

# Tight Sampling Bounds for Eigenvalue Approximation \*

William Swartworth  
Carnegie Mellon University

David P. Woodruff  
Carnegie Mellon University

## Abstract

We consider the problem of estimating the spectrum of a symmetric bounded entry (not necessarily PSD) matrix via entrywise sampling. This problem was introduced by [Bhattacharjee, Dexter, Drineas, Musco, Ray '22], where it was shown that one can obtain an  $\epsilon n$  additive approximation to all eigenvalues of  $A$  by sampling a principal submatrix of dimension  $\frac{\text{poly}(\log n)}{\epsilon^3}$ . We improve their analysis by showing that it suffices to sample a principal submatrix of dimension  $\tilde{O}(\frac{1}{\epsilon^2})$  (with no dependence on  $n$ ). This matches known lower bounds and therefore resolves the sample complexity of this problem up to  $\log \frac{1}{\epsilon}$  factors. Using similar techniques, we give a tight  $\tilde{O}(\frac{1}{\epsilon^2})$  bound for obtaining an additive  $\epsilon \|A\|_F$  approximation to the spectrum of  $A$  via squared row-norm sampling, improving on the previous best  $\tilde{O}(\frac{1}{\epsilon^8})$  bound. We also address the problem of approximating the top eigenvector for a bounded entry, PSD matrix  $A$ . In particular, we show that sampling  $O(\frac{1}{\epsilon})$  columns of  $A$  suffices to produce a unit vector  $u$  with  $u^T A u \geq \lambda_1(A) - \epsilon n$ . This matches what one could achieve via the sampling bound of [Musco, Musco'17] for the special case of approximating the top eigenvector, but does not require adaptivity.

As additional applications, we observe that our sampling results can be used to design a faster eigenvalue estimation sketch for dense matrices resolving a question of [Swartworth, Woodruff'23], and can also be combined with [Musco, Musco'17] to achieve  $O(1/\epsilon^3)$  (adaptive) sample complexity for approximating the spectrum of a bounded entry PSD matrix to  $\epsilon n$  additive error.

## 1 Introduction

Computing the spectrum of a matrix is a fundamental problem with many applications. There are well-known high-precision algorithms that run in polynomial time [Fra62; GV00], although any such algorithm is necessarily at least linear time in the input size. As data grows larger, even linear algorithms can be prohibitive. This has motivated a flurry of activity studying sublinear time estimation of problems in numerical linear algebra, for instance for low-rank approximation [MM17; MW17; BW18; BCW20], kernel density estimation [CS17; Sim+19; Cha+20], testing positive-semidefiniteness [BCJ20], and matrix sparsification [Bha+23; DZ11].

For eigenvalue estimation, variants of the power method have long been known to give good approximations to the top eigenvalues and eigenvectors of  $A \in \mathbb{R}^{n \times n}$  while revealing sublinear information about  $A$ , i.e., using  $o(n)$  matrix-vector queries [RST10; MM15]. However it was only recently asked in [Bha+24] whether there are spectral approximation algorithms for symmetric, but non-PSD matrices that run in sublinear time in the entry query model. This is perhaps the most natural model if one imagines having an extremely large matrix saved on disk for example. This may be in the form of a graph for instance, where one could be interested in obtaining spectral information about its Laplacian or adjacency matrix. One could also imagine having a large collection of data points with some kernel function that can be computed for pairs of points. Obtaining a rough spectral summary of the associated kernel matrix is a natural step for data analysis, for instance, to spot low-rank structure in the data. If data points are large or expensive to collect, or if kernel evaluation is expensive, it is natural to aim for minimizing entry queries to the kernel matrix.

Of course, it is not reasonable to ask for sublinear time spectral approximation algorithms, without some additional assumptions. For instance, our matrix  $A$  could contain all zeros but with a single large entry at indices  $(i, j)$  and  $(j, i)$ . Given only the ability to query entries, and no additional information, even distinguishing  $A$  from the all zeros matrix would take  $\Omega(n^2)$  queries.

We consider two assumptions that allow for improved guarantees. The first is an assumption on the structure of  $A$  called the *bounded entry model*, which assumes that  $A$  has entries bounded by 1 in magnitude. This condition

\*The authors thank support from a Simons Investigator Award, NSF CCF-2335412, and a Google Faculty Award.

was introduced in [Bal+19] and studied further by [BCJ20] who showed that it was sufficient in order to test for positive semi-definiteness with sublinear entry queries. Motivated by this result, [Bha+24] showed that all eigenvalues of a symmetric bounded-entry matrix can be approximated using a sublinear number of queries, simply by sampling a  $\text{poly}(\frac{\log n}{\epsilon})$  sized submatrix.

Another way of getting sublinear sample complexity is to give the sampler additional power. In our case, as in [Bha+24], we consider having access to a sampler that can produce a row index with probability proportional to its squared row norm. Such samplers have been increasingly studied under the guise of “quantum-inspired” machine-learning algorithms [Tan19; Che+22; GLT18; GST22]. Such samplers are practical to maintain when  $A$  is stored entrywise. For example by using an appropriate data structure, they can be built in  $\text{nnz}(A)$  time, admit  $O(\log n)$  time sampling, and can handle entry updates in  $O(\log n)$  time.

[Bha+24] showed that given a real symmetric matrix  $A$  with entries bounded by 1, one can sample a principal submatrix of  $A$  of dimensions  $O(\frac{\text{poly log } n}{\epsilon^3}) \times O(\frac{\text{poly log } n}{\epsilon^3})$  and then output an additive  $\epsilon n$  approximation to the entire spectrum, i.e., to all eigenvalues of  $A$ . On the other hand, the best lower bound states that a principal-submatrix algorithm must sample at least  $O(\frac{1}{\epsilon^4})$  entries of  $A$ .

There are two ways that one could hope to improve the sampling bound of [Bha+24]. First one could hope to improve the  $\epsilon$  dependence in the dimension from  $O(1/\epsilon^3)$  to  $O(1/\epsilon^2)$ . Several prior results suggested that this might be possible. For example [Bha+24] showed that their  $\epsilon$  dependence could be improved both when  $A$  is PSD, and when the spectrum of  $A$  is flat. Concurrently [SW23] showed that one can obtain an  $\epsilon \|A\|_F$  additive approximation to the spectrum of  $A$  by using a so-called bilinear sketch of  $A$  of dimensions  $O(1/\epsilon^2) \times O(1/\epsilon^2)$ . Such a sketch would give  $\epsilon n$  additive error for approximating all eigenvalues when  $A$  has bounded entries. Unfortunately this sketch is Gaussian, and it seems difficult to directly adapt its analysis to obtain a sampling bound instead.

In this paper we close the gap between sketching and sampling for bounded entry matrices by showing that uniformly sampling an  $\tilde{O}(1/\epsilon^2) \times \tilde{O}(1/\epsilon^2)$  principal submatrix of  $A$  suffices to approximate all eigenvalues of  $A$  up to  $\epsilon n$  additive error, even when  $A$  is not necessarily PSD. In addition to obtaining an optimal  $\epsilon$ -dependence, we note that our uniform sampling bound contains no dependence on  $n$ .

We also address the squared-row norm sampling model. Here we improve the analysis of [Bha+24] to show that it suffices to query a principal submatrix of size  $\tilde{O}(\frac{\text{poly log } n}{\epsilon^2})$ , compared to the  $\tilde{O}(\frac{\text{poly log } n}{\epsilon^8})$  dimensional principal submatrix required by [Bha+24].

**Our approximation model.** We note that all of the guarantees considered in our work and prior work focus on additive approximations to the spectrum. Ideally, one might like to aim for a relative error guarantee. However as pointed out, by [Bha+24] for example, this is not possible for entry queries. Such an algorithm would be able to distinguish the 0 matrix from a matrix with a single off-diagonal pair of nonzero entries, which clearly requires  $\Omega(n^2)$  samples. Indeed even with squared row-norm sampling, relative error is still too much to hope for. In fact, even for sketches, approximating the top eigenvalue to within a constant factor requires  $\Omega(n^2)$  sketching dimension in general [LW16; Woo+14]. One can always turn such a sketching lower bound into a sampling bound, even allowing for row-norm sampling - simply conjugate by a random orthogonal matrix to flatten all rows. Then squared row sampling is effectively uniform, and so a sampling algorithm could be used to construct a sketch of the same dimensions. These existing lower bounds are why we (as well as prior work) choose to focus on an additive approximation guarantee.

## 2 Our Results

As discussed above, we are interested in the same type of guarantee as considered in [Bha+24] and [SW23].

**DEFINITION 2.1.** *Let  $A \in \mathbb{R}^{n \times n}$  be a symmetric matrix. We say that a sequence  $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_n$  is an additive  $\alpha$ -approximation to the spectrum of  $A$  if  $|\hat{\lambda}_i(A) - \lambda_i(A)| \leq \alpha$  for all  $i \in [n]$ .*

We consider the sampling algorithms introduced by [Bha+24], with some technical modifications. Our novelty is an improved analysis which yields optimal bounds up to logarithmic factors. Our techniques apply both to uniform sampling for bounded entry matrices, as well as to squared row norm sampling for arbitrary matrices.

**Uniform sampling.** For uniform sampling, we show the following result.

**THEOREM 2.1.** *Let  $A$  be a symmetric matrix with all entries bounded by 1. Then Algorithm 1 with  $s \geq c \frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon}$  (where  $c$  is an absolute constant) outputs an additive  $\epsilon n$  approximation to the spectrum of  $A$ , with at least  $2/3$  probability.*

Eigenvalue Estimation		
Model/Error guarantee	Entry samples	Authors
Bounded entries/ Additive $\epsilon n$	$O\left(\frac{\text{poly log } n}{\epsilon^6}\right)$	[Bha+24]
	$O\left(\frac{\text{poly log } n}{\epsilon^5}\right)^*$	[Bha+24]
	$\tilde{O}\left(\frac{1}{\epsilon^4}\right)$	Ours
Row-norm sampling/ Additive $\epsilon \ A\ _F$	$O\left(\frac{\text{poly log } n}{\epsilon^{16}}\right)$	[Bha+24]
	$\tilde{O}\left(\frac{\text{poly log } n}{\epsilon^4}\right)$	Ours

Table 1: Comparison of our eigenvalue sampling results to prior work. The \* indicates that the algorithm does not query a principal submatrix. All algorithms here are non-adaptive.

---

**Algorithm 1** Uniform Sampling

---

1. Input: Symmetric matrix  $A \in \mathbb{R}^{n \times n}$  with  $\|A\|_\infty \leq 1$ , expected sample size  $s$
2. If  $s \geq n$ , let  $S = I_n$
3. Otherwise, let  $S \in \mathbb{R}^{k \times n}$  be a (rescaled) sampling matrix which samples each row of  $A$  independently with probability  $p = s/n$  and rescales each sampled row by  $\sqrt{\frac{1}{p}}$ .
4. Return the  $k$  eigenvalues of  $S^T AS$ , along with  $n - k$  additional 0's.

---

The sampling algorithm used here is perhaps as natural as possible. We simply sample a random principal submatrix of expected size  $s \times s$  and appropriately rescale. We then output the eigenvalues of the sampled submatrix, along with roughly  $n - s$  additional 0's to account for the remaining spectrum of  $A$ . Our sampling bound is an improvement over [Bha+24] which requires sampling a principal submatrix of dimension  $\frac{c}{\epsilon^3} \log \frac{1}{\epsilon} \log^3 n$  to obtain the same guarantee. We note that [Bha+24] was able to push their entry sample complexity down to  $O(1/\epsilon^5)$  by using a matrix sparsification result to subsample their principal submatrix. However the resulting sample is not a principal submatrix, and it is unclear how to decrease the sample complexity further.

[Bha+24] also asked whether one can obtain an additive  $\epsilon n$  guarantee to the spectrum of  $A$  with sample size independent of  $A$ . Our result shows that this is indeed possible. This may seem surprising since [Bha+24] uses a bound due to Tropp [RV07; Tro08], to control the contribution from the small-magnitude eigenvalues. This bound contains a  $\log n$  term that appears difficult to remove, and we use the same bound on the small-magnitude eigenvalues. Nonetheless, we give a very simple bootstrapping argument to show that this  $\log n$  dependence is non-essential. Specifically, we show that it is possible to replace the  $\log n$  with a  $\log \frac{1}{\epsilon}$ , thus removing the  $n$  dependence, at the cost of only additional  $\log \frac{1}{\epsilon}$  factors.

**Squared row-norm sampling.** For squared row-norm sampling, we use our analysis to again give nearly tight bounds.

**THEOREM 2.2.** *Let  $A$  be an arbitrary symmetric matrix. Algorithm 2 with  $s \geq \frac{1}{\epsilon^2} \text{poly log}(n/\epsilon)$  outputs an additive  $\epsilon \|A\|_F$  additive approximation to the spectrum of  $A$  with at least  $2/3$  probability.*

This is a substantial improvement over the sampling bound in [Bha+24], which requires sampling a principal submatrix of dimension  $\frac{c \log^{10} n}{\epsilon^8}$  to match our bound.

Again the most challenging part of our analysis is in controlling the outlying eigenvalues. Here [Bha+24] again requires roughly  $1/\epsilon^3$  samples, whereas we require  $1/\epsilon^2$  samples by essentially the same argument as for the case of uniform sampling. However on its own, this is enough to give improved bounds as the bottleneck in their argument is in controlling the middle eigenvalues. The authors do not attempt to give tight  $\epsilon$  bounds for this part of their analysis, and note that their bound is likely not tight. However we observe that their analysis here is surprisingly close to optimal. By optimizing their variance argument, we observe that their approach is

---

**Algorithm 2** Row-norm Sampling

---

1. Input: Symmetric matrix  $A \in \mathbb{R}^{n \times n}$  along with its row norms, expected sample size  $s$
2. Let  $S \in \mathbb{R}^{k \times n}$  be a (rescaled) sampling matrix which for all  $i \in [n]$  samples  $\text{Binomial}(s, \frac{\|A_i\|_2^2}{\|A\|_F^2})$  copies of row  $i$  and rescales those rows by  $\frac{1}{\sqrt{p_i}}$  where  $p_i := \frac{s\|A_i\|_2^2}{\|A\|_F^2}$ .
3. Implicitly form the matrix  $A' \in \mathbb{R}^{n \times n}$  defined by

$$(A')_{ij} = \begin{cases} 0 & i = j \\ 0 & i \neq j \text{ and } \|A_i\|_2^2 \|A_j\|_2^2 \leq \frac{\epsilon^2 \|A\|_F^2 |A_{ij}|^2}{c \log^4 n} \\ A_{ij} & \text{otherwise} \end{cases}$$

4. Return the  $k$  eigenvalues of  $S^T A' S$ , along with  $n - k$  additional 0's.

---

sufficient to control the middle eigenvalues with optimal  $\epsilon$  dependence.

We note that one cannot hope to do better than  $O(1/\epsilon^4)$  queries for non-adaptive entry sampling algorithms. Indeed a lower bound of [LW16] implies that learning the top singular value of a (non-symmetric) matrix  $G \in \mathbb{R}^{n \times n}$  to  $\epsilon \|G\|_F$  additive error requires  $\Omega(1/\epsilon^4)$  sketching dimension. One can symmetrize to  $G' = [0, G; G^T, 0]$  to show that learning the top eigenvalue of  $G'$  to  $\epsilon \|G'\|_F$  additive error requires  $\Omega(1/\epsilon^4)$  sketching dimension in general. The hard instance  $G$  is Gaussian, so its entries are bounded by  $\epsilon^2 \|G\|_F^2$  up to log factors implying an  $\tilde{\Omega}(1/\epsilon^4)$  lower bound in the bounded entry model, even for sketches. The rows of  $G$  have equal norms up to constant factors say, and so this implies that  $\tilde{\Omega}(1/\epsilon^4)$  queries are needed even given access to a row-norm sampling oracle that is accurate up to constant factors (which is all that our sampling algorithm requires).

Finally, note that while our results are not stated as a high-probability guarantee, it is easy to improve them. Simply run the algorithm  $O(\log \frac{1}{\delta})$  times and take the median estimate for each eigenvalue  $\lambda_i$ .

**Top eigenvector estimation.** While our main results focus on eigenvalue estimation, in many situations, one is interested not only in the top eigenvalues but also the associated top eigenvectors. As a step in this direction, we show that for bounded entry PSD matrices, one can obtain an  $\epsilon n$ -approximate top eigenvector by sampling just  $O(1/\epsilon)$  columns of  $A$ .

The same guarantee could be achieved by [MM17], by using  $\tilde{O}(n/\epsilon)$  *adaptive* entry samples to obtain an additive spectral approximation to  $A$ . Interestingly, in the bounded entry case, we show that no adaptivity is needed.

---

**Algorithm 3** Top Eigenvector Estimation

---

1. **Input:** Symmetric PSD matrix  $A \in \mathbb{R}^{n \times n}$  with  $\|A\|_\infty \leq 1$ , column sampling probability  $p$
2. **Output:** An approximate top eigenvector  $u \in \mathbb{R}^n$
3. Sample  $S \in \mathbb{R}^{n \times m}$ , a column sampling matrix that independently samples each column of  $A$  with probability  $p$
4. Given  $S$ , compute  $AS$  by sampling the requisite columns of  $A$
5. Use  $AS$  and  $S$  to compute the matrices  $S^T AS$  and  $S^T A^2 S$ .
6. Find  $x \in \mathbb{R}^n$  that maximizes  $\frac{x^T S^T A^2 S x}{x^T S^T A S x}$
7. Return  $\frac{ASx}{\|ASx\|_2}$

---

**THEOREM 2.3.** *Let  $A \in \mathbb{R}^{n \times n}$  be a symmetric PSD with  $\|A\|_\infty \leq 1$ . For  $p = \frac{c}{\epsilon n}$ , with  $3/4$  probability, Algorithm 3 returns a unit vector  $u \in \mathbb{R}^n$  satisfying*

$$u^T A u \geq \lambda_1 - \epsilon n,$$

where  $\lambda_1$  is the top eigenvalue of  $A$ .

**Faster sketching.** [SW23] studied the problem of sketching the eigenvalues of  $A$  to  $\epsilon \|A\|_F$  additive error, and showed that this can be accomplished with a Gaussian sketch of dimension  $O(1/\epsilon^4)$ . Their sketch takes roughly  $\frac{1}{\epsilon^2} n^2$  time to apply, so if  $\epsilon$  were say  $n^{-1/4}$ , this would be substantially slower than linear in the size of  $A$ . It was left open whether faster sketching is possible here when  $A$  is dense. We observe that our row-norm sampling analysis implies the existence of a sketch that runs in  $\tilde{O}(n^2)$  time with sketching dimension  $\tilde{O}(1/\epsilon^4)$ .

## 2.1 Our Techniques

**Bounded Entry Matrices.** We use the same algorithm that was introduced in [Bha+24], however we show how to reduce the sample size without sacrificing in terms of the approximation guarantee. At a high level we follow a similar approach of splitting the matrix  $A$  as  $A_o + A_m$  into its “outer” and “middle” eigenvalues. The matrix  $A_o$  will zero out all eigenvalues of  $A$  with magnitude smaller than  $\epsilon n$ . We will think of  $A_o$  as containing all the eigenvalues that we need to approximate, and  $A_m$  as adding some additional noise whose contribution we must bound. For a bounded entry matrix  $A$ , note that  $\|A\|_F^2 \leq n^2$ , so that there can be at most  $1/\epsilon^2$  nonzero eigenvalues in  $A_o$ . Thus it is at least reasonable to hope that an  $O(1/\epsilon^2)$  sized principal submatrix is sufficient.

If  $S$  is the (appropriately rescaled) sampling matrix, then we can write  $SAS^T = SA_o S^T + SA_m S^T$ . Note that  $SAS^T$  is simply our principal submatrix sample of  $A$ .

Given this decomposition, the analysis naturally decomposes into two parts: showing that the eigenvalues of  $SA_o S^T$  concentrate appropriately around the large eigenvalues, and showing that the perturbation from  $SA_m S^T$  is small. Prior work [Bha+24] handles the middle eigenvalues nearly optimally (although we give a small improvement here that we describe later). So our main technical innovation is to show that the eigenvalues of  $SA_o S^T$  concentrate to within  $\epsilon n$  of the eigenvalues of  $A_o$  when  $S$  samples only  $\tilde{O}(1/\epsilon^2)$  rows.

To see the main ideas, write the spectral decomposition of  $A_o$  as  $A_o = V_o \Lambda_o V_o^T$ , where  $V_o$  contains the top eigenvectors of  $A_o$  (where “top” eigenvectors have eigenvalues at least  $\epsilon n$  in magnitude). Consider the extreme case where  $A_o$  has a flat spectrum – say  $k$  copies of the eigenvalue  $\lambda$  so that  $A_o$  is the scaled projection  $\lambda V_o V_o^T$ . After applying  $S$ , we are left with the matrix  $\lambda S V_o V_o^T S^T$  whose spectrum coincides with  $\lambda V_o^T S^T S V_o$ .

This suggests showing that  $S$  is a subspace embedding for  $V_o$ . Indeed, if  $S$  distorts by at most  $(1 \pm \alpha)$  on  $V_o$ , then for all  $x$  we would have

$$x^T (\lambda V_o^T S^T S V_o) x = \lambda \|S V_o x\|_2^2 = (1 \pm \alpha) \lambda,$$

so that the quadratic form for  $S V_o S^T$  is always at most  $\lambda + \alpha \lambda$  on unit vectors  $x$ , and is at least  $\lambda - \alpha \lambda$  on a  $k$ -dimensional subspace. By Courant-Fischer this would be sufficient to show concentration of the  $k$  eigenvalues of  $S V_o S^T$  to within  $\epsilon n$  of those of  $V_o$  as long as  $\alpha \leq \frac{\epsilon n}{\lambda}$ .

In the worst case,  $\lambda$  could be as large as  $n$  if  $A$  were the all-ones matrix for example. Also since we allow negative eigenvalues it is possible to have up to  $1/\epsilon^2$  eigenvalues of magnitude  $\epsilon n$ . This suggests that we choose  $S$  so that it yields a  $(1 \pm \epsilon)$ -distortion subspace embedding over a  $1/\epsilon^2$  dimensional space. Unfortunately, this would require  $S$  to have at least  $1/\epsilon^3$  rows, which would only allow us to match the sampling bound of [Bha+24]. In order to do better we observe similarly to [SW23], that large eigenvalues cannot occur many times in the spectrum. Since  $A$  has bounded entries, its Frobenius norm is at most  $n$ , which means that there are at most  $n^2/\lambda^2$  eigenvalues of magnitude at least  $\lambda$ . So we could hope to only require a  $1 \pm (\epsilon n/\lambda)$  distortion subspace embedding over a space of dimension  $n^2/\lambda^2$ , which can plausibly be achieved by an  $S$  with roughly  $(n^2/\lambda^2) \cdot (\lambda/(\epsilon n))^2 = 1/\epsilon^2$  rows.

Indeed, by using the “incoherence bound” of [Bha+24], we observe that the leverage scores of  $V_o$  are sufficiently small so that uniform sampling approximates leverage score sampling. This allows to argue that sampling roughly  $1/\epsilon^2$  rows is sufficient to obtain a  $1 \pm (\epsilon n/\lambda)$  distortion subspace embedding on  $V_{\geq \lambda}$ , the eigenvectors with associated eigenvalue at least  $\lambda$  in magnitude. It is possible however that the spectrum of  $A$  is not completely flat, and contains a range of outlying eigenvalues between  $\epsilon n$  and  $n$ . To handle this, we ask for  $S$  to satisfy the following deterministic condition:

$$S \text{ is a } 1 \pm (\epsilon n/\lambda) \text{ distortion subspace embedding on } V_{\geq \lambda} \text{ for all } \lambda \geq \epsilon n.$$

Perhaps surprisingly, we show that this simple condition is enough to approximate all eigenvalues of  $A_o$  to within  $\epsilon n$  additive error.

While the above paragraphs motivate this condition, there are a few technical challenges. The first is handling negative eigenvalues. If  $A$  were PSD, then the subspace embedding argument given above would essentially be sufficient to show that  $SAS^T$  is at least  $\lambda_k - \epsilon n$  on a subspace of dimension  $k$ , which would yield the right lower bound on the  $k$ th eigenvalue of  $SAS^T$ . However we have to worry that the negative eigenvalues of  $A$  could bring down this eigenvalue estimate. We indeed find a  $k$ -dimensional subspace on which  $SAS^T$  is at least  $\lambda_k - \epsilon n$  which is mostly aligned with the top  $k$  eigenvectors of  $A$ . However we choose it carefully so as to avoid the large negative eigenvectors of  $A$ . This is similar to the approach suggested in [SW23], however their argument crucially relied on  $S$  being Gaussian so that the positive and negative eigen-spans of  $SAS^T$  are nearly orthogonal. Our argument shows that this is not necessary; the simple deterministic guarantee stated above suffices.

The other main challenge is upper bounding the top eigenvalue of  $SAS^T$  (which turns out to be sufficient for upper bounding all eigenvalues). [SW23] gave a bound on the top eigenvalue of a Gaussian matrix with non-identity covariance. Unfortunately in the sampling setting, we cannot rely on Gaussianity. Instead, we show that that the same deterministic subspace embedding condition given above that was natural for lower-bounding eigenvalues also suffices for obtaining the necessary upper bounds on the eigenvalues. The proof works by first partitioning eigenvalues into level sets. It is simple to prove the desired upper bound on the top eigenvalue of  $SAS^T$  if  $A$  were restricted to be just the top level set. Then by carefully keeping track of the interaction between each pair of level sets, we are able to show that the additional level sets do not add much mass to the top eigenvalue.

Finally we remove the extra  $\log n$  arising from the bound on the middle eigenvalues [Bha+24]. The idea is simply to observe that after applying the uniform sampling algorithm once, we are again left with a roughly  $s \times s$  matrix, whose spectrum we would like to approximate to additive  $O(\epsilon n)$  error. After rescaling, the entries of this matrix are bounded by  $n/s$ , and so we can now apply the uniform sampling algorithm again to obtain an  $O(\epsilon \cdot \frac{n}{s} \cdot s) = O(\epsilon n)$  additive approximation to the spectrum of the submatrix, this time by sampling  $O(\frac{1}{\epsilon^2} \text{poly log } s)$  rows. Note that  $\log n$  is now replaced with  $\log s$  which is potentially much smaller. If needed, this procedure can be repeated until the dependence on  $\log n$  is removed.

**Squared row-norm sampling.** As with uniform sampling, our algorithm is based off the algorithm of [Bha+24], although with a tighter analysis. Their result requires that one samples a principal submatrix of dimension roughly  $s \times s$  with  $s = O(\frac{1}{\epsilon^8} \text{poly log } n)$  in order to obtain an additive  $\epsilon \|A\|_F$  approximation to the spectrum of  $A$ . Notably, their algorithm does not simply return the spectrum of the sampled submatrix but instead judiciously zeros out certain entries in order to reduce variance. To see why this is necessary, it is instructive to consider the case where  $A$  is the identity matrix. If one samples a  $1/\epsilon^2$  sized principal submatrix and rescales it by  $n\epsilon^2$ , then the resulting matrix will have eigenvalues of size  $\epsilon^2 n$ . This would be fine if one wanted to achieve  $\epsilon n$  additive error. However, for  $\epsilon \|A\|_F = \epsilon \sqrt{n}$  additive error, this is unacceptable. Fortunately, [Bha+24] gives an entry zeroing procedure which zeros out most diagonal entries as well as entries in particularly sparse rows and columns. We are able to employ their zeroing procedure essentially as a black-box.

We do not modify their zeroing procedure, but instead improve their bounds for both the outer and middle eigenvalues of the sampled submatrix. We use the same decomposition  $SAS^T = SA_oS^T + SA_mS^T$  as for the uniform sampling analysis discussed above, where  $S$  now samples each row  $i$  with probability  $p_i$  proportional to its squared row norm, and rescales by  $1/\sqrt{p_i}$ . Our task is two-fold. We first show that the eigenvalues of  $SA_oS^T$  concentrate within  $\epsilon \|A\|_F$  to the eigenvalues of  $A_o$ . Then we bound the operator norm of  $SA_mS^T$  by  $\epsilon \|A\|_F$ , and so Weyl's inequality shows that the eigenvalues of  $SAS^T$  are within  $O(\epsilon n)$  of the eigenvalues of  $A$ .

Bounding the eigenvalues of  $SA_oS^T$  is similar to the uniform sampling case. Squared row-norm sampling provides an approximation to leverage score sampling, which is sufficient to obtain the subspace embedding guarantees that we require.

Bounding the operator norm of  $SA_mS^T$  is what required [Bha+24] to have a  $1/\epsilon^8$  dependence. It turns out that a fairly simple technical improvement to their argument improves this to the optimal  $1/\epsilon^2$  dependence, at least when  $p_i := \frac{s\|A_i\|_2^2}{\|A\|_F^2} \leq 1$ , so that the sampling probabilities for each row are all at most 1. Their argument contains a technical fix for the situation where  $p_i > 1$ , however we do not need this. Instead we can duplicate the rows with large norm, say  $N$  times, while scaling them down by a factor of  $1/\sqrt{N}$ . This does not change the spectrum, but reduces to the situation where all  $p_i$ 's are at most 1. The advantage of this is that rows of  $A$  with large norm can now be sampled multiple times if the value of  $p_i$  dictates that this should occur. We note that one could likely also sample  $s$  rows i.i.d. from the squared-row-norm distribution to achieve a similar result.

**Top eigenvector estimation.** For PSD  $A \in \mathbb{R}^{n \times n}$  with entries bounded by 1, we are interested in producing a vector  $u$  that nearly maximizes the Rayleigh quotient  $\frac{u^T A u}{u^T u}$ . In the same spirit as our eigenvalue estimation results, we are willing to tolerate  $\epsilon n$  additive error, so we would like to find  $u$  that satisfies  $\frac{u^T A u}{u^T u} \geq \lambda_1 - \epsilon n$ .

Perhaps the simplest attempt would be to choose a column sampling matrix  $S$  and form  $S^T A S$ . We now have access to the Rayleigh quotient on the image of  $S$ , and so we could choose  $u \in \text{Im}(S)$  to maximize  $\frac{u^T A u}{u^T u}$ . Unfortunately it is clear that this does not work. For example consider the situation where  $A$  is the all ones matrix and  $S$  samples  $k$  columns. Then  $\lambda_1 = n$ , but  $u$  is supported on  $k$  coordinates which means that the Rayleigh quotient is at most  $k$ . Even to obtain  $n/2$  additive error, we would need to set  $k = \Omega(n)$ , meaning that we would sample a constant fraction of  $A$ 's entries.

Searching for  $u \in \text{Im}(S)$  is not good enough. A better approach would be to optimize the Rayleigh quotient over  $\text{Im}(AS)$ . Indeed if we could implement this approach it would work. The issue is that computing the Rayleigh quotient for such  $U$  would require us to calculate expressions of the form

$$\frac{(ASx)^T A (ASx)}{(ASx)^T (ASx)} = \frac{x^T S^T A^3 S x}{x^T S^T A^2 S x}.$$

Unfortunately  $AS$  does not give us enough information to calculate the numerator. However it does give us enough information to calculate the denominator. This suggests that we should instead optimize  $u$  over  $\text{Im}(A^{1/2} S)$ . Indeed this only requires us to compute quantities of the form

$$\frac{x^T S^T A^2 S x}{x^T S^T A S x},$$

which we can do if we have  $AS$ . We show that this works. That is, for some  $x$ , the generalized Rayleigh quotient above is at least  $\lambda_1 - \epsilon n$ . Thus if  $x$  is optimal, then we can take  $A^{1/2} S x$  as our approximate eigenvector. Given  $x$ , we do not have a direct way to compute  $A^{1/2} S x$ . However we can compute  $ASx$ . This can be thought of as applying an additional half iteration of power method to  $A^{1/2} S x$ , and so as we observe, the Rayleigh quotation for  $ASx$  will be at least as large as for  $A^{1/2} S x$ . Thus our algorithm will ultimately return  $ASx$ .

**Fast eigenvalue sketching.** Previous work, [SW23] considered the problem of sketching the eigenvalues of a matrix  $A$  up to  $\epsilon \|A\|_F$  additive error. While this sketch achieved the optimal  $O(\frac{1}{\epsilon^4})$  sketching dimension, it requires roughly  $\frac{n^2}{\epsilon^2}$  time to apply when  $A$  is a dense  $n \times n$  matrix. We can use our sampling analysis to obtain a faster sketch.

The idea is to rotate by a random orthogonal matrix that supports fast matrix multiplication. Specifically we take a matrix  $U$  which is a Hadamard matrix composed with random sign flips and compute  $U^T A U$ . This flattens the row-norms of  $A$  so that one can simply take the sketch to be a random principle submatrix of  $U^T A U$ . Then applying our analysis for squared row-norm sampling shows that the sketch can be used to obtain an additive  $\epsilon \|A\|_F$  approximation to the spectrum of  $A$ .

## 2.2 Additional Related Work

**Spectral density estimation.** Recently, spectral density estimation was studied by [BKM22], for normalized graph Laplacians and adjacency matrices. This result measures error with respect to Wasserstein distance for the (normalized) spectral histogram, and queries  $\Omega(n/\text{poly}(\epsilon))$  entries of  $A$ , so is not directly comparable to our setting. A related work [Coh+18], achieves the same type of Wasserstein guarantee but with  $\exp(1/\epsilon)$  queries and access to a random walk on the underlying graph. Since we aim for  $\text{poly}(1/\epsilon)$  queries this is again not directly comparable to our setting.

**Matrix sparsification.** [Bha+23] studies the problem of deterministically constructing a sparsifier of bounded entry matrix  $A$  by using a sublinear number of queries. Their approximation result would yield an additive  $\epsilon n$  approximation to all eigenvalues but requires  $\Omega(n/\epsilon^2)$  entry queries. Similarly, for bounded entry PSD matrices they could achieve the same  $\Omega(n/\epsilon^2)$  bound for approximating the top eigenvector of a PSD matrix, but this is worse than what we require for sampling by a factor of  $1/\epsilon$ .

**Low rank approximation.** A line of work [MW17; BW18; BCW20] considered the problem of constructing low-rank approximations to PSD matrices  $A$  using a sublinear number of queries. Notably [BCW20] gave an algorithm with optimal  $O(kn/\epsilon)$  query complexity for obtaining a PSD rank  $k$  approximation  $\hat{A}$  to  $A$  satisfying

$$\|A - \hat{A}\|_F \leq (1 + \epsilon) \|A - A_{(k)}\|_F,$$

where  $A_{(k)}$  is the optimal rank  $k$  approximation of  $A$ . For the case  $k = 1$ , this is related to the top eigenvector approximation problem that we consider. In a similar spirit, [MM17] gives an additive spectral approximation to  $A$  which also implies our top eigenvector bound. However notably these algorithms are all adaptive, whereas we are interested in algorithms that sample non-adaptively. Thus the techniques involved are quite different from what we consider. Additionally these algorithms all require at least an extra  $O(\log \frac{1}{\epsilon})$  overhead compared to our result.

**Sketching.** In a somewhat different setting to sampling, eigenvalue sketching was previously studied by [SW23] where an optimal  $O(1/\epsilon^4)$  sketching dimension was given for obtaining  $\epsilon \|A\|_F$  additive error. Previously [AN13] gave a sketch for the top eigenvalues. By rearranging their bounds, this implies an  $O(1/\epsilon^6)$  for sketching all eigenvalues to  $\epsilon \|A\|_F$  additive error; however this only applies to PSD matrices. Finally, we note the question of sketching the operator norm was considered in [LW16]. These results imply an  $\Omega(n^2)$  lower bound for sketching the top eigenvalue to relative error, which justifies the focus on additive error for sketching and sampling algorithms.

### 3 Preliminaries

**3.1 Notation** Throughout, we use  $c$  to denote an absolute constant, which may change between uses. The notation  $\tilde{O}(f)$  means  $O(f \log^c f)$  for some absolute constant  $c$ .

Unless otherwise stated, the matrix  $A$  always refers to a symmetric (not necessarily PSD) matrix. We use the notation  $A_{(k)}$  to denote optimal rank  $k$  approximation to  $A$ . That is,  $A_{(k)}$  minimizes  $\|A - A_{(k)}\|_F$  over rank  $k$  matrices, where  $\|A\|_F$  is the Frobenius norm. In other words,  $A_{(k)}$  is  $A$  with all but its  $k$  largest magnitude eigenvalues zeroed out. The notation  $A_{(-k)}$  means  $A - A_{(k)}$ . For a matrix  $A$ , the notation  $A_i$  denotes the  $i$ th row of  $A$ . A norm without subscripts is always the  $\ell_2$  norm for vectors and the operator norm for matrices. We sometimes denote the eigenvalues of a matrix  $A \in \mathbb{R}^{n \times n}$  as  $\lambda_1(A) \geq \lambda_2(A) \geq \lambda_n(A)$ . The notation  $\lambda_{\max}(A)$  is equivalent to  $\lambda_1(A)$ . We use  $\|A\|_\infty$  to denote the largest magnitude of an entry in  $A$ . When we say that  $A$  has “bounded entries” we will always mean that  $\|A\|_\infty \leq 1$ .

**3.2 Basic Definitions** We recall the definition of a subspace embedding embedding in the form that we will use it here. This definition is standard in sketching literature (see [Woo+14] for example).

**DEFINITION 3.1.** *We say that  $S \in \mathbb{R}^{k \times n}$  is an  $1 \pm \epsilon$  distortion subspace embedding for  $X \in \mathbb{R}^{n \times m}$  if for all  $v \in \mathbb{R}^m$ ,*

$$(1 - \epsilon) \|Xv\|^2 \leq \|SXv\|^2 \leq (1 + \epsilon) \|Xv\|^2.$$

We will also use the notion of leverage scores early on in order to show that the subspace embedding guarantee that we require is satisfied. We recall the basic definition here.

**DEFINITION 3.2.** *Let  $X \in \mathbb{R}^{n \times d}$  be a matrix. The leverage score for row  $i$  of  $X$  is defined by*

$$\tau_i = e_i^T X (X^T X)^\dagger X^T e_i,$$

where  $(X^T X)^\dagger$  denotes the Moore-Penrose pseudo-inverse of  $X^T X$  and  $e_i \in \mathbb{R}^n$  is the  $i$ th standard basis vector.

**3.3 Incoherence Bound** We use an incoherence bound from [Bha+24] in order to control the leverage scores associated to the top eigenvectors. Roughly these bounds say that eigenvectors corresponding to large eigenvalues must be spread out, and moreover that the supports of such eigenvectors cannot overlap much with one another. For completeness, we supply a short proof here.

**LEMMA 3.1.** *Let  $A$  be a matrix and let  $A_o$  denote the projection of  $A$  onto the eigenvectors with eigenvalue at least  $\alpha$ . Write  $A_o = V_o \Lambda_o V_o^T$  where  $V_o$  has the eigenvectors of  $A$  with eigenvalue at least  $\alpha$  as its columns, and  $\Lambda_o$  is diagonal matrix containing the associated eigenvalues. Then*

$$\|(V_o)_i\|^2 \leq \frac{\|A_i\|^2}{\alpha^2}.$$

*Proof.* We have  $AV_o = V_o \Lambda_o$ . Therefore

$$\|(AV_o)_i\|^2 = \|(V_o \Lambda_o)_i\|^2 \geq \alpha^2 \|(V_o)_i\|^2,$$

since the diagonal entries of  $\Lambda_o$  are all at least  $\alpha$  in magnitude. On the other hand,

$$\|(AV_o)_i\|^2 = \|A_i V_o\|^2 \leq \|A_i\|^2,$$

since  $V_o$  has orthonormal columns. The lemma follows from combining these two bounds.  $\square$

**3.4 Leverage Score Sampling** We use the following leverage score sampling bound, which is a consequence of a matrix Chernoff bound [Kyn18; Tro12].

**THEOREM 3.1.** *Let  $V \in \mathbb{R}^{n \times d}$ , and form a sampling matrix  $S$  by including each index  $i$  with probability  $p_i$  and rescaling by  $\frac{1}{\sqrt{p_i}}$ .*

*Let  $\tau_i$  denote the leverage score for row  $i$  of  $V$ . If  $p_i \geq \min(\frac{2}{\epsilon^2} \tau_i \log \frac{d}{\delta}, 1)$ , then with probability at least  $1 - \delta$  we have  $(1 - \epsilon)V^T V \preceq (SV)^T (SV) \preceq (1 + \epsilon)V^T V$ .*

**3.5 Approximate matrix product** The following observation was Lemma 1 of [CNW15]. We supply a short proof here for completeness.

**LEMMA 3.2.** *Let  $A$  and  $B$  be real matrices and let  $[A|B]$  be the concatenation of  $A$  and  $B$ . Let  $S$  be an  $\epsilon$ -distortion subspace embedding for  $[A|B]$ . Then*

$$\|A^T S^T S B - A^T B\| \leq \epsilon \|A\| \|B\|.$$

*Proof.* Let  $x$  and  $y$  be unit vectors of appropriate dimension such that

$$\|A^T S^T S B - A^T B\| = x^T (A^T S^T S B - A^T B) y = \langle S A x, S B y \rangle - \langle A x, B y \rangle.$$

Since  $S$  is a subspace embedding for  $[A|B]$ , it preserves the inner product  $\langle A x, B y \rangle$  up to additive error:

$$|\langle S A x, S B y \rangle - \langle A x, B y \rangle| \leq \epsilon \|A x\| \|B y\| \leq \epsilon \|A\| \|B\|.$$

The lemma follows.  $\square$

Combining with Theorem 3.1 above immediately gives a generalization of Theorem 3.1 for constructing approximating matrix products. This will account for most of our use cases for leverage score sampling guarantees, so we record this fact here for easy reference.

### 3.6 Additional facts

**PROPOSITION 3.1.** *Let  $A$  be a symmetric matrix. Let  $P$  be a projection onto the span of a subset of  $A$ 's eigenvectors. Then  $\|(PAP^T)_k\| \leq \|A_k\|$  where  $A_k$  is the  $k$ th row of  $A$ .*

*Proof.* Using the spectral decomposition of  $A$ , we can write  $A = \sum_i u_i u_i^T$  and  $PAP^T = \sum_j v_j v_j^T$  where the  $u_i$ 's are mutually orthogonal, and where the  $v_j$ 's are a subset of the  $u_i$ 's. Then

$$\begin{aligned} \|A_k\|^2 &= \|A e_k\|^2 \\ &= \left\| \sum_i \langle u_i, e_k \rangle u_i \right\|^2 \\ &= \sum_i \|u_i\|^2 |\langle u_i, e_k \rangle|^2 \\ &\geq \sum_j \|v_j\|^2 |\langle v_j, e_k \rangle|^2 \\ &= \|(PAP^T)_k\|^2. \end{aligned}$$

$\square$

## 4 Our subspace embedding condition

In the next section we obtain both upper and lower bounds on the outlying eigenvalues. Importantly, we show how to obtain these bounds using only subspace embedding guarantees. The leverage score estimates will show that we can use a sampling procedure to construct the desired subspace embeddings.

We will let  $S$  be our sampling matrix throughout. The key subspace embedding property that we require is as follows.

**ASSUMPTION 4.1.** *For all  $\lambda \geq L$ ,  $S$  is a  $\min(L/\lambda, 1/10)$  distortion subspace embedding on the span of eigenvectors of  $A$  with associated eigenvalue at least  $\lambda$  in magnitude*

To streamline our analysis, we use this assumption in full generality. One should think of  $L$  as representing the threshold defining the outlying eigenvalues. We will later specialize to  $L = \epsilon n$  and  $L = \epsilon \|A\|_F$ . Our first step is to show that Assumption 4.1 is realized by our sampling schemes. Then we treat this property as a black-box for the remainder of the argument, and show that it implies an additive  $O(L)$  spectral approximation guarantee.

**4.1 Realizing Assumption 4.1** As a consequence of the incoherence bound, and leverage score sampling guarantee, we show that a version of Assumption 4.1 is achieved both for uniform sampling and row sampling. For uniform sampling we will assume that  $A$  has bounded entries and set  $L = \epsilon n$ . For row-norm sampling, we will make no assumptions on  $A$  and set  $L = \epsilon \|A\|_F$ . In both cases, we will need our sample to have size roughly  $\frac{1}{\epsilon^2}$ .

**LEMMA 4.1.** *Let  $A \in \mathbb{R}^{n \times n}$  have all entries bounded in magnitude by 1. Let  $S$  be a sampling matrix which samples each row of  $A$  with probability  $s/n$ . Then for  $s \geq \frac{c}{\epsilon^2} (\log \log \frac{1}{\epsilon} + \log \frac{1}{\epsilon^2 \delta})$ ,  $S$  satisfies Assumption 4.1 with  $L = \epsilon n$ .*

*Proof.* Consider a fixed value of  $\lambda$ , and let  $V_{\geq \lambda} \in \mathbb{R}^{n \times d}$  be an orthonormal matrix whose columns are the eigenvectors of  $A$  with associated eigenvalue at least  $\lambda$ .

By Lemma 3.1,

$$\|V_{\geq \lambda, i}\|^2 \leq \frac{1}{\lambda^2} \|A_i\|^2 \leq \frac{n}{\lambda^2},$$

since  $A$  has bounded entries. Since  $V_{\geq \lambda, i}$  has orthonormal columns, this says that all leverage scores of  $V_{\geq \lambda}$  are at least  $n/\lambda^2$ .

Let  $p = s/n$ , and set  $\epsilon = \frac{1}{2}(L/\lambda)$  in Theorem 3.1. This shows that if

$$p \geq 8 \frac{\lambda^2}{L^2} \left( \frac{n}{\lambda^2} \right) \log \frac{d}{\delta} = 8 \frac{n}{L^2} \log \frac{d}{\delta} = \frac{8}{\epsilon^2 n} \log \frac{d}{\delta},$$

then  $S$  is a  $\frac{1}{2}(L/\lambda)$  distortion subspace embedding for  $V_{\geq \lambda}$ , with probability at least  $1 - \delta$ . Note that  $\|A\|_F^2 \leq n^2$ , so  $A$  has at most  $1/\epsilon^2$  eigenvalues that are at least  $\epsilon n$ . This implies that  $d \leq \frac{1}{\epsilon^2}$  so it suffices to have  $p \geq \frac{8}{\epsilon^2 n} \log \frac{1}{\epsilon^2 \delta}$ , and hence to choose  $s \geq \frac{8}{\epsilon^2} \log \frac{1}{\epsilon^2 \delta}$ .

In order for Assumption 4.1 to hold, it suffices to have  $S$  be a  $1/10$  distortion embedding for  $V$ , a  $\frac{1}{2}$  distortion embedding for  $V_{\geq \epsilon n}$ , a  $\frac{1}{4}$  distortion embedding for  $V_{\geq 2\epsilon n}$ , and in general a  $\frac{1}{2^{r+1}}$  distortion embedding for  $V_{\geq 2^r \epsilon n}$  for  $r = 0, \dots, \lfloor \log_2 \frac{1}{\epsilon} \rfloor$ . By the same calculations as above, we can achieve the  $1/10$  distortion embedding for  $V$  using a sampling probability of  $p \geq \frac{c}{\epsilon^2 n} \log \frac{1}{\delta}$ . Then replacing  $\delta$  by  $\delta/\lfloor \log_2 \frac{1}{\epsilon} \rfloor$  and taking a union bound over all  $r$  yields the lemma.

□

The following lemma gives the analogous result for squared row-norm sampling. For technical reasons, we give a version that will allow some entries of  $A$  to be zeroed out prior to sampling, although this does not change much.

**LEMMA 4.2.** *Let  $A \in \mathbb{R}^{n \times n}$  be arbitrary. Let  $S$  be a sampling matrix which samples each row index  $i$  with probability  $p_i = \min \left( s \frac{\|A_i\|^2}{\|A\|_F^2}, 1 \right)$ . Let  $A'$  be the matrix  $A$ , but possibly with some of its entries replaced with 0. Then for  $s \geq \frac{c}{\epsilon^2} (\log \log \frac{1}{\epsilon} + \log \frac{1}{\epsilon^2 \delta})$ ,  $S$  satisfies Assumption 4.1 for  $A'$  with  $L = \epsilon \|A\|_F$ .*

*Proof.* Consider a fixed value of  $\lambda$ , and let  $V'_{\geq \lambda} \in \mathbb{R}^{n \times d}$  be an orthonormal matrix whose columns are the eigenvectors of  $A'$  with associated eigenvalue at least  $\lambda$ .

By Lemma 3.1,

$$\|V'_{\geq \lambda, i}\|^2 \leq \frac{1}{\lambda^2} \|A'_i\|^2 \leq \frac{1}{\lambda^2} \|A_i\|^2.$$

Set  $\epsilon = \frac{1}{2}(L/\lambda) = \frac{1}{2} \frac{\epsilon \|A\|_F}{\lambda}$  in Theorem 3.1. This shows that if

$$p_i \geq \min \left( 8 \frac{\lambda^2}{L^2} \left( \frac{\|A_i\|^2}{\lambda^2} \right) \log \frac{d}{\delta}, 1 \right) = \min \left( \frac{8 \|A_i\|^2}{\epsilon^2 \|A\|_F^2} \log \frac{d}{\delta}, 1 \right),$$

then  $S$  is a  $\frac{1}{2}(L/\lambda)$  distortion subspace embedding for  $V'_{\geq \lambda}$ , with probability at least  $1 - \delta$ . Note that  $A'$  has at most  $1/\epsilon^2$  eigenvalues that are at least  $\epsilon \|A\|_F$ . This implies that  $d \leq \frac{1}{\epsilon^2}$  so it suffices to have

$$p_i \geq \min \left( \frac{8 \|A_i\|^2}{\epsilon^2 \|A\|_F^2} \log \frac{1}{\epsilon^2 \delta}, 1 \right),$$

and hence to choose  $s \geq \frac{8}{\epsilon^2} \log \frac{1}{\epsilon^2 \delta}$ .

In order for Assumption 4.1 to hold, it suffices to have  $S$  be a  $1/10$  distortion embedding for  $V'$ , a  $\frac{1}{2}$  distortion embedding for  $V'_{\geq \epsilon n}$ , a  $\frac{1}{4}$  distortion embedding for  $V'_{\geq 2\epsilon n}$ , and in general a  $\frac{1}{2^{r+1}}$  distortion embedding for  $V'_{\geq 2^r \epsilon n}$  for  $r = 0, \dots, \lfloor \log_2 \frac{1}{\epsilon} \rfloor$ . By the same calculations as above, we can achieve the  $1/10$  distortion embedding for  $V'$  using a sampling probability of  $p \geq \frac{c}{\epsilon^2 n} \log \frac{1}{\delta}$ . Then replacing  $\delta$  by  $\delta / \lfloor \log_2 \frac{1}{\epsilon} \rfloor$  and taking a union bound over all  $r$  yields the lemma.  $\square$

## 5 Outlying Eigenvalue Bounds

Using Assumption 4.1 we show how to obtain both upper and lower bound on the outlying eigenvalues of  $A$ .

### 5.1 Lower Bounds

**PROPOSITION 5.1.** *Let  $W_1$  and  $W_2$  be subspaces of  $\mathbb{R}^n$ , and let  $P$  be the orthogonal projection onto  $W_2$ . Let  $x$  be a nonzero vector in  $W_1$ . Then*

$$\frac{\|Px\|}{\|x\|} \leq \max_{w_1 \in W_1, w_2 \in W_2} \frac{\langle w_1, w_2 \rangle}{\|w_1\| \|w_2\|}.$$

*Proof.* For  $w_1 \in W_1$  and  $w_2 \in W_2$  we have

$$\max_{w_1, w_2} \frac{\langle w_1, w_2 \rangle}{\|w_1\| \|w_2\|} = \max_{w_1} \frac{\langle w_1, Pw_1 \rangle}{\|w_1\| \|Pw_1\|} \geq \frac{\langle x, Px \rangle}{\|x\| \|Px\|} = \frac{\|Px\|}{\|x\|}.$$

$\square$

**PROPOSITION 5.2.** *Let  $A$  be symmetric and  $x$  and  $h$  be vectors. Suppose that  $\|x\| \geq 1 - \alpha$  and  $\|h\| \leq \alpha$  with  $\alpha \leq 1/4$ . Then*

$$\left| \frac{(x+h)^T A (x+h)}{\|x+h\|^2} - \frac{x^T A x}{\|x\|^2} \right| \leq 8 \|A\| \alpha.$$

*Proof.* Let  $F(x) = \frac{x^T A x}{x^T x}$  be the Rayleigh quotient. Recall that its gradient is given by

$$\nabla F(v) = \frac{2}{\|v\|^2} \left( A - \frac{v^T A v}{\|v\|^2} I \right) v.$$

Then

$$\|\nabla F(v)\| \leq \frac{2}{\|v\|^2} \left\| A - \frac{v^T A v}{\|v\|^2} I \right\| \|v\| \leq \frac{2}{\|v\|} \left( \|A\| + \left\| \frac{v^T A v}{\|v\|^2} I \right\| \right) \leq 4 \frac{\|A\|}{\|v\|}.$$

This together with the hypotheses implies that  $F$  is  $(8 \|A\|)$ -Lipschitz on a  $\|h\|$ -neighborhood of  $x$ , and the claim follows.  $\square$

LEMMA 5.1. *Let  $\Pi$  have orthonormal columns and suppose that  $S$  is an  $\alpha$ -distortion subspace embedding for  $\Pi$ . Then for all nonzero  $v$ ,*

$$\frac{(S\Pi v)^T S\Pi \Pi^T S^T (S\Pi v)}{\|S\Pi v\|^2} \geq 1 - \alpha.$$

*Proof.* The numerator above is equal to  $\|\Pi^T S^T S\Pi v\|^2$ . By Cauchy-Schwarz,

$$\|\Pi^T S^T S\Pi v\|^2 \geq \frac{\langle v, \Pi^T S^T S\Pi v \rangle^2}{\|v\|^2} = \frac{\|S\Pi v\|^4}{\|v\|^2}.$$

So the fraction in the lemma is at least  $\frac{\|S\Pi v\|^2}{\|v\|^2}$ . Since  $S$  is an  $\alpha$  distortion subspace embedding for  $\Pi$ , we have

$$\|S\Pi v\|^2 \geq (1 - \alpha) \|\Pi v\|^2 = (1 - \alpha) \|v\|^2,$$

and the lemma follows.  $\square$

LEMMA 5.2. *Let  $U$  and  $\Pi$  have mutually orthonormal columns and let  $M$  be symmetric. Suppose that  $S$  is an  $\alpha$ -distortion embedding for  $[U|\Pi]$ , where  $\alpha \leq 1/7$ . Let  $P$  be the orthogonal projection onto the column space of  $SU$  and let  $P^\perp = I - P$ . Then for all vectors  $x$ ,  $P^\perp S\Pi x$  is nonzero and*

$$\left| \frac{(P^\perp S\Pi x)^T M(P^\perp S\Pi x)}{\|P^\perp S\Pi x\|^2} - \frac{(S\Pi x)^T M(S\Pi x)}{\|S\Pi x\|^2} \right| \leq 10\alpha \|M\|.$$

*Proof.* We first show that  $P^\perp S\Pi x$  is close to  $S\Pi x$ . By Proposition 5.1 we have

$$\frac{\|PS\Pi x\|}{\|S\Pi x\|} \leq \max_{y,z} \frac{\langle S\Pi y, SUz \rangle}{\|S\Pi y\| \|SUz\|} = \max_{y,z} \frac{y^T \Pi^T S^T SUz}{\|S\Pi y\| \|SUz\|} \leq \max_{y,z} \frac{\|y\| \|z\| \|\Pi^T S^T SU\|}{\|S\Pi y\| \|SUz\|}.$$

By our hypothesis that  $S$  is a subspace embedding, along with Lemma 3.2, we bound the operator norm in the numerator by

$$\|\Pi^T S^T SU\| = \|\Pi^T S^T SU - \Pi^T U\| \leq \alpha \|\Pi\| \|U\| = \alpha.$$

For the denominator, the subspace embedding property implies that

$$(1 - \alpha)^{1/2} \|y\| \leq \|S\Pi y\| \leq (1 + \alpha)^{1/2} \|y\|$$

and similarly for  $\|SUz\|$  and  $\|S\Pi z\|$ . Plugging into the bound above gives that for all unit vectors  $x$ ,

$$\|PS\Pi x\| \leq \frac{\alpha}{(1 - \alpha)^{1/2}(1 - \alpha)^{1/2}} \|S\Pi x\| \leq \frac{\alpha}{1 - \alpha} (1 + \alpha)^{1/2} < 1.25\alpha,$$

for  $\alpha \leq 1/7$ . By the triangle inequality,

$$0 = \|PS\Pi x + P^\perp S\Pi x - S\Pi x\| \geq \|P^\perp S\Pi x - S\Pi x\| - \|PS\Pi x\|,$$

and so  $\|P^\perp S\Pi x - S\Pi x\| \leq 1.25\alpha$  for all unit vectors  $x$ .

Note that  $\|S\Pi x\| \geq (1 - \alpha)^{1/2} \|x\| \geq (1 - \alpha)$ . To apply Proposition 5.2 we need  $1.25\alpha \leq 1/4$ , which is true by hypothesis.

So by Proposition 5.2 we have

$$\left| \frac{(P^\perp S\Pi x)^T M(P^\perp S\Pi x)}{\|P^\perp S\Pi x\|^2} - \frac{(S\Pi x)^T M(S\Pi x)}{\|S\Pi x\|^2} \right| \leq (8 \cdot 1.25)\alpha \|M\| = 10\alpha \|M\|,$$

as desired.  $\square$

LEMMA 5.3. Let  $A \in \mathbb{R}^{n \times n}$  be symmetric, not necessarily PSD. Assume that all nonzero eigenvalues of  $A$  are at least  $L$  in magnitude. Suppose that  $S$  satisfies Assumption 4.1. Then for all  $k$  with  $\lambda_k(A) > 0$  we have  $\lambda_k(SAS^T) \geq \lambda_k(A) - 51L$ .

*Proof.* Suppose that the  $k$ th largest eigenvalue of  $A$  is  $\lambda > 0$ . We will show that  $SAS^T$  is at least  $\lambda - O(L)$  on a  $k$ -dimensional subspace  $W$ , which will imply that  $\lambda_k(SAS^T) \geq \lambda - O(L)$ .

Towards this end, decompose  $A = A_+ - A_-$  into its positive and negative parts. Since  $A_+$  has  $k$  eigenvalues that are at least  $\lambda$ , we can write  $A_+ \succeq \lambda \Pi \Pi^T$  where the columns of  $\Pi$  are the  $k$  largest eigenvectors of  $A_+$  (chosen to be orthonormal when there are duplicate eigenvalues). Let  $U_{\geq \lambda}$  have as its columns the eigenvectors of  $A_-$  with associated eigenvalue at least  $\lambda$ , and similarly for  $U_{< \lambda}$ . Let  $P$  be the orthogonal projection onto the column span of  $SU_{\geq \lambda}$  and set  $P^\perp = I - P$ . As suggested by the previous lemma, we choose our subspace  $W$  to be the column span of  $P^\perp \Pi$ .

Write  $f(\lambda) := \min(1/7, L/\lambda)$  and recall that  $S$  is an  $f(\lambda)$ -distortion subspace on the span of the eigenvectors of  $A$  with associated eigenvalue at least  $\lambda$  in magnitude, in other words on  $[\Pi|U_{\geq \lambda}]$ . Lemma 5.2 applies and gives that for all  $x$ ,

$$\begin{aligned} \frac{(P^\perp \Pi x)^T (S \Pi \Pi^T S^T) (P^\perp \Pi x)}{\|P^\perp \Pi x\|^2} &\geq \frac{(\Pi x)^T (S \Pi \Pi^T S^T) (\Pi x)}{\|\Pi x\|^2} - 10f(\lambda) \|S \Pi \Pi^T S^T\| \\ &\geq (1 - f(\lambda)) - 10f(\lambda) \|S \Pi \Pi^T S^T\| \\ &\geq 1 - 12f(\lambda), \end{aligned}$$

where we used Lemma 5.1 as well as the observation that

$$\|S \Pi \Pi^T S^T\| = \|\Pi^T S^T \Pi\| \leq (1 + f(\lambda))^{1/2} \leq (1 + 1/7)^{1/2} \leq 1.07,$$

by the subspace embedding property. Also note that the dimension of  $W$  is  $k$  since Lemma 5.2 also states that  $P^\perp \Pi$  has trivial kernel. Now for  $w \in W$ , with  $\|w\| = 1$  we have

$$w^T S A_+ S^T w \geq \lambda w^T S \Pi \Pi^T S^T w = \lambda \|\Pi^T S^T w\|^2 \geq \lambda(1 - f(\lambda)).$$

Next we study  $w^T A_- w$  where  $w \in W$  is a unit vector. We first split the spectral decomposition of  $A_-$  into two pieces as follows:

$$A_- = U_{\geq \lambda} \Lambda_{\geq \lambda} U_{\geq \lambda}^T + U_{< \lambda} \Lambda_{< \lambda} U_{< \lambda}^T.$$

Then we have

$$(5.1) \quad w^T (S A_- S^T) w = w^T S U_{\geq \lambda} \Lambda_{\geq \lambda} U_{\geq \lambda}^T S^T w + w^T S U_{< \lambda} \Lambda_{< \lambda} U_{< \lambda}^T S^T w$$

$$(5.2) \quad = w^T S U_{< \lambda} \Lambda_{< \lambda} U_{< \lambda}^T S^T w,$$

where the second inequality is because  $w$  is orthogonal to the image of  $SU_{\geq \lambda}$ , which comes from the definition of  $W$ .

We now partition the eigenvalues into approximate level sets. Let  $V_k$  have as its columns the columns of  $U_{< \lambda}$  with associated eigenvalue in  $[2^{k-1}L, 2^kL]$ . Then we have the spectral bound

$$U_{\leq \lambda} \Lambda_{\leq \lambda} U_{\leq \lambda}^T \preceq (2L)V_1 V_1^T + (4L)V_2 V_2^T + \dots + (2^r L)V_r V_r^T,$$

where  $r$  is  $\lceil \log_2(\lambda/L) \rceil$ .

We would like to bound  $w^T S V_k V_k^T S^T w$ . Write  $w = P^\perp \Pi x$  for some  $x$ . Then similar to our application of Lemma 5.2 above,

$$\frac{(P^\perp \Pi x)^T (S V_k V_k^T S^T) (P^\perp \Pi x)}{\|P^\perp \Pi x\|^2} \leq \frac{(\Pi x)^T (S V_k V_k^T S^T) (\Pi x)}{\|\Pi x\|^2} + 10f(\lambda) \|S V_k V_k^T S^T\|.$$

Since  $S$  a  $1/7$  distortion embedding for  $A$  and in particular  $V_k$ , we have

$$\|S V_k V_k^T S^T\| = \|V_k^T S^T S V_k\| = \max_{\|x\|=1} \|S V_k x\|^2 \leq \frac{8}{7}.$$

For the first term above,

$$\frac{(S\Pi x)^T(SV_k V_k^T S^T)(S\Pi x)}{\|S\Pi x\|^2} = \frac{\|V_k^T S^T S\Pi x\|^2}{\|S\Pi x\|^2} \leq \frac{\|V_k^T S^T S\Pi\|^2 \|x\|^2}{\|S\Pi x\|^2}.$$

Since  $S$  is an  $f(2^{k-1}L)$ -distortion subspace embedding for  $[V_k|\Pi]$ , we have

$$\|V_k^T S^T S\Pi\| \leq f(2^{k-1}L).$$

Also  $\|S\Pi x\|^2 \geq \frac{6}{7} \|x\|^2$  again since  $S$  is a subspace embedding for  $\Pi$ , and so

$$\frac{(P^\perp S\Pi x)^T(SV_k V_k^T S^T)(P^\perp S\Pi x)}{\|P^\perp S\Pi x\|^2} \leq \frac{7}{6} f(2^{k-1}L)^2 + 12f(\lambda).$$

In particular, for  $w \in W$  a unit vector,

$$w^T(SV_k V_k^T S^T)w \leq \frac{7}{6} f(2^k L)^2 + 12f(\lambda).$$

Plugging into (5.1) above gives

$$\begin{aligned} w^T(SA_- S^T)w &= w^T S U_{\leq \lambda} \Lambda_{\leq \lambda} U_{\leq \lambda}^T S^T w \\ &\leq w^T S \left( (2L)V_1 V_1^T + (4L)V_2 V_2^T + \dots + (2^r L)V_r V_r^T \right) S^T w \\ &= \sum_{k=1}^r (2^k L) w^T S V_k V_k^T S^T w \\ &\leq \sum_{k=1}^r (2^k L) \left( \frac{7}{6} f(2^k L)^2 + 12f(\lambda) \right) . \\ &\leq \sum_{k=1}^r (2^k L) \left( \frac{7}{6} \left( \frac{L}{2^k L} \right)^2 + 12 \frac{L}{\lambda} \right) \\ &= L \sum_{k=1}^r \left( \frac{7}{6} 2^{-k} + 2^k \frac{12L}{\lambda} \right) \\ &\leq \frac{7}{6} L + 12L \sum_{k=1}^r \frac{2^k L}{\lambda} \\ &\leq \frac{7}{6} L + 48L, \end{aligned}$$

since  $2^r L \leq 2\lambda$ .

Putting the pieces together,

$$\begin{aligned} w^T S A S^T w &= w^T S A_+ S^T w - w^T S A_- S^T w \\ &\geq \lambda(1 - f(\lambda)) - 50L \\ &\geq \lambda \left( 1 - \frac{L}{\lambda} \right) - 50L \\ &= \lambda - 51L, \end{aligned}$$

for all unit vectors  $w \in W$ . Since the dimension of  $W$  is  $k$ , this implies that  $\lambda_k(SAS^T) \geq \lambda - 51L$  as desired.  $\square$

## 5.2 Upper Bounds

LEMMA 5.4. *Suppose that the nonzero eigenvalues of  $A$  are all at least  $L$  and let  $\lambda_1 = \lambda_{\max}(A)$ . Suppose that  $S$  satisfies Assumption 4.1. Then  $\lambda_{\max}(SAS^T) \leq \lambda_1 + cL \log \frac{\lambda_1}{L}$ .*

*Proof.* Let  $A = U\Lambda U^T$  be the spectral decomposition of  $A$ .

We partition the columns of  $U$  based on the eigenvalues. For  $i = 1, \dots, r$  with  $r = \lceil \log_2 \frac{\lambda_1}{L} \rceil$ , let  $U^{(i)}$  consist of the columns of  $U$  corresponding to eigenvalues in  $(\lambda_1 2^{-i}, \lambda_1 2^{-i+1}]$ . Define  $\Lambda^{(i)}$  to be the diagonal matrix consisting of the associated eigenvalues.

Let  $j \leq i$  be an arbitrary index. Note that if  $S$  is a subspace embedding for  $U^{(i)}$  then  $S$  is also a subspace embedding for  $U^{(i)}\Lambda^{(i)}$  with the same parameters.

By Assumption 4.1,  $S$  is a  $L/(\lambda_1 2^{-i}) = 2^i L/\lambda_1$  distortion subspace for  $[U^{(i)}|U^{(j)}]$ . Therefore we have

$$\begin{aligned} \left\| \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( SU^{(j)}(\Lambda^{(j)})^{1/2} \right) - \left( U^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( U^{(j)}(\Lambda^{(j)})^{1/2} \right) \right\| &\leq \frac{2^i L}{\lambda_1} \left\| U^{(i)}(\Lambda^{(i)})^{1/2} \right\| \left\| U^{(j)}(\Lambda^{(j)})^{1/2} \right\| \\ &\leq \frac{2^i L}{\lambda_1} (\lambda_1 2^{-i+1})^{1/2} (\lambda_1 2^{-j+1})^{1/2} \\ &= 2L\sqrt{2}^{i-j}. \end{aligned}$$

For  $i > j$ ,  $U^{(i)}$  and  $U^{(j)}$  are orthogonal to one another, so

$$\left\| \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( SU^{(j)}(\Lambda^{(j)})^{1/2} \right) \right\| \leq 2L\sqrt{2}^{i-j}.$$

Therefore for all  $i \neq j$ ,

$$\left\| \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( SU^{(j)}(\Lambda^{(j)})^{1/2} \right) \right\| \leq 2L\sqrt{2}^{|i-j|}.$$

Similarly for the  $i = j$  case, we have

$$\left\| \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right) - \left( U^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( U^{(i)}(\Lambda^{(i)})^{1/2} \right) \right\| \leq 2L,$$

which by the triangle inequality implies that

$$\left\| \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right)^T \left( SU^{(i)}(\Lambda^{(i)})^{1/2} \right) \right\| \leq \lambda_1 2^{-i+1} + 2L.$$

Now let  $x$  be a unit vector, and partition the coordinates of  $x$  into pieces  $x^{(i)}$  corresponding to the decomposition of  $U$  into  $U^{(i)}$ 's. Then

$$\left\| SU\Lambda^{1/2}x \right\|^2 = \left\| \sum_i SU^{(i)}(\Lambda^{(i)})^{1/2}x^{(i)} \right\|^2 = \sum_i \left\| SU^{(i)}(\Lambda^{(i)})^{1/2}x^{(i)} \right\|^2 + 2 \sum_{i < j} \left\langle SU^{(i)}(\Lambda^{(i)})^{1/2}x^{(i)}, SU^{(j)}(\Lambda^{(j)})^{1/2}x^{(j)} \right\rangle$$

By the bounds above, the cross terms are all bounded by  $L\sqrt{2}^{|i-j|} \|x^{(i)}\| \|x^{(j)}\|$  and the diagonal terms are bounded by  $(\lambda_1 2^{-i+1} + L) \|x^{(i)}\|^2$ .

Thus

$$\begin{aligned} \left\| SU\Lambda^{1/2}x \right\|^2 &\leq \sum_{i=1}^r \left[ (\lambda_1 2^{-i+1} + L) \|x^{(i)}\|^2 \right] + 2 \sum_{i < j} L\sqrt{2}^{|i-j|} \|x^{(i)}\| \|x^{(j)}\| \\ &= L + \sum_{i=1}^r \lambda_1 2^{-i+1} \|x^{(i)}\|^2 + 2L \sum_{i < j} \sqrt{2}^{|i-j|} \|x^{(i)}\| \|x^{(j)}\|. \end{aligned}$$

Let  $M$  be an  $r \times r$  matrix defined by  $M_{1,1} = \lambda_1$ ,  $M_{i,i} = \lambda_1/2$  for  $i > 1$ , and for  $i \neq j$   $M_{i,j} = M_{j,i} = L\sqrt{2}^r := \alpha$ . Comparing the entries of  $M$  with the coefficients above gives

$$\left\| SU\Lambda^{1/2}x \right\|^2 \leq L + \|M\|.$$

To calculate  $\|M\|$ , note from the symmetry of the last  $r - 1$  coordinates that  $v^T M v$  is maximized for a  $v$  of the form  $[a, b/\sqrt{r-1}, b/\sqrt{r-1}, \dots, b/\sqrt{r-1}]$ , where  $a^2 + b^2 = 1$ . For this choice of  $v$ , we have

$$v^T M v = \lambda_1 a^2 + 2\alpha\sqrt{r-1}ab + (\lambda_1/2 + (r-2)\alpha)b^2.$$

So the operator norm of  $M$  is the same as the operator norm of

$$M' := \begin{bmatrix} \lambda_1 & \alpha\sqrt{r-1} \\ \alpha\sqrt{r-1} & \lambda_1/2 + \alpha(r-2) \end{bmatrix}.$$

Now suppose that  $\lambda_1 + T$  is an eigenvalue of  $M'$ . From the characteristic equation,

$$T(T + \lambda_1/2 - \alpha(r-2)) = \alpha^2(r-1).$$

We claim that  $\lambda_1/2 - \alpha(r-2) \geq \lambda_1/4$  when  $\lambda_1 \geq CL$  for an absolute constant  $C$ . To see this, note that the desired bound is implied by  $\lambda_1/4 \geq \alpha r$ . Note that  $\alpha = L\sqrt{2^r} \leq \sqrt{2}\sqrt{L\lambda_1}$ , so it suffices to check that  $\lambda_1/4 \geq \sqrt{2}\sqrt{L\lambda_1}(\log_2(\lambda_1/L) + 1)$ , or equivalently  $\sqrt{\lambda_1/L} \geq 4\sqrt{2}(\log_2(\lambda_1/L) + 1)$ . The latter is true for large enough  $C$ , since we have the real inequality  $\sqrt{x} \geq 4\sqrt{2}(\log_2(x) + 1)$  for large enough  $x$ .

Then for  $\lambda_1 \geq CL$  we have

$$T(T + \lambda_1/2 - \alpha(r-2)) \geq T(T + \lambda_1/4) \geq T\lambda_1/4,$$

while on the other hand  $\alpha^2(r-1) \leq 2L\lambda_1 r$ . So we conclude that

$$T\lambda_1/4 \leq 2L\lambda_1 r,$$

which implies that  $T \leq 8Lr = 8L\lceil\log_2(\frac{\lambda_1}{L})\rceil$ . This in turn gives the bound

$$\|SUA^{1/2}x\|^2 \leq \lambda_1 + 8L\lceil\log_2(\frac{\lambda_1}{L})\rceil.$$

But  $x$  was arbitrary so

$$\|SAS^T\| = \|\Lambda^{1/2}U^TS^TSUA^{1/2}\| = \|SUA^{1/2}\|^2 \leq \lambda_1 + 8L\lceil\log_2(\frac{\lambda_1}{L})\rceil$$

whenever  $\lambda_1 \geq CL$ . The lemma follows after adjusting constants.

□

LEMMA 5.5. *Let  $A \in \mathbb{R}^{n \times n}$  with  $\lambda_i(A) > 0$ . Suppose that  $S$  satisfies Assumption 4.1 for  $A$ . Then  $\lambda_i(SAS^T) \leq \lambda_i(A) + L \log \frac{\lambda_i(A)}{L}$ .*

*Proof.* First note that we may as well prove this result for PSD  $A$ , since  $\lambda_i(SAS^T) \leq \lambda_i(SA_+S^T)$ , where  $A_+$  is the PSD part of  $A$ . So we assume from now on that  $A$  is PSD. Let  $A_{\langle -i \rangle}$  denote  $A$  with the top  $i$  eigenvalues zeroed out, and let  $S_i$  denote the orthogonal complement of the span of the top  $i - 1$  eigenvectors of  $A$ . Note that  $S_i$  has dimension  $n - i + 1$ . By the min-max theorem,

$$\begin{aligned} \lambda_i(SAS^T) &= \lambda_i(A^{1/2}S^TSA^{1/2}) \\ &= \min_{\substack{T \subseteq \mathbb{R}^n \\ \dim(T) = n-i+1}} \max_{\substack{x \in T \\ \|x\|=1}} x^T A^{1/2} S^T S A^{1/2} x \\ &\leq \max_{\substack{x \in S_i, \|x\|=1}} x^T A^{1/2} S^T S A^{1/2} x \\ &= \max_{\substack{x \in S_i, \|x\|=1}} x^T A_{-i}^{1/2} S^T S A_{\langle -i \rangle}^{1/2} x \\ &\leq \max_{\|x\|=1} x^T A_{\langle -i \rangle}^{1/2} S^T S A_{-i}^{1/2} x \\ &= \lambda_{\max}(A_{\langle -i \rangle}^{1/2} S^T S A_{-i}^{1/2}) \\ &= \lambda_{\max}(S A_{\langle -i \rangle} S^T). \end{aligned}$$

It is clear that if Assumption 4.1 is satisfied for  $A$  then it is satisfied for  $A_{(-i)}$  as well. Thus the previous lemma applies and gives

$$\lambda_i(SAS^T) \leq \lambda_i(A) + L \log \frac{\lambda_i(A)}{L},$$

as desired.  $\square$

**5.3 Two-sided eigenvalue bound.** By simply combining our upper and lower eigenvalue bounds, we obtain a two-sided bound that we will use throughout.

**THEOREM 5.1.** *Suppose that  $S$  is such that Assumption 4.1 holds for  $A$  where all nonzero eigenvalues of  $A$  are at least  $L$  in magnitude. Then for all  $i$  with  $\lambda_i(A) \neq 0$  we have*

$$|\lambda_i(SAS^T) - \lambda_i(A)| \leq L \log \frac{\|A\|}{L}.$$

*Proof.* Lemma 5.5 and Lemma 5.3 yield the desired bound for all  $i$  with  $\lambda_i(A) > 0$ . To obtain the analogous bound for the negative eigenvalues, simply apply these Lemmas to  $-A$ .  $\square$

## 6 Bounded Entry Matrices

As in [Bha+24], we decompose  $A$  corresponding to its “middle” and “outlying” eigenvalues.

**DEFINITION 6.1.** *Let  $A$  be a symmetric matrix with spectral decomposition  $A = U\Lambda U^T$ . Let  $L$  be the parameter in Assumption 4.1. Set  $\Lambda_o$  to be  $\Lambda$  with all entries smaller than  $L$  zeroed out, and set  $\Lambda_m = \Lambda - \Lambda_o$ . Then we define  $A_o = U\Lambda_o U^T$  and  $A_m = U\Lambda_m U^T$ .*

Throughout we will slightly abuse notation by writing  $A_{o,i}$  to refer to row  $i$  of the matrix  $A_o$ , and similarly  $A_{o,i,j}$  to refer to the  $i,j$  entry of  $A_o$ . We use the following bound on the middle eigenvalues from [Bha+24].

**LEMMA 6.1.** *Let  $A \in \mathbb{R}^{n \times n}$  be symmetric with entries bounded by 1, and let  $A_m$  be  $A$  restricted to its middle eigenvalues as defined in Definition 6.1. Let  $S$  be a uniform sampling matrix with sampling probability at least  $\frac{c \log n}{\epsilon^2 \delta}$ . Then with probability at least  $1 - \delta$ ,  $\|SA_m S^T\|_2 \leq \epsilon n$ .*

*Proof.* This is a rephrasing of Lemma 4 of [Bha+24]. Note that our sampling matrix  $S$  already rescales by  $n/s$ , whereas the matrix in their theorem statement is prior to rescaling.  $\square$

**THEOREM 6.1.** *Let  $A \in \mathbb{R}^{n \times n}$  be symmetric, not necessarily PSD with all entries bounded by 1. Then when  $s \geq \frac{c}{\epsilon^2} \frac{\log n}{\delta}$ , observing  $SAS^T$  allows recovery of all eigenvalues of  $A$  to within  $\epsilon n$  additive error with probability at least  $1 - \delta$ .*

*Proof.* Following [Bha+24], write  $A = A_o + A_m$  where  $A_o$  has the eigenvalues of  $A$  that are at least  $\epsilon n$  in magnitude. For our sampling matrix  $S$ , Theorem 5.1 implies that the positive eigenvalues of  $SA_o S^T$  are  $\epsilon n$  additive approximations to the eigenvalues of  $A_o$ . Also Lemma 6.1 shows that  $\|SA_m S^T\| \leq O(\epsilon n)$  with 0.9 probability. So the result follows from Weyl’s inequality.  $\square$

**6.1 Removing  $\log n$  factors.** We show via a surprisingly simple trick that the  $\log n$  dependence can be replaced with a  $\log \frac{1}{\epsilon \delta}$  dependence, resolving a question left open by [Bha+24]. The idea is that we can directly apply our sampling result to get good eigenvalue approximations by sampling a principal submatrix of dimension  $\tilde{O}(\log n / \epsilon^2)$ . The new matrix still has a dimension depending on  $n$ , but is dramatically smaller. Thus we may apply our sampling procedure again to reduce the dimensions even further. It turns out that we can repeat this enough times to remove the  $n$  dependence entirely. Since the sampling procedure at each stage is uniform, the final sample is uniform as well. So this argument does not yield a new algorithm, but rather shows that the  $O(\log n / \epsilon)$  guarantee can be bootstrapped to achieve a tighter sampling bound.

**THEOREM 6.2.** *Algorithm 1 approximates all eigenvalues of  $A$  to within  $\epsilon n$  additive error, with probability at least  $1 - \delta$  when  $s = \frac{c}{\epsilon^2 \delta} \log \frac{1}{\epsilon \delta}$ .*

*Proof.* Say that a sampling algorithm has a  $(\epsilon, \delta)$  recovery guarantee, if it approximates all eigenvalues of  $A$  to within  $\epsilon n$  additive error with failure probability at most  $\delta$ . Fix  $n$ , and let  $N \geq 1$  be minimal such that for all  $\epsilon, \delta$  bounded by a sufficiently small constant, setting  $s = N/(\epsilon^2 \delta)$  in Algorithm 1 suffices for obtaining an  $(\epsilon, \delta)$  approximation guarantee. (We already know  $N$  can be taken to be  $c \log n$  by Theorem 6.1.)

By our definition of  $N$ , taking  $s = 12N/(\epsilon^2 \delta)$  suffices to obtain an  $(\epsilon/2, \delta/3)$  guarantee. Let  $S$  be the corresponding sampling matrix. Note that  $S$  takes  $s$  samples of expectation, and is unlikely to take many more. By a Chernoff bound, the probability that  $S$  takes at most  $2s$  samples is at least  $1 - \exp(-s/3)$ . Note that  $s \geq 1/\delta$ , so  $\exp(-s/3) \leq \delta$ , when  $\delta \leq 1/2$ . Also, note that  $SAS^T$  has entries bounded by  $n/s$  since the entries of  $A$  are assumed to be bounded by 1.

Now we consider sub-sampling again by a new sampling matrix  $S_2$ . By Theorem 6.1, if  $S_2$  is a uniform sampling matrix that takes  $\frac{c}{\delta \epsilon^2} \log(s/\epsilon)$  samples in expectation, then with failure probability at most  $\delta/3$ , the matrix  $S_2 SAS^T S_2^T$  yields an additive

$$(\epsilon/4)(2s) \|SAS^T\|_\infty = (\epsilon/2)s(n/s) = \epsilon n/2$$

approximation to the spectrum of  $SAS^T$ . Combining the bounds, we see that  $S_2 SAS^T S_2^T$  gives an additive  $\epsilon n$  spectral approximation to  $A$  with failure probability at most  $\delta$ .

Now, note that  $S_2 S$  is itself a sampling matrix that takes  $\frac{c}{\delta \epsilon^2} \log(s/\epsilon)$  samples in expectation and yields an  $(\epsilon, \delta)$  guarantee. By minimality of  $N$ , this means that

$$N \leq c \log \frac{s}{\epsilon}.$$

But  $s = \frac{N}{\epsilon^2 \delta}$ , so in fact

$$N \leq c \log \frac{N}{\epsilon^3 \delta},$$

which implies that  $N \leq c \log \frac{1}{\epsilon \delta}$ . Thus we conclude that a sampling matrix with  $s = \frac{c}{\epsilon^2 \delta} \log \frac{1}{\epsilon \delta}$  is sufficient to obtain an  $(\epsilon, \delta)$  guarantee.  $\square$

## 7 Squared row-norm Sampling

In this section we show how to modify the analysis of [Bha+24] to obtain improved bounds given access to a squared row-norm sampler.

We use the entry zeroing procedure of [Bha+24] (the same procedure used in Algorithm 2), which we now recall.

**DEFINITION 7.1.** *Given a symmetric matrix  $A \in \mathbb{R}^{n \times n}$  let  $A'$  be the matrix formed by zeroing out all entries  $A_{ij}$  satisfying one of the following conditions.*

1.  $i = j$  and  $\|A_i\|^2 < \frac{\epsilon^2}{4} \|A\|_F^2$
2.  $i \neq j$  and  $\|A_i\|^2 \|A_j\|^2 < \frac{\epsilon^2 \|A\|_F^2 |A_{ij}|^2}{c \log^4 n}$  for an absolute constant  $c$ .

We also use a lemma from [Bha+24] arguing that the spectrum of  $A'$  is close to the spectrum of  $A$ .

**LEMMA 7.1.** *For all  $i$ ,  $|\lambda_i(A') - \lambda_i(A)| \leq \epsilon \|A\|_F$ .*

We apply the following bound on the operator norm of the sampled matrix, which is given in the proof of Lemma 14 in [Bha+24].

**LEMMA 7.2.** *For  $s \geq \frac{\log n}{\epsilon^2}$ , we have the bound*

$$\mathbb{E}_2 \|SA'_m S\|_2 \leq 10\sqrt{\log n} \mathbb{E}_2 \left\| SH_m \hat{S} \right\|_{1 \rightarrow 2} + 15\epsilon \|A\|_F,$$

where  $H_m$  consists of the off-diagonal entries of  $A'_m$ ,  $\mathbb{E}_2(X) = \mathbb{E}(X^2)^{1/2}$  and where  $\|\cdot\|_{1 \rightarrow 2}$  denotes the maximum  $\ell_2$  norm of any column.

Our main task is to improve the bound on  $\mathbb{E}_2 \left\| SH_m \hat{S} \right\|_{1 \rightarrow 2}$ . To do this, we continue following the proof of [Bha+24], but improve on their variance calculation.

LEMMA 7.3. (From [Bha+24]) Assume that  $p_i \leq 1$  for all  $i$ . For a fixed  $i$ , define the random variable  $z_j$  to be  $\frac{1}{p_j} |A'_{m,i,j}|^2$  with probability  $p_j$ , and 0 otherwise. Then

$$\text{Var} \left( \sum_{j=1}^n z_j \right) \leq \sum_{j=1}^n |A'_{m,i,j}|^4 + \sum_{j=1}^n \frac{12 \|A\|_F^2}{s \|A_j\|_2^2} (|A'_{i,j}|^4 + |A'_{o,i,j}|^4)$$

LEMMA 7.4. Suppose that  $s \geq c \frac{\log^4 n}{\epsilon^2}$  for some absolute constant  $c$ . Fix  $i$  with  $p_i < 1$  and let  $z_j$  be as in the above lemma. We have

$$\text{Var} \left( \sum_{j=1}^n z_j \right) \leq 2 \|A_i\|_2^4.$$

*Proof.* From Lemma 7.3,

$$\begin{aligned} \text{Var} \left( \sum_{j=1}^n z_j \right) &\leq \sum_{j=1}^n |A'_{m,i,j}|^4 + \sum_{j=1}^n \frac{12 \|A\|_F^2}{s \|A_j\|_2^2} (|A'_{i,j}|^4 + |A'_{o,i,j}|^4) \\ &= \sum_{j=1}^n |A'_{m,i,j}|^4 + 12 \sum_{j=1}^n \frac{\|A\|_F^2}{s \|A_j\|_2^2} |A'_{i,j}|^4 + 12 \sum_{j=1}^n \frac{\|A\|_F^2}{s \|A_j\|_2^2} |A'_{o,i,j}|^4. \end{aligned}$$

Also from the argument of [Bha+24], we have the bound

$$|A'_{o,i,j}| \leq \frac{\|A_i\|_2 \|A_j\|_2}{\epsilon \|A\|_F}.$$

(Note that we take their  $\delta$  to be 1.)

Compared to their argument, our main improvement is that we use the bound

$$|A'_{o,i,j}|^4 \leq \frac{\|A_i\|_2^2 \|A_j\|_2^2}{\epsilon^2 \|A\|_F^2} |A'_{o,i,j}|^2.$$

This gives

$$\begin{aligned} \sum_{j=1}^n \frac{\|A\|_F^2}{s \|A_j\|_2^2} |A'_{o,i,j}|^4 &\leq \sum_{j=1}^n \frac{\|A\|_F^2}{s \|A_j\|_2^2} \left( \frac{\|A_i\|_2^2 \|A_j\|_2^2}{\epsilon^2 \|A\|_F^2} |A'_{o,i,j}|^2 \right) \\ &= \sum_{j=1}^n \frac{\|A_i\|_2^2}{s \epsilon^2} |A'_{o,i,j}|^2 \\ &= \frac{\|A_i\|_2^2}{s \epsilon^2} \|A'_{o,i}\|_2^2 \\ &\leq \frac{\|A_i\|_2^4}{s \epsilon^2}, \end{aligned}$$

where in the last line, we used that  $\|A'_{o,i}\|_2^2 \leq \|A'_i\|_2^2 \leq \|A_i\|_2^2$ .

Again following the proof of [Bha+24], we have

$$\sum_{j=1}^n |A'_{m,i,j}|^4 \leq \|A'_{m,i}\|_2^4 \leq \|A'_i\|_2^4 \leq \|A_i\|_2^4.$$

Continuing their proof, the thresholding procedure for  $A'$  implies that if  $i \neq j$  and  $A'_{ij} \neq 0$ , then

$$\frac{\|A_j\|_2^2}{|A'_{i,j}|^2} \geq \frac{\epsilon^2 \|A\|_F^2}{c \log^4 n \|A_i\|_2^2}.$$

Therefore,

$$\sum_{j=1}^n \frac{\|A\|_F^2}{s \|A_j\|_2^2} |A'_{i,j}|^4 = \sum_{j: A'_{ij} \neq 0} \frac{\|A\|_F^2}{s \|A_j\|_2^2} |A'_{i,j}|^4 \leq \sum_{j=1}^n \frac{c \log^4 n \|A_i\|_2^2}{s \epsilon^2} |A'_{ij}|^2 \leq \frac{c \log^4 n \|A_i\|_2^4}{s \epsilon^2}.$$

Combining the various bounds gives

$$\text{Var} \left( \sum_{j=1}^n z_j \right) \leq \|A_i\|_2^4 + \frac{c \log^4 n \|A_i\|_2^4}{s \epsilon^2}.$$

Thus for  $s \geq c \frac{\log^4 n}{\epsilon^2}$  we have

$$\text{Var} \left( \sum_{j=1}^n z_j \right) \leq 2 \|A_i\|_2^4,$$

as desired.  $\square$

With this improved variance bound, we obtain our improved middle eigenvalue bound by continuing the proof of Lemma 14 from [Bha+24].

LEMMA 7.5. *We have the operator norm bound  $\|SA'_m S^T\| \leq \epsilon \|A\|_F$  with 2/3 probability provided that  $s \geq c \frac{\log^4 n}{\epsilon^2}$  and  $p_i := \frac{s \|A_i\|_2}{\|A\|_F} \leq 1$  for all  $i$ .*

*Proof.* In their Lemma 14, [Bha+24] shows that  $|z_j| \leq 2 \|A_i\|_2^2$  for  $s \geq c \frac{\log^4 n}{\epsilon^2}$  if  $p_i < 1$ . Using our improved variance bound above, and applying Bernstein's inequality as in [Bha+24] gives

$$\Pr(\|(SA'_m)_{:,i}\|_2^2 \geq \mathbb{E} \|(SA'_m)_{:,i}\|_2^2 + t) \leq \left( \sum_{j=1}^n z_j \geq \|A_i\|_2^2 + t \right) \leq \exp \left( c \frac{-t^2}{\|A_i\|_2^4 + t \|A_i\|_2^2} \right).$$

Then setting  $t = c \log n \|A_i\|_2^2$  gives

$$\Pr(\|(SA'_m)_{:,i}\|_2^2 \geq \mathbb{E} \|(SA'_m)_{:,i}\|_2^2 + c \log n \|A_i\|_2^2) \leq 1/n^4,$$

for appropriate constants. It then follows that

$$\frac{1}{p_i} \|(SH_m)_{:,i}\|^2 \leq \frac{1}{p_i} \|(SA'_m)_{:,i}\|^2 \leq \frac{1}{p_i} (c \log n \|A_i\|_2^2) \leq \frac{\epsilon^2 \|A\|_F^2}{\log n},$$

for  $s \geq c \frac{\log^4 n}{\epsilon^2}$  with failure probability at most  $1/n^4$ . Thus with failure probability at most  $1/n^3$  this bound holds for all  $i$  simultaneously.

In order to apply this, we would instead like a second moment bound. Let  $E$  be the event that all  $p_i$  corresponding to sampled rows are bounded by  $1/n^2$ . Note that  $E$  occurs with probability at least  $1 - 1/n$ . From now on we condition on  $E$  with only a  $1/n$  loss in probability. Then we have

$$\frac{1}{\sqrt{p_i}} \|(SH_m)_{:,i}\| \leq n \|A\|_F$$

deterministically, and our bound above then implies that  $\mathbb{E}_2 \|SH_m \hat{S}\|_{1 \rightarrow 2}$  is bounded by  $\epsilon \|A\|_F$ .

Finally plugging into Lemma 7.2 we have

$$\mathbb{E}_2 \|SA'_m S\|_2 \leq c \epsilon \|A\|_F,$$

and the desired bound follows from Markov's inequality.  $\square$

---

**Algorithm 4** Restricted Row-norm Sampling Algorithm

---

1. Input: Symmetric matrix  $A \in \mathbb{R}^{n \times n}$  along with its row norms, sample size  $s$ . The inputs are restricted so that  $A$  and  $s$  must satisfy  $\frac{s\|A_i\|^2}{\|A\|^2} \leq 1$  for all  $i$ .
2. Let  $S \in \mathbb{R}^{k \times n}$  be a (rescaled) sampling matrix which samples each row  $i$  of  $A$  independently with probability  $\frac{s\|A_i\|^2}{\|A\|^2}$ .
3. Implicitly form the matrix  $A' \in \mathbb{R}^{n \times n}$  defined by
$$(A')_{ij} = \begin{cases} 0 & i = j \text{ and } \|A_i\|_2^2 \leq \frac{\epsilon^2}{4} \|A\|_F^2 \\ 0 & i \neq j \text{ and } \|A_i\|_2^2 \|A_j\|_2^2 \leq \frac{\epsilon^2 \|A\|_F^2 |A_{ij}|^2}{c \log^4 n} \\ A_{ij} & \text{otherwise} \end{cases}$$
4. Return the  $k$  eigenvalues of  $S^T A' S$ , along with  $n - k$  additional 0's.

---

This is sufficient to show correctness of the “restricted” sampling algorithm, Algorithm 4.

LEMMA 7.6. *Suppose that  $s \geq c \frac{\log^4 n}{\epsilon^2} \log^2 \frac{1}{\epsilon}$  and that  $A$  and  $s$  satisfy the condition of Algorithm 4 that the row sampling probabilities  $p_i$  are bounded by 1. Then with 2/3 probability, Algorithm 4 returns an additive  $\epsilon \|A\|_F$  approximation to the spectrum of  $A$ .*

*Proof.* As discussed above, write

$$S A' S^T = S A'_o S^T + S A'_m S^T.$$

By Lemma 4.2, Assumption 4.1 holds for  $A'_o$  and  $S$  with  $L = \frac{\epsilon}{\log \frac{1}{\epsilon}} \|A\|_F$ . As a result, by Theorem 5.1, the eigenvalues of  $S A'_o S^T$  are an additive  $c\epsilon \|A\|_F$  approximation to the eigenvalues of  $A'_o$ , which are in turn additive  $\epsilon \|A\|_F$  approximations to the eigenvalues of  $A_o$  by Lemma 7.1.

Also by Lemma 7.5 we have  $\|S A'_m S^T\| \leq c\epsilon \|A\|_F$ , so correctness follows from Weyl's inequality [Wey12].

□

Finally it remains to deal with the situation where the restricted condition of Algorithm 4 is not met. The idea is simple – for each row with  $p_i$  larger than 1, we simple split that row into several scaled-down rows. This will not affect the spectrum, but will allow the condition of Algorithm 4 to be met. In fact, for simplicity we split all rows in this way although this is not strictly necessary.

**THEOREM 7.1.** *Algorithm 2 yields an  $\epsilon \|A\|_F$  additive approximation to the spectrum of  $A$  with 2/3 probability when  $s \geq \frac{c \log^4 n}{\epsilon^2} \log^2 \frac{1}{\epsilon}$ .*

*Proof.* We consider an algorithm which will be equivalent to Algorithm 2. Set  $p_i = \frac{s\|A_i\|^2}{\|A\|_F^2}$ . Note that  $p_i \leq s$  for all  $i$ , so it will suffice to “split” each row into roughly  $s$  smaller rows. More formally, let  $U$  be a vertical stack of consisting of  $s$  copies of  $\frac{1}{\sqrt{s}} I_n$ . Note that  $U$  has orthonormal columns and that the image of  $U^T$  is  $\mathbb{R}^n$ . This means that the nonzero spectrum of  $U A U^T$  coincides with the nonzero spectrum of  $A$ . In particular we also have  $\|A\|_F = \|U A U^T\|_F$ . Note that the norm of row  $i$  of  $U A U^T$  is  $1/\sqrt{s^2}$  times the norm of a corresponding row  $i'$  of  $A$ .

Thus for all  $i$ ,

$$\frac{s \|(U A U^T)_i\|_2^2}{\|U A U^T\|_F^2} \leq \frac{1}{s} \frac{s \|A_{i'}\|_2^2}{\|A\|_F^2} = \frac{\|A_{i'}\|_2^2}{\|A\|_F^2} \leq 1.$$

Thus Lemma 7.6 establishes correctness of the sampling procedure for  $U A U^T$  (which has the same nonzero spectrum of  $A$ ). In order to achieve  $\epsilon \|U A U^T\|_F = \epsilon \|A\|_F$  additive error, we need  $s \geq \frac{c \log^4 (n s^2)}{\epsilon^2}$ , which is achieved for  $s \geq \frac{c \log^4 n}{\epsilon^2}$  (possibly after adjusting the constant  $c$ ).

Now to simulate Algorithm 4 on  $UAU^T$ , note that we can simply sample  $\text{Binomial}(s, \frac{\|A_i\|_2^2}{\|A\|_F^2})$  copies of each index  $i$ , as is done in Algorithm 2.

We finally note that for  $s \geq c/\epsilon^4$  that the diagonal zeroing condition for  $UAU^T$  in Algorithm 2 is always satisfied. Moreover the condition for zeroing out off-diagonal entries is unchanged since applying  $U$  to  $A$  scales row norms by  $1/\sqrt{s}$  but entries by  $1/s$ .  $\square$

## 8 Additional Applications

**8.1 Application to Sketching** [SW23] showed that a sketching dimension of  $\frac{1}{\epsilon^2}$  is sufficient to approximate all eigenvalues of  $A$  to within  $\epsilon\|A\|_F$  additive error. Unfortunately, the sketch consisted of a dense Gaussian matrix, which means that it requires roughly  $n^2/\epsilon^2$  time to apply to a dense matrix  $A$ <sup>1</sup>. As a consequence of our sampling results, we show that it is possible to match the optimal sketching dimension with a sketch that can be applied in time linear in the number of entries of  $A$ . We observe that conjugating by a Hadamard matrix is sufficient to flatten the row norms, so that our row-norm sampling procedure reduces to uniform sampling. The idea of applying Hadamard matrices to improve the runtimes of sketches is well-known in the literature (see for example [Woo+14; Tro11]). Importantly, if  $H$  is Hadamard, then a matrix-vector product  $Hv$  can be carried out in  $O(n \log n)$  time. Similarly if  $A \in \mathbb{R}^{n \times n}$  is a matrix, then  $HA$  can be carried out in  $O(n^2 \log n)$  by applying  $H$  columnwise.

**THEOREM 8.1.** *There is a bilinear sketch  $\mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{k \times k}$  with  $k = O(1/\epsilon^2)$  that can be applied in  $O(n^2)$  time, and allows all eigenvalues of  $A$  to be recovered to within  $\epsilon\|A\|_F$  additive error.*

*Proof.* Let  $U = HD \in \mathbb{R}^{n \times n}$  where  $D$  has i.i.d. random signs on its diagonal and  $H$  is a Hadamard matrix scaled by  $1/\sqrt{n}$ . Note that  $U$  is an orthogonal matrix, and so

$$\|(UAU^T)_i\|^2 = \|UAU^T e_i\|^2 = (U^T e_i)^T A^2 (U^T e_i).$$

Also note that  $U^T e_i = DHe_i$  which is distributed as a Rademacher random vector scaled by  $1/\sqrt{n}$ . Recalling that  $\mathbb{E}(x^T A^2 x) = \text{tr}(A^2) = \|A\|_F^2$ , when  $x$  is a random sign vector, by Hanson-Wright we have the bound

$$\Pr(|(UAU^T)_i - \frac{1}{n} \|A\|_F^2| \geq t) \leq 2 \exp(-c) \leq 2 \exp\left(-c \min\left(\frac{nt}{\|A^2\|_F}, \left(\frac{nt}{\|A^2\|_F}\right)^2\right)\right).$$

Setting  $t \geq c \frac{\|A\|_F^2}{n} \log \frac{n}{\delta}$  makes the the right hand side bounded by  $\frac{\delta}{n}$ , since  $\|A^2\|_F \leq \text{tr}(A^2) = \|A\|_F^2$ . By a union bound we have  $|(UAU^T)_i|^2 \leq c \frac{\|A\|_F^2}{n} \log \frac{n}{\delta}$  for all  $i$ , with probability at least  $1 - \delta$ .

It follows from Theorem 7.1 that it suffices to run Algorithm 2 by uniformly sampling each row with probability  $O(\frac{1}{n\epsilon^2} \text{poly log } \frac{n}{\epsilon})$ , resulting in a submatrix of dimension  $k \times k$ , where  $k = O(\frac{1}{\epsilon^2} \text{poly log } \frac{n}{\epsilon})$  with high probability. To implement this as a sketch, we can choose a uniformly random permutation matrix  $P$ , (implicitly) form  $(PHD)A(PHD)^T$  and then sample a leading principal minor of dimension  $O(\frac{1}{\epsilon^2} \text{poly log } \frac{n}{\epsilon})$ .

Finally, in parallel we can sketch the norms of the first  $1/\epsilon^2$  columns of  $(PHD)A(PHD)^T$  to within a constant factor using a Johnson-Lindenstrauss sketch (see [Woo+14] for example). This only takes an additional  $\frac{1}{\epsilon^2} \log \frac{1}{\epsilon\delta}$  space to succeed with  $1 - \delta$  probability on all of the first  $1/\epsilon^2$  rows, and can applied in  $O(\frac{n}{\epsilon^2} \text{poly log } \frac{n}{\epsilon})$  time. This allows us to run Algorithm 2 on the resulting matrix.  $\square$

**8.2 Improvements by combining with adaptive results** By combining with known adaptive sampling bounds of [MM17], we point out that one can approximate the eigenvalues of a bounded entry PSD  $A$  to within  $\epsilon n$  additive error using just  $\tilde{O}(1/\epsilon^3)$  entry queries. We note that the PSD sampling result that we use was already present in [Bha+24], although this improvement using adaptivity was not pointed out.

We borrow a result from [MM17], slightly restated for our setting.

---

<sup>1</sup>Some speedup is possible by using fast matrix multiplication, although these algorithms are impractical for reasonably sized inputs.

**THEOREM 8.2.** ([MM17] Theorem 3.) *Let  $A \in \mathbb{R}^{n \times n}$  be PSD and  $\lambda > 0$ . There is an algorithm that adaptively queries  $O(nd_\lambda \log d_\lambda)$  entries of  $A$  and with  $2/3$  probability produces a spectral approximation  $\hat{A}$  of  $A$  satisfying*

$$\|A - \hat{A}\| \leq \lambda.$$

Here  $d_\lambda := \text{tr}(A(A + \lambda I)^{-1})$  is the  $\lambda$ -effective dimension of  $A$ .

We observe that this result improves sampling guarantees using adaptivity.

**THEOREM 8.3.** *Let  $A$  be PSD with  $\|A\|_\infty \leq 1$ . Then there is an algorithm that adaptively queries  $\tilde{O}(1/\epsilon^3)$  entries of  $A$  and with  $2/3$  probability produces an additive  $\epsilon n$  approximation to the spectrum of  $A$ .*

*Proof.* For the first step we apply our Theorem 6.1 (or [Bha+24] since  $A$  is PSD), to (implicitly) produce a rescaled submatrix  $S^T AS$  of dimension  $\tilde{O}(1/\epsilon^2)$  whose eigenvalues contain additive  $\epsilon n$  approximations to all eigenvalues of  $A$  that are at least  $\epsilon n$ . Note that we do not actually form the matrix yet, however we have query access to its entries simply by querying entries of  $A$ . By our Theorem 7.1, outputting the spectrum of  $S^T AS$  along with additional zeros would suffice.

To achieve an  $O(\epsilon n)$  additive error overall, it suffices to approximate the eigenvalues of  $S^T AS$  to within  $O(\epsilon n)$  additive error. For this we apply Theorem 8.3 to obtain  $\hat{A}$  with  $\|A - \hat{A}\| \leq \epsilon n$ . Note that we have

$$d_\lambda = \text{tr}(A(A + \epsilon n I)^{-1}) \leq \sum_{i=1}^n \frac{\lambda_i}{\epsilon n} = \frac{\text{tr}(A)}{\epsilon n} \leq \frac{1}{\epsilon}.$$

Therefore only  $\tilde{O}(1/\epsilon^3)$  queries are needed to produce  $\hat{A}$ . Finally, by Weyl's inequality [Wey12],  $|\lambda_i(A) - \lambda_i(\hat{A})| \leq \epsilon n$  for all  $i$ . The result follows by replacing  $\epsilon$  with  $\epsilon/c$  for an appropriate constant.  $\square$

## 9 Top Eigenvector Estimation

In this section, we analyze Algorithm 3 for producing an approximate top eigenvector for PSD  $A$  with  $\|A\|_\infty \leq 1$ . To prove correctness, we begin with a few simple facts that we will need below.

**LEMMA 9.1.** *Suppose that  $A \in \mathbb{R}^{n \times n}$  is PSD with entries bounded by 1, and let  $v$  be an eigenvector of  $A$  with associated eigenvalue  $\lambda$ . Then  $\|v\|_\infty \leq \frac{1}{\sqrt{\lambda}}$*

*Proof.* Write the spectral decomposition of  $A$  as

$$A = \sum_{i=1}^n \lambda_i v_i v_i^T,$$

where each  $v_i$  is an eigenvector of  $A$  with associated eigenvalue  $\lambda_i$ . For arbitrary  $j$ , we then have

$$A_{jj} = \sum_{i=1}^n \lambda_i v_{ij}^2.$$

By the bounded entry hypothesis, this is bounded by 1, so in particular  $\lambda_1 v_{1j}^2 \leq 1$  since  $A$  is PSD. Since  $j$  was arbitrary, the claim follows.  $\square$

A related fact is that zeroing out eigenvalues of  $A$  only decreases the diagonal entries. We only need this fact for the top eigenvalue, but it is true more generally.

**LEMMA 9.2.** *Let  $A_{-1}$  denote  $A$  with its top eigenvalue zeroed out. Then  $(A_{-1})_{kk} \leq A_{kk}$  for all  $k$*

*Proof.* Write the spectral decomposition of  $A$  as

$$A = \sum_{i=1}^n \lambda_i v_i v_i^T.$$

Now simply note that

$$\lambda_1(v_1 v_1^T)_{kk} = \lambda_1 v_{1k}^2 \geq 0.$$

$\square$

The following fact states that applying a half iteration of power method only increases the Rayleigh quotient. (This statement is true more general, but we only need the 1/2 version.)

LEMMA 9.3. *Let  $A \in \mathbb{R}^{n \times n}$  be PSD. Then for all nonzero  $x \in \mathbb{R}^n$ ,*

$$\frac{x^T A^2 x}{x^T A x} \geq \frac{x^T A x}{x^T x}.$$

*Proof.* The desired inequality is equivalent to

$$\langle Ax, x \rangle^2 \leq \|Ax\|_2^2 \|x\|_2^2,$$

which is true by Cauchy-Schwarz.  $\square$

The following is our main lemma for top eigenvector approximation.

LEMMA 9.4. *Let  $S$  be a column-sampling matrix that samples each column of  $A$  independently with probability  $p = \frac{c}{\epsilon n}$  for an absolute constant  $c$ . Then with probability at least  $2/3$ , there is some  $x$  in the image of  $S$ , such that*

$$\frac{x^T A^2 x}{x^T A x} \geq \lambda_1 - \epsilon n.$$

*Proof.* If  $\lambda_1 < \epsilon n$  then the result is trivial. So we assume from now on that  $\lambda_1 \geq \epsilon n$ .

First, for an arbitrary fixed  $x \in \mathbb{R}^n$ , consider the generalized Rayleigh quotient  $\frac{x^T A^2 x}{x^T A x}$ . Let  $v^{(1)}, \dots, v^{(n)}$  be an orthonormal basis of eigenvectors of  $A$  with associated eigenvalues  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ . Then we have

$$\begin{aligned} \frac{x^T A^2 x}{x^T A x} &= \frac{\lambda_1^2 \langle x, v^{(1)} \rangle^2 + \lambda_2^2 \langle x, v^{(2)} \rangle^2 + \dots + \lambda_n^2 \langle x, v^{(n)} \rangle^2}{\lambda_1 \langle x, v^{(1)} \rangle^2 + \lambda_2 \langle x, v^{(2)} \rangle^2 + \dots + \lambda_n \langle x, v^{(n)} \rangle^2} \\ &\geq \frac{\lambda_1^2 \langle x, v^{(1)} \rangle^2}{\lambda_1 \langle x, v^{(1)} \rangle^2 + \lambda_2 \langle x, v^{(2)} \rangle^2 + \dots + \lambda_n \langle x, v^{(n)} \rangle^2} \\ &= \frac{\lambda_1^2}{\lambda_1 + \frac{1}{\langle x, v^{(1)} \rangle^2} (\lambda_2 \langle x, v^{(2)} \rangle^2 + \dots + \lambda_n \langle x, v^{(n)} \rangle^2)} \\ &\geq \lambda_1 - \frac{1}{\langle x, v^{(1)} \rangle^2} (\lambda_2 \langle x, v^{(2)} \rangle^2 + \dots + \lambda_n \langle x, v^{(n)} \rangle^2), \end{aligned}$$

where the last inequality follows from the difference-of-squares factorization. Now write  $A_{-1}$  to denote  $A$  with its top eigenvalue zeroed out. The term in parentheses above is  $x^T A_{-1} x$ , so we have

$$\frac{x^T A^2 x}{x^T A x} \geq \lambda_1 - \frac{x^T A_{-1} x}{\langle x, v \rangle^2}.$$

It therefore suffices to show that  $\frac{x^T A_{-1} x}{\langle x, v \rangle^2} \leq \epsilon n$  for some  $x$  in the image of  $S$ . We will take  $x = \Pi_S v$  where  $v = v^{(1)}$  is the top eigenvector of  $A$  and where  $\Pi$  is the orthogonal projection onto the image of  $S$ . We assume that  $v$  is normalized to have unit norm. For this choice of  $x$  we will bound the numerator and denominator separately.

**Denominator bound.** First note that

$$\langle x, v \rangle^2 = \langle \Pi v, v \rangle^2 = \langle \Pi v, \Pi v \rangle^2 = \|\Pi v\|_2^4,$$

so we simply need to bound  $\|\Pi v\|_2^2$ . We will prove the denominator bound under the assumption that  $p = \frac{10}{\epsilon n}$ . Clearly this is sufficient as the denominator for larger  $p$  majorizes the denominator for smaller  $p$ . By Lemma 9.1,

recall that  $|v_i|^2 \leq \frac{1}{\lambda_1} \leq \frac{1}{\epsilon n}$  for all  $i$ . Let  $\sigma_i$  be i.i.d. Bernoulli( $p$ ) random variables. We wish to obtain a lower bound on  $X := \sum_i \sigma_i v_i^2$ . Since  $v$  is a unit vector, by linearity of expectation  $\mathbb{E}X = p$ . For the second moment,

$$\begin{aligned}\mathbb{E}X^2 &= p(v_1^4 + \dots + v_n^4) + 2p^2 \sum_{i < j} v_i^2 v_j^2 \\ &= p(v_1^4 + \dots + v_n^4) + p^2(v_1^2 + \dots + v_n^2)^2 - p^2(v_1^4 + \dots + v_n^4) \\ &\leq p(v_1^4 + \dots + v_n^4) + p^2 \\ &\leq \frac{p}{\epsilon n}(v_1^2 + \dots + v_n^2) + p^2 \\ &= \frac{p}{\epsilon n} + p^2.\end{aligned}$$

Therefore  $\text{Var}(X) \leq \frac{p}{\epsilon n}$ , and so by Chebyshev's inequality,

$$\Pr(X \leq \frac{1}{\epsilon n}) \leq \Pr(|X - p| \geq \frac{9}{\epsilon n}) \leq \frac{\text{Var}(X)}{(9/(\epsilon n))^2} \leq \frac{\epsilon n p}{81} = \frac{10}{81}.$$

Thus we have  $\langle x, v \rangle^2 \geq \frac{1}{(\epsilon n)^2}$  with probability at least 7/8.

**Numerator bound.** Since  $A_{-1}$  is PSD, we can write  $A_{-1} = U^T U$  for some  $U \in \mathbb{R}^{n \times n}$ . By Lemma 9.2, the diagonal entries of  $A_{-1}$  are all bounded by 1, which means that the columns of  $U$  each have  $\ell_2$  norm at most 1. Also, since  $v$  is the top eigenvector of  $A$ , we have  $v^T A_{-1} v = 0$ , which implies that  $Uv = 0$ .

Let  $w^{(i)} = v_i U_i$  where  $v_i$  is the  $i$ th entry of  $v$  and  $U_i$  is the  $i$ th column of  $U$ . Also set  $p = \frac{1}{\epsilon n}$ .

Note that we have

$$\sum_{i=1}^n w^{(i)} = 0,$$

and

$$\sum_{i=1}^n \|w^{(i)}\|_2^2 \leq 1.$$

We wish to bound

$$\sum_{i=1}^n \|\sigma_i w^{(i)}\|^2,$$

where the  $\sigma_i$  are i.i.d. Bernoulli( $p$ ). Since the quantity we wish to bound is non-negative, it suffices to bound its expectation:

$$\begin{aligned}\mathbb{E} \left\| \sum_i \sigma_i w^{(i)} \right\|_2^2 &= p \sum_{i=1}^n \|w^{(i)}\|^2 + 2p^2 \sum_{i < j} \langle w^{(i)}, w^{(j)} \rangle \\ &\leq p + 2p^2 \sum_{i < j} \langle w^{(i)}, w^{(j)} \rangle.\end{aligned}$$

This last sum can be rewritten as

$$2 \sum_{i < j} \langle w^{(i)}, w^{(j)} \rangle = \|w^{(1)} + \dots + w^{(n)}\|_2^2 - \left( \|w^{(1)}\|_2^2 + \dots + \|w^{(n)}\|_2^2 \right) \leq 0,$$

since  $w^{(1)} + \dots + w^{(n)} = 0$ .

It follows that

$$\mathbb{E} \left\| \sum_i \sigma_i w^{(i)} \right\|_2^2 \leq p,$$

so by Markov's inequality

$$\left\| \sum_i \sigma_i w^{(i)} \right\|_2^2 \leq 10p,$$

with probability at least 9/10.

Combining the bound on the numerator and the denominator shows that  $\frac{x^T A_{-1} x}{x^T A x} \geq \lambda_1 - 100\epsilon n$ . The result follows by replacing  $\epsilon$  with  $\epsilon/100$ .  $\square$

Finally we prove correctness of Algorithm 3.

**THEOREM 9.1.** *Let  $A \in \mathbb{R}^{n \times n}$  be symmetric PSD with  $\|A\|_\infty \leq 1$ . For  $p = \frac{c}{\epsilon n}$ , with 3/4 probability, Algorithm 3 returns a unit vector  $u \in \mathbb{R}^n$  satisfying*

$$u^T A u \geq \lambda_1 - \epsilon n,$$

where  $\lambda_1$  is the top eigenvalue of  $A$ .

*Proof.* Consider the  $x \in \mathbb{R}^n$  produced in Algorithm 3 which maximizes  $\frac{x^T S^T A^2 S x}{x^T S^T A S x}$ . By Lemma 9.4, with 3/4 probability we have that

$$\frac{x^T S^T A^2 S x}{x^T S^T A S x} = \frac{(A^{1/2} S x)^T A (A^{1/2} S x)}{(A^{1/2} S x)^T (A^{1/2} S x)} \geq \lambda_1 - \epsilon n.$$

Now from Lemma 9.3, an additional half-iteration of power method only helps. In other words,

$$\frac{(A S x)^T A (A S x)}{(A S x)^T (A S x)} \geq \frac{(A^{1/2} S x)^T A (A^{1/2} S x)}{(A^{1/2} S x)^T (A^{1/2} S x)} \geq \lambda_1 - \epsilon n,$$

and the main result follows.  $\square$

## References

- [AN13] Alexandr Andoni and Huy L Nguyen. “Eigenvalues of a matrix in the streaming model”. In: *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2013, pp. 1729–1737.
- [Bal+19] Maria-Florina Balcan et al. “Testing matrix rank, optimally”. In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2019, pp. 727–746.
- [BCJ20] Ainesh Bakshi, Nadiia Chepurko, and Rajesh Jayaram. “Testing positive semi-definiteness via random submatrices”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 1191–1202.
- [BCW20] Ainesh Bakshi, Nadiia Chepurko, and David P Woodruff. “Robust and sample optimal algorithms for PSD low rank approximation”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 506–516.
- [Bha+23] Rajarshi Bhattacharjee et al. “Universal matrix sparsifiers and fast deterministic algorithms for linear algebra”. In: *arXiv preprint arXiv:2305.05826* (2023).
- [Bha+24] Rajarshi Bhattacharjee et al. “Sublinear time eigenvalue approximation via random sampling”. In: *Algorithmica* (2024), pp. 1–66.
- [BKM22] Vladimir Braverman, Aditya Krishnan, and Christopher Musco. “Linear and sublinear time spectral density estimation”. In: *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 2022.
- [BW18] Ainesh Bakshi and David Woodruff. “Sublinear time low-rank approximation of distance matrices”. In: *Advances in Neural Information Processing Systems* 31 (2018).
- [Cha+20] Moses Charikar et al. “Kernel density estimation through density constrained near neighbor search”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 172–183.
- [Che+22] Nadiia Chepurko et al. “Quantum-inspired algorithms from randomized numerical linear algebra”. In: *International Conference on Machine Learning*. PMLR. 2022, pp. 3879–3900.

[CNW15] Michael B Cohen, Jelani Nelson, and David P Woodruff. “Optimal approximate matrix product in terms of stable rank”. In: *arXiv preprint arXiv:1507.02268* (2015).

[Coh+18] David Cohen-Steiner et al. “Approximating the spectrum of a graph”. In: *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2018, pp. 1263–1271.

[CS17] Moses Charikar and Paris Siminelakis. “Hashing-based-estimators for kernel density in high dimensions”. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2017, pp. 1032–1043.

[DZ11] Petros Drineas and Anastasios Zouzias. “A note on element-wise matrix sparsification via a matrix-valued Bernstein inequality”. In: *Information Processing Letters* 111.8 (2011), pp. 385–389.

[Fra62] John GF Francis. “The QR transformation—part 2”. In: *The Computer Journal* 4.4 (1962), pp. 332–345.

[GLT18] András Gilyén, Seth Lloyd, and Ewin Tang. “Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension”. In: *arXiv preprint arXiv:1811.04909* (2018).

[GST22] András Gilyén, Zhao Song, and Ewin Tang. “An improved quantum-inspired algorithm for linear regression”. In: *Quantum* 6 (2022), p. 754.

[GV00] Gene H Golub and Henk A Van der Vorst. “Eigenvalue computation in the 20th century”. In: *Journal of Computational and Applied Mathematics* 123.1-2 (2000), pp. 35–65.

[Kyn18] Rasmus Kyng. “A tutorial on matrix approximation by row sampling”. In: (2018).

[LW16] Yi Li and David P Woodruff. “Tight bounds for sketching the operator norm, schatten norms, and subspace embeddings”. In: *International Conference on Approximation Algorithms for Combinatorial Optimization Problems and International Conference on Randomization and Computation*. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing. 2016.

[MM15] Cameron Musco and Christopher Musco. “Randomized block krylov methods for stronger and faster approximate singular value decomposition”. In: *Advances in neural information processing systems* 28 (2015).

[MM17] Cameron Musco and Christopher Musco. “Recursive sampling for the nyström method”. In: *Advances in neural information processing systems* 30 (2017).

[MW17] Cameron Musco and David P Woodruff. “Sublinear time low-rank approximation of positive semidefinite matrices”. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2017, pp. 672–683.

[RST10] Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. “A randomized algorithm for principal component analysis”. In: *SIAM Journal on Matrix Analysis and Applications* 31.3 (2010), pp. 1100–1124.

[RV07] Mark Rudelson and Roman Vershynin. “Sampling from large matrices: An approach through geometric functional analysis”. In: *Journal of the ACM (JACM)* 54.4 (2007), 21–es.

[Sim+19] Paris Siminelakis et al. “Rehashing kernel evaluation in high dimensions”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 5789–5798.

[SW23] William Swartworth and David P Woodruff. “Optimal Eigenvalue Approximation via Sketching”. In: *arXiv preprint arXiv:2304.09281* (2023).

[Tan19] Ewin Tang. “A quantum-inspired classical algorithm for recommendation systems”. In: *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*. 2019, pp. 217–228.

[Tro08] Joel A Tropp. “Norms of random submatrices and sparse approximation”. In: *Comptes Rendus. Mathématique* 346.23-24 (2008), pp. 1271–1274.

[Tro11] Joel A Tropp. “Improved analysis of the subsampled randomized Hadamard transform”. In: *Advances in Adaptive Data Analysis* 3.01n02 (2011), pp. 115–126.

[Tro12] Joel A Tropp. “User-friendly tail bounds for sums of random matrices”. In: *Foundations of computational mathematics* 12 (2012), pp. 389–434.

- [Wey12] Hermann Weyl. “The asymptotic distribution law for the eigenvalues of linear partial differential equations (with applications to the theory of black body radiation)”. In: *Math. Ann* 71.1 (1912), pp. 441–479.
- [Woo+14] David P Woodruff et al. “Sketching as a tool for numerical linear algebra”. In: *Foundations and Trends® in Theoretical Computer Science* 10.1–2 (2014), pp. 1–157.