

Efficient Binary Embedding of Categorical Data using BinSketch

Bhisham Dev Verma · Rameshwar Pratap ·
Debajyoti Bera *

Received: date / Accepted: date

Abstract In this work, we present a dimensionality reduction algorithm, *aka.* sketching, for categorical datasets. Our proposed sketching algorithm **Cabin** constructs low-dimensional binary sketches from high-dimensional categorical vectors, and our distance estimation algorithm **Cham** computes a close approximation of the Hamming distance between any two original vectors only from their sketches. The minimum dimension of the sketches required by **Cham** to ensure a good estimation theoretically depends only on the sparsity of the data points – making it useful for many real-life scenarios involving sparse datasets. We present a rigorous theoretical analysis of our approach and supplement it with extensive experiments on several high-dimensional real-world data sets, including one with over a million dimensions. We show that the **Cabin** and **Cham** duo is a significantly fast and accurate approach for tasks such as RMSE, all-pair similarity, and clustering when compared to working with the full dataset and other dimensionality reduction techniques.

Keywords Dimensionality Reduction · Sketching · Feature Hashing · Clustering · Categorical data.

* Corresponding author

Bhisham Dev Verma
School of Computing and Electrical Electrical Engineering,
IIT Mandi, H.P., India
E-mail: d18039@students.iitmandi.ac.in

Rameshwar Pratap
School of Computing and Electrical Electrical Engineering,
IIT Mandi, H.P., India
E-mail: rameshwar@iitmandi.ac.in

Debajyoti Bera
Indraprastha Institute of Information Technology Delhi,
New Delhi, India
E-mail: dbera@iiitd.ac.in

1 Introduction

Recent decades have witnessed the ability to generate a large volume of high-dimensional data arising out of the world-wide-web, IoT, various social media platforms, and applications of finance, biology, etc. Many of these datasets have large dimensions, sometimes in the order of millions [2,31]. A general observation is that many high-dimensional datasets are sparse in nature. Since computation on such large datasets is cumbersome, requires heavy computational machinery, and often suffers from the “curse of dimensionality” [39], one of the natural approaches to tackle this challenge is to compute a low-dimension representation (*aka sketch*) of each of their vectors which preserves inherent geometric properties of the corresponding full-dimensional datasets.

In this work, we focus on categorical datasets, where the feature values are from a finite set of categories, e.g., days in a week, colour, month, age group, gender, *etc.* Categorical features appear in many machine learning and data mining applications such as transaction datasets [15,23,3], images [22], in bio-informatics [29,36], in recommendation systems on click-stream data [37], *etc.* Further, Hamming distance (HD) appears to be the natural distance metric for categorical data points. Hamming distance between two n -dimensional categorical data points x and y is defined as:

$$HD(x, y) = \sum_{i=1}^n d(x_i, y_i), \text{ where } d(x_i, y_i) = \begin{cases} 1, & \text{if } x_i \neq y_i, \\ 0, & \text{otherwise.} \end{cases}$$

A common way to represent categorical data points is via *label-encoding*. If the number of possible categories is c , then we use an integer from $\{1, 2, \dots, c\}$ to represent a feature; 0 is used to represent a missing feature (e.g., if some data points do not have any attribute for “age group” while other points have that attribute, we say that the “age group” feature is missing in the former points and use 0 for the corresponding feature). In this paper, we focus on developing a sketching algorithm for high-dimensional and sparse categorical vectors. Sparsity of a feature vector denotes the percentage of zero (missing) entries in it, and the smallest sparsity across all the vectors in a dataset is defined as the sparsity of the dataset. To keep our analysis simple, we define the *density* of a label-encoded vector as its Hamming weight which is the number of non-missing features it has. Thus a data point with high sparsity (equivalently, low density) has lots of missing features.

Problem statement: Given high-dimensional and high-sparsity categorical data points, (i) develop an efficient dimensionality reduction algorithm that compresses the input points into low-dimensional binary vectors, and (ii) develop an efficient algorithm that can estimate the Hamming distances between the original data points from their sketches.

One of the naive approaches to perform dimensionality reduction for categorical data is to first represent it via binary vectors using *one-hot encoding*, where a feature value x is replaced by a $c + 1$ dimensional binary vector with 1 at the position x , and 0 otherwise. We can then further apply known dimensionality reduction

algorithms for binary data [33,28,35,34] on those binary vectors. However, this approach becomes impractical when the number of categories is large and may lead to an exponential blowup in the dimension of the resultant binary vector.

A quick browse through tutorials, forums and blogs on the Internet reveals a variety of alternatives [30, Tip 3], [13], mostly along the lines of the usual dimensionality reduction techniques like PCA, SVD, etc.; however, it is not clear why those techniques, originally developed for real-valued data, should work for discrete-valued data. Some of the techniques require matrix operations, solving an optimisation problem, or running a neural network, all of which are computationally costly. A few techniques are based on hashing or random projections; however, we did not find any theoretical guarantees they offer towards categorical vectors. It appears that many of the approaches used in practice were not designed for categorical vectors to start with, and have been merely re-purposed, possibly due to the lack of a sound alternative. To summarise, the above sketching algorithms are not specifically designed for categorical datasets as they do not offer any provable guarantee on the pairwise Hamming distance estimation from the compressed vectors. To address this gap, in this work we suggest a practical compression algorithm for such datasets that enables the aforesaid estimation with high accuracy.

Finally, some of the available methods do not produce compressed vectors in binary. We specifically want our compressed vectors to be binary for the following reasons. (i) Binary vectors are space-efficient as compared to the corresponding real-valued vectors of the same dimension – one feature of a real-valued vector requires 32 or 64 bits of memory (depending on word-size of a CPU), whereas one feature of a binary vector requires only one bit of storage. (ii) The binary vectors enable the possibility of using faster bitwise operators during training and inferencing steps which potentially can lead to several benefits such as less memory requirement, lower power consumption, and faster computation of machine learning tasks [20]. (iii) As our sketches are binary, and their pairwise Hamming distances are close to the original pairwise Hamming distance, this enables us to run the same machine learning algorithm on the sketch itself which we may not have been able to run on a categorical dataset due to its high dimensionality.

The major challenge that we faced is ensuring that the compressed vectors retain information about the pairwise distances between the original points. Even though there are several results known along this direction, e.g., the Johnson Lindenstrauss lemma [17], Feature Hashing [41], random projections for clustering [7], all of them deal with distances on the Euclidean space. Recently several methods were proposed for distances between discrete vectors, e.g., BCS [34], however, they were specifically designed for binary vectors without any scope for extension towards categorical vectors.

Our results: In this work, we present (i) the **Cabin** algorithm that compresses categorical vectors to succinct binary vectors, and (ii) the **Cham** algorithm to estimate the Hamming distance between the full-dimension vectors from their reduced-dimension embeddings.

Compared to the alternatives, our algorithms are designed specifically for categorical vectors, come with sound theoretical guarantees, and have excellent performance in practice. We discuss these benefits below.

- **Unsupervised:** Many feature-selection-based compression algorithms for categorical dataset such as χ^2 [27], mutual information-based feature selection methods [32] require labelled data. They essentially compute the correlation between the input features and their respective labels and discard those features whose correlations are smaller than a certain threshold. In contrast, our proposal is completely unsupervised and does not require any labels for compression.
- **Succinct and sparse embedding:** **Cabin** sketches have an attractive property — they are sparse when the original vectors are sparse. In fact, we prove in Lemma 4 that the number of ones in the **Cabin** embedding of a categorical vector u is at most half of the number of non-zero entries in u (in expectation). But there is a deeper connection to sparsity. Theoretically speaking, the required dimension of our sketch depends only on the density of input vectors and is independent of their original dimension. Thus, a vector with a large number of dimensions but, at the same time, of high sparsity can be compressed to an extremely small sketch yet retaining its useful theoretical properties. We observed that even smaller sketches sufficed in practice in our empirical exercises. Moreover, our algorithm outputs a binary-valued sketch which leads to additional space-saving as compared to a real-valued sketch of the same dimension. Finally, our succinct, sparse, and binary sketches not only provide saving in space required to store the sketch but also enable faster training and inference due to fewer arithmetic computations.
- **Superfast data analysis:** As desired above, our solution generates binary sketches bringing along with it significant advantages against real-valued sketches with respect to space, training, and inferencing time. Efficiency is further aided by the one-pass nature of our algorithms – **Cabin** and **Cham**. Further, the running time complexity of both **Cabin** and **Cham** is linear in their respective input size. Empirically, we obtained roughly $136\times$ speedup while generating a similarity heatmap of a Brain-Cell [42] dataset with 1.3 million features, and $112.3\times$ speedup in clustering of a 10^5 -dimensional NYTimes dataset [26] as compared to performing the tasks on the uncompressed full-dimensional datasets.
- **High accuracy:** The biggest advantage of our proposal is that of high-quality sketches. **Cabin** is able to compress very-high dimensional points to low-dimensional sketches such that on several downstream evaluation tasks, the results on the sketches closely approximate the corresponding results on the original input. We were able to theoretically explain this behaviour in the form of Theorem 2 which shows that the Hamming distance between two data vectors can be approximated with high accuracy from their sketches if we choose the embedding dimension as $\tilde{O}(s\sqrt{s})$ (ignoring $poly(\log)$ factors over the error probability) where s is an upper bound on the density of data. We observe much better compression and accuracy in practice. For example, **Cabin** is able to compress a Brain-Cell [42] dataset with 1.3 million features to only 1000-dimension binary vectors; furthermore, **Cham** ensures that we get almost identical looking heat-maps generated on both – full dimensional and compressed data (see Figure 11). Take another data analytic task, that of clustering. **Cabin** is able to compress a 10^5 -dimensional NYTimes dataset [26] to 1000-dimensions that still generate almost accurate results *vis-a-vis* the original dataset. The high accuracy of our proposal is also validated by the variance analysis experiments. We notice that the variance of the inaccuracies arising out of **Cham** is small,

which reaffirms our claim on accurate estimation of the Hamming distances of original data points (see Figures 4, 5). Even in our comparison with the related methods, **Cabin** is super-fast, yet offering a comparable performance with respect to heat-map generation, *root-mean-square error* (RMSE) test, and clustering.

Our **Cabin** and **Cham** algorithms follow a two-step approach. In the first step, we reduce the embedding problem over categorical vectors to the same over binary vectors, and in the second step, we solve the embedding problem over binary vectors. The two steps allow us to meet the twin requirements of *compressed* and *binary* embeddings. The first step is commonly performed in data analysis using one-hot encoding or other deterministic methods. Our proposal is to use a random binary encoding that not only retains the original dimension but also preserves the Hamming distance (see Lemma 1 and Lemma 2). For the second step, we observed that a few candidates BinSketch [33], BCS [34], MinHash [8], SimHash [9], OddSketch [28] are already available for embedding binary vectors. We decided to use BinSketch for the second step for a few important reasons. First is that we found BinSketch to perform well in our experiments (see Section 5). Secondly, some of the alternatives do not have the desired theoretical properties, namely, allow approximation of Hamming distance and generation of binary vectors. BinSketch not only generates binary sketches, but it also allows the approximation of Hamming as well as cosine, Jaccard, and inner product distances, making it a natural choice for us. Lastly, we were able to prove additional properties of BinSketch embeddings that played a crucial role in our main theoretical result which is summarised below.

Informal version of Theorem 2: *Let $\hat{u}, \hat{v} \in \{0, 1\}^d$ denote the **Cabin** embeddings of two categorical vectors $u, v \in \{0, c\}^n$, respectively, where c denote the number of categories, and d is a suitably chosen constant that depends only on the sparsity of u and v . Then the output of **Cham**(\hat{u}, \hat{v}) is close to the Hamming distance of u and v , with high probability.*

Our objective behind **Cabin** and **Cham** are two-fold — state-of-the-art performance in experiments and theoretical bounds to explain the behaviour. Some of the choices we made in the design of our method and empirical evaluation are driven by the second requirement which we believe is crucial for explainable data analysis. In particular, there could be a better alternative to BinSketch, and further, an integrated single-step solution for the whole approach. We leave this as a future research direction.

Organisation of the paper: The rest of the paper is organised as follows. We discuss several related works in Section 2. In Section 3, we briefly revisit the BinSketch algorithm and state other preliminary definitions. In Section 4, we present our algorithms **Cabin** and **Cham** and derive their theoretical bounds. In Section 5, we empirically compare the performance of our proposal on several end tasks with state-of-the-art algorithms. In Section 6, we conclude our discussion and state some potential open questions of the work.

2 Related work

Dimensionality reduction or compression of high-dimensional points is a well-studied phenomenon in data mining and machine learning. Such algorithms can be broadly grouped into the following four categories: (i) random projection, (ii) feature hashing, (iii) spectral projection, and (iv) locality sensitive hashing (LSH). At a high level, all such algorithms compress high dimensional vectors into low-dimensional vectors that maintain some measure of similarity among the input vectors. We discuss them below.

JL-lemma (or random projection) [17] is a seminal algorithm that popularised random projection for dimensionality reduction. Their algorithm essentially projects the input matrix on another matrix (called a projection matrix) each of whose entries is sampled from the normal distribution $\mathcal{N}(0,1)$. The *Euclidean distance and inner product of the* low-dimensional vectors obtained after random projection closely approximate the Euclidean distance and inner product, respectively, of the original data points. Achlioptas [1] improved over this work by suggesting a faster algorithm which is particularly suitable in database applications. Their contribution lies in suggesting a projection matrix whose each entry is sampled from the *Rademacher distribution* ($\{-1, +1\}$ with probability $1/2$), which leads to faster projection. Realising the need to generate sparse sketches, Li *et.al.* [25] modified random projection by carefully choosing the entries of the projection matrix to make the resultant sketch sparse.

Similar to random projection, feature hashing [41] also offers dimensionality reduction while approximating Euclidean distance and inner product. It is also known to maintain the sparsity of input in the sketch. In comparison with `Cabin`, it is unclear whether random projection or feature hashing algorithms can be applied on a categorical dataset for approximating Hamming distance. It can be noted that similar to feature hashing, `Cabin` also maintains the sparsity of the input data in the sketch (see Lemma 4).

Another recently proposed approach on the lines of random projection is `BinSketch` [33] that suggests an efficient compression algorithm for sparse binary datasets. It compresses high-dimensional binary vectors into low-dimensional binary vectors and includes a method that accurately approximates various similarity measures such as Hamming distance, cosine similarity, Jaccard similarity, and inner product. However, as `BinSketch` only operates on binary vectors, we can not directly use it for compressing categorical vectors. One naive approach is to first convert the categorical vectors into binary vectors using *one-hot encoding* and then use `BinSketch` on it to generate a binary sketch. However, the dimension of a binary embedding obtained via *one-hot encoding* is $d \times c$, where d and c are the original dimension and the number of categories, respectively. Therefore, this approach becomes impractical when the number of categories and the dimension of input vectors are large. The novelty in our approach is to suggest a *randomised* encoding of categorical data into binary vectors of the same dimension that allows us to approximate the original pairwise Hamming distances. `BinSketch` can be applied to this binary sketch to further compress it into low dimensional binary vectors; the original pairwise Hamming distances can then be approximated from those vectors. Note that there are other known compression algorithms for binary vectors such as BCS [34]. However, we prefer to use `BinSketch` as it offers both better theoretical as well as practical guarantees on the quality of its estimation. Nev-

ertheless, we included BCS as a baseline for empirical comparison and observed that our proposal is better in practice.

Principal component analysis (PCA) is another popular method for dimensionality reduction that creates uncorrelated features that successively maximise variance. Multiple Correspondence Analysis (MCA) [5] is analogous of PCA but designed for categorical datasets. However, our aim is to estimate the pairwise Hamming distances of the data points from their sketches, whereas that of MCA is to deduce uncorrelated features. Locality-sensitive hashing (LSH) [16,9,8,12] is another important line of dimensionality reduction algorithms that are primarily used for nearest neighbour search problems. In LSH, points are grouped in a way such that similar points are hashed into the same bucket and dissimilar points hashed into different buckets, with high probability. Thus, LSH algorithms like SimHash (SH), Hamming-LSH (H-LSH) do not provide explicit estimates of any type of similarities which we aim for in this work.

There are some learning-based sketching algorithms such as Latent Semantic Analysis (LSA) [11], Latent Dirichlet Allocation (LDA) [6], Non-negative Matrix Factorisation (NNMF) [24], Variational auto-encoder (VAE) [21], *etc.* that learn low-dimensional representations while maintaining some inherent properties of the input. However, we are not aware of any such technique designed for categorical datasets that maintains or approximates Hamming distances.

Outside of feature selection and dimensionality reduction, finding space-efficient sketches for approximating Hamming distance has been well studied in the streaming algorithms framework. Cormode et al. [10] described how to compute succinct sketches of large data streams which closely approximates the Hamming distance between the data streams. This result was subsequently improved by Kane et al. [18] by providing an optimal bound on the size of the sketch. These algorithms output real-valued sketches that are optimal only in the asymptotic sense, and are difficult to implement; in contrast to those, here we aim for a compression algorithm that outputs binary sketches and are also practically efficient.

To the best of our knowledge, there is no dimensionality reduction algorithm available that compresses high-dimensional categorical data into low-dimensional categorical data that can be used to closely approximate the original pairwise Hamming distances. Many sketching algorithms such as PCA, MCA, LSA, LDA, NNMF generate real-valued sketches whereas some others such as SH, H-LSH, Kendall rank correlation test (KT) output discrete sketches but they are not known to approximate pairwise Hamming distances. We nevertheless empirically compare our results with most of the methods mentioned above. We obtained significant speedup in dimensionality reduction time while being able to accurately estimate the Hamming distances from the low-dimensional vectors. We discuss these in detail in Section 5.

3 Background

3.1 Revisiting BinSketch algorithm [33]

We first recall BinSketch — an algorithm to compress sparse binary vectors into binary sketches that preserve several similarity measures such as Hamming distance, inner product, cosine, and Jaccard similarity.

Definition 1 (BinSketch [33]) Let ϕ be a random mapping from $\{1, \dots, n\}$ to $\{1, \dots, k\}$. Then a vector $a \in \{0, 1\}^n$ is compressed into a vector $a_s \in \{0, 1\}^k$ as

$$a_s[j] = \bigvee_{i:\phi(i)=j} a[i],$$

where $a[i]$ denotes the i -th index of the vector a , and \bigvee denotes the **bitwise-OR** operator.

Theorem 1 of [33] (restated below) delivers the guarantee of the estimate of inner product similarity from BinSketch sketches; BinSketch was shown to also approximate a few other similarity measures which are not useful for us.

Theorem 1 (Theorem 1 of [33] – Inner product estimation.) *Suppose we want to estimate the inner product of two d -dimensional binary vectors $a, b \in \{0, 1\}^n$, whose density is at most s , with probability at least $1 - \delta$. We can use BinSketch to construct their k -dimensional binary sketches, where $k = s\sqrt{s/2 \ln 1/\delta}$. Then using the sketches of a and b , the inner product between a and b can be estimated with accuracy $O(\sqrt{s \ln(1/\delta)})$.*

3.2 Evaluation metrics for clustering performance

Consider a clustering task on some dataset. Let m denote the number of data points, $\{\omega_1, \omega_2, \dots, \omega_k\}$ represent the ground-truth clustering results, and $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$ represent a clustering on the reduced dimension data. In what follows, we discuss some important metrics to estimate the quality of the clustering \mathcal{C} .

- **Purity index:** The *purity index* of \mathcal{C} is defined as

$$\frac{1}{m} \sum_{i=1}^k \max_{1 \leq j \leq k} |\omega_i \cap c_j|.$$

The *purity-index* takes a value between 0 and 1 — closer to 1 indicates better performance.

- **Normalised Mutual Information (NMI):** The *NMI* of \mathcal{C} is defined as

$$\sum_k \sum_j \frac{|\omega_k \cap c_j|}{m} \log \frac{m |\omega_k \cap c_j|}{|\omega_k| \cdot |c_j|}.$$

The *NMI* index takes a value between 0 and 1 — closer to 1 indicates better performance.

- **Adjusted Rand Index (ARI):** The *ARI* score of \mathcal{C} is defined as

$$\frac{\sum_{i,j} \binom{m_{ij}}{2} - \left[\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2} \right] / \binom{m}{2}}{\frac{1}{2} \left[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2} \right] - \left[\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2} \right] / \binom{m}{2}},$$

where $m_{ij} := |\omega_i \cap c_j|$, $a_i := \sum_j m_{ij}$, and $b_j := \sum_i m_{ij}$. The *ARI* takes a value between -1 and 1 — closer to 1 indicates better performance.




Fig. 1 An illustration of **Cabin** generating the binary embedding $\langle 100010 \rangle$ from a categorical point with the feature vector $\langle 40200102003104 \rangle$ (0 values in this indicate missing features). The random mappings that were used are shown on the left.

:q

4 Cabin and Cham- algorithm and analysis

In this section, we present our **Cabin** sketching algorithm and an algorithm to estimate a Hamming distance from the sketches.

Suppose we want to run **Cabin** on a dataset with n -dimensional vectors and each attribute of a vector could either be missing or must belong to at most c categories (for example, the 5^{th} attribute could be day and that can take a value from $\{\text{Sunday, Monday, } \dots, \text{Saturday}\}$). We will assume that the categories are represented by $\{1, 2, \dots, c\}$ by some data transformation; if some attribute, say the i -th one, is missing then the i -th coordinate of the vector will be assigned 0. It is entirely possible to have different sets of categories for each attribute (for example, day of week for the 5^{th} attribute and month of year as the 6^{th} attribute) as long as we have an upper bound on the largest number of categories of any attribute – this bound is denoted c . So at the end of the data-transformation, we end up with vectors from $\{0, 1, \dots, c\}^n$ which form the input to **Cabin**.

Cabin embeds an n -dimensional data vector to a d -dimensional binary vector where $d \ll n$ and d is chosen as $s\sqrt{\frac{s}{2} \ln \frac{6}{\delta}}$ in which s denotes an upper bound on the density of u and v and δ is the desired probability of error. It uses two uniformly random mappings,

1. category mapping $\psi : \{0, 1, \dots, c\} \xrightarrow{R} \{0, 1\}$
2. attribute mapping $\pi : \{1, 2, \dots, n\} \xrightarrow{R} \{1, 2, \dots, d\}$

and operates in the following two stages.

1. Generate an n -dimensional binary vector, say u' from an n -dimensional category vector, say u , using the **BinEm** algorithm. **BinEm** uses ψ .
2. Generate a d -dimension binary sketch from u' . For this step, we chose to use a recently proposed binary sketching technique named BinSketch [33]. BinSketch uses π . Note that BinSketch can be replaced with any other sketching algorithm for binary vectors that allows us to estimate a Hamming distance with theoretical guarantees.

We describe these two components in Algorithm 1; see Figure 1 for an illustration. We include the code for BinSketch for completeness.

Next, we discuss two important properties of the output vectors of **BinEm**. The first is that these vectors are extremely sparse, and this is the reason we chose BinSketch in stage-2 which is known to be highly efficient for sparse binary data.

Algorithm 1 Sketching algorithm for categorical vectors

```

1: function Cabin( $u$ )                                ▷  $u \in \{0, 1, \dots, c\}^n$ 
2:    $u' = \text{BinEm}(u)$                                 ▷  $u' \in \{0, 1\}^n$ 
3:    $\tilde{u} = \text{BinSketch}(u')$                           ▷  $\tilde{u} \in \{0, 1\}^d$ 
4:   return  $\tilde{u}$ 
5: end function
6: function BinEm( $u$ )                                ▷  $u \in \{0, 1, \dots, c\}^n$ 
7:    $u' = 0^n$ 
8:   for all non-empty attribute  $a$  in  $u$  do
9:     let  $i$  be the position of  $a$                     ▷  $u_i = a$ 
10:    set  $u'_i = \psi(a)$ 
11:  end for
12:  return  $u'$ 
13: end function
14: function BinSketch( $u'$ )                          ▷  $u' \in \{0, 1\}^n$ 
15:   $\tilde{u} = 0^d$ 
16:  for all non-zero bit in  $u'$  do
17:    let  $i$  be the position of the bit
18:    set  $\tilde{u}_{\pi(i)}$  to 1
19:  end for
20:  return  $\tilde{u}$ 
21: end function

```

Lemma 1 Consider any u with a non-zero attributes. Let a' denote the number of non-zero bits in $\text{BinEm}(u)$. Then the following three claims hold.

$$(a) a' \leq a, \quad (b) \mathbb{E}[a'] = \frac{a}{2}, \quad (c) \Pr[|a' - \frac{a}{2}| \geq \epsilon_1] \leq \exp(-\frac{2\epsilon_1^2}{a}).$$

Proof Fact (a) is obvious from the observation that if $u_i = 0$ then $u'_i = 0$ as well.

For facts (b) and (c), observe that a' can be treated as the sum of a independent and identically distributed binary random variables with 0.5 success probability; the statements follow from standard analysis of the number of heads among n tosses of a fair coin.

Let u' denote $\text{BinEm}(u)$ and v' denote $\text{BinEm}(v)$. The second property says that the Hamming distance between the outputs of BinEm is sufficient to compute the original distance.

Lemma 2 (a) $HD(u, v) = 2\mathbb{E}[HD(u', v')]$, and

$$(b) \Pr[|HD(u', v') - HD(u, v)/2| > \epsilon_2] \leq \exp(-\frac{2\epsilon_2^2}{HD(u, v)}).$$

Proof For every $i \in \{1, \dots, n\}$, let W_i be an indicator variable that is equal to 1 when $u_i \neq v_i$, and W'_i be an indicator variable that is equal to 1 when $u'_i \neq v'_i$. Observe that $HD(u, v) = \sum_i W_i$ and $HD(u', v') = \sum_i W'_i$.

We make two observations. First is that $u'_i = v'_i$ whenever $u_i = v_i$. This is true for both the cases of $u_i = v_i = 0$ and $u_i = v_i = a \in \{1, \dots, c\}$. From this observation we get that if $W_i = 0$ then $W'_i = 0$.

The second observation is that when $u_i \neq v_i$ then $\Pr[u'_i \neq v'_i] = \frac{1}{2}$ (or equivalently, $\mathbb{E}[W'_i] = \frac{1}{2}$ when $W_i = 1$). To see this consider these three cases.

1. $u_i = 0 \neq v_i$: For this case $u'_i = 0$ but $v'_i = 1$ with probability $\frac{1}{2}$.
2. $u_i \neq 0 = v_i$ (this is similar to the above case)

3. $u_i \neq 0 \neq v_i$: For this case, u'_i and v'_i are mapped randomly and independently to 0 or 1, so, $u'_i \neq v'_i$ with probability $\frac{1}{2}$.

Let h denote $HD(u, v)$. Without loss of generality, we can consider $W_i = 1$ for $i = 1 \dots h$, and $W_i = 0$ for all other values of i . Now from the above observations we can see that $HD(u', v')$ is essentially a sum of h Bernoulli random variables W'_1, \dots, W'_h , each having a probability $\frac{1}{2}$ of success. It immediately follows that $\mathbb{E}[HD(u', v')] = \mathbb{E}[\sum_{i=1}^h W'_i] = h/2 = HD(u, v)/2$ proving claim (a). Claim (b) simply says that this sum is tightly concentrated around its mean, and this can be obtained by a straight forward application of Chernoff-Hoeffding's bound for additive error.

Now we move to our algorithm for estimating Hamming distance between two vectors, say u and v solely from their **Cabin** sketches. We call it **Cham** and it is described in Algorithm 2. It employs an algorithm to estimate the Hamming distance between two binary vectors from their compressions generated by BinSketch [33, Algorithm 2] — we refer to this algorithm as **BinHamming** and include it for readability. We have used D to denote $1 - \frac{1}{d}$ in the algorithm.

Algorithm 2 Algorithm to estimate Hamming distance

```

1: function Cham( $\tilde{u}, \tilde{v}$ )                                     ▷  $\tilde{u}, \tilde{v} \in \{0, 1\}^d$ 
2:   Estimate  $\tilde{h} = \text{BinHamming}(\tilde{u}, \tilde{v})$                        ▷ [33, Algo. 2]
3:   return  $2\tilde{h}$ 
4: end function
5: function BinHamming( $\tilde{u}, \tilde{v}$ )
6:   Compute  $|\tilde{u}| = \text{Hamming weight of } \tilde{u}$ 
7:   Compute  $|\tilde{v}| = \text{Hamming weight of } \tilde{v}$ 
8:   Compute  $\langle \tilde{u}, \tilde{v} \rangle = \text{bitwise inner-product of } \tilde{u} \text{ and } \tilde{v}$ 
9:   Compute  $\tilde{h} = \frac{1}{\ln D} \left( D^{|\tilde{u}|} + D^{|\tilde{v}|} + \frac{\langle \tilde{u}, \tilde{v} \rangle}{d} - 1 \right)$ 
10:  return  $\tilde{h}$ 
11: end function

```

The fact that $\text{Cham}(\text{Cabin}(u), \text{Cabin}(v))$ returns a good estimator of $HD(u, v)$ should be apparent from the results given above and the properties of **BinHamming**. We briefly show how to derive a concentration bound on \tilde{h} output by **BinHamming** and then formally prove a concentration bound on **Cham** with additive accuracy.

Lemma 3 *Let h denote the Hamming distance of two d -dimensional binary vectors \tilde{u} and \tilde{v} , and \tilde{h} denote the output of $\text{BinHamming}(\tilde{u}, \tilde{v})$. Then, with probability at least $1 - \delta$, $|h - \tilde{h}| \leq 6\sqrt{\frac{s}{2} \ln \frac{6}{\delta}}$.*

Lemma 3 is not proved explicitly in the BinSketch paper. However, it is included as an intermediate step for proving Lemma 12 (refer to the upper bound on W in [33, Appendix B]). We will also require the fact that the density (number of ones) of \tilde{u} and \tilde{v} is at most s which follows from Lemma 1 claim (a).

Now we come to our main result stating the effectiveness of **Cham**. Fix a small number δ close to zero. Suppose vectors $u, v \in \{0, 1, \dots, c\}^n$ were compressed first to $u', v' \in \{0, 1\}^n$ using **BinEm** and then to $\tilde{u}, \tilde{v} \in \{0, 1\}^d$ using BinSketch. From

Lemma 3 we know that

$$\Pr \left[|2 \cdot HD(u', v') - 2 \cdot \text{BinHamming}(\tilde{u}, \tilde{v})| > 12\sqrt{\frac{s}{2} \ln \frac{6}{\delta}} \right] \leq \delta,$$

and using $\epsilon_2 = \sqrt{s \ln \frac{6}{\delta}}$ in Lemma 2, we get that

$$\begin{aligned} \Pr \left[|2 \cdot HD(u', v') - HD(u, v)| > 2\sqrt{s \ln \frac{6}{\delta}} \right] &\leq \exp\left(-\frac{2s}{HD(u, v)} \ln \frac{6}{\delta}\right). \\ &\leq \exp\left(-\ln \frac{6}{\delta}\right) = \frac{\delta}{6}. \end{aligned}$$

For the last inequality, we used the fact that $HD(u, v) \leq 2s$ which follows from the sparsity assumption of the input data. Combining these two inequalities, we arrive at

$$\begin{aligned} &\Pr \left[|2HD(u', v') - HD(u, v)| > 2\sqrt{s \ln \frac{6}{\delta}} \quad \text{OR} \right. \\ &\quad \left. |2HD(u', v') - 2 \cdot \text{BinHamming}(\tilde{u}, \tilde{v})| > 12\sqrt{\frac{s}{2} \ln \frac{6}{\delta}} \right] \leq \frac{\delta}{6} + \delta = \frac{7\delta}{6} \\ \implies &\Pr \left[|2HD(u', v') - HD(u, v)| \leq 2\sqrt{s \ln \frac{6}{\delta}} \quad \text{AND} \right. \\ &\quad \left. |2HD(u', v') - 2 \cdot \text{BinHamming}(\tilde{u}, \tilde{v})| \leq 12\sqrt{\frac{s}{2} \ln \frac{6}{\delta}} \right] \geq 1 - \frac{7\delta}{6} \\ \implies &\Pr \left[|HD(u, v) - 2 \cdot \text{BinHamming}(\tilde{u}, \tilde{v})| \leq \left(2 + \frac{12}{\sqrt{2}}\right)\sqrt{s \ln \frac{6}{\delta}} \right] \geq 1 - \frac{7\delta}{6} \\ \implies &\Pr \left[|HD(u, v) - 2 \cdot \text{BinHamming}(\tilde{u}, \tilde{v})| \leq 11\sqrt{s \ln \frac{6}{\delta}} \right] \geq 1 - \frac{7\delta}{6} \\ \implies &\Pr \left[|HD(u, v) - 2 \cdot \text{BinHamming}(\tilde{u}, \tilde{v})| \geq 11\sqrt{s \ln \frac{6}{\delta}} \right] \leq \frac{7\delta}{6}, \end{aligned}$$

and adjusting the probabilities further, we obtain our main result for **Cham**.

Theorem 2 *For a small $\delta \in (0, 1)$, and n -dimensional category vectors u, v , $|\text{Cham}(\tilde{u}, \tilde{v}) - HD(u, v)| > 11\sqrt{s \ln \frac{7}{\delta}}$ with probability at most δ , where \tilde{u} and \tilde{v} are the outputs of **Cabin**(u, v), respectively.*

Even though the theorem requires that the density of u, v be at most s and then stipulates that the estimate of the Hamming distance has at most $O(\sqrt{s})$ additive error, we should point out that these bounds are obtained using loose probability inequalities. We show in the next section that the performance is superior on real datasets.

We end this section with a quick fact that **Cabin** retains, in fact, improves sparsity. We prove this fact below in expectation.

Lemma 4 *Consider any $u \in \{0, 1, \dots, c\}^n$ with T non-zero values, and let \tilde{u} denote **Cabin**(u) with \tilde{T} non-zero values. Then, $\mathbb{E}[\tilde{T}] \leq T/2$.*

Proof Let $\text{BinEm}(u)$ be denoted u' and let T' be a random variable denoting the number of ones in u . We proved in Lemma 1 that $\mathbb{E}[T'] = \frac{T}{2}$. Now, $\mathbb{E}[\tilde{T}]$ can be expressed as the number of non-empty bins when T' balls are thrown uniformly into d bins; we obtain $\mathbb{E}[\tilde{T}|T'] = d - d(1 - \frac{1}{d})^{T'}$. Therefore, $\mathbb{E}[\tilde{T}] = \mathbb{E}[\mathbb{E}[\tilde{T}|T']] = d - d \cdot \mathbb{E}[(1 - \frac{1}{d})^{T'}]$. We can bound $\mathbb{E}[(1 - \frac{1}{d})^{T'}] \geq (1 - \frac{1}{d})^{\mathbb{E}[T']}$ using Jensen's inequality, and using Lemma 1 and a standard binomial inequality, prove that $\mathbb{E}[\tilde{T}] \leq d - d \cdot (1 - \frac{1}{d})^{T/2} \leq \frac{T}{2}$.

Table 1 Datasets used for empirical study, ordered according to dimension.

Datasets	Categories	Dimension	Sparsity	Density	Number of points
KOS blog entries [26]	42	6,906	93.38%	457	3,430
NIPS full papers [26]	132	12419	92.64%	914	1,500
Enron Emails [26]	150	28,102	92.81%	2,021	39,861
NYTimes articles[26]	114	102,660	99.15%	871	10,000
PubMed abstracts[26]	47	141,043	99.86%	199	10,000
Million Brain Cells, E18 Mice[42]	2,036	1,306,127	99.92%	1,051	2,000

Computational complexity of Cabin and Cham: The BinSketch and BinEm sub-routines stated in Cabin (Algorithm 1) are one-pass methods. The complexity of BinSketch is linear in the number of data points and the data dimension. The complexity of BinEm is linear in the number of data points and the sketch dimension. Note that the dimensionality of the sketch is $O(s^{3/2})$, which is independent of the original dimension of data, where s is the density of the input. Therefore, the overall running time complexity of Cabin remains linear in the number of data points and the data dimension.

Further, Cham (Algorithm 2) makes one pass over the sketch obtained via Algorithm 1 to estimate a Hamming distance. Therefore the complexity of Cham is linear in the number of data points and the sketch dimension.

5 Experiments

Hardware description: We performed our experiments on a Tyrone DS300TR-34 machine having the following configuration: CPU: Intel(R) Xeon(R) E5-2630 v4 CPU @ 2.20GHz x 20, 32 GB RAM, Windows 10 Professional OS.

Datasets: The efficacy of our proposal is best described for high-dimensional datasets. The categorical datasets that we found to be publicly available were mostly low-dimensional, therefore we considered several integer-valued freely available real-world datasets as categorical. We chose five datasets (see Table 1) for our experiments. The dimensions of the data in these datasets ranged from 6906 to 1.3 million, their sparsity ranged from 92.64% to 99.92%, and the number of categories varied from 42 to 2036. We observed that on certain experiments/datasets, many baselines suffered from “out-of-memory” (OOM) and “did not stop” (DNS) even after running them for a sufficiently long time. Therefore, for those instances we conducted the experiments on a randomly sampled subset of the corresponding dataset. The datasets are of two types.

- **BoW (Bag-of-words)** [26]: We consider the following five datasets – KOS blog entries, NIPS full papers, Enron Emails, NYTimes news articles, and PubMed abstracts – that have “BoW” (bag-of-words) representations of the corresponding text corpora. In all these datasets, the attributes represent the frequency of the words appearing in the documents. Since these frequencies take integer values, we consider them as categories.

- **1.3 Million Brain Cell Dataset** [42]: This dataset consists of the results of single-cell RNA-sequencing (scRNA-seq) of 1.3 million cells captured and sequenced from an E18.5 mouse brain, and made available in public by 10x Genomics[42]

¹. Each gene represents a data point and for every gene, the dataset stores the integer-valued read-count of that gene corresponding to each cell – these read-counts form our features.

Table 2 Baseline algorithms.

Binary Compression Scheme (BCS) [34] *
Hamming LSH (H-LSH) [12] *
Feature Hashing (FH) [41]
Signed-random projection/SimHash (SH)[9]
Kendall rank correlation coefficient (KT)[19]
Latent Semantic Analysis (LSA)[11]
Latent Dirichlet Allocation (LDA)[6]
Multiple Correspondence Analysis (MCA)[5]
Non-neg. Matrix Factorisation (NNMF)[24]
Variational auto-encoder (VAE) [21]
vanilla Principal component analysis (PCA)

* BCS and H-LSH are applied on a BinEm embedding

Baseline algorithms: The alternative approaches that we compare against are listed in Table 2. To the best of our knowledge, there is no off-the-shelf unsupervised dimensionality reduction method available which gives low-dimensional binary embedding of a categorical dataset while approximating Hamming distance. Therefore, we identified some of the state-of-the-art unsupervised dimensionality reduction algorithms for empirically evaluating our solution. Our method is completely unsupervised and doesn’t require labels of the data points for dimensionality reduction. Nevertheless, to evaluate quality of data-analytics tasks, we also included some supervised feature selection methods such as χ^2 [27] and mutual information-based [32] which compute the correlation of features and labels before performing feature selection.

Recall that our method works in a two-step manner: (i) we first compute the binary embedding of the given categorical data points, and, (ii) then we further compress binary vectors obtained in the first step using BinSketch [33]. The BinSketch algorithm due to Pratap *et.al.* [33] is a state-of-the-art algorithm for computing low-dimensional binary vectors for given high-dimensional binary vectors while simultaneously approximating Hamming distance, inner product, Jaccard, and cosine similarity in the same sketch. It is also possible to use Hamming-LSH [12] and BCS [34] for the second step instead of BinSketch, i.e., after generating a binary vector using BinEm. We treat these combinations of BinEm +Hamming-LSH (denoted H-LSH) and BinEm +BCS (denoted BCS) as two other possible techniques.

Reproducibility details: We implemented Feature Hashing (FH) [41], SimHash (SH)[9], BCS [34], Hamming-LSH [12], and BinSketch [33] algorithms on our own. Hamming-LSH is implemented by randomly sampling d features (d denotes the embedding dimension) from each data point, computing the Hamming distance restricted to the sampled features, and then scaling it appropriately for

¹ https://support.10xgenomics.com/single-cell-gene-expression/datasets/1.3.0/1M_neurons

Table 3 Speedup of *Cabin* *w.r.t.* baselines on the reduced dimension 1000.

Dataset	NNMF	MCA	VAE	LDA	LSA	PCA	FH	SH	KT	BCS	H-LSH
Brain Cell	DNS	OOM	OOM	251.76×	52.77×	40.08×	0.803×	3.80×	OOM	0.95×	0.0018×
PubMed	11250.6×	OOM	OOM	131.83×	30.42×	24.82×	0.752×	4.99×	OOM	1.05×	0.0009×
NYTimes	8061×	OOM	OOM	441.44×	65.56×	22.2×	1.864×	3.92×	OOM	1.12×	0.0007×
Enron	9956.8×	392.76×	OOM	195.46×	40.80×	15.67×	1.987×	3.02×	DNS	1.09×	0.0019×
NIPS	3477.9×	54.61×	OOM	358.4×	13.94×	5.84×	2.51×	1.02×	34493×	1.34×	0.0019×
KOS	4255.45×	37.06×	91.2×	282.33×	18.24×	4.72×	2.467×	1.28×	14932×	1.20×	0.0021×

OOM indicates “out-of-memory error” and DNS indicates “did not stop” even after 20 hours.

the full dimension. We made all these implementations freely available². We used `pandas DataFrame`³ for Kendall rank correlation coefficient [19]. For Latent Semantic Analysis (LSA) [11], Latent Dirichlet Allocation (LDA) [6], Non-negative Matrix Factorisation (NNMF) [24], Variational auto-encoder (VAE) [21], and vanilla Principal component analysis (PCA), we used their implementations available from the `sklearn.decomposition` library⁴. For Multiple Correspondence Analysis (MCA) [5], we used an existing Python implementation⁵.

We used `numpy` arrays for storing our vectors. We invoked `numpy.sum(u != v)` for computing the Hamming distance and `numpy.dot(u, v)` for computing the inner product between two data points u and v .

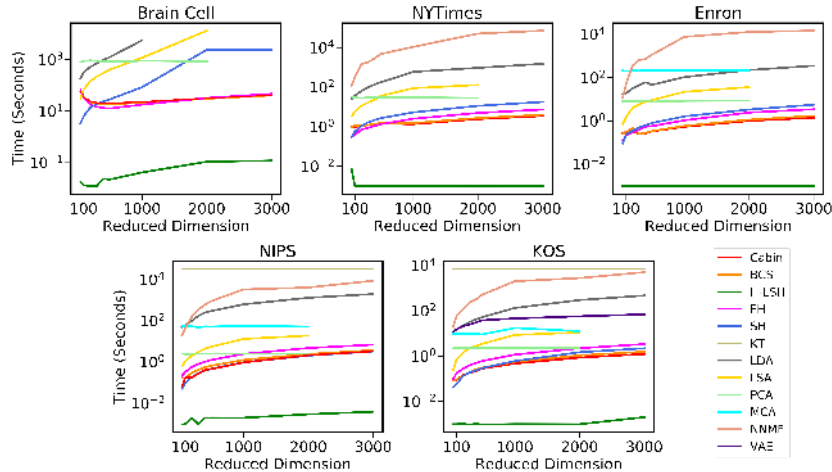


Fig. 2 Comparing the speed of dimensionality reduction. PCA, MCA, and LSA cannot compress beyond the number of data points and the original dimension, leading to missing values beyond a certain point. LDA did not stop after 20 hours for the reduced dimensions more than 1000. Note that many other baseline methods also did not finish or ran out of memory as the dimensions of the datasets increased (also refer to Table 3). Among those which finished successfully, *Cabin* is only slower than Hamming-LSH, but the latter displays a worse performance in the RMSE and the clustering experiments (see Figures 3, 6, 7, and 8).

² https://github.com/Vicky175/Cabin_Cham

³ <https://pandas.pydata.org/pandas-docs/stable/reference/api/pandas.DataFrame.corr.html>

⁴ <https://scikit-learn.org/stable/modules/classes.html#module-sklearn.decomposition>

⁵ <https://pypi.org/project/mca/>

5.1 Speed of dimensionality reduction

A comparison of the time taken for dimensionality reduction on the datasets mentioned above is illustrated in Figure 2 and Table 3. We notice that Hamming LSH is the fastest, however, its performance in the RMSE experiments (Subsection 5.2, Figure 3) and the other tasks (Subsection 5.4, Figure 6) is significantly worse. The speed of `Cabin` is comparable with that of Feature Hashing, SimHash, and BCS, however, the latter too suffer from an inaccurate estimation of Hamming distance leading to fairly poor performance according to the RMSE measure (Subsection 5.2) and the other tasks (Subsection 5.4). We want to draw attention to the fact that many baseline methods such as VAE, MCA, KT give *out-of-memory (OOM)* error or their reported running time is quite high, especially on high values of input or output dimensions. Therefore, we could not perform dimensionality reduction for all dimensions with some of the algorithms.

5.2 Quality of sketches using root mean square error

We evaluate the quality of sketches obtained using different approaches by comparing the error in Hamming distance estimation. For this, we define the Hamming error for a pair of points u, v as

$$HE(u, v) = \text{actual Hamming distance between } u \text{ and } v \\ - \text{estimated Hamming distance obtained from their sketches.}$$

To evaluate the quality of `Cabin` sketches, we compare its root-mean-squared-error (RMSE) with the relevant baseline algorithms. RMSE is defined as $\sqrt{\sum_{u,v} HE(u, v)^2 / N}$ where N represents the total number of pairs. It is a standard metric for comparing dimensionality reduction algorithms and a lower value indicates better performance. For this experiment we compare our solution with BCS, Hamming LSH, Kendall rank correlation coefficient, Feature Hashing, and SimHash. Feature Hashing and SimHash are known to approximate inner product and cosine similarity, respectively, which do not have a direct relation to Hamming distance; however, we include them in the comparison nonetheless since they output discrete sketches, and so, Hamming distance can be defined on them. We did not find it meaningful to compare with the methods that output real-valued sketches.

We performed the RMSE experiment on a sample of 2000 data points from each dataset. Note that this experiment requires generating all pairwise distances, which in our case amounts to $\binom{2000}{2} \approx 1,999,000$ pairs. Therefore including more than 2000 points would be tedious. The results of RMSE comparison are illustrated in Figure 3 where we notice that the RMSE of `Cabin` is the lowest, and rapidly reaches a very low value at reduced dimensions of a few hundred. This further indicates that `Cabin` can compress to much smaller dimensions with a negligible loss in quality compared to the other discrete-valued sketches.

We wish to point out an interesting trend for some of the hashing-based methods, namely, `Cabin`, Feature Hashing (FH), and BCS. They tend to perform better when there are few hash collisions, and this is pronounced when the embedding dimension is large and the input vectors are sparse. Indeed, all three show a remarkable trend of low RMSE as their embedding dimensions are increased for all

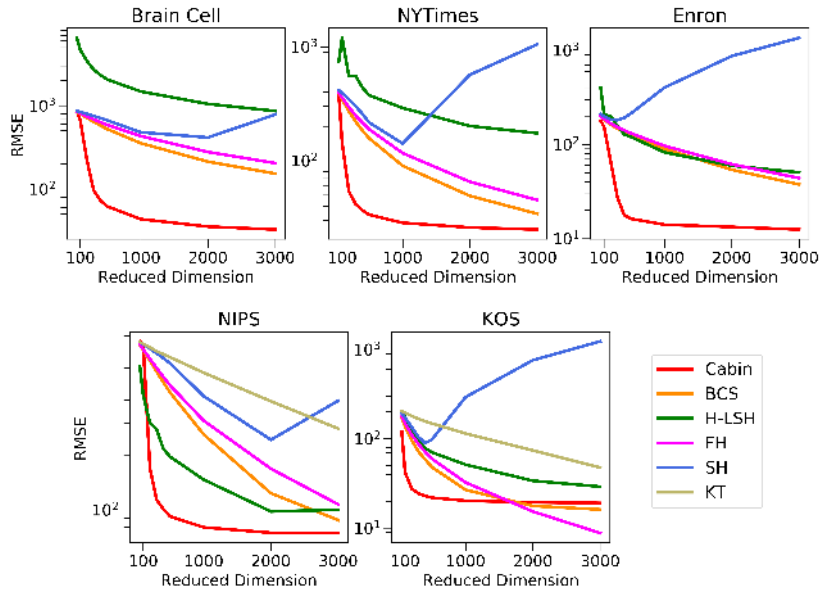


Fig. 3 Comparison on RMSE (root-mean-squared-error) among baselines. A lower value is an indication of better performance. On Enron dataset KT couldn't finish in 10 hours, and on Brain cell and NYTimes datasets, it gave an out-of-memory error.

the datasets. This effect gains prominence for KOS whose dimension is less than 7000 and whose sparsity is about 93.4%. Thus, as the embedding dimension is increased to 2000 (a very high value compared to the original dimension), BCS and FH outperform Cabin; however, the latter still remains the better choice for KOS when the embedding dimension is less than 1000.

5.3 Analysis of Cabin sketches

Recall that `Cabin` generates low-dimensional binary embeddings in a two-step process: (i) `BinEm` first computes the binary embedding of a given categorical vector, and then, (ii) `BinSketch` compresses that binary vector [33]. Intrigued by the low error in the Hamming distances estimated from the `Cabin` sketches, we decided to dive deeper and check the accuracy of both these steps.

5.3.1 Analysis of the first step: BinEm

We conducted two experiments to understand the efficacy of `BinEm`. In the first experiment we chose two random data points (say, u, v), obtained the difference between $HD(u, v)$ and $HD(\text{BinEm}(u), \text{BinEm}(v))$ and generated a box-plot of these errors obtained from 1000 independent trials for the same u, v , by generating random binary embeddings via `BinEm`. The purpose was to understand how much the random embeddings overestimate or underestimate a particular Hamming distance. It is evident from the first two plots of Figure 4 that `BinEm` embeddings




Fig. 4 Variance analysis of **BinEm** embeddings. The pair of box-plots in the first row displays the Hamming errors for a randomly sampled pair of data points, whereas the bottom pair of box-plots show the average of the absolute Hamming errors for all pairs of points.

preserve the actual Hamming distance pretty accurately and appears to be almost uniformly distributed around the actual value.

In the second experiment, we box-plot the average Hamming error

$$\sum_{u,v} |HD(u,v) - HD(\text{BinEm}(u), \text{BinEm}(v))|/N,$$

obtained from 1000 independent runs of **BinEm** (here N denotes the total number of pairs). We added only the absolute errors since the errors could be seen to be both positive and negative. The last two plots of Figure 4 make it clear that **BinEm** embeddings are highly consistent, with very little variance, and with a low average error. These two experiments explain that there is very little loss in the Hamming distance information in the first step of **Cabin**.

5.3.2 Analysis of the second step: BinSketch

Here we conducted an experiment to ask the question: Why BinSketch? Recall that we gave a theoretical justification of this choice in Section 4 that we now wanted to supplement. For the experiment, we chose a random pair of points (say u, v) from the Enron dataset, generated their binary embeddings u', v' via **BinEm**, and then compressed them using BinSketch and other discrete sketching methods – BCS, Hamming-LSH, Feature Hashing, and SimHash, to various reduced dimensions. As before, we computed the Hamming error between the pairs, repeated this process 1000 times, and generated the corresponding box-plots which we present in Figure 5. We observe that the expected value of the Hamming error is close to zero when BinSketch is employed, for all the dimensions. For BCS and Feature Hashing, the expected error is large at low dimensions. The behaviour of Hamming-LSH is close to that of **Cabin** with a slightly worse variance. In fact, the variance of the sketches obtained using BinSketch is the lowest among the alternatives, and moreover, the variance starts decreasing rapidly as the reduced dimension is increased.

So with these experiments, we have empirically validated that the Hamming distances estimates are close to the actual distance on an average, and further, have a low variance. That explains why **Cham** is able to estimate the pairwise Hamming distances from the **Cabin** sketches with high precision.

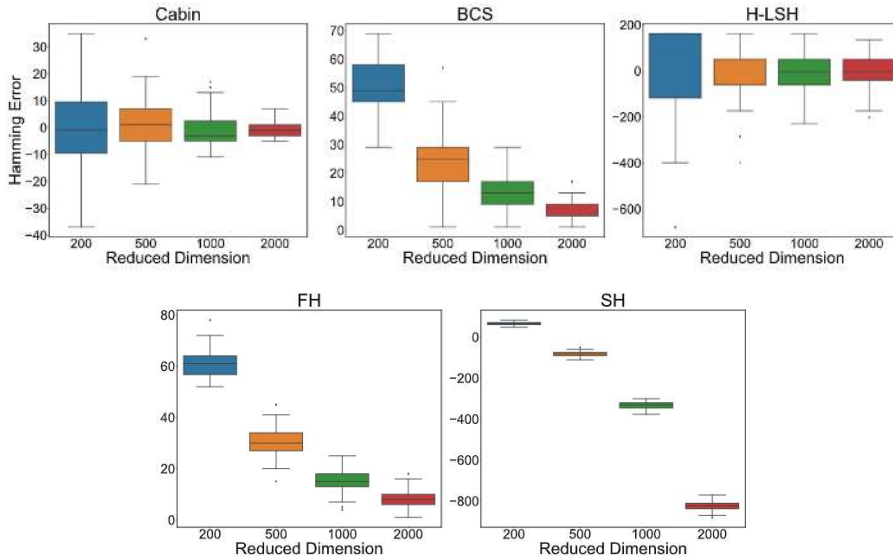


Fig. 5 Analysis of Hamming error in compressing the BinEm embeddings for a random pair from the Enron dataset.

5.4 Performance of clustering

Next, we conducted experiments to understand if the `Cabin` sketches are suitable for data analytic end tasks. The first task we chose was *clustering*, and the aim was to verify if those sketches can compete with the compressed vectors obtained from the standard dimensionality reduction approaches, when used for clustering. For the clustering experiments on NYTimes and PubMed, we used a random sample of 10,000 points since several baseline algorithms started throwing *OOM/DNS* on more points or on the entire dataset. KOS, Enron and NIPS datasets were considered in their entirety. We do not include clustering results on the Brain Cell dataset since the clustering process faced an “out-of-memory” error on the full-dimensional dataset.

We first generated the ground-truth clustering on the original dataset using the classical k -mode algorithm [14] for several values of k . The k -mode algorithm is analogous to k -means but for Hamming distance. We then performed dimensionality reduction using the available techniques for multiple values of reduced dimensions. Note that several baselines such as LDA, LSA, PCA, MCA, and NNMF generate real-valued sketches, therefore, instead of k -mode, we ran k -means (using k -means++ sampling distribution [4]) to generate a clustering. We compare the times taken for clustering and the quality of clustering on the reduced dimensions with the ground truth clustering results. For evaluation the quality of clustering we employed standard metrics such as *purity index*, *adjusted rand index (ARI)*, and *normalised mutual information (NMI)*; we have explained all these metrics in Subsection 3.2.

We used the same random seed for all the baselines to ensure that all of them are initialised with the same set of cluster centres in their first iteration. Thus their

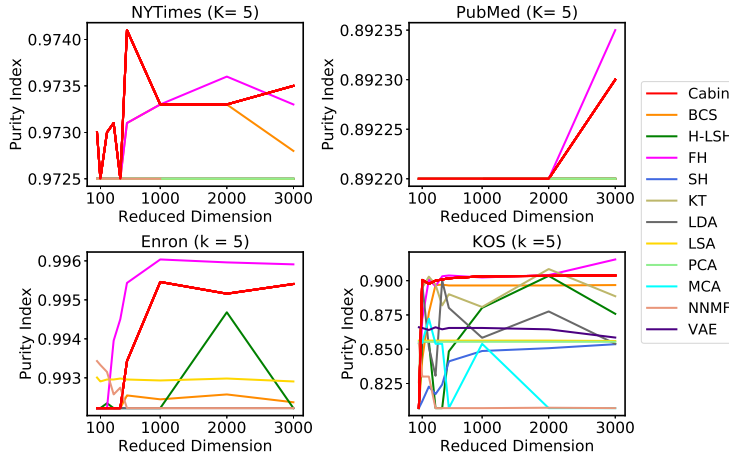


Fig. 6 Comparing the quality of clusters on the compressed datasets using *purity-index* metric. The *purity-index* takes a value between 0 and 1. A higher value indicates a better performance. Observe that the performance of Cabin is among the top few at 1000 or more embedding dimensions.




Fig. 7 Comparing the quality of clusters on the compressed datasets using *Normalised Mutual Information (NMI)* metric. The *NMI* takes a value between 0 and 1. A higher value indicates a better performance. Observe that the performance of Cabin is among the top few at 1000 or more embedding dimensions.

respective final clustering results ought to depend only on the quality of sketches, and not on the randomisation involved in the k initial cluster centres.

Observations: The comparisons of clustering quality on NYTimes, Enron, PubMed, and KOS are presented in Figures 6, 7, and 8 for *purity-index*, *NMI*, and *ARI* evaluation metrics, respectively. The same on the NIPS dataset is presented in Figure 9.

For all three evaluation metrics (*purity-index*, *ARI* and *NMI*), the clustering results achieved using our method is quite high; it consistently remains one of

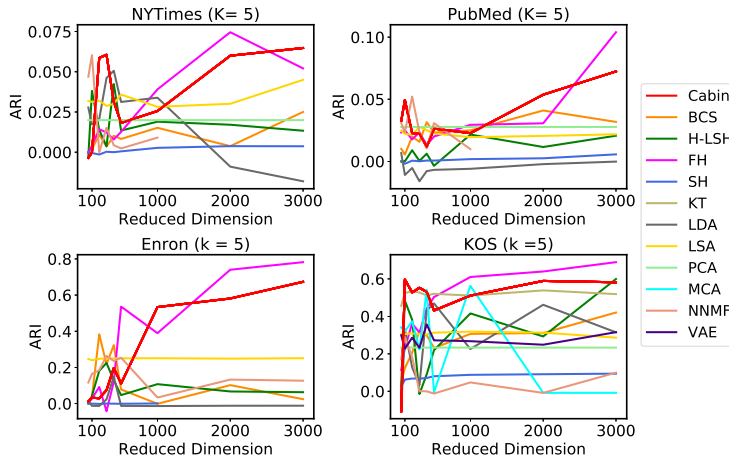


Fig. 8 Comparing the quality of clusters on the compressed datasets using *Adjusted Rand Index (ARI)* metric. The ARI takes a value between -1 and 1 . A higher value indicates a better performance. Observe that the performance of Cabin is among the top few at 1000 or more embedding dimensions.




Fig. 9 Comparing the quality of clusters on the compressed NIPS datasets using *purity-index*, *ARI*, *NMI* evaluation metrics. Observe that the performance of Cabin is among the top two on all metrics beyond 2000 dimensions.

the top two approaches across all the bag-of-words datasets and for all but very small dimensions. It should be noted that, theoretically, the embedding dimension should be above a certain minimum that depends upon the sparsity of a dataset. Nevertheless, we ran the experiments at very low dimensions, and were pleasantly surprised to see good clustering scores, e.g., for the KOS dataset even at a dimension as low as 100.

We explained earlier that succinct binary sketches offer significant advantage during data analytic tasks. To illustrate this advantage, we compare the speedups obtained by clustering 1000-dimensional Cabin sketches *vs.* the uncompressed data points belonging to the different datasets. The speedup is presented in Figure 10 and is significant, e.g., $112.3\times$ on the NYTimes dataset.

5.5 Generation of all-pairs similarity matrix

A common data analysis task is to generate a “all pairs similarity matrix”, *aka.* heatmap (also referred to as pairwise distance/similarity matrix in libraries and

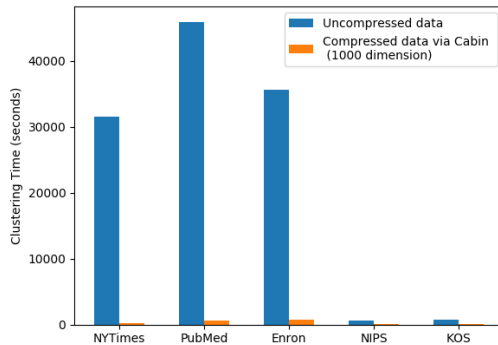
tool-kits) [40]. A heatmap refers to an $N \times N$ matrix, where N denotes the number of data points, whose (i, j) -th entry stores the (dis)similarity between the i -th and the j -th data points. Generating a heatmap of a high-dimensional categorical dataset is tricky since a significant amount of time is spent on computing the pairwise Hamming distances. Our earlier experiments show that **Cham** is able to estimate Hamming distances from low-dimensional binary sketches with high accuracy, so we conducted an experiment to evaluate how well and how fast our method can generate a heatmap.

For this experiment we took 2000 data points from the Brain cell data set and generated the pairwise Hamming distance matrix (left image in Figure 11). Then we reduced the dimension of those points to 1000 using **Cabin** and other baselines that output discrete sketches, namely, BCS, Hamming-LSH, SimHash, Feature Hashing (KT was not included due to OOM error). Heatmaps were generated from the respective sketches (the heatmap from **Cabin** is shown as the right image in Figure 11).

Table 4 Comparing Mean Absolute Hamming error (MAE) in Heat-maps.

Mean abs. Ham. error (MAE)	Cabin	BCS	H-LSH	FH	SH
	23.86	281.19	505.23	351.01	368.32

Observations: A quick glance shows that both the heatmaps in Figure 11 are visually quite similar and it appears that the heatmap from **Cabin** sketches can very well substitute for the latter. The visual similarity may not be very convincing, and further, to compare against the other methods, we generated a heatmap of the Hamming errors for all the appropriate methods; the heatmaps are displayed in Figure 12. We also tabulated the mean (absolute) Hamming errors (MAE) in



Brain Cell	NYTimes	PubMed	Enron	NIPS	KOS
OOM	112.3×	70.7×	45.1×	10.48×	6.5×

Fig. 10 Comparison of the running time of clustering on the uncompressed dataset with the 1000 dimensional sketch obtained via **Cabin**. Clustering on the full-dimension Brain Cell dataset could not be run on the server that we used for experiments.

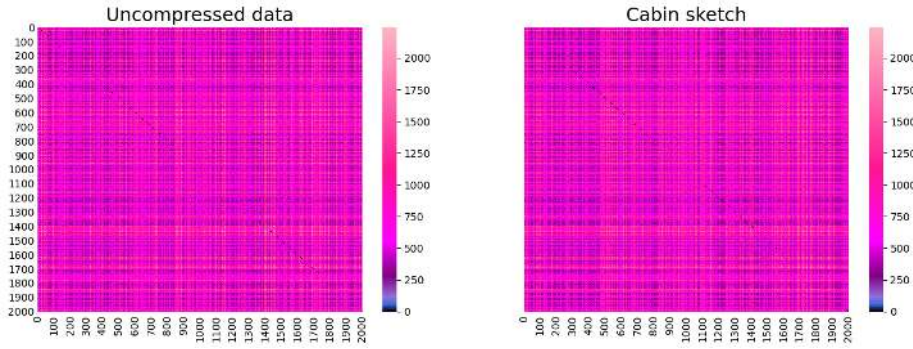


Fig. 11 Heat-maps for pairwise Hamming distance estimation. The left heat-map corresponds to pairwise Hamming distance on the full-dimensional Brain Cell dataset, and the right heat-map corresponds to pairwise Hamming distance estimation from a 1000 dimensional sketch of Brain Cell dataset obtained via *Cabin*.




Fig. 12 Heat-maps on pairwise Hamming error among the baselines. The darker colours correspond to zero error or close to zero error and indicate better performance. Also, see Table 4 in which we have compared the MAE as well.

Table 4. The MAE of our solution is less than 1/10-*th* of those arising out of the baseline approaches and the heatmap of the Hamming errors of *Cabin* sketches is markedly superior compared to the sketches of the other methods. It is clearly evident that the heatmap from *Cabin* sketches is a close approximation to the actual heatmap and the best, by far, among all the other alternatives.

Efficient algorithms for generating an all-pairs similarity matrix is an emerging area [38]. In this context, we found it quite surprising that *Cabin* is able to compress the Brain cell dataset from 1306127 dimensions to 1000 dimensions and is still able to accurately estimate the pairwise Hamming distances. Heatmap generation on the compressed data takes about 570 μ sec compared to about 78 msec on the uncompressed version for each entry of the matrix, leading to roughly 136 \times speedup.

Errors during dimensionality reduction: We noticed that several baseline methods give an *out-of-memory (OOM)* error or their reported running time is quite high, especially on high dimensional datasets. For example, VAE reported *OOM* error on all datasets except KOS, KT threw *OOM* error on NYTimes, PubMed, Brain Cell, and on Enron, it didn't stop even after 20 hrs. MCA also reported *OOM* error for NYTimes, PubMed, and Brain Cell datasets. Further, the dimensionality

reduction time for NNMF was quite high; on NYTimes, it took around 20 hrs to reduce to 3000 dimensions, and on PubMed and Brain cell dataset, it didn't stop even after 20 hrs. We, therefore, could not perform dimensionality reduction with some of the algorithms for all dimensions.

6 Conclusion

In this work, we propose an efficient dimensionality reduction algorithm for sparse categorical data that takes high-dimensional categorical vectors as input and outputs low-dimensional binary vectors. We also present an algorithm to closely approximate the Hamming distance between the original data from the low-dimensional vectors. With the help of theoretical analysis and extensive experiments, we establish that our approach can be used to reduce the dimension of high-dimensional sparse datasets before sending them for data analysis to speed up the tasks without hurting their results.

The novelty of our method is that it strives to be task agnostic. Based on the observations in the RMSE and MAE experiments, we can recommend that a low-dimensional Cabin-compressed dataset can safely be used in the place of a high-dimensional categorical dataset for any task that relies on the Hamming distances between the data points. Our experiments revealed that there could be alternatives that excel at some particular task (e.g., using the Kendall-tau rank correlation compression algorithm for clustering the “NIPS full paper” dataset), but if a go-to approach is required, then there is really no other alternative. To top it off, our solution comes with theoretical guarantees that explain when and why it can be really useful.

There are several interesting directions along which our approach can be further improved. Even though we use completely random maps (π and ψ) during hashing, it may be possible to design better hash functions with some knowledge of the distribution of sparsity beyond a simple upper bound on it that Cabin uses. There is also the theoretical question of using the best-possible and least-resource intensive hash functions, e.g., whether pairwise-uniform hash functions are suitable for Cabin. Even though our proposed solution is entirely unsupervised, we appreciate the benefits of supervised learning and we think that it may be possible to “learn” the right hash functions from training data while retaining the underlying properties of Cabin and Cham.

References

1. Dimitris Achlioptas. Database-friendly random projections. In Peter Buneman, editor, *Proceedings of the Twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 21-23, 2001, Santa Barbara, California, USA*. ACM, 2001.
2. Alekh Agarwal, Oliveier Chapelle, Miroslav Dudík, and John Langford. A reliable effective terascale linear learning system. *Journal of Machine Learning Research*, 15:1111–1133, 2014.
3. Rakesh Agrawal, Tomasz Imielinski, and Arun Swami. Mining association rules between sets of items in large databases. In *SIGMOD '93: Proceedings of the 1993 ACM SIGMOD international conference on Management of data*, pages 207–216, New York, NY, USA, 1993. ACM Press.

4. David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '07*, pages 1027–1035, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.
5. Jörg Blasius and Michael Greenacre. Multiple correspondence analysis and related methods. *Multiple Correspondence Analysis and Related Methods*, 06 2006.
6. David M. Blei, Andrew Y. Ng, Michael I. Jordan, and John Lafferty. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:2003, 2003.
7. Christos Boutsidis, Anastasios Zouzias, and Petros Drineas. Random projections for k -means clustering. *Advances in Neural Information Processing Systems*, 23:298–306, 2010.
8. Andrei Z. Broder, Moses Charikar, Alan M. Frieze, and Michael Mitzenmacher. Min-wise independent permutations (extended abstract). In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 327–336, 1998.
9. Moses Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 380–388, 2002.
10. Graham Cormode, Mayur Datar, Piotr Indyk, and S. Muthukrishnan. Comparing data streams using Hamming norms (how to zero in). *IEEE Trans. Knowl. Data Eng.*, 15(3):529–540, 2003.
11. Scott Deerwester, Susan T. Dumais, George W. Furnas, Thomas K. Landauer, and Richard Harshman. Indexing by latent semantic analysis. *JOURNAL OF THE AMERICAN SOCIETY FOR INFORMATION SCIENCE*, 41(6):391–407, 1990.
12. Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Similarity search in high dimensions via hashing. In *VLDB'99, Proceedings of 25th International Conference on Very Large Data Bases, September 7-10, 1999, Edinburgh, Scotland, UK*, pages 518–529, 1999.
13. Alexey Grigorev. *Mastering Java for data science: building data science applications in Java*. Packt Publishing, 2017.
14. Zhexue Huang. Extensions to the k -means algorithm for clustering large data sets with categorical values. *Data Mining and Knowledge Discovery*, 2(3):283–304, Sep 1998.
15. W. Hämmäläinen and M. Nykänen. Efficient discovery of statistically significant association rules. In *2008 Eighth IEEE International Conference on Data Mining*, pages 203–212, 2008.
16. Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 604–613, 1998.
17. W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Conference in modern analysis and probability (New Haven, Conn., 1982)*, *Amer. Math. Soc., Providence, R.I.*, pages 189–206, 1983.
18. Daniel M. Kane, Jelani Nelson, and David P. Woodruff. An optimal algorithm for the distinct elements problem. In *Proceedings of the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2010, June 6-11, 2010, Indianapolis, Indiana, USA*, pages 41–52, 2010.
19. M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
20. Minje Kim and Paris Smaragdis. Bitwise neural networks for efficient single-channel source separation. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pages 701–705. IEEE, 2018.
21. Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
22. Lukasz Kurgan, Krzysztof Cios, Ryszard Tadeusiewicz, Marek Ogiela, and Lucy Goodenday. Knowledge discovery approach to automated cardiac spect diagnosis. *Artificial intelligence in medicine*, 23:149–69, 11 2001.
23. Jennifer Lavergne, Ryan Benton, and Vijay V. Raghavan. Min-max itemset trees for dense and categorical datasets. In Li Chen, Alexander Felfernig, Jiming Liu, and Zbigniew W. Raś, editors, *Foundations of Intelligent Systems*, pages 51–60, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
24. Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. In Todd K. Leen, Thomas G. Dietterich, and Volker Tresp, editors, *NIPS*, pages 556–562. MIT Press, 2000.

25. Ping Li, Trevor J. Hastie, and Kenneth W. Church. Very sparse random projections. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, page 287–296, New York, NY, USA, 2006. Association for Computing Machinery.
26. M. Lichman. UCI machine learning repository, 2013.
27. Huan Liu and Rudy Setiono. Chi2: feature selection and discretization of numeric attributes. In *Seventh International Conference on Tools with Artificial Intelligence, ICTAI '95, Herndon, VA, USA, November 5-8, 1995*, pages 388–391, 1995.
28. Michael Mitzenmacher, Rasmus Pagh, and Ninh Pham. Efficient estimation for high similarities using odd sketches. In *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, pages 109–118, 2014.
29. J. Moody, D. Touretsky (eds, Morgan Kaufmann, Michiel O. Noordewier, Geoffrey G. Towell, and Jude W. Shavlik. Training knowledge-based neural networks to recognize genes in dna sequences, 1991.
30. Lan Huong Nguyen and Susan Holmes. Ten quick tips for effective dimensionality reduction. *PLOS Computational Biology*, 15(6):1–19, 06 2019.
31. M. M. A. Patwary, S. Byna, N. R. Satish, N. Sundaram, Z. Lukić, V. Roytershteyn, M. J. Anderson, Y. Yao, Prabhat, and P. Dubey. Bd-cats: big data clustering at trillion particle scale. In *SC '15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–12, 2015.
32. Hanchuan Peng, Fuhui Long, and Chris H. Q. Ding. Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(8):1226–1238, 2005.
33. Rameshwar Pratap, Debajyoti Bera, and Karthik Revanuru. Efficient sketching algorithm for sparse binary data. In *2019 IEEE International Conference on Data Mining, ICDM 2019, Beijing, China, November 8-11, 2019*, pages 508–517, 2019.
34. Rameshwar Pratap, Raghav Kulkarni, and Ishan Sohony. Efficient dimensionality reduction for sparse binary data. In *IEEE International Conference on Big Data, Big Data 2018, Seattle, WA, USA, December 10-13, 2018*, pages 152–157, 2018.
35. Rameshwar Pratap, Ishan Sohony, and Raghav Kulkarni. Efficient compression technique for sparse sets. In *Advances in Knowledge Discovery and Data Mining - 22nd Pacific-Asia Conference, PAKDD 2018, Melbourne, VIC, Australia, June 3-6, 2018, Proceedings, Part III*, pages 164–176, 2018.
36. Thorsteinn Rognvaldsson, Liwen You, and Daniel Garwicz. State of the art prediction of hiv-1 protease cleavage sites. *Bioinformatics (Oxford, England)*, 31, 12 2014.
37. Sumit Sidana, Charlotte Laclau, and Massih-Reza Amini. Learning to recommend diverse items over implicit feedback on pandor. pages 427–431, 09 2018.
38. Quico Pepijn Spaen. *Applications and Advances in Similarity-based Machine Learning*. PhD thesis, University of California, Berkeley, 2019.
39. Michael Steinbach, Levent Ertöz, and Vipin Kumar. *The Challenges of Clustering High Dimensional Data*, pages 273–309. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
40. Charlotte Wang, Wen-Hsin Kao, and Chuhsing Kate Hsiao. Using Hamming distance as information for snp-sets clustering and testing in disease association studies. *PLoS one*, 10:e0135918, 08 2015.
41. Kilian Q. Weinberger, Anirban Dasgupta, John Langford, Alexander J. Smola, and Josh Attenberg. Feature hashing for large scale multitask learning. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, pages 1113–1120, 2009.
42. Grace XY Zheng, Jessica M Terry, Phillip Belgrader, Paul Ryvkin, Zachary W Bent, Ryan Wilson, Solongo B Ziraldo, Tobias D Wheeler, Geoff P McDermott, Junjie Zhu, et al. Massively parallel digital transcriptional profiling of single cells. *Nature communications*, 8(1):1–12, 2017. Made available by 10x Genomics at https://support.10xgenomics.com/single-cell-gene-expression/datasets/1.3.0/1M_neurons.