



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

GPH: Similarity Search in Hamming Space

Citation for published version:

Qin, J, Wang, Y, Xiao, C, Wang, W, Lin, X & Ishikawa, Y 2018, GPH: Similarity Search in Hamming Space. in *34th IEEE International Conference on Data Engineering*. Institute of Electrical and Electronics Engineers (IEEE), Paris, France, pp. 1-14, 34th IEEE International Conference on Data Engineering, Paris, France, 16/04/18. <https://doi.org/10.1109/ICDE.2018.00013>

Digital Object Identifier (DOI):

[10.1109/ICDE.2018.00013](https://doi.org/10.1109/ICDE.2018.00013)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

34th IEEE International Conference on Data Engineering

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



GPH: Similarity Search in Hamming Space

Jianbin Qin[†] Yaoshu Wang[†] Chuan Xiao[‡] Wei Wang[†] Xuemin Lin[†] Yoshiharu Ishikawa[‡]

[†]University of New South Wales, Australia

{jqin, yaoshuw, weiw, lxue}@cse.unsw.edu.au

[‡] Nagoya University, Japan

chuanx@nagoya-u.jp ishikawa@i.nagoya-u.ac.jp

Abstract—A similarity search in Hamming space finds binary vectors whose Hamming distances are no more than a threshold from a query vector. It is a fundamental problem in many applications, including image retrieval, near-duplicate Web page detection, and machine learning. State-of-the-art approaches to answering such queries are mainly based on the *pigeonhole principle* to generate a set of candidates and then verify them.

We observe that the constraint based on the pigeonhole principle is *not always tight* and hence may bring about unnecessary candidates. We also observe that the distribution in real data is often *skew*, but most existing solutions adopt a simple equi-width partitioning and allocate the same threshold to all the partitions, and hence fail to exploit the data skewness to optimize the query processing. In this paper, we propose a new form of the pigeonhole principle which allows variable partition size and threshold. Based on the new principle, we first develop a tight constraint of candidates, and then devise cost-aware methods for dimension partitioning and threshold allocation to optimize query processing. Our evaluation on datasets with various data distributions shows the robustness of our solution and its superior query processing performance to the state-of-the-art methods.

I. INTRODUCTION

Finding similar objects is a fundamental problem in database research and has been studied for several decades [37]. Among many types of queries to find similar objects, Hamming distance search on binary vectors is an important one. Given a query q , a Hamming distance search finds all vectors in a database whose Hamming distances to q are no greater than a threshold τ . Answering such queries efficiently plays an important role in many applications. In addition, it will receive increasingly attention due to the continuous development and adoption of deep learning techniques.

- Various forms of Hamming similarity search are fundamental issues in computational geometry and theoretical computer science. For example, nearest neighbor [14] or range searches [9] can be reduced and solved in a Hamming cube using the Hamming distance.
- Hamming similarity searches have been used in various similarity search or retrieval applications. This is because one can use various hashing techniques (either locality sensitive hashing [39] or metric-learning or deep-learning based learned hash functions [38]) or manual feature exactions (e.g., in Cheminformatics [40]) to obtain a compact binary code for objects, and similarity search or retrieval can be performed efficiently in this Hamming cube.

We give several concrete and diverse application examples in the following three areas:

- For image retrieval, images are converted to compact binary vectors in traditional methods [34], [28]. Zhang *et al.* proposed to identify the vectors within a Hamming distance threshold of 16 as candidates for further image-level verification [42]. Recently, deep learning has become remarkably successful in image recognition. Learning to hash algorithms that utilize neural networks have been actively explored [17], [19], [7]. In these studies, images are represented by binary vectors and Hamming distance is utilized to capture the dissimilarity between images.
- To process text documents, state-of-the-art information retrieval methods [29], [8] represent documents by binary vectors through hashing. For Web page deduplication, Google uses SimHash to convert a Web page to a 64-bit vector, and then two pages are considered as near-duplicate if the Hamming distance of the vectors are within 3 [22].
- For scientific databases, a fundamental task in cheminformatics is to find similar molecules [12], [24]. In this task, molecules are converted into binary vectors, and the Tanimoto similarity is used to measure the similarity between molecules. This similarity constraint can be converted to an equivalent Hamming distance constraint [43].

The naïve algorithm to answer a Hamming distance search query requires access and computation of every vector in the database; hence it is expensive and does not scale well to large datasets. Therefore, there has been much interest in devising efficient indexes and algorithms. Many existing methods [1], [18], [43], [25] adopt the *filter-and-refine* framework to quickly find a set of candidates and then verify them. They are based on the naïve application of the *pigeonhole principle* to this problem: If the n dimensions are partitioned into m *equi-width* parts¹, then a necessary condition for the Hamming distance of two vectors to be within τ is that they must share a part in which the Hamming distance is within $\lfloor \frac{\tau}{m} \rfloor$. This leads to a *filtering condition*, and produces a set of candidate vectors, which are then verified by calculating the Hamming distances and comparing with the threshold. As a result, the efficiencies of these methods critically depend on the candidate size.

However, despite the success and prevalence of this framework, we identify that the filtering condition has two inherent major weaknesses: (1) **The threshold on each partition is not always tight**. Hence, many unnecessary candidates are included. For example, when $m = 3$, the filtering conditions for τ in

¹In this paper, we assume $n \bmod m = 0$.




Fig. 1. Skewness ($\frac{|\#1s - \#0s|}{\#data}$) by dimension of datasets in [16].

[9, 11] are the same (Hamming distance $\leq \lfloor \frac{\tau}{m} \rfloor = 3$), and hence will produce the same set of candidates.

(2) **The thresholds on the partitions are evenly distributed.** It assumes a uniform distribution and does not work well when the dataset is skewed. We found that many real datasets are skewed to varying degrees and complex correlations exist among dimensions. Fig. 1 shows that 8 out of 11 real datasets have dimensions with skewness greater than 0.3², and 5 out of the 8 datasets contain a vector whose frequency ≥ 0.1 on a partition, meaning that at least 1/10 data vectors become candidates if the query matches the data vector on this partition.

In this paper, we propose a novel method to answer the Hamming distance search problem and address the above-mentioned weaknesses. We propose a tight form of the pigeonhole principle named *general pigeonhole principle*. Based on the new principle, the thresholds of the m partitions sum up to $\tau - m + 1$, less than τ , thus yielding a stricter filtering condition than the existing methods. In addition, the threshold on each partition is a *variable* in the range of $[-1, \tau]$, where -1 indicates that this partition is ignored when generating candidates. This enables us to choose proper thresholds for different partitions in order to improve query processing performance. We prove that the candidate condition based on the general pigeonhole principle is *tight*; i.e., the threshold allocated to each partition cannot be further reduced. To tackle data skewness and dimension correlations, we first devise an online algorithm to allocate thresholds to partitions using a query processing cost model, and then devise an offline algorithm to optimize the partitioning of vectors by taking account of the distribution of dimensions. The proposed techniques constitute the GPH algorithm. Experiments are run on several real datasets with different data distributions. The results show that the GPH algorithm performs consistently well on all these datasets and is faster than state-of-the-art methods by up to two orders of magnitude.

Our contributions can be summarized as follows. (1) We propose a new form of the pigeonhole principle to obtain a tight filtering condition and enable flexible threshold allocation. (2) We propose an efficient online query optimization method to allocate thresholds on the basis of the new pigeonhole principle. (3) We propose an offline partitioning method to address the selectivity issue caused by data skewness and dimension correlations. (4) We

²To measure the skewness of the i -th dimension, we calculate the numbers of vectors whose values on the i -th dimension are 0 and 1, respectively, and then take the ratio of their difference and the total number of vectors.

conduct extensive experimental study on several real datasets to evaluate the proposed method. The results demonstrate the superiority of the proposed method over state-of-the-art methods.

II. PRELIMINARIES

A. Problem Definition

In this paper, we focus on the similarity search on binary vectors. We can view an object as an n -dimensional binary vector x . $x[i]$ denotes the value of the i -th dimension of x . Let $\Delta(x[i], y[i]) = 0$, if $x[i] = y[i]$; or 1, otherwise. The Hamming distance between two vectors x and y , denoted $H(x, y)$, is the number of dimensions on which x and y differ:

$$H(x, y) = \sum_{i=1}^n \Delta(x[i], y[i]).$$

Hamming distance is a symmetric measure. If we regard x (respectively, y) as a yardstick, we can also say that y (respectively, x) has $H(x, y)$ errors with respect to x (respectively, y).

Given a collection of data objects \mathcal{D} , a query object q , a Hamming distance search is to find all data objects whose Hamming distance to q is no greater than a threshold τ , i.e., $\{x \mid x \in \mathcal{D}, H(x, q) \leq \tau\}$.

B. Basic Pigeonhole Principle

Most exact solutions to Hamming distance search are based on the filter-and-refine framework to generate a set of candidates that satisfy a necessary condition of the Hamming distance constraint. The majority of these methods [1], [18], [43], [25] are based on the intuition that if two vectors are similar, there will be a pair of similar partitions from the two vectors. Hence the (basic) pigeonhole principle is utilized by these methods.

Lemma 1 (Basic Pigeonhole Principle): x and y are divided into m partitions. Each partition consists of $\frac{n}{m}$ dimensions. Let x_i and y_i ($1 \leq i \leq m$) denote each partition in x and y , respectively. If $H(x, y) \leq \tau$, there exists at least one partition i such that $H(x_i, y_i) \leq \lfloor \frac{\tau}{m} \rfloor$.

Any data object x satisfying the condition that $\exists i, H(x_i, q_i) \leq \lfloor \frac{\tau}{m} \rfloor$ is called a *candidate*. Since these candidates will be verified by computing the exact Hamming distance to the query, the query processing performance depends heavily on the number of candidates.

C. Overview of Existing Approaches

We briefly introduce a state-of-the-art method, Multi-index Hamming (MIH) [25]; other methods based on the basic pigeonhole principle work in a similar way. MIH partitions the n dimensions into m equi-width partitions. In each partition, based on basic pigeonhole principle, it performs Hamming distance search on $n' = \lfloor \frac{n}{m} \rfloor$ dimensions with a threshold $\tau' = \lfloor \frac{\tau}{m} \rfloor$. MIH builds an inverted index offline, mapping each partition of a data object to the object ID. For each partition of the query, it enumerates n' -dimensional vectors whose Hamming distances to the partition are within τ' . These vectors are called *signatures*. Then it looks up signatures in the index to find candidates and verifies them.

D. Weaknesses of Basic Pigeonhole Principle

Next we analyze the major drawbacks of the filtering condition based on the basic pigeonhole principle. Note that the filtering condition is uniquely characterized by a vector of thresholds allocated to each corresponding partition; we call the vector *threshold vector*, and denote the one used by the basic pigeonhole principle as $T_{\text{basic}} = [\lfloor \frac{\tau}{m} \rfloor, \dots, \lfloor \frac{\tau}{m} \rfloor]$. We also define the *dominance* relationship between threshold vectors. Let n_i denote the number of dimensions in the i -th partition. T_1 dominates T_2 , or $T_1 \prec T_2$, iff $\forall i \in \{1, \dots, m\}$, $T_1[i] \leq T_2[i]$ and $[T_1[i], T_2[i]] \cap [-1, n_i - 1] \neq \emptyset$, and $\exists i, T_1[i] < T_2[i]$.

- T_{basic} is **not always tight**. By the tightness of a threshold vector T , we mean that (1) (correctness) every vector whose Hamming distance to the query is within the threshold will be found by the filtering condition based on T , and (2) (minimality) there does not exist another vector T' that dominates T yet still guarantees correctness. As the candidate size is monotonic with respect to the threshold, an algorithm based on a threshold vector which dominates T_{basic} will generate fewer or at most equal number of candidates compared with an algorithm based on T_{basic} .

Example 1: Consider $\tau = 9$ and $m = 3$. The threshold vector T_{basic} is $[3, 3, 3]$. We can find a dominating threshold vector $T = [2, 2, 3]$ which is tight and guarantees both correctness and minimality. Note that there may be multiple tight threshold vectors for the same τ . E.g., another tight threshold vector for the example can be $[2, 3, 2]$ or $[4, 3, 0]$ ³.

- The filtering condition does **not adapt to the data distribution** in the partitions. Skewness and correlations among dimensions often exist in real data. Equal allocation of thresholds, as done in T_{basic} , may result in poor selectivity for some partitions, hence excessive number of candidates. Several recent studies recognized this issue and proposed several methods to either obtain relatively less skew partitions by partition rearrangement [43] or allocating varying thresholds heuristically to different partitions [11]. In contrast, we propose that skewed partitions can be beneficial and we can reduce the candidate size by judiciously allocating different thresholds to different partitions for *each* query to exploit such skewness, as shown in Example 2.

Example 2: Suppose $n = 8$, $m = 2$, and $\tau = 2$. Consider the four data vectors and the query, and two different partitioning schemes in Table I. Consider the first query and existing method will use $T_{\text{basic}} = [1, 1]$. This will result in all the four data vectors recognized as candidates, but only one (x^1) is the result. If we use the first six dimensions as one partition and the rest two dimensions as the other dimension, and use $T = [2, 0]$, the candidate size will be reduced to 2 (x^1 and x^2).

III. GENERAL PIGEONHOLE PRINCIPLE

In this section, we propose a general form of the pigeonhole principle which allows variable thresholds to guarantee the tightness of threshold vectors.

³Please refer to Section III for more explanation of tightness.

TABLE I
BENEFITS OF ADAPTIVE PARTITIONING AND THRESHOLDING

	Equi-width Partitioning		Variable Partitioning	
	Partition 1	Partition 2	Partition 1	Partition 2
$x^1 = 00000000$	0000	0000	000000	00
$x^2 = 00000111$	0000	0111	000001	11
$x^3 = 00001111$	0000	1111	000011	11
$x^4 = 10011111$	1001	1111	100111	11
$q^1 = 10000000$	1000 $\tau_1 = 1$	0000 $\tau_2 = 1$	100000 $\tau_1 = 2$	00 $\tau_2 = 0$

We begin with the allocation of thresholds. Given a threshold vector, we use the notation $\|T\|_1$ to denote the sum of thresholds in all the partitions, i.e., $\|T\|_1 = \sum_{i=1}^m T[i]$. The flexible pigeonhole principle is stated below.

Lemma 2 (Flexible Pigeonhole Principle): A partitioning \mathcal{P} divides a n -dimensional vector into m disjoint partitions. x and y are partitioned by \mathcal{P} . Consider a vector $T = [\tau_1, \dots, \tau_m]$ such that τ_i are integers and $\|T\|_1 = \tau$. If $H(x, y) \leq \tau$, there exists at least one partition i such that $H(x_i, y_i) \leq \tau_i$.

Proof: Assume that $\nexists i$ such that $H(x_i, y_i) \leq \tau_i$. Since partitions are disjoint, $H(x, y) = \sum_{i=1}^m H(x_i, y_i) > \sum_{i=1}^m \tau_i$. Hence $H(x, y) > \tau$, which contradicts that $H(x, y) \leq \tau$. ■

The principle stated by Lemma 2 is more flexible than the basic pigeonhole principle in the sense that we can choose arbitrary thresholds for different partitions. Intuitively, we may tolerate more errors for selective partitions and fewer errors for unselective partitions.

To achieve tightness, we first extend the threshold allocation from integers to real numbers.

Lemma 3: x and y are partitioned by \mathcal{P} into m disjoint partitions. Consider a vector $T = [\tau_1, \dots, \tau_m]$ in which the thresholds are real numbers. $\|T\|_1 = \tau$. If $H(x, y) \leq \tau$, there exists at least one partition i such that $H(x_i, y_i) \leq \lfloor \tau_i \rfloor$.

Proof: The proof of Lemma 2 also applies to real numbers. Therefore, if $\sum_{i=1}^m \tau_i = \tau$ and $H(x, y) \leq \tau$, then $\exists i$, $H(x_i, y_i) \leq \tau_i$. Because τ_i are real numbers and $H(x_i, y_i)$ are integers, $\exists i$, $H(x_i, y_i) \leq \lfloor \tau_i \rfloor$. ■

Definition 1 (Integer Reduction): Given a threshold vector $T = [\tau_1, \tau_2, \dots, \tau_m]$, we can reduce it to $T' = [\lfloor \tau_1 \rfloor, \lfloor \tau_2 \rfloor, \dots, \lfloor \tau_m \rfloor]$. This reduction is called *integer reduction*.

It is obvious that the candidate size does not change after an integer reduction, as the Hamming distances must be integers.

When we combine Lemma 3 and the integer reduction technique, they can produce a threshold vector which dominates T_{basic} , as shown in Example 3.

Example 3: Recall in Example 1, T_{basic} is $[3, 3, 3]$ using the basic pigeonhole principle.

To obtain a dominating vector, we can start with a possible threshold vector $T = [2.9, 2.9, 3.2]$. Then by the integer reduction technique, T is reduced to $T' = [2, 2, 3]$. To see this is correct, if $\nexists i$, $H(x_i, y_i) \leq T'[i]$, there will be $3 + 3 + 4 = 10$ errors between x and y . Compared to $[3, 3, 3]$, T' is a dominating threshold vector, and the constraints on the first two partitions are stricter.

The above example also shows that the sum of thresholds of partitions can be reduced. The following lemma and theorem show how they work in the general case and the tightness guarantee of the resulting threshold vectors.

Lemma 4 (General Pigeonhole Principle): x and y are partitioned by \mathcal{P} into m disjoint partitions. Consider a threshold vector T composed of integers. $\|T\|_1 = \tau - m + 1$. If $H(x, y) \leq \tau$, there exists at least one partition i such that $H(x_i, y_i) \leq \tau_i$.

Proof: Given a vector $T = [\tau_1, \dots, \tau_m]$ such that $\|T\|_1 = \tau - m + 1$, we consider another vector $T' = [\tau'_1, \dots, \tau'_m] = [\tau_1 + 1, \dots, \tau_{m-1} + 1, \tau_m]$; i.e., it equals to T on the last partition and is greater than T by 1 in the other $m - 1$ partitions. Because $\|T'\|_1 = \|T\|_1 + (m - 1) = \tau$, by Lemma 2, if $H(x, y) \leq \tau$, then $\exists i, H(x_i, y_i) \leq \tau'_i$.

For the first $(m - 1)$ dimensions in T' , we decrease each of their thresholds by a small positive real number ϵ , and for the last dimension, we increase the threshold by $(m - 1)\epsilon$; i.e., the sum of thresholds does not change. Hence we have a vector $T'' = [\tau''_1, \dots, \tau''_m] = [\tau_1 + 1 - \epsilon, \dots, \tau_{m-1} + 1 - \epsilon, \tau_m + (m - 1)\epsilon]$. Because $\|T''\|_1 = \|T'\|_1 = \tau$, by Lemma 3, if $H(x, y) \leq \tau$, then $\exists i, H(x_i, y_i) \leq \lfloor \tau''_i \rfloor$. Because

$$\lfloor \tau''_i \rfloor = \begin{cases} \lfloor \tau_i + 1 - \epsilon \rfloor = \tau_i, & \text{if } i < m; \\ \lfloor \tau_i + (m - 1)\epsilon \rfloor = \tau_i, & \text{if } i = m, \end{cases}$$

if $H(x, y) \leq \tau$, then $\exists i, H(x_i, y_i) \leq \tau_i$. ■

One may notice that in the above proof, the partitions we choose to decrease thresholds are not limited to the first $(m - 1)$ ones. Therefore, given a threshold vector T such that $\|T\|_1 = \tau$, we may choose any $(m - 1)$ partitions and decrease their thresholds by 1. For the resulting vector T' , $\|T'\|_1 = \tau - m + 1$. We may use it as a stricter condition to generate candidates and the correctness of the algorithm is still guaranteed. We call the process of converting T to T' ϵ -transformation.

Theorem 1: The filtering condition based on the general pigeonhole principle is tight.

Proof: The correctness is stated in Lemma 4. We prove the minimality. Given a threshold vector T based on the general pigeonhole principle, i.e., $\|T\|_1 = \tau - m + 1$, we consider a threshold vector T' which dominates T , i.e., $\forall i \in \{1, \dots, m\}$, $T'[i] \leq T[i]$ and $[T'[i], T[i]] \cap [-1, n_i - 1] \neq \emptyset$, and $\exists j \in \{1, \dots, m\}$, $T'[j] < T[j]$. Because $\forall i \in \{1, \dots, m\}$, $H(x_i, q_i) \in [0, n_i]$ and $[T'[i], T[i]] \cap [-1, n_i - 1] \neq \emptyset$, we may construct a vector x such that $\forall i \in \{1, \dots, m\}$, $H(x_i, q_i) = \max(0, T'[i] + 1)$. $\forall i \in \{1, \dots, m\}$, because $T'[i] \leq T[i]$ and $[T'[i], T[i]] \cap [-1, n_i - 1] \neq \emptyset$, $H(x_i, q_i) \leq T[i] + 1$. Because $\exists j \in \{1, \dots, m\}$, $T'[j] < T[j]$, $\exists j \in \{1, \dots, m\}$, $H(x_i, q_i) \leq T[j]$. Because $H(x_i, q_i) > T'[i]$ on all the partitions, x is not a candidate by T' . However, $H(x, q) = \sum_{i=1}^m H(x_i, q_i) \leq \|T\|_1 + m - 1 = \tau$, meaning that x is result of the query. Therefore, the filtering condition based on T' is incorrect, and thus the minimality of T is proved. ■

One surprising but beneficial consequence of the ϵ -transformation is that the resulting threshold of a partition may

become *negative*. For example, $[1, 0, 0]$ becomes $[0, 0, -1]^4$ if the first and third partitions are chosen to decrease thresholds. Since $H(x_i, y_i)$ is a non-negative integer, $H(x_i, y_i) \leq T[i]$ is always false if $T[i]$ is negative. This fact indicates that the partitions with negative thresholds can be **safely ignored** for candidate generation. As will be shown in the next section, this allows us to ignore partitions where the query and most of the data are identical. This endows our method the *unique* ability to handle highly skewed data or partitions.

Example 4: Consider the four data vectors and two queries in Table II. For q^1 , we show the threshold vectors based on the

TABLE II
THRESHOLD VECTOR AND THEIR CANDIDATE SIZES

	Partition 1	Partition 2
$x^1 = 00000000$	000000	00
$x^2 = 00000111$	000001	11
$x^3 = 00001111$	000011	11
$x^4 = 10011111$	100111	11
$q^1 = 10000000$	100000	00
$q^2 = 10000011$	100000	11

q^1	$T = [2, 0]$	$Cand = \{x^1, x^2\}$
	$T = [1, 0]$	$Cand = \{x^1\}$
q^2	$T = [1, 0]$	$Cand = \{x^1, x^2, x^3, x^4\}$
	$T = [2, -1]$	$Cand = \{x^1, x^2\}$

flexible pigeonhole principle and the general pigeonhole principle. The candidate sizes are 2 and 1, respectively. For q^2 , we show two different threshold vectors based on the general pigeonhole principle. The candidate sizes are 4 and 2, respectively.

IV. THRESHOLD ALLOCATION

To utilize the general pigeonhole principle to process queries, there are two key issues: (1) how to divide the n dimensions into m partitions, and (2) how to compute the threshold vector T such that $\|T\|_1 = \tau - m + 1$. We will tackle the first issue in Section V with an offline solution. Before that, we focus on the second issue in this section and propose an online algorithm.

A. Cost Model

To optimize the threshold allocation, we first analyze the query processing cost. Like MIH, we also build an inverted index offline to map each partition of a data object to the object ID. Then for each partition of the query, we enumerate signatures to generate candidates.

The query processing cost consists of three parts:

$$C_{query_proc}(q, T) = C_{sig_gen}(q, T) + C_{cand_gen}(q, T) + C_{verify}(q, T),$$

where C_{sig_gen} , C_{cand_gen} , and C_{verify} denote the costs of signature generation, candidate generation, and verification, respectively.

⁴Note that in our method, we only consider the case of -1 for the negative threshold of a partition since the other negative values are not necessary.

For each partition i , a signature is a vector whose Hamming distance is within τ_i to the i -th partition of query q . Since we enumerate all such vectors, the signature generation cost is

$$C_{sig_gen}(q, T) = \sum_{i=1}^m \binom{n_i}{\tau_i} \cdot c_{enum},$$

where n_i denotes the number of dimensions in the i -th partition, and c_{enum} is the cost of enumerating the value of a dimension in a given vector. If $\tau_i < 0$, the cost is 0 for the i -th partition.

Let S_{sig} denote the set of signatures generated. The candidate generation cost can be modeled by inverted index lookup:

$$C_{cand_gen}(q, T) = \sum_{s \in S_{sig}} |I_s| \cdot c_{access},$$

where $|I_s|$ denotes the length of the postings list of signature s , and c_{access} is the cost of accessing an entry in a postings list.

The verification cost is

$$C_{verify}(q, T) = |S_{cand}| \cdot c_{verify},$$

where S_{cand} is the set of candidates, and c_{verify} is the cost to check if two n -dimensional vectors' Hamming distance is within τ .

In practice, the signature generation cost is usually much less than the candidate generation cost and the verification cost (see Section VII-B for experiments). So we can ignore the signature generation cost when optimizing the threshold allocation. In addition, it is difficult to accurately estimate the size of S_{cand} using the lengths of postings lists, because it can be reduced from the minimal k -union problem [35], which is proved to be NP-hard. Nonetheless, $|S_{cand}|$ is upper-bounded by the sum of candidates generated in all the partitions, i.e., $\sum_{s \in S_{sig}} |I_s|$. Our experiments (Section VII-B) show that the ratio of $|S_{cand}|$ and this upper bound depends on data distribution and τ . Given a dataset, the ratio with respect to varying τ can be computed and recorded by generating a number of queries and processing them. Let α denote this ratio. We may rewrite the number of candidates in the form of $\alpha \cdot \sum_{i=1}^m CN(q_i, \tau_i)$, where $CN(q_i, \tau_i)$ is the number of candidates generated by the i -th partition of the query q with a threshold of τ_i (when $\tau_i = -1$, $CN(q_i, \tau_i) = 0$). Hence the query processing cost can be estimated as:

$$C_{query_proc}(q, T) = \sum_{i=1}^m CN(q_i, \tau_i) \cdot (c_{access} + \alpha \cdot c_{verify}). \quad (1)$$

With the above cost model, we can formulate the threshold allocation as an optimization problem.

Problem 1 (Threshold Allocation): Given a collection of data objects \mathcal{D} , a query q and a threshold τ , find the threshold vector T that minimizes the estimated query processing cost under the general pigeonhole principle; i.e.,

$$\arg \min_T C_{query_proc}(q, T), \quad \text{s.t. } \|T\|_1 = \tau - m + 1.$$

Algorithm 1: DPAllocate(q, m, τ)

```

1 for  $e = -1$  to  $\tau$  do
2    $OPT[1, e] \leftarrow CN(q_1, e)$ ,  $PATH[1, e] \leftarrow e$ ;
3 for  $i = 2$  to  $m$  do
4   for  $t = -i$  to  $\tau - i + 1$  do
5      $c_{min} = +\infty$ ;
6     for  $e = -1$  to  $t + i - 1$  do
7       if  $OPT[i - 1, t - e] + CN(q_i, e) < c_{min}$  then
8          $c_{min} \leftarrow OPT[i - 1, t - e] + CN(q_i, e)$ ;
9          $e_{min} \leftarrow e$ ;
10     $OPT[i, t] = c_{min}$ ,  $PATH[i, t] = e_{min}$ ;
11  $e \leftarrow \tau - m + 1$ ;
12 for  $i = m$  to 1 do
13    $T[i] \leftarrow PATH[i, e]$ ;
14    $e \leftarrow e - PATH[i, e]$ ;
15 return  $T$ ;
```

B. Threshold Allocation Algorithm

Since c_{access} , c_{verify} , and α are independent of $CN(q_i, \tau_i)$, we can omit the coefficient ($c_{access} + \alpha \cdot c_{verify}$) in Equation 1 and find the minimum query processing cost with only $CN(q_i, \tau_i)$. The computation of $CN(q_i, \tau_i)$ values will be introduced in Section IV-C. Here we treat $CN(q_i, \tau_i)$ as a black box with $O(1)$ time complexity and propose an online threshold allocation algorithm based on dynamic programming.

Let $OPT[i, t]$ record the minimum query processing cost (omitting the coefficient ($c_{access} + \alpha \cdot c_{verify}$)) for partitions $1, \dots, i$ with a sum of thresholds t . We have the following recursive formula:

$$OPT[i, t] = \begin{cases} \min_{e=-1}^{t+i-1} OPT[i-1, t-e] + CN(q_i, e), & \text{if } i > 1; \\ CN(q_i, t), & \text{if } i = 1. \end{cases}$$

With the recursive formula, we design a dynamic programming algorithm for threshold allocation, whose pseudo-code is shown in Algorithm 1. It first initializes the costs for the first partition (Lines 1 – 2), i.e., $OPT[1, -1], \dots, OPT[1, \tau]$. Then it iterates through the other partitions and compute the minimum costs (Lines 3 – 10). Note that the negative threshold -1 is also consider for each partition. Finally, we trace the path that reaches $OPT[m, \tau - m + 1]$ to obtain the threshold vector (Lines 11 – 14). The time complexity of the algorithm is $O(m \cdot (\tau + 1)^2)$.

Example 5: Consider a dataset of 100 binary vectors and we partition it into 4 partitions. Given a query q , for each partition i , suppose the numbers of candidates (denoted CN_i) under different thresholds are provided in the table below.

	$\tau_i = -1$	$\tau_i = 0$	$\tau_i = 1$	$\tau_i = 2$	$\tau_i = 3$	$\tau_i = 4$
CN_1	0	5	10	15	50	100
CN_2	0	10	80	90	95	100
CN_3	0	5	15	20	70	100
CN_4	0	10	70	80	95	100

We use Algorithm 1 to compute the threshold vector. The $OPT[i, t]$ values are given in the table below.

t =	i = 1	i = 2	i = 3	i = 4
-3	0	0	0	5
-2	0	0	5	10
-1	0	5	10	20
0	5	15	20	30
1	10	20	20	30
2	15	25	35	45
3	50	60	40	45
4	100	110	45	55

The minimum query processing cost $OPT[4, 4] = 55$. We trace the path (in boldface) that reaches this value and obtain the threshold vector $[2, 0, 2, 0]$.

C. Computing Candidate Numbers

In order to run the threshold allocation algorithms, we need to obtain the candidate numbers $CN(q_i, \tau_i)$ beforehand. An exact solution to computing $CN(q_i, \tau_i)$ is to enumerate all possible vectors for the i -th partition and then count how many vectors in \mathcal{D} has a Hamming distance within τ_i to the enumerated vector in this partition. These numbers are stored in a table. When processing the query, with the given q_i , the table is looked up for the corresponding entry $CN(q_i, \tau_i)$. The time complexity of this algorithm is $O(m \cdot 2^n \cdot 2^\tau)$, and the space complexity is $O(m \cdot 2^n)$. This method is only feasible when n and τ are small. To cope with large n and τ , we devise two approximation algorithms to estimate the number of candidates.

Sub-partitioning. The basic idea of the first approximation algorithm is splitting q_i into smaller equi-width sub-partitions and estimating $CN(q_i, \tau_i)$ with the candidate numbers of the sub-partitions. We divide q_i into m_i sub-partitions. Each sub-partition has a fixed number of dimensions so that its candidate number can be computed using the exact algorithm in reasonable amount of time and stored in main memory. For the thresholds of the sub-partitions, we may use the general pigeonhole principle and divide τ_i into m_i values such that they sum up to $\tau_i - m_i + 1$. Let q_{ij} denote a sub-partition of q_i and τ_{ij} denote its threshold. Let $G(m_i, \tau_i)$ be the set of threshold vectors of which the total thresholds sum up to no more than $\tau_i - m_i + 1$; i.e., $\{[\tau_{i1}, \dots, \tau_{im_i}] | \tau_{ij} \in [-1, \tau_i] \wedge \sum_{j=1}^{m_i} \tau_{ij} \leq \tau_i - m_i + 1\}$.

We offline compute all the $CN(q_{ij}, \tau_{ij})$ values for all $\tau_{ij} \in [-1, \tau_i]$ using the aforementioned exact algorithm; i.e., enumerate all possible query vectors and then count how many data vectors in \mathcal{D} has a Hamming distance within τ_{ij} to the enumerated vector in this sub-partition. We assume that the candidates in the m_i sub-partitions are independent. Then $CN(q_i, \tau_i)$ can be approximately estimated online with the following equation.

$$CN(\widehat{q_i, \tau_i}) = \sum_{g \in G(m_i, \tau_i)} \prod_{j=1}^{m_i} (CN(q_{ij}, g[j]) - CN(q_{ij}, g[j] - 1)).$$

Machine Learning. We may also use machine learning technique to predict the candidate number for a given $\langle q_i, \tau_i \rangle$. For each τ_i , we regard each dimension of q_i as a feature and randomly generate feature vectors $x_k = [b_1, \dots, b_{|q_i|}]$. The candidate number $CN(x_k, \tau_i)$ can be obtained by processing

x_k as a query with a threshold τ_i . Then we apply the regression model on the training data $\mathcal{T}_i = \{\langle x_k, CN(x_k, \tau_i) \rangle\}$.

Let $h_{\tau_i}(x_k, \theta_i)$ denote the machine learning model, where θ_i denotes its parameters. Traditional regression models utilize mean squared error as loss function. To reduce the impact of large $CN(x_k, \tau_i)$, we use relative error as our loss function: $J(\mathcal{T}_i, \theta_i) = \sum_{k=1}^{|\mathcal{T}_i|} \left\{ \frac{CN(x_k, \tau_i) - h_{\tau_i}(x_k, \theta_i)}{CN(x_k, \tau_i)} \right\}^2$. According to [27], we utilize the approximation $\ln(t) \approx t - 1$ to estimate $J(\mathcal{T}_i, \theta_i)$:

$$\begin{aligned} J(\mathcal{T}_i, \theta_i) &= \sum_{k=1}^{|\mathcal{T}_i|} \left\{ 1 - \frac{h_{\tau_i}(x_k, \theta_i)}{CN(x_k, \tau_i)} \right\}^2 \\ &\approx \sum_{i=1}^{|\mathcal{T}_i|} \left\{ \ln \frac{CN(x_k, \tau_i)}{h_{\tau_i}(x_k, \theta_i)} \right\}^2 \\ &= \sum_{i=1}^{|\mathcal{T}_i|} \{ \ln CN(x_k, \tau_i) - \ln h_{\tau_i}(x_k, \theta_i) \}^2. \end{aligned}$$

From the above equation, we can simply convert training data $\langle x_k, CN(x_k, \tau_i) \rangle$ into $\langle x_k, \ln CN(x_k, \tau_i) \rangle$ and then take mean squared error to train an SVM model with RBF kernel.

V. DIMENSION PARTITIONING

To deal with data skewness and dimension correlations, the existing methods for Hamming distance search resort to random shuffle [1] or dimension rearrangement [43], [36], [20]. All of them are aiming towards the direction that the dimensions in each partition or the signatures in the index are uniformly distributed, so as to reduce the candidates caused by frequent signatures. In this section, we present our method for dimension partitioning. We devise a cost model of dimension partitioning and convert the partitioning into an optimization problem to optimize query processing performance. Then we propose the algorithm to solve this problem.

A. Cost Model

Let P_i denote a set of dimensions in the range $[1, n]$. Our goal is to find a partitioning $\mathcal{P} = \{P_1, \dots, P_m\}$ such that $P_i \cap P_j = \emptyset$ if $i \neq j$, and $\cup_{i=1}^m P_i = \{1, \dots, n\}$. Given a query workload $\mathcal{Q} = \{\langle q^1, \tau^1 \rangle, \dots, \langle q^{|\mathcal{Q}|}, \tau^{|\mathcal{Q}|} \rangle\}$, the query processing cost of the workload is the sum of the costs of its constituent queries:

$$C_{workload}(\mathcal{Q}, \mathcal{P}) = \sum_{i=1}^{|\mathcal{Q}|} C_{\widehat{query_proc}}(q^i, \tau^i, \mathcal{P}), \quad (2)$$

where $C_{\widehat{query_proc}}(q^i, \tau^i, \mathcal{P})$ is the processing cost of query q^i with a threshold τ^i , which can be computed using the dynamic programming algorithm proposed in Section IV. Then we can formulate the dimension partitioning as an optimization problem.

Problem 2 (Dimension Partitioning): Given a collection of data objects \mathcal{D} , a query workload \mathcal{Q} , find the partitioning \mathcal{P} that minimizes the query processing cost of \mathcal{Q} under the general pigeonhole principle; i.e.,

$$\arg \min_{\mathcal{P}} C_{workload}(\mathcal{Q}, \mathcal{P}).$$

Lemma 5: The dimension partitioning problem is NP-hard.

Proof: We can reduce the dimension partitioning problem from the number partitioning problem [2], which is to partition a multiset of positive integers, S , into two subsets S_1 and S_2 such that the difference between the sums in two sets is minimized. Consider a special case of $m = 2$ and a \mathcal{Q} of only one query. Let S be a multiset of n positive integers, each representing a dimension in the dimension partitioning problem. Let $sum(S)$ denote the sum of numbers in S . For $i \in \{1, 2\}$, Let $CN(q_i, \tau_i) = sum(S_i)^2, \forall \tau_i \in [-1, \tau]$; i.e., the candidate number in partition i equals to the square of the sum of numbers in this partition. By Equations 1 and 2, $C_{workload}(\mathcal{Q}, \mathcal{P}) = (sum(S_1)^2 + sum(S_2)^2) \cdot (c_{access} + \alpha \cdot c_{verify})$. $C_{workload}$ is minimized when the difference between $sum(S_1)$ and $sum(S_2)$ is minimized. Hence the special case of dimension partitioning problem is reduced from the number partitioning problem. Because the number partitioning problem is NP-complete, the dimension partitioning is NP-hard. ■

B. Partitioning Algorithm

Seeing the difficulty of the dimension partitioning problem, we propose a heuristic algorithm to select a good partitioning: first generate an initial partitioning and then refine it.

Algorithm 2 captures the pseudo-code of the heuristic partitioning algorithm. It first generates an initial partitioning \mathcal{P} of m partitions (Line 1). The details of the initialization step will be introduced in Section V-C. Then the algorithm iteratively improves the current partitioning by selecting the best option of moving a dimension from one partition to another. In each iteration, we pick a dimension from a partition P_i (Line 8), try to move it to another partition $P_j, j \neq i$ (Line 10), and compute the resulting query processing cost of the workload. We try all possible combination of P_i and P_j , and the option that yields the minimum cost is taken as the move of this iteration (Line 16). The above steps repeat until the cost cannot be further improved by moving a dimension. The time complexity of the algorithm is $O(lmnc)$. l is the number of iterations. c is the time complexity of computing the cost of the workload, $O(|\mathcal{Q}| \cdot m \cdot (\tau + 1)^2)$. We also note that due to the replacement of dimensions, partitions may become empty in our algorithm. Hence it is not mandatory to output exactly m partitions for an input partition number m .

For the input query workload \mathcal{Q} , in case a historical query workload is unavailable, a sample of data objects can be used as a surrogate. Our experiments show that even if the distribution of real queries are different from the query workload that we use to compute the partitioning, our query processing algorithm still achieves good performance (Section VII-G). We also note that we may assign varying thresholds to the queries in the workload \mathcal{Q} . The benefit is that we can offline compute the partitioning using the workload which cover a wide range of thresholds, and then build an index without being aware of the thresholds of real queries beforehand.

C. Initial Partitioning

Since the dimension partitioning algorithm stops at a local optimum, we may achieve a better result with a carefully

Algorithm 2: HeuristicPartition($\mathcal{D}, \mathcal{Q}, m$)

```

1  $\mathcal{P} \leftarrow \text{InitialPartition}(\mathcal{D}, \mathcal{Q}, m)$ ;
2  $c_{\min} \leftarrow C_{workload}(\mathcal{Q}, \mathcal{P})$ ;
3  $f \leftarrow \text{true}$ ;
4 while  $f = \text{true}$  do
5    $f \leftarrow \text{false}$ ;
6   foreach  $P_i \in \mathcal{P}$  do
7     foreach  $d \in P_i$  do
8        $P'_i \leftarrow P_i \setminus \{d\}, \mathcal{P}' \leftarrow (\mathcal{P} \setminus P_i) \cup P'_i$ ;
9       foreach  $P_j \in \mathcal{P}, j \neq i$  do
10         $P'_j \leftarrow P_j \cup \{d\}, \mathcal{P}' \leftarrow (\mathcal{P}' \setminus P_j) \cup P'_j$ ;
11        if  $C_{workload}(\mathcal{Q}, \mathcal{P}') < c_{\min}$  then
12           $f \leftarrow \text{true}$ ;
13           $c_{\min} \leftarrow C_{workload}(\mathcal{Q}, \mathcal{P}')$ ;
14           $\mathcal{P}_{\min} \leftarrow \mathcal{P}'$ ;
15   if  $f = \text{true}$  then
16      $\mathcal{P} \leftarrow \mathcal{P}_{\min}$ ;
17 return  $\mathcal{P}$ ;

```

selected initial partitioning. The correlation of dimensions play an important role here. Unlike the existing methods which try to make dimensions in each partition uniformly distributed, our method aims at the opposite direction. We observe that the query processing performance is usually improved if highly correlated dimensions are put into the same partitions. This is because our threshold allocation algorithm works online and optimizes each query individually. When highly correlated dimensions are put together, more errors are likely to be identified in a partition, and thus our threshold allocation algorithm can assign a larger threshold to this partition and smaller thresholds to the other partitions; i.e., choosing proper thresholds for different partitions. If the dimensions are uniformly distributed, all the partitions will have the same distribution and there is little chance to optimize for specific partitions.

We may measure the correlation of dimensions with entropy. For a partition P_i , we project all the data objects in \mathcal{D} on the dimensions of P_i , and use \mathcal{D}_{P_i} to denote the set of the resulting vectors. The correlation of the dimensions of P_i is measured by:

$$H(\mathcal{D}_{P_i}) = - \sum_{X \in \mathcal{D}_{P_i}} P(X) \cdot \log P(X).$$

According to the definition of entropy, a smaller value of entropy indicates a higher correlation of the dimensions of P_i . The entropy of the partitioning \mathcal{P} is the sum of the entropies of its constituent partitions:

$$H(\mathcal{P}) = \sum_{i=1}^m H(\mathcal{D}_{P_i}).$$

Our goal is to find an initial partitioning \mathcal{P} to minimize $H(\mathcal{P})$. To achieve this, we generate an equi-width partitioning in a greedy manner: Starting with an empty partition, we select the dimension which yields the smallest entropy if it is put into this partition. This is repeated until a fixed partition size $\lfloor \frac{n}{m} \rfloor$ is reached, and thereby the first partition is obtained. Then we

repeat the above procedure on the unselected dimensions to generate the other $(m - 1)$ partitions.

VI. THE GPH ALGORITHM

Based on the general pigeonhole principle and the techniques proposed in Sections IV and V, we devise the **GPH** (short for the **General Pigeonhole** principle-based algorithm for **H**amming distance search) algorithm.

The GPH algorithm consists of two phases: indexing phase and query processing phase. In the indexing phase, it takes as input the dataset \mathcal{D} , the query workload \mathcal{Q} , and a tunable parameter m for the number of partitions. The partitioning \mathcal{P} is generated using the heuristic partitioning algorithm proposed in Section V. Then for each n -dimensional vector x in \mathcal{D} , we divided it by \mathcal{P} into m partitions. Then for the projection of x on each partition, the ID of vector x is inserted into the postings list of this projection. In the query processing phase, the query q and the threshold τ are input to the algorithm. It first partitions q by \mathcal{P} into m partitions. Then the threshold vector T is computed using the dynamic programming algorithm proposed in Section IV. For the projection of q on each partition, we enumerate the signatures whose Hamming distances to the projection do not exceed the allocated threshold. Then for each signature, we probe the inverted index to find the data objects that have this signature in the same partition, and insert the vector IDs into the candidate set. The candidates are finally verified using Hamming distance and the true results are returned. We omit the pseudo-code here in the interest of space.

We note that in case of a query service level agreement, with the data and query workloads provided by the user, we are able to estimate the query response time of GPH using the cost returned by the threshold allocation algorithm. Then we can either guarantee the number of queries that can be handled in a specific amount of time using current computing resources, or let the user know the amount of additional resources required to process the queries.

VII. EXPERIMENTS

We report experiment results and analyses in this section.

A. Experiments Setup

The following algorithms are compared in the experiment.

- **MIH** is a method based on the basic pigeonhole principle [25]. It divides vectors into m equi-width partitions and uses a threshold $\lfloor \frac{\tau}{m} \rfloor$ on all the partitions to generate candidates. Its filtering condition is not tight. Signatures are enumerated on the query side. We utilize the open source of MIH on GitHub⁵ and chose the fastest m setting on each dataset.
- **HmSearch** is a method based on the basic pigeonhole principle [43]. Vectors are divided into $\lfloor \frac{\tau+3}{2} \rfloor$ equi-width partitions. It has a filtering condition in multiple cases but not tight. The threshold of a partition is either 0 or 1. This is one of our previous work and we utilize the existing source code.
- **PartAlloc** is a method to solve the set similarity join problem [11]. It divides vectors into $\tau + 1$ equi-width partitions

and allocate thresholds to partitions with three options: -1 , 0 , and 1 . -1 means that the partition is disregarded for candidate generation. Its filtering condition is tight. Signatures are enumerated on both data and query vectors. The source code is received from the authors of [11]. We convert the Hamming distance constraint to an equivalent Jaccard similarity constraint [1], which is supported by the source code. The greedy method [11] is chosen to allocate thresholds. Positional filters (checking the number of dimensions whose values are 1 in each partition and discarding a candidate if the difference exceeds τ) have already been implemented in the provided source code and are invoked when generating candidates.

- **LSH** is an algorithm to retrieve approximate answers. We convert the Hamming distance constraint to an equivalent Jaccard similarity constraint and then use the minhash LSH [5]. The dimension which yields the minimum hash value is chosen as a minhash. k minhashes are concatenated into a single signature, and this is repeated l times to obtain l signatures. We set k to 3 and recall to 95%. $l = \lceil \log_{1-t^k}(1-r) \rceil$, where t is the Jaccard similarity threshold. The algorithm is implemented by ourselves.
- **GPH** is the method proposed in this paper. We implement it on top of the source code of MIH for fair comparison.

Other methods for Hamming distance search, e.g., [18], [15], [22], are not compared since prior work [43] showed they are outperformed by HmSearch. We do not consider the method in [30] because it focuses on small n (≤ 64) and small τ (≤ 4), and it is significantly slower than the other algorithms in our experiments. E.g., on GIST, when $\tau = 8$, its average query response time is 128 times longer than GPH. The approximate method proposed in [26] is only fast for small thresholds. On SIFT, when $\tau \geq 12$, it becomes slower than MIH even if the recall is set to 0.9 [26]. Due to its performance compared to MIH and the much larger threshold settings in our experiments, we do not compare with the method in [26].

We select five publicly available real datasets with different data distributions and application domains.

- **SIFT** is a set of 1 billion SIFT features from the BIGANN dataset⁶ [13]. We follow the method used in [25] to convert them into 128-dimensional binary vectors.
- **GIST** is a set of 80 million 256-dimensional GIST descriptors for tiny images⁷ [33].
- **PubChem** is a database of chemical molecules⁸. We sample 1 million entries, each of which is a 881-dimensional vector.
- **FastText** is a set of 1 million English word vectors trained on Wikipedia 2017⁹. We convert them into binary vectors of 128 dimensions by spectral hashing [39].
- **UQVideo** is a set of 3.3 million keyframes of Web videos from the UQ Video project¹⁰. Each keyframe is converted to a binary vector of 256 dimensions by multiple feature hashing [31].

⁶<http://corpus-texmex.irisa.fr/>


⁷<http://horatio.cs.nyu.edu/mit/tiny/data/index.html>

⁸<https://pubchem.ncbi.nlm.nih.gov/>

⁹<https://fasttext.cc/docs/en/english-vectors.html>

¹⁰http://staff.itee.uq.edu.au/shenht/UQ_VIDEO/

⁵<https://github.com/norouzi/mih>



(a) Response Time Decomposed (b) Compare $\sum_{s \in S_{sig}} |I_s|$ and S_{cand}

Fig. 2. Justification of Assumptions

SIFT has the smallest skewness among the five. GIST and UQVideo are medium skewed datasets. PubChem and FastText are highly skewed datasets. In addition to the five real datasets, we generate a synthetic dataset with varying skewness.

We sample a subset of 100 vectors from each dataset as the query workload for the partitioning of GPH. To generate real queries, for each dataset we sample 1,000 vectors (**differ** from the query workload for partitioning) and take the rest as data objects. We vary τ and measure the query response time averaged over 1,000 queries. For GPH and PartAlloc, threshold allocation time are also included. The τ settings are up to 32, 64, 32, 20, and 48 on the five real datasets, respectively. The reason why we set smaller thresholds on PubChem is that due to the skewness, more than 10% data objects are results when $\tau = 32$.

The experiments are carried out on a server with a Quad-Core Intel Xeon E3-1231 @3.4GHz Processor and 96GB RAM, running Debian 6.0. All the algorithms are implemented in C++ in a main memory fashion.

B. Justification of Assumptions

We first justify our assumptions for the cost model of threshold allocation. Fig. 2(a) shows the query processing time of GPH on SIFT, GIST, and PubChem (denoted S, G, and P, respectively). The time is decomposed into four parts: threshold allocation, signature enumeration, candidate generation, and verification. The figure is plotted in **logscale** so that threshold allocation and signature enumeration can be seen. Compared to candidate generation and verification, the time spent on threshold allocation and signature enumeration is negligible ($< 3\%$), meaning that we can ignore them when estimating the query processing cost. Fig. 2(b) shows the sum of candidates generated in all the partitions ($\sum_{s \in S_{sig}} |I_s|$, denoted **dataset-sum**) and the candidate sizes ($|S_{cand}|$, denoted **dataset-cand**) on the three datasets. It can be seen that $|S_{cand}|$ is upper-bounded by $\sum_{s \in S_{sig}} |I_s|$. The ratio of them varies from 0.69 to 0.98, depending on dataset and τ . The ratios on different datasets and τ settings are recorded as the value of α in Equation 1 for cost estimation.

C. Evaluation of Threshold Allocation

We evaluate threshold allocation by comparing with a baseline algorithm (denoted RR). RR allocates thresholds in a round robin manner, and the thresholds of all partitions sum up to $\tau - m + 1$. For a fair comparison, we randomly shuffle the dimensions and then use the equi-width partitioning (m is chosen for the best




Fig. 3. Evaluation of Threshold Allocation

performance) for the competitors in this set of experiments. Figs. 3(a), 3(c), and 3(e) show the query processing costs (in terms of candidate numbers) estimated by DP on SIFT, GIST, and PubChem. We also plot the costs of RR using our cost model. The corresponding query response times are shown in Figs. 3(b), 3(d), and 3(f). The trends of the cost and the time are similar, indicating that the cost model effectively estimates the query processing performance. DP is significantly faster than RR in query processing, and the gap is more remarkable on datasets with more skewness. On PubChem, the time of RR is close to sequential scan due to the skewness. With judicious threshold allocation, the time is reduced by nearly two orders of magnitude.

To evaluate the candidate number computation, we compare the sub-partitioning algorithm (denoted SP) and the machine learning algorithm based on SVM model (denoted SVM). To show why we choose SVM as the machine learning model, we also compare with two other learning models: random forest (RF) and a 3-layer deep neural network (DNN). The number of sub-partitions is 2. The size of the training data is 1,000 for the machine learning algorithms. Table III shows the relative errors with respect to the exact method and the times of candidate number computation (in microseconds). Since the performances on the real datasets are similar, we only show the results on the GIST dataset. The relative error of SVM is very small, and it is more accurate and faster than SP. To compare learning models, the relative error of RF is much higher than the other methods. Although DNN estimates candidate numbers slightly more accurately than SVM in some settings, their relative errors are both very small, and the running time of DNN is much more than SVM. In addition, we tried logistic regression and gradient boosting decision tree. Their relative errors are higher than the above methods and hence not shown here. Seeing these results,




Fig. 4. Evaluation of Dimension Partitioning

we choose the machine learning algorithm based on SVM model to estimate candidate numbers in the rest of the experiments.

TABLE III
ESTIMATION WITH VARIOUS MODELS ON GIST (EACH CELL SHOWS PERCENTAGE ERROR AND PREDICTION TIME (μ S), SEPARATED BY /)

τ	SP	SVM	RF	DNN
16	1.75%/0.47	1.64%/0.31	8.73%/0.40	1.78%/2.64
32	0.37%/0.77	0.28%/ 0.28	12.43%/0.39	0.19% /2.60
48	0.15%/2.67	0.10%/ 0.43	9.26%/0.73	0.08% /3.83
64	0.07%/3.45	0.06%/ 0.29	3.58%/0.44	0.03% /2.44

D. Evaluation of Dimension Partitioning

To evaluate the effect of partitioning, we compare our method (denoted GR) with the following competitors: (1) OR is to use the original unshuffled order of the dataset. (2) RS is to perform a random shuffle on the original order. (3) OS [43] and DD [36] are two dimension rearrangement methods to make dimensions in each partition uniformly distributed. We run GPH with the above partitioning methods and show the query response times in Figs. 4(a), 4(c), and 4(e). On SIFT, their performances are close. When the dataset has more skewness, the advantage of GR becomes remarkable. It is faster than the runner-up by up to 4 times on GIST and 8 times on PubChem.

To evaluate the effect of initial partitioning, we run our partitioning algorithm with three initial states: (1) the proposed method which tries to minimize entropy (denoted GreedyInit), (2) equi-width partitioning on the original unshuffled data (denoted OriginalInit), and (3) equi-width partitioning after random shuffle (denoted RandomInit). The corresponding query response times on the three datasets are plotted in Figs. 4(b), 4(d), and 4(f). The trends are similar to the previous set of experiments. On datasets with more skewness, GreedyInit is consistently




Fig. 5. Effect of Partition Number

faster than the other competitors, and the gap to the runner-up can be up to 2 times.

As for the query workload Q to compute dimension partitioning, our results show that the effect of its size on the query processing performance is not obvious. E.g., when $\tau = 64$, the average query processing times vary from 4.19 to 3.97 seconds on GIST, if we increase $|Q|$ from 100 to 1000. Thus we choose 100 as the size of Q in our experiments.

We also study the effect of partition number on the query processing performance. Figs. 5(a) – 5(c) show the query response times on SIFT, GIST, and PubChem by varying the number of partitions. The general trend is that a smaller m performs better under small τ settings. When τ increases, the best choice of m slightly increases. The reason is: (1) When τ is small, a small m is good enough. Dividing vectors into unnecessarily large number of partitions yields very small partitions and hence increases the frequency of signatures. (2) When τ is large, a small m means more thresholds will be allocated to a partition, and this results in more candidates. Hence a slightly larger m is better in this case. Based on the results, we suggest user choose $m \approx \frac{n}{24}$ for GPH for good query processing performance.

E. Comparison with Existing Methods

We compare GPH with alternative methods (equipped with the OS partitioning [43]) for Hamming distance search.

Index are compared first. Figs. 6(a) – 6(e) show the index sizes of the algorithms on the five datasets. LSH, HmSearch, and PartAlloc run out of memory for some τ settings on SIFT and GIST. We only show the points when the memory can hold their indexes. GPH consume more space than MIH due to the machine learning-based technique to estimate candidate numbers. Both algorithms consume less space than the other exact competitors. This is expected as GPH and MIH enumerate signatures on query vectors only. HmSearch and PartAlloc enumerate 1-deletion variants on data vectors; i.e., removing an arbitrary dimension from a partition and taking the rest as a signature. The variants are indexed and this will increase their index sizes. PartAlloc and LSH exhibit variable index sizes with respect to τ . LSH has the smallest index size on PubChem, but consumes much more space on the other datasets.




Fig. 6. Comparison with Alternatives - Index Size

TABLE IV
INDEX CONSTRUCTION TIME ON GIST (s)

τ	MIH	HmSearch	PartAlloc	LSH	GPH
16	481	1681	1736	583	5026 + 560
32	481	1689	3244	5221	5026 + 560
48	481	1711	7600	64256	5026 + 560
64	481	1747	9605	N/A	5026 + 560

The reason is that PubChem has much more dimensions than the other datasets. Hence given a τ , the equivalent Jaccard threshold is higher on PubChem, resulting in less number of signatures. The corresponding index construction times on GIST are shown in Table IV. LSH runs out of memory when $\tau = 64$, and thus is shown for the other τ settings. The time of GPH is decomposed into dimension partitioning and indexing. MIH spends the least amount of time on index construction. Despite more time consumption on partitioning, GPH spends less time indexing data objects than the other algorithms. We argue that the partitioning can be done offline and the time is affordable. Because the query workload \mathcal{Q} for partitioning computation consists of queries with varying thresholds, we can run the partitioning once and use the same partitioning for different τ settings in real queries. This is also the reason why GPH has constant partitioning and indexing time irrespective of τ .

The candidate numbers are plotted in Figs. 7(a), 7(c), 7(e), 7(g), and 7(i). The corresponding query response times are plotted in Figs. 7(b), 7(d), 7(f), 7(h), and 7(j). For all the algorithms, candidate numbers and running times increase when τ moves towards larger values, and their trends are similar. Thanks to the tight filtering condition and cost-aware partitioning and threshold allocation, GPH is consistently smaller than MIH and HmSearch in candidate size and faster than the two methods. The only exception is that HmSearch has smaller candidate

size when $\tau = 4$ on PubChem and $\tau \leq 8$ on UQVideo, but turns out to be slower than GPH. This is because HmSearch generates many signatures whose postings lists are empty, and this drastically increases signature enumeration and index lookup times. Although PartAlloc has a tight filtering condition and utilizes threshold allocation, it is not as fast as GPH, and even slower than MIH. This result showcases that PartAlloc’s partitioning and threshold allocation is not efficient for Hamming distance search, though it pays off on set similarity search. Another interesting observation is that LSH does not perform well on highly skewed data. The reason is that the hash functions may choose highly skewed and correlated dimensions, and thus the selectivity of the chosen signatures becomes very bad. On PubChem, LSH’s performance is close to a sequential scan. Overall, GPH is the fastest algorithm. The speed-ups against the runner-up algorithms on the five datasets are up to 22, 21, 135, 32, and 8 times, respectively. We also notice that when $\tau \geq 16$ on FastText, the speed-up of GPH against MIH becomes less remarkable. This is because more than 59% data objects are results, meaning that both MIH and GPH are ineffective in filtering. Nonetheless, GPH is faster than MIH by 1.3 times, and both methods are significantly faster than the other competitors.

F. Varying Number of Dimensions

We compare the five competitors to evaluate their performances when varying the number of dimensions. We sample 25%, 50%, 75%, and 100% dimensions from SIFT, GIST, and PubChem, and then run the experiment. $\tau = 12, 24, 12$ for the 100% sample on the three datasets, respectively, and we let τ change linearly with the number of sampled dimensions. Figs. 8(a) – 8(c) show the query response times of the algorithms on the three datasets. We observe that the times of all the algorithms increase with n . There are two factors: (1) Although τ and n increase proportionally, the number of results increases with n due to dimension correlations. Hence we have more candidates to verify. (2) The verification cost increases with n because more dimensions are compared. Nonetheless, GPH is always the fastest algorithm among the competitors, especially on the more skewed PubChem.

G. Varying Skewness

We study the performance by varying skewness¹¹. As seen from Fig. 1, the relationship between skewness and dimensions is approximately linear (except PubChem) on most datasets. On the basis of this observation, the synthetic dataset is generated as follows: The number of dimensions is 128. The mean skewness is controlled by a parameter γ , and the skewnesses of the 128 dimensions range from 0 to 2γ . We set $\tau = 12$. The query processing times are plotted in Fig. 8(d). The general trend is that all the algorithms become slower on more skewed data. This is expected as signatures become less selective. Nonetheless, thanks to variable partitioning and threshold allocation, GPH is the fastest among the five competitors.

To demonstrate the robustness of GPH, we show that even if the distribution of real queries is different from the sample to

¹¹See the footnote in Section I for the measurement of dataset skewness.




Fig. 7. Comparison with Alternatives - Candidate Number & Time

compute partitioning, our method retains good performance. We generate a synthetic dataset with a γ of 0.5, and then compute partitioning with two query workloads: $\gamma = 0.5$ (denoted GPH-0.5) and $\gamma = 0.1$ (denoted GPH-0.1), respectively. Then we run a set of queries with a γ of 0.1. The gap between GPH-0.5 and GPH-0.1 can be regarded as the extent to which GPH's performance deteriorates in the presence of a different query distribution. Then we set γ to 0.1 for the synthetic dataset and run the experiment again. Results are plotted in Figs. 8(e) – 8(f). It can be seen that although GPH computes partitioning with a workload whose distribution is different from real queries, the query processing performance is almost the same. A slight difference can be noticed only when τ is as large as 12, where the query processing speed drops by 11.1% and 4.4%, respectively.

VIII. RELATED WORK

The notion of Hamming distance search was first proposed in [23]. Due to its wide range of applications, the problem has received considerable attention in the last few decades.




Fig. 8. Varying Number of Dimensions and Skewness

A few studies focused on the special case when $\tau = 1$ [3], [4], [21], [41]. Among them, the method by [21] indexes all the 1-variants of the data vectors to answer the query in $O(1)$ time and $O(\binom{n}{\tau})$ space. A data structure was proposed in [4] to answer this special case in $O(1)$ time using $O(n \log m)$ space by a cell probe model with word size m .

For the general case of Hamming distance search, the method by [10] is able to answer Hamming distance search in $O(m + \log^\tau(nm) + occ)$ time and $O(n \log^\tau(nm))$ space, where occ is the number of results. In practice, many solutions are based on the pigeonhole principle to convert the problem to sub-problems with a threshold τ' , where $\tau' < \tau$. In [32], [18], [25], vectors are divided into a number of partitions such that query results must have at least one exact match with the query in one of the partitions. The idea of recursive partitioning was covered in [22]. Before that, a two-level partitioning idea was adopted by the PartEnum method [1]. Song *et al.* [30] proposed to enumerate the combinations within threshold τ' in each partition to avoid the verification of candidates. Ong and Bober [26] proposed an approximate method utilizing variable length hash keys. In [43], vectors are divided into $\lfloor \frac{\tau+3}{2} \rfloor$ partitions, and the threshold of a partition can be either 0 or 1. Deng *et al.* [11] also proposed to use different thresholds on partitions, and the thresholds are computed by the allocation algorithm. We briefly discuss the differences from our method: (1) The method in [11] targets the set similarity join problem. Although it can be converted to Hamming distance searches, the number of dimensions are large (usually > 10000) and the vectors are sparse (usually < 1000 non-zero values), which are different from the datasets in most applications of Hamming distance search. (2) Equi-width partitioning is used in [11], and the number of partitions is $\tau + 1$. In GPH, we use variable partition size and the number

of partitions is a tunable parameter m . (3) For the thresholds on partitions, there are three options in [11]: 0, 1, and skipped (equivalent to -1 in GPH). We do not have this limitation, and the thresholds may vary from -1 to τ . (4) To find candidates, signature enumerations are performed on both data and query vectors in [11], whereas we only enumerate on queries.

To handle the poor selectivity caused by data skewness and dimension correlations, existing work mainly focused on two strategies. The first is to perform a random shuffle [1] in original dimensions to avoid highly correlated dimensions in same partitions. The second is to perform a dimension rearrangement [43], [36], [20] to minimize the correlation between dimensions in each partition. These methods are able to answer queries efficiently on slightly skewed datasets, but the performances deteriorate on highly skewed datasets.

We note that a strong form of the pigeonhole principle was introduced in [6] which states that given n positive integers q_1, \dots, q_m , if $(\sum_{i=1}^m q_i - m + 1)$ objects are distributed into m boxes, then either the first box contains at least q_1 objects, \dots , or the n -th box contains at least q_n objects. Although the general pigeonhole principle proposed in this paper coincides with the above strong form, by integer reduction and ϵ -transformation, the general pigeonhole principle is not limited to positive integers (this is the reason why GPH performs well on skewed data) and the tightness of threshold allocation is proved, hence providing a deeper understanding of the pigeonhole principle.

IX. CONCLUSION AND FUTURE WORK

In this paper, we proposed a new approach to similarity search in Hamming space. Observing the major drawbacks of the basic pigeonhole principle adopted by many existing methods, we developed a new form of the pigeonhole principle, based on which the condition of candidate generation is tight. The cost of query processing was modeled, and then an offline dimension partitioning algorithm and an online threshold allocation algorithm were devised on top of the model. We conducted experiments on real datasets with various distributions, and showed that our approach performs consistently well on all these datasets and outperforms state-of-the-art methods.

Our future work includes extending general pigeonhole principle to other similarity constraints. Another direction is to explore the techniques to dealing with the parallel case.

Acknowledgements. J. Qin, Y. Wang, and W. Wang are partially supported by ARC DP170103710, and D2DCRC DC25002 and DC25003. C. Xiao and Y. Ishikawa are supported by JSPS Kakenhi 16H01722. X. Lin is supported by NSFC 61672235, ARC DP170101628 and DP180103096. We thank the authors of [11] for kindly providing their source codes.

REFERENCES

- [1] A. Arasu, V. Ganti, and R. Kaushik. Efficient exact set-similarity joins. In *VLDB*, 2006.
- [2] J. M. Borwein and D. H. Bailey. *Mathematics by experiment - plausible reasoning in the 21st century*. A K Peters, 2003.
- [3] G. S. Brodal and L. Gasieniec. Approximate dictionary queries. In *CPM*, pages 65–74, 1996.
- [4] G. S. Brodal and S. Venkatesh. Improved bounds for dictionary look-up with one error. *Inf. Process. Lett.*, 75(1-2):57–59, 2000.
- [5] A. Z. Broder. On the resemblance and containment of documents. In *SEQS*, 1997.
- [6] R. Brualdi. *Introductory Combinatorics*. Math Classics. Pearson, 2017.
- [7] Z. Cao, M. Long, J. Wang, and P. S. Yu. Hashnet: Deep learning to hash by continuation. In *ICCV*, pages 5609–5618, 2017.
- [8] S. Chaidaroon and Y. Fang. Variational deep semantic hashing for text documents. In *SIGIR Conference*, pages 75–84, 2017.
- [9] B. Chazelle, D. Liu, and A. Magen. Approximate range searching in higher dimension. *Comput. Geom.*, 39(1):24–29, 2008.
- [10] R. Cole, L.-A. Gottlieb, and M. Lewenstein. Dictionary matching and indexing with errors and don't cares. In *STOC*, pages 91–100, 2004.
- [11] D. Deng, G. Li, H. Wen, and J. Feng. An efficient partition based method for exact set similarity joins. *PVLDB*, 9(4):360–371, 2015.
- [12] D. R. Flower. On the properties of bit string-based measures of chemical similarity. *Journal of Chemical Information and Computer Sciences*, 38(3):379–386, 1998.
- [13] H. Jégou, R. Tavenard, M. Douze, and L. Amsaleg. Searching in one billion vectors: re-rank with source coding. *CoRR*, abs/1102.3828, 2011.
- [14] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 614–623, New York, NY, USA, 1998. ACM.
- [15] C. Li, J. Lu, and Y. Lu. Efficient merging and filtering algorithms for approximate string searches. In *ICDE*, 2008.
- [16] W. Li, Y. Zhang, Y. Sun, W. Wang, W. Zhang, and X. Lin. Approximate nearest neighbor search on high dimensional data - experiments, analyses, and improvement (v1.0). *CoRR*, abs/1610.02455, 2016.
- [17] K. Lin, H. Yang, J. Hsiao, and C. Chen. Deep learning of binary hash codes for fast image retrieval. In *CVPR Workshops*, pages 27–35, 2015.
- [18] A. X. Liu, K. Shen, and E. Torng. Large scale hamming distance query processing. In *ICDE*, pages 553–564, 2011.
- [19] H. Liu, R. Wang, S. Shan, and X. Chen. Deep supervised hashing for fast image retrieval. In *CVPR Conference*, pages 2064–2072, 2016.
- [20] Y. Ma, H. Zou, H. Xie, and Q. Su. Fast search with data-oriented multi-index hashing for multimedia data. *TIIS*, 9(7):2599–2613, 2015.
- [21] U. Manber and S. Wu. An algorithm for approximate membership checking with application to password security. *Inf. Process. Lett.*, 50(4):191–197, 1994.
- [22] G. S. Manku, A. Jain, and A. D. Sarma. Detecting near-duplicates for web crawling. In *WWW*, pages 141–150, 2007.
- [23] M. Minsky and S. Papert. *Perceptrons - an introduction to computational geometry*. MIT Press, 1987.
- [24] R. Nasr, R. Vernica, C. Li, and P. Baldi. Speeding up chemical searches using the inverted index: The convergence of chemoinformatics and text search methods. *J. Chem. Inf. Model*, 2012.
- [25] M. Norouzi, A. Punjani, and D. J. Fleet. Fast search in hamming space with multi-index hashing. In *CVPR*, pages 3108–3115, 2012.
- [26] E. Ong and M. Bober. Improved hamming distance search using variable length hashing. In *CVPR Conference*, pages 2000–2008, 2016.
- [27] H. Park and L. Stefanski. Relative-error prediction. *Statistics & Probability Letters*, 40(3):227 – 236, 1998.
- [28] M. Raginsky and S. Lazebnik. Locality-sensitive binary codes from shift-invariant kernels. In *NIPS Conference*, pages 1509–1517, 2009.
- [29] R. Salakhutdinov and G. E. Hinton. Semantic hashing. *Int. J. Approx. Reasoning*, 50(7):969–978, 2009.
- [30] J. Song, H. T. Shen, J. Wang, Z. Huang, N. Sebe, and J. Wang. A distance-computation-free search scheme for binary code databases. *IEEE Trans. Multimedia*, 18(3):484–495, 2016.
- [31] J. Song, Y. Yang, Z. Huang, H. T. Shen, and R. Hong. Multiple feature hashing for real-time large scale near-duplicate video retrieval. In *Proceedings of the 19th ACM International Conference on Multimedia*, MM '11, pages 423–432, New York, NY, USA, 2011. ACM.
- [32] Y. Tabei, T. Uno, M. Sugiyama, and K. Tsuda. Single versus multiple sorting in all pairs similarity search. *Journal of Machine Learning Research - Proceedings Track*, 13:145–160, 2010.
- [33] A. Torralba, R. Fergus, and W. T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(11):1958–1970, 2008.
- [34] A. Torralba, R. Fergus, and Y. Weiss. Small codes and large image databases for recognition. In *CVPR Conference*, 2008.
- [35] S. A. Vinterbo. A note on the hardness of the k-ambiguity problem. Technical report, Harvard Medical School, 06 2002.
- [36] J. Wan, S. Tang, Y. Zhang, L. Huang, and J. Li. Data driven multi-index hashing. In *ICIP Conference*, pages 2670–2673, 2013.

- [37] J. Wang, H. T. Shen, J. Song, and J. Ji. Hashing for similarity search: A survey. *CoRR*, abs/1408.2927, 2014.
- [38] J. Wang, T. Zhang, J. Song, N. Sebe, and H. T. Shen. A survey on learning to hash. *CoRR*, abs/1606.00185, 2016.
- [39] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *NIPS Conference*, pages 1753–1760, 2008.
- [40] P. Willett, J. M. Barnard, and G. M. Downs. Chemical similarity searching. *Journal of Chemical Information and Computer Sciences*, 38(6):983–996, 1998.
- [41] A. C.-C. Yao and F. F. Yao. Dictionary look-up with one error. *J. Algorithms*, 25(1):194–202, 1997.
- [42] W. Zhang, K. Gao, Y. Zhang, and J. Li. Efficient approximate nearest neighbor search with integrated binary codes. In *ICMM Conference*, pages 1189–1192, 2011.
- [43] X. Zhang, J. Qin, W. Wang, Y. Sun, and J. Lu. Hmsearch: an efficient hamming distance query processing algorithm. In *SSDBM*, page 19, 2013.