

Low Rank Matrix-valued Chernoff Bounds and Approximate Matrix Multiplication

Avner Magen *

Anastasios Zouzias †

Abstract

In this paper we develop algorithms for approximating matrix multiplication with respect to the spectral norm. Let $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{n \times p}$ be two matrices and $\varepsilon > 0$. We approximate the product $A^\top B$ using two sketches $\tilde{A} \in \mathbb{R}^{t \times m}$ and $\tilde{B} \in \mathbb{R}^{t \times p}$, where $t \ll n$, such that

$$\left\| \tilde{A}^\top \tilde{B} - A^\top B \right\|_2 \leq \varepsilon \|A\|_2 \|B\|_2$$

with high probability. We analyze two different sampling procedures for constructing \tilde{A} and \tilde{B} ; one of them is done by i.i.d. non-uniform sampling rows from A and B and the other by taking random linear combinations of their rows. We prove bounds on t that depend only on the intrinsic dimensionality of A and B , that is their rank and their stable rank.

For achieving bounds that depend on rank when taking random linear combinations we employ standard tools from high-dimensional geometry such as concentration of measure arguments combined with elaborate ε -net constructions. For bounds that depend on the smaller parameter of stable rank this technology itself seems weak. However, we show that in combination with a simple truncation argument it is amenable to provide such bounds. To handle similar bounds for row sampling, we develop a novel matrix-valued Chernoff bound inequality which we call low rank matrix-valued Chernoff bound. Thanks to this inequality, we are able to give bounds that depend only on the stable rank of the input matrices.

We highlight the usefulness of our approximate matrix multiplication bounds by supplying two applications. First we give an approximation algorithm for the ℓ_2 -regression problem that returns an approximate solution by randomly projecting the initial problem to dimensions linear on the rank of the constraint matrix. Second we give improved approximation algorithms for the low rank matrix approximation problem with respect to the spectral norm.

1 Introduction

In many scientific applications, data is often naturally expressed as a matrix, and computational problems on such data are reduced to standard matrix operations including matrix multiplication, ℓ_2 -regression, and low rank matrix approximation.

In this paper we analyze several approximation algorithms with respect to these operations. All of our algorithms share a common underlying framework which can be described as follows: Let A be an input matrix that we may want to apply a matrix computation on it to infer some useful information about the data that it represents. The main idea is to work with a sample of A (a.k.a. sketch), call it \tilde{A} , and hope that the obtained information from \tilde{A} will be in some sense close to the information that would have been extracted from A .

In this generality, the above approach (sometimes called “Monte-Carlo method for linear algebraic problems”) is ubiquitous, and is responsible for much of the development in fast matrix computations [FKV04, DKM06a, Sar06, DMM06, AM07, CW09, DR10].

As we sample A to create a sketch \tilde{A} , our goal is twofold: (i) guarantee that \tilde{A} resembles A in the relevant measure, and (ii) achieve such a \tilde{A} using as few samples as possible. The standard tool that provides a handle on these requirements when the objects are real numbers, is the Chernoff bound inequality. However, since we deal with matrices, we would like to have an analogous probabilistic tool suitable for matrices. Quite recently a non-trivial generalization of Chernoff bound type inequalities for matrix-valued random variables was introduced by Ahlswede and Winter [AW02]. Such inequalities are suitable for the type of problems that we will consider here. However, this type of inequalities and their variants that have been proposed in the literature [GLF⁺09, Rec09, Gro09, Tro10] all suffer from the fact that their bounds depend on the dimensionality of the samples. We argue that in a wide range of applications, this dependency can be quite detrimental.

Specifically, whenever the following two conditions hold we typically provide stronger bounds compared with the existing tools: (a) the input matrix has low intrinsic dimensionality such as rank or stable rank, (b) the matrix samples themselves have low rank. The validity of condition (a) is very common in applications from the simple fact that viewing data using matrices typically leads to redundant representations. Typical sampling methods tend to rely on extremely simple

*University of Toronto, Department of Computer Science, Email : avner@cs.toronto.edu.

†University of Toronto, Department of Computer Science, Email : zouzias@cs.toronto.edu.

sampling matrices, i.e., samples that are supported on only one entry [AHK06, AM07, DZ10] or samples that are obtained by the outer-product of the sampled rows or columns [DKM06a, RV07], therefore condition (b) is often natural to assume. By incorporating the rank assumption of the matrix samples on the above matrix-valued inequalities we are able to develop a “dimension-free” matrix-valued Chernoff bound. See Theorem 1.1 for more details.

Fundamental to the applications we derive, are two probabilistic tools that provide concentration bounds of certain random matrices. These tools are inherently different, where each pertains to a different sampling procedure. In the first, we multiply the input matrix by a random sign matrix, whereas in the second we sample rows according to a distribution that depends on the input matrix. In particular, the first method is oblivious (the probability space does not depend on the input matrix) while the second is not.

The first tool is the so-called subspace Johnson-Lindenstrauss lemma. Such a result was obtained in [Sar06] (see also [Cla08, Theorem 1.3]) although it appears implicitly in results extending the original Johnson Lindenstrauss lemma (see [Mag07]). The techniques for proving such a result with possible worse bound are not new and can be traced back even to Milman’s proof of Dvoretzky theorem [Mil71].

LEMMA 1.1. (*Subspace JL lemma [Sar06]*) Let $\mathcal{W} \subseteq \mathbb{R}^d$ be a linear subspace of dimension k and $\varepsilon \in (0, 1/3)$. Let R be a $t \times d$ random sign matrix rescaled by $1/\sqrt{t}$, namely $R_{ij} = \pm 1/\sqrt{t}$ with equal probability. Then

$$\mathbb{P} \left((1 - \varepsilon) \|w\|_2^2 \leq \|Rw\|_2^2 \leq (1 + \varepsilon) \|w\|_2^2, \forall w \in \mathcal{W} \right) \geq 1 - c_2^k \cdot \exp(-c_1 \varepsilon^2 t), \quad (1.1)$$

where $c_1 > 0, c_2 > 1$ are constants.

The importance of such a tool, is that it allows us to get bounds on the necessary dimensions of the random sign matrix in terms of the *rank* of the input matrices, see Theorem 3.2 (i.a).

While the assumption that the input matrices have low rank is a fairly reasonable assumption, one should be a little cautious as the property of having low rank is not robust. Indeed, if random noise is added to a matrix, even if low rank, the matrix obtained will have full rank almost surely. On the other hand, it can be shown that the added noise cannot distort the Frobenius and operator norm significantly; which makes the notion of *stable rank* robust and so the assumption of low stable rank on the input is more applicable than the low rank assumption.

Given the above discussion, we resort to a different methodology, called matrix-valued Chernoff bounds. These are non-trivial generalizations of the standard Chernoff bounds over the reals and were first introduced in [AW02]. Part of the contribution of the current work is to show that such inequalities, similarly to their real-valued ancestors, provide powerful tools to analyze randomized algorithms. There is a rapidly growing line of research exploiting the power of such inequalities including matrix approximation by sparsification [AM07, DZ10]; analysis of algorithms for matrix completion and decomposition of low rank matrices [CR07, Gro09, Rec09]; and semi-definite relaxation and rounding of quadratic maximization problems [Nem07, So09a, So09b].

The quality of these bounds can be measured by the number of samples needed in order to obtain small error probability. The original result of [AW02, Theorem 19] shows that¹ if M is distributed according to some distribution over $n \times n$ matrices with zero mean², and if M_1, \dots, M_t are independent copies of M then for any $\varepsilon > 0$,

$$\mathbb{P} \left(\left\| \frac{1}{t} \sum_{i=1}^t M_i \right\|_2 > \varepsilon \right) \leq n \exp \left(-C \frac{\varepsilon^2 t}{\gamma^2} \right), \quad (1.2)$$

where $\|M\|_2 \leq \gamma$ holds almost surely and $C > 0$ is an absolute constant.

Notice that the number of samples in Ineq. (1.2) depends logarithmically in n . In general, unfortunately, such a dependency is inevitable: take for example a diagonal random sign matrix of dimension n . The operator norm of the sum of t independent samples is precisely the maximum deviation among n independent random walks of length t . In order to achieve a fixed bound on the maximum deviation with constant probability, it is easy to see that t should grow logarithmically with n in this scenario.

In their seminal paper, Rudelson and Vershynin provide a matrix-valued Chernoff bound that avoids the dependency on the dimensions by assuming that the matrix samples are the *outer product* $x \otimes x$ of a randomly distributed vector x [RV07]. It turns out that this assumption is too strong in most applications, such as the ones we study in this work, and so we wish to relax it without increasing the bound significantly. In the following theorem we replace this assumption with that of having *low rank*. We should note that we

¹For ease of presentation we actually provide the restatement presented in [WX08, Theorem 2.6], which is more suitable for this discussion.

²Zero mean means that the (matrix-valued) expectation is the zero $n \times n$ matrix.

are not aware of a simple way to extend Theorem 3.1 of [RV07] to the low rank case, even constant rank. The main technical obstacle is the use of the powerful Rudelson selection lemma, see [Rud99] or Lemma 3.5 of [RV07], which applies only for Rademacher sums of outer product of vectors. We bypass this obstacle by proving a more general lemma, see Lemma 5.2. The proof of Lemma 5.2 relies on the non-commutative Khintchine moment inequality [LP86, Buc01] which is also the backbone in the proof of Rudelson’s selection lemma. With Lemma 5.2 at our disposal, the proof techniques of [RV07] can be adapted to support our more general condition.

THEOREM 1.1. *Let $0 < \varepsilon < 1$ and M be a random symmetric real matrix with $\|\mathbb{E} M\|_2 \leq 1$ and $\|M\|_2 \leq \gamma$ almost surely. Assume that each element on the support of M has at most rank r . Set $t = \Omega(\gamma \log(\gamma/\varepsilon^2)/\varepsilon^2)$. If $r \leq t$ holds almost surely, then*

$$\mathbb{P} \left(\left\| \frac{1}{t} \sum_{i=1}^t M_i - \mathbb{E} M \right\|_2 > \varepsilon \right) \leq \frac{1}{\text{poly}(t)}.$$

where M_1, M_2, \dots, M_t are i.i.d. copies of M .

Proof. See Appendix, page 12.

REMARK 1. (OPTIMALITY) *The above theorem cannot be improved in terms of the number of samples required without changing its form, since in the special case where the rank of the samples is one it is exactly the statement of Theorem 3.1 of [RV07], see [RV07, Remark 3.4].*

We highlight the usefulness of the above main tools by first proving a “dimension-free” approximation algorithm for matrix multiplication with respect to the spectral norm (Section 3.1). Utilizing this matrix multiplication bound we get an approximation algorithm for the ℓ_2 -regression problem which returns an approximate solution by randomly projecting the initial problem to dimensions linear on the rank of the constraint matrix (Section 3.2). Finally, in Section 3.3 we give improved approximation algorithms for the low rank matrix approximation problem with respect to the spectral norm, and moreover answer in the affirmative a question left open by the authors of [NDT09].

2 Preliminaries and Definitions

The next discussion reviews several definitions and facts from linear algebra; for more details, see [SS90, GV96, Bha96]. We abbreviate the terms independently and identically distributed and almost surely with i.i.d. and a.s., respectively. We let $\mathbb{S}^{n-1} := \{x \in \mathbb{R}^n \mid \|x\|_2 =$

$1\}$ be the $(n - 1)$ -dimensional sphere. A *random Gaussian* matrix is a matrix whose entries are i.i.d. standard Gaussians, and a *random sign* matrix is a matrix whose entries are independent Bernoulli random variables, that is they take values from $\{\pm 1\}$ with equal probability. For a matrix $A \in \mathbb{R}^{n \times m}$, $A_{(i)}$, $A^{(j)}$, denote the i ’th row, j ’th column, respectively. For a matrix with rank r , the Singular Value Decomposition (SVD) of A is the decomposition of A as $U\Sigma V^\top$ where $U \in \mathbb{R}^{n \times r}$, $V \in \mathbb{R}^{m \times r}$ where the columns of U and V are orthonormal, and $\Sigma = \text{diag}(\sigma_1(A), \dots, \sigma_r(A))$ is $r \times r$ diagonal matrix. We further assume $\sigma_1 \geq \dots \geq \sigma_r > 0$ and call these real numbers the *singular values* of A . By $A_k = U_k \Sigma_k V_k^\top$ we denote the best rank k approximation to A , where U_k and V_k are the matrices formed by the first k columns of U and V , respectively. We denote by $\|A\|_2 = \max\{\|Ax\|_2 \mid \|x\|_2 = 1\}$ the spectral norm of A , and by $\|A\|_F = \sqrt{\sum_{i,j} A_{ij}^2}$ the Frobenius norm of A . We denote by A^\dagger the Moore-Penrose pseudo-inverse of A , i.e., $A^\dagger = V\Sigma^{-1}U^\top$. Notice that $\sigma_1(A) = \|A\|_2$. Also we define by $\text{sr}(A) := \|A\|_F^2 / \|A\|_2^2$ the *stable rank* of A . Notice that the inequality $\text{sr}(A) \leq \text{rank}(A)$ always holds. The orthogonal projector of a matrix A onto the row-space of a matrix C is denoted by $P_C(A) = AC^\dagger C$. By $P_{C,k}(A)$ we define the best rank- k approximation of the matrix $P_C(A)$.

3 Applications

All the proofs of this section have been deferred to Section 4.

3.1 Matrix Multiplication The seminal research of [FKV04] focuses on using non-uniform row sampling to speed-up the running time of several matrix computations. The subsequent developments of [DKM06a, DKM06b, DKM06c] also study the performance of Monte-Carlo algorithms on primitive matrix algorithms including the matrix multiplication problem with respect to the Frobenius norm. Sarlos [Sar06] extended (and improved) this line of research using random projections. Most of the bounds for approximating matrix multiplication in the literature are mostly with respect to the Frobenius norm [DKM06a, Sar06, CW09]. In some cases, the techniques that are utilized for bounding the Frobenius norm also imply *weak* bounds for the spectral norm, see [DKM06a, Theorem 4] or [Sar06, Corollary 11] which is similar with part (*i.a*) of Theorem 3.2.

In this section we develop approximation algorithms for matrix multiplication with respect to the spectral norm. The algorithms that will be presented in this section are based on the tools mentioned in Section 1.

Variants of Matrix-valued Inequalities				
Assumption on the sample M	# of samples (t)	Failure Prob.	References	Comments
$\ M\ _2 \leq \gamma$ a.s.	$\Omega(\gamma^2 \log(n)/\varepsilon^2)$	$1/\text{poly}(n)$	[WX08]	Hoeffding
$\ M\ _2 \leq \gamma$ a.s., $\ \mathbb{E}M^2\ _2 \leq \rho^2$	$\Omega((\rho^2 + \gamma\varepsilon/3) \log(n)/\varepsilon^2)$	$1/\text{poly}(n)$	[Rec09]	Bernstein
$\ M\ _2 \leq \gamma$ a.s., $M = x \otimes x$, $\ \mathbb{E}M\ _2 \leq 1$	$\Omega(\gamma \log(\gamma/\varepsilon^2)/\varepsilon^2)$	$\exp(-\Omega(\varepsilon^2 t/(\gamma \log t)))$	[RV07]	Rank one
$\ M\ _2 \leq \gamma$, $\text{rank}(M) \leq t$ a.s., $\ \mathbb{E}M\ _2 \leq 1$	$\Omega(\gamma \log(\gamma/\varepsilon^2)/\varepsilon^2)$	$1/\text{poly}(t)$	Theorem 1.1	Low rank

Table 1: Summary of matrix-valued Chernoff bounds. M is a probability distribution over symmetric $n \times n$ matrices. M_1, \dots, M_t are i.i.d. copies of M .

Before stating our main dimension-free matrix multiplication theorem (Theorem 3.2), we discuss the best possible bound that can be achieved using the current known matrix-valued inequalities (to the best of our knowledge). Consider a direct application of Ineq. (1.2), where a similar analysis with that in proof of Theorem 3.2 (ii) would allow us to achieve a bound of $\Omega(\tilde{r}^2 \log(m+p)/\varepsilon^2)$ on the number of samples (details omitted). However, as the next theorem indicates (proof omitted) we can get linear dependency on the stable rank of the input matrices gaining from the “variance information” of the samples; more precisely, this can be achieved by applying the matrix-valued Bernstein Inequality see e.g. [GLF⁺09], [Rec09, Theorem 3.2] or [Tro10, Theorem 2.10].

THEOREM 3.1. *Let $0 < \varepsilon < 1/2$ and let $A \in \mathbb{R}^{n \times m}$, $B \in \mathbb{R}^{n \times p}$ both having stable rank at most \tilde{r} . The following hold:*

- (i) *Let R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$. Denote by $\tilde{A} = RA$ and $\tilde{B} = RB$. If $t = \Omega(\tilde{r} \log(m+p)/\varepsilon^2)$ then*

$$\mathbb{P}\left(\left\|\tilde{A}^\top \tilde{B} - A^\top B\right\|_2 \leq \varepsilon \|A\|_2 \|B\|_2\right) \geq 1 - \frac{1}{\text{poly}(\tilde{r})}.$$

- (ii) *Let $p_i = \|A_{(i)}\|_2 \|B_{(i)}\|_2 / S$, where $S = \sum_{i=1}^n \|A_{(i)}\|_2 \|B_{(i)}\|_2$ be a probability distribution over $[n]$. If we form a $t \times m$ matrix \tilde{A} and a $t \times p$ matrix \tilde{B} by taking $t = \Omega(\tilde{r} \log(m+p)/\varepsilon^2)$ i.i.d. (row indices) samples from p_i , then*

$$\mathbb{P}\left(\left\|\tilde{A}^\top \tilde{B} - A^\top B\right\|_2 \leq \varepsilon \|A\|_2 \|B\|_2\right) \geq 1 - \frac{1}{\text{poly}(\tilde{r})}.$$

Notice that the above bounds depend linearly on the stable rank of the matrices and logarithmically on their dimensions. As we will see in the next theorem we can remove the dependency on the dimensions, and replace it with the stable rank. Recall that in most cases matrices *do* have low stable rank, which is much smaller than their dimensionality.

THEOREM 3.2. *Let $0 < \varepsilon < 1/2$ and let $A \in \mathbb{R}^{n \times m}$, $B \in \mathbb{R}^{n \times p}$ both having rank and stable rank at most r and \tilde{r} , respectively. The following hold:*

- (i) *Let R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$. Denote by $\tilde{A} = RA$ and $\tilde{B} = RB$.*

- (a) *If $t = \Omega(r/\varepsilon^2)$ then*

$$\mathbb{P}(\forall x \in \mathbb{R}^m, y \in \mathbb{R}^p, |x^\top (\tilde{A}^\top \tilde{B} - A^\top B)y| \leq \varepsilon \|Ax\|_2 \|By\|_2) \geq 1 - e^{-\Omega(r)}.$$

- (b) *If $t = \Omega(\tilde{r}/\varepsilon^4)$ then*

$$\mathbb{P}\left(\left\|\tilde{A}^\top \tilde{B} - A^\top B\right\|_2 \leq \varepsilon \|A\|_2 \|B\|_2\right) \geq 1 - e^{-\Omega(\frac{\tilde{r}}{\varepsilon^2})}.$$

- (ii) *Let $p_i = \|A_{(i)}\|_2 \|B_{(i)}\|_2 / S$, where $S = \sum_{i=1}^n \|A_{(i)}\|_2 \|B_{(i)}\|_2$ be a probability distribution over $[n]$. If we form a $t \times m$ matrix \tilde{A} and a $t \times p$ matrix \tilde{B} by taking $t = \Omega(\tilde{r} \log(\tilde{r}/\varepsilon^2)/\varepsilon^2)$ i.i.d. (row indices) samples from p_i , then*

$$\mathbb{P}\left(\left\|\tilde{A}^\top \tilde{B} - A^\top B\right\|_2 \leq \varepsilon \|A\|_2 \|B\|_2\right) \geq 1 - \frac{1}{\text{poly}(\tilde{r})}.$$

REMARK 2. *In part (ii), we can actually achieve the stronger bound of $t = \Omega(\sqrt{\text{sr}(A) \text{sr}(B)} \log(\text{sr}(A) \text{sr}(B)/\varepsilon^4)/\varepsilon^2)$ (see proof). However, for ease of presentation and comparison we give the above displayed bound.*

Part (i.b) follows from (i.a) via a simple truncation argument. This was pointed out to us by Mark Rudelson (personal communication). To understand the significance and the differences between the different components of this theorem, we first note that the probabilistic event of part (i.a) is superior to the probabilistic event of (i.b) and (ii). Indeed, when $B = A$ the former implies that $|x^\top (\tilde{A}^\top \tilde{A} - A^\top A)x| < \varepsilon \cdot x^\top A^\top A x$ for every x , which is stronger than $\left\|\tilde{A}^\top \tilde{A} - A^\top A\right\|_2 \leq \varepsilon \|A\|_2^2$. We will *heavily* exploit this fact in Section 4.3 to prove Theorem 3.4 (i.a) and (ii). Also notice that part (i.b) is

essential computationally inferior to (ii) as it gives the same bound while it is more expensive computationally to multiply the matrices by random sign matrices than just sampling their rows. However, the advantage of part (i) is that the sampling process is *oblivious*, i.e., does not depend on the input matrices.

3.2 ℓ_2 -regression In this section we present an approximation algorithm for the least-squares regression problem; given an $n \times m$, $n > m$, real matrix A of rank r and a real vector $b \in \mathbb{R}^n$ we want to compute $x_{\text{opt}} = A^\dagger b$ that minimizes $\|Ax - b\|_2$ over all $x \in \mathbb{R}^m$. In their seminal paper [DMM06], Drineas et al. show that if we non-uniformly sample $t = \Omega(m^2/\varepsilon^2)$ rows from A and b , then with high probability the optimum solution of the $t \times d$ sampled problem will be within $(1 + \varepsilon)$ close to the original problem. The main drawback of their approach is that finding or even approximating the sampling probabilities is computationally intractable. Sarlos [Sar06] improved the above to $t = \Omega(m \log m/\varepsilon^2)$ and gave the first $o(nm^2)$ relative error approximation algorithm for this problem.

In the next theorem we eliminate the extra $\log m$ factor from Sarlos bounds, and more importantly, replace the dimension (number of variables) m with the rank r of the constraints matrix A . We should point out that independently, the same bound as our Theorem 3.3 was recently obtained by Clarkson and Woodruff [CW09] (see also [DMMS09]). The proof of Clarkson and Woodruff uses heavy machinery and a completely different approach. In a nutshell they manage to improve the matrix multiplication bound with respect to the Frobenius norm. They achieve this by bounding higher moments of the Frobenius norm of the approximation viewed as a random variable instead of bounding the *local* differences for each coordinate of the product. To do so, they rely on intricate moment calculations spanning over four pages, see [CW09] for more. On the other hand, the proof of the present ℓ_2 -regression bound uses only basic matrix analysis, elementary deviation bounds and ε -net arguments. More precisely, we argue that Theorem 3.2 (i.a) immediately implies that by randomly-projecting to dimensions linear in the intrinsic dimensionality of the constraints, i.e., the rank of A , is sufficient as the following theorem indicates.

THEOREM 3.3. *Let $A \in \mathbb{R}^{n \times m}$ be a real matrix of rank r and $b \in \mathbb{R}^n$. Let $\min_{x \in \mathbb{R}^m} \|b - Ax\|_2$ be the ℓ_2 -regression problem, where the minimum is achieved with $x_{\text{opt}} = A^\dagger b$. Let $0 < \varepsilon < 1/3$, R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$ and $\tilde{x}_{\text{opt}} = (RA)^\dagger Rb$.*

- If $t = \Omega(r/\varepsilon)$, then with high probability,

$$(3.3) \quad \|b - A\tilde{x}_{\text{opt}}\|_2 \leq (1 + \varepsilon) \|b - Ax_{\text{opt}}\|_2.$$

- If $t = \Omega(r/\varepsilon^2)$, then with high probability,

$$(3.4) \quad \|x_{\text{opt}} - \tilde{x}_{\text{opt}}\|_2 \leq \frac{\varepsilon}{\sigma_{\min}(A)} \|b - Ax_{\text{opt}}\|_2.$$

REMARK 3. *The above result can be easily generalized to the case where b is an $n \times p$ matrix B of rank at most r (see proof). This is known as the generalized ℓ_2 -regression problem in the literature, i.e., $\arg \min_{X \in m \times p} \|AX - B\|_2$ where B is an $n \times p$ rank r matrix.*

3.3 Spectral Low Rank Matrix Approximation

A large body of work on low rank matrix approximations [DK03, FKV04, DRVW06, Sar06, RV07, AM07, RST09, CW09, NDT09, HMT09] has been recently developed with main objective to develop more efficient algorithms for this task. Most of these results study approximation algorithms with respect to the Frobenius norm, except for [RV07, NDT09] that handle the spectral norm.

In this section we present two $(1 + \varepsilon)$ -relative-error approximation algorithms for this problem with respect to the spectral norm, i.e., given an $n \times m$, $n > m$, real matrix A of rank r , we wish to compute $A_k = U_k \Sigma_k V_k^\top$, which minimizes $\|A - X_k\|_2$ over the set of $n \times m$ matrices of rank k , X_k . The first additive bound for this problem was obtained in [RV07]. To the best of our knowledge the best relative bound was recently achieved in [NDT09, Theorem 1]. The latter result is not directly comparable with ours, since it uses a more restricted projection methodology and so their bound is weaker compared to our results. The first algorithm randomly projects the rows of the input matrix onto t dimension. Here, we set t to be either $\Omega(r/\varepsilon^2)$ in which case we get an $(1 + \varepsilon)$ error guarantee, or to be $\Omega(k/\varepsilon^2)$ in which case we show a $(2 + \varepsilon\sqrt{(r-k)/k})$ error approximation. In both cases the algorithm succeeds with high probability. The second approximation algorithm samples non-uniformly $\Omega(r \log(r/\varepsilon^2)/\varepsilon^2)$ rows from A in order to satisfy the $(1 + \varepsilon)$ guarantee with high probability.

The following lemma (Lemma 3.1) is essential for proving both relative error bounds of Theorem 3.4. It gives a sufficient condition that any matrix \tilde{A} should satisfy in order to get a $(1 + \varepsilon)$ spectral low rank matrix approximation of A for every k , $1 \leq k \leq \text{rank}(A)$.

LEMMA 3.1. *Let A be an $n \times m$ matrix and $\varepsilon > 0$. If there exists a $t \times m$ matrix \tilde{A} such that for every $x \in \mathbb{R}^m$, $(1 - \varepsilon)x^\top A^\top Ax \leq x^\top \tilde{A}^\top \tilde{A} x \leq (1 + \varepsilon)x^\top A^\top Ax$, then*

$$\|A - P_{\tilde{A},k}(A)\|_2 \leq (1 + \varepsilon) \|A - A_k\|_2,$$

for every $k = 1, \dots, \text{rank}(A)$.

The theorem below shows that it's possible to satisfy the conditions of Lemma 3.1 by randomly projecting A onto $\Omega(r/\varepsilon^2)$ or by non-uniform sampling i.i.d. $\Omega(r \log(r/\varepsilon^2)/\varepsilon^2)$ rows of A as described in parts (i.a) and (ii), respectively.

THEOREM 3.4. *Let $0 < \varepsilon < 1/3$ and let $A = U\Sigma V^\top$ be a real $n \times m$ matrix of rank r with $n \geq m$.*

(i) (a) *Let R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$ and set $\tilde{A} = RA$. If $t = \Omega(r/\varepsilon^2)$, then with high probability*

$$\|A - P_{\tilde{A},k}(A)\|_2 \leq (1 + \varepsilon) \|A - A_k\|_2,$$

for every $k = 1, \dots, r$.

(b) *Let R be a $t \times n$ random Gaussian matrix rescaled by $1/\sqrt{t}$ and set $\tilde{A} = RA$. If $t = \Omega(k/\varepsilon^2)$, then with high probability*

$$\|A - P_{\tilde{A},k}(A)\|_2 \leq (2 + \varepsilon \sqrt{\frac{r-k}{k}}) \|A - A_k\|_2.$$

(ii) *Let $p_i = \|U_{(i)}\|_2^2/r$ be a probability distribution over $[n]$. Let \tilde{A} be a $t \times m$ matrix that is formed (row-by-row) by taking t i.i.d. samples from p_i and rescaled appropriately. If $t = \Omega(r \log(r/\varepsilon^2)/\varepsilon^2)$, then with high probability*

$$\|A - P_{\tilde{A},k}(A)\|_2 \leq (1 + \varepsilon) \|A - A_k\|_2,$$

for every $k = 1, \dots, r$.

We should highlight that in part (ii) the probability distribution p_i is in general hard to compute. Indeed, computing $\|U_{(i)}\|_2^2$ requires computing the SVD of A . In general, these values are known as statistical leverage scores [DM10]. In the special case where A is an edge-vertex matrix of an undirected weighted graph then p_i , the probability distribution over edges (rows), corresponds to the effective-resistance of the i -th edge [SS08].

Theorem 3.4 gives an $(1 + \varepsilon)$ approximation algorithm for the special case of low rank matrices. However, as discussed in Section 1 such an assumption is too restrictive for most applications. In the following theorem, we make a step further and relax the rank condition with a condition that depends on the stable rank of the residual matrix $A - A_k$. More formally, for an integer $k \geq 1$, we say that a matrix A has a k -low stable rank tail iff $k \geq \text{sr}(A - A_k)$.

Notice that the above definition is useful since it contains the set of matrices whose spectrum follows

a power-law distribution and those with exponentially decaying spectrum. Therefore the following theorem combined with the remark below (partially) answers in the affirmative the question posed by [NDT09]: Is there a relative error approximation algorithm with respect to the spectral norm when the spectrum of the input matrix decays in a power law?

THEOREM 3.5. *Let $0 < \varepsilon < 1/3$ and let A be a real $n \times m$ matrix with a k -low stable rank tail. Let R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$ and set $\tilde{A} = RA$. If $t = \Omega(k/\varepsilon^4)$, then with high probability*

$$\|A - P_{\tilde{A},k}(A)\|_2 \leq (2 + \varepsilon) \|A - A_k\|_2.$$

REMARK 4. *The $(2 + \varepsilon)$ bound can be improved to a relative $(1 + \varepsilon)$ error bound if we return as the approximate solution a slightly higher rank matrix, i.e., by returning the matrix $P_{\tilde{A}}(A)$, which has rank at most $t = \Omega(k/\varepsilon^4)$ (see [HMT09, Theorem 9.1]).*

4 Proofs

4.1 Proof of Theorem 3.2 (Matrix Multiplication)

Random Projections - Part (i)

Part (a): In this section we show the first, to the best of our knowledge, non-trivial spectral bound for matrix multiplication. Although the proof is an immediate corollary of the subspace Johnson-Lindenstrauss lemma (Lemma 1.1), this result is powerful enough to give, for example, tight bounds for the ℓ_2 regression problem. We prove the following more general theorem from which Theorem 3.2 (i.a) follows by plugging in $t = \Omega(r/\varepsilon^2)$.

THEOREM 4.1. *Let $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{n \times p}$. Assume that the ranks of A and B are at most r . Let R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$. Denote by $\tilde{A} = RA$ and $\tilde{B} = RB$. The following inequality holds*

$$\begin{aligned} \mathbb{P} \left(\forall x \in \mathbb{R}^m, y \in \mathbb{R}^p, \quad |x^\top (\tilde{A}^\top \tilde{B} - A^\top B)y| \leq \varepsilon \|Ax\|_2 \|By\|_2 \right) \\ \geq 1 - c_2^x \exp(-c_1 \varepsilon^2 t), \end{aligned}$$

where $c_1 > 0, c_2 > 1$ are constants.

Proof. (of Theorem 4.1) Let $A = U_A \Sigma_A V_A^\top$, $B = U_B \Sigma_B V_B^\top$ be the singular value decomposition of A and B respectively. Notice that $U_A \in \mathbb{R}^{n \times r_A}$, $U_B \in \mathbb{R}^{n \times r_B}$, where r_A and r_B is the rank of A and B , respectively.

Let $x_1 \in \mathbb{R}^m, x_2 \in \mathbb{R}^p$ two arbitrary unit vectors. Let $w_1 = Ax_1$ and $w_2 = Bx_2$. Recall that

$$\|A^\top R^\top RB - A^\top B\|_2 =$$

$$\sup_{x_1 \in \mathbb{S}^{m-1}, x_2 \in \mathbb{S}^{p-1}} |x_1^\top (A^\top R^\top R B - A^\top B) x_2|.$$

We will bound the last term for any arbitrary vector. Denote with \mathcal{V} the subspace³ $\text{colspan}(U_A) \cup \text{colspan}(U_B)$ of \mathbb{R}^n . Notice that the size of $\dim(\mathcal{V}) \leq r_A + r_B \leq 2r$. Applying Lemma 1.1 to \mathcal{V} , we get that with probability at least $1 - c_2^2 \exp(-c_1 \varepsilon^2 t)$ that

$$(4.5) \quad \forall v \in \mathcal{V} : \quad | \|Rv\|_2^2 - \|v\|_2^2 | \leq \varepsilon \|v\|_2^2.$$

Therefore we get that for any unit vectors $v_1, v_2 \in \mathcal{V}$:

$$\begin{aligned} (Rv_1)^\top Rv_2 &= \frac{\|Rv_1 + Rv_2\|_2^2 - \|Rv_1 - Rv_2\|_2^2}{4} \\ &\leq \frac{(1 + \varepsilon) \|v_1 + v_2\|_2^2 - (1 - \varepsilon) \|v_1 - v_2\|_2^2}{4} \\ &= \frac{\|v_1 + v_2\|_2^2 - \|v_1 - v_2\|_2^2}{4} \\ &+ \varepsilon \frac{\|v_1 + v_2\|_2^2 + \|v_1 - v_2\|_2^2}{4} \\ &= v_1^\top v_2 + \varepsilon \frac{\|v_1\|_2^2 + \|v_2\|_2^2}{2} = v_1^\top v_2 + \varepsilon, \end{aligned}$$

where the first equality follows from the Parallelogram law, the first inequality follows from Equation (4.5), and the last inequality since v_1, v_2 are unit vectors. By similar considerations we get that $(Rv_1)^\top Rv_2 \geq v_1^\top v_2 - \varepsilon$. By linearity of R , we get that

$$\forall v_1, v_2 \in \mathcal{V} : \quad |(Rv_1)^\top Rv_2 - v_1^\top v_2| \leq \varepsilon \|v_1\|_2 \|v_2\|_2.$$

Notice that $w_1, w_2 \in \mathcal{V}$, hence $|w_1^\top R^\top R w_2 - w_1^\top w_2| \leq \varepsilon \|w_1\|_2 \|w_2\|_2 = \varepsilon \|A x_1\|_2 \|B x_2\|_2$.

Part (b): We start with a technical lemma that bounds the spectral norm of any matrix A when it's multiplied by a random sign matrix rescaled by $1/\sqrt{t}$.

LEMMA 4.1. *Let A be an $n \times m$ real matrix, and let R be a $t \times n$ random sign matrix rescaled by $1/\sqrt{t}$. If $t \geq sr(A)$, then*

$$(4.6) \quad \mathbb{P}(\|RA\|_2 \geq 4\|A\|_2) \leq 2e^{-t/2}.$$

Proof. Without loss of generality assume that $\|A\|_2 = 1$. Then $\|A\|_F = \sqrt{sr(A)}$. Let G be a $t \times n$ Gaussian matrix. Then by the Gordon-Chevèet inequality⁴

$$\begin{aligned} \mathbb{E} \|GA\|_2 &\leq \|I_t\|_2 \|A\|_F + \|I_t\|_F \|A\|_2 \\ &= \|A\|_F + \sqrt{t} \leq 2\sqrt{t}. \end{aligned}$$

³We denote by $\text{colspan}(A)$ the subspace generated by the columns of A , and $\text{rowspan}(A)$ the subspace generated by the rows of A .

⁴For example, set $S = I_t, T = A$ in [HMT09, Proposition 10.1, p. 54].

The Gaussian distribution is symmetric, so G_{ij} and $\sqrt{t}R_{ij} \cdot |G_{ij}|$, where G_{ij} is a Gaussian random variable have the same distribution. By Jensen's inequality and the fact that $\mathbb{E}|G_{ij}| = \sqrt{2/\pi}$, we get that $\sqrt{2/\pi} \mathbb{E} \|RA\|_2 \leq \mathbb{E} \|GA\|_2 / \sqrt{t}$. Define the function $f : \{\pm 1\}^{t \times n} \rightarrow \mathbb{R}$ by $f(S) = \left\| \frac{1}{\sqrt{t}} SA \right\|_2$. The calculation above shows that $\text{median}(f) \leq \sqrt{2\pi}$. Since f is convex and $(1/\sqrt{t})$ -Lipschitz as a function of the entries of S , Talagrand's measure concentration inequality for convex functions yields

$$\mathbb{P}(\|RA\|_2 \geq \text{median}(f) + \delta) \leq 2 \exp(-\delta^2 t / 2).$$

Setting $\delta = 1$ in the above inequality implies the lemma.

Now using the above Lemma together with Theorem 3.2 (i.a) and a simple truncation argument we can prove part (i.b).

Proof. (of Theorem 3.2 (i.b)) Without loss of generality assume that $\|A\|_2 = \|B\|_2 = 1$. Set $r = \lfloor \frac{1600 \max\{sr(A), sr(B)\}}{\varepsilon^2} \rfloor$. Set $\hat{A} = A - A_r, \hat{B} = B - B_r$. Since $\|A\|_F^2 = \sum_{j=1}^{\text{rank}(A)} \sigma_j(A)^2$,

$$\|\hat{A}\|_2 \leq \frac{\|A\|_F}{\sqrt{r}} \leq \frac{\varepsilon}{40}, \quad \text{and} \quad \|\hat{B}\|_2 \leq \frac{\|B\|_F}{\sqrt{r}} \leq \frac{\varepsilon}{40}.$$

By triangle inequality, it follows that

$$\begin{aligned} \|\tilde{A}^\top \tilde{B} - A^\top B\|_2 & \leq \|A_r^\top R^\top R B_r - A_r^\top B_r\|_2 \\ & + \|\hat{A}^\top R^\top R B_r\|_2 \\ & + \|A_r^\top R^\top R \hat{B}\|_2 + \|\hat{A}^\top R^\top R \hat{B}\|_2 \\ & + \|\hat{A}^\top B_r\|_2 + \|A_r^\top \hat{B}\|_2 + \|\hat{A}^\top \hat{B}\|_2. \end{aligned} \tag{4.7-4.9}$$

Choose a constant in Theorem 3.2 (i.a) so that the failure probability of the right hand side of (4.7) does not exceed $\exp(-c\varepsilon^2 t)$, where $c = c_1/32$. The same argument shows that $\mathbb{P}(\|RA_r\|_2 \geq 1 + \varepsilon) \leq \exp(-c\varepsilon^2 t)$ and $\mathbb{P}(\|RB_r\|_2 \geq 1 + \varepsilon) \leq \exp(-c\varepsilon^2 t)$. This combined with Lemma 4.1 applied on \hat{A} and \hat{B} yields that the sum in (4.8) is less than $2(1 + \varepsilon)\varepsilon/10 + \varepsilon^2/100$. Also, since $\|A_r\|_2, \|B_r\|_2 \leq 1$, the sum in (4.9) is less than $2\varepsilon/10 + \varepsilon^2/100$. Combining the bounds for (4.7), (4.8) and (4.9) concludes the claim.

Row Sampling - Part (ii): By homogeneity normalize A and B such that $\|A\|_2 = \|B\|_2 = 1$. Notice that $A^\top B = \sum_{i=1}^n A_{(i)}^\top B_{(i)}$. Define $p_i = \frac{\|A_{(i)}^\top\|_2 \|B_{(i)}\|_2}{S}$, where $S = \sum_{i=1}^n \|A_{(i)}^\top\|_2 \|B_{(i)}\|_2$. Also

define a distribution over matrices in $\mathbb{R}^{(m+p) \times (m+p)}$ with n elements by

$$\mathbb{P} \left(M = \frac{1}{p_i} \begin{bmatrix} 0 & B_{(i)}^\top A_{(i)} \\ A_{(i)}^\top B_{(i)} & 0 \end{bmatrix} \right) = p_i.$$

First notice that

$$\begin{aligned} \mathbb{E} M &= \sum_{i=1}^n \frac{1}{p_i} \begin{bmatrix} 0 & B_{(i)}^\top A_{(i)} \\ A_{(i)}^\top B_{(i)} & 0 \end{bmatrix} \cdot p_i \\ &= \sum_{i=1}^n \begin{bmatrix} 0 & B_{(i)}^\top A_{(i)} \\ A_{(i)}^\top B_{(i)} & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & B^\top A \\ A^\top B & 0 \end{bmatrix}. \end{aligned}$$

This implies that $\|\mathbb{E} M\|_2 = \|A^\top B\|_2 \leq 1$. Next notice that the spectral norm of the random matrix M is upper bounded by $\sqrt{\text{sr}(A) \text{sr}(B)}$ almost surely. Indeed,

$$\begin{aligned} \|M\|_2 &\leq \sup_{i \in [n]} \left\| \frac{A_{(i)}^\top B_{(i)}}{p_i} \right\|_2 \\ &= S \sup_{i \in [n]} \left\| \frac{A_{(i)}^\top}{\|A_{(i)}\|_2} \frac{B_{(i)}}{\|B_{(i)}\|_2} \right\|_2 = S \cdot 1 \\ &= \sum_{i=1}^n \|A_{(i)}\|_2 \|B_{(i)}\|_2 \leq \|A\|_F \|B\|_F \\ &= \sqrt{\text{sr}(A) \text{sr}(B)} \leq (\text{sr}(A) + \text{sr}(B))/2, \end{aligned}$$

by definition of p_i , properties of norms, Cauchy-Schwartz inequality, and arithmetic/geometric mean inequality. Notice that this quantity (since the spectral norms of both A, B are one) is at most \tilde{r} by assumption. Also notice that every element on the support of the random variable M , has rank at most two. It is easy to see that, by setting $\gamma = \tilde{r}$, all the conditions in Theorem 1.1 are satisfied, and hence we get i_1, i_2, \dots, i_t indices from $[n]$, $t = \Omega(\tilde{r} \log(\tilde{r}/\varepsilon^2)/\varepsilon^2)$, such that with high probability

$$\left\| \frac{1}{t} \sum_{j=1}^t \begin{bmatrix} 0 & \frac{1}{p_{i_j}} B_{(i_j)}^\top A_{(i_j)} \\ \frac{1}{p_{i_j}} A_{(i_j)}^\top B_{(i_j)} & 0 \end{bmatrix} - \begin{bmatrix} 0 & B^\top A \\ A^\top B & 0 \end{bmatrix} \right\|_2 \leq \varepsilon.$$

The first sum can be rewritten as $\tilde{A}^\top \tilde{B}$ where $\tilde{A} = \frac{1}{\sqrt{t}} \begin{bmatrix} \frac{1}{\sqrt{p_{i_1}}} A_{(i_1)}^\top & \frac{1}{\sqrt{p_{i_2}}} A_{(i_2)}^\top & \dots & \frac{1}{\sqrt{p_{i_t}}} A_{(i_t)}^\top \end{bmatrix}^\top$ and $\tilde{B} = \frac{1}{\sqrt{t}} \begin{bmatrix} \frac{1}{\sqrt{p_{i_1}}} B_{(i_1)}^\top & \frac{1}{\sqrt{p_{i_2}}} B_{(i_2)}^\top & \dots & \frac{1}{\sqrt{p_{i_t}}} B_{(i_t)}^\top \end{bmatrix}^\top$. This concludes the theorem.

4.2 Proof of Theorem 3.3 (ℓ_2 -regression)

Proof. (of Theorem 3.3) Similarly as the proof in [Sar06]. Let $A = U\Sigma V^\top$ be the SVD of A . Let $b = Ax_{opt} + w$, where $w \in \mathbb{R}^n$ and $w \perp \text{colspan}(A)$. Also let $A(\tilde{x}_{opt} - x_{opt}) = Uy$, where $y \in \mathbb{R}^{\text{rank}(A)}$. Our goal is to bound this quantity

$$\begin{aligned} \|b - A\tilde{x}_{opt}\|_2^2 &= \|b - A(\tilde{x}_{opt} - x_{opt}) - Ax_{opt}\|_2^2 \\ &= \|w - Uy\|_2^2 \\ &= \|w\|_2^2 + \|Uy\|_2^2, \quad \text{since } w \perp \text{colspan}(U) \\ (4.10) \quad &= \|w\|_2^2 + \|y\|_2^2, \quad \text{since } U^\top U = I. \end{aligned}$$

It suffices to bound the norm of y , i.e., $\|y\|_2 \leq 3\varepsilon \|w\|_2$. Recall that given A, b the vector w is uniquely defined. On the other hand, vector y depends on the random projection R . Next we show the connection between y and w through the ‘‘normal equations’’.

$$\begin{aligned} RA\tilde{x}_{opt} &= Rb + w_2 \implies \\ RA\tilde{x}_{opt} &= R(Ax_{opt} + w) + w_2 \implies \\ RA(\tilde{x}_{opt} - x_{opt}) &= R w + w_2 \implies \\ U^\top R^\top R U y &= U^\top R^\top R w + U^\top R^\top w_2 \implies \\ (4.11) \quad U^\top R^\top R U y &= U^\top R^\top R w, \end{aligned}$$

where $w_2 \perp \text{colspan}(R)$, and used this fact to derive Ineq. (4.11). A crucial observation is that the $\text{colspan}(U)$ is perpendicular to w . Set $A = B = U$ in Theorem 3.2, and set $\varepsilon' = \sqrt{\varepsilon}$, and $t = \Omega(r/\varepsilon'^2)$. Notice that $\text{rank}(A) + \text{rank}(B) \leq 2r$, hence with constant probability we know that $1 - \varepsilon' \leq \sigma_i(RU) \leq 1 + \varepsilon'$. It follows that $\|U^\top R^\top R U y\|_2 \geq (1 - \varepsilon')^2 \|y\|_2$. A similar argument (set $A = U$ and $B = w$ in Theorem 3.2) guarantees that $\|U^\top R^\top R w\|_2 = \|U^\top R^\top R w - U^\top w\|_2 \leq \varepsilon' \|U\|_2 \|w\|_2 = \varepsilon' \|w\|_2$. Recall that $\|U\|_2 = 1$, since $U^\top U = I_n$ with high probability. Therefore, taking Euclidean norms on both sides of Equation (4.11) we get that

$$\|y\|_2 \leq \frac{\varepsilon'}{(1 - \varepsilon')^2} \|w\|_2 \leq 4\varepsilon' \|w\|_2.$$

Summing up, it follows from Equation (4.10) that, with constant probability, $\|b - A\tilde{x}_{opt}\|_2^2 \leq (1 + 16\varepsilon'^2) \|b - Ax_{opt}\|_2^2 = (1 + 16\varepsilon) \|b - Ax_{opt}\|_2^2$. This proves Ineq. (3.3).

Ineq. (3.4) follows directly from the bound on the norm of y repeating the above proof for $\varepsilon' \leftarrow \varepsilon$. First recall that x_{opt} is in the row span of A , since $x_{opt} = V\Sigma^{-1}U^\top b$ and the columns of V span the row space of A . Similarly for \tilde{x}_{opt} since the row span of $R \cdot A$ is contained in the row-span of A . Indeed, $\varepsilon \|w\|_2 \geq \|y\|_2 = \|Uy\|_2 = \|A(x_{opt} - \tilde{x}_{opt})\|_2 \geq \sigma_{\min}(A) \|x_{opt} - \tilde{x}_{opt}\|_2$.

4.3 Proof of Theorems 3.4, 3.5 (Spectral Low Rank Matrix Approximation)

Proof. (of Lemma 3.1) By the assumption and using Lemma 5.1 we get that

$$(4.12) \quad (1 - \varepsilon)\sigma_i(A^\top A) \leq \sigma_i(\tilde{A}^\top \tilde{A}) \leq (1 + \varepsilon)\sigma_i(A^\top A)$$

for all $i = 1, \dots, \mathbf{rank}(A)$. Let $\tilde{\Pi}_k$ be the projection matrix onto the first k right singular vectors of \tilde{A} , i.e., $(\tilde{A}_k)^\dagger \tilde{A}_k$. It follows that for every $k = 1, \dots, \mathbf{rank}(A)$

$$\begin{aligned} \|A - P_{\tilde{A},k}(A)\|_2 &\leq \|A - A\tilde{\Pi}_k\|_2^2 \\ &= \sup_{x \in \mathbb{R}^m, \|x\|=1} \|A(I - \tilde{\Pi}_k)x\|_2^2 \\ &= \sup_{x \in \ker \tilde{\Pi}_k, \|x\|=1} \|Ax\|_2^2 \\ &= \sup_{x \in \ker \tilde{\Pi}_k, \|x\|=1} x^\top A^\top Ax \\ &\leq (1 + \varepsilon) \sup_{x \in \ker \tilde{\Pi}_k, \|x\|_2=1} x^\top \tilde{A}^\top \tilde{A}x \\ &= (1 + \varepsilon)\sigma_{k+1}(\tilde{A}^\top \tilde{A}) \\ &\leq (1 + \varepsilon)^2\sigma_{k+1}(A^\top A) \\ &= (1 + \varepsilon)^2\|A - A_k\|_2^2, \end{aligned}$$

using that $x \perp \ker \tilde{\Pi}_k$ implies $\tilde{\Pi}_k x = x$, left side of the hypothesis, Courant-Fischer on $\tilde{A}^\top \tilde{A}$ (see Eqn. (5.17)), Eqn. (4.12), and properties of singular values, respectively.

Proof of Theorem 3.4 (i):

Part (a): Now we are ready to prove our first corollary of our matrix multiplication result to the problem of computing an approximate low rank matrix approximation of a matrix with respect to the spectral norm (Theorem 3.4).

Proof. Set $\tilde{A} = \frac{1}{\sqrt{t}}RA$ where R is a $\Omega(r/\varepsilon^2) \times n$ random sign matrix. Apply Theorem 3.2 (i.a) on A we have with high probability that

$$(4.13) \quad \forall x \in \mathbb{R}^n, (1 - \varepsilon)x^\top A^\top Ax \leq x^\top \tilde{A}^\top \tilde{A}x \leq (1 + \varepsilon)x^\top A^\top Ax.$$

Combining Lemma 3.1 with Ineq. (4.13) concludes the proof.

Part (b): The proof is based on the following lemma which reduces the problem of low rank matrix approximation to the problem of bounding the norm of a random matrix. We restate it here for reader's convenience and completeness [NDT09, Lemma 8], (see also [HMT09, Theorem 9.1] or [BMD09]).

LEMMA 4.2. Let $A = A_k + U_{r-k}\Sigma_{r-k}V_{r-k}^\top$, $H_k = U_{r-k}\Sigma_{r-k}$ and R be any $t \times n$ matrix. If the matrix (RU_k) has full column rank, then the following inequality holds,

$$(4.14) \quad \|A - P_{(RA),k}(A)\|_2 \leq 2\|A - A_k\|_2 + \|(RU_k)^\dagger RH_k\|_2.$$

Notice that the above lemma, reduces the problem of spectral low rank matrix approximation to a problem of approximation the spectral norm of the random matrix $(RU_k)^\dagger RH_k$.

First notice that by setting $t = \Omega(k/\varepsilon^2)$ we can guarantee that the matrix (RU_k) will have full column rank with high probability. Actually, we can say something much stronger; applying Theorem 3.2 (i.a) with $A = U_k$ we can guarantee that all the singular values are within $1 \pm \varepsilon$ with high probability. Now by conditioning on the above event ((RU_k) has full column rank), it follows from Lemma 4.2 that

$$\begin{aligned} \|A - P_{(RA),k}(A)\|_2 &\leq 2\|A - A_k\|_2 + \|(RU_k)^\dagger RH_k\|_2 \\ &\leq 2\|A - A_k\|_2 + \|(RU_k)^\dagger\|_2 \|RH_k\|_2 \\ &\leq 2\|A - A_k\|_2 + \frac{1}{1 - \varepsilon} \|RH_k\|_2 \\ &\leq 2\|A - A_k\|_2 + \frac{3}{2} \|RU_{r-k}\|_2 \|\Sigma_{r-k}\|_2 \end{aligned}$$

using the sub-multiplicative property of matrix norms, and that $\varepsilon < 1/3$. Now, it suffices to bound the norm of $W := RU_{r-k}$. Recall that $R = \frac{1}{\sqrt{t}}G$ where G is a $t \times n$ random Gaussian matrix, It is well-known that the distribution of the random matrix GU_{r-k} (by rotational invariance of the Gaussian distribution) has entries which are also i.i.d. Gaussian random variables. Now, we can use the following fact about random sub-Gaussian matrices to give a bound on the spectral norm of W . Indeed, we have the following

THEOREM 4.2. [RV09, Proposition 2.3] Let W be a $t \times (r - k)$ random matrix whose entries are independent mean zero Gaussian random variables. Assume that $r - k \geq t$, then

$$(4.15) \quad \mathbb{P}\left(\|W\|_2 \geq \delta\sqrt{r - k}\right) \leq e^{-c_0\delta^2\sqrt{r - k}}.$$

for any $\delta > \delta_0$, where δ_0 is a positive constant.

Apply union bound on the above theorem with δ be a sufficient large constant and on the conditions of Lemma 4.2, we get that with high probability, $\|W\|_2 \leq C_3\sqrt{r - k}$ and $\sigma_{\min}((RU_k)^\dagger) \leq 1/(1 - \varepsilon)$. Hence, Lemma 4.2 combined with the above discussion implies that

$$\begin{aligned}
\|A - P_{(RA),k}(A)\|_2 &\leq 2\|A - A_k\|_2 \\
&+ 3/2\|RU_{r-k}\|_2\|A - A_k\|_2 \\
&= 2\|A - A_k\|_2 \\
&+ \frac{3}{2\sqrt{t}}\|GU_{r-k}\|_2\|A - A_k\|_2 \\
&\leq \left(2 + c_4\varepsilon\sqrt{\frac{r-k}{k}}\right)\|A - A_k\|_2,
\end{aligned}$$

where $c_4 > 0$ is an absolute constant. Rescaling ε by c_4 concludes Theorem 3.4 (i.b).

Proof of Theorem 3.4 (ii) Here we prove that we can achieve the same relative error bound as with random projections by just sampling rows of A through a judiciously selected distribution. However, there is a price to pay and that's an extra logarithmic factor on the number of samples, as is stated in Theorem 3.4, part (ii).

Proof. (of Theorem 3.4 (ii)) The proof follows closely the proof of [SS08]. Similar with the proof of part (a). Let $A = U\Sigma V^\top$ be the singular value decomposition of A . Define the projector matrix $\Pi = UU^\top$ of size $n \times n$. Clearly, the rank of Π is equal to the rank of A and Π has the same image with A since every element in the image of A and Π is a linear combination of columns of U . Recall that for any projection matrix, the following holds $\Pi^2 = \Pi$ and hence $\text{sr}(\Pi) = \text{rank}(A) = r$. Moreover, $\sum_{i=1}^n \|U_{(i)}\|_2^2 = \text{tr}(UU^\top) = \text{tr}(\Pi) = \text{tr}(\Pi^2) = r$. Let $p_i = \Pi(i, i)/r = \|U_{(i)}\|_2^2/r$ be a probability distribution on $[n]$, where U_i is the i -th row of U .

Define a $t \times n$ random matrix S as follows: Pick t samples from p_i ; if the i -th sample is equal to $j \in [n]$ then set $S_{ij} = 1/\sqrt{tp_j}$. Notice that S has exactly one non-zero entry in each row, hence it has t non-zero entries. Define $\tilde{A} = SA$.

It is easy to verify that $\mathbb{E}_S \Pi S^\top S \Pi = \Pi^2 = \Pi$. Apply Theorem 1.1 (alternatively we can use [RV07, Theorem 3.1], since the matrix samples are rank one) on the matrix Π , notice that $\|\Pi\|_F^2 = r$ and $\|\Pi\|_2 = 1$, $\|\mathbb{E}_S \Pi S^\top S \Pi\|_2 \leq 1$, hence the stable rank of Π is r . Therefore, if $t = \Omega(r \log(r/\varepsilon^2)/\varepsilon^2)$ then with high probability

$$(4.16) \quad \|\Pi S^\top S \Pi - \Pi\|_2 \leq \varepsilon.$$

It suffices to show that Ineq. (4.16) is equivalent with

the condition of Lemma 3.1. Indeed,

$$\begin{aligned}
\sup_{x \in \mathbb{R}^n, x \neq 0} \left| \frac{x^\top (\Pi S^\top S \Pi - \Pi) x}{x^\top x} \right| &\leq \varepsilon \Leftrightarrow \\
\sup_{x \notin \ker \Pi, x \neq 0} \left| \frac{x^\top (\Pi S^\top S \Pi - \Pi) x}{x^\top x} \right| &\leq \varepsilon \Leftrightarrow \\
\sup_{y \in \text{Im}(A), y \neq 0} \left| \frac{y^\top (\Pi S^\top S \Pi - \Pi) y}{y^\top y} \right| &\leq \varepsilon \Leftrightarrow \\
\sup_{x \in \mathbb{R}^m, Ax \neq 0} \left| \frac{x^\top A^\top (\Pi S^\top S \Pi - \Pi) Ax}{x^\top A^\top Ax} \right| &\leq \varepsilon \Leftrightarrow \\
\sup_{x \in \mathbb{R}^m, Ax \neq 0} \left| \frac{x^\top (A^\top S^\top S A - A^\top A) x}{x^\top A^\top Ax} \right| &\leq \varepsilon \Leftrightarrow \\
\sup_{x \in \mathbb{R}^m, Ax \neq 0} \left| \frac{x^\top (\tilde{A}^\top \tilde{A} - A^\top A) x}{x^\top A^\top Ax} \right| &\leq \varepsilon,
\end{aligned}$$

since $x \notin \ker \Pi$ implies $x \in \text{Im}(A)$, $\text{Im}(A) \equiv \text{Im}(\Pi)$, and $\Pi A = A$. By re-arranging terms we get Equation (4.13) and so the claim follows.

Proof of Theorem 3.5: Similarly with the proof of Theorem 3.4 (i.b). By following the proof of part (i.b), conditioning on the event that (RU_k) has full column rank in Lemma 4.2, we get with high probability that

$$\|A - P_{\tilde{A},k}(A)\|_2 \leq 2\|A - A_k\|_2 + \frac{\|U_k^\top R^\top R H_k\|_2}{(1-\varepsilon)^2}$$

using the fact that if (RU_k) has full column rank then $(RU_k)^\dagger = ((RU_k)^\top RU_k)^{-1} U_k^\top R^\top$ and $\|((RU_k)^\top RU_k)^{-1}\|_2 \leq 1/(1-\varepsilon)^2$. Now observe that $U_k^\top H_k = 0$. Since $\text{sr}(H_k) \leq k$, using Theorem 3.2 (i.b) with $t = \Omega(k/\varepsilon^4)$, we get that $\|U_k^\top R^\top R H_k\|_2 = \|U_k^\top R^\top R H_k - U_k^\top H_k\|_2 \leq \varepsilon \|U_k\|_2 \|H_k\|_2 = \varepsilon \|A - A_k\|_2$ with high probability. Rescaling ε concludes the proof.

5 Acknowledgments

Many thanks go to Petros Drineas for many helpful discussions and pointing out the connection of Theorem 3.2 with the ℓ_2 -regression problem. The second author would like to thank Mark Rudelson for his valuable comments on an earlier draft and also for sharing with us the proof of Theorem 3.2 (i.b).

References

- [AHK06] S. Arora, E. Hazan, and S. Kale. A Fast Random Sampling Algorithm for Sparsifying Matrices. In *Proceedings of the International Workshop on Randomization and Approximation Techniques (RANDOM)*, pages 272–279, 2006. (Cited on page 2)
- [AM07] D. Achlioptas and F. Mcsherry. Fast Computation of Low-rank Matrix Approximations. *Journal of the ACM (JACM)*, 54(2):9, 2007. (Cited on pages 1, 2 and 5)
- [AW02] R. Ahlswede and A. Winter. Strong Converse for Identification via Quantum Channels. *IEEE Transactions on Information Theory*, 48(3):569–579, 2002. (Cited on pages 1 and 2)
- [Bha96] R. Bhatia. *Matrix Analysis*, volume 169. Graduate Texts in Mathematics, Springer, First edition, 1996. (Cited on pages 3 and 13)
- [BMD09] C. Boutsidis, M. W. Mahoney, and P. Drineas. An Improved Approximation Algorithm for the Column Subset Selection Problem. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 968–977, 2009. (Cited on page 9)
- [Buc01] A. Buchholz. Operator Khintchine Inequality in Non-commutative Probability. *Mathematische Annalen*, 319(1-16):1–16, January 2001. (Cited on pages 3 and 13)
- [Cla08] K. L. Clarkson. Tighter Bounds for Random Projections of Manifolds. In *Proceedings of the ACM Symposium on Computational Geometry (SoCG)*, pages 39–48, 2008. (Cited on page 2)
- [CR07] E. Candès and J. Romberg. Sparsity and Incoherence in Compressive Sampling. *Inverse Problems*, 23(3):969, 2007. (Cited on page 2)
- [CW09] K. L. Clarkson and D. P. Woodruff. Numerical Linear Algebra in the Streaming Model. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 205–214, 2009. (Cited on pages 1, 3 and 5)
- [DK03] P. Drineas and R. Kannan. Pass Efficient Algorithms for Approximating Large Matrices. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 223–232, 2003. (Cited on page 5)
- [DKM06a] P. Drineas, R. Kannan, and M. W. Mahoney. Fast Monte Carlo Algorithms for Matrices I: Approximating Matrix Multiplication. *Journal of the ACM (JACM)*, 36(1):132–157, 2006. (Cited on pages 1, 2 and 3)
- [DKM06b] P. Drineas, R. Kannan, and M. W. Mahoney. Fast Monte Carlo Algorithms for Matrices II: Computing a Low-Rank Approximation to a Matrix. *Journal of the ACM (JACM)*, 36(1):158–183, 2006. (Cited on page 3)
- [DKM06c] P. Drineas, R. Kannan, and M. W. Mahoney. Fast Monte Carlo Algorithms for Matrices III: Computing a Compressed Approximate Matrix Decomposition. *Journal of the ACM (JACM)*, 36(1):184–206, 2006. (Cited on page 3)
- [DM10] P. Drineas and M. W. Mahoney. Effective Resistances, Statistical Leverage, and Applications to Linear Equation Solving. Available at arxiv:1005.3097, May 2010. (Cited on page 6)
- [DMM06] P. Drineas, M. W. Mahoney, and S. Muthukrishnan. Sampling Algorithms for ℓ_2 -regression and Applications. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1127–1136, 2006. (Cited on pages 1 and 5)
- [DMMS09] P. Drineas, M. W. Mahoney, S. Muthukrishnan, and T. Sarlos. Faster Least Squares Approximation. Available at arxiv:0710.1435, May 2009. (Cited on page 5)
- [DR10] A. Deshpande and L. Rademacher. Efficient Volume Sampling for Row/column Subset Selection. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, 2010. (Cited on page 1)
- [DRVW06] A. Deshpande, L. Rademacher, S. Vempala, and G. Wang. Matrix Approximation and Projective Clustering via Volume Sampling. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1117–1126, 2006. (Cited on page 5)
- [DZ10] P. Drineas and A. Zouzias. A Note on Element-wise Matrix Sparsification via Matrix-valued Chernoff Bounds. Available at arxiv:1006.0407, June 2010. (Cited on page 2)
- [FKV04] A. Frieze, R. Kannan, and S. Vempala. Fast Monte-carlo Algorithms for Finding Low-rank Approximations. *Journal of the ACM (JACM)*, 51(6):1025–1041, 2004. (Cited on pages 1, 3 and 5)
- [GLF⁺09] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum State Tomography via Compressed Sensing. Available at 0909.3304, September 2009. (Cited on pages 1 and 4)
- [Gro09] D. Gross. Recovering Low-rank Matrices from Few Coefficients in any Basis. Available at arxiv:0910.1879, December 2009. (Cited on pages 1 and 2)
- [GV96] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Third edition, October 1996. (Cited on pages 3 and 12)
- [HMT09] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding Structure with Randomness: Stochastic Algorithms for Constructing Approximate Matrix Decompositions. Available at arxiv:0909.4061, Sep. 2009. (Cited on pages 5, 6, 7 and 9)
- [LP86] F. Lust-Piquard. Inégalités de Khintchine dans C_p ($1 < p < \infty$). *C. R. Acad. Sci. Paris Sér. I Math.*, 303(7):289–292, 1986. (Cited on pages 3 and 13)
- [LPP91] F. Lust-Piquard and G. Pisier. Non Commutative Khintchine and Paley Inequalities. *Arkiv för Matematik*, 29(1-2):241–260, December 1991. (Cited on page 13)
- [LT91] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer-Verlag, 1991. Isoperimetry and Processes. (Cited on page 13)
- [Mag07] A. Magen. Dimensionality Reductions in ℓ_2 that Preserve Volumes and Distance to Affine Spaces. *Discrete & Computational Geometry*, 38(1):139–153, 2007. (Cited on page 2)
- [Mil71] V.D. Milman. A new Proof of A. Dvoretzky’s The-

orem on Cross-sections of Convex Bodies. *Funkcional. Anal. i Prilozhen.*, 5(4):28–37, 1971. (Cited on page 2)

[NDT09] N. H. Nguyen, T. T. Do, and T. D. Tran. A Fast and Efficient Algorithm for Low-rank Approximation of a Matrix. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 215–224, 2009. (Cited on pages 3, 5, 6, 9 and 13)

[Nem07] A. Nemirovski. Sums of Random Symmetric Matrices and Quadratic Optimization under Orthogonality Constraints. *Mathematical Programming*, 109(2):283–317, 2007. (Cited on page 2)

[Rec09] B. Recht. A Simpler Approach to Matrix Completion. Available at arxiv:0910.0651, October 2009. (Cited on pages 1, 2 and 4)

[RST09] V. Rokhlin, A. Szlam, and M. Tygert. A Randomized Algorithm for Principal Component Analysis. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1100–1124, 2009. (Cited on page 5)

[Rud99] M. Rudelson. Random Vectors in the Isotropic Position. *J. Funct. Anal.*, 164(1):60–72, 1999. (Cited on pages 3 and 13)

[RV07] M. Rudelson and R. Vershynin. Sampling from Large Matrices: An Approach through Geometric Functional Analysis. *Journal of the ACM (JACM)*, 54(4):21, 2007. (Cited on pages 2, 3, 4, 5, 10 and 13)

[RV09] M. Rudelson and R. Vershynin. The Smallest Singular Value of a Random Rectangular Matrix. *Communications on Pure and Applied Mathematics*, 62(1-2):1707–1739, 2009. (Cited on page 9)

[Sar06] T. Sarlos. Improved Approximation Algorithms for Large Matrices via Random Projections. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 143–152, 2006. (Cited on pages 1, 2, 3, 5 and 8)

[So09a] A. Man-Cho So. Improved Approximation Bound for Quadratic Optimization Problems with Orthogonality Constraints. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1201–1209, 2009. (Cited on page 2)

[So09b] A. Man-Cho So. Moment Inequalities for sums of Random Matrices and their Applications in Optimization. *Mathematical Programming*, December 2009. (Cited on page 2)

[SS90] G. W. Stewart and J. G. Sun. *Matrix Perturbation Theory (Computer Science and Scientific Computing)*. Academic Press, June 1990. (Cited on page 3)

[SS08] D. A. Spielman and N. Srivastava. Graph Sparsification by Effective Resistances. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 563–568, 2008. (Cited on pages 6 and 10)

[Tro10] J. A. Tropp. User-Friendly Tail Bounds for Sums of Random Matrices. Available at arxiv:1004.4389, April 2010. (Cited on pages 1 and 4)

[WX08] A. Wigderson and D. Xiao. Derandomizing the Ahlswede-Winter Matrix-valued Chernoff Bound using Pessimistic Estimators, and Applications. *Theory of Computing*, 4(1):53–76, 2008. (Cited on pages 2 and 4)

Appendix

The next lemma states that if a symmetric positive semi-definite matrix \tilde{A} approximates the Rayleigh quotient of a symmetric positive semi-definite matrix A , then the eigenvalues of \tilde{A} also approximate the eigenvalues of A .

LEMMA 5.1. *Let $0 < \varepsilon < 1$. Assume A, \tilde{A} are $n \times n$ symmetric positive semi-definite matrices, such that the following inequality holds*

$$(1 - \varepsilon)x^\top Ax \leq x^\top \tilde{A}x \leq (1 + \varepsilon)x^\top Ax, \quad \forall x \in \mathbb{R}^n.$$

Then, for $i = 1, \dots, n$ the eigenvalues of A and \tilde{A} are the same up-to an error factor ε , i.e.,

$$(1 - \varepsilon)\lambda_i(A) \leq \lambda_i(\tilde{A}) \leq (1 + \varepsilon)\lambda_i(A).$$

Proof. The proof is an immediate consequence of the Courant-Fischer’s characterization of the eigenvalues. First notice that by hypothesis, A and \tilde{A} have the same null space. Hence we can assume without loss of generality, that $\lambda_i(A), \lambda_i(\tilde{A}) > 0$ for all $i = 1, \dots, n$. Let $\lambda_i(A)$ and $\lambda_i(\tilde{A})$ be the eigenvalues (in non-decreasing order) of A and \tilde{A} , respectively. The Courant-Fischer min-max theorem [GV96, p. 394] expresses the eigenvalues as

$$(5.17) \quad \lambda_i(A) = \min_{S^i} \max_{x \in S^i} \frac{x^\top Ax}{x^\top x},$$

where the minimum is over all i -dimensional subspaces S^i . Let the subspaces S_0^i and S_1^i where the minimum is achieved for the eigenvalues of A and \tilde{A} , respectively. Then, it follows that

$$\lambda_i(\tilde{A}) = \min_{S^i} \max_{x \in S^i} \frac{x^\top \tilde{A}x}{x^\top x} \leq \max_{x \in S_0^i} \frac{x^\top \tilde{A}x}{x^\top x} \frac{x^\top Ax}{x^\top Ax} \leq (1 + \varepsilon)\lambda_i(A).$$

and similarly,

$$\lambda_i(A) = \min_{S^i} \max_{x \in S^i} \frac{x^\top Ax}{x^\top x} \leq \max_{x \in S_1^i} \frac{x^\top Ax}{x^\top Ax} \frac{x^\top \tilde{A}x}{x^\top \tilde{A}x} \leq \frac{\lambda_i(\tilde{A})}{1 - \varepsilon}.$$

Therefore, it follows that for $i = 1, \dots, n$,

$$(1 - \varepsilon)\lambda_i(A) \leq \lambda_i(\tilde{A}) \leq (1 + \varepsilon)\lambda_i(A).$$

Proof of Theorem 1.1 For notational convenience, let $Z = \left\| \frac{1}{t} \sum_{i=1}^t M_i - \mathbb{E}M \right\|_2$ and define $E_p := \mathbb{E}_{M_1, M_2, \dots, M_t} Z^p$. Moreover, let X_1, X_2, \dots, X_n be copies of a (matrix-valued) random variables X , we will denote $\mathbb{E}_{X_1, X_2, \dots, X_n}$ by $\mathbb{E}_{X_{[n]}}$. Our goal is to give sharp bounds on the moments of the non-negative random

variable Z and then using the moment method to give concentration result for Z .

First we give a technical lemma of independent interest that bounds the p -th moments of Z as a function of p, r (the rank of the samples), and the $p/2$ -th moment of the random variable $\left\| \sum_{j=1}^t M_j^2 \right\|_2$. More formally, we have the following

LEMMA 5.2. *Let M_1, \dots, M_t be i.i.d. copies of M , where M is a symmetric matrix-valued random variable that has rank at most r almost surely. Then for every $p \geq 2$*

$$(5.18) \quad E_p \leq r t^{1-p} (2B_p)^p \mathbb{E}_{M_{[t]}} \left\| \sum_{j=1}^t M_j^2 \right\|_2^{p/2},$$

where B_p is a constant that depends on p .

We need a non-commutative version of Khintchine inequality due to F. Lust-Piquard [LP86], see also [LPP91] and [Buc01, Theorem 5]. We start with some preliminaries; let $A \in \mathbb{R}^{n \times n}$ and denote by C_p^n the p -th Schatten norm space—the Banach space of linear operators (or matrices in our setting) in \mathbb{R}^n —equipped with the norm

$$(5.19) \quad \|A\|_{C_p^n} := \left(\sum_{i=1}^n \sigma_i(A)^p \right)^{1/p},$$

where $\sigma_i(A)$ are the singular values of A , see [Bha96, Chapter IV, p.92] for a discussion on Schatten norms. Notice that $\|A\|_2 = \sigma_1(A)$, hence we have the following inequality

$$(5.20) \quad \|A\|_2 \leq \|A\|_{C_p^n} \leq (\mathbf{rank}(A))^{1/p} \|A\|_2,$$

for any $p \geq 1$. Notice that when $p = \log_2(\mathbf{rank}(A))$, then $\mathbf{rank}(A)^{1/\log_2(\mathbf{rank}(A))} = 2$. Therefore, in this case, the Schatten norm is essentially the spectral norm. We are now ready to state the matrix-valued Khintchine inequality. See e.g. [Rud99] or [NDT09, Lemma 8].

THEOREM 5.1. *Assume $2 \leq p < \infty$. Then there exists a constant B_p such that for any sequence of t symmetric matrices M_1, \dots, M_t , with $M_i \in C_p^n$, such that the following inequalities hold*

$$(5.21) \quad \left(\mathbb{E}_{\varepsilon_{[t]}} \left\| \sum_{i=1}^t \varepsilon_i M_i \right\|_{C_p^n}^p \right)^{1/p} \leq B_p \left\| \left(\sum_{i=1}^t M_i^2 \right)^{1/2} \right\|_{C_p^n}$$

where for every $i \in [t]$, ε_i is a Bernoulli random variable. Moreover, B_p is at most⁵ $2^{-1/4} \sqrt{\pi/e} \sqrt{p}$.

⁵See Eqn. (17) in [NDT09] or [Buc01].

Now we are ready to prove Lemma 5.2.

Proof. (of Lemma 5.2) The proof is inspired from [RV07, Theorem 3.1]. Let $p \geq 2$. First, apply a standard symmetrization argument (see [LT91]), which gives that

$$\left(\mathbb{E}_{M_{[t]}} \left\| \frac{1}{t} \sum_{i=1}^t M_i - \mathbb{E} M \right\|_2^p \right)^{1/p} \leq 2 \left(\mathbb{E}_{M_{[t]}} \mathbb{E}_{\varepsilon_{[t]}} \left\| \frac{1}{t} \sum_{i=1}^t \varepsilon_i M_i \right\|_2^p \right)^{1/p}.$$

Indeed, let $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_t$ denote independent Bernoulli variables. Let $M_1, \dots, M_t, \widetilde{M}_1, \dots, \widetilde{M}_t$ be independent copies of M . We essential estimate the p -th root of E_p ,

$$(5.22) \quad E_p^{1/p} = \left(\mathbb{E}_{M_{[t]}} \left\| \frac{1}{t} \sum_{i=1}^t M_i - \mathbb{E} M \right\|_2^p \right)^{1/p}$$

Notice that $\mathbb{E} \widetilde{M} = \mathbb{E}_{\widetilde{M}_{[t]}} \left(\frac{1}{t} \sum_{i=1}^t \widetilde{M}_i \right)$. We plug this into (5.22) and apply Jensen's inequality,

$$\begin{aligned} E_p^{1/p} &= \left(\mathbb{E}_{M_{[t]}} \left\| \mathbb{E}_{\widetilde{M}_{[t]}} \frac{1}{t} \sum_{i=1}^t M_i - \frac{1}{t} \sum_{i=1}^t \widetilde{M}_i \right\|_2^p \right)^{1/p} \\ &\leq \left(\mathbb{E}_{M_{[t]}} \mathbb{E}_{\widetilde{M}_{[t]}} \left\| \frac{1}{t} \sum_{i=1}^t M_i - \frac{1}{t} \sum_{i=1}^t \widetilde{M}_i \right\|_2^p \right)^{1/p}. \end{aligned}$$

Now, notice that $M_i - \widetilde{M}_i$ is a symmetric matrix-valued random variable for every $i \in [t]$, i.e., it is distributed identically with $\varepsilon_i(M_i - \widetilde{M}_i)$. Thus

$$E_p^{1/p} \leq \left(\mathbb{E}_{M_{[t]}} \mathbb{E}_{\widetilde{M}_{[t]}} \mathbb{E}_{\varepsilon_{[t]}} \left\| \frac{1}{t} \sum_{i=1}^t \varepsilon_i (M_i - \widetilde{M}_i) \right\|_2^p \right)^{1/p}.$$

Denote $Y = \frac{1}{t} \sum_{i=1}^t \varepsilon_i M_i$ and $\widetilde{Y} = \frac{1}{t} \sum_{i=1}^t \varepsilon_i \widetilde{M}_i$. Then $\|Y - \widetilde{Y}\|^p \leq (\|Y\| + \|\widetilde{Y}\|)^p \leq 2^p (\|Y\|^p + \|\widetilde{Y}\|^p)$, and $\mathbb{E} \|Y\|^p = \mathbb{E} \|\widetilde{Y}\|^p$. Thus, we obtain that

$$(5.23) \quad E_p^{1/p} \leq 2 \left(\mathbb{E}_{M_{[t]}} \mathbb{E}_{\varepsilon_{[t]}} \left\| \frac{1}{t} \sum_{i=1}^t \varepsilon_i M_i \right\|_2^p \right)^{1/p}.$$

Now by the Khintchine's inequality the following holds

for any fixed symmetric matrices M_1, M_2, \dots, M_t .

$$\begin{aligned}
 \left(\mathbb{E}_{\varepsilon_{[t]}} \left\| \frac{1}{t} \sum_{j=1}^t \varepsilon_j M_j \right\|_2^p \right)^{\frac{1}{p}} &\leq \frac{1}{t} \left(\mathbb{E}_{\varepsilon_{[t]}} \left\| \sum_{j=1}^t \varepsilon_j M_j \right\|_{C_p}^p \right)^{\frac{1}{p}} \\
 &\leq \frac{1}{t} B_p \left\| \left(\sum_{j=1}^t M_j^2 \right)^{1/2} \right\|_{C_p} \\
 &\leq \frac{(rt)^{1/p} B_p}{t} \left\| \left(\sum_{j=1}^t M_j^2 \right)^{1/2} \right\|_2 \\
 (5.24) \qquad &= \frac{(rt)^{1/p} B_p}{t} \left\| \sum_{j=1}^t M_j^2 \right\|_2^{\frac{1}{2}},
 \end{aligned}$$

taking $1/t$ outside the expectation and using the left part of Ineq. (5.20), Ineq. (5.21), the right part of Ineq. (5.20) and the fact that the matrix $\left(\sum_{j=1}^t M_j^2\right)^{1/2}$ has rank at most rt .

Now raising Ineq. (5.24) to the p -th power on both sides and then take expectation with respect to M_1, \dots, M_t , it follows from Ineq. (5.23) that

$$E_p \leq 2^p \cdot \frac{rt}{t^p} B_p^p \mathbb{E}_{M_{[t]}} \left\| \sum_{j=1}^t M_j^2 \right\|_2^{p/2}.$$

This concludes the proof of Lemma 5.2.

Now we are ready to prove Theorem 1.1. First we can assume without loss of generality that $M \succeq 0$ almost surely losing only a constant factor in our bounds. Indeed, by the spectral decomposition theorem any symmetric matrix can be written as $M = \sum_j \lambda_j u_j u_j^\top$. Set $M_+ = \sum_{\lambda_j \geq 0} \lambda_j u_j u_j^\top$ and $M_- = M - M_+$. It is clear that $\|M_+\|_2, \|M_-\|_2 \leq \|M\|_2, \|M_+\|_F, \|M_-\|_F \leq \|M\|_F$ and $\mathbf{rank}(M_+), \mathbf{rank}(M_-) \leq \mathbf{rank}(M)$. Triangle inequality tells us that

$$\begin{aligned}
 \left\| \frac{1}{t} \sum_{i=1}^t M_j - \mathbb{E} M \right\|_2 &\leq \left\| \frac{1}{t} \sum_{i=1}^t (M_j)_+ - \mathbb{E} M_+ \right\|_2 \\
 &\quad + \left\| \frac{1}{t} \sum_{i=1}^t (M_j)_- - \mathbb{E} M_- \right\|_2
 \end{aligned}$$

and one can bound each term of the right hand side separately. Hence, from now on we assume that $M \succeq 0$ a.s.. Now use the fact that for every $j \in [t]$, $M_j^2 \preceq \gamma \cdot M_j$ since M_j 's are positive semi-definite and $\|M\|_2 \leq \gamma$

almost surely. Summing up all the inequalities we get that

$$(5.25) \qquad \left\| \sum_{j=1}^t M_j^2 \right\|_2 \leq \gamma \left\| \sum_{j=1}^t M_j \right\|_2.$$

It follows that

$$\begin{aligned}
 E_p &\leq rt^{1-p} (2B_p)^p \mathbb{E}_{M_{[t]}} \left\| \sum_{j=1}^t M_j^2 \right\|_2^{p/2} \\
 &\leq rt^{1-p} (2B_p)^p \gamma^{p/2} \mathbb{E}_{M_{[t]}} \left\| \sum_{j=1}^t M_j \right\|_2^{p/2} \\
 &= \frac{rt(2B_p\sqrt{\gamma})^p}{t^{p/2}} \mathbb{E}_{M_{[t]}} \left\| \frac{1}{t} \sum_{j=1}^t M_j \right\|_2^{p/2} \\
 &= \frac{rt(2B_p\sqrt{\gamma})^p}{t^{p/2}} \mathbb{E}_{M_{[t]}} \left\| \frac{1}{t} \sum_{j=1}^t M_j - \mathbb{E} M + \mathbb{E} M \right\|_2^{p/2} \\
 &\leq \frac{rt(2B_p\sqrt{\gamma})^p}{t^{p/2}} \left(\left(\mathbb{E} \left\| \frac{1}{t} \sum_{j=1}^t M_j - \mathbb{E} M \right\|_2^{\frac{2}{p}} \right)^{\frac{p}{2}} + 1 \right)^{\frac{p}{2}} \\
 &\leq \frac{rt(2B_p\sqrt{\gamma})^p}{t^{p/2}} \left(\left(\mathbb{E} \left\| \frac{1}{t} \sum_{j=1}^t M_j - \mathbb{E} M \right\|_2^p \right)^{\frac{1}{p}} + 1 \right)^{\frac{p}{2}} \\
 &= \frac{rt(2B_p\sqrt{\gamma})^p}{t^{p/2}} \left(E_p^{1/p} + 1 \right)^{p/2},
 \end{aligned}$$

using Lemma 5.2, Ineq. (5.25), Minkowski's inequality, Jensen's inequality, definition of E_p and the assumption $\|\mathbb{E} M\|_2 \leq 1$. This implies the following inequality

$$(5.26) \qquad E_p^{1/p} \leq \frac{2B_p\sqrt{\gamma}(rt)^{1/p}}{\sqrt{t}} (E_p^{1/p} + 1),$$

using that $\sqrt{1+x} \leq 1+x, x \geq 0$. Let $a_p = \frac{4B_p\sqrt{\gamma}(rt)^{1/p}}{\sqrt{t}}$.

Then it follows from the above inequality that $E_p^{1/p} \leq \frac{a_p}{2} (E_p^{1/p} + 1)$. It follows that⁶ $\min\{E_p^{1/p}, 1\} \leq a_p$. Also notice that

$$(5.27) \qquad (\mathbb{E} \min\{Z, 1\}^p)^{1/p} \leq \min(E_p^{1/p}, 1).$$

Now for any $0 < \varepsilon < 1$,

$$\mathbb{P}(Z > \varepsilon) = \mathbb{P}(\min\{Z, 1\} > \varepsilon).$$

⁶Indeed, if $E_p^{1/p} < 1$, then $E_p^{1/p} < a_p$. Otherwise $1 \leq a_p$.

By the moment method we have that

$$\begin{aligned}
\mathbb{P}(\min\{Z, 1\} > \varepsilon) &= \mathbb{P}(\min\{Z, 1\}^p > \varepsilon^p) \\
&\leq \inf_{p \geq 2} \left(\frac{\mathbb{E} \min\{Z, 1\}^p}{\varepsilon^p} \right) \\
&\leq \inf_{p \geq 2} \left(\frac{\min\{E_p^{1/p}, 1\}^p}{\varepsilon^p} \right) \quad (5.27) \\
&\leq \inf_{p \geq 2} \left(\frac{a_p}{\varepsilon} \right)^p \\
&= \inf_{p \geq 2} \left(\frac{4B_p \sqrt{\gamma}(rt)^{1/p}}{\varepsilon \sqrt{t}} \right)^p \\
&= \inf_{p \geq 2} \left(C_2 \frac{\sqrt{p\gamma}(rt)^{1/p}}{\varepsilon \sqrt{t}} \right)^p,
\end{aligned}$$

where $C_2 > 0$ is an absolute constant.

Now assume that $r \leq t$ and then set $p = c_2 \log t$, where $c_2 > 0$ is a sufficient large constant, at the infimum expression in the above inequality, it follows that

$$\mathbb{P} \left(\left\| \frac{1}{t} \sum_{i=1}^t M_i - \mathbb{E} M \right\|_2 > \varepsilon \right) \leq \left(C \frac{\sqrt{\gamma \log t} (rt)^{\frac{1}{\log t}}}{\varepsilon \sqrt{t}} \right)^{c_2 \log t}$$

We want to make the base of the above exponent smaller than one. It is easy to see that this is possible if we set $t = C_0 \gamma / \varepsilon^2 \log(C_0 \gamma / \varepsilon^2)$ where C_0 is sufficiently large absolute constant. Hence it implies that the above probability is at most $1/\text{poly}(t)$. This concludes the proof.