Chapter 5

Low-rank approximation with subsampled unitary transformations

5.1 Introduction

In this chapter, we analyze the theoretical performance of a randomized low-rank approximation algorithm introduced in [WLRT08] and analyzed in [WLRT08, HMT11, NDT09]. Our analysis often provides sharper approximation bounds than those in [WLRT08, HMT11, NDT09]. We provide bounds on the residual and forward errors of this approximation algorithm in the spectral and Frobenius norms, and provide experimental evidence that this low-rank approximation algorithm performs as well as a more expensive low-rank approximation algorithm based upon projections onto uniformly distributed random subspaces[1]. Further, we provide approximation bounds for a variant of the algorithm that returns approximations with even lower rank.

The setting is as follows: fix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a target rank $k \leq \min\{m, n\}$. We would like to approximate $\mathbf{A}$ with a matrix $\mathbf{X}$ that has rank close to k, and we would like $\|\mathbf{A} - \mathbf{X}\|_\xi$ to be within a small multiplicative factor of the smallest error achievable when approximating $\mathbf{A}_k$ with a rank-k matrix, for $\xi = 2, F$.

It is well-known that the rank-k matrix $\mathbf{A}_k$ that minimizes both the Frobenius and the

[1]The content of this chapter is adapted from the article [GB12] co-authored with Christos Boutsidis.

spectral-norm approximation errors can be calculated using the singular value decomposition (SVD) in $O(mn \min\{m,n\})$ arithmetic operations, using classical so-called *direct* algorithms such as QR iteration or Jacobi iteration [GV96]. Computing the full SVD is expensive when $\mathbf{A}$ is a large matrix. In this case, it is often more efficient to use *iterative* projection methods (e.g. Krylov subspace methods) to obtain approximations to $\mathbf{A}_k$. It is difficult to state a precise guarantee for the number of arithmetic operations carried out by Krylov methods, but one iteration of a Krylov method requires $\Omega(mnk)$ operations (assuming $\mathbf{A}$ has no special structure which can be exploited to speed up the computation of matrix–vector products). To obtain even an accurate rank-1 approximation requires $O(\log n)$ iterations [KW92]. Thus, an optimistic estimate for the number of operations required to compute approximate rank-k truncated SVDs using a Krylov method is $\Omega(mnk \log n)$.

Our discussion thus far has concerned only the arithmetic cost of computing truncated SVDs, but an equally or more important issue is that of the communication costs: bandwidth costs (proportional to the amount of times storage is accessed) and latency costs (proportional to the cost of transferring the information over a network or through the levels of a hierarchical memory system) [BDHS11]. If the algorithm is to be parallelized, then the complexity of the required information interchange must also be taken into account.

The randomized algorithms considered in this chapter, Algorithms 5.1 and 5.2, are of interest because they yield low-rank approximations after $\Omega(mnk \max\{\log n, \log k\})$ arithmetic operations and have low communication costs. In particular, each element of $\mathbf{A}$ is accessed only twice, and the algorithms are simple enough that they are amenable to straightforward parallelization. The guarantees provided are probabilistic, and allow one to trade off between the operation count of the algorithms and the accuracy and failure probabilities of the algorithms.

Both of the algorithms considered in this chapter are based on the intuition that, when

Algorithm 5.1: Randomized approximate truncated SVD

Input: an $m \times n$ matrix $\mathbf{A}$ and an $n \times \ell$ matrix $\mathbf{S}$, where ℓ is an integer in $[1, n]$.
Output: matrices $\tilde{\mathbf{U}}, \tilde{\mathbf{\Sigma}}, \tilde{\mathbf{V}}$ constituting the SVD of $\mathbf{P}_{\mathbf{AS}}\mathbf{A} = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.

1: Let $\mathbf{Y} = \mathbf{AS}$.
2: Compute the QR decomposition $\mathbf{Y} = \mathbf{QR}$.
3: Compute the SVD of $\mathbf{Q}^T\mathbf{A} = \mathbf{W}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.
4: Set $\tilde{\mathbf{U}} = \mathbf{QW}$.

Algorithm 5.2: Rank-k randomized approximate truncated SVD

Input: an $m \times n$ matrix $\mathbf{A}$, integers ℓ and k that satisfy $\ell > k$ and $k \in [1, n]$, and an $n \times \ell$ matrix $\mathbf{S}$.
Output: matrices $\tilde{\mathbf{U}}, \tilde{\mathbf{\Sigma}}, \tilde{\mathbf{V}}$ constituting the SVD of $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A}) = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.

1: Let $\mathbf{Y} = \mathbf{AS}$.
2: Compute the QR decomposition $\mathbf{Y} = \mathbf{QR}$.
3: Compute the rank-k truncated SVD of $\mathbf{Q}^T\mathbf{A}$ to obtain $(\mathbf{Q}^T\mathbf{A})_k = \mathbf{W}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.
4: Set $\tilde{\mathbf{U}} = \mathbf{QW}$.

$\mathbf{S} \in \mathbb{R}^{n \times \ell}$ is randomly selected and ℓ is sufficiently larger than k, the range of the matrix $\mathbf{AS}$ "captures" the top k-dimensional left singular space of $\mathbf{A}$. When this phenomenon occurs, the low-rank matrix formed by projecting $\mathbf{A}$ onto the range of $\mathbf{AS}$ should be almost as accurate an approximation of $\mathbf{A}$ as is the optimal approximation $\mathbf{A}_k$:

$$\|\mathbf{A} - \mathbf{P}_{\mathbf{AS}}\mathbf{A}\|_\xi \approx \|\mathbf{A} - \mathbf{A}_k\|_\xi \quad \text{for } \xi = 2, \mathrm{F}.$$

Algorithm 5.1 computes exactly this approximation, $\mathbf{P}_{\mathbf{AS}}\mathbf{A}$. Note that this approximation may have rank up to ℓ, which may be much larger than k. Algorithm 5.2 instead returns the approximation $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})$, which is guaranteed to have rank at most k.

Unlike classical iterative methods for approximating the trucated SVD, which use as many iterations as necessary to satisfy some convergence condition, Algorithms 5.1 and 5.2 use only one matrix–matrix product $\mathbf{AS}$ to generate an approximate basis for the top left singular space of $\mathbf{A}$. Accordingly, the quality of approximations obtained using either of these algorithms is

more dependent on the properties of $\mathbf{A}$ itself and the sampling matrix $\mathbf{S}$ than is the quality of approximations derived from classical iterative methods. Thus it is important to supply theoretical guarantees on the errors of the algorithms that identify which properties of $\mathbf{A}$ affect the quality of the approximations, as well as to carry to empirical studies investigating the influence of the choice of $\mathbf{S}$.

Recent years have produced a large body of research on designing random sampling matrices $\mathbf{S}$. Some proposals for $\mathbf{S}$ include: (i) every entry of $\mathbf{S}$ takes the values $+1, -1$ with equal probability [CW09, MZ11]; (ii) the entries of $\mathbf{S}$ are i.i.d. Gaussian random variables with zero mean and unit variance [HMT11]; (iii) the columns of $\mathbf{S}$ are chosen independently from the columns of the $m \times m$ identity matrix with probabilities that are proportional to the Euclidean length of the columns of $\mathbf{A}$ [FKV98, DKM06b]; and (iv) $\mathbf{S}$ is designed carefully such that $\mathbf{AS}$ can be computed in at most $O(\mathrm{n}nz(\mathbf{A}))$ arithmetic operations, where $\mathrm{n}nz(\mathbf{A})$ denotes the number of non-zero entries in $\mathbf{A}$ [CW12].

In this chapter we take $\mathbf{S}$ to be a subsampled randomized Hadamard transform (SRHT) matrix, i.e. $\mathbf{S}$ comprises a subset of the columns of a randomized Hadamard matrix (see Definitions 5.1 and 5.2 below). This choice for $\mathbf{S}$ was introduced in [AC06].

Definition 5.1 (Normalized Walsh–Hadamard Matrix)**.** *Fix an integer $n = 2^p$, for $p = 1, 2, 3, \ldots$. The (non-normalized) $n \times n$ matrix of the Hadamard-Walsh transform is defined recursively as,*

$$\mathbf{H}_n = \begin{bmatrix} \mathbf{H}_{n/2} & \mathbf{H}_{n/2} \\ \mathbf{H}_{n/2} & -\mathbf{H}_{n/2} \end{bmatrix}, \qquad \textit{with} \qquad \mathbf{H}_2 = \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix}.$$

The $n \times n$ normalized matrix of the Walsh–Hadamard transform is equal to $\mathbf{H} = n^{-\frac{1}{2}}\mathbf{H}_n \in \mathbb{R}^{n \times n}$.

Definition 5.2 (Subsampled Randomized Hadamard Transform (SRHT) matrix)**.** *Fix integers ℓ*

and $n = 2^p$ *with* $\ell < n$ *and* $p = 1, 2, 3, \ldots$. *An SRHT matrix is an* $\ell \times n$ *matrix of the form*

$$\boldsymbol{\Theta} = \sqrt{\frac{n}{\ell}} \cdot \mathbf{RHD};$$

- $\mathbf{D} \in \mathbb{R}^{n \times n}$ *is a random diagonal matrix whose entries are independent random signs, i.e. random variables uniformly distributed on* $\{\pm 1\}$.
- $\mathbf{H} \in \mathbb{R}^{n \times n}$ *is a normalized Walsh–Hadamard matrix.*
- $\mathbf{R} \in \mathbb{R}^{\ell \times n}$ *is a subset of* ℓ *rows from the* $n \times n$ *identity matrix, where the rows are chosen uniformly at random and without replacement.*

The choice of $\mathbf{S}$ as an SRHT matrix is particularly practical because the highly structured nature of $\mathbf{S}$ can be exploited to reduce the time of computing $\mathbf{AS}$ from $\mathrm{O}(mn\ell)$ to $\mathrm{O}(mn \log_2 \ell)$.

Lemma 5.3 (Fast Matrix–Vector Multiplication, Theorem 2.1 in [AL08])**.** *Given* $\mathbf{x} \in \mathbb{R}^n$ *and* $\ell < n$, *one can construct* $\boldsymbol{\Theta} \in \mathbb{R}^{\ell \times n}$ *and compute* $\boldsymbol{\Theta}\mathbf{x}$ *in at most* $2n \log_2(\ell + 1))$ *operations.*

Beyond the SRHT. The SRHT is defined only when the matrix dimension is a power of two. An alternative option is to use other structured orthonormal randomized transforms such as the real Fourier transform (DFT), the discrete cosine transform (DCT) or the discrete Hartley transform (DHT) [WLRT08, NDT09, RT08, AMT10], whose entries are on the order of $n^{-1/2}$. None of these transforms place restrictions on the size of the matrix being approximated. With minimal effort, the results of this chapter can be extended to encompass these transforms. Specifically, the statements of Lemma 5.5 and Lemma 5.8 in this chapter would need to be modified slightly to account for the difference in the transform; essentially, the constants present in the statements of the Lemmas would change. These two lemmas isolate the effects of the particular choice of

S from the remainder of the arguments used in this chapter, so the changes would propagate, *mutatis mutandis*, throughout the remaining results in this chapter.

We note, further, that Algorithms 5.1 and 5.2 can be modified to use $\mathbf{Y} = (\mathbf{A}\mathbf{A}^T)^p\mathbf{A}\mathbf{S}$, where $p \geq 1$ is an integer, as an approximate basis for the top left singular space of **A**. Approximations to $\mathbf{A}_k$ generated using this choice of **Y** are more accurate than those generated using our choice of **AS**, but one loses the speed conferred by taking **S** to be an SRHT matrix: after the first multiplication **AS**, all the matrix multiplications required to form **Y** are dense and unstructured.

Outline. In Section 5.2, we present a portion of our results on the quality of SRHT low-rank approximations and compare them to prior results in the literature. Section 5.3 presents new results on the application of SRHTs to general matrices and the approximation of matrix multiplication using SRHTs under the Frobenius norm. Section 5.4 contains the statements and proofs of our main results. We conclude the chapter with an experimental evaluation of the SRHT low-rank approximation algorithms in Section 5.5.

5.2 Low-rank matrix approximation using SRHTs

Using an SRHT matrix (see Definition 5.2), one can quickly construct low-rank approximations of a given matrix **A** using Algorithms 5.1 and 5.2. Our main results, Theorems 5.13 and 5.14, given in Section 5.4, respectively provide theoretical guarantees on the spectral and Frobenius-norm residual and forward errors of these approximations. To facilitate the comparison of our results with prior work, we highlight our residual error guarantees for Algorithm 5.1.

Theorem 5.4. *Assume n is a power of 2. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ have rank ρ and fix an integer k satisfying $2 \leq k < \rho$. Let $0 < \epsilon < 1/3$ be an accuracy parameter, $0 < \delta < 1$ be a failure probability, and*

where the notation Avg$(\cdot)$ indicates the average of a quantity over all the entries of $\mathbf{A}$.

Their choice of p_{jk}, in particular the insertion of the $(8\log n)^4/n$ factor, is an artifact of their method of proof. Instead, we consider a scheme which compares the magnitudes of a_{jk} and b to determine p_{jk}. Introduce the quantity $R = \max_{a_{jk}\neq 0} b/|a_{jk}|$ to measure the spread of the entries in $\mathbf{A}$, and take

$$X_{jk} \sim \begin{cases} \frac{a_{jk}}{p_{jk}} \text{ Bern}(p_{jk}), \text{ where } p_{jk} = \frac{pa_{jk}^2}{pa_{jk}^2+b^2}, & a_{jk} \neq 0 \\ 0, & a_{jk} = 0. \end{cases}$$

With this scheme, $\text{Var}(X_{jk}) = 0$ when $a_{jk} = 0$, otherwise $\text{Var}(X_{jk}) = b^2/p$. Likewise, $\mathbb{E}(X_{jk} - a_{jk})^4 = 0$ if $a_{jk} = 0$, otherwise

$$\mathbb{E}(X_{jk} - a_{jk})^4 \leq \text{Var}(X_{jk}) \left\|X_{jk} - a_{jk}\right\|_\infty^2 = \frac{b^2}{p} \max\left\{|a_{jk}|, |a_{jk}|\left(\frac{pa_{jk}^2 + b^2}{pa_{jk}^2} - 1\right)\right\}^2 \leq \frac{b^4}{p^2}R^2,$$

so

$$\mathbb{E}\,\|\mathbf{A} - \mathbf{X}\| \leq \text{C}\left(b\sqrt{\frac{n}{p}} + b\sqrt{\frac{m}{p}} + b\sqrt{\frac{R}{p}}\sqrt[4]{mn}\right) \leq \text{C}(2+\sqrt{R})b\sqrt{\frac{n}{p}}.$$

Applying Corollary 3.16, we find that the error satisfies

$$\mathbb{P}\left\{\|\mathbf{A} - \mathbf{X}\| > \text{C}(2+\sqrt{R})b\sqrt{\frac{n}{p}}(\epsilon+1)\right\} \leq e^{-\epsilon^2\text{C}^2(2+\sqrt{R})^2pn/16},$$

with probability at least $1 - \exp(-\text{C}^2(2+\sqrt{R})^2pn/16)$,

$$\|\mathbf{A} - \mathbf{X}\| \leq 2\text{C}(2+\sqrt{R})b\sqrt{\frac{n}{p}}.$$

Thus, Theorem 3.14 and Achlioptas and McSherry's scheme-specific analysis yield results of the same order in n and p. As before, we see that our bound holds with higher probability and over

a larger range of n. Furthermore, since the expected number of nonzero entries in $\mathbf{X}$ satisfies

$$\sum_{jk} p_{jk} = \sum_{jk} \frac{pa_{jk}^2}{pa_{jk}^2 + b^2} \leq pnm \times \text{Avg}\left[\left(\frac{a_{jk}}{b}\right)^2\right],$$

we have established a smaller limit on the expected number of nonzero entries.

3.6.1.3 A scheme which simultaneously sparsifies and quantizes

Finally, we use Theorem 3.15 to estimate the error of the scheme from [AHK06] which simultaneously quantizes and sparsifies. Fix $\delta > 0$ and consider

$$X_{jk} = \begin{cases} \text{sgn}(a_{jk}) \frac{\delta}{\sqrt{n}} \text{ Bern}\left(\frac{|a_{jk}|\sqrt{n}}{\delta}\right), & |a_{jk}| \leq \frac{\delta}{\sqrt{n}} \\ a_{jk}, & \text{otherwise.} \end{cases}$$

Then $\text{Var}(X_{jk}) = 0$ if $|a_{jk}| \geq \delta/\sqrt{n}$, otherwise

$$\text{Var}(X_{jk}) = |a_{jk}|^3 \frac{\sqrt{n}}{\delta} - 2a_{jk}^2 + |a_{jk}| \frac{\delta}{\sqrt{n}} \leq \frac{\delta^2}{n}.$$

The fourth moment term is zero when $|a_{jk}| \geq \delta/\sqrt{n}$, and when $|a_{jk}| < \delta/\sqrt{n}$,

$$\mathbb{E}(X_{jk} - a_{jk})^4 = |a_{jk}|^5 \frac{\sqrt{n}}{\delta} - 4a_{jk}^4 + 6|a_{jk}|^3 \frac{\delta}{\sqrt{n}} - 4a_{jk}^2 \frac{\delta^2}{n} + |a_{jk}| \left(\frac{\delta}{\sqrt{n}}\right)^3 \leq 8\frac{\delta^4}{n^2}.$$

This gives the estimates

$$\mathbb{E}\,\|\mathbf{A} - \mathbf{X}\| \leq C\left(\sqrt{n}\frac{\delta}{\sqrt{n}} + \sqrt{m}\frac{\delta}{\sqrt{n}} + 2\frac{\delta}{\sqrt{n}}\sqrt[4]{mn}\right) \leq 4C\delta$$

and

$$\mathbb{P}\left\{\|\mathbf{A}-\mathbf{X}\| > 4C\delta(\gamma+1)\right\} \leq \mathrm{e}^{-\gamma^2 C^2 n}.$$

Taking $\gamma = 1$, we see that with probability at least $1 - \exp(-C^2 n)$,

$$\|\mathbf{A}-\mathbf{X}\| \leq 8C\delta.$$

Let $S = \sum_{j,k} |A_{jk}|$, then appealing to Lemma 1 in [AHK06], we find that $\mathbf{X}$ has $\mathrm{O}\left(\frac{\sqrt{n}S}{\gamma}\right)$ nonzero entries with probability at least $1 - \exp\left(-\Omega\left(\frac{\sqrt{n}S}{\gamma}\right)\right)$.

Arora, Hazan, and Kale establish that this scheme guarantees $\|\mathbf{A}-\mathbf{X}\| = \mathrm{O}(\delta)$ with probability at least $1 - \exp(-\Omega(n))$, so we see that our general bound recovers a bound of the same order.

3.7 Comparison with later bounds

The papers [NDT10, DZ11, AKL13], written after the results in this chapter were obtained, present alternative schemes for sparsification and quantization.

The scheme presented in [NDT10] sparsifies a matrix by zeroing out all sufficiently small entries of $\mathbf{A}$, keeping all sufficiently large entries, and randomly sampling the remaining entries of the matrix with a probability depending on their magnitudes. More precisely, given a parameter $s > 0$, it generates an approximation whose entries are distributed as

$$X_{jk} = \begin{cases} 0, & a_{jk}^2 \leq (\log^2(n)/n)\|\mathbf{A}\|_{\mathrm{F}}^2/s \\ a_{jk} & a_{jk}^2 \geq \|\mathbf{A}\|_{\mathrm{F}}^2/s \\ (a_{jk}/p_{jk})\mathrm{Bern}(p_{jk}), & \text{otherwise, where } p_{jk} = s a_{jk}^2/\|\mathbf{A}\|_{\mathrm{F}}^2. \end{cases}$$

The analysis offered guarantees that if $s = \Omega(\epsilon^{-2} n \log^3 n)$, then with probability at least $1 - n^{-1}$, $\|\mathbf{A} - \mathbf{X}\|_2 \leq \epsilon$ and, in expectation, $\mathbf{X}$ has less than $2s$ nonzero entries. It is not clear whether or not this scheme can be analyzed using Theorem 3.14. It is straightforward to establish that $\mathrm{Var}(X_{jk}) \leq \epsilon^2/(n \log^3 n)$ for this scheme, but obtaining a sufficiently small upper bound on the fourth moment $\mathbb{E}(X_{jk} - a_{jk})^4$ is challenging. In particular, the estimate

$$\mathbb{E}(X_{jk} - a_{jk})^4 \leq \mathrm{Var}(X_{jk}) \|X_{jk} - a_{jk}\|_\infty$$

gives an upper bound on the order of $\epsilon a_{jk}^2 n / \log^5 n$, which is sufficient only to establish a much weaker guarantee on the error $\mathbb{E}\|\mathbf{A} - \mathbf{X}\|_2$ than the guarantee given in [NDT10].

The scheme introduced in [DZ11] first zeroes out all entries of $\mathbf{A} \in \mathbb{R}^{n \times n}$ of sufficiently small magnitude, then samples elements from $\mathbf{A}$ in s i.i.d. trials with replacement. The elements are selected with probabilities proportional to their squared magnitudes. Thus, the approximant can be written in the form

$$\mathbf{X} = \frac{1}{s} \sum_{t=1}^{s} \frac{a_{j_t k_t}}{p_{j_t k_t}} \boldsymbol{e}_{j_t} \boldsymbol{e}_{k_t}^T,$$

where (j_t, k_t) is the index of the element of $\mathbf{A}$ selected in the tth trial, $p_{jk} = a_{jk}^2 / \|\mathbf{A}\|_{\mathrm{F}}^2$ is the probability that the entry a_{jk} is selected, and $\boldsymbol{e}_j$ denotes the jth standard basis vector in n. Clearly $\mathbf{X}$ has at most s nonzero entries. Let $s = \Omega(\epsilon^{-2} n \log(n) \|\mathbf{A}\|_{\mathrm{F}}^2)$. Then the authors show that, with probability at least $1 - n^{-1}$, the error of the approximation satisfies $\|\mathbf{A} - \mathbf{X}\|_2 \leq \epsilon$. This scheme is not easily analyzable using our Theorem 3.14. Since the approximant $\mathbf{X}$ is a sum of rank-one matrices, it is most natural to analyze its approximation error using tail bounds for sums of independent random matrices. Indeed, the authors of [DZ11] use a matrix Bernstein inequality to provide their results.

Finally, the scheme presented in [AKL13] computes an approximation of the same form as

the scheme introduced in [DZ11], but samples entries of $\mathbf{A}$ with probabilities proportional their absolute values. That is,

$$\mathbf{X} = \frac{1}{s} \sum_{t=1}^{s} \frac{a_{j_t k_t}}{p_{j_t k_t}} \boldsymbol{e}_{j_t} \boldsymbol{e}_{k_t}^T,$$

where $p_{jk} = |a_{jk}| / \sum_{pq} |a_{pq}|$. Again, this scheme is not amenable to analysis using Theorem 3.14. Recall that $\mathbf{A}^{(k)}$ denotes the kth row of $\mathbf{A}$. The authors establish that, when

$$s = \Omega\left(\epsilon^{-2} \log(n/\delta) \left(\sum_{jk} |A_{jk}|\right) \max_k \|\mathbf{A}^{(k)}\|_1\right),$$

the error bound $\|\mathbf{A} - \mathbf{X}\|_2 \leq \epsilon$ is satisfied with probability at least $1 - \delta$. The approximant $\mathbf{X}$ has, in expectation, at most $2s$ nonzero entries.

Comparing the extents to which we were able to reproduce the guarantees of the sparsification schemes introduced in [AM01, AHK06, AM07, NDT10, DZ11, AKL13], we see that Theorem 3.14 sometimes can recover competitive guarantees on the approximation errors of element-wise sparsification schemes in which X_{jk} is directly related to a_{jk} through a simple expression. When $\mathbf{X}$ is more naturally represented as a sum of rank-1 matrices, Theorem 3.14 is not easily applicable.

Chapter 4

Preliminaries for the investigation of low-rank approximation algorithms

This chapter consolidates probabilistic and linear algebraic tools used in Chapters 5 and 6. We also establish two lemmas of independent interest: the first, Lemma 4.3, is an exponential tail bound on the Frobenius-norm error incurred when approximating the product of two matrices using randomized column and row sampling without replacement; the second, Lemma 4.9, is a deterministic bound on the forward errors of column-based low-rank approximations.

4.1 Probabilistic tools

In this section, we review several tools that are used to deal with random matrices and more generally, random processes.

4.1.1 Concentration of convex functions of Rademacher variables

Rademacher random variables take the values ± 1 with equal probability. Rademacher vectors are vectors of i.i.d. Rademacher random variables. Rademacher vectors often play a crucial role in the construction of dimension reduction maps, an area where the strong measure concentration properties of Rademacher sums are often exploited. The following result states a large-deviation

property of convex Lipschitz functions of Rademacher vectors: namely, these functions tend to be not much larger than their expectations.

Lemma 4.1 (A large deviation result for convex Lipschitz functions of Rademacher random variables [Corollary 1.3 ff. in [Led96]])**.** *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is a convex function that satisfies the Lipschitz bound*

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq L \left\|\mathbf{x} - \mathbf{y}\right\|_2 \quad \text{for all } \mathbf{x}, \mathbf{y}.$$

Let $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ be a Rademacher vector. For all $t \geq 0$,

$$\mathbb{P}\left\{f(\boldsymbol{\varepsilon}) \geq \mathbb{E}\left[f(\boldsymbol{\varepsilon})\right] + Lt\right\} \leq \mathrm{e}^{-t^2/8}.$$

4.1.2 Chernoff bounds for sums of random matrices sampled without replacement

Classical Chernoff bounds provide tail bounds for sums of nonnegative random variables. Their matrix analogs provide tail bounds on the eigenvalues and singular values of sums of positive-semidefinite random matrices. Matrix Chernoff bounds are particularly useful for analyzing algorithms involving randomized column-sampling. Most matrix Chernoff bounds available in the literature require the summands to be independent. Indeed, the Chernoff bounds developed in Chapter 2 bound the eigenvalues of a sum of independent random Hermitian matrices. However, occasionally one desires Chernoff bounds that do not require the summands to be independent. The following Chernoff bounds are useful in the case where the summands are drawn without replacement from a set of bounded random matrices.

Lemma 4.2 (Matrix Chernoff Bounds, Theorem 2.2 in [Tro11b])**.** *Let $\mathscr{X}$ be a finite set of*

positive-semidefinite matrices with dimension k, and suppose that

$$\max_{\mathbf{X} \in \mathcal{X}} \lambda_{\max}(\mathbf{X}) \leq B.$$

Sample $\{\mathbf{X}_1, \dots, \mathbf{X}_\ell\}$ uniformly at random from $\mathcal{X}$ without replacement. Compute

$$\mu_{\max} = \ell \cdot \lambda_1(\mathbb{E}\mathbf{X}_1) \quad \text{and} \quad \mu_{min} = \ell \cdot \lambda_k(\mathbb{E}\mathbf{X}_1).$$

Then

$$\mathbb{P}\left\{\lambda_1\left(\sum\nolimits_j \mathbf{X}_j\right) \geq (1+\nu)\mu_{\max}\right\} \leq k \cdot \left[\frac{\mathrm{e}^{\nu}}{(1+\nu)^{1+\nu}}\right]^{\mu_{\max}/B} \quad \text{for } \nu \geq 0, \text{ and}$$
$$\mathbb{P}\left\{\lambda_k\left(\sum\nolimits_j \mathbf{X}_j\right) \leq (1-\nu)\mu_{\min}\right\} \leq k \cdot \left[\frac{\mathrm{e}^{-\nu}}{(1-\nu)^{1-\nu}}\right]^{\mu_{\min}/B} \quad \text{for } \nu \in [0,1).$$

We also use the following standard simplification of the lower Chernoff bound, which holds under the setup of Lemma 4.2:

$$\mathbb{P}\left\{\lambda_k\left(\sum\nolimits_j \mathbf{X}_j\right) \leq \varepsilon\mu_{\min}\right\} \leq k \cdot \mathrm{e}^{-(1-\varepsilon)^2\mu_{\min}/(2B)} \quad \text{for } \varepsilon \in [0,1]. \tag{4.1.1}$$

4.1.3 Frobenius-norm error bounds for matrix multiplication

We now establish a tail bound on the Frobenius-norm error of a simple approximate matrix multiplication scheme based upon randomized column and row sampling. This simple approximate multiplication scheme is a staple in randomized numerical linear algebra, and variants have been analyzed multiple times [DK01, DKM06a, Sar06]. The result derived here differs in that it applies to the sampling without replacement model, and it provides bounds on the error that hold with high probability, rather than simply an estimate of the expected error.

Lemma 4.3 (Matrix Multiplication)**.** *Let* $\mathbf{X} \in \mathbb{R}^{m\times n}$ *and* $\mathbf{Y} \in \mathbb{R}^{n\times p}$*. Fix* $\ell \leq n$*. Select uniformly at random and without replacement* ℓ *columns from* $\mathbf{X}$ *and the corresponding rows from* $\mathbf{Y}$ *and multiply the selected columns and rows with* $\sqrt{n/\ell}$*. Let* $\hat{\mathbf{X}} \in \mathbb{R}^{m\times \ell}$ *and* $\hat{\mathbf{Y}} \in \mathbb{R}^{\ell\times p}$ *contain the scaled columns and rows, respectively. Choose*

$$\sigma^2 \geq \frac{4n}{\ell}\sum\nolimits_{i=1}^{n} \|\mathbf{X}_{(i)}\|_2^2 \|\mathbf{Y}^{(i)}\|_2^2 \quad \text{and} \quad B \geq \frac{2n}{\ell}\max_i \|\mathbf{X}_{(i)}\|_2 \|\mathbf{Y}^{(i)}\|_2.$$

Then if $0 \leq t \leq \sigma^2/B$,

$$\mathbb{P}\left\{\left\|\hat{\mathbf{X}}\hat{\mathbf{Y}} - \mathbf{X}\mathbf{Y}\right\|_{\mathrm{F}} \geq t + \sigma\right\} \leq \exp\left(-\frac{t^2}{4\sigma^2}\right).$$

To prove Lemma 4.3, we use the following vector Bernstein inequality for sampling without replacement in Banach spaces; this result follows directly from a similar inequality for sampling with replacement established by Gross in [Gro11]. Again, vector Bernstein inequalities have been derived by multiple authors [LT91, BLM03, Rec11, Tro12, CP11, Gro11]; the value of this specific result is that it applies to the sampling without replacement model.

Lemma 4.4. *Let* $\mathcal{V}$ *be a collection of* n *vectors in a Hilbert space with norm* $\|\cdot\|_2$*. Choose* $\mathbf{V}_1, \ldots, \mathbf{V}_\ell$ *from* $\mathcal{V}$ *uniformly at random* without *replacement. Choose* $\mathbf{V}_1', \ldots, \mathbf{V}_\ell'$ *from* $\mathcal{V}$ *uniformly at random* with *replacement. Let*

$$\mu = \mathbb{E}\left\|\sum\nolimits_{i=1}^{\ell}(\mathbf{V}_i' - \mathbb{E}\mathbf{V}_i')\right\|_2$$

and set

$$\sigma^2 \geq 4\ell\mathbb{E}\left\|\mathbf{V}_1'\right\|_2^2 \quad \text{and} \quad B \geq 2\max_{\mathbf{V}\in\mathcal{V}}\|\mathbf{V}\|_2.$$

If $0 \le t \le \sigma^2/B$, *then*

$$\mathbb{P}\left\{\left\|\sum_{i=1}^{\ell} \mathbf{V}_i - \ell\mathbb{E}\mathbf{V}_1\right\|_2 \ge \mu + t\right\} \le \exp\left(-\frac{t^2}{4\sigma^2}\right).$$

Proof. We proceed by developing a bound on the moment generating function (mgf) of

$$\left\|\sum_{i=1}^{\ell} \mathbf{V}_i - \ell\mathbb{E}\mathbf{V}_1\right\|_2 - \mu.$$

This mgf is controlled by the mgf of a similar sum where the vectors are sampled with replacement. That is, for $\lambda \ge 0$,

$$\mathbb{E}\exp\left(\lambda \cdot \left\|\sum_{i=1}^{\ell} \mathbf{V}_i - \ell\mathbb{E}\mathbf{V}_1\right\|_2 - \lambda\mu\right) \le \mathbb{E}\exp\left(\lambda \cdot \left\|\sum_{i=1}^{\ell} \mathbf{V}_i' - \ell\mathbb{E}\mathbf{V}_1\right\|_2 - \lambda\mu\right). \tag{4.1.2}$$

This follows from a classical observation due to Hoeffding [Hoe63] that for any convex real-valued function g,

$$\mathbb{E}g\left(\sum_{i=1}^{\ell} \mathbf{V}_i\right) \le \mathbb{E}g\left(\sum_{i=1}^{\ell} \mathbf{V}_i'\right).$$

The paper [GN10] provides an alternate exposition of this fact. Specifically, take $g(\mathbf{V}) = \exp\left(\lambda\left\|\mathbf{V} - \ell\mathbb{E}\mathbf{V}_1\right\|_2 - \lambda\mu\right)$ to obtain the inequality of mgfs asserted in (4.1.2).

In the proof of Theorem 12 in [Gro11], Gross establishes that any random variable Z whose mgf is less than the righthand side of (4.1.2) satisfies a tail inequality of the form

$$\mathbb{P}\{Z \ge \mu + t\} \le \exp\left(-\frac{t^2}{4s^2}\right) \tag{4.1.3}$$

when $t \le s^2/M$, where

$$s^2 \ge \sum_{i=1}^{\ell} \mathbb{E}\left\|\mathbf{V}_i' - \mathbb{E}\mathbf{V}_1'\right\|_2^2$$

and $\left\|\mathbf{V}_i' - \mathbb{E}\mathbf{V}_1'\right\|_2 \le M$ almost surely for all $i = 1, \dots, \ell$. To apply this result, note that for all $i = 1, \dots, \ell$,

$$\left\|\mathbf{V}_i' - \mathbb{E}\mathbf{V}_1'\right\|_2 \le 2 \max_{\mathbf{V} \in \mathscr{V}} \|\mathbf{V}\|_2 = B.$$

Take $\mathbf{V}_1''$ to be an i.i.d. copy of $\mathbf{V}_1'$ and observe that, by Jensen's inequality,

$$\begin{aligned}
\sum_{i=1}^{\ell} \mathbb{E}\left\|\mathbf{V}_i' - \mathbb{E}\mathbf{V}_1'\right\|_2^2 &= \ell \mathbb{E}\left\|\mathbf{V}_1' - \mathbb{E}\mathbf{V}_1'\right\|_2^2 \\
&\le \ell \mathbb{E}\left\|\mathbf{V}_1' - \mathbf{V}_1''\right\|_2^2 \le \ell \mathbb{E}(\left\|\mathbf{V}_1'\right\|_2 + \left\|\mathbf{V}_1''\right\|_2)^2 \\
&\le 2\ell \mathbb{E}\left\|\mathbf{V}_1'\right\|_2^2 + \left\|\mathbf{V}_1''\right\|_2^2 \\
&= 4\ell \mathbb{E}\left\|\mathbf{V}_1'\right\|_2^2 \le \sigma^2.
\end{aligned}$$

The bound given in the statement of Lemma 4.4 when we take $s^2 = \sigma^2$ and $M = B$ in (4.1.3). □

With this Bernstein bound in hand, we proceed to the proof of Lemma 4.3. Let $\mathrm{vec} : \mathbb{R}^{m \times n} \to \mathbb{R}^{mn}$ denote the operation of vectorization, which stacks the columns of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ to form the vector $\mathrm{vec}(\mathbf{A})$.

Proof of Lemma 4.3. Let $\mathscr{V}$ be the collection of vectorized rank-one products of columns of $\sqrt{n/\ell} \cdot \mathbf{X}$ and rows of $\sqrt{n/\ell} \cdot \mathbf{Y}$. That is, take

$$\mathscr{V} = \left\{ \frac{n}{\ell} \mathrm{vec}(\mathbf{X}_{(i)} \mathbf{Y}^{(i)}) \right\}_{i=1}^{n}.$$

Sample $\mathbf{V}_1, \dots, \mathbf{V}_\ell$ uniformly at random from $\mathscr{V}$ without replacement, and observe that $\mathbb{E}\mathbf{V}_i = \ell^{-1} \mathrm{vec}(\mathbf{XY})$. With this notation,

$$\left\|\hat{\mathbf{X}}\hat{\mathbf{Y}} - \mathbf{XY}\right\|_{\mathrm{F}} \sim \left\|\sum_{i=1}^{\ell} (\mathbf{V}_i - \mathbb{E}\mathbf{V}_i)\right\|_2,$$

where $\sim$ refers to identical distributions. Therefore any probabilistic bound developed for the right-hand side quantity holds for the left-hand side quantity. The conclusion of the lemma follows when we apply Lemma 4.4 to bound the right-hand side quantity.

We calculate the variance-like term in Lemma 4.4,

$$4\ell\mathbb{E}\left\|\mathbf{V}_1\right\|_2^2 = 4\ell\frac{1}{n}\sum_{i=1}^{n}\frac{n^2}{\ell^2}\|\mathbf{X}_{(i)}\|_2^2\|\mathbf{Y}^{(i)}\|_2^2 = 4\frac{n}{\ell}\sum_{i=1}^{n}\|\mathbf{X}_{(i)}\|_2^2\|\mathbf{Y}^{(i)}\|_2^2 \leq \sigma^2.$$

Now we consider the expectation

$$\mu = \mathbb{E}\left\|\sum_{i=1}^{\ell}(\mathbf{V}_i' - \mathbb{E}\mathbf{V}_i')\right\|_2.$$

In doing so, we will use the notation $\mathbb{E}[C\,|A,B,\ldots]$ to denote the conditional expectation of a random variable C with respect to the random variables $A,B,\ldots$. Recall that a Rademacher vector is a random vector whose entries are independent and take the values ± 1 with equal probability. Let $\boldsymbol{\varepsilon}$ be a Rademacher vector of length ℓ and sample $\mathbf{V}_1',\ldots,\mathbf{V}_\ell'$ and $\mathbf{V}_1'',\ldots,\mathbf{V}_\ell''$ uniformly at random from $\mathcal{V}$ with replacement. Now μ can be bounded as follows:

$$\begin{aligned}
\mu &= \mathbb{E}\left\|\sum_{i=1}^{\ell}(\mathbf{V}_i' - \mathbb{E}\mathbf{V}_i')\right\|_2 \\
&\leq \mathbb{E}\left[\left\|\sum_{i=1}^{\ell}(\mathbf{V}_i' - \mathbf{V}_i'')\right\|_2 \,\middle|\, \{\mathbf{V}_i'\},\{\mathbf{V}_i''\}\right] \\
&= \mathbb{E}\left[\left\|\sum_{i=1}^{\ell}\varepsilon_i(\mathbf{V}_i' - \mathbf{V}_i'')\right\|_2 \,\middle|\, \{\mathbf{V}_i'\},\{\mathbf{V}_i''\},\boldsymbol{\varepsilon}\right] \\
&\leq 2\mathbb{E}\left[\left\|\sum_{i=1}^{\ell}\varepsilon_i\mathbf{V}_i'\right\|_2 \,\middle|\, \{\mathbf{V}_i'\},\boldsymbol{\varepsilon}\right] \\
&\leq 2\sqrt{\mathbb{E}\left[\left\|\sum_{i=1}^{\ell}\varepsilon_i\mathbf{V}_i'\right\|_2^2 \,\middle|\, \{\mathbf{V}_i'\},\boldsymbol{\varepsilon}\right]} \\
&= 2\sqrt{\mathbb{E}\left[\mathbb{E}\left[\sum_{i,j=1}^{\ell}\varepsilon_i\varepsilon_j{\mathbf{V}_i'}^T\mathbf{V}_j' \,\middle|\, \boldsymbol{\varepsilon}\right] \,\middle|\, \{\mathbf{V}_i'\}\right]} \\
&= 2\sqrt{\mathbb{E}\sum_{i=1}^{\ell}\left\|\mathbf{V}_i'\right\|_2^2}.
\end{aligned}$$

The first inequality is Jensen's, and the following equality holds because the components of the sequence $\{\mathbf{V}_i' - \mathbf{V}_i''\}$ are symmetric and independent. The next two manipulations are the triangle inequality and Jensen's inequality. This stage of the estimate is concluded by conditioning and using the orthogonality of the Rademacher variables. Next, the triangle inequality and the fact that $\mathbb{E}\left\|\mathbf{V}_1'\right\|_2^2 = \mathbb{E}\left\|\mathbf{V}_1\right\|_2^2$ allow us to further simplify the estimate of μ :

$$\mu \leq 2\sqrt{\mathbb{E}\sum_{i=1}^{\ell}\left\|\mathbf{V}_i'\right\|_2^2} = 2\sqrt{\ell\mathbb{E}\left\|\mathbf{V}_1\right\|_2^2} \leq \sigma.$$

We also calculate the quantity

$$2\max_{\mathbf{V}\in\mathscr{V}}\|\mathbf{V}\|_2 = \frac{2n}{\ell}\max_i \|\mathbf{X}_{(i)}\|_2\|\mathbf{Y}^{(i)}\|_2 \leq B.$$

The tail bound given in the statement of the lemma follows from applying Lemma 4.4 with our estimates for B, σ^2, and μ. □

4.2 Linear Algebra notation and results

In subsequent chapters, we use the following partitioned compact SVD to state results for rectangular matrices $\mathbf{A}$ with $\text{rank}(\mathbf{A}) = \rho$:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \begin{array}{c} \begin{matrix} k & \rho-k \end{matrix} \\ \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \end{bmatrix} \end{array} \begin{array}{c} \begin{matrix} k & \rho-k \end{matrix} \\ \begin{bmatrix} \mathbf{\Sigma}_1 & \\ & \mathbf{\Sigma}_2 \end{bmatrix} \end{array} \begin{bmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \end{bmatrix}. \tag{4.2.1}$$

Here, $\mathbf{\Sigma}_1$ contains the k largest singular values of $\mathbf{A}$ and the columns of $\mathbf{U}_1$ and $\mathbf{V}_1$ respectively span top k-dimensional left and right singular spaces of $\mathbf{A}$. The matrix $\mathbf{A}_k = \mathbf{U}_1\mathbf{\Sigma}_1\mathbf{V}_1^T$ is the optimal rank-k approximation to $\mathbf{A}$, and $\mathbf{A}_{\rho-k} = \mathbf{A} - \mathbf{A}_k = \mathbf{U}_2\mathbf{\Sigma}_2\mathbf{V}_2^T$. The Moore-Penrose

pseudoinverse of $\mathbf{A}$ is denoted by $\mathbf{A}^\dagger$.

When $\mathbf{A}$ is a positive-semidefinite matrix, $\mathbf{U} = \mathbf{V}$ and (4.2.1) becomes the following partitioned eigenvalue decomposition:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T = \begin{array}{c} \begin{matrix} k & \rho - k \end{matrix} \\ \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \end{bmatrix} \end{array} \begin{array}{c} \begin{matrix} k & \rho - k \end{matrix} \\ \begin{bmatrix} \mathbf{\Sigma}_1 & \\ & \mathbf{\Sigma}_2 \end{bmatrix} \end{array} \begin{bmatrix} \mathbf{U}_1^T \\ \mathbf{U}_2^T \end{bmatrix}. \tag{4.2.2}$$

The eigenvalues of an $n \times n$ symmetric matrix $\mathbf{A}$ are ordered $\lambda_1(\mathbf{A}) \geq \cdots \geq \lambda_n(\mathbf{A})$.

The orthoprojector onto the column space of a matrix $\mathbf{A}$ is written $\mathbf{P}_\mathbf{A}$ and satisfies

$$\mathbf{P}_\mathbf{A} = \mathbf{A}\mathbf{A}^\dagger = \mathbf{A}(\mathbf{A}^T\mathbf{A})^\dagger\mathbf{A}^T.$$

Let $\mathscr{S}$ be a k-dimensional subspace of $\mathbb{R}^n$ and $\mathbf{P}_\mathscr{S}$ denote the projection onto $\mathscr{S}$. Then the *coherence* of $\mathscr{S}$ is

$$\mu(\mathscr{S}) = \frac{n}{k} \max_i (\mathbf{P}_\mathscr{S})_{ii}.$$

The coherence of a matrix $\mathbf{U} \in \mathbb{R}^{n \times k}$ with orthonormal columns is the coherence of the subspace $\mathscr{S}$ which it spans:

$$\mu(\mathbf{U}) := \mu(\mathscr{S}) = \frac{n}{k} \max_i (\mathbf{P}_\mathscr{S})_{ii} = \frac{n}{k} \max_i (\mathbf{U}\mathbf{U}^T)_{ii}.$$

The kth column of the matrix $\mathbf{A}$ is denoted by $\mathbf{A}_{(k)}$; the jth row is denoted by $\mathbf{A}^{(j)}$. The vector $\boldsymbol{e}_i$ is the ith element of the standard Euclidean basis (whose dimensionality will be clear from the context).

We often compare SPSD matrices using the semidefinite ordering. In this ordering, $\mathbf{A}$ is greater than or equal to $\mathbf{B}$, written $\mathbf{A} \succeq \mathbf{B}$ or $\mathbf{B} \preceq \mathbf{A}$, when $\mathbf{A} - \mathbf{B}$ is positive semidefinite. Each SPSD matrix $\mathbf{A}$ has a unique square root $\mathbf{A}^{1/2}$ that is also SPSD, has the same eigenspaces as

$\mathbf{A}$, and satisfies $\mathbf{A} = \left(\mathbf{A}^{1/2}\right)^2$. The eigenvalues of an SPSD matrix $\mathbf{A}$ are arranged in weakly decreasing order: $\lambda_{\max}(\mathbf{A}) = \lambda_1(\mathbf{A}) \geq \lambda_2(\mathbf{A}) \geq \cdots \geq \lambda_n(\mathbf{A}) = \lambda_{\min}(\mathbf{A})$. Likewise, the singular values of a rectangular matrix $\mathbf{A}$ with rank ρ are ordered $\sigma_{\max}(\mathbf{A}) = \sigma_1(\mathbf{A}) \geq \sigma_2(\mathbf{A}) \geq \cdots \geq \sigma_\rho(\mathbf{A}) = \sigma_{\min}(\mathbf{A})$. The spectral norm of a matrix $\mathbf{B}$ is written $\|\mathbf{B}\|_2$; its Frobenius norm and trace are written $\|\mathbf{B}\|_{\mathrm{F}}$ and $\mathrm{Tr}(\mathbf{B})$, respectively. The notation $\|\cdot\|_\xi$ indicates that an expression holds for both $\xi = 2$ and $\xi = \mathrm{F}$.

4.2.1 Column-based low-rank approximation

The remainder of this thesis concerns low-rank matrix approximation algorithms: Chapter 5 provides bounds on the approximation errors of low-rank approximations that are formed using fast orthonormal transformations, and Chapter 6 provides bounds on the approximation errors of a class of low-rank approximations to SPSD matrices.

Both of these low-rank approximation schemes are amenable to interpretation as schemes wherein a matrix is projected onto a subspace spanned by some linear combination of its columns. The problem of providing a general framework for studying the error of these projection schemes is well studied [BMD09, HMT11, BDMI11]. The authors of these works have provided a set of so-called *structural* results: deterministic bounds on the spectral and Frobenius-norm approximation errors incurred by these projection schemes. Structural results allow us to relate the errors of low-rank approximations formed using projection schemes to the optimal errors $\|\mathbf{A} - \mathbf{A}_k\|_\xi$ for $\xi = 2, \mathrm{F}$.

Before stating the specific structural results that are used in the sequel, we review the necessary background material on low-rank matrix approximations that are restricted to lie within a particular subspace.

4.2.1.1 Matrix Pythagoras and generalized least-squares regression

Lemma 4.5 is the analog of Pythagoras' theorem in the matrix setting. A proof of this lemma can be found in [BDMI11]. Lemma 4.6 is an immediate corollary that generalizes the Eckart–Young theorem.

Lemma 4.5. *If* $\mathbf{X}\mathbf{Y}^T = \mathbf{0}$ *or* $\mathbf{X}^T\mathbf{Y} = \mathbf{0}$, *then*

$$\|\mathbf{X}+\mathbf{Y}\|_{\mathrm{F}}^2 = \|\mathbf{X}\|_{\mathrm{F}}^2 + \|\mathbf{Y}\|_{\mathrm{F}}^2$$

and

$$\max\{\|\mathbf{X}\|_2^2, \|\mathbf{Y}\|_2^2\} \leq \|\mathbf{X}+\mathbf{Y}\|_2^2 \leq \|\mathbf{X}\|_2^2 + \|\mathbf{Y}\|_2^2.$$

Lemma 4.6. *Given* $\mathbf{A} \in \mathbb{R}^{m\times n}$ *and* $\mathbf{C} \in \mathbb{R}^{m\times \ell}$, *for all* $\mathbf{X} \in \mathbb{R}^{\ell\times n}$

$$\|\mathbf{A} - \mathbf{P}_{\mathbf{C}}\mathbf{A}\|_\xi^2 \leq \|\mathbf{A} - \mathbf{C}\mathbf{X}\|_\xi^2$$

for both $\xi = 2$ *and* $\xi = \mathrm{F}$.

Proof. Write

$$\mathbf{A} - \mathbf{C}\mathbf{X} = (\mathbf{I} - \mathbf{P}_{\mathbf{C}})\mathbf{A} + (\mathbf{P}_{\mathbf{C}}\mathbf{A} - \mathbf{C}\mathbf{X})$$

and observe that

$$((\mathbf{I} - \mathbf{P}_{\mathbf{C}})\mathbf{A})^T(\mathbf{P}_{\mathbf{C}}\mathbf{A} - \mathbf{C}\mathbf{X}) = \mathbf{0},$$

so by Lemma 4.5,

$$\|\mathbf{A} - \mathbf{C}\mathbf{X}\|_\xi^2 \geq \|(\mathbf{I} - \mathbf{P}_{\mathbf{C}})\mathbf{A}\|_\xi^2.$$

□

4.2.1.2 Low-rank approximations restricted to subspaces

Given $\mathbf{A} \in \mathbb{R}^{m \times n}$; a target rank $k < n$; another matrix $\mathbf{Y} \in \mathbb{R}^{m \times \ell}$, where $\ell > k$; and a choice of norm ξ ($\xi = 2$ or $\xi = \mathrm{F}$), we use the notation $\mathbf{\Pi}^{\xi}_{\mathbf{Y},k}(\mathbf{A})$ to refer to the matrix that lies in the column span of $\mathbf{Y}$, has rank k or less, and minimizes the ξ-norm error in approximating $\mathbf{A}$. More concisely, $\mathbf{\Pi}^{\xi}_{\mathbf{Y},k}(\mathbf{A}) = \mathbf{Y}\mathbf{X}^{\xi}$, where

$$\mathbf{X}^{\xi} = \underset{\mathbf{X} \in \mathbb{R}^{\ell \times n} : \mathrm{rank}(\mathbf{X}) \leq k}{\arg\min} \|\mathbf{A} - \mathbf{Y}\mathbf{X}\|_{\xi}^{2}.$$

The approximation $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{Y},k}(\mathbf{A})$ can be computed using the following three-step procedure:

1: Orthonormalize the columns of $\mathbf{Y}$ to construct a matrix $\mathbf{Q} \in \mathbb{R}^{m \times \ell}$.

2: Compute $\mathbf{X}_{\mathrm{opt}} = \arg\min_{\mathbf{X} \in \mathbb{R}^{\ell \times n},\ \mathrm{rank}(\mathbf{X}) \leq k} \left\| \mathbf{Q}^{T}\mathbf{A} - \mathbf{X} \right\|_{\mathrm{F}}$.

3: Compute and return $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{Y},k}(\mathbf{A}) = \mathbf{Q}\mathbf{X}_{\mathrm{opt}} \in \mathbb{R}^{m \times n}$.

There does not seem to be a similarly efficient algorithm for computing $\mathbf{\Pi}^{2}_{\mathbf{Y},k}(\mathbf{A})$.

The following result, which appeared as Lemma 18 in [BDMI11], both verifies the claim that this algorithm computes $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{Y},k}(\mathbf{A})$ and shows that $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{Y},k}(\mathbf{A})$ is a constant factor approximation to $\mathbf{\Pi}^{2}_{\mathbf{Y},k}(\mathbf{A})$.

Lemma 4.7. *[Lemma 18 in [BDMI11]] Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{Y} \in \mathbb{R}^{m \times \ell}$, and an integer $k \leq \ell$, the matrix $\mathbf{Q}\mathbf{X}_{\mathrm{opt}} \in \mathbb{R}^{m \times n}$ described above satisfies $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{Y},k}(\mathbf{A}) = \mathbf{Q}\mathbf{X}_{\mathrm{opt}}$, can be computed in $\mathrm{O}(mn\ell + (m+n)\ell^{2})$ time, and satisfies*

$$\left\| \mathbf{A} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{Y},k}(\mathbf{A}) \right\|_{2}^{2} \leq 2 \left\| \mathbf{A} - \mathbf{\Pi}^{2}_{\mathbf{Y},k}(\mathbf{A}) \right\|_{2}^{2}.$$

4.2.2 Structural results for low-rank approximation

The following result, which appears as Lemma 7 in [BMD09], provides an upper bound on the residual error of the low-rank matrix approximation obtained via projections onto subspaces. The paper [HMT11] also supplies an equivalent result.

Lemma 4.8. *[Lemma 7 in [BMD09]] Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ have rank ρ. Fix k satisfying $0 \le k \le \rho$. Given a matrix $\mathbf{S} \in \mathbb{R}^{n\times \ell}$, with $\ell \ge k$, construct $\mathbf{Y} = \mathbf{AS}$. If $\mathbf{V}_1^T\mathbf{S}$ has full row-rank, then, for $\xi = 2, \mathrm{F}$,*

$$\|\mathbf{A} - \mathbf{P}_\mathbf{Y}\mathbf{A}\|_\xi^2 \le \|\mathbf{A} - \mathbf{\Pi}_{\mathbf{Y},k}^{\xi}(\mathbf{A})\|_\xi^2 \le \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 + \|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\|_\xi^2. \tag{4.2.3}$$

In addition to this bound on the residual error, we use the following novel structural bound on the forward errors of low-rank approximants.

Lemma 4.9. *Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ have rank ρ. Fix k satisfying $0 \le k \le \rho$. Given a matrix $\mathbf{S} \in \mathbb{R}^{n\times \ell}$, where $\ell \ge k$, construct $\mathbf{Y} = \mathbf{AS}$. If $\mathbf{V}_1^T\mathbf{S}$ has full row-rank, then, for $\xi = 2, \mathrm{F}$,*

$$\|\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}\|_\xi^2 \le \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 + \|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\|_\xi^2. \tag{4.2.4}$$

Proof. Observe that

$$(\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}_k)^T(\mathbf{P}_\mathbf{Y}\mathbf{A}_{\rho-k}) = \mathbf{0},$$

so Lemma 4.5 implies that

$$\|\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}\|_\xi^2 = \|\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}_{\rho-k}\|_\xi^2 \le \|\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}_k\|_\xi^2 + \|\mathbf{A}_{\rho-k}\|_\xi^2.$$

Applying Lemma 4.6 with $\mathbf{X} = (\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T$, we see that

$$\begin{aligned}\|\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}\|_\xi^2 &\leq \|\mathbf{A}_k - \mathbf{Y}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\|_\xi^2 + \|\mathbf{A}_{\rho-k}\|_\xi^2 \\ &= \|\mathbf{A}_k - \mathbf{A}_k\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T + \mathbf{A}_{\rho-k}\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\|_\xi^2 + \|\mathbf{A}_{\rho-k}\|_\xi^2 \\ &= \|\mathbf{A}_k - \mathbf{U}_1\mathbf{\Sigma}_1\mathbf{V}_1^T\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T + \mathbf{A}_{\rho-k}\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\|_\xi^2 + \|\mathbf{A}_{\rho-k}\|_\xi^2.\end{aligned}$$

Since $\mathbf{V}_1^T\mathbf{S}$ has full row rank, $(\mathbf{V}_1^T\mathbf{S})(\mathbf{V}_1^T\mathbf{S})^\dagger = \mathbf{I}_k$. Recall that $\mathbf{A}_k = \mathbf{U}_1\mathbf{\Sigma}_1\mathbf{V}_1^T$ and $\mathbf{A}_{\rho-k} = \mathbf{U}_2\mathbf{\Sigma}_2\mathbf{V}_2^T$. Consequently, the above inequality reduces neatly to the desired inequality

$$\begin{aligned}\|\mathbf{A}_k - \mathbf{P}_\mathbf{Y}\mathbf{A}\|_\xi^2 &\leq \|\mathbf{A}_k - \mathbf{U}_1\mathbf{\Sigma}_1\mathbf{V}_1^T + \mathbf{A}_{\rho-k}\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\|_\xi^2 + \|\mathbf{A}_{\rho-k}\|_\xi^2 \\ &= \|\mathbf{A}_{\rho-k}\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\|_\xi^2 + \|\mathbf{A}_{\rho-k}\|_\xi^2 \\ &= \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 + \|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\|_\xi^2.\end{aligned}$$

□

4.2.2.1 A geometric interpretation of the sampling interaction matrix

Let $\mathbf{\Omega}_1 = \mathbf{V}_1^T\mathbf{S}$ and $\mathbf{\Omega}_2 = \mathbf{V}_2^T\mathbf{S}$ denote the interaction of the sampling matrix $\mathbf{S}$ with the top and bottom right-singular spaces of $\mathbf{A}$. It is evident from Lemmas 4.8 and 4.9 that the quality of the low-rank approximations depend upon the norm of the *sampling interaction matrix*

$$\mathbf{V}_2^T\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger = \mathbf{\Omega}_2\mathbf{\Omega}_1^\dagger.$$

The smaller the spectral norm of the $\mathbf{\Omega}_2\mathbf{\Omega}_1^\dagger$ the more effective $\mathbf{S}$ is as a sampling matrix. To give the sampling interaction matrix a geometric interpretation, we first recall the definition of the

sine between the range spaces of two matrices $\mathbf{M}_1$ and $\mathbf{M}_2$:

$$\sin^2(\mathbf{M}_1, \mathbf{M}_2) = \|(\mathbf{I} - \mathbf{P}_{\mathbf{M}_1})\mathbf{P}_{\mathbf{M}_2}\|_2.$$

Note that this quantity is *not* symmetric: it measures how well the range of $\mathbf{M}_1$ captures that of $\mathbf{M}_2$ [GV96, Chapter 12].

Lemma 4.10. *Fix $\mathbf{A} \in \mathbb{R}^{m\times n}$, a target rank k, and $\mathbf{S} \in \mathbb{R}^{n\times \ell}$ where $\ell > k$. Assume $\mathbf{S}$ has orthonormal columns. Define*

$$\boldsymbol{\Omega}_1 = \mathbf{V}_1^T\mathbf{S} \quad \text{and} \quad \boldsymbol{\Omega}_2 = \mathbf{V}_2^T\mathbf{S}.$$

Then, if $\boldsymbol{\Omega}_1$ has full row-rank,

$$\|\boldsymbol{\Omega}_2\boldsymbol{\Omega}_1^\dagger\|_2 = \tan^2(\mathbf{S}, \mathbf{V}_1).$$

Proof. Since $\mathbf{V}_1$ and $\mathbf{S}$ have orthonormal columns, we see that

$$\begin{aligned}
\sin^2(\mathbf{S}, \mathbf{V}_1) &= \left\|(\mathbf{I} - \mathbf{S}\mathbf{S}^T)\mathbf{V}_1\mathbf{V}_1^T\right\|_2^2 \\
&= \left\|\mathbf{V}_1^T(\mathbf{I} - \mathbf{S}\mathbf{S}^T)\mathbf{V}_1\right\|_2 \\
&= \left\|\mathbf{I} - \mathbf{V}_1^T\mathbf{S}\mathbf{S}^T\mathbf{V}_1\right\|_2 \\
&= 1 - \lambda_k(\mathbf{V}_1^T\mathbf{S}\mathbf{S}^T\mathbf{V}_1) \\
&= 1 - \|\boldsymbol{\Omega}_1^\dagger\|^{-2}.
\end{aligned}$$

The second to last equality holds because $\mathbf{V}_1^T\mathbf{S}$ has k rows and we assumed it has full row-rank. Accordingly,

$$\tan^2(\mathbf{S}, \mathbf{V}_1) = \frac{\sin^2(\mathbf{S}, \mathbf{V}_1)}{1 - \sin^2(\mathbf{S}, \mathbf{V}_1)} = \|\boldsymbol{\Omega}_1^\dagger\|_2^2 - 1.$$

Now observe that

$$\begin{aligned}\|\mathbf{\Omega}_2\mathbf{\Omega}_1^\dagger\|_2^2 &= \left\|(\mathbf{S}^T\mathbf{V}_1)^\dagger\mathbf{S}^T\mathbf{V}_2\mathbf{V}_2^T\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\right\|_2 \\ &= \left\|(\mathbf{S}^T\mathbf{V}_1)^\dagger(\mathbf{I}-\mathbf{S}^T\mathbf{V}_1\mathbf{V}_1^T\mathbf{S})(\mathbf{V}_1^T\mathbf{S})^\dagger\right\|_2 \\ &= \left\|(\mathbf{S}^T\mathbf{V}_1)^\dagger\right\|_2^2 - 1 \\ &= \tan^2(\mathbf{S},\mathbf{V}_1).\end{aligned}$$

The second to last equality holds because of the fact that, for any matrix $\mathbf{M}$,

$$\left\|\mathbf{M}^\dagger(\mathbf{I}-\mathbf{M}\mathbf{M}^T)(\mathbf{M}^T)^\dagger\right\|_2 = \left\|\mathbf{M}^\dagger\right\|_2^2 - 1;$$

this identity can be established with a routine SVD argument. □

Thus, when $\mathbf{S}$ has orthonormal columns and $\mathbf{V}_1^T\mathbf{S}$ has full row-rank, $\|\mathbf{\Omega}_2\mathbf{\Omega}_1^\dagger\|_2$ is the tangent of the largest angle between the range of $\mathbf{S}$ and the top right singular space spanned by $\mathbf{V}_1$. If $\mathbf{V}_1^T\mathbf{S}$ does not have full row-rank, then our derivation above shows that $\sin^2(\mathbf{S},\mathbf{V}_1) = 1$, meaning that there is a vector in the eigenspace spanned by $\mathbf{V}_1$ which has no component in the space spanned by the sketching matrix $\mathbf{S}$.

We note that $\tan(\mathbf{S},\mathbf{V}_1)$ also arises in the classical bounds on the convergence of the orthogonal iteration algorithm for approximating the top k-dimensional singular spaces of a matrix (see, e.g. [GV96, Theorem 8.2.2]).

Chapter 5

Low-rank approximation with subsampled unitary transformations

5.1 Introduction

In this chapter, we analyze the theoretical performance of a randomized low-rank approximation algorithm introduced in [WLRT08] and analyzed in [WLRT08, HMT11, NDT09]. Our analysis often provides sharper approximation bounds than those in [WLRT08, HMT11, NDT09]. We provide bounds on the residual and forward errors of this approximation algorithm in the spectral and Frobenius norms, and provide experimental evidence that this low-rank approximation algorithm performs as well as a more expensive low-rank approximation algorithm based upon projections onto uniformly distributed random subspaces[1]. Further, we provide approximation bounds for a variant of the algorithm that returns approximations with even lower rank.

The setting is as follows: fix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a target rank $k \leq \min\{m, n\}$. We would like to approximate $\mathbf{A}$ with a matrix $\mathbf{X}$ that has rank close to k, and we would like $\|\mathbf{A} - \mathbf{X}\|_\xi$ to be within a small multiplicative factor of the smallest error achievable when approximating $\mathbf{A}_k$ with a rank-k matrix, for $\xi = 2, F$.

It is well-known that the rank-k matrix $\mathbf{A}_k$ that minimizes both the Frobenius and the

[1]The content of this chapter is adapted from the article [GB12] co-authored with Christos Boutsidis.

spectral-norm approximation errors can be calculated using the singular value decomposition (SVD) in $O(mn \min\{m,n\})$ arithmetic operations, using classical so-called *direct* algorithms such as QR iteration or Jacobi iteration [GV96]. Computing the full SVD is expensive when $\mathbf{A}$ is a large matrix. In this case, it is often more efficient to use *iterative* projection methods (e.g. Krylov subspace methods) to obtain approximations to $\mathbf{A}_k$. It is difficult to state a precise guarantee for the number of arithmetic operations carried out by Krylov methods, but one iteration of a Krylov method requires $\Omega(mnk)$ operations (assuming $\mathbf{A}$ has no special structure which can be exploited to speed up the computation of matrix–vector products). To obtain even an accurate rank-1 approximation requires $O(\log n)$ iterations [KW92]. Thus, an optimistic estimate for the number of operations required to compute approximate rank-k truncated SVDs using a Krylov method is $\Omega(mnk \log n)$.

Our discussion thus far has concerned only the arithmetic cost of computing truncated SVDs, but an equally or more important issue is that of the communication costs: bandwidth costs (proportional to the amount of times storage is accessed) and latency costs (proportional to the cost of transferring the information over a network or through the levels of a hierarchical memory system) [BDHS11]. If the algorithm is to be parallelized, then the complexity of the required information interchange must also be taken into account.

The randomized algorithms considered in this chapter, Algorithms 5.1 and 5.2, are of interest because they yield low-rank approximations after $\Omega(mnk \max\{\log n, \log k\})$ arithmetic operations and have low communication costs. In particular, each element of $\mathbf{A}$ is accessed only twice, and the algorithms are simple enough that they are amenable to straightforward parallelization. The guarantees provided are probabilistic, and allow one to trade off between the operation count of the algorithms and the accuracy and failure probabilities of the algorithms.

Both of the algorithms considered in this chapter are based on the intuition that, when

Algorithm 5.1: Randomized approximate truncated SVD

Input: an $m \times n$ matrix $\mathbf{A}$ and an $n \times \ell$ matrix $\mathbf{S}$, where ℓ is an integer in $[1, n]$.
Output: matrices $\tilde{\mathbf{U}}, \tilde{\mathbf{\Sigma}}, \tilde{\mathbf{V}}$ constituting the SVD of $\mathbf{P}_{\mathbf{AS}}\mathbf{A} = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.

1: Let $\mathbf{Y} = \mathbf{AS}$.
2: Compute the QR decomposition $\mathbf{Y} = \mathbf{QR}$.
3: Compute the SVD of $\mathbf{Q}^T\mathbf{A} = \mathbf{W}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.
4: Set $\tilde{\mathbf{U}} = \mathbf{QW}$.

Algorithm 5.2: Rank-k randomized approximate truncated SVD

Input: an $m \times n$ matrix $\mathbf{A}$, integers ℓ and k that satisfy $\ell > k$ and $k \in [1, n]$, and an $n \times \ell$ matrix $\mathbf{S}$.
Output: matrices $\tilde{\mathbf{U}}, \tilde{\mathbf{\Sigma}}, \tilde{\mathbf{V}}$ constituting the SVD of $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A}) = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.

1: Let $\mathbf{Y} = \mathbf{AS}$.
2: Compute the QR decomposition $\mathbf{Y} = \mathbf{QR}$.
3: Compute the rank-k truncated SVD of $\mathbf{Q}^T\mathbf{A}$ to obtain $(\mathbf{Q}^T\mathbf{A})_k = \mathbf{W}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T$.
4: Set $\tilde{\mathbf{U}} = \mathbf{QW}$.

$\mathbf{S} \in \mathbb{R}^{n \times \ell}$ is randomly selected and ℓ is sufficiently larger than k, the range of the matrix $\mathbf{AS}$ "captures" the top k-dimensional left singular space of $\mathbf{A}$. When this phenomenon occurs, the low-rank matrix formed by projecting $\mathbf{A}$ onto the range of $\mathbf{AS}$ should be almost as accurate an approximation of $\mathbf{A}$ as is the optimal approximation $\mathbf{A}_k$:

$$\|\mathbf{A} - \mathbf{P}_{\mathbf{AS}}\mathbf{A}\|_\xi \approx \|\mathbf{A} - \mathbf{A}_k\|_\xi \quad \text{for } \xi = 2, \mathrm{F}.$$

Algorithm 5.1 computes exactly this approximation, $\mathbf{P}_{\mathbf{AS}}\mathbf{A}$. Note that this approximation may have rank up to ℓ, which may be much larger than k. Algorithm 5.2 instead returns the approximation $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})$, which is guaranteed to have rank at most k.

Unlike classical iterative methods for approximating the trucated SVD, which use as many iterations as necessary to satisfy some convergence condition, Algorithms 5.1 and 5.2 use only one matrix–matrix product $\mathbf{AS}$ to generate an approximate basis for the top left singular space of $\mathbf{A}$. Accordingly, the quality of approximations obtained using either of these algorithms is

more dependent on the properties of $\mathbf{A}$ itself and the sampling matrix $\mathbf{S}$ than is the quality of approximations derived from classical iterative methods. Thus it is important to supply theoretical guarantees on the errors of the algorithms that identify which properties of $\mathbf{A}$ affect the quality of the approximations, as well as to carry to empirical studies investigating the influence of the choice of $\mathbf{S}$.

Recent years have produced a large body of research on designing random sampling matrices $\mathbf{S}$. Some proposals for $\mathbf{S}$ include: (i) every entry of $\mathbf{S}$ takes the values $+1, -1$ with equal probability [CW09, MZ11]; (ii) the entries of $\mathbf{S}$ are i.i.d. Gaussian random variables with zero mean and unit variance [HMT11]; (iii) the columns of $\mathbf{S}$ are chosen independently from the columns of the $m \times m$ identity matrix with probabilities that are proportional to the Euclidean length of the columns of $\mathbf{A}$ [FKV98, DKM06b]; and (iv) $\mathbf{S}$ is designed carefully such that $\mathbf{AS}$ can be computed in at most $O(\mathrm{n}nz(\mathbf{A}))$ arithmetic operations, where $\mathrm{n}nz(\mathbf{A})$ denotes the number of non-zero entries in $\mathbf{A}$ [CW12].

In this chapter we take $\mathbf{S}$ to be a subsampled randomized Hadamard transform (SRHT) matrix, i.e. $\mathbf{S}$ comprises a subset of the columns of a randomized Hadamard matrix (see Definitions 5.1 and 5.2 below). This choice for $\mathbf{S}$ was introduced in [AC06].

Definition 5.1 (Normalized Walsh–Hadamard Matrix)**.** *Fix an integer $n = 2^p$, for $p = 1, 2, 3, \ldots$. The (non-normalized) $n \times n$ matrix of the Hadamard-Walsh transform is defined recursively as,*

$$\mathbf{H}_n = \begin{bmatrix} \mathbf{H}_{n/2} & \mathbf{H}_{n/2} \\ \mathbf{H}_{n/2} & -\mathbf{H}_{n/2} \end{bmatrix}, \qquad \textit{with} \qquad \mathbf{H}_2 = \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix}.$$

The $n \times n$ normalized matrix of the Walsh–Hadamard transform is equal to $\mathbf{H} = n^{-\frac{1}{2}}\mathbf{H}_n \in \mathbb{R}^{n \times n}$.

Definition 5.2 (Subsampled Randomized Hadamard Transform (SRHT) matrix)**.** *Fix integers ℓ*

and $n = 2^p$ *with* $\ell < n$ *and* $p = 1, 2, 3, \ldots$. *An SRHT matrix is an* $\ell \times n$ *matrix of the form*

$$\boldsymbol{\Theta} = \sqrt{\frac{n}{\ell}} \cdot \mathbf{RHD};$$

- $\mathbf{D} \in \mathbb{R}^{n \times n}$ *is a random diagonal matrix whose entries are independent random signs, i.e. random variables uniformly distributed on* $\{\pm 1\}$.
- $\mathbf{H} \in \mathbb{R}^{n \times n}$ *is a normalized Walsh–Hadamard matrix.*
- $\mathbf{R} \in \mathbb{R}^{\ell \times n}$ *is a subset of* ℓ *rows from the* $n \times n$ *identity matrix, where the rows are chosen uniformly at random and without replacement.*

The choice of $\mathbf{S}$ as an SRHT matrix is particularly practical because the highly structured nature of $\mathbf{S}$ can be exploited to reduce the time of computing $\mathbf{AS}$ from $\mathrm{O}(mn\ell)$ to $\mathrm{O}(mn \log_2 \ell)$.

Lemma 5.3 (Fast Matrix–Vector Multiplication, Theorem 2.1 in [AL08])**.** *Given* $\mathbf{x} \in \mathbb{R}^n$ *and* $\ell < n$, *one can construct* $\boldsymbol{\Theta} \in \mathbb{R}^{\ell \times n}$ *and compute* $\boldsymbol{\Theta}\mathbf{x}$ *in at most* $2n \log_2(\ell + 1))$ *operations.*

Beyond the SRHT. The SRHT is defined only when the matrix dimension is a power of two. An alternative option is to use other structured orthonormal randomized transforms such as the real Fourier transform (DFT), the discrete cosine transform (DCT) or the discrete Hartley transform (DHT) [WLRT08, NDT09, RT08, AMT10], whose entries are on the order of $n^{-1/2}$. None of these transforms place restrictions on the size of the matrix being approximated. With minimal effort, the results of this chapter can be extended to encompass these transforms. Specifically, the statements of Lemma 5.5 and Lemma 5.8 in this chapter would need to be modified slightly to account for the difference in the transform; essentially, the constants present in the statements of the Lemmas would change. These two lemmas isolate the effects of the particular choice of

rank at most k, while the bounds in Theorem 5.4 are for approximations using Algorithm 5.2, which may result in an approximation with a rank larger than k. However, the comparisons above also hold, as stated, for approximations $\tilde{\mathbf{A}}$ generated using Algorithm 5.2, as can be verified directly from Theorems 5.13 and 5.14.

Halko et al. [HMT11]. Halko et al. [HMT11] consider the performance of Algorithm 5.1 when $\mathbf{S} = \mathbf{\Theta}^T$ is an SRHT matrix, and conclude that if ℓ satisfies

$$4\left[\sqrt{k} + \sqrt{8\log(kn)}\right]^2 \log(k) \leq \ell \leq n, \tag{5.2.2}$$

then, for both $\xi = 2, \mathrm{F}$,

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_\xi \leq \left(1 + \sqrt{7n/\ell}\right) \cdot \|\mathbf{A} - \mathbf{A}_k\|_\xi,$$

with probability at least $1 - \mathrm{O}(1/k)$. Our Frobenius-norm bound is always tighter than the Frobenius-norm bound given here. To compare the spectral-norm bounds, note that our spectral-norm bound is on the order of

$$\max\left\{\sqrt{\frac{\log(\rho/\delta)\log(n/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2, \sqrt{\frac{\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_\mathrm{F}\right\}. \tag{5.2.3}$$

If the residual spectrum $\mathbf{A}$ (the set of singular values smaller than $\sigma_k(\mathbf{A})$) is constant, or more generally decays slowly, then the spectral-norm result in [HMT11] is perhaps optimal. But when $\mathbf{A}$ is rank-deficient or the singular values of $\mathbf{A}$ decay fast, the spectral-norm bound in Theorem 5.4 is more useful. Specifically, if

$$\left\|\mathbf{A} - \mathbf{A}_k\right\|_\mathrm{F} \ll \sqrt{\frac{n}{\log(\rho/\delta)}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2,$$

then when ℓ is chosen according to Theorem 5.4, the quantity in (5.2.3) is much smaller than $\sqrt{7n/\ell} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2$.

We were able to obtain this improved bound by using the results in Section 5.3.1, which allow one to take into account decay in the spectrum of $\mathbf{A}$. Finally, notice that our theorem makes explicit the intuition that the probability of failure can be driven to zero independently of the target rank k by increasing ℓ.

Two alternative approximate SVD algorithms. Instead of an SRHT matrix, one can take $\mathbf{S}$ in Algorithms 5.1 and 5.2 to be a matrix of i.i.d. standard Gaussian random variables. One gains theoretically and often empirically better worse-case tradeoffs between the number of samples taken, the failure probability, and the error guarantees. The SRHT algorithms are still faster, since a matrix multiplication with a Gaussian matrix requires $O(mn\ell)$ time (assuming $\mathbf{A}$ is dense and unstructured). One can also take $\mathbf{S}$ to be a matrix of i.i.d. random signs (± 1 with equal probability). In many ways, this is analogous to the Gaussian algorithm: in both cases $\mathbf{S}$ is a matrix of i.i.d. subgaussian random variables. Consequently, we expect this algorithm to have the same advantages and disadvantages relative to the SRHT algorithm. We now compare the best available performance bounds for these schemes to our SRHT performance bounds.

We use the notion of the stable rank of a matrix,

$$\mathrm{sr}(\mathbf{A}) = \|\mathbf{A}\|_{\mathrm{F}}^2 / \|\mathbf{A}\|_2^2,$$

to capture the decay of the singular values of $\mathbf{A}$. As can be seen by considering a matrix with a flat singular spectrum, in general the stable rank is no smaller than the rank.

When $\ell > k + 4$, Theorem 10.7 and Corollary 10.9 in [HMT11] imply that, when using

Gaussian sampling in Algorithm 5.1, with probability at least $1 - 2 \cdot 32^{-(\ell-k)} - \mathrm{e}^{\frac{-(\ell-k+1)}{2}}$,

$$\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_{\mathrm{F}} \leq \left(1 + 32\frac{\sqrt{3k} + \mathrm{e}\sqrt{\ell}}{\sqrt{\ell - k + 1}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}$$

and with probability at least $1 - 3e^{-(\ell-k)}$,

$$\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_2 \leq \left(1 + 16\sqrt{1 + \frac{k}{\ell - k}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \frac{8\sqrt{\ell}}{\ell - k + 1} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}.$$

Comparing to the guarantees of Theorem 5.4 we see that the two bounds just stated suggest that with the same number of samples, Gaussian low-rank approximations outperform SRHT-based low-rank approximations. In particular, the spectral-norm bound guarantees that if $\mathrm{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq k$, that is, $\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \leq \sqrt{k}\left\|\mathbf{A} - \mathbf{A}_k\right\|_2$, then the Gaussian version of Algorithm 5.1 requires $\mathrm{O}(k/\epsilon^2)$ samples to return a $17 + \epsilon$ constant factor spectral-norm error approximation with high probability. Similarly, the Frobenius-norm bound guarantees that the same number of samples returns a $1 + 32\epsilon$ constant factor Frobenius-norm error approximation with high probability. Neither the spectral nor Frobenius bounds given in Theorem 5.4 for SRHT-based low-rank approximations apply for this few samples.

The paper [MZ11] does not consider the Frobenius-norm error of the random sign low-rank approximation algorithm, but Remark 4 in [MZ11] shows that when $\ell = \mathrm{O}(k/\epsilon^4 \log(1/\delta))$, for $0 < \delta < 1$, and $\mathrm{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq k$, Algorithm 5.1 ensures that with probability at least $1 - \delta$,

$$\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_2 \leq (1 + \epsilon)\left\|\mathbf{A} - \mathbf{A}_k\right\|_2.$$

To compare our results with those stated in [HMT11, MZ11] we assume that $k \gg \log(n/\delta)$ so that $\ell > k \log k$ suffices for Theorem 5.4 to apply. Then, in order to acquire a $4 + \epsilon$ relative

error bound from Theorem 5.4, it suffices that

$$\ell \geq \mathrm{C}' \epsilon^{-2} k \log(\rho/\delta) \quad \text{and} \quad \operatorname{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq \mathrm{C}' k,$$

where C' is an explicit constant no larger than 6.

We see, therefore, that the Gaussian and random sign approximation versions of Algorithm 5.1 return $17 + \epsilon$ and $1 + \epsilon$ relative spectral error approximations, respectively, when ℓ is on the order of k and the relatively weak spectral decay condition $\operatorname{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq k$ is satisfied, while our bound for the SRHT version of Algorithm 5.1 requires $\ell > k \log(\rho/\delta)$ and the spectral decay condition

$$\operatorname{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq \mathrm{C}' k$$

to ensure a $6 + \epsilon$ relative spectral error approximation. We note that the SRHT algorithm can be used to obtain relative spectral error approximations of matrices with arbitrary stable rank at the cost of increasing ℓ; the same is of course true for the Gaussian and random sign algorithms.

The bounds for the SRHT, Gaussian, and random sign low-rank approximation algorithms differ in two significant ways. First, there are logarithmic factors in the spectral-norm error bound for the SRHT algorithm that are absent from the corresponding bounds for the Gaussian and random sign algorithms. Second, the spectral-norm bound for the SRHT algorithm applies only when $\ell > k \log(\rho/\delta)$, while the corresponding bounds for the Gaussian and random sign algorithms apply when ℓ is on the order of k. These disparities may reflect a fundamental tradeoff between the structure and randomness of the sampling matrix $\mathbf{S}$. Assuming $\mathbf{A}$ is dense and unstructured, the highly structured nature of SRHT matrices makes it possible to calculate $\mathbf{AS}$ much faster than when Gaussian or random sign sampling matrices are used, but this moves us away from the very nice isotropic randomness present in the Gaussian $\mathbf{S}$ and the similarly nice

properties of a matrix of i.i.d subgaussian random variables, thus resulting in slacker bounds which require more samples.

5.3 Matrix computations with SRHT matrices

An important ingredient in analyzing the performance of Algorithms 5.1 and 5.2 is understanding how an SRHT changes the spectrum of a matrix after postmultiplication: given a matrix $\mathbf{A}$ and an SRHT matrix $\mathbf{\Theta}$, how are the singular values of $\mathbf{A}$ and $\mathbf{A\Theta}^T$ related?

To be more precise, Lemma 4.8 suggests that one path towards establishing the efficacy of SRHT-based low-rank approximations lies in understanding how the SRHT perturbs the singular values of matrices. To see this, we repeat the statement of the lemma here. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ have rank ρ and recall the following partitioning of its SVD:

$$\mathbf{A} = \mathbf{U\Sigma V}^T = \begin{array}{c} \begin{matrix} k & \rho-k \end{matrix} \\ \begin{bmatrix} & \\ \mathbf{U}_1 & \mathbf{U}_2 \end{bmatrix} \end{array} \begin{array}{c} \begin{matrix} k & \rho-k \end{matrix} \\ \begin{bmatrix} \mathbf{\Sigma}_1 & \\ & \mathbf{\Sigma}_2 \end{bmatrix} \end{array} \begin{bmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \end{bmatrix}. \tag{5.3.1}$$

Fix k satisfying $0 \le k \le \rho$. Given a matrix $\mathbf{S} \in \mathbb{R}^{n\times \ell}$, with $\ell \ge k$, Lemma 4.8 states that if $\mathbf{V}_1^T\mathbf{S}$ has full row-rank, then for $\xi = 2, \mathrm{F}$,

$$\|\mathbf{A} - \mathbf{P}_{\mathbf{AS}}\mathbf{A}\|_\xi^2 \le \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 + \|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{S}\left(\mathbf{V}_1^T\mathbf{S}\right)^\dagger\|_\xi^2.$$

Note that $\mathbf{P}_{\mathbf{AS}}\mathbf{A}$ is exactly the low-rank approximation of $\mathbf{A}$ returned by Algorithm 5.1.

Now take $\mathbf{S} = \mathbf{\Theta}^T$ to be an SRHT matrix and observe that if the product $\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\left(\mathbf{V}_1^T\mathbf{\Theta}^T\right)^\dagger$ has small norm, then the residual error of the approximant $\mathbf{P}_{\mathbf{A\Theta}^T}\mathbf{A}$ is small. The norm of this

product is small when the norm of the perturbed matrix $\left(\mathbf{V}_1^T\boldsymbol{\Theta}^T\right)^\dagger$ is small, because

$$\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\|_\xi^2 \leq \|\boldsymbol{\Sigma}_2\|_\xi^2\|\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_2^2\left\|(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\right\|_2^2. \tag{5.3.2}$$

The matrix $(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger$ has small norm precisely when the singular values of $\mathbf{V}_1^T\boldsymbol{\Theta}^T$ are close to those of $\mathbf{V}_1$. This strategy is developed in [HMT11] to supply bounds for the SRHT low-rank approximation algorithm, and relies upon knowledge of how multiplication with SRHT matrices perturbs the singular values of matrices with orthonormal rows.

The main contribution of this chapter is the realization that one can take advantage of the decay in the singular values of $\mathbf{A}$ encoded in $\boldsymbol{\Sigma}_2$ to obtain sharper results. In view of the fact that

$$\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\|_\xi^2 \leq \|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_\xi^2\left\|(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\right\|_2^2, \tag{5.3.3}$$

we can additionally consider the behavior of the singular values of $\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T$. It is clear that (5.3.3) provides a tighter bound than (5.3.2): for example, when $\xi = \mathrm{F}$,

$$\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_\mathrm{F}^2 \leq \|\boldsymbol{\Sigma}_2\|_\mathrm{F}^2\|\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_2^2,$$

and the quantity on the left-hand side is potentially much smaller than that on the right-hand side.

In the remainder of this section, we refer to matrices with more columns than rows as “fat”; similarly, we refer to matrices with more rows than columns as “tall.”

$C \geq 1$ be any specified constant. If $\boldsymbol{\Theta} \in \mathbb{R}^{\ell \times n}$ is an SRHT matrix and ℓ satisfies

$$6C^2\epsilon^{-1}\left[\sqrt{k} + \sqrt{8\log(n/\delta)}\right]^2 \log(k/\delta) \leq \ell \leq n, \tag{5.2.1}$$

then the approximation $\tilde{\mathbf{A}}$ generated by Algorithm 5.1 with $\mathbf{S} = \boldsymbol{\Theta}^T$ satisfies

$$\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A}\boldsymbol{\Theta}^T}\mathbf{A}\right\|_{\mathrm{F}} \leq (1 + 11\epsilon)\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}$$

with probability at least $1 - \delta^{C^2\log(k/\delta)/4} - 7\delta$, and

$$\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A}\boldsymbol{\Theta}^T}\mathbf{A}\right\|_2 \leq \left(4 + \sqrt{\frac{3\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{3\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}$$

with probability at least $1 - 5\delta$.

The Frobenius-norm bound in this theorem is slightly stronger than the best bound appearing in prior efforts [NDT09]. The spectral-norm bound on the residual error is often much smaller than the bounds presented in prior work and sheds light on an open question mentioned in [NDT09] and [HMT11]. We do not, however, claim that the error bounds provided are the tightest possible. Certainly the specific constants (11, 4, etc.) in the error estimates are not optimized.

We now present a detailed comparison of the guarantees given in Theorem 5.4 with those available in the existing literature.

5.2.1 Detailed comparison with prior work

The subsampled randomized Fourier transform (SRFT). The algorithm in Section 5.2 of [WLRT08], which was the first to use the idea of employing subsampled randomized orthogonal

transforms to compute low-rank approximations to matrices, provides a spectral-norm error bound but replaces the SRHT with an SRFT, i.e. the matrix $\mathbf{H}$ of Definition 5.2 is replaced by a matrix where the (p,q)th entry is $\mathbf{H}_{pq} = \mathrm{e}^{-2\pi \mathrm{i}(p-1)(q-1)/n}$, where $\mathrm{i} = \sqrt{-1}$, i.e. $\mathbf{H}$ is the unnormalized discrete Fourier transform. Woolfe et al. [WLRT08, Equation 190] argue that, for any $\alpha > 1$, $\beta > 1$, if

$$\ell \geq \alpha^2 \beta \, (\alpha - 1)^{-1} \, (2k)^2,$$

then with probability at least $1 - 3/\beta$ ($\omega = \max\{m, n\}$),

$$\left\| \mathbf{A} - \tilde{\mathbf{U}}_k \tilde{\mathbf{\Sigma}}_k \tilde{\mathbf{V}}_k^T \right\|_2 \leq 2 \left(\sqrt{2\alpha - 1} + 1 \right) \cdot \left(\sqrt{\alpha\omega + 1} + \sqrt{\alpha\omega} \right) \cdot \left\| \mathbf{A} - \mathbf{A}_k \right\|_2 .$$

Here, $\tilde{\mathbf{U}}_k \in \mathbb{R}^{m \times k}$ contains orthonormal columns, as does $\tilde{\mathbf{V}}_k \in \mathbb{R}^{n \times k}$, while $\tilde{\mathbf{\Sigma}}_k \in \mathbb{R}^{k \times k}$ is diagonal with nonegative entries. These matrices can be computed deterministically from $\mathbf{A}\mathbf{\Theta}^T$ in $\mathrm{O}\left(k^2(m+n) + k\ell^2 \log \ell\right)$ time. Also, computing $\mathbf{Y} = \mathbf{A}\mathbf{\Theta}^T$ takes $\mathrm{O}\left(mn \log \ell\right)$ time.

The analysis of [WLRT08] applies when $\ell = \Omega(k^2)$, while the spectral-norm guarantee of Theorem 5.4 applies for potentially much smaller values of $\ell = \Omega(\max\{k \log k, \log(n) \log k\})$.

Nguyen et al. [NDT09]. An analysis of the Frobenius-norm error of an SRHT-based low-rank matrix approximation algorithm appeared in Nguyen et al. [NDT09]. Let δ be a probability parameter with $0 < \delta < 1$ and ϵ be an accuracy parameter with $0 < \epsilon < 1$. Then, Nguyen et al. show that in order to get a rank-k matrix $\tilde{\mathbf{A}}_k$ satisfying

$$\left\| \mathbf{A} - \tilde{\mathbf{A}}_k \right\|_{\mathrm{F}} \leq (1 + \epsilon) \cdot \left\| \mathbf{A} - \mathbf{A}_k \right\|_{\mathrm{F}}$$

and

$$\left\| \mathbf{A} - \tilde{\mathbf{A}}_k \right\|_2 \leq \left(2 + \sqrt{2n/\ell} \right) \cdot \left\| \mathbf{A} - \mathbf{A}_k \right\|_2$$

with probability of success at least $1 - 5\delta$, one requires

$$\ell = \Omega\left(\epsilon^{-1} \max\{k, \sqrt{k}\log(2n/\delta)\} \cdot \max\{\log k, \log(3/\delta)\}\right).$$

Theorem 5.4 gives a tighter spectral-norm error bound when $\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \ll (n/\log(\rho/\delta))^{1/2} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2$. It also provides an equivalent Frobenius-norm error bound with a comparable failure probability for a smaller number of samples. Specifically, if

$$\begin{aligned}\ell &\geq 528\epsilon^{-1}[\sqrt{k} + \sqrt{8\log(8n/\delta)}]^2 \log(8k/\delta) \\ &= \Omega\left(\epsilon^{-1} \max\{k, \log(n/\delta)\} \cdot \max\{\log k, \log(1/\delta)\}\right),\end{aligned}$$

then the Frobenius-norm bound in Theorem 5.4 ensures that, with probability at least $1 - 8\delta$, the approximation satisfies $\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_{\mathrm{F}} \leq (1+\epsilon) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}$.

In [HMT11] and [NDT09], the authors left as a subject for future research the explanation of a curious experimental phenomenon: when the singular values decay according to power laws, the SRHT low-rank approximation algorithm empirically achieves relative-error spectral norm approximations. Our spectral norm result provides an explanation of this phenomenon: when the singular values of $\mathbf{A}$ decay fast enough, as in power law decay, one has $\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} = \Theta(1) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2$. In this case, when ℓ is chosen to satisfy

$$24\epsilon^{-1}\left[\sqrt{k} + \sqrt{8\log(n/\delta)}\right]^2 \log(k/\delta)\log(n/\delta) \leq \ell \leq n$$

our spectral-norm bound assures us that $\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_2 \leq \mathrm{O}(1) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2$ with probability of at least $1 - 8\delta$, thus predicting the observed empirical behavior of the algorithm.

The approximation scheme addressed in [NDT09] generates low-rank approximations with

rank at most k, while the bounds in Theorem 5.4 are for approximations using Algorithm 5.2, which may result in an approximation with a rank larger than k. However, the comparisons above also hold, as stated, for approximations $\tilde{\mathbf{A}}$ generated using Algorithm 5.2, as can be verified directly from Theorems 5.13 and 5.14.

Halko et al. [HMT11]. Halko et al. [HMT11] consider the performance of Algorithm 5.1 when $\mathbf{S} = \mathbf{\Theta}^T$ is an SRHT matrix, and conclude that if ℓ satisfies

$$4\left[\sqrt{k} + \sqrt{8\log(kn)}\right]^2 \log(k) \leq \ell \leq n, \tag{5.2.2}$$

then, for both $\xi = 2, \mathrm{F}$,

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_\xi \leq \left(1 + \sqrt{7n/\ell}\right) \cdot \|\mathbf{A} - \mathbf{A}_k\|_\xi,$$

with probability at least $1 - \mathrm{O}(1/k)$. Our Frobenius-norm bound is always tighter than the Frobenius-norm bound given here. To compare the spectral-norm bounds, note that our spectral-norm bound is on the order of

$$\max\left\{\sqrt{\frac{\log(\rho/\delta)\log(n/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2, \sqrt{\frac{\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_\mathrm{F}\right\}. \tag{5.2.3}$$

If the residual spectrum $\mathbf{A}$ (the set of singular values smaller than $\sigma_k(\mathbf{A})$) is constant, or more generally decays slowly, then the spectral-norm result in [HMT11] is perhaps optimal. But when $\mathbf{A}$ is rank-deficient or the singular values of $\mathbf{A}$ decay fast, the spectral-norm bound in Theorem 5.4 is more useful. Specifically, if

$$\left\|\mathbf{A} - \mathbf{A}_k\right\|_\mathrm{F} \ll \sqrt{\frac{n}{\log(\rho/\delta)}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2,$$

then when ℓ is chosen according to Theorem 5.4, the quantity in (5.2.3) is much smaller than $\sqrt{7n/\ell} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2$.

We were able to obtain this improved bound by using the results in Section 5.3.1, which allow one to take into account decay in the spectrum of $\mathbf{A}$. Finally, notice that our theorem makes explicit the intuition that the probability of failure can be driven to zero independently of the target rank k by increasing ℓ.

Two alternative approximate SVD algorithms. Instead of an SRHT matrix, one can take $\mathbf{S}$ in Algorithms 5.1 and 5.2 to be a matrix of i.i.d. standard Gaussian random variables. One gains theoretically and often empirically better worse-case tradeoffs between the number of samples taken, the failure probability, and the error guarantees. The SRHT algorithms are still faster, since a matrix multiplication with a Gaussian matrix requires $O(mn\ell)$ time (assuming $\mathbf{A}$ is dense and unstructured). One can also take $\mathbf{S}$ to be a matrix of i.i.d. random signs (± 1 with equal probability). In many ways, this is analogous to the Gaussian algorithm: in both cases $\mathbf{S}$ is a matrix of i.i.d. subgaussian random variables. Consequently, we expect this algorithm to have the same advantages and disadvantages relative to the SRHT algorithm. We now compare the best available performance bounds for these schemes to our SRHT performance bounds.

We use the notion of the stable rank of a matrix,

$$\mathrm{sr}(\mathbf{A}) = \|\mathbf{A}\|_{\mathrm{F}}^2 / \|\mathbf{A}\|_2^2,$$

to capture the decay of the singular values of $\mathbf{A}$. As can be seen by considering a matrix with a flat singular spectrum, in general the stable rank is no smaller than the rank.

When $\ell > k + 4$, Theorem 10.7 and Corollary 10.9 in [HMT11] imply that, when using

Gaussian sampling in Algorithm 5.1, with probability at least $1 - 2 \cdot 32^{-(\ell-k)} - \mathrm{e}^{\frac{-(\ell-k+1)}{2}}$,

$$\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_{\mathrm{F}} \leq \left(1 + 32\frac{\sqrt{3k} + \mathrm{e}\sqrt{\ell}}{\sqrt{\ell - k + 1}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}$$

and with probability at least $1 - 3e^{-(\ell-k)}$,

$$\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_2 \leq \left(1 + 16\sqrt{1 + \frac{k}{\ell - k}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \frac{8\sqrt{\ell}}{\ell - k + 1} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}.$$

Comparing to the guarantees of Theorem 5.4 we see that the two bounds just stated suggest that with the same number of samples, Gaussian low-rank approximations outperform SRHT-based low-rank approximations. In particular, the spectral-norm bound guarantees that if $\mathrm{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq k$, that is, $\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \leq \sqrt{k}\left\|\mathbf{A} - \mathbf{A}_k\right\|_2$, then the Gaussian version of Algorithm 5.1 requires $\mathrm{O}(k/\epsilon^2)$ samples to return a $17 + \epsilon$ constant factor spectral-norm error approximation with high probability. Similarly, the Frobenius-norm bound guarantees that the same number of samples returns a $1 + 32\epsilon$ constant factor Frobenius-norm error approximation with high probability. Neither the spectral nor Frobenius bounds given in Theorem 5.4 for SRHT-based low-rank approximations apply for this few samples.

The paper [MZ11] does not consider the Frobenius-norm error of the random sign low-rank approximation algorithm, but Remark 4 in [MZ11] shows that when $\ell = \mathrm{O}(k/\epsilon^4 \log(1/\delta))$, for $0 < \delta < 1$, and $\mathrm{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq k$, Algorithm 5.1 ensures that with probability at least $1 - \delta$,

$$\left\|\mathbf{A} - \tilde{\mathbf{A}}\right\|_2 \leq (1 + \epsilon)\left\|\mathbf{A} - \mathbf{A}_k\right\|_2.$$

To compare our results with those stated in [HMT11, MZ11] we assume that $k \gg \log(n/\delta)$ so that $\ell > k \log k$ suffices for Theorem 5.4 to apply. Then, in order to acquire a $4 + \epsilon$ relative

error bound from Theorem 5.4, it suffices that

$$\ell \geq \mathrm{C}' \epsilon^{-2} k \log(\rho/\delta) \quad \text{and} \quad \operatorname{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq \mathrm{C}' k,$$

where C' is an explicit constant no larger than 6.

We see, therefore, that the Gaussian and random sign approximation versions of Algorithm 5.1 return $17 + \epsilon$ and $1 + \epsilon$ relative spectral error approximations, respectively, when ℓ is on the order of k and the relatively weak spectral decay condition $\operatorname{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq k$ is satisfied, while our bound for the SRHT version of Algorithm 5.1 requires $\ell > k \log(\rho/\delta)$ and the spectral decay condition

$$\operatorname{sr}\left(\mathbf{A} - \mathbf{A}_k\right) \leq \mathrm{C}' k$$

to ensure a $6 + \epsilon$ relative spectral error approximation. We note that the SRHT algorithm can be used to obtain relative spectral error approximations of matrices with arbitrary stable rank at the cost of increasing ℓ; the same is of course true for the Gaussian and random sign algorithms.

The bounds for the SRHT, Gaussian, and random sign low-rank approximation algorithms differ in two significant ways. First, there are logarithmic factors in the spectral-norm error bound for the SRHT algorithm that are absent from the corresponding bounds for the Gaussian and random sign algorithms. Second, the spectral-norm bound for the SRHT algorithm applies only when $\ell > k \log(\rho/\delta)$, while the corresponding bounds for the Gaussian and random sign algorithms apply when ℓ is on the order of k. These disparities may reflect a fundamental tradeoff between the structure and randomness of the sampling matrix $\mathbf{S}$. Assuming $\mathbf{A}$ is dense and unstructured, the highly structured nature of SRHT matrices makes it possible to calculate $\mathbf{AS}$ much faster than when Gaussian or random sign sampling matrices are used, but this moves us away from the very nice isotropic randomness present in the Gaussian $\mathbf{S}$ and the similarly nice

properties of a matrix of i.i.d subgaussian random variables, thus resulting in slacker bounds which require more samples.

5.3 Matrix computations with SRHT matrices

An important ingredient in analyzing the performance of Algorithms 5.1 and 5.2 is understanding how an SRHT changes the spectrum of a matrix after postmultiplication: given a matrix $\mathbf{A}$ and an SRHT matrix $\mathbf{\Theta}$, how are the singular values of $\mathbf{A}$ and $\mathbf{A\Theta}^T$ related?

To be more precise, Lemma 4.8 suggests that one path towards establishing the efficacy of SRHT-based low-rank approximations lies in understanding how the SRHT perturbs the singular values of matrices. To see this, we repeat the statement of the lemma here. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ have rank ρ and recall the following partitioning of its SVD:

$$\mathbf{A} = \mathbf{U\Sigma V}^T = \begin{array}{c} \begin{matrix} k & \rho-k \end{matrix} \\ \begin{bmatrix} & \\ \mathbf{U}_1 & \mathbf{U}_2 \end{bmatrix} \end{array} \begin{array}{c} \begin{matrix} k & \rho-k \end{matrix} \\ \begin{bmatrix} \mathbf{\Sigma}_1 & \\ & \mathbf{\Sigma}_2 \end{bmatrix} \end{array} \begin{bmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \end{bmatrix}. \tag{5.3.1}$$

Fix k satisfying $0 \le k \le \rho$. Given a matrix $\mathbf{S} \in \mathbb{R}^{n\times \ell}$, with $\ell \ge k$, Lemma 4.8 states that if $\mathbf{V}_1^T\mathbf{S}$ has full row-rank, then for $\xi = 2, \mathrm{F}$,

$$\|\mathbf{A} - \mathbf{P}_{\mathbf{AS}}\mathbf{A}\|_\xi^2 \le \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 + \|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{S}\left(\mathbf{V}_1^T\mathbf{S}\right)^\dagger\|_\xi^2.$$

Note that $\mathbf{P}_{\mathbf{AS}}\mathbf{A}$ is exactly the low-rank approximation of $\mathbf{A}$ returned by Algorithm 5.1.

Now take $\mathbf{S} = \mathbf{\Theta}^T$ to be an SRHT matrix and observe that if the product $\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\left(\mathbf{V}_1^T\mathbf{\Theta}^T\right)^\dagger$ has small norm, then the residual error of the approximant $\mathbf{P}_{\mathbf{A\Theta}^T}\mathbf{A}$ is small. The norm of this

product is small when the norm of the perturbed matrix $\left(\mathbf{V}_1^T\boldsymbol{\Theta}^T\right)^\dagger$ is small, because

$$\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\|_\xi^2 \leq \|\boldsymbol{\Sigma}_2\|_\xi^2\|\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_2^2\left\|(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\right\|_2^2. \tag{5.3.2}$$

The matrix $(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger$ has small norm precisely when the singular values of $\mathbf{V}_1^T\boldsymbol{\Theta}^T$ are close to those of $\mathbf{V}_1$. This strategy is developed in [HMT11] to supply bounds for the SRHT low-rank approximation algorithm, and relies upon knowledge of how multiplication with SRHT matrices perturbs the singular values of matrices with orthonormal rows.

The main contribution of this chapter is the realization that one can take advantage of the decay in the singular values of $\mathbf{A}$ encoded in $\boldsymbol{\Sigma}_2$ to obtain sharper results. In view of the fact that

$$\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\|_\xi^2 \leq \|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_\xi^2\left\|(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\right\|_2^2, \tag{5.3.3}$$

we can additionally consider the behavior of the singular values of $\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T$. It is clear that (5.3.3) provides a tighter bound than (5.3.2): for example, when $\xi = \mathrm{F}$,

$$\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_\mathrm{F}^2 \leq \|\boldsymbol{\Sigma}_2\|_\mathrm{F}^2\|\mathbf{V}_2^T\boldsymbol{\Theta}^T\|_2^2,$$

and the quantity on the left-hand side is potentially much smaller than that on the right-hand side.

In the remainder of this section, we refer to matrices with more columns than rows as “fat”; similarly, we refer to matrices with more rows than columns as “tall.”

5.3.1 SRHTs applied to orthonormal matrices

In this subsection, we collect known results on how the singular values of a matrix with orthonormal rows are affected by postmultiplication by an SRHT matrix. These results will be used to control the quantity $\left\|(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\right\|_2^2$ in (5.3.3).

It has recently been shown by Tropp [Tro11b] that, if the SRHT matrix is of sufficiently large dimensions, post-multiplying a fat matrix with orthonormal rows by an SRHT matrix preserves the singular values of the orthonormal matrix, with high probability, up to a small multiplicative factor. The following lemma is essentially a reparametrization of Theorem 3.1 in [Tro11b], but we include a full proof for completeness.

Lemma 5.5 (The SRHT preserves geometry)**.** *Assume n is a power of 2. Let $\mathbf{V}^T \in \mathbb{R}^{k\times n}$ have orthonormal rows. Choose parameters $0 < \epsilon < 1/3$ and $0 < \delta < 1$. Construct an SRHT matrix $\boldsymbol{\Theta} \in \mathbb{R}^{\ell\times n}$ with ℓ satisfying*

$$6\epsilon^{-1}\left[\sqrt{k} + \sqrt{8\log(n/\delta)}\right]^2 \log(k/\delta) \le \ell \le n. \tag{5.3.4}$$

Then, with probability at least $1 - 3\delta$, for all $i = 1, \ldots, k$,

$$\sqrt{1-\sqrt{\epsilon}} \le \sigma_i(\mathbf{V}^T\boldsymbol{\Theta}^T) \le \sqrt{1+\sqrt{\epsilon}}$$

and

$$\left\|(\mathbf{V}^T\boldsymbol{\Theta}^T)^\dagger - (\mathbf{V}^T\boldsymbol{\Theta}^T)^T\right\|_2 \le 1.54\sqrt{\epsilon}.$$

By the definition of an SRHT matrix, $\mathbf{V}^T\boldsymbol{\Theta}^T = \sqrt{n/\ell}\cdot(\mathbf{RHDV})^T$. Tropp [Tro11b] argues that the above lemma follows from a more fundamental fact: if $\mathbf{V}$ has orthonormal columns, then

the rows of the product **HDV** all have roughly the same norm. That is, premultiplication by **HD** equalizes the row norms of an orthonormal matrix. See also [AC06].

Lemma 5.6 (Row norms, Lemma 3.3 in [Tro11b]). *Assume n is a power of 2. Let $\mathbf{V} \in \mathbb{R}^{n\times k}$ have orthonormal columns; let $\mathbf{H} \in \mathbb{R}^{n\times n}$ be a normalized Hadamard matrix; and let $\mathbf{D} \in \mathbb{R}^{n\times n}$ be a diagonal matrix of independent Rademacher random variables. Choose a failure probability $0 < \delta < 1$. Then, with probability at least $1 - \delta$,*

$$\max_{i=1,\dots,n} \left\| (\mathbf{HDV})^{(i)} \right\|_2 \le \sqrt{\frac{k}{n}} + \sqrt{\frac{8\log(n/\delta)}{n}}.$$

Recall that $(\mathbf{HDV})^{(i)}$ denotes the ith row of the matrix $\mathbf{HDV} \in \mathbb{R}^{n\times k}$.

To prove Lemma 5.5 we need one more result on uniform random sampling (without replacement) of rows from thin matrices with orthonormal columns. The following lemma is a reparameterization of Lemma 3.4 of [Tro11b].

Lemma 5.7 (Uniform sampling without replacement from an orthonormal matrix). *Let $\mathbf{W} \in \mathbb{R}^{n\times k}$ have orthonormal columns. Let $0 < \epsilon < 1$ and $0 < \delta < 1$. Let $M := n \cdot \max_{i=1,\dots,n} \left\| \mathbf{W}^{(i)} \right\|_2^2$. Let ℓ be an integer such that*

$$6\epsilon^{-2} M \log(k/\delta) \le \ell \le n\,. \tag{5.3.5}$$

Let $\mathbf{R} \in \mathbb{R}^{\ell\times n}$ be a matrix which consists of a subset of ℓ rows from $\mathbf{I}_n$ where the rows are chosen uniformly at random and without replacement. Then, with probability at least $1 - 2\delta$, for $i \in [k]$:

$$\sqrt{\frac{\ell}{n}} \cdot \sqrt{1-\epsilon} \le \sigma_i(\mathbf{RW}) \le \sqrt{1+\epsilon} \cdot \sqrt{\frac{\ell}{n}}.$$

Proof. Apply Lemma 3.4 of [Tro11b] with the following choice of parameters: $\ell = \alpha M \log(k/\delta)$, $\alpha = 6/\epsilon^2$, and $\delta_{\text{tropp}} = \eta = \epsilon$. Here, ℓ, α, M, k, η are the variables named in Lemma 3.4

of [Tro11b], and δ_{tropp} plays the role of an error parameter named δ in Lemma 3.4 of [Tro11b]. The variables ϵ and δ are from our Lemma. The choice of ℓ proportional to $\log(k/\delta)$ rather than proportional to $\log(k)$, as in the original statement of Lemma 3.4, is what results in a probability proportional to δ instead of k; this can easily be seen by tracing the modified choice of ℓ through the proof of Lemma 3.4. □

Proof of Lemma 5.5. To obtain the bounds on the singular values, we combine Lemmas 5.6 and 5.7. More specifically, apply Lemma 5.7 with $\mathbf{W} = \mathbf{HDV}$ and use the bound for M from Lemma 5.6. Then, the bound on ℓ in (5.3.5), the bound on the singular values in Lemma 5.7, and a union bound together establish that, with probability at least $1 - 3\delta$,

$$\sqrt{\frac{\ell}{n}} \cdot \sqrt{1-\epsilon} \le \sigma_i(\mathbf{RHDV}) \le \sqrt{1+\epsilon} \cdot \sqrt{\frac{\ell}{n}} \quad \text{for all } i \in [k].$$

Now, multiply this inequality with $\sqrt{n/\ell}$ and recall the definition $\mathbf{\Theta} = \sqrt{n/\ell} \cdot \mathbf{RHD}$ to obtain

$$\sqrt{1-\epsilon} \le \sigma_i(\mathbf{\Theta V}) \le \sqrt{1+\epsilon} \quad \text{for all } i \in [k]. \tag{5.3.6}$$

Replacing ϵ with $\sqrt{\epsilon}$ and using the bound on ℓ given in (5.3.4) concludes the proof of the first inequality in Lemma 5.5.

The second bound in the lemma follows from the first bound after a simple algebraic manipulation. Let $\mathbf{X} = \mathbf{V}^T\mathbf{\Theta}^T \in \mathbb{R}^{k\times\ell}$ have SVD $\mathbf{X} = \mathbf{U}_\mathbf{X}\mathbf{\Sigma}_\mathbf{X}\mathbf{V}_\mathbf{X}^T$ where $\mathbf{\Sigma}_\mathbf{X} \in \mathbb{R}^{k\times k}$ is invertible, then

$$\left\|(\mathbf{V}^T\mathbf{\Theta}^T)^\dagger - (\mathbf{V}^T\mathbf{\Theta}^T)^T\right\|_2 = \left\|\mathbf{V}_\mathbf{X}\mathbf{\Sigma}_\mathbf{X}^{-1}\mathbf{U}_\mathbf{X}^T - \mathbf{V}_\mathbf{X}\mathbf{\Sigma}_\mathbf{X}\mathbf{U}_\mathbf{X}^T\right\|_2 = \left\|\mathbf{V}_\mathbf{X}(\mathbf{\Sigma}_\mathbf{X}^{-1} - \mathbf{\Sigma}_\mathbf{X})\mathbf{U}_\mathbf{X}^T\right\|_2 = \left\|\mathbf{\Sigma}_\mathbf{X}^{-1} - \mathbf{\Sigma}_\mathbf{X}\right\|_2,$$

by unitary invariance of the spectral norm. Let $\mathbf{Y} = \mathbf{\Sigma}_\mathbf{X}^{-1} - \mathbf{\Sigma}_\mathbf{X} \in \mathbb{R}^{k\times k}$. Then, for all $i = 1, \ldots, k$,

$\mathbf{Y}_{ii} = (1 - \sigma_i^2(\mathbf{X}))/\sigma_i(\mathbf{X})$. We conclude the proof with the series of estimates

$$\begin{aligned}\|\mathbf{Y}\|_2 &= \max_{1\le i\le k} \left|\mathbf{Y}_{ii}\right| = \max_{1\le i\le k} \left|\frac{1-\sigma_i^2(\mathbf{X})}{\sigma_i(\mathbf{X})}\right| \\ &= \max\left\{\frac{\left|1-\sigma_1^2(\mathbf{X})\right|}{\sigma_1(\mathbf{X})}, \frac{\left|1-\sigma_k^2(\mathbf{X})\right|}{\sigma_k(\mathbf{X})}\right\} \\ &\le \frac{\sqrt{\epsilon}}{\sqrt{1-\sqrt{\epsilon}}} \le 1.54\sqrt{\epsilon}.\end{aligned}$$

The second to last inequality follows from the lower bound in (5.3.6). □

5.3.2 SRHTs applied to general matrices

Recall our observation that the inequality

$$\|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T(\mathbf{V}_1^T\mathbf{\Theta}^T)^\dagger\|_\xi^2 \le \|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\|_\xi^2 \left\|(\mathbf{V}_1^T\mathbf{\Theta}^T)^\dagger\right\|_2^2, \tag{5.3.7}$$

together with Lemma 4.8, allows us to bound the errors of SRHT-based low-rank approximations. In the previous subsection, we collected the results necessary to bound the term $\left\|(\mathbf{V}_1^T\mathbf{\Theta}^T)^\dagger\right\|_2^2$. In this subsection, we present new results on the perturbative effects of SRHT multiplication. These allow us to bound the quantities $\|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\|_\xi^2$.

Our main tool is a generalization of Lemma 5.6 that states that the maximum column norm of a matrix to which an SRHT has been applied is, with high probability, not much larger than the root mean square of the column norms of the original matrix.

Lemma 5.8 (SRHT equalization of column norms)**.** *Suppose that* $\mathbf{A}$ *is a matrix with* n *columns, where* n *is a power of 2. Let* $\mathbf{H} \in \mathbb{R}^{n\times n}$ *be a normalized Walsh–Hadamard matrix, and* $\mathbf{D} \in \mathbb{R}^{n\times n}$ *a*

diagonal matrix of Rademacher random variables. Then for every $t \geq 0$,

$$\mathbb{P}\left\{\max_{j=1,\ldots,n} \|(\mathbf{A}\mathbf{D}\mathbf{H}^T)_{(j)}\|_2 \leq \frac{1}{\sqrt{n}}\|\mathbf{A}\|_{\mathrm{F}} + \frac{t}{\sqrt{n}}\|\mathbf{A}\|_2\right\} \geq 1 - n \cdot \mathrm{e}^{-t^2/8}.$$

Recall that $(\mathbf{A}\mathbf{D}\mathbf{H}^T)_{(j)}$ *denotes the* j*th column of* $\mathbf{A}\mathbf{D}\mathbf{H}^T$.

Proof. Our proof of Lemma 5.8 is essentially that of Lemma 5.6 in [Tro11b], with attention paid to the fact that $\mathbf{A}$ is no longer assumed to have orthonormal columns.

Lemma 5.8 follows immediately from the observation that the norm of any one column of $\mathbf{A}\mathbf{D}\mathbf{H}^T$ is a convex Lipschitz function of a Rademacher vector. Consider the norm of the jth column of $\mathbf{A}\mathbf{D}\mathbf{H}^T$ as a function of $\boldsymbol{\varepsilon}$, where $\mathbf{D} = \mathrm{diag}(\boldsymbol{\varepsilon})$:

$$f_j(\boldsymbol{\varepsilon}) = \|\mathbf{A}\mathbf{D}\mathbf{H}^T\mathbf{e}_j\| = \left\|\mathbf{A}\,\mathrm{diag}(\boldsymbol{\varepsilon})\mathbf{h}_j\right\|_2 = \left\|\mathbf{A}\,\mathrm{diag}(\mathbf{h}_j)\boldsymbol{\varepsilon}\right\|_2,$$

where $\mathbf{h}_j$ denotes the jth column of $\mathbf{H}^T$. Evidently f_j is convex. Furthermore,

$$|f_j(\mathbf{x}) - f_j(\mathbf{y})| \leq \left\|\mathbf{A}\,\mathrm{diag}(\mathbf{h}_j)(\mathbf{x}-\mathbf{y})\right\|_2 \leq \|\mathbf{A}\|_2 \left\|\mathrm{diag}(\mathbf{h}_j)\right\|_2 \left\|\mathbf{x}-\mathbf{y}\right\|_2 = \frac{1}{\sqrt{n}}\|\mathbf{A}\|_2 \left\|\mathbf{x}-\mathbf{y}\right\|_2,$$

where we used the triangle inequality and the fact that $\left\|\mathrm{diag}(\mathbf{h}_j)\right\|_2 = \|\mathbf{h}_j\|_\infty = n^{-1/2}$. Thus f_j is convex and Lipschitz with Lipschitz constant of at most $n^{-1/2}\|\mathbf{A}\|_2$.

We calculate

$$\begin{aligned}
\mathbb{E}\left[f_j(\boldsymbol{\varepsilon})\right] \leq \mathbb{E}\left[f_j(\boldsymbol{\varepsilon})^2\right]^{1/2} &= \left[\mathrm{Tr}\left(\mathbf{A}\,\mathrm{diag}(\mathbf{h}_j)\,\mathbb{E}\left[\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}^*\right]\mathrm{diag}(\mathbf{h}_j)\mathbf{A}^T\right)\right]^{1/2} \\
&= \left[\mathrm{Tr}\left(\frac{1}{n}\mathbf{A}\mathbf{A}^T\right)\right]^{1/2} \\
&= \frac{1}{\sqrt{n}}\|\mathbf{A}\|_{\mathrm{F}}.
\end{aligned}$$

It now follows from Lemma 4.1 that, for all $j = 1, 2, \ldots, n$, the norm of the jth column of $\mathbf{ADH}^T$ satisfies the tail bound

$$\mathbb{P}\left\{ \left\| \mathbf{ADH}^T \mathbf{e}_j \right\|_2 \geq \frac{1}{\sqrt{n}} \|\mathbf{A}\|_{\mathrm{F}} + \frac{t}{\sqrt{n}} \|\mathbf{A}\|_2 \right\} \leq \mathrm{e}^{-t^2/8}.$$

Taking a union bound over all columns of $\mathbf{ADH}^T$, we conclude that

$$\mathbb{P}\left\{ \max_{j=1,\ldots,n} \left\| (\mathbf{ADH}^T)_{(j)} \right\|_2 \geq \frac{1}{\sqrt{n}} \|\mathbf{A}\|_{\mathrm{F}} + \frac{t}{\sqrt{n}} \|\mathbf{A}\|_2 \right\} \leq n \cdot \mathrm{e}^{-t^2/8}.$$

□

Our next lemma shows that the SRHT does not substantially increase the spectral norm of a matrix.

Lemma 5.9 (SRHT-based subsampling in the spectral norm)**.** *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ have rank ρ, and assume that n is a power of 2. For some $r < n$, let $\mathbf{\Theta} \in \mathbb{R}^{\ell \times n}$ be an SRHT matrix. Fix a failure probability $0 < \delta < 1$. Then,*

$$\mathbb{P}\left\{ \left\| \mathbf{A}\mathbf{\Theta}^T \right\|_2^2 \leq 5 \|\mathbf{A}\|_2^2 + \frac{\log(\rho/\delta)}{\ell} \left(\|\mathbf{A}\|_{\mathrm{F}} + \sqrt{8 \log(n/\delta)}\, \|\mathbf{A}\|_2 \right)^2 \right\} \geq 1 - 2\delta.$$

To establish Lemma 5.9, we use the upper Chernoff bound for sums of matrices stated in Lemma 4.2.

Proof of Lemma 5.9. Write the SVD of $\mathbf{A}$ as $\mathbf{U\Sigma V}^T$, where $\mathbf{\Sigma} \in \mathbb{R}^{\rho \times \rho}$, and observe that the spectral norm of $\mathbf{A\Theta}^T$ is the same as that of $\sqrt{n/\ell} \cdot \mathbf{\Sigma V}^T \mathbf{\Theta}^T$.

We control the norm of $\mathbf{\Sigma V}^T \mathbf{\Theta}^T$ by considering the maximum singular value of its Gram matrix. Define $\mathbf{M} = \mathbf{\Sigma V}^T \mathbf{DH}^T$, so that $\mathbf{\Sigma V}^T \mathbf{\Theta}^T = \sqrt{n/\ell} \cdot \mathbf{MR}^T$, and let $\mathbf{G}$ be the $\rho \times \rho$ Gram

matrix of $\mathbf{MR}^T$:

$$\mathbf{G} = \mathbf{MR}^T(\mathbf{MR}^T)^T.$$

Evidently

$$\lambda_{\max}(\mathbf{G}) = \frac{\ell}{n}\left\|\mathbf{\Sigma V}^T\mathbf{\Theta}^T\right\|_2^2. \tag{5.3.8}$$

Recall that $\mathbf{M}_{(j)}$ denotes the jth column of $\mathbf{M}$. If we denote by C the random set of r coordinates to which $\mathbf{R}$ restricts, then

$$\mathbf{G} = \sum\nolimits_{j\in C}\mathbf{M}_{(j)}\mathbf{M}_{(j)}^T.$$

Thus $\mathbf{G}$ is a sum of r random matrices $\mathbf{X}_1,\ldots,\mathbf{X}_r$ sampled without replacement from the set $\mathscr{X} = \{\mathbf{M}_{(j)}\mathbf{M}_{(j)}^T : j = 1,2,\ldots,n\}$. There are two sources of randomness in $\mathbf{G}$: the subsampling matrix $\mathbf{R}$ and the Rademacher random variables on the diagonal of $\mathbf{D}$.

Set

$$B = \frac{1}{n}\left(\|\mathbf{\Sigma}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{\Sigma}\|_2\right)^2$$

and let E be the event

$$\max\nolimits_{j=1,\ldots,n}\|\mathbf{M}_{(j)}\|_2^2 \le B.$$

By Lemma 5.8, E occurs with probability at least $1-\delta$. When E holds, for all $j = 1,2,\ldots,n$,

$$\lambda_{\max}\left(\mathbf{M}_{(j)}\mathbf{M}_{(j)}^T\right) = \|\mathbf{M}_{(j)}\|_2^2 \le B,$$

so $\mathbf{G}$ is a sum of random positive-semidefinite matrices each of whose norms is bounded by B. Note that the event E is entirely determined by the random matrix $\mathbf{D}$; in particular E is independent of $\mathbf{R}$.

Conditioning on E, the randomness in $\mathbf{R}$ allows us to use the upper matrix Chernoff bound of Lemma 4.2 to control the maximum eigenvalue of $\mathbf{G}$. We observe that

$$\mu_{\max} = \ell \cdot \lambda_{\max}\left(\mathbb{E}\left[\mathbf{X}_1\right]\right) = \frac{\ell}{n}\lambda_{\max}\left(\sum\nolimits_{j=1}^{n} \mathbf{M}_{(j)}\mathbf{M}_{(j)}^T\right) = \frac{\ell}{n}\left\|\mathbf{\Sigma}\right\|_2^2.$$

Take the parameter ν in Lemma 4.2 to be

$$\nu = 4 + \frac{B}{\mu_{\max}}\log(\rho/\delta)$$

to obtain the relation

$$\begin{aligned}\mathbb{P}\left\{\lambda_{\max}(\mathbf{G}) \geq 5\mu_{\max} + B\log(\rho/\delta) \mid E\right\} &\leq (\rho - k)\cdot \mathrm{e}^{[\delta - (1+\nu)\log(1+\nu)](\mu_{\max}/B)} \\ &\leq \rho \cdot \mathrm{e}^{(1-(5/4)\log 5)\delta(\mu_{\max}/B)} \\ &\leq \rho \cdot \mathrm{e}^{(-(5/4)\log 5 - 1)\log(\rho/\delta)} < \delta. \end{aligned} \tag{5.3.9}$$

The second inequality holds because $\nu \geq 4$ implies that $(1+\nu)\log(1+\nu) \geq \nu \cdot (5/4)\log 5$.

We have conditioned on E, the event that the squared norms of the columns of $\mathbf{M}$ are all smaller than B. Thus, substituting the values of B and $\mu_{\max}$ into (5.3.9), we find that

$$\mathbb{P}\left\{\lambda_{\max}(\mathbf{G}) \geq \frac{\ell}{n}\left(5\left\|\mathbf{\Sigma}\right\|_2^2 + \frac{\log(\rho/\delta)}{\ell}\left(\left\|\mathbf{\Sigma}\right\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\left\|\mathbf{\Sigma}\right\|_2\right)^2\right)\right\} \leq 2\delta.$$

Use equation (5.3.8) to wrap up. □

Similarly, the SRHT is unlikely to substantially increase the Frobenius norm of a matrix.

Lemma 5.10 (SRHT-based subsampling in the Frobenius norm)**.** *Assume n is a power of 2. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$, and let $\mathbf{\Theta} \in \mathbb{R}^{\ell\times n}$ be an SRHT matrix for some $\ell < n$. Fix a failure probability $0 < \delta < 1$.*

Then, for any $\eta \geq 0$,

$$\mathbb{P}\left\{\left\|\mathbf{A}\boldsymbol{\Theta}^T\right\|_{\mathrm{F}}^2 \leq (1+\eta)\left\|\mathbf{A}\right\|_{\mathrm{F}}^2\right\} \geq 1 - \left[\frac{\mathrm{e}^{\eta}}{(1+\eta)^{1+\eta}}\right]^{\ell/\left(1+\sqrt{8\log(n/\delta)}\right)^2} - \delta.$$

Proof. Let $c_j = (n/\ell)\cdot\left\|(\mathbf{A}\mathbf{D}\mathbf{H}^T)_{(j)}\right\|_2^2$ denote the squared norm of the jth column of $\sqrt{n/\ell}\cdot\mathbf{A}\mathbf{D}\mathbf{H}^T$. Then, since right multiplication by $\mathbf{R}^T$ samples columns uniformly at random without replacement,

$$\left\|\mathbf{A}\boldsymbol{\Theta}^T\right\|_{\mathrm{F}}^2 = \frac{n}{\ell}\left\|\mathbf{A}\mathbf{D}\mathbf{H}^T\mathbf{R}^T\right\|_{\mathrm{F}}^2 = \sum\nolimits_{i=1}^{\ell} X_i \tag{5.3.10}$$

where the random variables X_i are chosen randomly without replacement from the set $\{c_j\}_{j=1}^n$. There are two independent sources of randomness in this sum: the choice of summands, which is determined by $\mathbf{R}$, and the magnitudes of the $\{c_j\}$, which are determined by $\mathbf{D}$.

To bound this sum, we first condition on the event E that each c_j is bounded by a quantity B which depends only on the random matrix $\mathbf{D}$. Then

$$\mathbb{P}\left\{\sum\nolimits_{i=1}^{\ell} X_i \geq (1+\eta)\sum\nolimits_{i=1}^{\ell}\mathbb{E}X_i\right\} \leq \mathbb{P}\left\{\sum\nolimits_{i=1}^{\ell} X_i \geq (1+\eta)\sum\nolimits_{i=1}^{\ell}\mathbb{E}X_i \mid E\right\} + \mathbb{P}(E^c).$$

To select B, we observe that Lemma 5.8 implies that with probability $1-\delta$, the entries of $\mathbf{D}$ are such that

$$\max\nolimits_j c_j \leq \frac{n}{\ell}\cdot\frac{1}{n}(\|\mathbf{A}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{A}\|_2)^2 \leq \frac{1}{\ell}(1+\sqrt{8\log(n/\delta)})^2\,\|\mathbf{A}\|_{\mathrm{F}}^2.$$

Accordingly, we take

$$B = \frac{1}{\ell}(1+\sqrt{8\log(n/\delta)})^2\,\|\mathbf{A}\|_{\mathrm{F}}^2,$$

thereby arriving at the bound

$$\mathbb{P}\left\{\sum_{i=1}^{\ell} X_i \geq (1+\eta)\sum_{i=1}^{\ell} \mathbb{E}X_i\right\} \leq \mathbb{P}\left\{\sum_{i=1}^{\ell} X_i \geq (1+\eta)\sum_{i=1}^{\ell} \mathbb{E}X_i \mid E\right\} + \delta. \quad (5.3.11)$$

After conditioning on $\mathbf{D}$, we observe that the randomness remaining on the right-hand side of (5.3.11) is due to the choice of the summands X_i, which is determined by $\mathbf{R}$. We address this randomness by applying a scalar Chernoff bound (Lemma 4.2 with $k = 1$). To do so, we need $\mu_{\max}$, the expected value of the sum; this is an elementary calculation:

$$\mathbb{E}X_1 = n^{-1}\sum_{j=1}^{n} c_j = \frac{1}{\ell}\|\mathbf{A}\|_{\mathrm{F}}^2,$$

so $\mu_{\max} = \ell\mathbb{E}X_1 = \|\mathbf{A}\|_{\mathrm{F}}^2$.

Applying Lemma 4.2 conditioned on E, we conclude that

$$\mathbb{P}\left\{\left\|\mathbf{A}\mathbf{\Theta}^T\right\|_{\mathrm{F}}^2 \geq (1+\eta)\|\mathbf{A}\|_{\mathrm{F}}^2 \mid E\right\} \leq \left[\frac{\mathrm{e}^{\eta}}{(1+\eta)^{1+\eta}}\right]^{\ell/(1+\sqrt{8\log(n/\delta)})^2} + \delta$$

for $\eta \geq 0$. □

Finally, we prove a novel result on approximate matrix multiplication involving SRHT matrices.

Lemma 5.11 (SRHT for approximate matrix multiplication)**.** *Assume n is a power of 2. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ and $\mathbf{B} \in \mathbb{R}^{n\times p}$. For some $\ell < n$, let $\mathbf{\Theta} \in \mathbb{R}^{\ell\times n}$ be an SRHT matrix. Fix a failure probability $0 < \delta < 1$. Assume R satisfies $0 \leq R \leq \sqrt{\ell}/(1+\sqrt{8\log(n/\delta)})$. Then,*

$$\mathbb{P}\left\{\left\|\mathbf{A}\mathbf{\Theta}^T\mathbf{\Theta}\mathbf{B} - \mathbf{A}\mathbf{B}\right\|_{\mathrm{F}} \leq 2(R+1)\frac{\|\mathbf{A}\|_{\mathrm{F}}\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{A}\|_{\mathrm{F}}\|\mathbf{B}\|_2}{\sqrt{\ell}}\right\} \geq 1 - \mathrm{e}^{-R^2/4} - 2\delta.$$

Remark 5.12. Recall that the stable rank $\mathrm{sr}(\mathbf{A}) = \|\mathbf{A}\|_{\mathrm{F}}^2 / \|\mathbf{A}\|_2^2$ reflects the decay of the spectrum of the matrix $\mathbf{A}$. The event under consideration in Lemma 5.11 can be rewritten as a bound on the relative error of the approximation $\mathbf{A}\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{B}$ to the product $\mathbf{AB}$:

$$\frac{\left\|\mathbf{A}\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{B} - \mathbf{AB}\right\|_{\mathrm{F}}}{\|\mathbf{AB}\|_{\mathrm{F}}} \leq \frac{\|\mathbf{A}\|_{\mathrm{F}}\|\mathbf{B}\|_{\mathrm{F}}}{\|\mathbf{AB}\|_{\mathrm{F}}} \cdot \frac{R+2}{\sqrt{\ell}} \cdot \left(1 + \frac{\sqrt{8\log(n/\delta)}}{\mathrm{sr}(\mathbf{B})}\right).$$

In this form, we see that the relative error is controlled by the deterministic condition number for the matrix multiplication problem as well as the stable rank of $\mathbf{B}$ and the number of column samples ℓ. Since the roles of $\mathbf{A}$ and $\mathbf{B}$ in this bound can be interchanged, in fact we have the bound

$$\frac{\left\|\mathbf{A}\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{B} - \mathbf{AB}\right\|_{\mathrm{F}}}{\|\mathbf{AB}\|_{\mathrm{F}}} \leq \frac{\|\mathbf{A}\|_{\mathrm{F}}\|\mathbf{B}\|_{\mathrm{F}}}{\|\mathbf{AB}\|_{\mathrm{F}}} \cdot \frac{R+2}{\sqrt{\ell}} \cdot \left(1 + \frac{\sqrt{8\log(n/\delta)}}{\max(\mathrm{sr}(\mathbf{B}), \mathrm{sr}(\mathbf{A}))}\right).$$

To prove the lemma, we use Lemma 4.3, our result for approximate matrix multiplication via uniform sampling (without replacement) of the columns and the rows of the two matrices involved in the product. Lemma 5.11 is simply a specific application of this generic result. We mention that Lemma 3.2.8 in [Dri02] gives a similar result for approximate matrix multiplication which, however, gives a bound on the expected value of the error term, while our Lemma 5.11 gives a comparable bound that holds with high probability.

Proof of Lemma 5.11. Let $\mathbf{X} = \mathbf{ADH}^T$ and $\mathbf{Y} = \mathbf{HDB}$ and form $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ according to Lemma 4.3. Then, $\mathbf{XY} = \mathbf{AB}$ and

$$\left\|\mathbf{A}\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{B} - \mathbf{AB}\right\|_{\mathrm{F}} = \left\|\hat{\mathbf{X}}\hat{\mathbf{Y}} - \mathbf{XY}\right\|_{\mathrm{F}}.$$

To apply Lemma 4.3, we first condition on the event that the SRHT equalizes the column norms

of our matrices. Namely, we observe that, from Lemma 5.8, with probability at least $1 - 2\delta$,

$$\max_i \|\mathbf{X}_{(i)}\|_2 \le \frac{1}{\sqrt{n}}(\|\mathbf{A}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{A}\|_2), \text{ and} \tag{5.3.12}$$
$$\max_i \|\mathbf{Y}^{(i)}\|_2 \le \frac{1}{\sqrt{n}}(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{B}\|_2).$$

We choose the parameters σ and B in Lemma 4.3. Set

$$\sigma^2 = \frac{4}{\ell}(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{B}\|_2)^2\,\|\mathbf{A}\|_{\mathrm{F}}^2. \tag{5.3.13}$$

In view of (5.3.12),

$$\sigma^2 = 4\frac{n}{\ell}\cdot\frac{(\|\mathbf{Y}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{Y}\|_2)^2}{n}\,\|\mathbf{X}\|_{\mathrm{F}}^2 \ge 4\frac{n}{\ell}\sum\nolimits_{i=1}^{n}\|\mathbf{X}_{(i)}\|_2^2\|\mathbf{Y}^{(i)}\|_2^2$$

so this choice of σ satisfies the inequality condition of Lemma 4.3. Next we choose

$$B = \frac{2}{\ell}(\|\mathbf{A}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{A}\|_2)(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)}\,\|\mathbf{B}\|_2).$$

Again, because of (5.3.12), B satisfies the requirement $B \ge \frac{2n}{\ell}\max_i \|\mathbf{X}_{(i)}\|_2\|\mathbf{Y}^{(i)}\|_2$.

For simplicity, abbreviate $\gamma = 8\log(n/\delta)$. With these choices for σ^2 and B,

$$\begin{aligned}\frac{\sigma^2}{B} &= \frac{2\,\|\mathbf{A}\|_{\mathrm{F}}^2(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{B}\|_2)^2}{(\|\mathbf{A}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{A}\|_2)(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{B}\|_2)} \\ &\ge \frac{2\,\|\mathbf{A}\|_{\mathrm{F}}^2(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{B}\|_2)^2}{(\|\mathbf{A}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{A}\|_{\mathrm{F}})(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{B}\|_2)} \\ &= \frac{2\,\|\mathbf{A}\|_{\mathrm{F}}(\|\mathbf{B}\|_{\mathrm{F}} + \sqrt{\gamma}\,\|\mathbf{B}\|_2)}{1 + \sqrt{\gamma}}.\end{aligned}$$

Now, referring to (5.3.13), identify the numerator as $\sqrt{\ell}\sigma$ to see that

$$\frac{\sigma^2}{B} \geq \frac{\sqrt{\ell}\sigma}{1+\sqrt{8\log(n/\delta)}}.$$

Apply Lemma 4.3 to see that, when (5.3.12) holds and $0 \leq R\sigma \leq \sigma^2/B$,

$$\mathbb{P}\left\{\left\|\mathbf{A}\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{B}-\mathbf{A}\mathbf{B}\right\|_{\mathrm{F}} \geq (R+1)\sigma\right\} \leq \exp\left(-\frac{R^2}{4}\right).$$

From our lower bound on σ^2/B, we know that the condition $R\sigma \leq \sigma^2/B$ is satisfied when

$$R \leq \sqrt{\ell}/(1+\sqrt{8\log(n/\delta)}).$$

We established above that (5.3.12) holds with probability at least $1-2\delta$. From these two facts, it follows that when $0 \leq R \leq \sqrt{\ell}/(1+\sqrt{8\log(n/\delta)})$,

$$\mathbb{P}\left\{\left\|\mathbf{A}\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{B}-\mathbf{A}\mathbf{B}\right\|_{\mathrm{F}} \geq (R+1)\sigma\right\} \leq \exp\left(-\frac{R^2}{4}\right)+2\delta.$$

The tail bound given in the statement of Lemma 5.11 follows when we substitute our estimate of σ. □

5.4 Proof of the quality of approximation guarantees

With the necessary preliminaries in hand, we now proceed to the proof of our main results: bounds on the spectral and Frobenius-norm residual and forward errors of low-rank approximations generated using Algorithms 5.1 and 5.2 with SRHT sampling matrices. We note that prior works have provided only residual error bounds [WLRT08, HMT11, NDT09].

Our first result bounds the spectral-norm errors.

Theorem 5.13. *Assume n is a power of 2. Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ have rank ρ and fix an integer k satisfying $2 \le k < \rho$. Let $0 < \epsilon < 1/3$ be an accuracy parameter, $0 < \delta < 1$ be a failure probability, and $C \ge 1$ be any specified constant. If $\mathbf{\Theta} \in \mathbb{R}^{\ell \times n}$ is an SRHT matrix and ℓ satisfies*

$$6C^2\epsilon^{-1}\left[\sqrt{k} + \sqrt{8\log(n/\delta)}\right]^2 \log(k/\delta) \le \ell \le n, \tag{5.4.1}$$

then the approximation $\mathbf{P}_{\mathbf{A}\mathbf{\Theta}^T}\mathbf{A}$ generated by Algorithm 5.1 with $\mathbf{S} = \mathbf{\Theta}^T$ satisfies the residual error bound

$$\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A}\mathbf{\Theta}^T}\mathbf{A}\right\|_2 \le \left(4 + \sqrt{\frac{3\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{3\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \quad \text{(i)}$$

and the forward error bound

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{A}\mathbf{\Theta}^T}\mathbf{A}\right\|_2 \le \left(4 + \sqrt{\frac{3\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{3\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \quad \text{(ii)}$$

simultaneously, with probability at least $1 - 5\delta$.

Likewise, the approximation $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})$ generated by Algorithm 5.2 with $\mathbf{S} = \mathbf{\Theta}^T$ satisfies the residual error bound

$$\left\|\mathbf{A} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\right\|_2 \le \left(6 + \sqrt{\frac{6\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{6\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \tag{iii}$$

and the forward error bound

$$\left\|\mathbf{A}_k - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\right\|_2 \leq \left(7 + \sqrt{\frac{6\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{6\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \tag{iv}$$

simultaneously, with probability at least $1 - 5\delta$.

Proof. First we derive the residual error bounds. Lemma 5.5 implies that, when ℓ satisfies (5.4.1),

$$\left\|(\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger\right\|_2^2 \leq (1 - \sqrt{\epsilon})^{-1}$$

with probability at least $1 - 3\delta$. Consequently, $\mathbf{V}_1^T \mathbf{\Theta}^T$ has full row-rank and Lemma 4.8 with $\mathbf{S} = \mathbf{\Theta}^T \in \mathbb{R}^{n\times\ell}$ and $\xi = 2$ applies with the same probability, yielding a bound on the residual error of the approximation $\mathbf{P}_{\mathbf{AS}}\mathbf{A}$:

$$\begin{aligned}
\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A}\mathbf{\Theta}^T}\mathbf{A}\right\|_2^2 &\leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + \|\mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T (\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger\|_2^2 \\
&\leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + \|\mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T\|_2^2 \|(\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger\|_2^2 \\
&\leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + (1 - \sqrt{\epsilon})^{-1} \left\|\mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T\right\|_2^2 .
\end{aligned} \tag{5.4.2}$$

We now provide an upper bound for the quantity

$$Z = \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + (1 - \sqrt{\epsilon})^{-1} \left\|\mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T\right\|_2^2 .$$

After applying Lemma 5.9 to estimate the term $\left\|\mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T\right\|_2^2$, we see that the estimate

$$Z \leq \left(1 + \frac{5}{1 - \sqrt{\epsilon}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + \frac{\log(\rho/\delta)}{(1 - \sqrt{\epsilon})\ell} \left(\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)} \left\|\mathbf{A} - \mathbf{A}_k\right\|_2\right)^2$$

holds with probability at least $1 - 5\delta$. Our assumption that $\epsilon < 1/3$ ensures that $(1 - \sqrt{\epsilon})^{-1} < 3$, so

$$Z \leq 16 \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + \frac{3\log(\rho/\delta)}{\ell} \left(\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} + \sqrt{8\log(n/\delta)} \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 \right)^2. \tag{5.4.3}$$

Introduce this estimate for Z into (5.4.2), use the subadditivity of the square-root function, and rearrange the spectral and Frobenius norm terms to arrive at Eqn. (i) in the theorem:

$$\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A}\boldsymbol{\Theta}^{\mathrm{T}}}\mathbf{A}\right\|_2 \leq \left(4 + \sqrt{\frac{3\log(n/\delta)\log(\rho/\delta)}{\ell}} \right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{3\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}.$$

We now establish the residual error bound for the approximation $\boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})$. We begin by recalling that Lemma 4.7 states that

$$\|\mathbf{A} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\|_2^2 \leq 2\|\mathbf{A} - \boldsymbol{\Pi}^{2}_{\mathbf{AS},k}(\mathbf{A})\|_2^2.$$

Lemma 4.8 can be used to bound the right-hand side quantity:

$$2\|\mathbf{A} - \boldsymbol{\Pi}^{2}_{\mathbf{AS},k}(\mathbf{A})\|_2^2 \leq 2\left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + 2\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\|_2^2.$$

We have already encountered the right-hand side of this expression, without the factor of two, in (5.4.2). It follows that

$$\|\mathbf{A} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\|_2^2 \leq 2Z.$$

Introduce our earlier estimate for Z, given in (5.4.3), into this inequality; apply the submulti-

plicativity of the square-root function; and rearrange terms to obtain Eqn. (iii) in the theorem:

$$\left\|\mathbf{A} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\right\|_2 \leq \left(6 + \sqrt{\frac{6\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{6\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}.$$

Once again, this bound holds with probability at least $1 - 5\delta$.

The forward error bounds follow in a similar manner. To establish Eqn. (ii) in the theorem, observe that Lemma 4.9 gives the bound

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{A}\boldsymbol{\Theta}^T}\mathbf{A}\right\|_2^2 \leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_2^2 + \|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T(\mathbf{V}_1^T\boldsymbol{\Theta}^T)^\dagger\|_2^2.$$

Once again, we observe that we encountered the right-hand side of this expression in (5.4.2), where we argued that it is bounded by Z, so

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{A}\boldsymbol{\Theta}^T}\mathbf{A}\right\|_2^2 \leq Z.$$

Introduce into this inequality the estimate for Z given in (5.4.3), apply the submultiplicativity of the square-root function, and rearrange terms to obtain Eqn. (ii) in the theorem:

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{A}\boldsymbol{\Theta}^T}\mathbf{A}\right\|_2^2 \leq \left(4 + \sqrt{\frac{3\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{3\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}.$$

To establish Eqn. (iv) in the theorem, observe that

$$\begin{aligned}
\|\mathbf{A}_k - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\|_2 &\leq \|\mathbf{A}_k + \mathbf{A}_{\rho-k} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A}) + \mathbf{A}_{\rho-k}\|_2 \\
&\leq \|\mathbf{A} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\|_2 + \|\mathbf{A}_{\rho-k}\|_2 \\
&= \|\mathbf{A} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\|_2 + \|\mathbf{A} - \mathbf{A}_{\rho-k}\|_2.
\end{aligned}$$

The first term on the right-hand side of this inequality is simply the forward error of the approximation $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})$. Introduce our bound on this error, given by Eqn. (iii) of the theorem, into this inequality to obtain the desired bound:

$$\|\mathbf{A}_k - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\|_2 \leq \left(7 + \sqrt{\frac{6\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_2 + \sqrt{\frac{6\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}.$$

This bound holds with probability at least $1 - 5\delta$. □

Our second result bounds the Frobenius-norm errors of the SRHT-based low-rank approximation algorithms.

Theorem 5.14. *Assume n is a power of 2. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ have rank ρ and fix an integer k satisfying $2 \leq k < \rho$. Let $0 < \epsilon < 1/3$ be an accuracy parameter, $0 < \delta < 1$ be a failure probability, and $C \geq 1$ be any specified constant. If $\mathbf{\Theta} \in \mathbb{R}^{\ell\times n}$ is an SRHT matrix and ℓ satisfies*

$$6C^2\epsilon^{-1}\left[\sqrt{k} + \sqrt{8\log(n/\delta)}\right]^2 \log(k/\delta) \leq \ell \leq n, \tag{5.4.4}$$

then the approximation $\mathbf{P}_{\mathbf{A\Theta}^T}\mathbf{A}$ generated by Algorithm 5.1 with $\mathbf{S} = \mathbf{\Theta}^T$ satisfies the residual error bound

$$\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A\Theta}^T}\mathbf{A}\right\|_{\mathrm{F}} \leq (1 + 11\epsilon) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \tag{i}$$

and the forward error bound

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{A\Theta}^T}\mathbf{A}\right\|_{\mathrm{F}} \leq (1 + 11\epsilon) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \tag{ii}$$

simultaneously, with probability at least $1 - \delta^{C^2\log(k/\delta)/4} - 7\delta$.

Likewise, the approximation $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})$ *generated by Algorithm 5.2 with* $\mathbf{S} = \mathbf{\Theta}^T$ *satisfies the residual error bound*

$$\left\|\mathbf{A} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\right\|_{\mathrm{F}} \leq (1 + 11\epsilon) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \tag{iii}$$

and the forward error bound

$$\left\|\mathbf{A}_k - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{AS},k}(\mathbf{A})\right\|_{\mathrm{F}} \leq (2 + 11\epsilon) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} \tag{iv}$$

simultaneously, with probability at least $1 - \delta^{C^2 \log(k/\delta)/4} - 7\delta$.

Proof. We first establish the residual error bounds. Because ℓ satisfies (5.4.4), Lemma 5.5 implies that with probability at least $1 - 3\delta$,

$$\operatorname{rank}(\mathbf{V}_1^T \mathbf{\Theta}^T) = k;$$

so, Lemma 4.8 applies with the same probability, yielding

$$\left\|\mathbf{A} - \mathbf{P}_{\mathbf{A}\mathbf{\Theta}^T}\mathbf{A}\right\|_{\mathrm{F}}^2 \leq \left\|\mathbf{A} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{A}\mathbf{\Theta}^T,k}(\mathbf{A})\right\|_{\mathrm{F}}^2 \leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}^2 + \left\|\mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T (\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger\right\|_{\mathrm{F}}^2. \tag{5.4.5}$$

We complete the estimate by bounding the second term in the right hand side of the above

inequality. Justifications appear below.

$$\begin{aligned} S &:= \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T (\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger \right\|_{\mathrm{F}}^2 \\ &\leq 2 \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T \mathbf{\Theta} \mathbf{V}_1 \right\|_{\mathrm{F}}^2 + 2 \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T ((\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger - (\mathbf{V}_1^T \mathbf{\Theta}^T)^T) \right\|_{\mathrm{F}}^2 \qquad (5.4.6) \\ &\leq 2 \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T \mathbf{\Theta} \mathbf{V}_1 \right\|_{\mathrm{F}}^2 + 2 \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{\Theta}^T \right\|_{\mathrm{F}}^2 \left\| (\mathbf{V}_1^T \mathbf{\Theta}^T)^\dagger - (\mathbf{V}_1^T \mathbf{\Theta}^T)^T \right\|_2^2 \\ &\leq 8\epsilon \cdot \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \right\|_{\mathrm{F}}^2 + 2 \cdot \left(\frac{11}{4} \left\| \mathbf{\Sigma}_2 \mathbf{V}_2^T \right\|_{\mathrm{F}}^2 \right) \cdot (2.38\epsilon) \qquad (5.4.7) \\ &\leq 22\epsilon \cdot \left\| \mathbf{\Sigma}_2 \right\|_{\mathrm{F}}^2 . \qquad (5.4.8) \end{aligned}$$

In (5.4.6) we used the fact that $\|\mathbf{X}+\mathbf{Y}\|_{\mathrm{F}}^2 \leq 2\|\mathbf{X}\|_{\mathrm{F}}^2 + 2\|\mathbf{Y}\|_{\mathrm{F}}^2$ for any two matrices $\mathbf{X}$ and $\mathbf{Y}$. The first estimate in (5.4.7) is justified by an application of Lemma 5.11 on SRHT-based approximate matrix multiplication; the second estimate is an application of Lemma 5.10, which predicts the effect that postmultiplication by an SRHT matrix has upon the Frobenius norm of a matrix; and the third estimate is an application of Lemma 5.5, which predicts the effect of postmultiplication by an SRHT matrix on the singular values of a matrix with orthonormal rows. We postpone a detailed discussion of the calculations involved in these estimations until after we have established the residual error bounds.

Combining (5.4.5) with the bound on S, we obtain

$$\left\| \mathbf{A} - \mathbf{P}_{\mathbf{AS}} \mathbf{A} \right\|_{\mathrm{F}}^2 \leq \| \mathbf{A} - \mathbf{\Pi}_{\mathbf{AS},k}^{\mathrm{F}}(\mathbf{A}) \|_{\mathrm{F}}^2 \leq (1 + 22\epsilon) \cdot \left\| \mathbf{A} - \mathbf{A}_k \right\|_{\mathrm{F}}^2 . \qquad (5.4.9)$$

Since $1 + 2x \leq (1+x)^2$ when x is positive, it follows that $\sqrt{1+2x} \leq 1 + x$ when x is positive. In particular, $\sqrt{1+22\epsilon} \leq 1 + 11\epsilon$. Introduce this observation into (5.4.9) to conclude that

$$\left\| \mathbf{A} - \mathbf{P}_{\mathbf{AS}} \mathbf{A} \right\|_{\mathrm{F}} \leq \| \mathbf{A} - \mathbf{\Pi}_{\mathbf{AS},k}^{\mathrm{F}}(\mathbf{A}) \|_{\mathrm{F}} \leq (1 + 11\epsilon) \cdot \left\| \mathbf{A} - \mathbf{A}_k \right\|_{\mathrm{F}} .$$

Thus we have established Eqns. (i) and (iii) in the theorem.

We now supply the details of the manipulations in (5.4.7). To justify the first estimate, notice that $\mathbf{V}_2^T\mathbf{V}_1 = \mathbf{0}$. Next use Lemma 5.11 with $R = C\sqrt{\log(k/\delta)}$. From the lower bound (5.4.4) on ℓ, we have that

$$\frac{\sqrt{\ell}}{1+\sqrt{8\log(n/\delta)}} \geq \sqrt{6\epsilon^{-1}} \cdot \frac{\sqrt{k}+\sqrt{8\log(n/\delta)}}{1+\sqrt{8\log(n/\delta)}} \cdot C\sqrt{\log(k/\delta)} > R > 0,$$

so this choice of R satisfies the requirements of Lemma 5.11. Apply Lemma 5.11 to obtain

$$\mathbb{P}\left\{ \left\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{V}_1\right\|_{\mathrm{F}}^2 \leq 4(R+1)^2\frac{(\sqrt{k}+\sqrt{8\log(n/\delta)})^2}{\ell}\left\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\right\|_{\mathrm{F}}^2 \right\} \geq 1 - \mathrm{e}^{-R^2/4} - 2\delta.$$

Recall that $R = C\sqrt{\log(k/\delta)}$. Use the lower bound (5.4.4) on ℓ to justify the estimate

$$\begin{aligned} 4(R+1)^2\frac{\left[\sqrt{k}+\sqrt{8\log(n/\delta)}\right]^2}{\ell} &\leq 4(R+1)^2\frac{\left[\sqrt{k}+\sqrt{8\log(n/\delta)}\right]^2}{6C^2\epsilon^{-1}\left[\sqrt{k}+\sqrt{8\log(n/\delta)}\right]^2\log(k/\delta)} \\ &= \frac{2\epsilon}{3}\cdot\frac{(C\sqrt{\log(k/\delta)}+1)^2}{C^2\log(k/\delta)} \\ &\leq \frac{2\epsilon}{3}\left(1+\frac{1}{C\sqrt{\log(k/\delta)}}\right)^2. \end{aligned}$$

This estimate implies that

$$\mathbb{P}\left\{ \left\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{V}_1\right\|_{\mathrm{F}}^2 \leq \frac{2\epsilon}{3}\left(1+\frac{1}{C\sqrt{\log(k/\delta)}}\right)^2\left\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\right\|_{\mathrm{F}}^2 \right\} \geq 1 - \delta^{C^2\log(k/\delta)/4} - 2\delta.$$

Since $C > 1$ and $k \geq 2$, a simple numerical bound allows us to state that, more simply,

$$\mathbb{P}\left\{ \left\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\boldsymbol{\Theta}^T\boldsymbol{\Theta}\mathbf{V}_1\right\|_{\mathrm{F}}^2 \leq 4\epsilon\left\|\boldsymbol{\Sigma}_2\mathbf{V}_2^T\right\|_{\mathrm{F}}^2 \right\} \geq 1 - \delta^{C^2\log(k/\delta)/4} - 2\delta.$$

This bound on $\left\|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\mathbf{\Theta}\mathbf{V}_1\right\|_{\mathrm{F}}^2$ is used to estimate the first term in (5.4.7). The remaining estimates in (5.4.7) follow from applying Lemma 5.5 (keeping in mind the lower bound (5.4.4) on ℓ) to obtain

$$\mathbb{P}\left\{\left\|(\mathbf{V}_1^T\mathbf{\Theta}^T)^\dagger - (\mathbf{V}_1^T\mathbf{\Theta}^T)^T\right\|_2^2 \le 2.38\epsilon\right\} \ge 1 - 3\delta.$$

and Lemma 5.10 with $\eta = 7/4$ to obtain

$$\mathbb{P}\left\{\left\|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\right\|_{\mathrm{F}}^2 \le \frac{11}{4}\left\|\mathbf{\Sigma}_2\mathbf{V}_2^T\right\|_{\mathrm{F}}^2\right\} \ge 1 - \left(\frac{\mathrm{e}^{7/4}}{(1+7/4)^{1+7/4}}\right)^{\ell/(1+\sqrt{8\log(n/\delta)})^2} - \delta.$$

We have the estimate

$$\frac{\mathrm{e}^{7/4}}{(1+7/4)^{1+7/4}} < \frac{1}{\mathrm{e}},$$

so in fact

$$\begin{aligned}\mathbb{P}\left\{\left\|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T\right\|_{\mathrm{F}}^2 \le \frac{11}{4}\left\|\mathbf{\Sigma}_2\mathbf{V}_2^T\right\|_{\mathrm{F}}^2\right\} &\ge 1 - \mathrm{e}^{-\ell/(1+\sqrt{8\log(n/\delta)})^2} - \delta \\ &\ge 1 - \mathrm{e}^{-6C^2\epsilon^{-1}\log(k/\delta)} - \delta \\ &\ge 1 - \mathrm{e}^{-\log(k/\delta)} - \delta \\ &\ge 1 - 2\delta.\end{aligned}$$

Adding up the failure probabilities of the three estimates used in (5.4.7), we conclude that the bound on S given in (5.4.8) holds with probability at least $1 - \delta^{C^2\log(k/\delta)/4} - 7\delta$. Thus Eqns. (i) and (iii) hold with this probability.

Next we establish the forward error bounds. Lemma 4.9 with $\mathbf{S} = \mathbf{\Theta}^T \in \mathbb{R}^{n\times\ell}$ gives

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{AS}}\mathbf{A}\right\|_{\mathrm{F}}^2 \le \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}^2 + \left\|\mathbf{\Sigma}_2\mathbf{V}_2^T\mathbf{\Theta}^T(\mathbf{V}_1^T\mathbf{\Theta}^T)^\dagger\right\|_{\mathrm{F}}^2.$$

Identify the second term on the right-hand side as S and introduce the estimate for S given in (5.4.8) to see that

$$\left\|\mathbf{A}_k - \mathbf{P}_{\mathbf{AS}}\mathbf{A}\right\|_{\mathrm{F}}^2 \leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}^2 + 22\epsilon \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}^2 = (1 + 22\epsilon) \cdot \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}^2.$$

This bound holds with probability at least $1 - \delta^{C^2 \log(k/\delta)/4} - 7\delta$. Taking the square-roots of both sides and using the fact that $\sqrt{1 + 22\epsilon} \leq 1 + 22\epsilon$ gives Eqn. (iii).

Finally, we prove Eqn. (iv):

$$\begin{aligned}\left\|\mathbf{A}_k - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{A}\mathbf{\Theta}^T,k}(\mathbf{A})\right\|_{\mathrm{F}} &= \left\|\mathbf{A} - \mathbf{A}_k - (\mathbf{A} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{A}\mathbf{\Theta}^T,k}(\mathbf{A}))\right\|_{\mathrm{F}} \\ &\leq \left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}} + \left\|\mathbf{A} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{A}\mathbf{\Theta}^T,k}(\mathbf{A})\right\|_{\mathrm{F}} \leq (2 + 11\epsilon)\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}},\end{aligned}$$

where the first inequality follows by the triangle inequality and the second follows from Eqn. (ii) in the theorem. This bound holds with probability at least $1 - \delta^{C^2 \log(k/\delta)/4} - 7\delta$.

□

5.5 Experiments

In this section, we experimentally investigate the tightness of the residual and forward error bounds provided in Theorems 5.13 and 5.14 for the spectral and Frobenius-norm approximation errors of SRHT low-rank approximations of the forms $\mathbf{P}_{\mathbf{A}\mathbf{\Theta}^T}\mathbf{A}$ and $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{A}\mathbf{\Theta}^T,k}(\mathbf{A})$. Additionally, we experimentally verify that the SRHT algorithms are not significantly less accurate than the Gaussian low-rank approximation algorithms.

5.5.1 The test matrices

Let $n = 1024$, and consider the following three test matrices:

1. Matrix $\mathbf{A} \in \mathbb{R}^{(n+1)\times n}$ is given by

$$\mathbf{A} = [100\mathbf{e}_1 + \mathbf{e}_2, 100\mathbf{e}_1 + \mathbf{e}_3, \ldots, 100\mathbf{e}_1 + \mathbf{e}_{n+1}],$$

   where $\mathbf{e}_i \in \mathbb{R}^{n+1}$ are the standard basis vectors.

2. Matrix $\mathbf{B} \in \mathbb{R}^{n\times n}$ is diagonal with entries $(\mathbf{B})_{ii} = 100(1 - (i-1)/n)$.

3. Matrix $\mathbf{C} \in \mathbb{R}^{n\times n}$ has the same singular values as $\mathbf{B}$, but its singular spaces are sampled from the uniform measure on the set of orthogonal matrices. More precisely, $\mathbf{C} = \mathbf{U}\mathbf{B}\mathbf{V}^T$, where $\mathbf{G} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ is the SVD of an $n \times n$ matrix whose entries are standard Gaussian random variables.

These three matrices exhibit properties that, judging from the bounds in Theorems 5.13 and 5.14, could challenge the SRHT approximation algorithm. Matrix $\mathbf{A}$ is approximately rank one—there is a large spectral gap after the first singular value—but the residual spectrum is flat, so for $k \geq 1$, the $\left\|\mathbf{A} - \mathbf{A}_k\right\|_{\mathrm{F}}$ terms in the spectral norm bounds of Theorem 5.13 are quite large compared to the $\left\|\mathbf{A} - \mathbf{A}_k\right\|_2$ terms. Matrices $\mathbf{B}$ and $\mathbf{C}$ both have slowly decaying spectrums, so one again has a large Frobenius term present in the spectral norm error bound.

Matrices $\mathbf{B}$ and $\mathbf{C}$ were chosen to have the same singular values but different singular spaces to reveal any effect that the structure of the singular spaces of the matrix has on the quality of SRHT approximations. The coherence of their right singular spaces provides a summary of the relevant difference in the singular spaces of $\mathbf{B}$ and $\mathbf{C}$. Recall that the coherence of a

k-dimensional subspace $\mathscr{S}$ is defined as

$$\mu(\mathscr{S}) = \frac{n}{k} \max_i \mathbf{P}_{ii},$$

where $\mathbf{P}$ is the projection onto $\mathscr{S}$; the coherence of $\mathscr{S}$ is always between 1 and n/k [CR09]. It is clear that all the right singular spaces of $\mathbf{B}$ are maximally coherent, and it is known that with high probability the dominant right k-dimensional singular space of $\mathbf{C}$ is quite incoherent, with coherence on the order of $k \log n$ [CR09].

To gain an intuition for the potential significance of this difference in coherence, consider a randomized column sampling approach to forming low-rank approximants; that is, consider approximating $\mathbf{M}_k$ with a matrix $\mathbf{P}_\mathbf{Y}\mathbf{M}$ where $\mathbf{Y}$ comprises randomly sampled columns of $\mathbf{M}$. It is known that such approximations are quite inaccurate unless the dominant k-dimensional right singular space of $\mathbf{M}$ is incoherent (see, e.g., Chapter 6 or [TR10]). One could interpret SRHT approximation algorithms as consisting of a rotation of the right singular spaces of $\mathbf{M}$ by multiplying from the right with $\mathbf{DH}^T$ followed by forming a column sample-based approximation. The rotation lowers the coherence of the right singular spaces and thereby increases the probability of obtaining an accurate low-rank approximation. One expects that if $\mathbf{M}$ has highly coherent right singular spaces then the right singular spaces of $\mathbf{MDH}^T$ will be less coherent. Thus we compare the performance of the SRHT approximations on $\mathbf{B}$, which has maximally coherent right singular spaces, to their performance on $\mathbf{C}$, which has almost maximally incoherent right singular spaces.

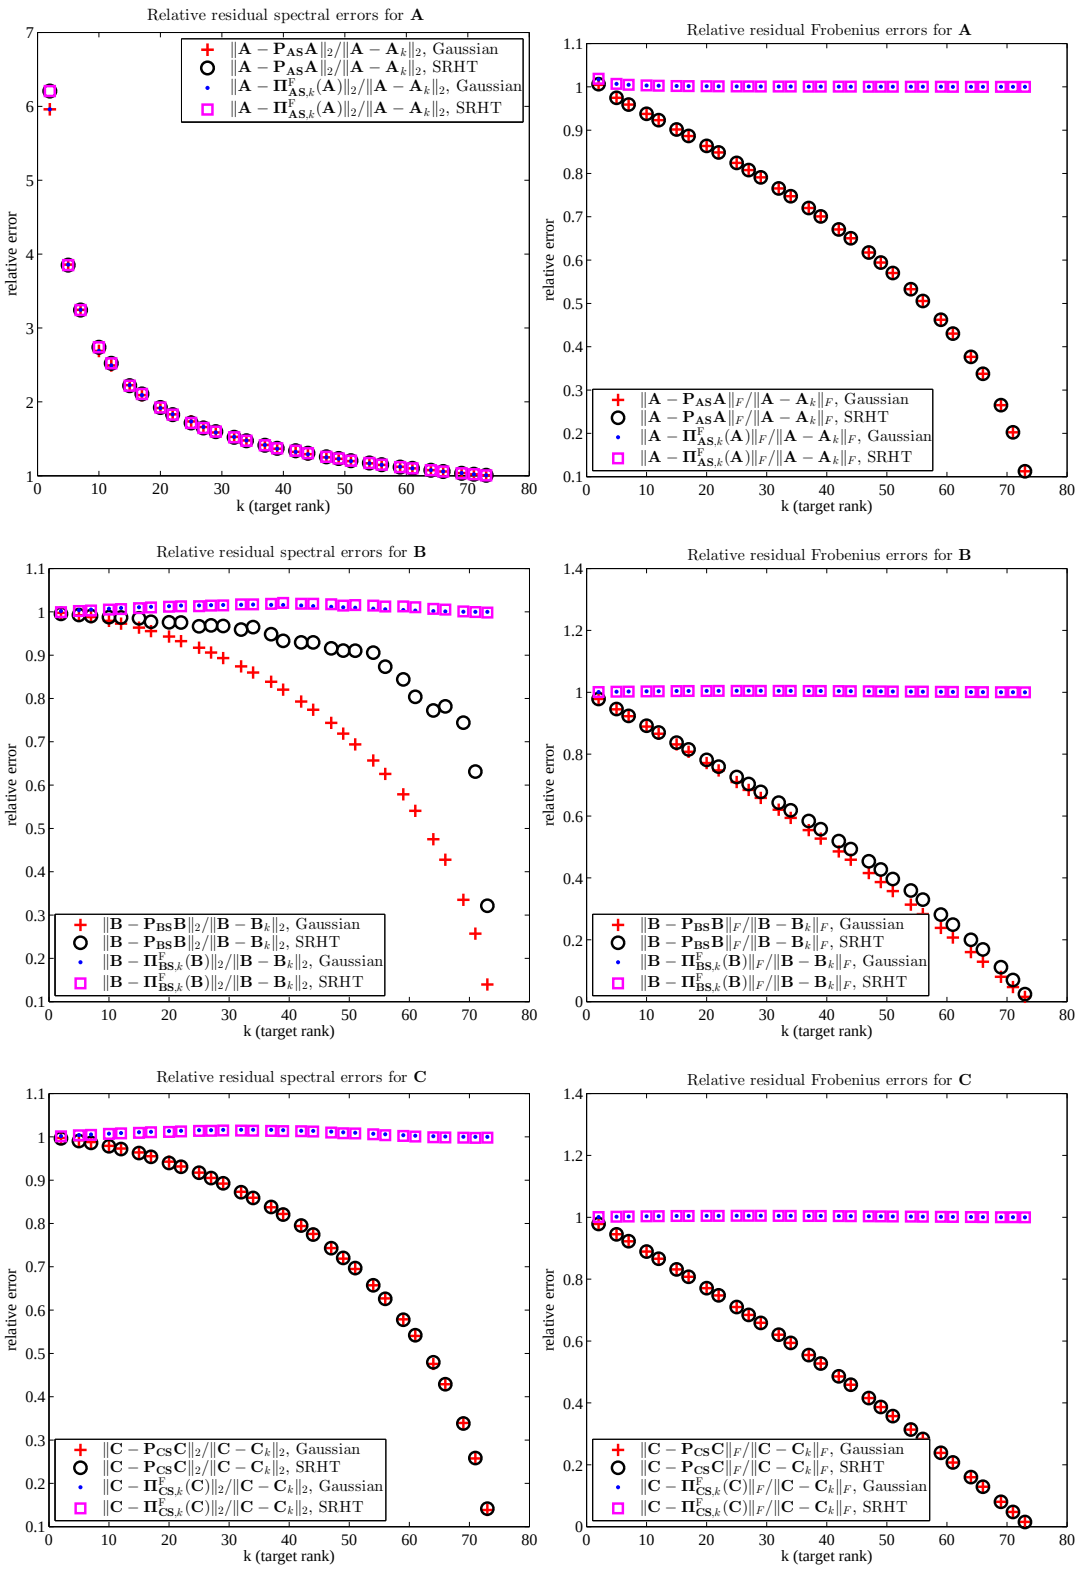

Figure 5.1: RESIDUAL ERRORS OF LOW-RANK APPROXIMATION ALGORITHMS. Relative spectral and Frobenius-norm residual errors of the SRHT and Gaussian low-rank approximation algorithms ($\|\mathbf{M} - \mathbf{P}_{\mathbf{MS}}\mathbf{M}\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ and $\|\mathbf{M} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{MS},k}(\mathbf{M})\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ for $\xi = 2, \mathrm{F}$) as a function of the target rank k for the three matrices $\mathbf{M} = \mathbf{A}, \mathbf{B}, \mathbf{C}$. Each point is the average error observed over 30 trials. In each trial, $\ell = \lceil 2k \log n \rceil$ column samples were used.

5.5.2 Empirical comparison of the SRHT and Gaussian algorithms

Figure 5.1 depicts the relative residual errors of the Gaussian and SRHT algorithms for approximations generated using Algorithms 5.1 and 5.2: $\mathbf{P}_{\mathbf{MS}}\mathbf{M}$ and $\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{MS},k}(\mathbf{M})$, which we shall hereafter refer to respectively as the non-rank-restricted and rank-restricted approximations. Here the matrix $\mathbf{M}$ is used to refer interchangeably to $\mathbf{A}, \mathbf{B}$, and $\mathbf{C}$. The relative residual errors ($\|\mathbf{M} - \mathbf{P}_{\mathbf{MS}}\mathbf{M}\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ and $\|\mathbf{M} - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{MS},k}(\mathbf{M})\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ for $\xi = 2, \mathrm{F}$) shown in this figure for each value of k were obtained by taking the average of the relative residual errors observed over 30 trials of low-rank approximations, each formed using $\ell = \lceil 2k \log n \rceil$ samples.

With the exception of the residual spectral errors on $\mathbf{A}$, which range from between two and nine times the size of the optimal rank-k spectral residual error for $k < 20$, we see that the residual errors for all three matrices are less than 1.1 times the residual error of $\mathbf{M}_k$, if not significantly smaller. Specifically, the relative residual errors of the restricted-rank approximations remain less than 1.1 over the entire range of k while the relative residual errors of the non-rank-restricted approximations actually decrease as k increases. Note that, because $\ell > k$, the relative errors of the non-rank-restricted approximations are often smaller than 1, while those of the restricted-rank approximations are never smaller than 1.

Since the matrices $\mathbf{B}$ and $\mathbf{C}$ have the same singular values, but the singular spaces of $\mathbf{C}$ are less coherent, the difference in the residual errors of the approximations of $\mathbf{B}$ and $\mathbf{C}$ is evidence that the spectral-norm accuracy of the SRHT approximations is increased on less coherent datasets; the same is true for the Frobenius norm accuracy to a lesser extent. The Gaussian approximations seem insensitive to the level of coherence. Only on the highly coherent matrix $\mathbf{B}$ do we see a notable decrease in the residual errors when Gaussian sampling is used rather than an SRHT; however, even in this case the residual errors of the SRHT approximations are

comparable with that of $\mathbf{B}_k$. In all, Figure 5.1 suggests that the gain in computational efficiency provided by the SRHT does not come at the cost of a significant loss in accuracy and that taking $\ell = \lceil 2k \log n \rceil$ samples suffices to obtain approximations with small residual errors relative to those of the optimal rank-k approximations. Up to the specific value of the constant, this latter observation coincides with the conclusions of Theorems 5.13 and 5.14.

Figure 5.2 depicts the relative forward errors of the Gaussian and SRHT algorithms ($\|\mathbf{M}_k - \mathbf{P}_{\mathbf{MS}}\mathbf{M}\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ and $\|\mathbf{M}_k - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{MS},k}(\mathbf{M})\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ for $\xi = 2, \mathrm{F}$) for the non-rank-restricted and rank-restricted approximations. The error shown for each k is the average relative forward error observed over 30 trials of low-rank approximations each formed using $\ell = \lceil 2k \log n \rceil$ samples. We observe that the forward errors of both algorithms for both choices of sampling matrices are on the scale of the norm of $\mathbf{M}_k$. By looking at the relative spectral-norm forward errors we see that in this norm, perhaps contrary to intuition, the rank-restricted approximation does not provide a more accurate approximation to $\mathbf{M}_k$ than does the non-rank-restricted approximation. However the rank-restricted approximation clearly provides a more accurate approximation to $\mathbf{M}_k$ than the non-rank-restricted approximation in the Frobenius norm. A rather unexpected observation is that the rank-restricted approximations are more accurate in the spectral norm for highly coherent matrices (**B**) than they are for matrices which are almost minimally coherent (**C**). Overall, Figure 5.2 suggests that the SRHT low-rank approximation algorithms provide accurate approximations to $\mathbf{M}_k$ when ℓ is in the regime suggested by Theorems 5.13 and 5.14.

5.5.3 Empirical evaluation of our error bounds

Figures 5.1 and 5.2 show that when $\ell = \lceil 2k \log n \rceil$ samples are taken, the SRHT low-rank approximation algorithms both provide approximations to $\mathbf{M}$ that are within a factor of $1 + \epsilon$ as accurate in the Frobenius norm as $\mathbf{M}_k$, as Theorem 5.14 suggests should be the case. More

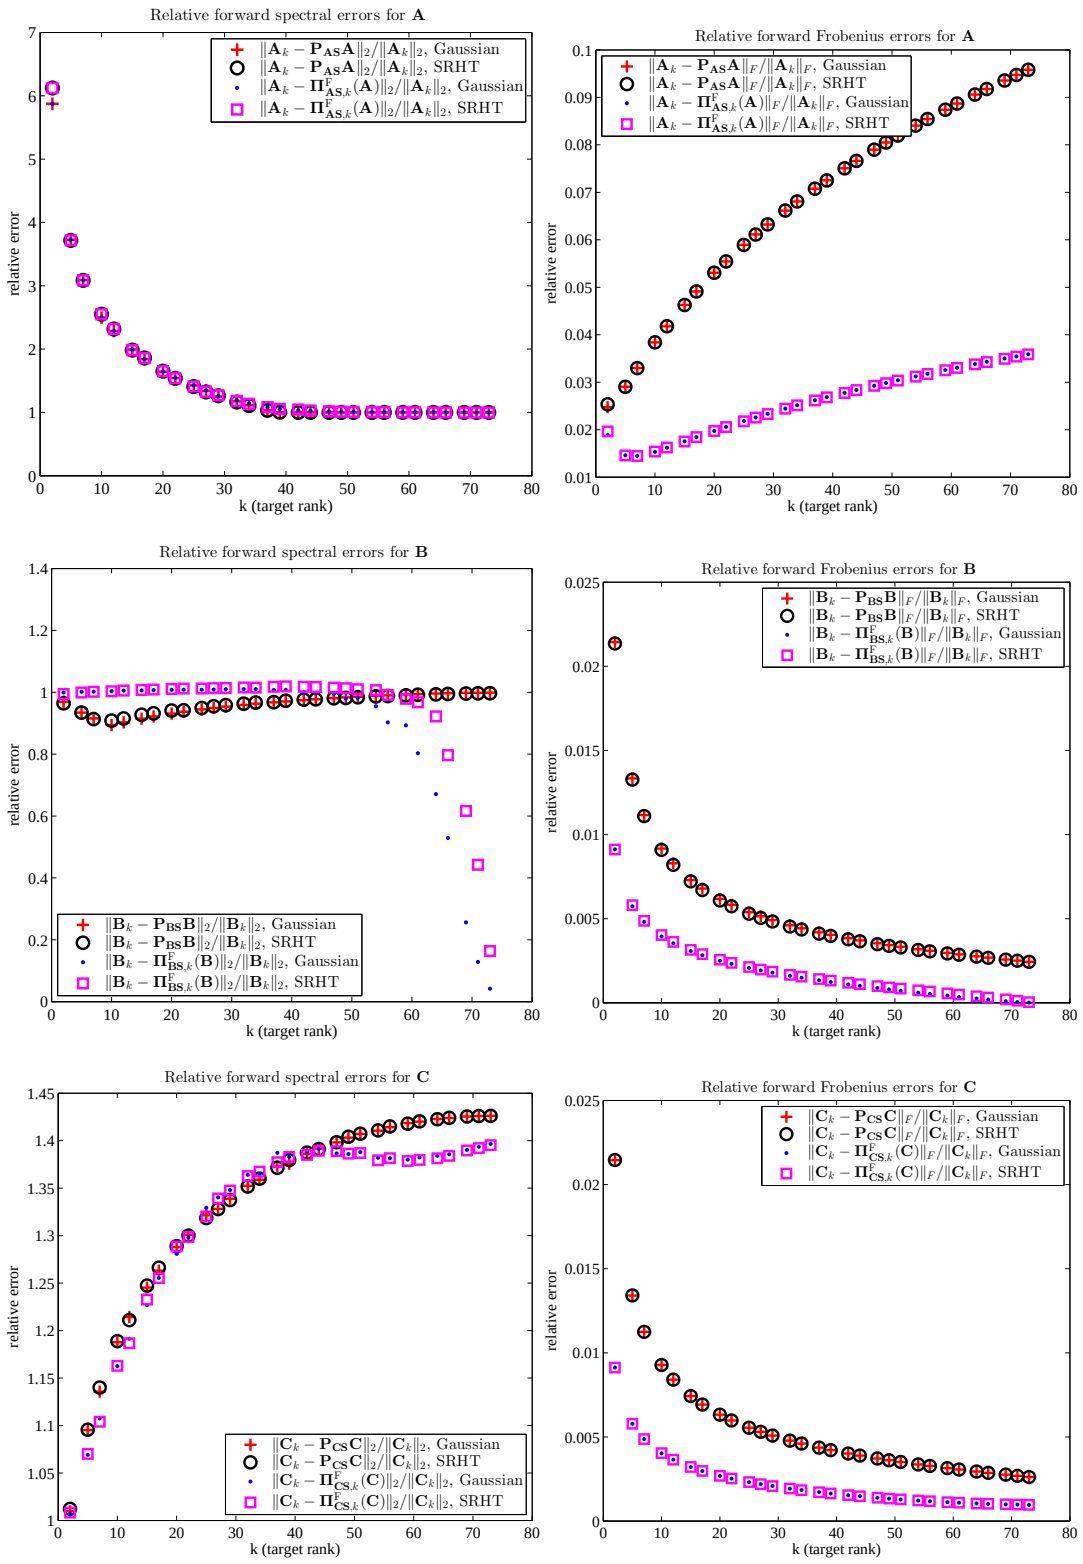

Figure 5.2: FORWARD ERRORS OF LOW-RANK APPROXIMATION ALGORITHMS. The relative spectral and Frobenius-norm forward errors of the SRHT and Gaussian low-rank approximation algorithms ($\|\mathbf{M}_k - \mathbf{P}_{\mathbf{MS}}\mathbf{M}\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ and $\|\mathbf{M}_k - \mathbf{\Pi}^{\mathrm{F}}_{\mathbf{MS},k}(\mathbf{M})\|_\xi / \|\mathbf{M} - \mathbf{M}_k\|_\xi$ for $\xi = 2, \mathrm{F}$) as a function of the target rank k for the three matrices $\mathbf{M} = \mathbf{A}, \mathbf{B}, \mathbf{C}$. Each point is the average of the errors observed over 30 trials. In each trial, $\ell = \lceil 2k \log n \rceil$ column samples were used.

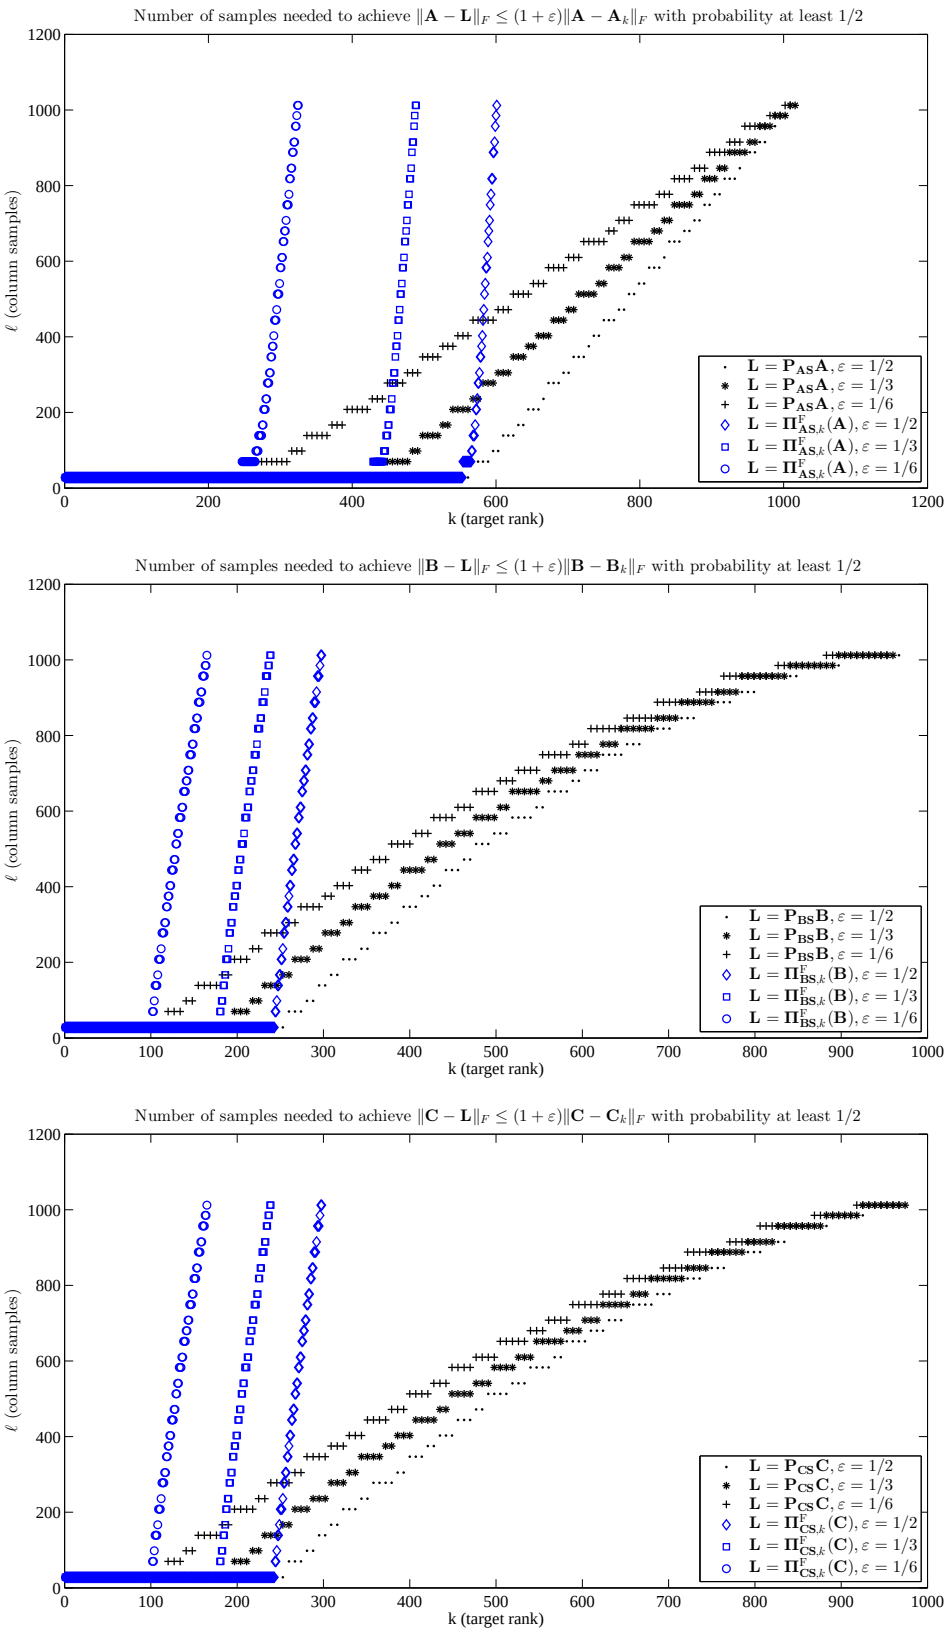

Figure 5.3: THE NUMBER OF COLUMN SAMPLES REQUIRED FOR RELATIVE ERROR FROBENIUS-NORM APPROXIMATIONS. The value of ℓ empirically necessary to ensure that, with probability at least one-half, approximations generated by the SRHT algorithms satisfy $\left\|\mathbf{M}-\mathbf{P}_{\mathbf{M}\mathbf{\Theta}^T}\mathbf{M}\right\|_{\mathrm{F}} \leq (1+\epsilon)\left\|\mathbf{M}-\mathbf{M}_k\right\|_{\mathrm{F}}$ and $\left\|\mathbf{M}-\mathbf{\Pi}^{\mathrm{F}}_{\mathbf{M}\mathbf{\Theta}^T,k}(\mathbf{M})\right\|_{\mathrm{F}} \leq (1+\epsilon)\left\|\mathbf{M}-\mathbf{M}_k\right\|_{\mathrm{F}}$ (for $\mathbf{M}=\mathbf{A},\mathbf{B},\mathbf{C}$).

precisely, Theorem 5.14 assures us that $528\epsilon^{-1}[\sqrt{k}+\sqrt{8\log(8n/\delta)}]^2\log(8k/\delta)$ column samples are sufficient to ensure that, with at least probability $1-\delta$, $\mathbf{P}_{\mathbf{M}\boldsymbol{\Theta}^T}\mathbf{M}$ and $\boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{M}\boldsymbol{\Theta}^T,k}(\mathbf{M})$ have Frobenius norm residual and forward error within $1+\epsilon$ of that of $\mathbf{M}_k$. The factor 528 can certainly be reduced by optimizing the numerical constants given in Theorem 5.14. But what is the smallest ℓ that ensures the Frobenius norm residual error bounds $\left\|\mathbf{M}-\mathbf{P}_{\mathbf{M}\boldsymbol{\Theta}^T}\mathbf{M}\right\|_{\mathrm{F}} \leq (1+\epsilon)\left\|\mathbf{M}-\mathbf{M}_k\right\|_{\mathrm{F}}$ and $\left\|\mathbf{M}-\boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{M}\boldsymbol{\Theta}^T,k}(\mathbf{M})\right\|_{\mathrm{F}} \leq (1+\epsilon)\left\|\mathbf{M}-\mathbf{M}_k\right\|_{\mathrm{F}}$ are satisfied with some fixed probability? To investigate, in Figure 5.3 we plot the values of ℓ determined empirically to be sufficient to obtain $(1+\epsilon)$ Frobenius norm residual errors relative to the optimal rank-k approximation; we fix the failure probability $\delta = 1/2$ and vary ϵ. Specifically, the ℓ plotted for each k is the smallest number of samples for which $\left\|\mathbf{M}-\mathbf{P}_{\mathbf{M}\boldsymbol{\Theta}^T}\mathbf{M}\right\|_{\mathrm{F}} \leq (1+\epsilon)\left\|\mathbf{M}-\mathbf{M}_k\right\|_{\mathrm{F}}$ or $\left\|\mathbf{M}-\boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{M}\boldsymbol{\Theta}^T,k}(\mathbf{M})\right\|_{\mathrm{F}} \leq (1+\epsilon)\left\|\mathbf{M}-\mathbf{M}_k\right\|_{\mathrm{F}}$ in at least 15 out of 30 trials.

It is clear that, for fixed k and ϵ, the number of samples ℓ required to form a non-rank-restricted approximation to $\mathbf{M}$ with $1+\epsilon$ relative residual error is smaller than the ℓ required to form a rank-restricted approximation with $1+\epsilon$ relative residual error. Note that for small values of k, the ℓ necessary for relative residual error to be achieved is actually smaller than k for all three datasets. This is a reflection of the fact that when $k_1 < k_2$ are small, the ratio $\left\|\mathbf{M}-\mathbf{M}_{k_2}\right\|_{\mathrm{F}} / \left\|\mathbf{M}-\mathbf{M}_{k_1}\right\|_{\mathrm{F}}$ is very close to one. Outside of the initial flat regions, the empirically determined value of r seems to grow linearly with k; this matches with the observation of Woolfe et al. that taking $\ell = k+8$ suffices to consistently form accurate low-rank approximations using the SRFT scheme, which is very similar to the SRHT scheme [WLRT08]. We also note that this matches with Theorem 5.14, which predicts that the necessary ℓ grows at most linearly with k with a slope like $\log n$.

Finally, Theorem 5.13 does *not* guarantee that $1+\epsilon$ spectral-norm relative residual errors can be achieved. Instead, it provides bounds on the spectral-norm residual errors achieved in terms of

$\left\|\mathbf{M} - \mathbf{M}_k\right\|_2$ and $\left\|\mathbf{M} - \mathbf{M}_k\right\|_{\mathrm{F}}$ that are guaranteed to hold when ℓ is sufficiently large. In Figure 5.4 we compare the spectral-norm residual error guarantees of Theorem 5.13 to what is achieved in practice. To do so, we take the optimistic viewpoint that the constants in Theorem 5.13 can be optimized to unity. Under this view, if more columns than $\ell_2 = \epsilon^{-1}[\sqrt{k} + \sqrt{\log(n/\delta)}]^2 \log(k/\delta)$ are used to construct the SRHT approximations, then the spectral-norm residual error is no larger than

$$b_2 = \left(1 + \sqrt{\frac{\log(n/\delta)\log(\rho/\delta)}{\ell}}\right) \cdot \left\|\mathbf{M} - \mathbf{M}_k\right\|_2 + \sqrt{\frac{\log(\rho/\delta)}{\ell}} \cdot \left\|\mathbf{M} - \mathbf{M}_k\right\|_{\mathrm{F}},$$

where ρ is the rank of $\mathbf{M}$, with probability greater than $1 - \delta$. Our comparison consists of using ℓ_2 samples to construct the SRHT approximations and then comparing the predicted upper bound on the spectral-norm residual error, b_2, to the empirically observed spectral-norm residual errors. Figure 5.4 shows, for several values of k, the upper bound b_2 and the observed relative spectral-norm residual errors, with precision parameter $\epsilon = 1/2$ and failure parameter $\delta = 1/2$. For each value of k, the empirical spectral-norm residual error plotted is the average of the errors over 30 trials of low-rank approximations. Note from Figure 5.4 that with this choice of ℓ, the spectral-norm residual errors of the rank-restricted and non-rank-restricted SRHT approximations are essentially the same.

Judging from Figures 5.3 and 5.4, even when we assume the constants present can be optimized away, the bounds given in Theorems 5.13 and 5.14 are pessimistic: it seems that in fact approximations with Frobenius-norm residual error within $1 + \epsilon$ of the error of the optimal rank-k approximation can be obtained with ℓ linear in k, and the spectral-norm residual errors are smaller than the supplied upper bounds. Thus there is still room for improvement in our understanding of the SRHT low-rank approximation algorithm, but as explained in Section 5.2.1,

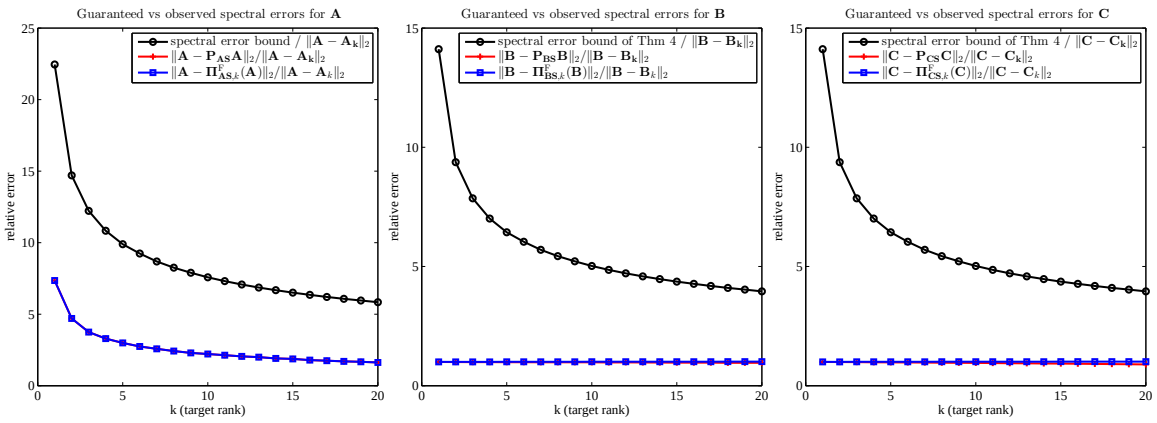

Figure 5.4: EMPIRICAL VERSUS PREDICTED SPECTRAL-NORM RESIDUAL ERRORS OF LOW-RANK APPROXIMATIONS. The empirical spectral-norm residual errors relative to those of the optimal rank-k approximants ($\left\|\mathbf{M} - \mathbf{P}_{\mathbf{M}\boldsymbol{\Theta}^T}\mathbf{M}\right\|_2 / \|\mathbf{M} - \mathbf{M}_k\|_2$ and $\|\mathbf{M} - \boldsymbol{\Pi}^{\mathrm{F}}_{\mathbf{M}\boldsymbol{\Theta}^T,k}(\mathbf{M})\|_2 / \|\mathbf{M} - \mathbf{M}_k\|_2$) plotted alongside the same ratio for the bounds given in Theorem 5.13, when $\ell = \lceil 2[\sqrt{k} + \sqrt{\log(2n)}]^2 \log(2k) \rceil$ (for $\mathbf{M} = \mathbf{A}, \mathbf{B}, \mathbf{C}$). On the scale shown, the errors of the two SRHT-based approximation algorithms are essentially identical.

ignoring constants, the bounds of Theorem 5.13 are often tighter than those obtained in earlier works.

To bring perspective to this discussion, consider that even if one limits consideration to deterministic algorithms, the known error bounds for the Gu–Eisenstat rank-revealing QR—a popular and widely used algorithm for low-rank approximation—are quite pessimistic and do not reflect the excellent accuracy that is seen in practice [GE96]. Regardless, we do not advocate using these approximation schemes for applications in which highly accurate low-rank approximations are needed. Rather, Theorems 5.13 and 5.14 and our numerical experiments suggest that they are appropriate in situations where one is willing to trade some accuracy for a gain in computational efficiency.