

---

# Large-Scale Distributed Learning via Private On-Device Locality-Sensitive Hashing

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1       Locality-sensitive hashing (LSH) based frameworks have been used efficiently to  
2       select weight vectors in a dense hidden layer with high cosine similarity to an input,  
3       enabling dynamic pruning. While this type of scheme has been shown to improve  
4       computational training efficiency, existing algorithms require repeated randomized  
5       projection of the full layer weight, which is impractical for computational- and  
6       memory-constrained devices. In a distributed setting, deferring LSH analysis to  
7       a centralized host is (i) slow if the device cluster is large and (ii) requires access  
8       to input data which is forbidden in a federated context. Using a new family of  
9       hash functions, we develop the first private, personalized, and memory-efficient  
10       on-device LSH framework. Our framework enables privacy and personalization  
11       by allowing each device to generate hash tables, without the help of a central host,  
12       using device-specific hashing hyper-parameters (*e.g.* number of hash tables or hash  
13       length). Hash tables are generated with a compressed set of the full weights, and  
14       can be serially generated and discarded if the process is memory-intensive. This  
15       allows devices to avoid maintaining (i) the fully-sized model and (ii) large amounts  
16       of hash tables in local memory for LSH analysis. We prove several statistical  
17       and sensitivity properties of our hash functions, and experimentally demonstrate  
18       that our framework is competitive in training large scale recommender networks  
19       compared to other LSH frameworks which assume unrestricted on-device capacity.

## 20   1 Introduction

21       Locality-sensitive hashing (LSH) has proven to be a remarkably effective tool for memory- and  
22       computationally-efficient data clustering and nearest neighbor search [6, 15, 2]. LSH algorithms  
23       such as SimHash [6] can be used to search for vectors in collection  $W \subset \mathbb{R}^d$  of massive cardinality  
24       which will form a large inner product with a reference vector  $x \in \mathbb{R}^d$ . This procedure, known as  
25       maximum inner product search (MIPS) [27], has been applied to neural network (NN) training. In  
26       NN training, the weights of a dense layer that are estimated to produce a large inner product with the  
27       input (thereby, a large softmax, for example) are activated while the remainder are dropped out.

28       While LSH-based pruning greatly reduces training costs associated with large-scale models, popular  
29       frameworks such as SLIDE [8] and Mongoose [7] cannot be deployed in distributed settings over  
30       memory-constrained devices such as GPUs or mobile phones for the following reasons: **(a)** required  
31       maintenance of a large target layer in memory and **(b)** access to the input is needed to conduct LSH.

32       With many modern NN architectures reaching billions of parameters in size, requiring resource-  
33       constrained devices to conduct LSH analysis over even part of such a large model is infeasible  
34       as it requires many linear projections of massive weights. The hope of offloading this memory-  
35       and computationally-intensive task to a central host in the distributed setting is equally fruitless.  
36       LSH-based pruning cannot be conducted by a central host as it requires access to either local client

37 data or hashed mappings of such data. Both of these violate the fundamental host-client privacy  
38 contract especially in a federated setting [20]. Therefore, in order to maintain privacy, devices are  
39 forced to conduct LSH themselves, returning us back to our original drawback in a vicious circle. We  
40 raise the following question then:

41 *Can a resource-constrained device conduct LSH-like pruning of a large dense layer without ever*  
42 *needing to see the entirety of its underlying weight?*

43 This work makes the following contributions to positively resolve this question:

44 (1) Introduce a novel family of hash functions, PGHash, for detection of high cosine similarity  
45 amongst vectors, which improves upon the efficiency of SimHash by comparing binarized random  
46 projections over vector hashings instead. We prove several statistical properties about PGHash,  
47 including that it is an LSH family, and angle/norm distortion bounds.

48 (2) Present an algorithmic LSH framework, leveraging our hash functions, which allows for private,  
49 personalized, and memory-efficient distributed/federated training of large scale recommender net-  
50 works via dynamic pruning.

51 (3) Showcase experimentally that our PGHash-based framework is able to efficiently train large-scale  
52 recommender networks. Our approach is competitive against a distributed implementation of SLIDE  
53 [8] using full-scale LSH. Furthermore, where entry-magnitude similarity is desired over angular  
54 similarity (training over Amazon-670K, for example), we empirically demonstrate that using our  
55 DWTA [9] variant of PGHash, PGHash-D, matches the performance of using full-scale DWTA.

## 56 2 Related work

57 **LSH Families.** Locality-sensitive hashing families have been used to efficiently solve the ap-  
58 proximate nearest neighbors problem [16, 15, 2]. SimHash [6], based on randomized hyperplane  
59 projections, is used to estimate cosine similarity. Each SimHash function requires a significant  
60 number of random bits if the dimensionality of each target point is large. However, bit reduction  
61 using Nisan’s pseudorandom generator [23] is often suggested [11, 14]. MinHash [5], a competitor to  
62 SimHash [28], measures Jaccard similarity between binary vectors and has been used for document  
63 classification. The Winner Take All (WTA) hash [32] compares the ordering of entries by magnitude  
64 (corresponding to Kendall-Tau similarity); such comparative reasoning has proven popular in vision  
65 applications [25]. However, it was observed that WTA was ineffective at differentiating highly-sparse  
66 vectors leading to the development of Densified WTA (DWTA) [9]. Since MinHash, WTA, and  
67 DWTA are better suited for binary vector comparison, and we require comparison over real-valued  
68 vectors, PGHash is founded on SimHash.

69 **Hash-based Pruning.** One of the earlier proposals of pruning based on input-neuron angular  
70 similarity via LSH tables is in [29], where a scheme for asynchronous gradient updates amongst  
71 multiple threads training over a batch along with hashed backpropagation are also outlined. These  
72 principles are executed to great effect in both the SLIDE [8] and Mongoose [7] frameworks for  
73 training extreme scale recommender systems. Mongoose improves SLIDE by using an adaptive  
74 scheduler to determine when to re-run LSH over a layer weight, and by utilizing learnable hash  
75 functions. Both works demonstrated that a CPU using LSH-based dynamic dropout could achieve  
76 competitive training complexity against a GPU conducting fully-dense training. Reformer [22] uses  
77 LSH to reduce the memory complexity of self-attention layers.

78 **Distributed Recommender Systems.** Several works which prune according to input-neuron  
79 angular similarity estimations via LSH utilize multiple workers on a single machine [29, 24, 7, 8].  
80 Federated training of recommender systems is an emerging topic of interest, with particular interest  
81 in personalized training [18, 26] malicious clients [30, 36], and wireless unreliability [1]. D-SLIDE  
82 [33], which is the federated version of SLIDE, eases local on-device memory and computational  
83 requirements by sharding the network across clients. However, in the presence of low client numbers,  
84 the proportion of the model owned per device can still be taxing, whereas our compression is  
85 independent of the number of federated agents. In [31], clients query the server for weights based  
86 off the results of LSH conducted using server-provided hash functions. We regard this complete  
87 server control over the hashing family, and therefore access to hash-encoding of local client data, as  
88 non-private and potentially open to honest-but-curious attacks.

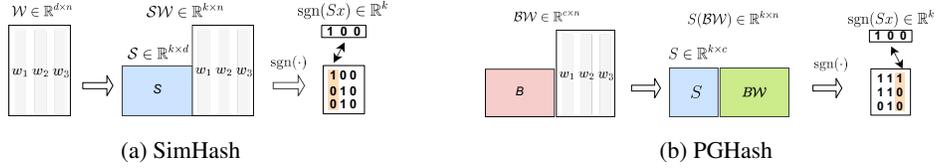


Figure 1: **Hash-based dropout.** Two procedures for conducting cosine similarity estimation between full weight matrix  $W \in \mathbb{R}^{d \times n}$  (column  $w_i$  denotes the weight of neurons  $i$ ) and an input  $x \in \mathbb{R}^d$ . **(a)** SimHash generates a hash table via left-multiplication by a randomized rectangular Gaussian matrix  $S \sim \mathcal{N}(0, I_d)$  onto the *fully-sized*  $W$ . **(b)** PGHash generates a hash table via left-multiplication by a randomized square Gaussian matrix  $S \sim \mathcal{N}(0, I_c)$  onto a base projection  $BW \in \mathbb{R}^{c \times n}$  of  $W$ . In both procedures, weight  $w_i$  is selected for activation if its signed hash matches the signed hash of  $x$ .

### 89 3 Preliminaries

90 Let  $W = \{w_1, w_2, \dots, w_n\} \subset \mathbb{R}^d$  be the weights of a dense hidden layer and  $x \in \mathbb{R}^d$  be the input.  
 91 For brevity, we refer to  $W \in \mathbb{R}^{d \times n}$  as the *weight* of the layer. In particular,  $w_i \in \mathbb{R}^d$  corresponds to  
 92 the  $i^{\text{th}}$  neuron of the layer. We assume that the layer contains  $dn$  parameters. Within our work we  
 93 perform MIPS, as we select weights which produce large inner products with  $x$ . Mathematically, we  
 94 can begin to define this by first letting  $p = \max w_i^\top x$  for  $1 \leq i \leq n$ . For  $0 < \epsilon < 1$  we are interested  
 95 in selecting  $S \subset W$ , such that for  $\forall w_i \in S, w_i^\top x > \epsilon p$ . The weights of  $S$  will pass through activation  
 96 while the rest are dropped out, reducing the computational complexity of the forward and backward  
 97 pass through this layer. As detailed in Section 4 and illustrated in Figure 1, we will determine  $S$  by  
 98 estimating angles with a projection  $BW = \{Bw_i\}_{i=1}^n$ , with  $B \in \mathbb{R}^{c \times d}$  such that  $c \ll d$ .

99 **Locality-sensitive Hashing (LSH).** LSH [2] is an efficient framework for solving the  $\epsilon$ -  
 100 approximate nearest neighbor search (NNS) problem:

101 **Definition 1** ( $\epsilon$ -NNS). *Given a set of points  $P = \{p_1, p_2, \dots, p_n\}$  in a metric space  $\mathcal{M}$  and similarity*  
 102 *function  $Sim : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$  over this space, find a point  $p \in P$ , such that for a query point  $q \in X$*   
 103 *and all  $p' \in P$ , we have that  $Sim(p, q) \leq (1 + \epsilon)Sim(p', q)$ .*

104 It is important to note that the similarity function  $Sim(\cdot)$  need not be a distance metric, but rather  
 105 any general comparison mapping. Popular choices include Euclidean distance and cosine similarity,  
 106 the latter of which is the primary focus of this paper. The cosine similarity for  $x, y \in \mathbb{R}^d$  is defined as  
 107  $\cos(x, y) \triangleq x^\top y / (||x||_2 ||y||_2)$ . We can frame the MIPS problem described previously as an  $\epsilon$ -NNS  
 108 one if we assume that the weights  $w_i \in W$  are of unit length. Thus, we are searching for  $\epsilon$ -nearest  
 109 neighbors in  $W$  of the query  $x$  according to cosine similarity.

110 Consider a family  $\mathcal{H}$  containing hash functions of the form  $h : \mathcal{M} \rightarrow \mathcal{S}$ , where  $\mathcal{S}$  is a co-domain  
 111 with significantly lower feature dimensionality than  $\mathcal{M}$ . We say that  $\mathcal{H}$  is locality-sensitive if the  
 112 hashes of a pair of points  $x, y$  in  $\mathcal{M}$ , computed by an  $h \in \mathcal{H}$  (selected uniformly at random), have a  
 113 higher collision (matching) probability in  $\mathcal{S}$  the more similar  $x$  and  $y$  are according to  $Sim$ . We now  
 114 formally define this notion following [15].

115 **Definition 2** (Locality-sensitive Hashing). *A family  $\mathcal{H}$  is called  $(S_0, \epsilon S_0, p_1, p_2)$ -sensitive if for any*  
 116 *two points  $x, y \in \mathbb{R}^d$  and  $h$  chosen uniformly at random from  $\mathcal{H}$  satisfies the following,*  
 117 *1. if  $Sim(x, y) \geq S_0 \Rightarrow Pr(h(x) = h(y)) \geq p_1$ , 2. if  $Sim(x, y) \leq \epsilon S_0 \Rightarrow Pr(h(x) = h(y)) \leq p_2$ .*

118 For an effective LSH,  $p_1 < p_2$  and  $\epsilon < 1$  is required. An LSH family allows us to conduct a similarity  
 119 search over a collection of vectors through comparison of their hashed mappings. Of course, locality  
 120 loss is inevitable if  $Sim$  is a dimension-lowering projection. Through a mixture of increased precision  
 121 (raising the output dimension of  $\mathcal{H}$ ) and repeated trials (running several trials over independently  
 122 chosen  $h$ ), we may tighten the correspondence between  $Sim$  and matches over  $\mathcal{H}$ , following the  
 123 spirit of the Johnson-Lindenstrauss Lemma [19].

124 **SimHash.** A popular LSH algorithm for estimating cosine similarity is SimHash, which uses  
 125 signed random projections [6] as its hash functions. Specifically, for a collection of vectors  $W \subset \mathbb{R}^d$ ,  
 126 the SimHash family  $\mathcal{H}^{Sim}$  consists of hash functions  $h_v$ , each indexed by a random Gaussian vector  
 127  $v \sim \mathcal{N}(0, I_n)$ , i.e., an  $n$ -dimensional vector with iid entries drawn from  $\mathcal{N}(0, 1)$ . For  $x \in \mathbb{R}^n$ ,  
 128 we define the hash mapping  $h_v(x) := \text{sgn}(v^\top x)$ . Here, we modify  $\text{sgn}(x)$  to return 1 if  $x > 0$ ,

129 else it returns 0. For Gaussian  $v$  chosen uniformly at random and fixed  $x, y \in \mathbb{R}^d$ , we have  
 130  $Pr(h_v(x) = h_v(y)) = 1 - \frac{\arccos(\frac{x^\top y}{\|x\| \|y\|})}{\pi}$ . This hashing scheme was popularized in [12] as part  
 131 of a randomized approximation algorithm for solving MAX-CUT. Notice that the probability of  
 132 a hashed pair matching is monotonically increasing with respect to the cosine similarity of  $x$  and  
 133  $y$ , satisfying Definition 2. More precisely, if we set  $S_0 = \cos(x, y)$  then  $\mathcal{H}^{Sim}$  is  $\left(S_0, \epsilon S_0, \left(1 - \arccos\left(\frac{S_0}{\pi}\right)\right), \left(1 - \arccos\left(\frac{\epsilon S_0}{\pi}\right)\right)\right)$ -sensitive [6, 28]. The above discussion considers the sign of a  
 134 single random projection, but in practice we will perform multiple projections.

136 **Definition 3** (Hash table). Let  $X = \{x_1, \dots, x_n\} \subset \mathbb{R}^d$  and  $V = \{v_1^\top, \dots, v_k^\top\} \subset \mathcal{N}(0, I_d)$ , where  
 137  $k$  is the **hash length**. Define  $h_V : \mathbb{R}^d \rightarrow \mathbb{R}^k$  by  $[h_V(x)]_i = h_{v_i}(x)$  where  $h_{v_i} \in \mathcal{H}^{Sim}$  for  $1 \leq i \leq k$ .  
 138 For fixed  $V$ , the **hash table**  $h_V(X) \in \mathbb{R}^{k \times n}$  is a binary matrix with columns  $h_V(x_j)$  for  $1 \leq j \leq n$ .

139 Following the notation above, we may estimate similarity between an input  $q \in \mathbb{R}$  and a collection of  
 140 vectors  $X \subset \mathbb{R}^d$  by measuring the Hamming distances (or exact sign matches) between  $h_V(q)$  and  
 141 columns of  $h_V(X)$ . SimHash is now more discriminatory, as  $\mathcal{H}^{Sim}$  can separate  $\mathbb{R}^d$  into  $2^k$  buckets  
 142 corresponding to all possible length  $k$  binary vectors (which we refer to as **hash codes**). Finally,  
 143 counting the frequency of exact matches or computing the average Hamming distance over several  
 144 independently generated hash tables further improves our estimation of closeness. Implementations  
 145 of the well-known SLIDE framework [8, 24], which utilize SimHash for LSH-based weight pruning,  
 146 require upwards of 50 tables.

147 **DWTA**. Another popular similarity metric is to measure how often high-magnitude entries between  
 148 two vectors occur at the exact same positions. The densified winner-take-all (DWTA) LSH family  
 149 [10] estimates this similarity by uniformly drawing  $k$  random coordinates over  $W$  and recording the  
 150 position of the highest-magnitude entry. Similar to SimHash, this process is repeated several times,  
 151 and vectors with the highest frequency of matches are expected to have similar magnitude ordering.  
 152 This type of comparative reasoning is useful for computer vision applications [37].

## 153 4 PGHash

154 In this section, we develop a family of hash functions  $\mathcal{H}^{PG}$  which allow for memory-efficient serial  
 155 generation of hash tables using a single dimensionally-reduced sketch of  $W$ . This is in contrast  
 156 to traditional LSH frameworks, which produce hash tables via randomized projections over the  
 157 entirety of  $W$ . We first present an overview of the Periodic Gaussian Hash (PGHash) followed by its  
 158 algorithm for distributed settings, an exploration of several statistical properties regarding the local  
 159 sensitivity of  $\mathcal{H}^{PG}$ .

160 **PGHash Motivation**. As detailed in Section 3, our goal is to efficiently estimate cosine similarity  
 161 between an input to a layer  $x \in \mathbb{R}^d$  and the columns of a large weight matrix  $W \in \mathbb{R}^{d \times n}$ . SimHash  
 162 performs hash table generation by first multiplying a matrix of uniformly drawn Gaussian hyperplanes  
 163  $T_V \in \mathbb{R}^{c \times d}$  with  $W$ . The full hash table is computed as  $h_V(W) = \text{sgn}(T_V W)$ . Then, the  $j^{\text{th}}$  neuron  
 164 is activated if  $\text{sgn}(T_V x) = h_V(w_j)$  for a layer input  $x$ .

165 One can immediately notice that generation of a new hash table  $h_{\tilde{V}}(W)$  requires both **(i)** computation  
 166 of  $T_{\tilde{V}} W$  which requires access to the fully-sized weights and **(ii)** the storage of  $T_{\tilde{V}}$  to compute  
 167  $\text{sgn}(T_{\tilde{V}} x')$  for further inputs  $x'$ . This is problematic for a memory-constrained device, as it would  
 168 need to maintain both  $W$  and  $T_{\tilde{V}}$  to generate further tables and perform dynamic pruning. To solve  
 169 this issue, we introduce a family of hash functions generated from a single projection  $BW$  of  $W$ .

### 170 4.1 PGHash theory

171 **Definition 4** (Periodic Gaussian Hash). Assume **sketch dimension**  $c \ll d$  divides  $d$  for simplicity.  
 172 Let  $B = [I_c | I_c | \dots | I_c] \in \mathbb{R}^{c \times d}$ , where  $I_c$  is the  $c \times c$  identity matrix and  $|$  denotes  $\frac{d}{c}$  concatenations.  
 173 Let  $S \in \mathbb{R}^{k \times c}$  be a random Gaussian matrix with iid entries drawn from  $\mathcal{N}(0, 1)$ . We may define  
 174 a **Periodic Gaussian Hash (PGHash)** function  $h_S^{PG} : \mathbb{R}^d \rightarrow \mathbb{R}^k$  by  $[h_S^{PG}(x)]_i = [\text{sgn}(SBx)]_i$  for  
 175  $1 \leq i \leq k$ . We denote the family of all such hash functions as  $\mathcal{H}^{PG}(c, d)$ .

176 We use the term “periodic” to describe the hash functions described in Definition 4, since unlike  
 177 SimHash which projects a point via a fully random Gaussian vector as in SimHash, our projection

178 is accomplished using a repeating concatenation of a length  $c$  Gaussian vector. Furthermore, for  
 179  $h_S^{PG} \subset H^{PG}(c, d)$ , the matrix representation  $SB$  is a tiling of a single Gaussian matrix. Notice that  
 180 we may easily extend the notion of a PGHash of one vector to an entire hash table over multiple vectors  
 181 following Definition 3. In this manner, we may generate a sequence of hash tables  $\{h_{S_i}^{PG}(W)\}_{i=1}^\tau$   
 182 over a weight matrix  $W$  simply by drawing random Gaussian matrices  $S_i \in \mathbb{R}^{k \times c}$  for  $1 \leq i \leq \tau$   
 183 (where  $k$  is the hash length) and computing  $\text{sgn}(S_i BW)$ .

184 **Extension to DWTA (PGHash-D) Remark.** When DWTA (described in Section 3) is preferred  
 185 over SimHash for similarity estimation, we may modify Definition 4 as follows:  $B = D_\perp P$  where  $P$   
 186 is a  $d \times d$  random permutation matrix, and  $D_\perp$  is a  $c \times d$  rectangular diagonal matrix with  $D_{ii} = 1$   
 187 for  $1 \leq i \leq c$ . We denote our hash functions as  $h_S^{PG-D} : \mathbb{R}^d \rightarrow \mathbb{R}^k$  by  $h_S^{PG-D}(x) = \max_i [SBx]_i$   
 188 for  $1 \leq i \leq k$ , with  $k \leq c$ , where  $S$  is now a rectangular permutation matrix which selects  $k$  rows of  
 189  $Bx$  at random. We refer to this scheme as PGHash-D, whereas PGHash refers to Definition 4.

190 **Local Memory Complexity Remark.** When generating a new table using PGHash, a device  
 191 maintains  $S$  and needs access to just  $BW$ , which cost  $\mathcal{O}(kc)$  and  $\mathcal{O}(cn)$  space complexity respectively.  
 192 This is much smaller than the  $\mathcal{O}(kd)$  and  $\mathcal{O}(dn)$  local memory requirements of SimHash.

193 **Sensitivity of  $\mathcal{H}^{PG}$ .** In this section, we will explore the sensitivity of  $\mathcal{H}^{PG}(c, d)$ .

194 **Definition 5.** Let  $x \in \mathbb{R}^d$  and  $c \in \mathbb{R}$  such that  $c|d$ . Define the  $(d, c)$ -**folding**  $x_c \in \mathbb{R}^c$  of  $x$  as  
 195  $[x_c]_i = \sum_{j=1}^{\frac{d}{c}} [x]_{i+j \cdot \frac{d}{c}}$ . Equivalently,  $x_c = Bx$ , with  $B$  as specified in Definition 4.

196 **Theorem 1.** Let  $x, y \in \mathbb{R}^d$ . Define the following similarity function  $\text{Sim}_c^d(x, y) \triangleq \cos(x_c, y_c)$ ,  
 197 where  $x_c, y_c$  are  $(d, c)$ -foldings of  $x, y$ .  $\mathcal{H}^{PG}(c, d)$  is an LSH family with respect to  $\text{Sim}_c^d$ .

198 *Proof.* Let  $h_v \in \mathcal{H}^{PG}$ . This means that for a randomly chosen  $v' \sim \mathcal{N}(0, I_{\frac{d}{c}})$ ,  $v$  is a  $c$ -times  
 199 concatenation of  $v$ . We see that  $\text{sgn}(v^\top x) = \text{sgn}\left(\frac{v^\top x}{\|v\| \cdot \|x\|}\right) = \text{sgn}\left(\sqrt{\frac{c}{d}} \frac{\|x_c\|}{\|x\|} \frac{v'^\top x_c}{\|v'\| \|x_c\|}\right)$ . Since  $\text{sgn}$  is  
 200 unaffected by the positive multiplicative factors, we conclude that  $\text{sgn}(v^\top x) = \text{sgn}(v'^\top x_c)$ . Through  
 201 symmetric argument, we find  $\text{sgn}(v^\top y) = \text{sgn}(v'^\top y_c)$ . Since  $v' \sim \mathcal{N}(0, I_c)$ , comparing the sign of  
 202  $v^\top x$  to  $v^\top y$  is equivalent to a standard SimHash over  $x_c$  and  $y_c$ , i.e., estimation of  $\cos(x_c, y_c)$ .  $\square$

203 **Corollary 1.** Let  $x, y \in \mathbb{R}^d$ , then  $\mathcal{H}^{PG}(c, d)$  is  $\left(S_c, \epsilon S_c, \left(1 - \arccos\left(\frac{S_c}{\pi}\right), \left(1 - \arccos\left(\frac{\epsilon S_c}{\pi}\right)\right)\right)\right)$ -  
 204 sensitive where  $S_c = \cos(x_c, y_c)$ .

205 *Proof.* This follows directly from the well-known sensitivity of SimHash [6].  $\square$

206 We see that  $\mathcal{H}^{PG}$  is LSH with respect to the angle between  $(d, c)$ -foldings of vectors. The use  
 207 of periodic Gaussian vectors restricts the degrees of freedom (from  $d$  to  $d/c$ ) of our projections.  
 208 However, the usage of pseudo-random and/or non-iid hash tables has been observed to perform  
 209 well in certain regimes [3, 35]. *Although  $\mathcal{H}^{PG}(c, d)$  is LSH, is  $\cos(x_c, y_c)$  necessarily an acceptable*  
 210 *proxy for  $\cos(x, y)$ , in particular, for high angular similarity?* Heuristically, yes, for highly-cosine  
 211 similar vectors: assuming  $x$  and  $y$  are both unit (since scaling does not affect angle) then we have that  
 212  $\|x - y\|^2 = 2 - 2\cos(x, y)$ . If  $\ell_2$  similarity between  $x$  and  $y$  is already high, then the  $\ell_2$  similarity  
 213 of their (normalized)  $(d, c)$ -foldings will also be high, and thus their cosine similarity as well. We  
 214 now provide a characterization on the angle distortion of a  $(d, c)$ -folding.

215 **Theorem 2.** Let  $x, y \in \mathbb{S}^{d-1}$ . Assume that neither  $x$  nor  $y$  vanish under multiplication by  $B$  and that  
 216 the set  $S_{x,y} = \{v \in \text{span}(x, y) : \|v\| = 1\}$  does not contain a 0-eigenvector of  $B^\top B$ . We denote the  
 217 following quantities:  $\lambda = d/c$ ,  $\theta := \frac{1}{2} \arccos(\cos(x, y))$ ,  $\alpha = \inf\{c > 0 \mid \|Bv\| \geq c, \forall v \in S_{x,y}\}$ ,  
 218 and  $\beta = \alpha/\lambda$ . ( $\alpha > 0$  since  $S_{x,y}$  does not contain any 0-eigenvectors.) Then  $\cos(x_c, y_c)$  lives  
 219 between  $\frac{1-\beta^2 \tan^2 \theta}{1+\beta^2 \tan^2 \theta}$  and  $-\frac{\tan^2 \theta^2 - \beta^2}{\tan^2 \theta^2 + \beta^2}$ .

220 *Proof sketch.* Consider the unit circle  $S_{x,y}$  contained in  $\text{span}(x, y)$  (Let us assume  $x$  and  $y$  are unit,  
 221 WLOG). The linear distortion  $BS_{x,y}$  is an ellipse containing  $x_c = Bx$  and  $y_c = By$ . The length of  
 222 the axes of this ellipse are determined by the eigenvalues of  $B^\top B$ . The bounds follow from further  
 223 trigonometric arguments, by considering when the axes of  $BS_{x,y}$  are maximally stretched and shrunk  
 224 respectively. These distortions are strongly related to  $\lambda$  and  $\beta$ .

225 We can see that as  $\cos(x, y) \rightarrow 0$  we  
 226 have  $\cos(x_c, y_c) \rightarrow 0$ . It is natural to  
 227 consider the distribution of  $\alpha$  in Theorem 2 as how extreme shrinking by  $B$   
 228 (the folding matrix) can greatly distort the angle. We can characterize this statistical  
 229 distribution exactly.

232 **Proposition 1.** *Let  $u \in S^{d-1}$ ,  
 233 drawn uniformly at random. Then  $\|Bu\|^2 \sim \text{Beta}(\frac{\xi}{2}, \frac{d-c}{2}, 0, \frac{d}{c})$ , the four  
 234 parameter Beta distribution with pdf  
 235  $f(x) = \frac{(2x/c)^{c/2-1}(1-2x/c)^{(d-c)/2-1}}{(d/c)\beta(c/2, (d-c)/2)}$   
 236 and  $\mathbb{E}\|Bu\|^2 = 1$ .*

238 We defer proof of Proposition 1 to the  
 239 Appendix D. Since the folded magnitude  
 240  $\|Bu\|^2$  is unit in expectation, the distortion  
 241 term  $\alpha^2/\lambda^2$  in Theorem 2 will often  
 242 be close to 1, greatly tightening the angle  
 243 distortion bounds.

## 244 4.2 PGHash algorithm

245 Below, we detail our protocol for deploy-  
 246 ing PGHash in a centralized distributed  
 247 setting (presented algorithmically in Al-  
 248 gorithm 1). Over a network of  $N$  de-  
 249 vices, the central host identifies a target layer  $P$  whose weight  $W_t^P \in \mathbb{R}^{d \times n}$  (at iteration  $t$ ) is too  
 250 expensive for memory-constrained devices to fully train or host in local memory. Neurons (columns  
 251 of  $W_t^P$ ) are pruned by devices according to its estimated cosine similarity to the output  $x^{P-1}$  of the  
 252 previous layer.

253 The central host begins each round of dis-  
 254 tributed training  $t$  by sending each device (1)  
 255 all weights  $\{W_t^\ell\}_{\ell=1}^{P-1}$  required to generate the  
 256 input  $x^{P-1}$  and (2) the compressed target layer  
 257  $BW_t^P$ . Using these weights, each device con-  
 258 ducts PGHash analysis (via Algorithm 2) us-  
 259 ing its current batch of local data to deter-  
 260 mine its activated neurons. The central host  
 261 sends each device  $i$  their set of activated neu-  
 262 rons  $W_t^P(\Theta_i)$ , and each device performs a  
 263 standard gradient update on their new model  
 264  $\{W_t^\ell\}_{\ell=1}^{P-1} \cup W_t^P(\Theta_i) \cup \{W_t^\ell\}_{\ell=P+1}^L$ . Finally,  
 265 the central host receives updated models from  
 266 each device and averages only the weights which  
 267 are activated during training.

268 The on-device PGHash analysis (Algorithm 2)  
 269 consists of first running a forward pass up to the  
 270 target layer to generate the input  $x^{P-1}$ . Devices  
 271 generate personal hash tables by performing left-  
 272 multiplication of  $BW_t^P$  and  $x^{P-1}$  by a random  
 273 Gaussian matrix  $S \sim \mathcal{N}(0, I_c)$ , as described in  
 274 Section 3. The number of tables  $\tau$  is specified  
 275 by the user. A neuron  $j$  is marked as active if the input hash code  $\text{sgn}(Sx^{P-1})$  is identical to  $j$ -th  
 276 weight’s hash code  $\text{sgn}(SB[W_t^P]_{:,j})$ . In Appendix A.2, we detail how Hamming distance can also  
 277 be used for neuron selection (by selecting neurons which have the lowest average distance).

---

### Algorithm 1 Distributed PGHash

---

**Require:** weights  $\{W_0^\ell\}_{\ell=1}^L$ , objective  $\mathcal{F}$ ,  $N$  devices, tar-  
 get layer  $P$ ,  $T$  iterations, folding matrix  $B$ ,  $\tau$  hash  
 tables, hash length  $k$ , sketch dim  $c$ , comp. ratio  $CR$ .

- 1: **Server Executes:**
- 2: **for**  $t = 1, 2, \dots, T$  **do**
- 3:   Compile pre-target weights  $W_A = \{W_t^\ell\}_{\ell=1}^{P-1}$
- 4:   Compile post-target weights  $W_B = \{W_t^\ell\}_{\ell=P+1}^L$
- 5:   **for each device**  $i$  **in parallel do**
- 6:      $\Theta_i \leftarrow \text{DeviceLSH}(i, W_A, BW_t^P, CR, \tau, k, c)$
- 7:      $W_{t+1,i} \leftarrow \text{DeviceUpdate}(i, W_t^P(\Theta_i) \cup W_B)$
- 8:   **end for**
- 9:    $W_{t+1} \leftarrow$  Average the active weights  $W_{t+1,i}$   
 across all devices
- 10: **end for**
- 11: **return**  $\{W_T^\ell\}_{\ell=1}^L$
- 12: **DeviceLSH**( $i, W_A, BW, CR, \tau, k, c$ ):
- 13: Sample input  $\xi$  from local data distribution  $\mathcal{D}_i$
- 14:  $\xi^{P-1} = \mathcal{F}(W_A, \xi)$   $W_A$  is stored locally
- 15:  $\Theta_i \leftarrow \text{PGHash}(\xi^{P-1}, BW, CR, \tau, k, c)$
- 16: **return**  $\Theta_i$
- 17: **DeviceUpdate**( $i, W_B$ ):
- 18:  $W \leftarrow W_A \cup W_B$
- 19: **return**  $W - \eta \nabla_W \mathcal{F}(W; \xi)$

---



---

### Algorithm 2 PGHash

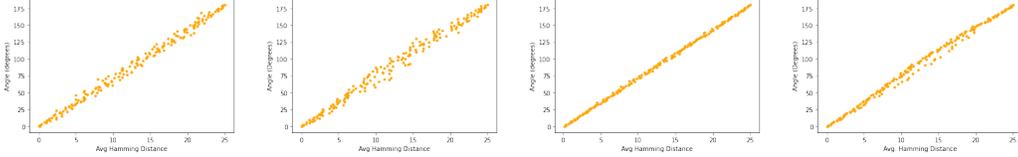
---

**Require:** batched input  $X \in \mathbb{R}^{d \times M}$ , projected  
 weight  $BW \in \mathbb{R}^{c \times n}$ , compression rate  $CR$ ,  $\tau$   
 hash tables, hash length  $k$ , sketch dim  $c$ .

- 1: Set  $\Theta = [ ]$
- 2: **for**  $i = 1, 2, \dots, \tau$  **do**
- 3:   Draw random Gaussian  $S \in \mathbb{R}^{k \times c}$
- 4:   **for each sample**  $x$  in  $X$  **do**
- 5:     **for**  $j = 1, 2, \dots, n$  **do**
- 6:       **if**  $\text{sgn}(SBx) = [\text{sgn}(SBW)]_{:,j}$  **then**
- 7:          $\Theta.\text{add}(j)$
- 8:       **end if**
- 9:     **end for**
- 10:   **end for**
- 11:   **if**  $|\Theta| > \lfloor CR \cdot n \rfloor$  **then**
- 12:     Randomly remove selected neurons  
 from table- $i$  until  $|\Theta| = \lfloor CR \cdot n \rfloor$
- 13:   **return**  $\Theta$
- 14:   **end if**
- 15: **end for**
- 16: **return**  $\Theta$

---

the input hash code  $\text{sgn}(Sx^{P-1})$  is identical to  $j$ -th  
 weight’s hash code  $\text{sgn}(SB[W_t^P]_{:,j})$ . In Appendix A.2, we detail how Hamming distance can also  
 be used for neuron selection (by selecting neurons which have the lowest average distance).



(a) SimHash (10 tables). (b) PGHash (10 tables). (c) SimHash (100 tables). (d) PGHash (100 tables).  
**Figure 2: Correlation between angle and Hamming distance.** We plot the average Hamming distance (x-axis) between a PG/SimHash hashed fixed vector  $x$  and collection of vectors  $\mathcal{V}$  versus their true angles (y-axis). Vectors are unit, length 100, and hashed down to dimension  $k = 25$  binary vectors according to  $\tau = 10$  or 100 hash tables. PGHash has a sketch dimension of  $c = 25$ . Both PGHash and SimHash show strong correlation between Hamming distance and angular similarity.

278 **Computational Complexity Remark.** Through the use of dynamic pruning, PGHash significantly  
 279 reduces both the forward and backward training computational complexities. PGHash activates *at*  
 280 *most*  $CR \cdot n$  neurons per sample as opposed to  $n$  for full training. In practice, PGHash activates only  
 281 a fraction of the  $CR \cdot n$  neurons (as shown in Figure 6a). Therefore, the number of floating point  
 282 operations within forward and backward training is dramatically reduced.

283 **Communication Complexity Remark.** By reducing the size of the model needed in local  
 284 memory and subsequently requesting a pruned version of the architecture we improve communication  
 285 efficiency. For a fixed number of rounds  $T$  and target weight size  $dn$ , the total communication  
 286 complexity, with respect to this data structure, is  $\mathcal{O}(T \cdot CR \cdot dn)$ , which significantly less bits than  
 287 the vanilla  $\mathcal{O}(Tdn)$  communication cost of vanilla federated training. In Section 5, we show that  
 288 PGHash achieves near state-of-the-art results with only  $CR = 0.1$  (10% of a massive weight matrix).

## 289 5 Experiments

290 In this section, we (1) gauge the sensitivity of PGHash and (2) analyze the performance of PGHash  
 291 and our own DWTA variant (PGHash-D) in training large-scale recommender systems. PGHash and  
 292 PGHash-D require only 6.25% ( $c = 8$ ) of the final layer sent by the server to perform on-device  
 293 LSH in our experiments. In PGHash, devices receive the compressed matrix  $BW \in \mathbb{R}^{c \times n}$  via  
 294 the procedure outlined in Section 4. In PGHash-D, devices receive  $c$  out of  $d$  randomly selected  
 295 coordinates for all  $n$  neurons in the final layer weight. Using  $k$  of the  $c$  coordinates (ensuring privacy  
 296 since the server is unaware of the coordinates used for LSH), PGHash-D selects neurons which, within  
 297 the  $k$  coordinates, share the same index of highest-magnitude entry between the input and weight.  
 298 We employ PGHash for Delicious-200K and PGHash-D for Amazon-670K and WikiLSHTC-325K.

299 **PGHash Sensitivity Analysis.** Our first experiment measures the ability of  $\mathcal{H}^{PG}(c, d)$  to estimate  
 300 cosine similarity. We produce a fixed unit vector  $x \in \mathbb{R}^{100}$  and set of 180 vectors  $\{v_i\}_{i=1}^{180}$  of the same  
 301 dimension. Both the Gaussian vector  $x$  and collection of vectors  $V$  are fed through varying numbers  
 302 of SimHash and PGHash tables. We produce a scatter plot measuring the correlation between angle  
 303 and average Hamming distance. PGHash, as seen in Figure 2, is an effective estimator of cosine  
 304 similarity. We observe that PGHash, like SimHash, successfully produces low average Hamming  
 305 distances for vectors that are indeed close in angle. This provides evidence that selecting neurons  
 306 with exact hash code matches (vanilla sampling) is effective for choosing neurons which are close in  
 307 angle to the input vector. Finally, we find increasing the number of hash tables helps reduce variance.

308 **Large-Scale Recommender System Training.** Our second experiment tests how well PGHash(-D)  
 309 can train large-scale recommender systems. We train these networks efficiently by utilizing dynamic  
 310 neuronal dropout as done in [8]. We use three extreme multi-label datasets for training recommender  
 311 systems: Delicious-200K, Amazon-670K, and WikiLSHTC-325K. These datasets come from the  
 312 Extreme Classification Repository [4]. The dimensionality of these datasets is large: 782,585/205,443  
 313 (Delicious-200K), 135,909/670,091 (Amazon-670K), and 1,617,899/325,056 (WikiLSHTC-325K)  
 314 features/labels. Due to space, Wiki results are found in Appendix A.3.

315 The feature and label sets of these datasets are extremely sparse. Akin to [8, 7, 34], we train a  
 316 recommender system using a fully-connected neural network with a single hidden layer of size 128.  
 317 Therefore, for Amazon-670K, our two dense layers have weight matrices of size  $(135,909 \times 128)$  and  
 318  $(128 \times 670,091)$ . The final layer weights output logits for label prediction, and we use PGHash(-D)  
 319 to prune its size to improve computational efficiency during training.

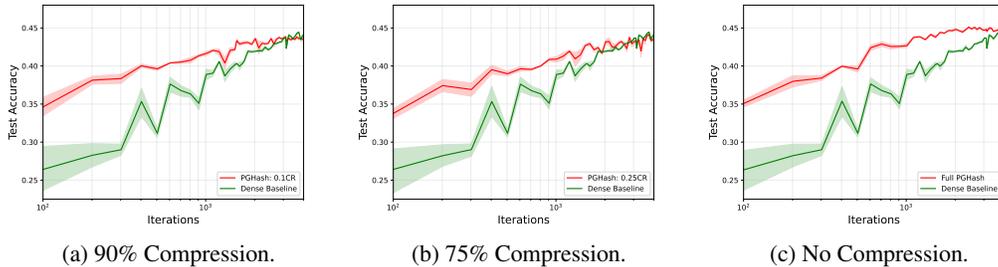


Figure 3: **Compressed PGHash.** We record model accuracy of a large recommendation system on an extreme classification task (Delicious-200K) using PGHash for varying compression rates ( $CR$ ). Compressed PGHash, even at 90% compression, is competitive with full training (*without even including effects of sparsity-induced neuronal drop out*). Hyperparameters are in Appendix A.1.

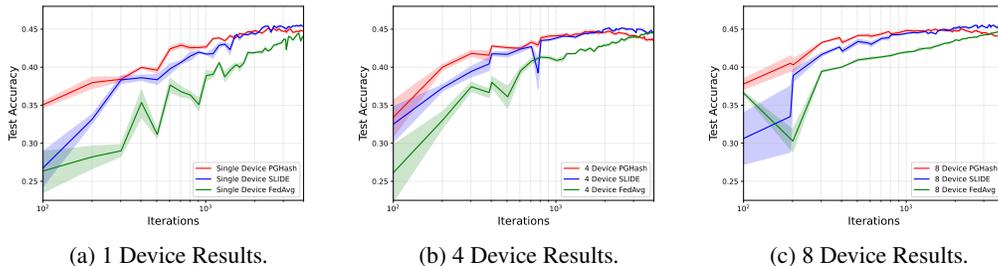


Figure 4: **Federated Delicious-200K PGHash.** We record model accuracy of a large recommendation system on an extreme classification task (Delicious-200K) trained in a federated setting. PGHash achieves competitive accuracies compared with Federated SLIDE and FedAvg. In fact, PGHash converges quicker to near-optimal accuracy. Hyperparameters are in Appendix A.1.

320 Unlike [8, 7, 34], PGHash(-D) can be deployed in a federated setting. Within our experiments, we  
 321 show the efficacy of PGHash for both single- and multi-device settings. Training in the federated  
 322 setting (following the protocols of Algorithm 1) allows each device to rapidly train portions of the  
 323 entire neural network in tandem. We partition data evenly (in an IID manner) amongst devices.  
 324 Finally, we train our neural network using TensorFlow. We use the Adam [21] optimizer with an  
 325 initial learning rate of  $1e-4$ . A detailed list of the hyper-parameters we use in our experiments can be  
 326 found in Appendix A.1. Accuracy in our figures refers to the  $P@1$  metric, which measures whether  
 327 the predicted label with the highest probability is within the true list of labels. These experiments are  
 328 run on a cloud cluster using Intel Xeon Silver 4216 processors with 128GB of total memory.

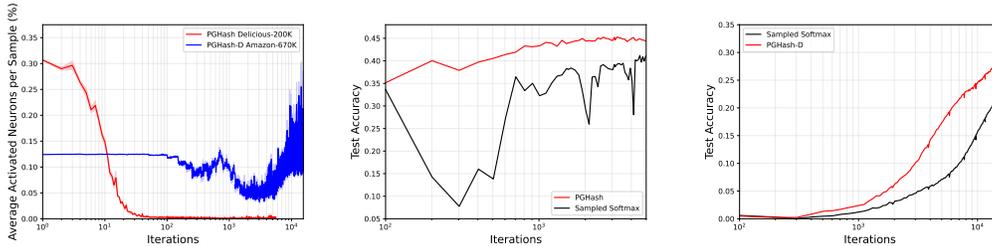
329 **Sampling Strategy.** One important aspect of training is how we select activated neurons for each  
 330 sample through LSH. Like [8], we utilize vanilla sampling. In our vanilla sampling protocol, a total  
 331 of  $CR \cdot n$  neurons are selected across the entire sampled batch of data. As detailed in Section 4 and  
 332 Algorithm 2, a neuron is selected when its hash code exactly matches the hash code of the input. We  
 333 retrieve neurons until either  $CR \cdot n$  are selected or all  $\tau$  tables have been looked up.

334 **Compression Efficacy.** We begin by analyzing how PGHash performs when varying the compression  
 335 rate  $CR$ . Figure 3 showcases how PGHash performs for compression rates of 75% and 90% as  
 336 well as no compression. Interestingly, PGHash reaches near-optimal accuracy even when compressed.  
 337 This shows the effectiveness of PGHash at accurately selecting fruitful active neurons given a batch  
 338 of data. The difference between the convergence of PGHash for varying compression rates lies within  
 339 the volatility of training. As expected, PGHash experiences more volatile training (Figures 3a and  
 340 3b) when undergoing compression as compared to non-compressed training (Figure 3c).

341 **Distributed Efficacy.** In Figures 4 and 5, we analyze how well PGHash(-D) performs in a federated  
 342 setting. We compare PGHash(-D) to a federated version of SLIDE [8] that we implemented (using  
 343 respectively, a full SimHash or DWTA), as well as fully-dense Federated Averaging (FedAvg) for  
 344 Delicious-200K. One can immediately see in Figures 4 and 5 that PGHash(-D) performs identically  
 345 to, or better than, Federated SLIDE. In fact, for Delicious-200K, PGHash and Federated SLIDE  
 346 outperform the dense baseline (FedAvg). In Appendix A.3 we detail the difficulties of PGHash-D  
 347 and SLIDE in matching the dense baseline as well as the failure of SimHash to achieve performance  
 348 akin to DWTA for Amazon-670K.

349 PGHash(-D) and Federated SLIDE smartly train  
 350 portions of the network related to each batch  
 351 of local device data, via LSH, in order to make  
 352 up for the lack of a full output layer. However,  
 353 unlike Federated SLIDE, PGHash(-D) can per-  
 354 form on-device LSH *using as little as 6.25% of*  
 355 *the full weight  $W$  ( $c = 8$ )* for both Delicious-  
 356 200K and Amazon-670K experiments. Further-  
 357 more, for Delicious-200K, PGHash generates a  
 358 dense Gaussian that is only 6.25% ( $c = 8$ ) the  
 359 size of that for Federated SLIDE. In summary,  
 360 PGHash(-D) attains similar performance to Fed-  
 361 erated SLIDE while storing less than a tenth of  
 362 the parameters.

363 **Induced Sparsity.** PGHash(-D) induces a large amount of sparsity through its LSH process. This  
 364 is especially prevalent in large-scale recommender systems, where the number of labels for each data  
 365 point is a miniscule fraction of the total output layer size (*e.g.* Delicious-200K has on average only  
 366 75.54 labels per point). PGHash(-D) performs well at identifying this small subset of neurons as  
 367 training progresses. As one can see in Figure 6a, even when PGHash is allowed to select all possible  
 368 neurons (*i.e.*, no compression  $CR = 1$ ), it still manages to *select fewer than 1% of the total neurons*  
 369 *after only 50 iterations of training over Delicious-200K*. For Amazon-670K, PGHash-D requires  
 370 less than 30% of the total neurons for the majority of training even. Therefore, PGHash(-D) greatly  
 371 increases the amount of sparsity within the NN, subsequently improving the computational efficiency  
 372 of the algorithm by reducing the number of floating point operations required in the forward and  
 373 backward training.



(a) 1 Device Results. (b) 4 Device Results.  
 Figure 5: **Federated Amazon-670K PGHash-D.** We record model accuracy of a large recommendation system on Amazon-670K trained in a federated setting. PGHash-D matches the convergence of Federated SLIDE without requiring LSH to be performed by the central server.



(a) Average Activated Neurons. (b) Sampled Softmax (Delicious). (c) Sampled Softmax (Amazon).  
 Figure 6: **PGHash(-D) Computational Efficiency.** In Figure 6a, we showcase that PGHash(-D) activates only a fraction of the total final layer neurons *even without compression*. Through this induced sparsity, PGHash(-D) greatly reduces the computational complexity of forward and backward training compared to full training. In Figures 6b and 6c, we compare PGHash(-D) with the Sampled Softmax heuristic [17] (randomly sampling 10% of the total neurons) for efficiently training recommender systems. PGHash(-D) outperforms the Sampled Softmax baseline, as it selects a better set of activated neurons via LSH to more efficiently train the recommender system.

## 374 6 Conclusion

375 In this work, we present a new hashing family, PGHash, which enables the generation of multiple  
 376 LSH hash tables using a single base projection of a massive target weight. These hash tables can be  
 377 used to dynamically select for neurons which are similar to the layer input. This alleviates memory,  
 378 communication, and privacy costs associated with conventional LSH-training approaches. As a proof  
 379 of concept, we demonstrate that (i) the PGHash family is effective at mimicking SimHash and (ii) our  
 380 framework is competitive against other, memory-inefficient, LSH-based federated training baselines  
 381 of large-scale recommender networks. For future work, we intend to explore how multi-layer PGHash  
 382 pruning affects model performance and incorporate learnable hashes as in the Mongoose [7] pipeline.

383 **Limitations.** Our theory indicates that PGHash is useful for detecting high angular similarity, but  
 384 could prove unreliable for differentiating between intermediately dissimilar vectors. Additionally,  
 385 LSH-based pruning has only shown success on large classification layers or attention layers in  
 386 transformers [22]. When considering broader impacts, large-scale recommender systems, and any  
 387 subsequent improvements to their design, can be used for strategically negative advertising purposes.

388 **References**

- 389 [1] Zareen Alamgir, Farwa K Khan, and Saira Karim. Federated recommenders: methods, chal-  
390 lenges and future. *Cluster Computing*, 25(6):4075–4096, 2022.
- 391 [2] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest  
392 neighbor in high dimensions. *Communications of the ACM*, 51(1):117–122, 2008.
- 393 [3] Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya Razenshteyn, and Ludwig Schmidt.  
394 Practical and optimal lsh for angular distance. *Advances in neural information processing*  
395 *systems*, 28, 2015.
- 396 [4] Kush Bhatia, Kunal Dahiya, Himanshu Jain, Anshul Mittal, Yashoteja Prabhu, and Manik Varma.  
397 The extreme classification repository: Multi-label datasets and code. URL <http://manikvarma.org/downloads/XC/XMLRepository.html>, 2016.
- 399 [5] Andrei Z Broder. On the resemblance and containment of documents. In *Proceedings. Com-*  
400 *pression and Complexity of SEQUENCES 1997 (Cat. No. 97TB100171)*, pages 21–29. IEEE,  
401 1997.
- 402 [6] Moses S Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings*  
403 *of the thirty-fourth annual ACM symposium on Theory of computing*, pages 380–388, 2002.
- 404 [7] Beidi Chen, Zichang Liu, Binghui Peng, Zhaozhuo Xu, Jonathan Lingjie Li, Tri Dao, Zhao  
405 Song, Anshumali Shrivastava, and Christopher Re. Mongoose: A learnable lsh framework for  
406 efficient neural network training. In *International Conference on Learning Representations*,  
407 2020.
- 408 [8] Beidi Chen, Tharun Medini, James Farwell, Charlie Tai, Anshumali Shrivastava, et al. Slide: In  
409 defense of smart algorithms over hardware acceleration for large-scale deep learning systems.  
410 *Proceedings of Machine Learning and Systems*, 2:291–306, 2020.
- 411 [9] Beidi Chen and Anshumali Shrivastava. Revisiting winner take all (wta) hashing for sparse  
412 datasets. *arXiv preprint arXiv:1612.01834*, 2016.
- 413 [10] Beidi Chen and Anshumali Shrivastava. Densified winner take all (wta) hashing for sparse  
414 datasets. In *Uncertainty in artificial intelligence*, 2018.
- 415 [11] Lars Engebretsen, Piotr Indyk, and Ryan O’Donnell. Derandomized dimensionality reduction  
416 with applications. 2002.
- 417 [12] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum  
418 cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*,  
419 42(6):1115–1145, 1995.
- 420 [13] Arjun K Gupta and Saralees Nadarajah. *Handbook of beta distribution and its applications*.  
421 CRC press, 2004.
- 422 [14] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream  
423 computation. *Journal of the ACM (JACM)*, 53(3):307–323, 2006.
- 424 [15] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the  
425 curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of*  
426 *computing*, pages 604–613, 1998.
- 427 [16] Piotr Indyk, Rajeev Motwani, Prabhakar Raghavan, and Santosh Vempala. Locality-preserving  
428 hashing in multidimensional spaces. In *Proceedings of the twenty-ninth annual ACM symposium*  
429 *on Theory of computing*, pages 618–625, 1997.
- 430 [17] Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. On using very large  
431 target vocabulary for neural machine translation. *arXiv preprint arXiv:1412.2007*, 2014.
- 432 [18] Junjie Jia and Zhipeng Lei. Personalized recommendation algorithm for mobile based on  
433 federated matrix factorization. In *Journal of Physics: Conference Series*, volume 1802, page  
434 032021. IOP Publishing, 2021.

- 435 [19] William B Johnson. Extensions of lipschitz mappings into a hilbert space. *Contemp. Math.*,  
436 26:189–206, 1984.
- 437 [20] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Ar-  
438 jun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings,  
439 et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine*  
440 *Learning*, 14(1–2):1–210, 2021.
- 441 [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*  
442 *arXiv:1412.6980*, 2014.
- 443 [22] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer.  
444 *arXiv preprint arXiv:2001.04451*, 2020.
- 445 [23] Noam Nisan. Pseudorandom generators for space-bounded computations. In *Proceedings of*  
446 *the twenty-second annual ACM symposium on Theory of computing*, pages 204–212, 1990.
- 447 [24] Zaifeng Pan, Feng Zhang, Hourun Li, Chenyang Zhang, Xiaoyong Du, and Dong Deng. G-slide:  
448 A gpu-based sub-linear deep learning engine via lsh sparsification. *IEEE Transactions on*  
449 *Parallel and Distributed Systems*, 33(11):3015–3027, 2021.
- 450 [25] Devi Parikh and Kristen Grauman. Relative attributes. In *2011 International Conference on*  
451 *Computer Vision*, pages 503–510. IEEE, 2011.
- 452 [26] Chengshuai Shi, Cong Shen, and Jing Yang. Federated multi-armed bandits with personalization.  
453 In *International Conference on Artificial Intelligence and Statistics*, pages 2917–2925. PMLR,  
454 2021.
- 455 [27] Anshumali Shrivastava and Ping Li. Asymmetric lsh (alsh) for sublinear time maximum inner  
456 product search (mips). *Advances in neural information processing systems*, 27, 2014.
- 457 [28] Anshumali Shrivastava and Ping Li. In defense of minhash over simhash. In *Artificial Intelli-*  
458 *gence and Statistics*, pages 886–894. PMLR, 2014.
- 459 [29] Ryan Spring and Anshumali Shrivastava. Scalable and sustainable deep learning via randomized  
460 hashing. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge*  
461 *Discovery and Data Mining*, pages 445–454, 2017.
- 462 [30] Chuhan Wu, Fangzhao Wu, Tao Qi, Yongfeng Huang, and Xing Xie. Fedattack: Effective and  
463 covert poisoning attack on federated recommendation via hard sampling. In *Proceedings of the*  
464 *28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4164–4172,  
465 2022.
- 466 [31] Zhaozhuo Xu, Luyang Liu, Zheng Xu, and Anshumali Shrivastava. Adaptive sparse federated  
467 learning in large output spaces via hashing. In *Workshop on Federated Learning: Recent*  
468 *Advances and New Challenges (in Conjunction with NeurIPS 2022)*.
- 469 [32] Jay Yagnik, Dennis Strelow, David A Ross, and Rwei-sung Lin. The power of comparative  
470 reasoning. In *2011 International Conference on Computer Vision*, pages 2431–2438. IEEE,  
471 2011.
- 472 [33] Minghao Yan, Nicholas Meisburger, Tharun Medini, and Anshumali Shrivastava. Distributed  
473 slide: Enabling training large neural networks on low bandwidth and simple cpu-clusters via  
474 model parallelism and sparsity. *arXiv preprint arXiv:2201.12667*, 2022.
- 475 [34] Ian En-Hsu Yen, Satyen Kale, Felix Yu, Daniel Holtmann-Rice, Sanjiv Kumar, and Pradeep  
476 Ravikumar. Loss decomposition for fast learning in large output spaces. In *International*  
477 *Conference on Machine Learning*, pages 5640–5649. PMLR, 2018.
- 478 [35] Felix Yu, Sanjiv Kumar, Yunchao Gong, and Shih-Fu Chang. Circulant binary embedding. In  
479 *International conference on machine learning*, pages 946–954. PMLR, 2014.
- 480 [36] Shijie Zhang, Hongzhi Yin, Tong Chen, Zi Huang, Quoc Viet Hung Nguyen, and Lizhen Cui.  
481 Pipattack: Poisoning federated recommender systems for manipulating item promotion. *arXiv*  
482 *preprint arXiv:2110.10926*, 2021.

483 [37] Andrew Ziegler, Eric Christiansen, David Kriegman, and Serge Belongie. Locally uniform  
484 comparison image descriptor. *Advances in Neural Information Processing Systems*, 25, 2012.

485 **Paper 1682: Supplementary Material**  
 486 **Large-Scale Distributed Learning via Private**  
 487 **On-Device LSH**

488 **A Experiment details**

489 In this section, we provide deeper background into how our experiments were run as well as some  
 490 additional results and observations. We first detail the hyper-parameters we used in order to reproduce  
 491 our results. Then, we provide additional comments and details into our sampling approach. Finally,  
 492 we describe some of the interesting observations we encountered while solving the Amazon-670K  
 493 and Wiki-325K recommender system problems.

494 **A.1 Experiment hyper-parameters**

Below, we detail the hyper-parameters we used when running our federated experiments.

Table 1: **Hyper-parameters for Federated Experiments (PGHash and Federated SLIDE).**

Dataset	Algorithm	Hash Type	LR	Batch Size	Steps per LSH	$k$	$c$	Tables	$CR$
Delicious-200K	PGHash	PGHash	1e-4	128	1	8	8	50	1
Delicious-200K	SLIDE	SimHash	1e-4	128	1	8	N/A	50	1
Amazon-670K	PGHash	PGHash-D	1e-4	256	50	8	8	50	1
Amazon-670K	SLIDE	DWTA	1e-4	256	50	8	N/A	50	1
Wiki-325K	PGHash	PGHash-D	1e-4	256	50	5	16	50	1
Wiki-325K	SLIDE	DWTA	1e-4	256	50	5	N/A	50	1

495  
 496 What one can immediately see from Table 1, is that we use a Densified Winner Take All (DWTA)  
 497 variant of PGHash for the larger output datasets Amazon-670K and Wiki-325K. As experienced in  
 498 [8, 7, 24], SimHash fails to perform well on these larger datasets. We surmise that SimHash fails  
 499 due in part to its inability to select a large enough number of neurons per sample (we observed this  
 500 dearth of activated neurons empirically). Reducing the hash length  $k$  does increase the number of  
 501 neurons selected, however this decreases the accuracy. Therefore, DWTA is used because it utilizes  
 502 more neurons per sample on these larger problems and also still achieves good accuracy.

Table 2: **Hyper-parameters for Compression Experiments (PGHash).**

Dataset	Algorithm	Hash Type	LR	Batch Size	Steps per LSH	$k$	$c$	Tables	$CR$
Delicious-200K	PGHash	PGHash	1e-4	128	1	8	8	50	0.1/0.25/1

503 As a quick note, we record test accuracy every so often (around 100 iterations for Delicious-200K  
 504 and Amazon-670K). Similar to [8], to reduce the test accuracy computations (as the test sets are very  
 505 large) we compute the test accuracy of 30 randomly sampled large batches of test data.

506 **A.2 Neuron sampling**

507 **Speed of Neuron Sampling.** In Table 3 we display the time it takes to perform LSH for PGHash  
 508 given a set number of tables. These times were collected locally during training. The entries in Table  
 509 3 denote the time it takes to compute hashing of the final layer weights  $w_i$  and each sample  $x$  in batch  
 510  $M$  as well as vanilla-style matching (neuron selection) for each sample.

Table 3: **Average LSH time for PGHash over a range of tables.** We compute the average  $\mu$  time  
 (and standard deviation  $\sigma$ ) it takes for PGHash to perform *vanilla sampling* (exact matches) between  
 the hash codes of sample  $x$  and each weight  $w_i$  in the final dense layer. Times are sampled for  
 PGHash on Delicious-200k for batch size  $M = 128$ ,  $k = 9$ , and  $c = 8$  for one device.

Method	1 table (seconds)	50 tables (seconds)	100 tables (seconds)
PGHash	$\mu = 0.0807, \sigma = 0.0076$	$\mu = 3.1113, \sigma = 0.0555$	$\mu = 6.2091, \sigma = 0.1642$
SLIDE	$\mu = 0.0825, \sigma = 0.0099$	$\mu = 3.2443, \sigma = 0.1671$	$\mu = 6.2944, \sigma = 0.0689$

511 We find in Table 3 that PGHash achieves near sub-linear speed with respect to the number of tables  $\tau$   
 512 and slightly outperforms SLIDE. PGHash edges out SLIDE due to the smaller matrix multiplication

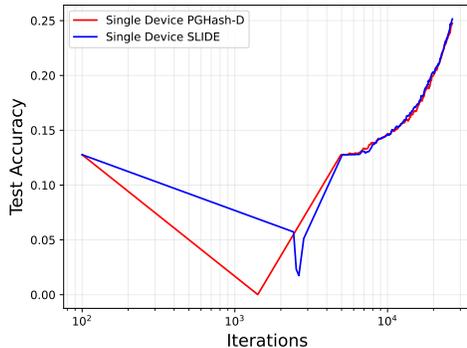


Figure 7: **Wiki-325K PGHash-D.** We record model accuracy of a large recommendation system on Wiki-325K. PGHash-D matches the convergence of SLIDE without requiring LSH to be performed by the central server. We note that test accuracy is determined by testing 30 randomly sampled large batches of test data (and not the full test data). We saw that the true full test accuracy (which we compute after each epoch) ran about 5% greater than the sampled batches.

513 cost, as PGHash utilizes a smaller random Gaussian matrix (size  $c \times c$ ). The speed-up over SLIDE will  
 514 become more significant when the input layer is larger (as  $d = 128$  in our experiments). Therefore,  
 515 PGHash obtains superior sampling performance to SLIDE.

516 **Hamming Distance Sampling.** An alternative method to vanilla sampling is to instead select final  
 517 layer weights (neurons)  $w_i$  which have a small Hamming distance relative to a given sample  $x$ . As a  
 518 refresher, the Hamming distance simply computes the number of non-matching entries between two  
 519 binary codes (strings). If two binary codes match exactly, then the Hamming distance is zero. In this  
 520 sampling routine, either (i) the top- $k$  weights  $w_i$  with the smallest Hamming distance to sample  $x$   
 521 are selected to be activated or (ii) all weights  $w_i$  with a Hamming distance of  $\beta$  or smaller to sample  
 522  $x$  are selected to be activated. Interestingly, the vanilla-sampling approach we use in our work is  
 523 equivalent to using  $\beta = 0$  in (ii).

524 In either of the scenarios listed above, hash codes for  $w_i$  and  $x$  are computed as done in PGHash(-D).  
 525 From there, however, the hash code for  $x$  is compared to the hash codes for all final layer weights  
 526 in order to compute the Hamming distance for each  $w_i$ . The process of computing  $n$  Hamming  
 527 distances for each sample  $x$  is very expensive (much harder than just finding exact matches). That is  
 528 why our work, as well as [8, 7], use vanilla sampling instead of other methods.

### 529 A.3 Amazon-670K and Wiki-325K experiment analysis

530 **Sub-par SimHash Performance.** SimHash is known to perform worse than DWTA on Amazon-  
 531 670K and Wiki-325K. Utilizing SimHash for these experiments is unfair as it is shown by [8, 7],  
 532 for example, that DWTA achieves much higher performance on Amazon-670K. For this reason,  
 533 DWTA is the chosen hash function in [8] for Amazon-670K experiments. To verify this observation,  
 534 we performed experiments on Amazon-670K with PGHash (not PGHash-D) and SLIDE (with a  
 535 SimHash hash function). Table 4 displays the SimHash approach for Amazon-670K.

Table 4: **PGHash and SLIDE performance on Amazon-670K using SimHash.** Accuracy across the first 5,000 iterations for a single device. Batch size  $M = 1024$ ,  $k = 8$ , and  $c = 8$ .

Iteration	SLIDE	PGHash
1,000	10.82%	10.04%
2,000	18.27%	15.99%
3,000	21.83%	19.51%
4,000	23.72%	21.65%
5,000	25.08%	23.38%

536 As shown in Table 4, even with a much larger batch size, SLIDE and PGHash are unable to crack  
 537 30% on Amazon-670K. We would like to note that using a smaller batch size (like the  $M = 256$

538 value we use in our Amazon-670K experiments) resulted in an even further drop in accuracy. These  
 539 empirical results back-up the notion that SimHash is ill-fit for Amazon-670K.

540 **Wiki-325K Performance.** In Figure 7, we showcase how PGHash-D performs on Wiki-325K.  
 541 Quite similar to the Amazon-670K results (shown in Figure 5), PGHash-D almost exactly matches up  
 542 with SLIDE. In order to map how well our training progresses, we periodically check test accuracies.  
 543 However, since the test set is very large, determining test accuracies over the entire test set is infeasible  
 544 due to time constraints on the cluster. Therefore, we determine test accuracies over 30 batches of  
 545 test data as a substitute as is done in [8, 7]. For Delicious-200K and Amazon-670K the entire  
 546 test set accuracies matched the randomly sampled batches, however the randomly sampled batches  
 547 underestimate the true test accuracies for Wiki-325K. For Wiki-325K, the true test accuracy ran about  
 548 5% greater than the sampled test accuracy values.

549 **Matching Full-Training Performance.** Along with the failure for SimHash to perform well on  
 550 Amazon-670K and Wiki-325K, SLIDE and PGHash(-D) are unable to match the performance of  
 551 full-training on these data-sets. This is observed empirically for Amazon-670K by GResearch in  
 552 the following article <https://www.gresearch.co.uk/blog/article/implementing-slide/>.  
 553 We surmise that the failure of SLIDE and PGHash(-D) to match full-training performance on Amazon-  
 554 670K and Wiki-325K arises due to the small average labels per point in these two data-sets (5.45  
 555 and 3.19 respectively). Early on in training, SLIDE and PGHash(-D) do not utilize enough activated  
 556 neurons. This is detrimental to performance when there are only a few labels per sample, as the  
 557 neurons corresponding to the true label are rarely selected at the beginning of training (and these  
 558 final layer weights are tuned much slower). In full-training, the true neurons are always selected and  
 559 therefore the final layer weights are better adjusted from the beginning. We also note that [33] requires  
 560 a hidden layer size of 1024 for a distributed version of SLIDE to achieve improved test accuracies for  
 561 Amazon-670K. Thus, increasing the hidden layer size may have improved our performance (we kept  
 562 it as 128 to match the original SLIDE paper [8]).

## 563 B PGHash: angle versus Hamming distance

564 In this section, we visually explore the degree to which PGHash is a consistent estimator of angular  
 565 similarity. Specifically, let  $x, y \in \mathbb{R}^d$ : then we know by Theorem 1 that  $\mathcal{H}^{PG}(c, d)$  is an LSH for  
 566  $\cos(x_c, y_c)$ . We demonstrate that in the unit vector regime,  $\theta_c = \arccos(\cos(x_c, Y_c))$  is an acceptable  
 567 surrogate for  $\theta = \arccos(x, Y)$ , where  $Y = \{y^i\}_{i=1}^N$  and  $Y_c = \{y_c^i\}_{i=1}^N$ .

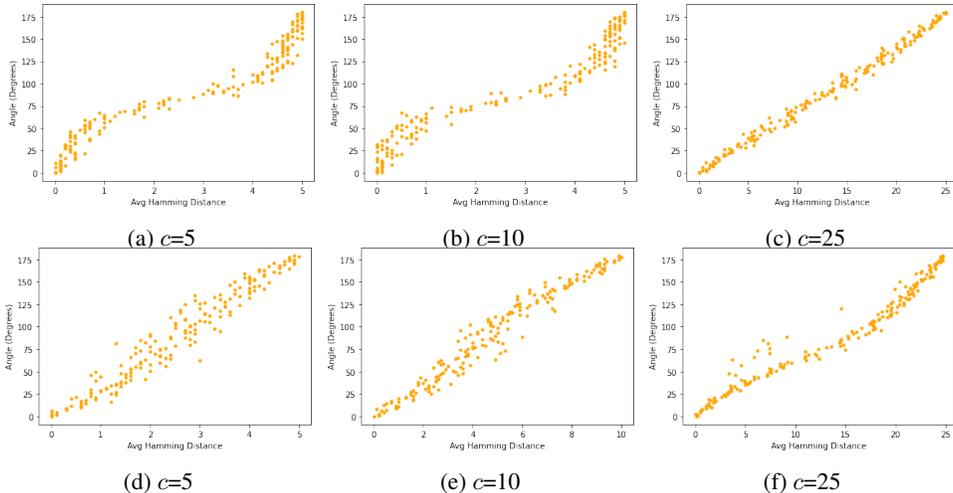


Figure 8: **Angle/Hamming Distance as a function of sketch dimension.** The average Hamming distance between a PGHashed fixed unit vector  $x \in \mathbb{R}^{100}$  and a collection of vectors  $y_i \in \mathbb{R}^{100}$  which form different angles with  $x$ . Increasing sketch dimension  $c$  smooths and reduces the variance of the scatter towards linear correlation. Furthermore, the Hamming scales linearly with  $c$ , improving discernibility. (a)-(c) & (d)-(f) are independent series.

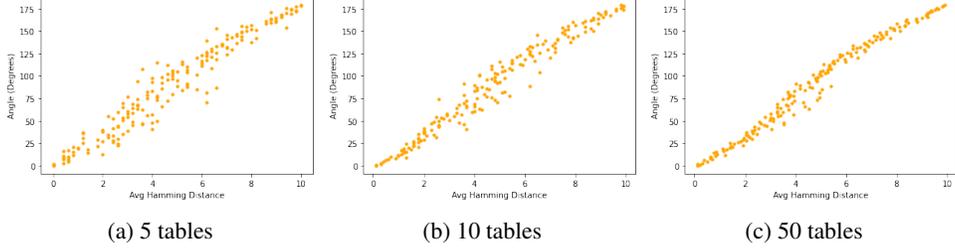


Figure 9: **Angle/Hamming Distance as a function of tables** The average Hamming distance between a PGHashed fixed unit vector  $x \in \mathbb{R}^{100}$  and a collection of vectors  $y_i \in \mathbb{R}^{100}$  which form different angles with  $x$  and fixed sketch dimension  $c = 10$ . Increasing the number of tables reduces variance.

## 568 C Additional proofs

569 **Fact 3.** Let  $x, y, e_1, e_2$  be  $d$ -dimensional unit vectors such that that the  $e_i$  lie on the unit circle  
570 contained with the plane spanned by  $x$  and  $y$  (denoted as  $S_{x,y}$ ) and  $e_1 \perp e_2$ . Consider the point  
571  $v$  on  $S_{x,y}$  such that the line through it bisects the angle of the lines passing through  $x$  and  $y$ . Let  
572  $\eta = \arccos(\cos(v, e_1))$ . Denote  $\theta = \frac{1}{2} \arccos(\cos(x, y))$ . Then we may write  $x = \cos(\eta + \theta)e_1 +$   
573  $\sin(\eta + \theta)e_2$  and  $y = \cos(\eta - \theta)e_1 + \sin(\eta - \theta)e_2$ .

### 574 C.1 Proof of Theorem 2

575 *Proof.* Let  $\theta = \frac{1}{2} \arccos(\cos(x, y))$  where  $x, y \in \mathbb{S}^{d-1}$  and  $A = B^\top B$ . The cosine similarity  
576 between  $x_c = Bx$  and  $y_c = By$  (for  $B$  correspondent to a  $(d, c)$ -folding), is expressible as

$$\cos(x_c, y_c) = \frac{x^\top B^\top B y}{(x^\top B^\top B x)(y^\top B^\top B y)} = \frac{x^\top A y}{\sqrt{(x^\top A x)(y^\top A y)}}. \quad (1)$$

577 Consider the SVD  $B = UDV^\top$  where  $U$  and  $V$  are orthogonal and  $D$  is  $c \times d$  rectangular diagonal  
578 matrix. We have then that  $A = B^\top B = V\hat{D}^2V^\top$ . (Here  $\hat{D}$  is now a square diagonal matrix  
579 containing squared  $D_{ii}$  along the diagonal and 0 everywhere else.) Notice that choice of  $U$  nor the  
580 ordering of columns  $v_i$  of  $V$  affects the angle calculation in Equation 1. First, we re-order the columns  
581 of  $V$  so as to order the diagonal entries  $d_i$  of  $D$  (i.e., the squared singular values) in decreasing order,  
582 and as an abuse of notation set  $B = \frac{1}{d_1}DV^\top$ . Denoting  $\hat{\lambda}_i = d_i/d_1$  for  $1 \leq i \leq n$ , we have that  
583  $Bv_i = \hat{\lambda}_i e_i$ . (By construction of  $B$  we have that  $d_i \in \{\frac{d}{c}, 0\}$ , therefore,  $\hat{\lambda}_i \in \{1, 0\}$ )

584 Consider  $B$  acting on  $S^{d-1}$ : it scales each dimension by  $\hat{\lambda}_i$ , thus (as with any linear transformation  
585 of a sphere), transforms it into an ellipsoid, with  $c$  principal axes determined by the  $v_i$ . The  
586 greatest possible distance from the origin to the ellipsoid  $BS^{d-1}$  is 1 while the shortest possible  
587 distance is 0. Now consider the unit circle  $S_{x,y} = \{v \in \text{span}(x, y) : \|v\| = 1\}$ . We have that  
588  $BS_{x,y} \subset BS^{d-1} \cap BU$  is an ellipse (since the intersection of an ellipsoid and plane is always an  
589 ellipse).

590 Choose unit  $w_1$  and  $w_2$  belonging to  $S_{x,y}$  such that  $w_1 \perp w_2$ . by Fact 3, we may parameterize  
591 our vectors as  $x = \cos(\eta - \theta)w_1 + \sin(\eta - \theta)w_2$  and  $y = \sin(\eta + \theta)w_1 + \sin(\eta + \theta)w_2$ , where  $\eta$   
592 is the angle made with  $w_1$  with the bisector of  $x$  and  $y$ . By assumption,  $\|Bw\| \geq \alpha$  (the minimal  
593 shrinking factor of  $B$  on  $S_{x,y}$ ), so denoting  $\lambda = \frac{d}{c}$  (the maximal stretching factor of  $B$  on  $S_{x,y}$ ), we  
594 have that the angle between  $Bx$  and  $By$  is upper-bounded by

$$f(\eta) = \arctan\left(\frac{\alpha}{\lambda} \tan(\eta + \theta)\right) - \arctan\left(\frac{\alpha}{\lambda} \tan(\eta - \theta)\right) \quad (2)$$

595 .

596 The numerator of  $\frac{df}{d\eta}$  is  $\beta(1 - \beta)(1 + \beta) \sin(2\theta) \sin(2\eta)$  where  $\beta = \alpha/\lambda$ . The derivative is trivially  
597 0 if (1)  $\beta = 0$ , (2)  $\beta = 1$ , or (3)  $\theta = 0$ . (1) will not occur as we assume that  $S_{x,y}$  does not contain a  
598 0-eigenvector of  $A = B^\top B$ . (2) can only occur if  $A$  is a multiple of the identity matrix (which it is  
599 not by construction), and (3) implies that  $x$  and  $y$  are parallel, in which case their angle will not be

600 distorted. Aside from these pathological cases, the critical points occur at  $\eta = 0, \pi/2$ . We have then  
 601 that  $\cos(Bx, By)$  lives between  $\cos(f(0)) = \frac{1-\beta^2 \tan^2 \theta}{1+\beta^2 \tan^2 \theta}$  and  $\cos(f(\pi/2)) = -\frac{\tan^2 \theta^2 - \beta^2}{\tan^2 \theta^2 + \beta^2}$ .

602 □

603 **Remark.** The constant  $\beta$  has an enormous influence on the bounds in Theorem 2. The smaller the  $\alpha$   
 604 (i.e., shrinking of  $\|w\|$ ), the greater the bounds on distortion. Although we have imposed constraints  
 605 on  $x, y$ , if we treat them as any possible pair of random unit vectors, then the  $w$  in  $S_{x,y}$  effectively  
 606 becomes a random unit vector as well. We can exactly characterize the distribution of  $\|BX\|$  where  
 607  $X$  denotes a random variable which selects a  $d$ -dimensional unit vector uniformly at random.

## 608 C.2 Proof of Proposition 1

609 *Proof.* We can sample a  $d$ -dimensional vector uniformly at random from the unit sphere  $S^{d-1}$  by  
 610 drawing a  $d$ -dimensional Gaussian vector with iid entries and normalizing. Let us represent this  
 611 as the random variable  $X = Z'/\|Z'\|$  where  $Z' \sim \mathcal{N}(0, I_d)$ . Consider a  $(c, d)$ -folding matrix  
 612  $B$ , i.e., a  $d/c$  horizontal stack of  $c \times c$  identity matrices (let us assume  $c|d$ ). We are interested in  
 613 determining the distribution of  $\|BX\|^2$ . For ease of notation, consider the permutation  $Z$  of  $Z'$   
 614 where  $Z_i = Z'_{(\lfloor \frac{d}{c} \rfloor - 1) * (d/c) + i \pmod{d/c}}$ . Since this permutation is representable as an orthogonal  
 615 matrix  $P$  (and multi-variate Gaussians are invariant in distribution under orthogonal transformations),  
 616 we may instead consider  $X := P(Z'/\|Z'\|) = Z/\|Z\|$ . We may write the norm-squared as

$$\|BX\|^2 = \frac{(Z_1 + \dots + Z_{d/c})^2}{\|Z\|^2} + \frac{(Z_{d/c+1} + \dots + Z_{2d/c})^2}{\|Z\|^2} + \dots + \frac{(Z_{(c-1)(d/c)+1} + \dots + Z_d)^2}{\|Z\|^2}. \quad (3)$$

617 Consider the first term  $\frac{(Z_1 + \dots + Z_{d/c})^2}{\|Z\|^2}$ . First note that for any unit vector  $u$ , the distribution of  $\frac{(u^\top Z)^2}{\|Z\|^2}$   
 618 does not depend on choice of  $u$ . Consider the unit vector  $u'$  then which contains  $\sqrt{d/c}$  in the first  $d/c$   
 619 entries and 0 otherwise. Then  $\frac{(u'^\top Z)^2}{\|Z\|^2}$  is equivalent to  $d/c$  times our first term. Of course, since  $\frac{(e_1^\top Z)^2}{\|Z\|^2}$   
 620 has the same distribution as  $\frac{(u'^\top Z)^2}{\|Z\|^2}$ , we have by transitivity that  $\frac{Z_1^2}{\|Z\|^2} \stackrel{d}{=} (n/q) \frac{(Z_1 + \dots + Z_{d/c})^2}{\|Z\|^2}$ .

621 By extending the discussion above to the other terms, and by their independence with respect to  
 622 rotation of  $Z$  (since their numerators contain squared sums of mutually disjoint  $Z$  coordinates), we  
 623 have that

$$\|BX\|^2 \stackrel{d}{=} \frac{d}{c} \cdot \frac{Z_1^2 + Z_{d/c}^2 + Z_{2d/c}^2 + \dots + Z_d^2}{\|Z\|^2}. \quad (4)$$

624 The distribution of  $\frac{Z_1^2 + Z_{d/c}^2 + Z_{2d/c}^2 + \dots + Z_d^2}{\|Z\|^2}$  is well-known to follow a  $\text{Beta}(\frac{c}{2}, \frac{d-c}{2})$  distribution [13].

625 In totality,  $\|BX\|^2 \stackrel{d}{=} \frac{d}{c} \text{Beta}(\frac{c}{2}, \frac{d-c}{2})$ . However, we will move to the four parameter description of  
 626 this scaled Beta distribution which is  $\text{Beta}(\frac{c}{2}, \frac{d-c}{2}, 0, \frac{d}{c})$ . The pdf and expected value follows by the  
 627 usual statistical descriptions of this distribution, which can also be found in [13]. □

628 Figure 10 depicts how  $(d, c)$ -foldings affect the norms of unit vectors.

## 629 D Additional theory

630 In this section, we provide additional theory relevant to SimHash.

631 We present several well-known results regarding SimHash.

632 **Proposition 2** (SimHash estimation). *Let  $x, y \in \mathbb{S}$ , i.e., unit  $d$ -dimensional vectors. Denote  $\theta =$   
 633  $\arccos(|\cos(x, y)|)$ . Let  $v \in S^d$  be a unit vector drawn uniformly at random (according to the Haar  
 634 measure, for example). Then,*

$$\Pr[\text{sgn}(v^\top x) \neq \text{sgn}(v^\top y)] = \frac{\theta}{\pi}. \quad (5)$$

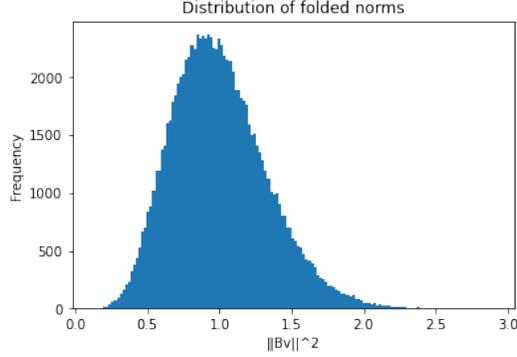


Figure 10: **Distribution of folded norms.** 100k randomly drawn unit vectors ( $d = 128$ ) are folded down to length 16 by are usual  $(d, c)$ -folding procedure. Depicted is a binned histogram of the norms. As predicted by the statistical description of  $\|BX\|^2$ , where  $X$  is a randomly drawn unit vector, the mass is centered at 1, i.e., most norms are preserved. Empirically we observe that folded rarely exceed  $\sqrt{12816}$ , although the theoretical support is  $[0, 8]$ : this concurs with the pdf.

635 *Proof.* We reproduce the argument of [12]. We have by symmetry that  $\Pr[\text{sgn}(v^\top x) \neq \text{sgn}(v^\top y)] =$   
 636  $2\Pr[v^\top x > 0, v^\top y < 0]$ . The set  $\mathcal{U} = \{v \in S^d : v^\top x > 0, v^\top y \leq 0\}$  corresponds to the  
 637 intersection of two half-spaces whose dihedral angle (i.e., angle between the normals of both spaces)  
 638 is exactly  $\theta$ . Intersecting with the  $d$ -dimensional unit sphere produces gives a subspace of measure  
 639  $\frac{\theta}{2\pi}$ , therefore,  $2\Pr[v^\top x > 0, v^\top y < 0] = \frac{\theta}{\pi}$ , completing the argument.  $\square$

640 **Corollary 2.** Let  $v$  instead be a  $d$ -dimensional random Gaussian vector with iid entries  $\sim \mathcal{N}(0, 1)$ .  
 641 Then for  $x, y \in \mathbb{R}^d$ ,

$$\Pr[\text{sgn}(v^\top x) \neq \text{sgn}(v^\top y)] = \frac{\theta}{\pi} \quad (6)$$

642 *Proof.* Randomly drawn, normalized Gaussian vectors are well-known to be uniformly distributed  
 643 on the unit sphere.  $\square$

644 In the setup as above, let the  $X$  be a random variable which returns 1 if  $x$  and  $y$  have differing signs  
 645 when taking the standard inner product with a randomly drawn Gaussian  $v$ . Let  $X_1, X_2, \dots, X_n$   
 646 represent a sequence of independent  $X$  events. Then,

647 **Proposition 3.**  $\mathbb{E}[\frac{1}{n} \sum_{i=1}^n X_i] = 1 - \frac{\theta}{\pi}$  and  $\mathbb{V}[X] = \frac{1}{N} \frac{\theta}{\pi} (1 - \frac{\theta}{\pi})$ .

648 Given that PGHash is equivalent to a SimHash over  $(d, c)$ -foldings of  $R^d$ , the variance reduction we  
 649 observe by using multiple tables (Figure 9 is explainable by Proposition 3.