Bound for the L_2 Norm of Random Matrix and Succinct Matrix Approximation

Rong Liu, Nian Yan, Yong Shi, and Zhengxin Chen

Research Center on Fictitious Economy and Data Science,
Chinese Academy of Sciences, Beijing 100080, China
School of Mathematical Science, Graduate University of
Chinese Academy of Sciences, Beijing 100049, China
College of Information Science and Technology,
University of Nebraska at Omaha, Omaha NE 68182, USA
liu.rong@163.com, yshi@gucas.ac.cn, {nyan,zchen}@mail.unomaha.edu

Abstract. This work furnished a sharper bound of exponential form for the L_2 norm of an arbitrary shaped random matrix. Based on the newly elaborated bound, a non-uniform sampling method was developed to succinctly approximate a matrix with a sparse binary one and hereby to relieve the computation loads in both time and storage. This method is not only pass-efficient but query-efficient also since the whole process can be completed in one pass over the input matrix and the sampling and quantizing are naturally combined in a single step.

Keywords: Sampling, Matrix Approximation, Data Reduction, Random Matrix, L_2 Norm.

1 Introduction

The computation for matrix by vector product is the most time and storage consuming step in many numeric algorithms. Typical complexity of this computation is $O(mn)$ in both time and storage for a matrix size of $m \times n$. In real world problems such as data mining, machine learning, computer vision, and information retrieval, computations soon become infeasible as the size of matrices gets large since there are so many operations of this type involved. In data mining and machine learning, low rank approximations provide compact representations of the data with limited loss of information and hereby relieve the curse of dimensionality. A well known technique for low rank approximations is the Singular Value Decomposition (SVD), also called Latent Semantic Indexing (LSI) in information retrieval [7] [4], which minimize the residual among all approximations with the same rank. Traditional SVD computation requires $O(mn\min\{m,n\})$ time and $O(mn)$ space. To speed up the SVD-based low rank approximations and save storage, sampling methods are often used to sparsifying the input matrices because the time and space complexity for a sparsified matrix of N non-zero entries are reduced to $O(N\min\{m,n\})$ and $O(N)$ respectively. If a quantizing (rounding) approach is combined, the computation loads can be relieved further.

M. Bubak et al. (Eds.): ICCS 2008, Part II, LNCS 5102, pp. 426–435, 2008.

Attracted by the computational benefits from sparsifying matrices, numerous algorithms for massive data problems adopt sampling methods and different error bounds are presented along with them. Frieze et al. (2004) [12] developed a fast Monte-Carlo algorithm which find an approximation $\widetilde{\boldsymbol{D}}$ in the vector space spanned by some sampled rows of the input matrix $\boldsymbol{D}$ such that $P(\|\boldsymbol{D} - \widetilde{\boldsymbol{D}}\|_F \leq \epsilon\|\boldsymbol{D}\|_F) \geq 1 - \delta$ in two passes through $\boldsymbol{D}$. Since many rows are discarded in the random sampling process, the output matrix can be interpreted as a spare representation of the input one. In low rank approximations, Deshpande and Vempala (2006) [8] further presented an adaptive sampling method to approximate the volume sampling with a probability at least $3/4$ that $\|\boldsymbol{D} - \widetilde{\boldsymbol{D}}_k\|_F \leq (1+\epsilon)\|\boldsymbol{D} - \boldsymbol{D}_k\|_F$ ($\boldsymbol{D}_k$ consists of the first k terms in the SVD of $\boldsymbol{D}$). This algorithm requires more passes over the input matrix. Passes through a large matrix can be extremely costive if the matrix can not be fed in Random Access Memory and must be stored in external storages such as Hard Disk. A alternative entrywise sampling method which uses only one pass was given by Achlioptas and McSherry (2001) [1]. Where the reconstruction error has a significant better bound in L_2 norm, a more effective indicator for the trends in a matrix than Frobenius norm, as $P(\|\boldsymbol{D} - \widetilde{\boldsymbol{D}}\|_2 \geq 5b\sqrt{s(m+n)}) \leq 1/(m+n)$ with certain regularity. The theorem [2] which bounds the deviations of eigenvalues of a random symmetric matrix from their medians is used to prove this bound. In [1] they also includes detailed comparisons of their application to fast SVD with the methods in [9] and [11].

There are also some achievements in non-sampling based methods. GLRAM proposed by Ye (2005) [16] is a such one. Instead of the linear transformation used in SVD, the GLRAM applies a bilinear mapping on the data because each point is represented by a matrix other than a vector. The algorithm intends to minimize $\sum_{i=1}^{n} \|\boldsymbol{A}_i - \boldsymbol{L}\boldsymbol{M}_i\boldsymbol{R}^T\|_F^2$ with the same $\boldsymbol{L}$ and $\boldsymbol{R}$ for all $\boldsymbol{A}_i$ (the data points) and hereby relieve computation loads. Their result is appealing but does not admit a closed form solution in general, hence no error bound presented.

From another point of view, Bar-Yossef (2003) [5] pointed out based on information theory that at least $O(mn)$ queries are necessary to produce a (ϵ, δ) approximation for a given matrix.

The review paper by Mannila (2002) [14] mentioned that methods of [1] worked very well in low rank approximation. However, [14] also suggested the necessity of further research for their boundaries. Besides, the following drawbacks are found in these methods under our own investigation. Firstly, the prior (uniform sampling) in [1] is flat which incorporates no knowledge about the data. Secondly, a zero entry can even be round to non-zero. Thirdly, the rounding may change the sign of an entry. Among them, the last two are extremely harmful to data mining tasks. The main weakness of the method in [3] is that matrices are required to be symmetric. In addition, this method still demands large amounts of computation and storage.

As a continuing study for data reduction, the objective of this paper is to develop an entrywise non-uniform sampling based data reduction (matrix approximation) method with lower reconstruction error.

2 Method and Analysis

The data matrix under consideration in this paper is $\boldsymbol{D} \in \mathbb{R}^{m \times n}$. Here, m is the number of data points, and n the number of attributes or the dimension of the points. The method sparsifying $\boldsymbol{D}$ is delivered and analyzed in the following.

2.1 Sampling Method and Its Bound

The sampling method we present to approximate a matrix is

$$\tilde{d}_{ij} = \begin{cases} 0 & w.p.\ 1 - |d_{ij}|/c \\ c \cdot \operatorname{sgn}(d_{ij}) & w.p.\ |d_{ij}|/c \end{cases} \tag{1}$$

Here, $w.p.$ means "with the probability of". $c = b \cdot s$, $b = \max_{i,j} |d_{ij}|$, and $s \geq 1$ is a controller for the sparsity. Actually, this method directly omit an entry or round it to $\pm c$ in one query unlike the method in [1] where an extra quantizing step is needed. Moreover, the matrix obtained can be succinctly represented by a binary matrix corresponding to the sign of the entries, enabling addition in place of multiplication, with a final scaling of the result by c. One of the most important capability indicators is the residual or the so called reconstruction error. There different metrics for the residual. Among them, matrix norm is a widely used one. For the insensitivity of F norm to the matrix structure, L_2 norm is a preferred metric in classification, LSI, and other data mining tasks [1]. Therefore, the error of the sampling method are estimated in L_2 norm as the theorem follows.

Theorem 1. *If $\widetilde{\boldsymbol{D}}$ is the matrix obtained from the sampling process in Equation (1), then $P(\|\boldsymbol{D} - \widetilde{\boldsymbol{D}}\|_2 \leq \epsilon) \geq 1 - e^{-2[(1-t)^4\epsilon^2/c^2 - (2+1/t)(m+n)]}$, for all $\epsilon > 0$ and some fixed $t \in (0, 1)$.*

The proof of Theorem 1. are given below along with the motivation of this method.

2.2 Motivation and Analysis

Computing the L_2 norm of an arbitrary matrix is a quite difficult task. The already known results to this problem are mainly involved in symmetric matrices. For an arbitrary matrix $\boldsymbol{B} \in \mathbb{R}^{m \times n}$, the following Lemma can be used to associate $\|\boldsymbol{B}\|_2$ with the L_2 norm of a symmetric matrix

$$\boldsymbol{A} = \begin{bmatrix} \mathbf{0} & \boldsymbol{B}^T \\ \boldsymbol{B} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{(m+n)\times(m+n)} \tag{2}$$

Lemma 1. *Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{m+n}$ be eigenvalues of $\boldsymbol{A}$, then $\|\boldsymbol{B}\|_2 = \|\boldsymbol{A}\|_2 = \lambda_1 = -\lambda_{m+n}$.*

Proof: Let $\boldsymbol{x}$ be the eigenvector belongs to the eigenvalue λ. We decompose the vector into two part as $\boldsymbol{x} = [\boldsymbol{x}_1^T, \boldsymbol{x}_2^T]^T$ according to the structure of $\boldsymbol{A}$. Then $\boldsymbol{Ax} = \lambda \boldsymbol{x}$ can be written as

$$\begin{bmatrix} \mathbf{0} & \boldsymbol{B}^T \\ \boldsymbol{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \end{bmatrix} = \lambda \begin{bmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \end{bmatrix}$$

Let $\boldsymbol{y}_1 = \boldsymbol{x}_1, \boldsymbol{y}_2 = -\boldsymbol{x}_2$, this becomes

$$\begin{bmatrix} \mathbf{0} & \boldsymbol{B}^T \\ \boldsymbol{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{y}_1 \\ \boldsymbol{y}_2 \end{bmatrix} = -\lambda \begin{bmatrix} \boldsymbol{y}_1 \\ \boldsymbol{y}_2 \end{bmatrix} \quad i.e. \quad \boldsymbol{Ay} = -\lambda \boldsymbol{y}$$

so $-\lambda$ is also an eigenvalue of $\boldsymbol{A}$. This means that the eigenvalues of $\boldsymbol{A}$ appear in a pairwise way. Consequently, the L_2 norm of the symmetric matrix $\boldsymbol{A}$ is given by $\|\boldsymbol{A}\|_2 = \max_i |\lambda_i| = \lambda_1 = -\lambda_{m+n}$. To prove that the L_2 norms of $\boldsymbol{B}$ and $\boldsymbol{A}$ are identical, let us consider the eigenfunction of $\boldsymbol{A}^T\boldsymbol{A}$

$$0 = |\boldsymbol{A}^T\boldsymbol{A} - \lambda \boldsymbol{I}| = \begin{vmatrix} \boldsymbol{B}^T\boldsymbol{B} - \lambda \boldsymbol{I} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{B}\boldsymbol{B}^T - \lambda \boldsymbol{I} \end{vmatrix} = |\boldsymbol{B}^T\boldsymbol{B} - \lambda \boldsymbol{I}| \, |\boldsymbol{B}\boldsymbol{B}^T - \lambda \boldsymbol{I}|$$

According to the Sylvester's theorem [6] $\boldsymbol{B}^T\boldsymbol{B}$ and $\boldsymbol{B}\boldsymbol{B}^T$ have the same none-zero eigenvalues. The equation above further reveals the truth that $\boldsymbol{A}^T\boldsymbol{A}$, $\boldsymbol{B}^T\boldsymbol{B}$, and $\boldsymbol{B}\boldsymbol{B}^T$ are all share the same non-zero eigenvalues. This lead us to the claim that $\|\boldsymbol{B}\|_2 = \|\boldsymbol{A}\|_2$. □

Now we focus on the L_2 norm of a symmetric matrix since the Lemma given above points out that one can always construct a symmetric matrix to make it has the same L_2 norm as the given matrix of arbitrary shape. Symmetric matrices possess many appealing properties and some of them are very helpful in estimating the L_2 norm. The Rayleigh quotient defined below is especially important among them owing to providing the information about the eigenvalues of a symmetric matrix.

$$R(\boldsymbol{x}) = \frac{\boldsymbol{x}^T\boldsymbol{Ax}}{\boldsymbol{x}^T\boldsymbol{x}}, \boldsymbol{x} \neq \mathbf{0}$$

The next Lemma is an improvement for the Theorem in [6] which associates the L_2 norm of a symmetric matrix with the Rayleigh quotient.

Lemma 2. *If $\boldsymbol{A}$ is a symmetric matrix, then $\|\boldsymbol{A}\|_2 = \max_{\boldsymbol{x}} |R(\boldsymbol{x})|$; furthermore, if the matrix takes the form given in Equation (2), then $\|\boldsymbol{A}\|_2 = \lambda_1 = \max_{\boldsymbol{x}} R(\boldsymbol{x})$.*

Proof: The best known property of Rayleigh quotient (cited in [6]) is

$$\max_{\boldsymbol{x} \neq \mathbf{0}} R(\boldsymbol{x}) = \lambda_1, \quad \min_{\boldsymbol{x} \neq \mathbf{0}} R(\boldsymbol{x}) = \lambda_n$$

Here, $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ are eigenvalues of $\boldsymbol{A}$. This property gives the first part of the Lemma straightforwardly since $\|\boldsymbol{A}\|_2 = \max\{|\lambda_1|, |\lambda_n|\} = \max_{\boldsymbol{x}} |R(\boldsymbol{x})|$.

According to Lemma 1., $\|\boldsymbol{A}\|_2 = \lambda_1 = \max_{\boldsymbol{x}} R(\boldsymbol{x})$ if $\boldsymbol{A}$ takes the form in Equation (2). □

The Rayleigh quotient can be defined equivalently as $R(\boldsymbol{x}) = \sum_{i,j}^{n} A_{ij}x_i x_j$ on a unit vector $\|\boldsymbol{x}\|_2 = 1$.

$R(\boldsymbol{x})$ is a random variable, if $\boldsymbol{A}$ is a random symmetric matrix. Inequalities are of great use in bounding quantities that might otherwise be hard to compute. But directly using Rayleigh quotient combined with a random inequality to bound the L_2 norm of $\boldsymbol{A}$ will encounter great difficulty because $P(\|\boldsymbol{A}\|_2 \geq \epsilon) \geq P(|\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}| > \epsilon)$ (recall that $\|\boldsymbol{A}\|_2 = \max_{\|\boldsymbol{x}\|_2=1} |\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}|$). Fortunately, [10] introduced a method that makes it sufficient to consider only vectors in a discrete space. We strengthen their result ([10], [3]) here as a Lemma that

Lemma 3. *(Reduction to Discrete Space) (Reduction to Discrete Space) Let* $S = \{\boldsymbol{z} : \boldsymbol{z} \in \frac{t}{\sqrt{n}}\mathbb{Z}^n, \|\boldsymbol{z}\|_2 \leq 1\}$ *for some fixed* $t \in (0,1)$*. If for every* $\boldsymbol{v}_i \in S$ *we have* $\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \leq \epsilon$*, then* $\boldsymbol{u}^T\boldsymbol{A}\boldsymbol{u} \leq \epsilon/(1-t)^2$ *for every unit vector* $\boldsymbol{u}$*. And the size of* S *(denoted by* $|S|$*) is at most* $e^{(2+1/t)n}$*.*

Proof: Let $\boldsymbol{y}$ be a vector of length at most $(1-t)$ and C be the specific hypercube from the grid $\frac{t}{\sqrt{n}}\mathbb{Z}^n$ in which $\boldsymbol{y}$ lies. Any two points inside C (include $\boldsymbol{y}$) are t close since the side length of C is $\frac{t}{\sqrt{n}}$. As the maximum length of $\boldsymbol{y}$ is assumed to be $(1-t)$, all vertices of C are within the distance of 1 from the origin. Therefore, all vertices of C belong to set S. The convex combination representability of $\boldsymbol{y}$ in the vertices of C can then be extended to S. Namely, $\boldsymbol{y} = \sum_i a_i\boldsymbol{v}_i$ ($a_i \geq 0$, $\sum_i a_i = 1$). In this way,

$$\begin{aligned}\boldsymbol{y}^T\boldsymbol{A}\boldsymbol{y} &= \sum_i a_i\boldsymbol{v}_i^t\boldsymbol{A}\sum_i a_j\boldsymbol{v}_j = \sum_i\sum_j a_i a_j \boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \\ &\leq \sum_i\sum_j a_i a_j \epsilon = \epsilon\end{aligned}$$

The inequality is obtained as we assumed that $\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \leq \epsilon$ for every $\boldsymbol{v}_i \in S$. Resultingly,

$$\boldsymbol{u}^T\boldsymbol{A}\boldsymbol{u} = \frac{\boldsymbol{y}^T\boldsymbol{A}\boldsymbol{y}}{(1-t)^2} \leq \frac{\epsilon}{(1-t)^2}$$

for any unit vector $\boldsymbol{u}$.

Map every point $\boldsymbol{z} \in S$ in a 1-1 correspondence with a n-dimensional hypercube of side length $\frac{t}{\sqrt{n}}$ centered on itself as $\boldsymbol{z} \mapsto C_{\boldsymbol{z}}(= \{\boldsymbol{z} \pm \boldsymbol{w} : \|\boldsymbol{w}\|_\infty \leq \frac{t}{2\sqrt{n}}\})$. The length of $\boldsymbol{w}$ satisfies the following inequality.

$$\|\boldsymbol{w}\|_2^2 = \sum_{i=1}^{n} w_i^2 \leq \sum_{i=1}^{n} \frac{t^2}{4n} = \frac{t^2}{4}$$

Then,

$$\|\boldsymbol{z}+\boldsymbol{w}\|_2 \leq \|\boldsymbol{z}\|_2 + \|\boldsymbol{w}\|_2 \leq 1 + t/2$$

The inequality above indicates the length of any vector in $C_{\boldsymbol{z}}$ is bounded by $1+t/2$. The union of these cubes is thus contained in a n-dimensional ball B of radius $(1+t/2)$. According to the volume relationship between the union and the ball,

$$|S| \cdot \left(\frac{t}{\sqrt{n}}\right)^n = \sum_{\boldsymbol{z} \in C_{\boldsymbol{z}}} \mathbf{Vol}(C_{\boldsymbol{z}}) \leq \mathbf{Vol}(B) = \frac{\pi^{n/2}}{\Gamma(n/2+1)}(1+t/2)^n$$

Consequently,

$$|S| \leq \frac{\pi^{n/2}}{\Gamma(n/2+1)} \left(\frac{(1+t/2)\sqrt{n}}{t}\right)^n \leq e^{(2+1/t)n}$$

That's all for the proof. $\square$

In respect that the probability relationship between $P(\|\boldsymbol{A}\|_2 \leq \epsilon)$, $P(\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} \leq \epsilon \mid \|\boldsymbol{x}\|_2 = 1)$, and $P(\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S)$ is rather involving, we give an additional Lemma to describe them.

Lemma 4. *If $\boldsymbol{A}$ takes the form in Equation (2) and $P(\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S) \leq \xi$, then $P(\|\boldsymbol{A}\|_2 \leq \epsilon) \geq 1 - \xi e^{2(2+1/t)(m+n)}$.*

Proof: The logic relationship of the propositions above is

$$\forall \boldsymbol{v}_i \in S, \boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \leq (1-t)^2\epsilon \Rightarrow \forall \|\boldsymbol{x}\|_2 = 1, \boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} \leq \epsilon \Rightarrow \|\boldsymbol{A}\|_2 \leq \epsilon$$

As a result,

$$\begin{aligned}
P(\|\boldsymbol{A}\|_2 \leq \epsilon) &\geq P(\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} \leq \epsilon \mid \|\boldsymbol{x}\|_2 = 1) \\
&\geq P(\bigcap_{i=1}^{|S|}\bigcap_{j=1}^{|S|} \boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \leq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S)
\end{aligned}$$

Note that

$$\begin{aligned}
&P(\bigcap_{i=1}^{|S|}\bigcap_{j=1}^{|S|} \boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \leq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S) \\
&= 1 - P(\bigcup_{i=1}^{|S|}\bigcup_{j=1}^{|S|} \boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S) \\
&\geq 1 - \sum_{i=1}^{|S|}\sum_{j=1}^{|S|} P(\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S) \\
&\geq 1 - \sum_{i=1}^{|S|}\sum_{j=1}^{|S|} \xi = 1 - |S|^2\xi \geq 1 - \xi e^{2(2+1/t)(m+n)}
\end{aligned}$$

The last inequality results from that Lemma 3. points out the maximum size of T is $e^{(2+1/t)(m+n)}$ (note that the dimension of $\boldsymbol{A}$ is $(m+n)$). Then the lemma follows. $\square$

Now we focus on furnishing $P(\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq \epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S)$ with an upper bound via probability inequalities. Markov's inequality obtained by tailing the expectation of a random variable is a widely used one in real world applications. Hoeffding's inequality is similar in spirit to Markov's inequality, but it is a sharper inequality owing to that it makes use of a Taylor expansion of second order. Furthermore, if there are large number of samples, the bound from Hoeffind's inequality is smaller than the bound from Chebyshev's inequality. Before applying Hoeffding's inequality to get a bound for the sampling method, we give a refinement for it as:

Lemma 5. *Let $Y_1, \cdots, Y_n$ be independent observation such that $E(Y_i) = 0$ and $\alpha_i \leq Y_i \leq \beta_i$. Then, for any $\epsilon > 0$,*

$$P\left(\sum_{i=1}^{n} Y_i \geq \epsilon\right) \leq e^{-2\epsilon^2/\sum_{i=1}^{n}(\beta_i-\alpha_i)^2}$$

Proof: The original Hoeffding's inequality given in [15] on the same condition is

$$P\left(\sum_{i=1}^{n} Y_i \geq \epsilon\right) \leq e^{-t\epsilon}\prod_{i=1}^{n} e^{t^2(\beta_i-\alpha_i)^2/8},\ \forall t > 0$$

When $t = 4\epsilon/\sum_{i=1}^{n}(\beta_i-\alpha_i)^2$, the right hand side of the inequality above archives its minmum $e^{-2\epsilon^2/\sum_{i=1}^{n}(\beta_i-\alpha_i)^2}$. Thus the claim. $\square$

A bound for $P(\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq \epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S)$ can then be obtain by applying the results obtained above. We state it as one of our main results as the Theorem follows.

Theorem 2. *Let $\boldsymbol{B} \in \mathbb{R}^{m\times n}$ be a random matrix with independent entries of zero mean. If $\alpha_{ij} \leq b_{ij} \leq \beta_{ij}$, then $P(\|\boldsymbol{B}\|_2 \leq \epsilon) \geq 1-\delta$ for $\forall\epsilon > 0$. Where $\delta = e^{-2[(1-t)^4\epsilon^2/\sum_{i=1}^{n}\sum_{j=1}^{m}(x_iw_j+u_iy_j)^2(\beta_{ij}-\alpha_{ij})^2-(2+1/t)(m+n)]}$ for some fixed $t \in (0,1)$, $[\boldsymbol{x}^T, \boldsymbol{y}^T]^T = \boldsymbol{v}_i \in S$, and $[\boldsymbol{u}^T, \boldsymbol{w}^T]^T = \boldsymbol{v}_j \in S$.*

Proof: We construct a symmetric matrix $\boldsymbol{A}$ by $\boldsymbol{B}$ as Equation (2), then $\|\boldsymbol{B}\|_2 = \|\boldsymbol{A}\|_2$ according to Lemma 1. and Lemma 2.. Lemma 4. reveals further that $P(\|\boldsymbol{B}\|_2 \leq \epsilon) = P(\|\boldsymbol{A}\|_2 \leq \epsilon) \geq 1 - \xi e^{2(2+1/t)(m+n)}$ if $P(\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j \geq (1-t)^2\epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S) \leq \xi$. For the matrix $\boldsymbol{A}$,

$$\begin{aligned}\boldsymbol{v}_i^T\boldsymbol{A}\boldsymbol{v}_j &= \begin{bmatrix}\boldsymbol{x}^T & \boldsymbol{y}^T\end{bmatrix}\begin{bmatrix}\boldsymbol{0} & \boldsymbol{B}^T\\ \boldsymbol{B} & \boldsymbol{0}\end{bmatrix}\begin{bmatrix}\boldsymbol{u}\\ \boldsymbol{w}\end{bmatrix}\\ &= \boldsymbol{y}^T\boldsymbol{B}\boldsymbol{u} + \boldsymbol{x}^T\boldsymbol{B}^T\boldsymbol{w} = \sum_{i=1}^{n}\sum_{j=1}^{m}(x_iw_j + u_iy_j)b_{ij}\end{aligned}$$

Since b_{ij} are independent, $(x_i w_j + u_i y_j)b_{ij}$ are also independent and zero mean. The requirements for Lemma 5. are thus satisfied and the range of random variable $(x_i w_j + u_i y_j)b_{ij}$ is from $(x_i w_j + u_i y_j)\alpha_{ij}$ to $(x_i w_j + u_i y_j)\beta_{ij}$ or conversely. As a result,

$$\begin{aligned}&P(\boldsymbol{v}_i^T \boldsymbol{A} \boldsymbol{v}_j \geq (1-t)^2 \epsilon \mid \boldsymbol{v}_i, \boldsymbol{v}_j \in S) \\ &\quad \leq e^{-2(1-t)^4 \epsilon^2 / \sum_{i=1}^{n} \sum_{j=1}^{m} (x_i w_j + u_i y_j)^2 (\beta_{ij} - \alpha_{ij})^2} = \xi\end{aligned}$$

Then the theorem follows. □

At this time we are prepared to prove Theorem 1.. But before going further, we present the motivation of the sampling method.

A general approach for sparsifying a dense matrix is

$$\tilde{d}_{ij} = \begin{cases} 0 & w.p.\ 1 - p_{ij} \\ \gamma_{ij} & w.p.\ p_{ij} \end{cases}$$

In many cases, the method that makes the expectation of the sampled matrix equal to the input one is a reasonable strategy. We follow this too, that is $d_{ij} = E(\tilde{d}_{ij}) = 0 \cdot (1 - p_{ij}) + \gamma_{ij} \cdot p_{ij} = \gamma_{ij} \cdot p_{ij}$. So our method is

$$\tilde{d}_{ij} = \begin{cases} 0 & w.p.\ 1 - p_{ij} \\ d_{ij}/p_{ij} & w.p.\ p_{ij} \end{cases}$$

Define $\boldsymbol{B} = \boldsymbol{D} - \widetilde{\boldsymbol{D}}$ as the error or difference matrix of the sampling. Obviously, $E(b_{ij}) = 0$ and b_{ij} are independent due to that $\tilde{d}_{ij}$ are sampled independently. The requirements for Theorem 2. are hereby satisfied. The range for the random variable $\tilde{b}_{ij}$ is from d_{ij} to $(1-1/p_{ij})d_{ij}$ or the converse. According to Theorem 2., the components that contain the ranges are then become as $\sum_{i=1}^{n} \sum_{j=1}^{m} (x_i w_j + u_i y_j)^2 (d_{ij}/p_{ij})^2$. Now the last free parameter need to be settled is the probability p_{ij} which is essential for the success of many sampling based methods. From the analysis above, we know that with the increase of p_{ij} the δ in the bound of $\|\boldsymbol{B}\|_2$ decreases but the density of the matrix $\widetilde{\boldsymbol{D}}$ increases. In applications, compromise between them should be take into consideration. Note that there is no priori reason to omit each entry in the input matrix $\boldsymbol{D}$ with the same probability, we set $p_{ij} = |d_{ij}|/c$ (c is a constant positive real number) so that all entries of $\boldsymbol{B}$ are contained in segments of equal length c and the bound for $\|\boldsymbol{B}\|_2$ are simplified further. In addition, that the larger the absolute value of an entry is, the more likely it retained is desired in many real world applcations. Intend to guarantee that p_{ij} is a well defined probability that contained in the segment of $[0, 1]$, the scale factor c is set to be $s \cdot b$. As a result, sparsifying (omitting) and quantizing (rounding) are combined naturally without any extra steps. The sampling method given in Equation (1) comes into being as thus.

After explaining the motivation, Theorem 1. can be proofed straightly.

Proof of Theorem 1.: For $p_{ij} = |d_{ij}|/c$, $\sum_{i=1}^{n}\sum_{j=1}^{m} (x_i w_j + u_i y_j)^2(\beta_{ij} - \alpha_{ij})^2 = c^2 \sum_{i=1}^{n}\sum_{j=1}^{m}(x_i w_j + u_i y_j)^2$. And we have

$$
\begin{aligned}
&\sum_{i=1}^{n}\sum_{j=1}^{m}(x_i w_j + u_i y_j)^2 \\
&\quad= \sum_{i=1}^{n}\sum_{j=1}^{m}(x_i^2 w_j^2 + u_i^2 y_j^2) + 2\sum_{i=1}^{n} x_i u_i \sum_{j=1}^{m} w_j y_j \\
&\quad\leq \sum_{i=1}^{n}\sum_{j=1}^{m}(x_i^2 w_j^2 + u_i^2 y_j^2) + (\sum_{i=1}^{n} x_i u_i)^2 + (\sum_{j=1}^{m} w_j y_j)^2 \\
&\quad\leq \sum_{i=1}^{n}\sum_{j=1}^{m}(x_i^2 w_j^2 + u_i^2 y_j^2) + \sum_{i=1}^{n} x_i^2 \sum_{i=1}^{n} u_i^2 + \sum_{j=1}^{m} w_j^2 \sum_{j=1}^{m} y_j^2 \\
&\quad= (\sum_{i=1}^{n} x_i^2 + \sum_{j=1}^{m} y_j^2)(\sum_{i=1}^{n} u_i^2 + \sum_{j=1}^{m} w_j^2) \leq 1
\end{aligned}
$$

The second inequality is obtained from Cauchy's inequality, and the last is from the fact that $[\boldsymbol{x}^T, \boldsymbol{y}^T]^T = \boldsymbol{v}_i \in S, [\boldsymbol{u}^T, \boldsymbol{w}^T]^T = \boldsymbol{v}_j \in S$. Accordingly, $\delta \leq e^{-2[(1-t)^4\epsilon^2/c^2-(2+1/t)(m+n)]}$. Applying Theorem 2., $P(\|\boldsymbol{B}\|_2 \leq \epsilon) \geq 1 - e^{-2[(1-t)^4\epsilon^2/c^2-(2+1/t)(m+n)]}$ comes forth. □

Our method is hereby much sharper in error bound with larger compression ratio than that of [1]. The bound from [3] is restricted to symmetric data matrices, moreover, their δ is two times worse than ours.

3 Discussion

In this paper, we furnish a bound of exponential form for the L_2 norm of a random matrix of arbitrary shape. The method we use differs from the one that bounds the deviations of eigenvalues of a random symmetric matrix from their medians and thus gives a bound of polynomial form. Through the Rayleigh quotient of the matrix, our work associates the L_2 norm of an input matrix with the inner product weighted by a constructed symmetric matrix in a reduced discrete space. As a result, an exponential bound for the input matrix is obtained by applying Hoeffding's inequality to the weighted inner product. Moreover, the upper probability bound is further improved by two times owing to utilizing the speciality of the constructed symmetric matrix. These findings shed some new light on the understanding of the L_2 norm of random matrices and bound them more tightly. Motivated by the exponential bound we found, a non-uniform sampling based sparse binary matrix approximation is presented to accelerate computations and save storages for massive data matrices. In the sampling method, omitting and rounding are naturally combined together without the need for extra steps because the sampling probabilities we chosen tailor the segments containing the retained entries into equal length. Consequently, the implementation requires only one pass over the input matrix and the bound is simplified further.

Acknowledgments. This work was partially supported by 973 Project of Chinese Ministry of Science and Technology (Grant No.2004CB720103), National Natural Science Foundation of China (Grant No.70621001, 70531040, 70501030, 10601064, 70472074), National Natural Science Foundation of Beijing (Grant No.9073020), and BHP Billiton Cooperation of Australia.

References

1. Achlioptas, D., McSherry, F.: Fast Computation of Low Rank Matrix Approximations. In: STOC 2001: Proceedings of the 32nd annual ACM symposium on Theory of computing, pp. 611–618 (2001)
2. Alon, N., Krivelevich, M., Vu, V.H.: On the Concentration of Eigenvalues of Random Symmetric Matrices. Tech. Report 60, Microsoft Research (2000)
3. Arora, S., Hazan, E., Kale, S.: A Fast Random Sampling Algorithm for Sparsifying Matrices. In: Díaz, J., Jansen, K., Rolim, J.D.P., Zwick, U. (eds.) APPROX 2006 and RANDOM 2006. LNCS, vol. 4110, pp. 272–279. Springer, Heidelberg (2006)
4. Berry, M., Dumais, S., O'Brie, G.: Using Linear Algebra for Intelligent Information Retrieval. SIAM Review 37, 573–595 (1995)
5. Bar-Yossef, Z.: Sampling Lower Bounds via Information Theory. In: Proceedings of the 35th annual ACM symposium on Theory of computing, pp. 335–344 (2003)
6. Chen, Y.P.: Matrix Theories, 2nd edn. NorthWestern Polytechnical University Press, Xian (1998)
7. Deerwester, S., Dumais, S., Furnas, G., Landauer, T., Harshman, R.: Indexing by Latent Semantic Anlysis. J. American Society for information science 41, 391–407 (1990)
8. Deshpande, A., Vempala, S.: Adaptive Sampling and Fast Low-Rank Matrix Approximation. In: Díaz, J., Jansen, K., Rolim, J.D.P., Zwick, U. (eds.) APPROX 2006 and RANDOM 2006. LNCS, vol. 4110, pp. 292–303. Springer, Heidelberg (2006)
9. Drineas, P., Frieze, A., Kannan, R., Vempala, S., Vinay, V.: Clustering in large graphs and matrices. In: Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 291–299 (1999)
10. Feige, U., Ofek, E.: Spectral Techniques Applied to Sparse random graphs. Random Structures and Algorithms 27(2), 251–275 (2005)
11. Frieze, A., Kannan, R., Vempala, S.: Fast Monte-Carlo algorithms for finding low-rank approximations. In: 39th Annual Symposium on Foundations of Computer Science, pp. 370–378 (1998)
12. Frieze, A., Kannan, R., Vempala, S.: Fast Monte-Carlo algorithms for finding low-rank approximations. Journal of the ACM 51(6), 1025–1041 (2004)
13. Golub, G.H., Van Loan, C.F.: Matrix Computations, 3rd edn. Johns Hopkins University Press, Maryland (1996)
14. Mannila, H.: Local and Global Methods in Data Mining: Basic Techniques and Open Problems. In: Proceedings of the 29th International Colloquium on Automata, Languages and Programming, pp. 57–68 (2002)
15. Wasserman, L.: All of Statistics: A Concise Course in Statistical Inference. Springer, New York (2004)
16. Ye, J.P.: Generalized Low Rank Matrix Approximations of Matrices. Machine Learning 61(1-3), 167–191 (2005)