defined as follows: with probability $1 - \gamma$, $A_{ij} = 0$; with probability γ , $A_{i,j}$ is generated according to an arbitrary continuous distribution with zero mean and finite variance. γ -sparsified random matrices are efficient for sparsity estimation, as proven by [31]; their use is the core of SparseHash, and is the basis to provide a rigorous analysis of its efficiency. SparseHash can be used to evaluate the similarity of sets of generic objects because, as already mentioned, a set $S \subseteq \Omega = \{1, \dots, n\}$ that contains $k \ll n$ objects can be represented by a sparse signal $u \in \{0, 1\}^n$, and set similarity can be interpreted as support overlap.

3.1 Computation via hashing

From a practical viewpoint, computing the measurements in SparseHash does not require to explicitly generate A and perform the matrix-vector product. An efficient implementation is possible by using hash functions, see [11], which map an index in S to a uniformly distributed value over the output range of the hash function (e.g., $b = 64$ bits integers). Let f_i be a hash function that maps its input to an integer in the range $[0, 2^b - 1]$ and define a threshold τ , where γ is as defined in the previous section. A measurement y_i is zero if and only if the hash function returns a value below τ for all the indexes in the support. By defining $\tau = \gamma(2^b - 1)$, measurements have the same probability to be zero as in the formulation $y = |\text{sign}(Au)|$ illustrated in the previous paragraph.

Multiple hash functions are used to generate m measurements. This typically involves randomizing a seed of the hash function. Algorithm 1 summarizes all the steps required to generate the SparseHash measurements y .

Algorithm 1 Computing SparseHash

Inputs:

$$\gamma, S = \text{supp}(u) = \{s_j\}_{j=1}^k$$

Initialize:

$$\tau \leftarrow \gamma(2^b - 1); \quad y_i \leftarrow 0, i = 1, \dots, m$$

for $i = 1, \dots, m$ **do**

for $j = 1, \dots, k$ **do**

$$h_{ij} \leftarrow f_i(s_j)$$

if $h_{ij} < \tau$ **then**

$$y_i \leftarrow 1$$

break

end if

end for

end for

Algorithm 1 is equivalent to computing $y = |\text{sign}(Au)|$ in the sense that it approximately yields the same probability to get a nonzero measurement y_i

as function of the size of the support. The equivalence would be exact if the hash function could generate output values that are truly uniformly distributed and whose range is large enough that the quantization of probabilities is negligible. However, popular functions, e.g., [29], are designed to be as uniform as possible and work using 32 bits or more as output range, which is large enough to be a good approximation. Finally, the procedure is repeatable, i.e., an hash function returns the same output to the same input.




Figure 1: Example of Fast SparseHash: a binary tree is generated from $m = 7$ random windows, denoted as A to G. The membership of the $k = 6$ hashes in the windows are determined by traversing the tree. The corresponding SparseHash measurements are 1101010. The tree structure must have a RIGHT and a LEFT pointers, and an integer MEASINDEX in $[0, m - 1]$ storing the index of the measurement.

3.2 Fast SparseHash

The technique to implement SparseHash just described requires to compute $\mathcal{O}(km)$ hashes (by hash, we mean the output of the hash function applied to a single entry of the support), and perform $\mathcal{O}(km)$ comparisons with the threshold τ . The complexity is thus equivalent to that of MinHash, which computes $\mathcal{O}(km)$ hashes and performs $\mathcal{O}(km)$ operations to find the m minima. However, there are variants of MinHash, in particular the bottom- m variant by [14], which reduce the complexity. SparseHash can be modified so as to require only $\mathcal{O}(k)$ hashes and $\mathcal{O}(k \log m)$ comparisons. We now illustrate this variant, that we call Fast SparseHash. The main idea behind Algorithm 2 is to compute only one hash per support entry and to check if it falls inside one of m randomly drawn windows of width τ instead of falling below a fixed threshold. A single hash is computed for each of the k elements in the support and m random windows of width $\tau = \gamma(2^b - 1)$ are drawn over the output range of the hash

Algorithm 2 Computing Fast SparseHash

Inputs: $\gamma, S = \text{supp}(u) = \{s_j\}_{j=1}^k$ **Initialize:** $\tau \leftarrow \gamma(2^b - 1); y_i \leftarrow 0, i = 1, \dots, m$ **for** $j = 1, \dots, |S|$ **do** $h_j \leftarrow f(s_j)$

▷ Compute hashes

end for**for** $i = 1, \dots, m$ **do**

Generate random bot[i]

▷ Generate bottom values

end for

Sort bot

head \leftarrow buildTree(bot)**for** $j = 1, \dots, k$ **do**ptr \leftarrow head**while** ptr \neq NULL **do****if** $h_j <$ ptr.bottomValue **then**ptr \leftarrow ptr.left**else****if** $h_j <$ ptr.topValue **then** $i \leftarrow$ ptr.measIndex; $y_i \leftarrow 1$ **if** ptr.measIndex \neq $m - 1$ **then**

▷ Check right

 $p \leftarrow$ ptr.measIndex + 1**while** $p < m$ and $h_j \geq$ bot[p] **do** $y_p \leftarrow 1; p \leftarrow p + 1$ **end while****end if****if** ptr.measIndex \neq 0 **then**

▷ Check left

 $p \leftarrow$ ptr.measIndex - 1**while** $p \geq 0$ and $h_j <$ bot[p] + τ **do** $y_p \leftarrow 1; p \leftarrow p - 1$ **end while****end if****break**

▷ No more measurements can collide

elseptr \leftarrow ptr.right**end if****end if****end while****end for**

Table 1: Fast SparseHash v. SparseHash - Runtime (sec.)

		10^4	$\overset{m}{10^5}$	10^6
k	10^4	0.007 / 1.337	0.040 / 12.36	0.194 / 118.8
	10^5	0.023 / 11.26	0.061 / 115.5	0.280 / 1119
	10^6	0.136 / 126.8	0.205 / 1163	0.476 / 11129

function. Notice that the width τ of the window is exactly the same as in Algorithm 1. A naive solution to compute the measurements would consist in setting a measurement to 1 if at least one of the hashes falls inside the window drawn for that measurement. However, this solution would still require $\mathcal{O}(km)$ comparisons. A better solution is to use a binary search tree to store the windows by means of their sorted bottom values so that the value of the measurement can be determined by traversing the tree, yielding a logarithmic complexity in the number of measurements. More precisely, the tree stores the measurement number $i \in [0, m - 1]$, the value of the bottom of the corresponding window as well as pointers to the two children. The tree is created in the following way. The bottom values are sorted in increasing order and their median value is inserted as root of the tree. Then the tree is recursively created following the rule that the left (respectively, right) subtree includes the windows with bottom values smaller (respectively, larger) than the parent. To determine whether the measurements are zero or nonzero, each hash of the set traverses the tree. The hash is first compared to the value of the window bottom stored in the root node to determine if it falls inside that window. If it does, the measurement corresponding to the index stored in that node is set to 1 and the next hash is examined. Otherwise, the right or left children is examined depending if the hash is below or above the bottom of the current window. This is repeated until a leaf is reached, or a measurement is set to 1. Partially overlapping windows are handled with a local search: if a hash is determined to fall inside a window, the windows whose bottoms are immediately smaller or larger are checked to determine if the hash also falls inside them. This is repeated until no neighboring windows report the hash falling inside. Algorithm 2 shows the whole procedure to compute measurements with Fast SparseHash, while Figure 1 shows the process in a pictorial fashion. Table 1 reports an experimental comparison of the runtime for SparseHash and Fast SparseHash for various values of k and m : Fast SparseHash is significantly faster and has a sublinear increase in runtime for increasing m while SparseHash has a linear increase.

4 Analytical results

In the following, we present two metrics to compute set similarities from measurements in the reduced domain. Given $u, v \in \mathbb{R}^n$ and their measurements $y, z \in \{0, 1\}^m$ obtained via SparseHash, we are interested in defining a similarity metric between y and z that approximately embeds the Jaccard coefficient $J(S_u, S_v) = |S_u \cap S_v| / |S_u \cup S_v|$ of $S_u = \text{supp}(u)$ and $S_v = \text{supp}(v)$. To simplify the notation, let $J_{u,v} := J(S_u, S_v)$.

4.1 Jaccard coefficient

Let $y, z \in \{0, 1\}^m$. We define

$$\begin{aligned} \text{sim}_{\cup}(y, z) &:= \frac{1}{m} \sum_{i=1}^m \mathbf{1}(\{y_i = 0, z_i = 0\}), \\ \text{sim}_{\cap}(y, z) &:= \frac{\sum_{i=1}^m \mathbf{1}(\{y_i = 0\}) \sum_{j=1}^m \mathbf{1}(\{z_j = 0\})}{m \sum_{i=1}^m \mathbf{1}(\{y_i = 0, z_i = 0\})} \end{aligned}$$

where $\mathbf{1}(\{A\})$ is the indicator function which returns 1 when A is true, while $\mathbf{1}(\{A, B\})$ returns 1 when both A and B are true. Then, comparison in the reduced space can be done with the following similarity index:

$$\text{sim}_{\text{sh}}(y, z) := \frac{\log(\text{sim}_{\cap}(y, z))}{\log(\text{sim}_{\cup}(y, z))}. \quad (1)$$

We notice that γ must be designed so that this formula has small probability to be undefined. In the following theorem, we state that sim_{sh} is a random variable that concentrates around the Jaccard coefficient between the supports of the original signals. We emphasize that Theorem 1 is a more refined version of Proposition 1 in [37]. More precisely, a deeper theoretical analysis leads to a new estimation of the performance that is more explicit in terms of the main parameters $m, \gamma, k_{\min}, k_{\max}$ and N .

Theorem 1. *Let $\mathcal{X}_N = \{x_i \in \mathbb{R}^n : |\text{supp}(x_i)| \in [k_{\min}, k_{\max}]\}_{i=1}^N$ be a set of N sparse vectors. For any $\epsilon > 0, \beta > 2$ and any integer n , let m be a positive integer such that $m > 32 \frac{\log 4 + \beta \log N}{\gamma^2 k_{\min}^2 e^{-\gamma k_{\max}} \epsilon^2}$. Then,*

$$\mathbb{P} \left[\bigcup_{(u,v) \in \mathcal{X}_N} \{|\text{sim}_{\text{sh}}(Au, Av) - J_{u,v}| > \epsilon\} \right] \leq N^{-\beta+2}. \quad (2)$$

Sketch of the proof For brevity, we report only the key steps of the proof; details are provided in the supplementary material. Let $u \in \Sigma_{k_1}$ and $v \in \Sigma_{k_2}$. Let us define: $\zeta := (1 - \gamma)^{\frac{k_1 + k_2}{1 + J_{u,v}}}$. Exploiting the Hoeffding's inequality (see [18]), we can prove the following inequalities:

$$\begin{aligned} \mathbb{P}(|\text{sim}_{\cup}(Au, Av) - \zeta| > \epsilon) &\leq 2e^{-2\epsilon^2 m} \\ \mathbb{P}(|\text{sim}_{\cap}(Au, Av) - \zeta^{J_{u,v}}| > \epsilon) &\leq 6e^{-m \min\left\{(1-\gamma)^{8k_{\max}} \frac{\epsilon^2}{8}, \frac{(1-\gamma)^{4k}}{2}\right\}}. \end{aligned} \quad (3)$$

By using the fact that for any positive random variable X such that $\mathbb{P}(|X - \mu_X| > \epsilon) \leq p_X(\epsilon)$ with $\mu_X > 0$, it holds that $\mathbb{P}(|\log(X) - \log(\mu_X)| > \epsilon) \leq p_X(\epsilon \mu_X)$, we can deduce the following inequalities from (3):

$$\begin{aligned} \mathbb{P}(|\log(\text{sim}_{\cap}(Au, Av)) - J_{u,v} \log \zeta| > \epsilon) &\leq 6e^{-m(1-\gamma)^{12k_{\max}} \frac{\epsilon^2}{8}}; \\ \mathbb{P}(|\log(\text{sim}_{\cup}(Au, Av)) - \log \zeta| > \epsilon) &\leq 2e^{-2m\epsilon^2(1-\gamma)^{4k_{\max}}}. \end{aligned} \quad (4)$$

Moreover, we can prove that, given $X_u = \frac{\sum_{i=1}^m \mathbb{1}(\{(Au)_i=0\})}{m}$,

$$\mathbb{P}(|X_u X_z - (1 - \gamma)^{k_1+k_2}| > \epsilon) \leq 4e^{-2m \min\left\{\frac{\epsilon^2}{9}, \frac{(1-\gamma)^{2k_{\max}}}{4}\right\}}. \quad (5)$$

Finally, by merging (4) and (5), under the given assumption on m , the thesis is easily derived. \square

For large n , k , and m , γk represents the average number of nonzeros in each row of A that align with the support of cardinality k . This observation reveals three regimes, corresponding to the scaling of γ and k : if $\gamma k_{\min} = \Theta(1)$, then $m = O(\log N)$ is sufficient to get the bound in Theorem 1. In sharp contrast, if $\gamma k_{\max} \rightarrow \infty$ or $\gamma k_{\min} \rightarrow 0$, then m must increase to guarantee the concentration with high probability.

4.2 Hamming distance and LSH

As for MinHash, signals can be compared in the reduced space using the Hamming distance d_H between two hash codes. From the law of large numbers we have the following properties. Given $u \in \Sigma_{k_1}$ and $v \in \Sigma_{k_2}$

$$E_{\text{sh}} := \frac{\mathbb{E}[d_H(Au, Av)]}{m} = (1 - \gamma)^{k_1} + (1 - \gamma)^{k_2} - 2(1 - \gamma)^{\frac{k_1+k_2}{1+J_{u,v}}} \quad (6)$$

and $\widehat{V} := \text{Var}[d_H(Au, Av)/m] = (1 - E_{\text{sh}})E_{\text{sh}}/m$. For signals with similar sparsity degree $k_1 \approx k_2 = k$, by setting $(1 - \gamma)^k = 1/2$ in order to maximize the entropy of the binary measurements, we obtain: $E_{\text{sh}} \approx 1 - 2^{\frac{J-1}{1+J}}$. The characterization of the relationship between the Hamming distance of the hashes and the original Jaccard coefficient derived in (6) is important in the context of LSH. In a nutshell, LSH allows us to approximate nearest neighbor database searches with sublinear complexity (see [2]), without scanning all the entries in the database. Let $\mathcal{X} \subset \mathbb{R}^n$ be a set of points with distance measure $d_{\mathcal{X}}$, and consider the (R, c) -NN problem where one is concerned with retrieving all the neighbors of the query point within a distance R , while discarding the points at distances greater than cR . An LSH family is defined as follows.

Definition 1. Let $p_1 > p_2$ and $r_1 < r_2$. A family $\mathcal{H} : \{h : \mathcal{X} \rightarrow \mathcal{U}\}$ is called (r_1, r_2, p_1, p_2) -sensitive for $d_{\mathcal{X}}$ if for any $x, \xi \in \mathcal{X}$ the following fact hold: if $d_{\mathcal{X}}(x, \xi) \leq r_1$ then $\mathbb{P}(h(\xi) = h(x)) \geq p_1$; if $d_{\mathcal{X}}(x, \xi) \geq r_2$ then $\mathbb{P}(h(\xi) = h(x)) \leq p_2$.

In the (c, R) -NN problem, we set $r_1 = R$ and $r_2 = cR$ and we define a new family of functions $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathcal{U}^m\}$ such that $f(x) = (h_1(x), \dots, h_m(x))^{\top}$, where $h_i \in \mathcal{H}$ are chosen independently uniformly at random from \mathcal{H} . We notice that $f \in \mathcal{F}$ is (r_1, r_2, p_1^m, p_2^m) -sensitive for $d_{\mathcal{X}}$. Fixed $L > 0$, we define a new family \mathcal{G} of hash functions g constructed from L random functions f_1, \dots, f_L from \mathcal{F} . We say that $g(\xi) = g(x)$ if $f_i(\xi) = f_i(x)$ for at least one $i \in \{1, \dots, L\}$. Since the members of \mathcal{F} are independently chosen for any $g \in \mathcal{G}$, \mathcal{G} is a $(r_1, r_2, 1 - (1 - p_1^m)^L, 1 - (1 - p_2^m)^L)$ -sensitive family.

During preprocessing, L hash tables, each corresponding to a different hash function g_i , are constructed, by storing each $x \in \mathcal{X}$ in the bucket $g_i(x)$. Given a query item ξ , we retrieve first $3L$ data points that are hashed to the same bucket $g_i(\xi)$ with $i = 1, \dots, L$ and if there is a point x^* within distance cR from ξ we return 'yes' and x^* , else we return 'no'. If $m = \log N / \log(1/p_2)$, $L = N^\rho$, $\rho = \frac{\log(1/p_1)}{\log(1/p_2)}$ then the algorithm is successful with constant probability and the algorithm has the following properties: (a) preprocessing time is $O(N^{1+\rho}mT)$, where T is the time to evaluate a function $h \in \mathcal{H}$ on an item; (b) storage is of order $O(NL + Nm) = O(N^{1+\rho} + Nm)$; (c) query time is $O(L(mT + nNp_2^m)) = O(N^\rho(mT + n))$. Notice that savings in terms of storage can be achieved when 1-bit measurements are used in place of real-valued measurements.

We now study the performance of LSH in terms of storage and time requirements to respond to a query, using MinHash and SparseHash as embeddings. Such performance metrics are entirely governed by embedding, more precisely, by the function that maps the Jaccard coefficient to the probability of having two equal bits. This function is linear in MinHash and nonlinear in SparseHash. In order to make a fair comparison, we use 1-bit MinHash, so that the storage complexity is equalized between the two approaches. Notice that m is chosen so as to minimize L . If a pair of signals in \mathcal{X} has Jaccard coefficient J , then the probability that their hashes computed with SparseHash become a candidate pair is given by: $P_{sh} = 1 - (1 - p_{sh}^{m_{sh}})^{L_{sh}}$ with $p_{sh} = 2^{\frac{J-1}{1+J}}$, while for 1-bit MinHash, $P_{mh} = 1 - (1 - p_{mh}^{m_{mh}})^{L_{mh}}$ with $p_{mh} = (J+1)/2$. The following proposition states that, with equal m , SparseHash requires less tables, which enables shorter query times and lower storage requirements.

Proposition 1. *If $P_{mh} = P_{sh}$, $m_{mh} = m_{sh} = m$, then $L_{sh} \leq L_{mh}$.*

Proof. We have

$$\frac{L_{sh}}{L_{mh}} = \frac{\log(1 - ((J+1)/2)^m)}{\log(1 - 2^{m \frac{J-1}{1+J}})}.$$

We then obtain $L_{sh} \leq L_{mh}$ if $2^{\frac{J-1}{1+J}} \geq (J+1)/2$, which holds for any $J \in [0, 1]$. More details are provided in the supplementary material. \square

5 Numerical experiments

In this section, we propose numerical experiments that validate the theoretical results and show the effective performance of SparseHash, both on synthetic and real datasets. The code for these experiments is available in [34].

5.1 Numerical validation

In [37], a numerical validation on the concentration of SparseHash around the true Jaccard coefficient was proposed. We retrieve the same experiment to validate the concentration on the Hamming distance. We randomly generate a




Figure 2: Numerical validation: Hamming distance (mean and variance).

large number of signals with different amount of support overlap and compute their random projections via γ -sparsified matrices $A \in \mathbb{R}^{m,n}$. We set $n = 1000$, $m = 50$, and sparsity level $k = 230$. Mean and variance are evaluated over 500 runs. γ is set as the value that maximizes the entropy of the binary measurements, i.e. generates zero or nonzero measurements with equal probability. Since $\mathbb{P}(f_i(u) = 0) = (1 - \gamma)^k$, we set: $\gamma = 1 - 2^{-\frac{1}{k}} \approx 3 \cdot 10^{-3}$. In Figure 2, we depict the Hamming distance. The dashed cyan line is the theoretical mean P_{sh} computed from the true Jaccard coefficient defined in (6). As expected, the experimental mean overlies the theoretical mean.

5.2 Similar text documents

As discussed in [37], the problem of finding near-duplicate or similar documents in an archive of text data is outstanding (see [7, 9, 17]). Documents can be represented with bag-of-words models. Given a vocabulary of n words, we can associate a document with a $u \in \mathbb{R}^n$, where u_i counts the occurrences of the i th word of the vocabulary in the document. Bag-of-words models typically yield sparse signals, as the number of different words appearing in a single document is usually much smaller than the size of the vocabulary. We now retrieve the experiment proposed in [26], where the effects of quantization are evaluated on MinHash. We use the UCI dataset of New York Times articles (see [27]), composed of about 300000 news articles, with a bag-of-words representation given for each article. The vocabulary contains $n = 102660$ words. The mean number of different words used in each article is $k = 232$; we use this value to assess γ as set in Section 5.1.




Figure 3: Experiment on similar text documents: precision and recall, threshold 0.5.




Figure 4: Experiment on similar text documents: precision and recall, threshold 0.6.

We compare the performance of SparseHash (with both Jaccard and Hamming metrics) and 1-bit MinHash, in terms of precision and recall. Specifically, we define as similar the documents with Jaccard coefficient larger than a certain threshold, and we try to detect them. In figures 3 and 4, we set the threshold to 0.5 and 0.6, respectively, and we show precision and recall as functions of the number of measurements m , and precision as a function of recall. We see that the precision of SparseHash with Jaccard metric outperforms 1-bit MinHash, in particular when $m < 100$. For larger m , both methods are efficient, with precision close to 1. The recall is close to 1 for both methods. The gain obtained by SparseHash is well visualized also in the precision-recall curves, depicted for $m = 48$ and $m = 96$.

5.3 Metagenome clustering

Metagenome clustering is concerned with detecting communities of microorganisms starting from genomic sequences. A genomic sequence can be seen as a long string of A, C, T or G characters representing the four types of nucleotides. In the following, we deal with assembled sequences, which are reconstructed from overlapping partial reads produced by the sequencing instruments. Metagenome clustering is formulated in terms of pairwise distances between sequences. The distance metric of interest is 1-ANI: ANI is the average nucleotide identity, i.e., the percentage of unchanged nucleotides in the two sequences. Since the genomic sequences can be extremely long, dimensionality reduction methods are essential to efficiently compute distances. In [30], the so-called MASH algorithm uses MinHash to this purpose. The authors split a sequence into substrings, called κ -mers, using a sliding window approach. In their experiments, $\kappa = 21$. Each genomic sequence is then represented by the set of its κ -mers and the Jaccard coefficient between such sets correlates with the expected ANI. By using MinHash to compress the set of κ -mers, the authors achieve a significant dimensionality reduction.

In this section, we test SparseHash for the same purpose, and compare it to MinHash. The performance is evaluated on the capacity of preservation of the Jaccard coefficient. The dataset from [35] is used, which contains $N = 747$ sequences of various length. A sliding sequence of stride equal to 1 is used to generate the set of κ -mers and, by keeping only the unique κ -mers, the sets present different cardinalities, ranging from $k_{min} = 4002$ to $k_{max} = 219972647$. These sets are then embedded with MinHash or SparseHash and the full matrix of all pairwise similarities is generated by computing distances on the measurement vectors. Due to the computational infeasibility of computing exact Jaccard values to estimate the quality of the approximation produced by MinHash and SparseHash, we approximate them using a large number of SparseHash measurements ($m = 5 \cdot 10^6$) and MinHash measurements ($m = 10^6$) and then average the two estimates. This is done to avoid any bias for a specific algorithm. We remark that we use the MASH code provided by [30] to implement MinHash. Due to the large values of k and m , that code uses the bottom- m MinHash sketch to accelerate the computation of the sketches. However, it is not amenable to bi-

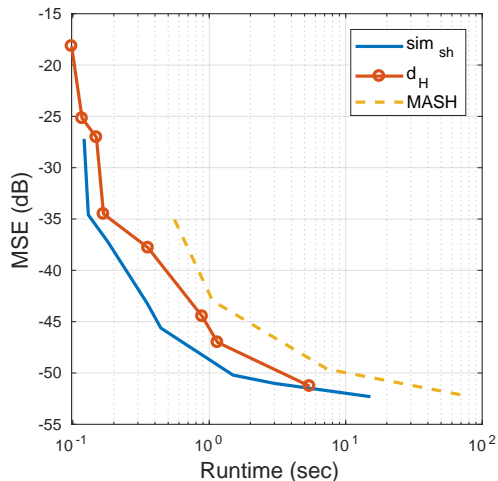


Figure 5: Metagenome clustering: MSE on pairwise Jaccard matrix with computation time.

narization: each MinHash measurement requires 64 bits against 1 bit required by SparseHash. Figure 5 shows the mean square error (MSE) on the matrix with all Jaccard pairwise similarities with respect to the true Jaccard values as function of the computation time required to compute all pairwise distances in the embedded space. It can be noticed that SparseHash outperforms MASH thanks to the 1-bit measurements and the faster sim_{sh} and Hamming distance metrics that can be implemented with bitwise operations.

6 Conclusion

In this letter, we have analyzed SparseHash, a novel embedding technique for dimensionality reduction of sets. Exploiting the concept of sparsity and concentration results, SparseHash is proven to preserve Jaccard metric. Efficient implementations and numerical experiments show that SparseHash outperforms MinHash in different applications. Future work will envisage the comparison with other strategies, e.g., Bloom filters.

References

- [1] D. Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. Syst. Sci.*, 66(4):671–687, 2003.
- [2] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM*, 51(1):117–122, 2008.

- [3] V. Bioglio, T. Bianchi, and E. Magli. On the fly estimation of the sparsity degree in compressed sensing using sparse sensing matrices. In *Proc. IEEE ICASSP*, pages 3801–3805, 2015.
- [4] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, 1970.
- [5] P. Boufounos and S. Rane. Secure binary embeddings for privacy preserving nearest neighbors. In *IEEE WIFS*, pages 1–6, 2011.
- [6] P. T. Boufounos and S. Rane. Efficient coding of signal distances using universal quantized embeddings. In *Data Compression Conference (DCC), 2013*, pages 251–260, 2013.
- [7] A. Broder. On the resemblance and containment of documents. In *Proc. Compression and Complexity of Sequences*, pages 21–29, 1997.
- [8] A. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. Min-wise independent permutations. *Journal Comput. Syst. Sci.*, 60(3):630 – 659, 2000.
- [9] A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig. Syntactic clustering of the web. *Comput. Netw. ISDN Syst.*, 29(8-13):1157–1166, 1997.
- [10] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inf. Theory*, 52(12):5406–5425, 2006.
- [11] J. Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143 – 154, 1979.
- [12] M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proc. ACM STOC*, pages 380–388, 2002.
- [13] Tobias Christiani and Rasmus Pagh. Set similarity search beyond MinHash. In *Proc. ACM STOC*, pages 1094–1107, 2017.
- [14] Edith Cohen and Haim Kaplan. Summarizing data using bottom-k sketches. In *Proc. ACM PODC*, pages 225–234, 2007.
- [15] Søren Dahlgaard, Mathias Bæk Tejs Knudsen, and Mikkel Thorup. Fast similarity sketching. *CoRR*, abs/1704.04370, 2017.
- [16] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006.
- [17] M. Henzinger. Finding near-duplicate web pages: A large-scale evaluation of algorithms. In *Proc. ACM SIGIR*, pages 284–291, 2006.

- [18] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, March 1963.
- [19] Anna Huang. Similarity measures for text document clustering. In *Proc. New Zealand computer science research student conf.*, pages 49–56, 2008.
- [20] P. Indyk. A small approximately min-wise independent family of hash functions. In *Proc. ACM-SIAM SODA*, pages 454–456, 1999.
- [21] Sergey Ioffe. Improved consistent sampling, weighted minhash and l1 sketching. In *Proc. IEEE Int. Conf. Data Mining*, pages 246–255, 2010.
- [22] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inf. Theory*, 59(4):2082–2102, 2013.
- [23] Jiyong Jang, David Brumley, and Shobha Venkataraman. Bitshred: feature hashing malware for scalable triage and semantic analysis. In *Proc. ACM CCS*, pages 309–320, 2011.
- [24] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary Mathematics*, 26, 1984.
- [25] P. Li, T. J. Hastie, and K. W. Church. Very sparse random projections. In *Proc. ACM SIGKDD*, pages 287–296, 2006.
- [26] P. Li and C. König. b-bit minwise hashing. In *Proceedings of the 19th International Conference on World Wide Web*, pages 671–680, 2010.
- [27] M. Lichman. UCI machine learning repository. <http://archive.ics.uci.edu/ml>, 2013.
- [28] L. Luo, Deke Guo, Richard TB Ma, Ori Rottenstreich, and Xueshan Luo. Optimizing bloom filter: Challenges, solutions, and comparisons. *arXiv preprint arXiv:1804.04777*, 2018.
- [29] MurmurHash3. Code. <https://github.com/aappleby/smhasher/blob/master/src/MurmurHash3.cpp>, 2016.
- [30] Brian D. Ondov, Todd J. Treangen, Páll Melsted, Adam B. Mallonee, Nicholas H. Bergman, Sergey Koren, and Adam M. Phillippy. Mash: fast genome and metagenome distance estimation using minhash. *Genome Biology*, 17(1):132, 2016.
- [31] C. Ravazzi, S. M. Fosson, T. Bianchi, and E. Magli. Signal sparsity estimation from compressive noisy projections via γ -sparsified random matrices. In *Proc. IEEE ICASSP*, pages 4029–4033, 2016.

- [32] C. Ravazzi, S. M. Fosson, T. Bianchi, and E. Magli. Sparsity estimation from compressive projections via sparse random matrices. *EURASIP J. Adv. Signal Process.*, 2018(56):1–18, 2018.
- [33] Anshumali Shrivastava. Exact weighted minwise hashing in constant time. *CoRR*, abs/1602.08393, 2016.
- [34] SparseHash. Code. <https://github.com/diegovalsesia/sparsehash>, 2017.
- [35] The Human Microbiome Project Consortium. Structure, function and diversity of the healthy human microbiome. *Nature*, 486(7402):207–214, 2012.
- [36] D. Valsesia, G. Coluccia, T. Bianchi, and E. Magli. Large-scale image retrieval based on compressed camera identification. *IEEE Trans. Multim.*, 17(9):1439–1449, 2015.
- [37] D. Valsesia, S. M. Fosson, C. Ravazzi, T. Bianchi, and E. Magli. Sparse-hash: Embedding Jaccard coefficient between supports of signals. In *IEEE ICMEW*, pages 1–6, 2016.
- [38] D. Valsesia and E. Magli. Binary adaptive embeddings from order statistics of random projections. *IEEE Signal Process. Lett.*, 24(1):111–115, 2017.
- [39] J. Wang, T. Zhang, J. song, N. Sebe, and H. T. Shen. A survey on learning to hash. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):769–790, 2018.
- [40] W.L. Zhao, H. Jégou, and G. Gravier. Sim-Min-Hash: An efficient matching technique for linking large image collections. In *Proc. ACM MM*, pages 577–580, 2013.

A Appendix

In this Appendix, we provide the details of the proofs of Theorem 1 and Proposition 1. Here, we report all the steps.

A.1 Notation

To facilitate the reading, we recall the notation used in the manuscript.

- $u, v \in \mathbb{R}^n$;
- $A \in \mathbb{R}^{m,n}$ is a γ -sparsified random matrix, $\gamma \in (0, 1)$, that is: each $A_{i,j}$ is zero with probability $1 - \gamma$, while $A_{i,j}$ is generated according to given distribution with zero mean and finite variance with probability γ ;
- $\mathbb{1}(\{A\})$ is the indicator function, which returns 1 when A is true;
- $\mathbb{1}(\{A, B\})$ is the indicator function which returns 1 when both A and B are true;

- $S_u = \{i \in \{1, \dots, n\} \text{ s.t. } u_i \neq 0\}$
- $y, z \in \{0, 1\}^m$
- $\text{sim}_{\cup}(y, z) := \frac{1}{m} \sum_{i=1}^m \mathbf{1}(\{y_i = 0, z_i = 0\})$;
- $\text{sim}_{\cap}(y, v) := \frac{\sum_{i=1}^m \mathbf{1}(\{y_i=0\}) \sum_{j=1}^m \mathbf{1}(\{z_j=0\})}{m \sum_{i=1}^m \mathbf{1}(\{y_i=0, z_i=0\})}$;
- $J_{u,v}$ (or J) = Jaccard coefficient of the pair (u, v) .

A.2 Proof of Theorem 1

We start with some preliminary properties.

Lemma 1. *Let $u \in \Sigma_{k_1}, v \in \Sigma_{k_2}$,*

$$\mathbb{P}\left(\left|\text{sim}_{\cup}(Au, Av) - (1 - \gamma)^{\frac{k_1+k_2}{1+J_{u,v}}}\right| > \epsilon\right) \leq 2e^{-2\epsilon^2 m}. \quad (7)$$

Proof. $S_u = \text{supp}(u)$ and $S_v = \text{supp}(v)$ It should be noticed that $\mathbf{1}(\{y_i = 0, z_i = 0\})$ is a Bernoulli random variable with

$$\begin{aligned} \mathbb{E}[\text{sim}_{\cup}(Au, Av)] &= (1 - \gamma)^{|S_u \cup S_v|} \\ &= (1 - \gamma)^{(k_1+k_2)/(1+J_{u,v})}, \end{aligned}$$

where the last equality is obtained using the relation

$$J_{u,v} = \frac{|S_u| + |S_v|}{|S_u \cup S_v|} - 1.$$

By Hoeffding's inequality (see [18]) we have

$$\mathbb{P}(|\text{sim}_{\cup}(Au, Av) - \mathbb{E}[\text{sim}_{\cup}(Au, Av)]| > \epsilon) \leq 2e^{-2\epsilon^2 m}.$$

□

Lemma 2. *Let X and Z be two random variables such that*

$$\begin{aligned} \mathbb{P}(|X - \mu_X| > \epsilon) &\leq p_X(\epsilon) \\ \mathbb{P}(|Z - \mu_Z| > \epsilon) &\leq p_Z(\epsilon) \end{aligned}$$

with $\mu_X > 0, \mu_Z > 0$, then

$$\begin{aligned} \mathbb{P}(|XZ - \mu_X \mu_Z| > \epsilon) &\leq p_X\left(\frac{\epsilon}{3\mu_Z}\right) \\ &\quad + p_Z\left(\min\left\{\frac{\epsilon}{2\mu_X}, \frac{\mu_Z}{2}\right\}\right) \end{aligned}$$

Proof. Let us define the following events

$$E_X = \left\{ |X - \mu_X| \leq \frac{\epsilon}{3\mu_Z} \right\}$$

$$E_Z = \left\{ |Z - \mu_Z| \leq \min \left\{ \frac{\epsilon}{2\mu_X}, \frac{\mu_Z}{2} \right\} \right\}.$$

From the law of total probability the following series of inequalities follows:

$$\begin{aligned} \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon) &= \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon | E_X) \mathbb{P}(E_X) \\ &\quad + \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon | E_X^c) \mathbb{P}(E_X^c) \\ &\leq \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon | E_X) + \mathbb{P}(E_X^c) \end{aligned}$$

Repeating the same argument, we get

$$\begin{aligned} \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon) &\leq \\ &\leq \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon | E_X \cap E_Z) + \mathbb{P}(E_X^c) + \mathbb{P}(E_Z^c) \end{aligned}$$

It should be noticed that

$$\begin{aligned} |XZ - \mu_X\mu_Z| &= |XZ - \mu_X Z + \mu_X Z - \mu_X\mu_Z| \\ &\leq |Z| |X - \mu_X| + \mu_X |Z - \mu_Z|. \end{aligned}$$

If E_Z, E_X hold then $\mu_Z/2 \leq Z \leq 3\mu_Z/2$ and

$$|XZ - \mu_X\mu_Z| \leq \epsilon$$

from which

$$\mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon | E_X \cap E_Z) = 0.$$

We conclude that

$$\begin{aligned} \mathbb{P}(|XZ - \mu_X\mu_Z| > \epsilon) &\leq \mathbb{P}(E_X^c) + \mathbb{P}(E_Z^c) \\ &= p_X \left(\frac{\epsilon}{3\mu_Z} \right) + p_Z \left(\min \left\{ \frac{\epsilon}{2\mu_X}, \frac{\mu_Z}{2} \right\} \right). \end{aligned}$$

□

From Lemma 2, we can derive the following result.

Corollary 1. Let $u \in \Sigma_{k_1}, v \in \Sigma_{k_2}$, and denote

$$X = \frac{\sum_{i=1}^m \mathbf{1}(\{(Au)_i = 0\})}{m}$$

$$Z = \frac{\sum_{i=1}^m \mathbf{1}(\{(Av)_i = 0\})}{m}$$

then

$$\mathbb{P}(|XZ - (1-\gamma)^{k_1+k_2}| > \epsilon) \leq 4e^{-2m \min \left\{ \frac{\epsilon^2}{9}, \frac{(1-\gamma)^{2k_{\max}}}{4} \right\}} \quad (8)$$

Proof. It should be noticed that $\mathbf{1}(\{(Au)_i = 0\})$ is a Bernoulli random variable with $\mathbb{P}(\{(Au)_i = 0\}) = (1 - \gamma)^{k_1} := \mu_X$. Then from the Hoeffding's inequality [18] we get

$$\mathbb{P}(|X - \mu_X| > \epsilon) \leq 2e^{-2\epsilon^2 m}.$$

and, analogously denoting $\mu_Z = (1 - \gamma)^{k_2}$,

$$\mathbb{P}(|Z - \mu_Z| > \epsilon) \leq 2e^{-2\epsilon^2 m}.$$

Applying Lemma 2 and computing

$$\mathbb{P}\left(|X - \mu_X| > \frac{\epsilon}{3\mu_Z}\right) \leq 2e^{-\frac{2\epsilon^2 m}{9(1-\gamma)^{2k_2}}} \leq 2e^{-\frac{2\epsilon^2 m}{9}} \quad (9)$$

and, similarly,

$$\begin{aligned} & \mathbb{P}\left(|Z - \mu_Z| > \min\left\{\frac{\epsilon}{2\mu_X}, \frac{\mu_Z}{2}\right\}\right) \\ & \leq 2e^{-2m \min\left\{\frac{\epsilon^2}{4(1-\gamma)^{k_1}}, \frac{(1-\gamma)^{2k_2}}{4}\right\}} \\ & \leq 2e^{-2m \min\left\{\frac{\epsilon^2}{4}, \frac{(1-\gamma)^{2k_{\max}}}{4}\right\}}. \end{aligned} \quad (10)$$

we obtain (8). □

Lemma 3. *Let X, Z be random variables such that*

$$\mathbb{P}(|X - \mu_X| > \epsilon) \leq p_X(\epsilon), \quad \mathbb{P}(|Z - \mu_Z| > \epsilon) \leq p_Z(\epsilon)$$

with $\mu_X > 0, \mu_Z > 0$, then

$$\begin{aligned} & \mathbb{P}\left(\left|\frac{X}{Z} - \frac{\mu_X}{\mu_Z}\right| > \epsilon\right) \\ & \leq p_X\left(\frac{\epsilon\mu_Z}{4}\right) + p_Z\left(\min\left\{\frac{\epsilon}{4}\frac{\mu_Z^2}{\mu_X}, \frac{\mu_Z}{2}\right\}\right), \end{aligned} \quad (11)$$

Proof. Let us define the following events

$$\begin{aligned} E_X &= \left\{|X - \mu_X| \leq \frac{\epsilon\mu_Z}{4}\right\} \\ E_Z &= \left\{|Z - \mu_Z| \leq \min\left\{\frac{\epsilon}{4}\frac{\mu_Z^2}{\mu_X}, \frac{\mu_Z}{2}\right\}\right\}. \end{aligned}$$

From the law of total probability the following series of inequalities follows

$$\begin{aligned}
\mathbb{P}\left(\left|\frac{X}{Z} - \frac{\mu_X}{\mu_Z}\right| > \epsilon\right) &= \mathbb{P}\left(\frac{|X\mu_Z - Z\mu_X|}{|Z|\mu_Z} > \epsilon\right) \\
&= \mathbb{P}\left(\frac{|X\mu_Z - \mu_X\mu_Z + \mu_X\mu_Z - Z\mu_X|}{|Z|\mu_Z} > \epsilon\right) \\
&\leq \mathbb{P}\left(\frac{\mu_Z|X - \mu_X| + \mu_X|Z - \mu_Z|}{|Z|\mu_Z} > \epsilon\right) \\
&\leq \mathbb{P}\left(\frac{\mu_Z|X - \mu_X| + \mu_X|Z - \mu_Z|}{|Z|\mu_Z} > \epsilon \mid E_X \cap E_Z\right) \\
&\quad + \mathbb{P}(E_X^c) + \mathbb{P}(E_Z^c)
\end{aligned}$$

If E_Z holds then $Z \geq \mu_Z/2$ and

$$\mathbb{P}\left(\frac{\mu_Z|X - \mu_X| + \mu_X|Z - \mu_Z|}{|Z|\mu_Z} > \epsilon \mid E_X, E_Z\right) = 0,$$

we have

$$\mathbb{P}\left(\left|\frac{X}{Z} - \frac{\mu_X}{\mu_Z}\right| > \epsilon\right) \leq p_X\left(\frac{\epsilon\mu_Z}{4}\right) + p_Z\left(\min\left\{\frac{\epsilon}{4}\frac{\mu_Z^2}{\mu_X}, \frac{\mu_Z}{2}\right\}\right).$$

□

From Lemma 3 and Corollary 1, we obtain the following result.

Corollary 2. *Let $u \in \Sigma_{k_1}, v \in \Sigma_{k_2}$, then*

$$\begin{aligned}
&\mathbb{P}\left(\left|\text{sim}_\cap(Au, Av) - (1 - \gamma)^{\frac{(k_1+k_2)J_{u,v}}{1+J_{u,v}}}\right| > \epsilon\right) \\
&\leq 6 \exp\left(-m \min\left\{(1 - \gamma)^{8k_{\max}} \frac{\epsilon^2}{8}, \frac{(1 - \gamma)^{4k}}{2}\right\}\right)
\end{aligned} \tag{12}$$

Proof. The inequality in (12) follows directly from Lemma 3 by setting

$$\begin{aligned}
X &= \frac{\sum_{i=1}^m \mathbf{1}(\{(Au)_i = 0\}) \sum_{i=1}^m \mathbf{1}(\{(Av)_i = 0\})}{m^2}, \\
Z &= \text{sim}_\cup(Au, Av), \\
\mu_X &= (1 - \gamma)^{k_1+k_2}, \quad \mu_Z = (1 + \gamma)^{(k_1+k_2)/(1+J_{u,v})}
\end{aligned}$$

and estimating the probabilities using Lemma 3, Corollary 1, $k_1 + k_2 \leq 2k_{\max}$ and $J \geq 0$

$$p_X\left(\frac{\epsilon\mu_Z}{4}\right) \leq 4e^{-m \min\left\{\frac{\epsilon^2}{9.8}(1-\gamma)^{4k_{\max}}, \frac{(1-\gamma)^{2k_{\max}}}{2}\right\}}$$

and

$$\begin{aligned}
& p_Z \left(\min \left\{ \frac{\epsilon \mu_Z^2}{4 \mu_X}, \frac{\mu_Z}{2} \right\} \right) = \\
& = 2 \exp \left(-2m \left(\min \left\{ \frac{\epsilon(1-\gamma)^{2(k_1+k_2)/(1+J)}}{4(1-\gamma)^{k_1+k_2}}, \frac{(1-\gamma)^{(k_1+k_2)/(1+J)}}{2} \right\} \right)^2 \right) \\
& \leq 2 \exp \left(-2m \left(\min \left\{ (1-\gamma)^{4k_{\max}} \frac{\epsilon}{4}, \frac{(1-\gamma)^{2k_{\max}}}{2} \right\} \right)^2 \right) \\
& = 2 \exp \left(-m \min \left\{ (1-\gamma)^{8k_{\max}} \frac{\epsilon^2}{8}, \frac{(1-\gamma)^{4k_{\max}}}{2} \right\} \right).
\end{aligned}$$

We conclude that

$$\begin{aligned}
& \mathbb{P} \left(\left| \text{sim}_{\cap}(Au, Av) - (1-\gamma)^{\frac{(k_1+k_2)J_{u,v}}{1+J_{u,v}}} \right| > \epsilon \right) \\
& \leq p_X \left(\frac{\epsilon \mu_Z}{4} \right) + p_Z \left(\min \left\{ \frac{\epsilon \mu_Z^2}{4 \mu_X}, \frac{\mu_Z}{2} \right\} \right) \\
& \leq 6 \exp \left(-m \min \left\{ (1-\gamma)^{8k_{\max}} \frac{\epsilon^2}{8}, \frac{(1-\gamma)^{4k_{\max}}}{2} \right\} \right).
\end{aligned}$$

□

Lemma 4. *Let X be a positive random variable such that*

$$\mathbb{P}(|X - \mu_X| > \epsilon) \leq p_X(\epsilon),$$

with $\mu_X > 0$, then

$$\mathbb{P}(|\log(X) - \log(\mu_X)| > \epsilon) \leq p_X(\epsilon \mu_X).$$

Proof. Noticing that $|\log(1+|t|)| \leq |t|$ for any t , the following series of relations holds

$$\begin{aligned}
& \mathbb{P}(|\log(X) - \log(\mu_X)| > \epsilon) \leq \mathbb{P} \left(\left| \log \left(1 + \left| \frac{X - \mu_X}{\mu_X} \right| \right) \right| > \epsilon \right) \\
& \leq \mathbb{P} \left(\frac{|X - \mu_X|}{\mu_X} > \epsilon \right) = p_X(\epsilon \mu_X).
\end{aligned}$$

□

Conclusion of the proof of Theorem 1

Applying Lemma 4 to the results obtained in Lemma 1 and 2 we obtain:

$$\begin{aligned}
& \mathbb{P} \left(\left| \log(\text{sim}_\cap(Au, Av)) - \frac{J_{u,v}(k_1 + k_2) \log(1 - \gamma)}{1 + J_{u,v}} \right| > \epsilon \right) \\
& \leq 6 \exp \left(-m \min \left\{ (1 - \gamma)^{8k_{\max}} \frac{\epsilon^2 \mu_Z^2}{8}, \frac{(1 - \gamma)^{4k_{\max}}}{2} \right\} \right) \\
& = 6 \exp \left(-m \min \left\{ (1 - \gamma)^{8k_{\max}} \frac{\epsilon^2 (1 - \gamma)^{2(k_1 + k_2)J/(1+J)}}{8}, \frac{(1 - \gamma)^{4k_{\max}}}{2} \right\} \right) \\
& \leq 6 \exp \left(-m \min \left\{ (1 - \gamma)^{12k_{\max}} \frac{\epsilon^2}{8}, \frac{(1 - \gamma)^{4k_{\max}}}{2} \right\} \right) \\
& = 6 \exp \left(-m (1 - \gamma)^{12k_{\max}} \frac{\epsilon^2}{8} \right); \\
& \mathbb{P} \left(\left| \log(\text{sim}_\cup(Au, Av)) - \frac{(k_1 + k_2) \log(1 - \gamma)}{1 + J_{u,v}} \right| > \epsilon \right) \\
& \leq 2 \exp \left(-2m\epsilon^2 (1 - \gamma)^{2(k_1 + k_2)} \right) \leq 2 \exp \left(-2m\epsilon^2 (1 - \gamma)^{4k_{\max}} \right).
\end{aligned}$$

Finally, using Lemma 3 and the assumption $m > 32 \frac{\log 4 + \beta \log N}{\gamma^2 k_{\min}^2 e^{-\gamma k_{\max} \epsilon^2}}$, we get

$$\begin{aligned}
& \mathbb{P} \left(\left| \frac{\log(\text{sim}_\cap(Au, Av))}{\log(\text{sim}_\cup(Au, Av))} - J_{u,v} \right| > \epsilon \right) \leq \\
& \leq 8 \exp \left(-m \frac{\epsilon^2}{128} (1 - \gamma)^{12k_{\max}} 4k_{\min}^2 \log^2(1 - \gamma) \right) \leq N^{-\beta+2}.
\end{aligned}$$

The thesis follows by applying the union bound on the last expression.

A.3 Proof of Proposition 1

If a pair of signals in \mathcal{X} has Jaccard coefficient $J \in [0, 1]$, then the probability that their hashes computed with SparseHash become a candidate pair is given by:

$$P_{\text{sh}} = 1 - (1 - p_{\text{sh}}^{m_{\text{sh}}})^{L_{\text{sh}}}$$

with $p_{\text{sh}} = 2^{\frac{J-1}{1+J}}$, while for 1-bit MinHash,

$$P_{\text{mh}} = 1 - (1 - p_{\text{mh}}^{m_{\text{mh}}})^{L_{\text{mh}}}$$

with $p_{\text{mh}} = (J + 1)/2$.

Let $m_{\text{sh}} = m_{\text{mh}} = m$. From the previous equations, it is straightforward to prove that:

$$\frac{L_{\text{sh}}}{L_{\text{mh}}} = \frac{\log(1 - ((J + 1)/2)^m)}{\log\left(1 - 2^{m \frac{J-1}{1+J}}\right)}.$$

We then obtain $L_{sh} \leq L_{mh}$ if

$$\frac{\log(1 - ((J+1)/2)^m)}{\log(1 - 2^{m \frac{J-1}{1+J}})} \leq 1.$$

Since $\log(1 - 2^{m \frac{J-1}{1+J}}) < 0$,

$$\begin{aligned} \log(1 - ((J+1)/2)^m) &\geq \log(1 - 2^{m \frac{J-1}{1+J}}) \\ (1 - ((J+1)/2)^m) &\geq (1 - 2^{m \frac{J-1}{1+J}}) \\ 1 - (J+1)/2 &\geq 1 - 2^{\frac{J-1}{1+J}} \\ (J+1)/2 &\leq 2^{\frac{J-1}{1+J}}. \end{aligned} \tag{13}$$

We can then numerically verify that the last inequality is always true (see Figure 6), in particular the strict inequality holds for any $J \in (0, 1)$.




Figure 6: Proof of Proposition 1