



Sublinear Time Eigenvalue Approximation via Random Sampling

Rajarshi Bhattacharjee¹ · Gregory Dexter² · Petros Drineas² ·
Cameron Musco¹ · Archan Ray¹

Received: 26 October 2023 / Accepted: 12 January 2024 / Published online: 12 February 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

We study the problem of approximating the eigenspectrum of a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with bounded entries (i.e., $\|\mathbf{A}\|_{\infty} \leq 1$). We present a simple sublinear time algorithm that approximates all eigenvalues of $\mathbf{A}$ up to additive error $\pm \epsilon n$ using those of a randomly sampled $\tilde{O}\left(\frac{\log^3 n}{\epsilon^3}\right) \times \tilde{O}\left(\frac{\log^3 n}{\epsilon^3}\right)$ principal submatrix. Our result can be viewed as a concentration bound on the complete eigenspectrum of a random submatrix, significantly extending known bounds on just the singular values (the magnitudes of the eigenvalues). We give improved error bounds of $\pm \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ and $\pm \epsilon \|\mathbf{A}\|_F$ when the rows of $\mathbf{A}$ can be sampled with probabilities proportional to their sparsities or their squared ℓ_2 norms respectively. Here $\text{nnz}(\mathbf{A})$ is the number of non-zero entries in $\mathbf{A}$ and $\|\mathbf{A}\|_F$ is its Frobenius norm. Even for the strictly easier problems of approximating the singular values or testing the existence of large negative eigenvalues (Bakshi, Chepurko, and Jayaram, FOCS '20), our results are the first that take advantage of non-uniform sampling to give improved error bounds. From a technical perspective, our results require several new eigenvalue concentration and

An extended abstract of this work appeared in ICALP 2023.

✉ Gregory Dexter
gdexter@purdue.edu
Rajarshi Bhattacharjee
rbhattacharj@cs.umass.edu
Petros Drineas
pdrineas@purdue.edu
Cameron Musco
cmusco@cs.umass.edu
Archan Ray
ray@cs.umass.edu

¹ Manning College of Information and Computer Sciences, University of Massachusetts, Amherst, MA, USA

² Department of Computer Science, Purdue University, West Lafayette, IN, USA

perturbation bounds for matrices with bounded entries. Our non-uniform sampling bounds require a new algorithmic approach, which judiciously zeroes out entries of a randomly sampled submatrix to reduce variance, before computing the eigenvalues of that submatrix as estimates for those of $\mathbf{A}$. We complement our theoretical results with numerical simulations, which demonstrate the effectiveness of our algorithms in practice.

Keywords Eigenvalue approximation · Sublinear algorithms · Randomized algorithms · Symmetric matrices

1 Introduction

Approximating the eigenvalues of a symmetric matrix is a fundamental problem—with applications in engineering, optimization, data analysis, spectral graph theory, and beyond. For an $n \times n$ matrix, all eigenvalues can be computed to high accuracy using direct eigendecomposition in $O(n^\omega)$ time, where $\omega \approx 2.37$ is the exponent of matrix multiplication [1, 2]. When just a few of the largest magnitude eigenvalues are of interest, the power method and other iterative Krylov methods can be applied [3]. These methods repeatedly multiply the matrix of interest by query vectors, requiring $O(n^2)$ time per multiplication when the matrix is dense and unstructured.

For large n , it is desirable to have even faster eigenvalue approximation algorithms, running in $o(n^2)$ time—i.e., sublinear in the size of the input matrix. Unfortunately, for general matrices, no non-trivial approximation can be computed in $o(n^2)$ time: without reading $\Omega(n^2)$ entries, it is impossible to distinguish with reasonable probability if all entries (and hence all eigenvalues) are equal to zero, or if there is a single pair of arbitrarily large entries at positions (i, j) and (j, i) , leading to a pair of arbitrarily large eigenvalues. Given this, we seek to address the following question:

Under what assumptions on a symmetric $n \times n$ input matrix, can we compute non-trivial approximations to its eigenvalues in $o(n^2)$ time?

It is well known that $o(n^2)$ time eigenvalue computation is possible for highly structured inputs, like tridiagonal or Toeplitz matrices [4]. For sparse or structured matrices that admit fast matrix vector multiplication, one can compute a small number of the largest in magnitude eigenvalues in $o(n^2)$ time using iterative methods. Through the use of robust iterative methods, fast top eigenvalue estimation is also possible for matrices that admit fast *approximate* matrix–vector multiplication, such as kernel similarity matrices [5–7]. Our goal is to study simple, sampling-based sublinear time algorithms that work under much weaker assumptions on the input matrix.

1.1 Our Contributions

Our main contribution is to show that a very simple algorithm can be used to approximate *all eigenvalues* of any symmetric matrix with *bounded entries*. In particular, for any $\mathbf{A} \in \mathbb{R}^{n \times n}$ with maximum entry magnitude $\|\mathbf{A}\|_\infty \leq 1$, sampling an $s \times s$ princi-

pal submatrix $\mathbf{A}_S$ of $\mathbf{A}$ with $s = \tilde{O}\left(\frac{\log^3 n}{\epsilon^3}\right)$ and scaling its eigenvalues by n/s yields a $\pm\epsilon n$ additive error approximation to all eigenvalues of $\mathbf{A}$ with good probability.¹ This result is formally stated below, where $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$.

Theorem 1 (Sublinear Time Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including each index independently with probability s/n as in Algorithm 1. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, with eigenvalues $\lambda_1(\mathbf{A}_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}_S)$.*

For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all other $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$. If $s \geq \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \delta}$, for large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$,

$$\lambda_i(\mathbf{A}) - \epsilon n \leq \tilde{\lambda}_i(\mathbf{A}) \leq \lambda_i(\mathbf{A}) + \epsilon n.$$

See Fig. 1 for an illustration of how the $|S|$ eigenvalues of $\mathbf{A}_S$ are mapped to estimates for all n eigenvalues of $\mathbf{A}$. Since $\mathbf{A}$ satisfies $\|\mathbf{A}\|_\infty \leq 1$, it must hold that $\|\mathbf{A}\|_F^2 = \sum_{i=1}^n \lambda_i^2(\mathbf{A}) \leq n^2$, so there are at most $1/\epsilon^2$ eigenvalues greater than ϵn . Therefore, we may set $\tilde{\lambda}_i(\mathbf{A})$ to zero for all but $1/\epsilon^2$ eigenvalues and still satisfy the stated error guarantee. Note that the principal submatrix $\mathbf{A}_S$ sampled in Theorem 1 will have $O(s) = \tilde{O}\left(\frac{\log^3 n}{\epsilon^3 \delta}\right)$ rows/columns with high probability. Thus, with high probability, the algorithm reads just $\tilde{O}\left(\frac{\log^6 n}{\epsilon^6 \delta^2}\right)$ entries of $\mathbf{A}$ and runs in $\text{poly}(\log n, 1/\epsilon, 1/\delta)$ time. Standard matrix concentration bounds imply that one can sample $O\left(\frac{s \log(1/\delta)}{\epsilon^2}\right)$ random entries from the $O(s) \times O(s)$ random submatrix $\mathbf{A}_S$ and preserve its eigenvalues to error $\pm\epsilon s$ with probability $1 - \delta$ [8]. See “Appendix F” for a proof. This can be directly combined with Theorem 1 to give improved sample complexity:

Corollary 1 (Improved Sample Complexity via Entrywise Sampling) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. For any $\epsilon, \delta \in (0, 1)$, there is an algorithm that reads $\tilde{O}\left(\frac{\log^3 n}{\epsilon^3 \delta}\right)$ entries of $\mathbf{A}$ and returns, with probability at least $1 - \delta$, $\tilde{\lambda}_i(\mathbf{A})$ for each $i \in [n]$ satisfying $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \epsilon n$.*

Observe that the dependence on δ in Theorem 1 and Corollary 1 can be improved via standard arguments: running the algorithm with failure probability $\delta' = 2/3$, repeating $O(\log(1/\delta))$ times, and taking the median estimate for each $\lambda_i(\mathbf{A})$. This guarantees that the algorithm will succeed with probability at most $1 - \delta$ at the expense of a $\log(1/\delta)$ dependence in the complexity.

Comparison to known bounds Theorem 1 can be viewed as a concentration inequality on the full eigenspectrum of a random principal submatrix $\mathbf{A}_S$ of $\mathbf{A}$. This significantly

¹ Here and throughout, $\tilde{O}(\cdot)$ hides logarithmic factors in the argument. Note that by scaling, our algorithm gives a $\pm\epsilon n \cdot \|\mathbf{A}\|_\infty$ approximation for any $\mathbf{A}$.




Fig. 1 Alignment of eigenvalues in Theorem 1 and Algorithm 1. We illustrate how the eigenvalues of $\mathbf{A}_S$, scaled by $\frac{n}{s}$, are used to approximate all eigenvalues of $\mathbf{A}$. If $\mathbf{A}_S$ has $p - 1$ positive eigenvalues, they are set to the top $p - 1$ eigenvalue estimates. Its $|S| - p + 1$ negative eigenvalues are set to the bottom eigenvalue estimates. All remaining eigenvalues are simply approximated as zero

extends prior work, which was able to bound just the spectral norm (i.e., the magnitude of the top eigenvalue) of a random principal submatrix [9, 10]. Bakshi et al. [11] recently identified developing such full eigenspectrum concentration inequalities as an important step in expanding our knowledge of sublinear time property testing algorithms for bounded entry matrices.

Standard matrix concentration bounds [12] can be used to show that the *singular values* of $\mathbf{A}$ (i.e., the magnitudes of its eigenvalues) are approximated by those of a $O\left(\frac{\log n}{\epsilon^2}\right) \times O\left(\frac{\log n}{\epsilon^2}\right)$ random submatrix (see “Appendix G”) with independently sampled rows and columns. However, such a random matrix will not be symmetric or even have real eigenvalues in general as the sampling of the rows is independent of the columns, and thus no analogous bounds were previously known for the eigenvalues themselves.

Recently, Bakshi et al. [11] studied the closely related problem of testing positive semidefiniteness in the bounded entry model. They show how to test whether the minimum eigenvalue of $\mathbf{A}$ is either greater than 0 or smaller than $-\epsilon n$ by reading just $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$ entries. They show that this result is optimal in terms of query complexity, up to logarithmic factors. Like our approach, their algorithm is based on random principal submatrix sampling. Our eigenvalue approximation guarantee strictly strengthens the testing guarantee—given $\pm \epsilon n$ approximations to all eigenvalues, we immediately solve the testing problem. Thus, our query complexity is tight up to a $\text{poly}(\log n, 1/\epsilon)$ factor. It is open if our higher sample complexity is necessary to solve the harder full eigenspectrum estimation problem. See Sect. 1.4 for further discussion.

Improved bounds for non-uniform sampling Our second main contribution is to show that, when it is possible to efficiently sample rows/columns of $\mathbf{A}$ with probabilities proportional to their sparsities or their squared ℓ_2 norms, significantly stronger eigenvalue estimates can be obtained. In particular, letting $\text{nnz}(\mathbf{A})$ denote the number of nonzero entries in $\mathbf{A}$ and $\|\mathbf{A}\|_F$ denote its Frobenius norm, we show that sparsity-based sampling yields eigenvalue estimates with error $\pm \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ and norm-based sampling gives error $\pm \epsilon \|\mathbf{A}\|_F$. See Theorems 2 and 3 for formal statements. Observe that when $\|\mathbf{A}\|_\infty \leq 1$, its eigenvalues are bounded in magnitude by $\|\mathbf{A}\|_2 \leq \|\mathbf{A}\|_F \leq \sqrt{\text{nnz}(\mathbf{A})} \leq n$. Thus, Theorems 2 and 3 are natural strengthen-

ings of Theorem 1. Row norm-based sampling (Theorem 3) additionally removes the bounded entry requirement of Theorems 1 and 2.

As discussed in Sect. 1.3.1, sparsity-based sampling can be performed in sublinear time when $\mathbf{A}$ is stored in a slightly augmented sparse matrix format or when $\mathbf{A}$ is the adjacency matrix of a graph accessed in the standard graph query model of the sublinear algorithms literature [13]. Norm-based sampling can also be performed efficiently with an augmented matrix format, and is commonly studied in randomized and ‘quantum-inspired’ algorithms for linear algebra [14, 15].

Theorem 2 (Sparse Matrix Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including the i^{th} index independently with probability $p_i = \min\left(1, \frac{s \cdot \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}\right)$ as in Algorithm 2. Here $\text{nnz}(\mathbf{A}_i)$ is the number of non-zero entries in the i^{th} row of $\mathbf{A}$. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, and let $\tilde{\lambda}_i(\mathbf{A})$ be the estimate of $\lambda_i(\mathbf{A})$ computed from $\mathbf{A}_S$ as in Algorithm 2. If $s \geq \frac{c \log^8 n}{\epsilon^8 \delta^4}$, for large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$, $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$.*

Theorem 3 (Row Norm Based Matrix Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including the i^{th} index independently with probability $p_i = \min\left(1, \frac{s \|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{n^2}\right)$ as in Algorithm 3. Here $\|\mathbf{A}_i\|_2$ is the ℓ_2 norm of the i^{th} row of $\mathbf{A}$. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, and let $\tilde{\lambda}_i(\mathbf{A})$ be the estimate of $\lambda_i(\mathbf{A})$ computed from $\mathbf{A}_S$ as in Algorithm 3. If $s \geq \frac{c \log^{10} n}{\epsilon^8 \delta^4}$, for large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$, $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \epsilon \|\mathbf{A}\|_F$.*

The above non-uniform sampling theorems immediately yield algorithms for testing the presence of a negative eigenvalue with magnitude at least $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$ or $\epsilon \|\mathbf{A}\|_F$ respectively, strengthening the testing results of [11], which require eigenvalue magnitude at least ϵn . In the graph property testing literature, there is a rich line of work exploring the testing of bounded degree or sparse graphs [13, 16]. Theorem 2 can be thought of as a first step in establishing a related theory of sublinear time approximation algorithms and property testers for sparse matrices.

Surprisingly, in the non-uniform sampling case, the eigenvalue estimates derived from $\mathbf{A}_S$ cannot simply be its scaled eigenvalues, as in Theorem 1. E.g., when $\mathbf{A}$ is the identity, our row sampling probabilities are uniform in all cases. However, the scaled submatrix $\frac{n}{s} \cdot \mathbf{A}_S$ will be a scaled identity, and have eigenvalues equal to n/s —failing to give a $\pm \epsilon \sqrt{\text{nnz}(\mathbf{A})} = \pm \epsilon \|\mathbf{A}\|_F = \pm \epsilon \sqrt{n}$ approximation to the true eigenvalues (all of which are 1) unless $s \gtrsim \frac{\sqrt{n}}{\epsilon}$. To handle this, and related cases, we must argue that selectively zeroing out entries in sufficiently low probability rows/columns of $\mathbf{A}$ (see Algorithms 2 and 3) does not significantly change the spectrum, and ensures concentration of the submatrix eigenvalues. It is not hard to see that simple random submatrix sampling fails even for the easier problem of singular value estimation. Theorems 2 and 3 give the first results of their kinds for this problem as well.

1.2 Related Work

Eigenspectrum estimation is a key primitive in numerical linear algebra, typically known as *spectral density estimation*. The eigenspectrum is viewed as a distribution with mass $1/n$ at each of the n eigenvalues, and the goal is to approximate this distribution [17, 18]. Applications include identifying motifs in social networks [19], studying Hessian and weight matrix spectra in deep learning [20–22], ‘spectrum splitting’ in parallel eigensolvers [23], and the study of many systems in experimental physics and chemistry [24–26].

Recent work has studied sublinear time spectral density estimation for graph structured matrices—Braverman et al. [27] show that the spectral density of a normalized graph adjacency or Laplacian matrix can be estimated to ϵ error in the Wasserstein distance in $\tilde{O}(n/\text{poly}(\epsilon))$ time. Cohen-Steiner, Kong, Sohler, and Valiant study a similar setting, giving runtime $2^{O(1/\epsilon)}$ [28]. We note that the additive error eigenvalue approximation result of Theorem 1 (analogously Theorems 2 and 3) directly gives an ϵn approximation to the spectral density in the Wasserstein distance—extending the above results to a much broader class of matrices. When $\|\mathbf{A}\|_\infty \leq 1$, $\mathbf{A}$ can have eigenvalues as large as n , while the normalized adjacency matrices studied in [27, 28] have eigenvalues in $[-1, 1]$. So, while the results are not directly comparable, our Wasserstein error can be thought as on order of their error of ϵ after scaling.

Our work is also closely related to a line of work on sublinear time property testing for bounded entry matrices, initiated by Balcan et al. [29]. In that work, they study testing of rank, Schatten- p norms, and several other global spectral properties. Sublinear time testing algorithms for the rank and other properties have also been studied under low-rank and bounded row norm assumptions on the input matrix [30, 31]. Recent work studies positive semidefiniteness testing and eigenvalue estimation in the matrix–vector query model, where each query computes $\mathbf{A}\mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^{n \times n}$. As in Theorem 3, $\pm\epsilon\|\mathbf{A}\|_F$ eigenvalue estimation can be achieved with $\text{poly}(\log n, 1/\epsilon)$ queries in this model [32]. Finally, several works study streaming algorithms for eigenspectrum approximation [33–35]. These algorithms are not sublinear time—they require at least linear time to process the input matrix. However, they use sublinear working memory. Note that Theorem 1 immediately gives a sublinear space streaming algorithm for eigenvalue estimation. We can simply store the sampled submatrix $\mathbf{A}_S$ as its entries are updated.

1.3 Technical Overview

In this section, we overview the main techniques used to prove Theorem 1, and then how these techniques are extended to prove Theorems 2 and 3. We start by defining a decomposition of any symmetric $\mathbf{A}$ into the sum of two matrices containing its large and small magnitude eigendirections which correspond to the “outlying” and “middle” eigenvalues respectively.

Definition 1.1 (Eigenvalue Split) Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric. For any $\epsilon, \delta \in (0, 1)$, let $\mathbf{A}_o = \mathbf{V}_o \mathbf{\Lambda}_o \mathbf{V}_o^T$ where $\mathbf{\Lambda}_o$ is diagonal, with the eigenvalues of $\mathbf{A}$ with magnitude $\geq \epsilon\sqrt{\delta}n$ on its diagonal, and $\mathbf{V}_o$ has the corresponding eigenvectors as columns.

Similarly, let $\mathbf{A}_m = \mathbf{V}_m \mathbf{\Lambda}_m \mathbf{V}_m^T$ where $\mathbf{\Lambda}_m$ has the eigenvalues of $\mathbf{A}$ with magnitude $< \epsilon\sqrt{\delta}n$ on its diagonal and $\mathbf{V}_m$ has the corresponding eigenvectors as columns. Then, $\mathbf{A}$ can be decomposed as

$$\mathbf{A} = \mathbf{A}_o + \mathbf{A}_m = \mathbf{V}_o \mathbf{\Lambda}_o \mathbf{V}_o^T + \mathbf{V}_m \mathbf{\Lambda}_m \mathbf{V}_m^T.$$

Any principal submatrix of $\mathbf{A}$, $\mathbf{A}_S$, can be similarly written as

$$\mathbf{A}_S = \mathbf{A}_{o,S} + \mathbf{A}_{m,S} = \mathbf{V}_{o,S} \mathbf{\Lambda}_o \mathbf{V}_{o,S}^T + \mathbf{V}_{m,S} \mathbf{\Lambda}_m \mathbf{V}_{m,S}^T,$$

where $\mathbf{V}_{o,S}$, $\mathbf{V}_{m,S}$ are the corresponding submatrices obtained by sampling rows of $\mathbf{V}_o$, $\mathbf{V}_m$.

Since $\mathbf{A}_S$, $\mathbf{A}_{m,S}$ and $\mathbf{A}_{o,S}$ are all symmetric, we can use Weyl's eigenvalue perturbation theorem [36] to show that for all eigenvalues of $\mathbf{A}_S$,

$$|\lambda_i(\mathbf{A}_S) - \lambda_i(\mathbf{A}_{o,S})| \leq \|\mathbf{A}_{m,S}\|_2. \quad (1)$$

We will argue that the eigenvalues of $\mathbf{A}_{o,S}$ approximate those of $\mathbf{A}_o$ —i.e. all eigenvalues of $\mathbf{A}$ with magnitude $\geq \epsilon\sqrt{\delta}n$. Further, we will show that $\|\mathbf{A}_{m,S}\|_2$ is small with good probability. Thus, via (1), the eigenvalues of $\mathbf{A}_S$ approximate those of $\mathbf{A}_o$. In the estimation procedure of Theorem 1, all other small magnitude eigenvalues of $\mathbf{A}$ are estimated to be 0, which will immediately give our $\pm\epsilon n$ approximation bound when the original eigenvalue has magnitude $\leq \epsilon n$.

Bounding the eigenvalues of $\mathbf{A}_{o,S}$. The first step is to show that the eigenvalues of $\mathbf{A}_{o,S}$ well-approximate those of $\mathbf{A}_o$. As in [11], we critically use that the eigenvectors corresponding to large eigenvalues are *incoherent*—intuitively, since $\|\mathbf{A}\|_\infty$ is bounded, their mass must be spread out in order to witness a large eigenvalue. Specifically, [11] shows that for any eigenvector $\mathbf{v}$ of $\mathbf{A}$ with corresponding eigenvalue $\geq \epsilon\sqrt{\delta}n$, $\|\mathbf{v}\|_\infty \leq \frac{1}{\epsilon\sqrt{\delta}\sqrt{n}}$. We give related bounds on the Euclidean norms of the rows of $\mathbf{V}_o$ (the *leverage scores* of $\mathbf{A}_o$), and on these rows after weighting by $\mathbf{\Lambda}_o$.

Using these incoherence bounds, we argue that the eigenvalues of $\mathbf{A}_{o,S}$ approximate those of $\mathbf{A}_o$ up to $\pm\epsilon n$ error. A key idea is to bound the eigenvalues of $\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,S}^T \mathbf{V}_{o,S} \mathbf{\Lambda}_o^{1/2}$, which are identical to the non-zero eigenvalues of $\mathbf{A}_{o,S} = \mathbf{V}_{o,S} \mathbf{\Lambda}_o \mathbf{V}_{o,S}^T$. Via a matrix Bernstein bound and our incoherence bounds on $\mathbf{V}_o$, we show that this matrix is close to $\mathbf{\Lambda}_o$ with high probability. However, since $\mathbf{\Lambda}_o^{1/2}$ may be complex, the matrix is *not necessarily Hermitian* and standard perturbation bounds [37, 38] do not apply. Thus, to derive an eigenvalue bound, we apply a perturbation bound from Kahan [39, 40], which generalizes Weyl's inequality to the non-Hermitian case, with a $\log n$ factor loss. To the best of our knowledge, this is the first time that perturbation theory bounds for non-Hermitian matrices have been used to prove improved algorithmic results in the theoretical computer science literature.

We note that in “Appendix B”, we give an alternate bound, which instead analyzes the Hermitian matrix $(\mathbf{V}_{o,S}^T \mathbf{V}_{o,S})^{1/2} \mathbf{\Lambda}_o (\mathbf{V}_{o,S}^T \mathbf{V}_{o,S})^{1/2}$, whose eigenvalues are again identical to those of $\mathbf{A}_{o,S}$. This approach only requires Weyl's inequality, and yields

an overall bound of $s = O\left(\frac{\log n}{\epsilon^4 \delta}\right)$, improving the $\log n$ factors of Theorem 1 at the cost of worse ϵ dependence.

Bounding the spectral norm of $\mathbf{A}_{m,S}$. The next step is to show that all eigenvalues of $\mathbf{A}_{m,S}$ are small provided a sufficiently large submatrix is sampled. This means that the “middle” eigenvalues of $\mathbf{A}$, i.e. those with magnitude $\leq \epsilon\sqrt{\delta}n$ do not contribute much to any eigenvalue $\lambda_i(\mathbf{A}_S)$. To do so, we apply a theorem of [9, 10] which shows concentration of the spectral norm of a uniformly random submatrix of an entrywise bounded matrix. Observe that while $\|\mathbf{A}\|_\infty \leq 1$, such a bound will not in general hold for $\|\mathbf{A}_m\|_\infty$. Nevertheless, we can use the incoherence of $\mathbf{V}_o$ to show that $\|\mathbf{A}_o\|_\infty$ is bounded, which via triangle inequality, yields a bound on $\|\mathbf{A}_m\|_\infty \leq \|\mathbf{A}\|_\infty + \|\mathbf{A}_o\|_\infty$. In the end, we show that if $s \geq O\left(\frac{\log n}{\epsilon^2 \delta}\right)$, with probability at least $1 - \delta$, $\|\mathbf{A}_{m,S}\|_2 \leq \epsilon s$. After the n/s scaling in the estimation procedure of Theorem 1, this spectral norm bound translates into an additive ϵn error in approximating the eigenvalues of $\mathbf{A}$.

Completing the argument Once we establish the above bounds on $\mathbf{A}_{o,S}$ and $\mathbf{A}_{m,S}$, Theorem 1 is essentially complete. Any eigenvalue in $\mathbf{A}$ with magnitude $\geq \epsilon n$ will correspond to a nearby eigenvalue in $\frac{n}{s} \cdot \mathbf{A}_{o,S}$ and in turn, $\frac{n}{s} \cdot \mathbf{A}_S$ given our spectral norm bound on $\mathbf{A}_{m,S}$. An eigenvalue in $\mathbf{A}$ with magnitude $\leq \epsilon n$ may or may not correspond to a nearby eigenvalue in $\mathbf{A}_{o,S}$ (it will only if it lies in the range $[\epsilon\sqrt{\delta}n, \epsilon n]$). In any case however, in the estimation procedure of Theorem 1, such an eigenvalue will either be estimated using a small eigenvalue of $\mathbf{A}_S$, or be estimated as 0. In both instances, the estimate will give $\pm \epsilon n$ error, as required.

Can we beat additive error? It is natural to ask if our approach can be improved to yield sublinear time algorithms with stronger relative error approximation guarantees for $\mathbf{A}$ ’s eigenvalues. Unfortunately, this is not possible—consider a matrix with just a single pair of entries $\mathbf{A}_{i,j}, \mathbf{A}_{j,i}$ set to 1. To obtain relative error approximations to the two non-zero eigenvalues, we must find the pair (i, j) , as otherwise we cannot distinguish $\mathbf{A}$ from the all zeros matrix. This requires reading a $\Omega(n^2)$ of $\mathbf{A}$ ’s entries. More generally, consider $\mathbf{A}$ with a random $n/t \times n/t$ principal submatrix populated by all 1s, and with all other entries equal to 0. $\mathbf{A}$ has largest eigenvalue n/t . However, if we read $s \ll t^2$ entries of $\mathbf{A}$, with good probability, we will not see even a single one, and thus we will not be able to distinguish $\mathbf{A}$ from the all zeros matrix. This example establishes that any sublinear time algorithm with query complexity s must incur additive error at least $\Omega(n/\sqrt{s})$.

1.3.1 Improved Bounds via Non-uniform Sampling

We now discuss how to give improved approximation bounds via non-uniform sampling. We focus on the $\pm \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ bound of Theorem 2 using sparsity-based sampling. The proof of Theorem 3 for row norm sampling follows the same general ideas, but with some additional complications.

Theorem 2 requires sampling a submatrix $\mathbf{A}_S$, where each index i is included in S with probability $p_i = \min(1, \frac{s \cdot \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})})$. We reweight each sampled row/column by $\frac{1}{\sqrt{p_i}}$. Thus, if entry $\mathbf{A}_{ij}$ is sampled, it is scaled by $\frac{1}{\sqrt{p_i \cdot p_j}}$. When the rows have uniform

sparsity (so all $p_i = s/n$), this ensures that the full submatrix is scaled by n/s , as in Theorem 1.

The proof of Theorem 2 follows the same outline as that of Theorem 1: we first argue that the outlying eigenvectors in $\mathbf{V}_o$ are incoherent, giving a bound on the norm of each row of $\mathbf{V}_o$ in terms of $\text{nnz}(\mathbf{A}_i)$. We then apply a matrix Bernstein bound and Bhatia's non-Hermitian eigenvalue perturbation bound to show that the eigenvalues of $\mathbf{A}_{o,S}$ approximate those of $\mathbf{A}_o$ up to $\pm \epsilon \sqrt{\text{nnz}(\mathbf{A})}$.

Bounding the spectral norm of $\mathbf{A}_{m,S}$. The major challenge is showing that the subsampled middle eigendirections do not significantly increase the approximation error, which we accomplish by bounding $\|\mathbf{A}_{m,S}\|_2$ by $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$. This is difficult since the indices in $\mathbf{A}_{m,S}$ are sampled nonuniformly, so existing bounds [10] on the spectral norm of uniformly random submatrices do not apply. We extend these bounds to the non-uniform sampling case, but still face an issue due to the rescaling of entries by $\frac{1}{\sqrt{p_i p_j}}$. In fact, without additional algorithmic modifications, $\|\mathbf{A}_{m,S}\|_2$ is simply not bounded by $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$! For example, as already discussed, if $\mathbf{A} = \mathbf{I}$ is the identity matrix, we get $\mathbf{A}_{m,S} = \frac{n}{s} \cdot \mathbf{I}$ and so $\|\mathbf{A}_{m,S}\|_2 = \frac{n}{s} > \epsilon \sqrt{\text{nnz}(\mathbf{A})}$, assuming $s < \frac{\sqrt{n}}{\epsilon}$. Relatedly, suppose that $\mathbf{A}$ is tridiagonal, with zeros on the diagonal and ones on the first diagonals above and below the main diagonal. Then, if $s \geq \sqrt{n}$, with constant probability, one of the ones will be sampled and scaled by $\frac{n}{s}$. Thus, we will again have $\|\mathbf{A}_{m,S}\|_2 \geq \frac{n}{s} \geq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$, assuming $s < \frac{\sqrt{n}}{2\epsilon}$. Observe that this implies $\sigma_1(\mathbf{A}_{m,S})$ is a poor estimate for $\sigma_1(\mathbf{A})$. Thus, while an analogous bound to the uniform sampling result of Theorem 1 can easily be given for singular value estimation via matrix concentration inequalities (see “Appendix G”), to the best of our knowledge, Theorems 2 and 3 are the first of their kind even for singular value estimation.

Zeroing out entries in sparse rows/columns To handle the above cases, we prove a novel perturbation bound, arguing that if we zero out any entry $\mathbf{A}_{ij}$ of $\mathbf{A}$ where $\sqrt{\text{nnz}(\mathbf{A}_i) \cdot \text{nnz}(\mathbf{A}_j)} \leq \frac{\epsilon \sqrt{\text{nnz}(\mathbf{A})}}{c \log n}$, then the eigenvalues of $\mathbf{A}$ are not perturbed by more than $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$. This can be thought of as a strengthening of Girshgorin's circle theorem, which would ensure that zeroing out entries in rows/columns with $\text{nnz}(\mathbf{A}_i) \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ does not perturb the eigenvalues by more than $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$. Armed with this perturbation bound, we argue that if we zero out the appropriate entries of $\mathbf{A}_S$ before computing its eigenvalues, then since we have removed entries in very sparse rows and columns which would be scaled by a large $\frac{1}{\sqrt{p_i p_j}}$ factor in $\mathbf{A}_S$, we can bound $\|\mathbf{A}_{m,S}\|_2$. This requires relating the magnitudes of the entries in $\mathbf{A}_{m,S}$ to those in $\mathbf{A}_S$ using the incoherence of the top eigenvectors, which gives bounds on the entries of $\mathbf{A}_{o,S} = \mathbf{A}_S - \mathbf{A}_{m,S}$.

Sampling model We note that the sparsity-based sampling of Theorem 2 can be efficiently implemented in several natural settings. Given a matrix stored in sparse format, i.e., as a list of nonzero entries, we can easily sample a row with probability $\frac{\text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}$ by sampling a uniformly random non-zero entry and looking at its corresponding row. Via standard techniques, we can convert several such samples into a sampled set S close in distribution to having each $i \in [n]$ included independently with probability

$\min \left(1, \frac{s \operatorname{nnz}(\mathbf{A}_i)}{\operatorname{nnz}(\mathbf{A})} \right)$. If we store the values of $\operatorname{nnz}(\mathbf{A})$, $\operatorname{nnz}(\mathbf{A}_1)$, $\dots$, $\operatorname{nnz}(\mathbf{A}_n)$, we can also efficiently access each p_i , which is needed for rescaling and zeroing out entries. Also observe that if $\mathbf{A}$ is the adjacency matrix of a graph, in the standard graph query model [13], it is well known how to approximately count edges and sample them uniformly at random, i.e., compute $\operatorname{nnz}(\mathbf{A})$ and sample its nonzero entries, in sublinear time [41, 42]. Further, it is typically assumed that one has access to the node degrees, i.e., $\operatorname{nnz}(\mathbf{A}_1)$, $\dots$, $\operatorname{nnz}(\mathbf{A}_n)$. Thus, our algorithm can naturally be used to estimate spectral graph properties in sublinear time.

The ℓ_2 norm-based sampling of Theorem 3 can also be performed efficiently using an augmented data structure for storing $\mathbf{A}$. Such data structures have been used extensively in the literature on quantum-inspired algorithms, and require just $O(\operatorname{nnz}(\mathbf{A}))$ time to construct, $O(\operatorname{nnz}(\mathbf{A}))$ space, and $O(\log n)$ time to update given an update to an entry of $\mathbf{A}$ [15, 43].

1.4 Towards Optimal Query Complexity

As discussed, Bakshi et al. [11] show that any algorithm which can test with good probability whether $\mathbf{A}$ has an eigenvalue $\leq -\epsilon n$ or else has all non-negative eigenvalues must read $\tilde{\Omega}\left(\frac{1}{\epsilon^2}\right)$ entries of $\mathbf{A}$. This testing problem is strictly easier than outputting $\pm\epsilon n$ error estimates of all eigenvalues, so gives a lower bound for our setting. If the queried entries are restricted to fall in a submatrix, [11] shows that this submatrix must have dimensions $\Omega\left(\frac{1}{\epsilon^2}\right) \times \Omega\left(\frac{1}{\epsilon^2}\right)$, giving total query complexity $\Omega\left(\frac{1}{\epsilon^4}\right)$. Closing the gap between our upper bound of $\tilde{O}\left(\frac{\log^3 n}{\epsilon^3}\right) \times \tilde{O}\left(\frac{\log^3 n}{\epsilon^3}\right)$ and the lower bound of $\Omega\left(\frac{1}{\epsilon^2}\right) \times \Omega\left(\frac{1}{\epsilon^2}\right)$ for submatrix queries is an intriguing open question.

We show in “Appendix A” that this gap can be easily closed via a surprisingly simple argument if $\mathbf{A}$ is positive semidefinite (PSD). In that case, $\mathbf{A} = \mathbf{B}\mathbf{B}^T$ with $\mathbf{B} \in \mathbb{R}^{n \times n}$. Writing $\mathbf{A}_S = \mathbf{S}^T \mathbf{A} \mathbf{S}$ for a sampling matrix $\mathbf{S} \in \mathbb{R}^{n \times |S|}$, the non-zero eigenvalues of $\mathbf{A}_S$ are identical to those of $\mathbf{B}\mathbf{S}\mathbf{S}^T \mathbf{B}^T$. Via a standard approximate matrix multiplication analysis [44], one can then show that, for $s \geq \frac{1}{\epsilon^2 \delta}$, with probability at least $1 - \delta$, $\|\mathbf{B}\mathbf{B}^T - \mathbf{B}\mathbf{S}\mathbf{S}^T \mathbf{B}^T\|_F \leq \epsilon n$. Via Weyl’s inequality, this shows that the eigenvalues of $\mathbf{B}\mathbf{S}\mathbf{S}^T \mathbf{B}^T$, and hence $\mathbf{A}_S$, approximate those of $\mathbf{A}$ up to $\pm\epsilon n$ error.²

Unfortunately, this approach breaks down when $\mathbf{A}$ has negative eigenvalues, and so cannot be factored as $\mathbf{B}\mathbf{B}^T$ for real $\mathbf{B} \in \mathbb{R}^{n \times n}$. This is more than a technical issue: observe that when $\mathbf{A}$ is PSD and has $\|\mathbf{A}\|_\infty \leq 1$, it can have at most $1/\epsilon$ eigenvalues larger than ϵn —since its trace, which is equal to the sum of its eigenvalues, is bounded by n , and since all eigenvalues are non-negative. When $\mathbf{A}$ is not PSD, it can have $\Omega(1/\epsilon^2)$ eigenvalues with magnitude larger than ϵn . In particular, if $\mathbf{A}$ is the tensor product of a $1/\epsilon^2 \times 1/\epsilon^2$ random ± 1 matrix and the $\epsilon^2 n \times \epsilon^2 n$ all ones matrix, the bulk of its eigenvalues (of which there are $1/\epsilon^2$) will concentrate around $1/\epsilon \cdot \epsilon^2 n = \epsilon n$. As a result it remains unclear whether we can match the $1/\epsilon^2$ dependence of the PSD case, or if a stronger lower bound can be shown for indefinite matrices.

² In fact, via more refined eigenvalue perturbation bounds [40] one can show an ℓ_2 norm bound on the eigenvalue approximation errors, which can be much stronger than the ℓ_∞ norm bound of Theorem 1.

Outside the ϵ dependence, it is unknown if full eigenspectrum approximation can be performed with sample complexity independent of the matrix size n . [11] achieve this for the easier positive semidefiniteness testing problem, giving sample complexity $\tilde{O}(1/\epsilon^2)$. However our bounds have additional $\log n$ factors. As discussed, in “Appendix B” we give an alternate analysis for Theorem 1, which shows that sampling a $O\left(\frac{\log n}{\epsilon^4 \delta}\right) \times O\left(\frac{\log n}{\epsilon^4 \delta}\right)$ submatrix suffices for $\pm \epsilon n$ eigenvalue approximation, saving a $\log^2 n$ factor at the cost of worse ϵ dependence. However, removing the final $\log n$ seems difficult—it arises when bounding $\|\mathbf{A}_{m,S}\|_2$ via bounds on the spectral norms of random principal submatrices [9]. Removing it seems as though it would require either improving such bounds, or taking a different algorithmic approach.

Also note that our $\log n$ and ϵ dependencies for non-uniform sampling (Theorems 2 and 3) are likely not tight. It is not hard to check that the lower bounds of [11] still hold in these settings. For example, in the sparsity-based sampling setting, by simply having the matrix entirely supported on a $\sqrt{\text{nnz}(\mathbf{A})} \times \sqrt{\text{nnz}(\mathbf{A})}$ submatrix, the lower bounds of [11] directly carry over. Giving tight query complexity bounds here would also be interesting. Finally, it would be interesting to go beyond principal submatrix based algorithms, to achieve improved query complexity, as in Corollary 1. Finding an algorithm matching the $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$ overall query complexity lower bound of [11] is open even in the much simpler PSD setting.

2 Roadmap

We introduce notation and preliminary theorems in Sect. 3. In Sect. 4, we show that using a simple algorithm, we can approximate all the eigenvalues of a bounded-entry symmetric matrix. In Sect. 5, we improve the error of eigenvalue approximation to $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$ when rows and columns are sampled with probability proportional to their sparsities. In Sect. 6, we demonstrate the empirical performance of the sampling algorithms given in Sects. 4 and 5 to approximate all the eigenvalues of several synthetic and real world matrices. In “Appendix A”, we prove optimal sampling bounds for eigenvalue approximation of PSD matrices using uniform sampling. In “Appendix B”, we present an alternate proof which improves the $\log n$ dependence of the upper bounds for sampling complexity for the uniform sampling algorithm at the cost of $O(1/\epsilon)$ worse dependence. In “Appendix C”, we give tighter approximation factors of the algorithm in Sect. 4 under certain assumptions. In “Appendix D”, we give improved spectral norm bounds for non-uniform random submatrices. In “Appendix E”, we achieve analogous bounds to Sect. 5 with error $\epsilon \|\mathbf{A}\|_F$, using an algorithm that samples rows and columns with probability proportional to the ℓ_2 -norm of each row. We improve the sampling complexity for the uniform sampling algorithm in Sect. 4 by using entrywise sampling in “Appendix F” by a factor of $O(\epsilon/\log^3 n)$. Finally in “Appendix G”, we show that all the singular values of a bounded-entry matrix can be approximated using $O(\log n)$ -factor optimal random sampling of rows and columns of the matrix.

3 Notation and Preliminaries

We now define notation and foundational results that we use throughout our work. For any integer n , let $[n]$ denote the set $\{1, 2, \dots, n\}$. We write matrices and vectors in bold literals—e.g., $\mathbf{A}$ or $\mathbf{x}$. We denote the eigenvalues of a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ by $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$, in decreasing order. A symmetric matrix is positive semidefinite if all its eigenvalues are non-negative. For two matrices $\mathbf{A}, \mathbf{B}$, we let $\mathbf{A} \succeq \mathbf{B}$ denote that $\mathbf{A} - \mathbf{B}$ is positive semidefinite. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $i \in [n]$, we let $\mathbf{A}_i$ denote the i^{th} row of $\mathbf{A}$, $\text{nnz}(\mathbf{A}_i)$ denote the number of non-zero elements in this row, and $\|\mathbf{A}_i\|_2$ denote its ℓ_2 norm. We let $\text{nnz}(\mathbf{A})$ denote the total number of non-zero elements $\mathbf{A}$. For a vector $\mathbf{x}$, we let $\|\mathbf{x}\|_2$ denote its Euclidean norm. For a matrix $\mathbf{A}$, we let $\|\mathbf{A}\|_\infty$ denote the largest magnitude of an entry, $\|\mathbf{A}\|_2 = \max_{\mathbf{x}} \frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2}$ denote the spectral norm, $\|\mathbf{A}\|_F = (\sum_{i,j} \mathbf{A}_{ij}^2)^{1/2}$ denote the Frobenius norm, and $\|\mathbf{A}\|_{1 \rightarrow 2}$ denote the maximum Euclidean norm of a column. For $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $S \subseteq [n]$ we let $\mathbf{A}_S$ denote the principal submatrix corresponding to S . We let $\mathbb{E}_2$ denote the L_2 norm of a random variable, $\mathbb{E}_2[X] = (\mathbb{E}[X^2])^{1/2}$, where $\mathbb{E}[\cdot]$ denotes expectation.

We use the following basic facts and identities on eigenvalues throughout our proofs.

Fact 1 (Eigenvalue of Matrix Product) *For any two matrices $\mathbf{A} \in \mathbb{C}^{n \times m}$, $\mathbf{B} \in \mathbb{C}^{m \times n}$, the non-zero eigenvalues of $\mathbf{AB}$ are identical to those of $\mathbf{BA}$.*

Fact 2 (Girshgorin’s circle theorem [45]) *Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ with entries $\mathbf{A}_{ij}$. For $i \in [n]$, let $\mathbf{R}_i$ be the sum of absolute values of non-diagonal entries in the i^{th} row. Let $D(\mathbf{A}_{ii}, \mathbf{R}_i)$ be the closed disc centered at $\mathbf{A}_{ii}$ with radius $\mathbf{R}_i$. Then every eigenvalue of $\mathbf{A}$ lies within one of the discs $D(\mathbf{A}_{ii}, \mathbf{R}_i)$.*

Fact 3 (Weyl’s Inequality [36]) *For any two Hermitian matrices $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{n \times n}$ with $\mathbf{A} - \mathbf{B} = \mathbf{E}$,*

$$\max_i |\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})| \leq \|\mathbf{E}\|_2.$$

Weyl’s inequality ensures that a small Hermitian perturbation of a Hermitian matrix will not significantly change its eigenvalues. The bound can be extended to the case when the perturbation is not Hermitian, with a loss of an $O(\log n)$ factor; to the best of our knowledge this loss is necessary:

Fact 4 (Non-Hermitian perturbation bound (Problem VI.8.6 in [40], Claim (i) in [39])) *Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ be Hermitian and $\mathbf{B} \in \mathbb{C}^{n \times n}$ be any matrix whose eigenvalues are $\lambda_1(\mathbf{B}), \dots, \lambda_n(\mathbf{B})$ such that $\text{Re}(\lambda_1(\mathbf{B})) \geq \dots \geq \text{Re}(\lambda_n(\mathbf{B}))$ (where $\text{Re}(\lambda_i(\mathbf{B}))$ denotes the real part of $\lambda_i(\mathbf{B})$). Let $\mathbf{A} - \mathbf{B} = \mathbf{E}$. For some universal constant C ,*

$$\max_i |\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})| \leq C \log n \|\mathbf{E}\|_2.$$

Beyond the above facts, we use several theorems to obtain eigenvalue concentration bounds. We first state a theorem from [10], which bounds the spectral norm of a principal submatrix sampled uniformly at random from a bounded entry matrix. We build on this to prove the full eigenspectrum concentration result of Theorem 1.

Theorem 4 (Random principal submatrix spectral norm bound [9, 10]) *Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ be Hermitian, decomposed into diagonal and off-diagonal parts: $\mathbf{A} = \mathbf{D} + \mathbf{H}$. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be a diagonal sampling matrix with the j^{th} diagonal entry set to 1 independently with probability s/n and 0 otherwise. Then, for some universal constant C ,*

$$\begin{aligned} & \mathbb{E}_2 \|\mathbf{SAS}\|_2 \\ & \leq C \left[\log n \cdot \mathbb{E}_2 \|\mathbf{SHS}\|_\infty + \sqrt{\frac{s \log n}{n}} \cdot \mathbb{E}_2 \|\mathbf{HS}\|_{1 \rightarrow 2} + \frac{s}{n} \cdot \|\mathbf{H}\|_2 \right] + \mathbb{E}_2 \|\mathbf{SDS}\|_2. \end{aligned}$$

For Theorems 2 and 3, we need an extension of Theorem 4 to the setting where rows are sampled non-uniformly. We will use two bounds here. The first is a decoupling and recoupling result for matrix norms. One can prove this lemma following an analogous result in [10] for sampling rows/columns uniformly. The proof is almost identical so we omit it.

Lemma 1 (Decoupling and recoupling) *Let $\mathbf{H}$ be a Hermitian matrix with zero diagonal. Let δ_j be a sequence of independent random variables such that $\delta_j = \frac{1}{\sqrt{p_j}}$ with probability p_j and 0 otherwise. Let $\mathbf{S}$ be a square diagonal sampling matrix with j^{th} diagonal entry set to δ_j . Then:*

$$\mathbb{E}_2 \|\mathbf{SHS}\|_2 \leq 2\mathbb{E}_2 \|\mathbf{SH}\hat{\mathbf{S}}\|_2 \quad \text{and} \quad \mathbb{E}_2 \|\mathbf{SH}\hat{\mathbf{S}}\|_\infty \leq 4\mathbb{E}_2 \|\mathbf{SHS}\|_\infty,$$

where $\hat{\mathbf{S}}$ is an independent diagonal sampling matrix drawn from the same distribution as $\mathbf{S}$.

The second theorem bounds the spectral norm of a non-uniform random column sample of a matrix. We give a proof in “Appendix D”, again following a theorem in [46] for uniform sampling.

Theorem 5 (Non-uniform column sampling—spectral norm bound) *Let $\mathbf{A}$ be an $m \times n$ matrix with rank r . Let δ_j be a sequence of independent random variables such that $\delta_j = \frac{1}{\sqrt{p_j}}$ with probability p_j and 0 otherwise. Let $\mathbf{S}$ be a square diagonal sampling matrix with j^{th} diagonal entry set to δ_j .*

$$\mathbb{E}_2 \|\mathbf{AS}\|_2 \leq 5\sqrt{\log r} \cdot \mathbb{E}_2 \|\mathbf{AS}\|_{1 \rightarrow 2} + \|\mathbf{A}\|_2$$

We use a standard Matrix Bernstein inequality to bound the spectral norm of random submatrices.

Theorem 6 (Matrix Bernstein [47]) *Consider a finite sequence $\{\mathbf{S}_k\}$ of independent random matrices in $\mathbb{R}^{d \times d}$. Assume that for all k , $\mathbb{E}[\mathbf{S}_k] = \mathbf{0}$ and $\|\mathbf{S}_k\|_2 \leq L$. Let $\mathbf{Z} = \sum_k \mathbf{S}_k$ and let $\mathbf{V}_1, \mathbf{V}_2$ be semidefinite upper-bounds for the matrix valued variances $\mathbf{Var}_1(\mathbf{Z})$ and $\mathbf{Var}_2(\mathbf{Z})$:*

$$\mathbf{V}_1 \succeq \mathbf{Var}_1(\mathbf{Z}) \stackrel{\text{def}}{=} \mathbb{E}(\mathbf{ZZ}^T) = \sum_k \mathbb{E}(\mathbf{S}_k \mathbf{S}_k^T), \quad \text{and}$$

$$\mathbf{V}_2 \succeq \mathbf{Var}_2(\mathbf{Z}) \stackrel{\text{def}}{=} \mathbb{E}(\mathbf{Z}^T \mathbf{Z}) = \sum_k \mathbb{E}(\mathbf{S}_k^T \mathbf{S}_k).$$

Then, letting $v = \max(\|\mathbf{V}_1\|_2, \|\mathbf{V}_2\|_2)$, for any $t \geq 0$,

$$\mathbb{P}(\|\mathbf{Z}\|_2 \geq t) \leq 2d \cdot \exp\left(\frac{-t^2/2}{v + Lt/3}\right).$$

For real valued random variables, we use the standard Bernstein inequality.

Theorem 7 (Bernstein inequality [48]) *Let $\{z_j\}$ for $j \in [n]$ be independent random variables with zero mean such that $|z_j| \leq M$ for all j . Then for all positive t ,*

$$\mathbb{P}\left(\left|\sum_{j=1}^n z_j\right| \geq t\right) \leq \exp\left(\frac{-t^2/2}{\sum_{i=1}^n \mathbb{E}[z_i^2] + Mt/3}\right).$$

4 Sublinear Time Eigenvalue Estimation Using Uniform Sampling

We now prove our main eigenvalue estimation result—Theorem 1. We give the pseudocode for our principal submatrix based estimation procedure in Algorithm 1. We will show that any positive or negative eigenvalue of $\mathbf{A}$ with magnitude $\geq \epsilon n$ will appear as an approximate eigenvalue in $\mathbf{A}_S$ with good probability. Thus, in step 5 of Algorithm 1, the positive and negative eigenvalues of $\mathbf{A}_S$ are used to estimate the outlying largest and smallest eigenvalues of $\mathbf{A}$. All other interior eigenvalues of $\mathbf{A}$ are estimated to be 0, which will immediately give our $\pm \epsilon n$ approximation bound when the original eigenvalue has magnitude $\leq \epsilon n$.

Algorithm 1 Eigenvalue estimator using uniform sampling

- 1: **Input:** Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$, Accuracy $\epsilon \in (0, 1)$, failure prob. $\delta \in (0, 1)$.
 - 2: Fix $s = \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \delta}$ where c is a sufficiently large constant.
 - 3: Add each index $i \in [n]$ to the sample set S independently with probability $\frac{s}{n}$. Let the principal submatrix of $\mathbf{A}$ corresponding S be $\mathbf{A}_S$.
 - 4: Compute the eigenvalues of $\mathbf{A}_S$: $\lambda_1(\mathbf{A}_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}_S)$.
 - 5: For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all remaining $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$.
 - 6: **Return:** Eigenvalue estimates $\tilde{\lambda}_1(\mathbf{A}) \geq \dots \geq \tilde{\lambda}_n(\mathbf{A})$.
-

Running time Observe that the expected number of indices chosen by Algorithm 1 is $s = \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \delta}$. A standard concentration bound can be used to show that with high probability $(1 - 1/\text{poly}(n))$, the number of sampled entries is $O(s)$. Thus, the algorithm reads a total of $O(s^2)$ entries of $\mathbf{A}$ and runs in $O(s^\omega)$ time—the time to compute a full eigendecomposition of $\mathbf{A}_S$.

4.1 Outer and Middle Eigenvalue Bounds

Recall that we will split $\mathbf{A}$ into two symmetric matrices (Definition 1.1): $\mathbf{A}_o = \mathbf{V}_o \mathbf{\Lambda}_o \mathbf{V}_o^T$ which contains its large magnitude (outlying) eigendirections with eigenvalue magnitudes $\geq \epsilon\sqrt{\delta}n$ and $\mathbf{A}_m = \mathbf{V}_m \mathbf{\Lambda}_m \mathbf{V}_m^T$ which contains its small magnitude (middle) eigendirections.

We first show that the eigenvectors in $\mathbf{V}_o$ are *incoherent*. That is, that their (eigenvalue weighted) squared row norms are bounded. This ensures that the outlying eigenspace of $\mathbf{A}$ is well-approximated via uniform sampling.

Lemma 2 (Incoherence of outlying eigenvectors) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$. Let $\mathbf{V}_o$ be as in Definition 1.1. Let $\mathbf{V}_{o,i}$ denote the i^{th} row of $\mathbf{V}_o$. Then,*

$$\|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i}\|_2^2 \leq \frac{1}{\epsilon\sqrt{\delta}} \quad \text{and} \quad \|\mathbf{V}_{o,i}\|_2^2 \leq \frac{1}{\epsilon^2 \delta n}.$$

Proof Observe that $\mathbf{A}\mathbf{V}_o = \mathbf{V}_o \mathbf{\Lambda}_o$. Let $[\mathbf{A}\mathbf{V}_o]_i$ denote the i^{th} row of the $\mathbf{A}\mathbf{V}_o$. Then we have

$$\|[\mathbf{A}\mathbf{V}_o]_i\|_2^2 = \|[\mathbf{V}_o \mathbf{\Lambda}_o]_i\|_2^2 = \sum_{j=1}^r \lambda_j^2 \cdot \mathbf{V}_{o,i,j}^2, \quad (2)$$

where $r = \text{rank}(\mathbf{A}_o)$, $\mathbf{V}_{o,i,j}$ is the $(i, j)^{\text{th}}$ element of $\mathbf{V}_o$ and $\lambda_j = \mathbf{\Lambda}_o(j, j)$. $\|\mathbf{A}\|_\infty \leq 1$ by assumption and since $\mathbf{V}_o$ has orthonormal columns, its spectral norm is bounded by 1, thus we have

$$\|[\mathbf{A}\mathbf{V}_o]_i\|_2^2 = \|\mathbf{A}\|_i \|\mathbf{V}_o\|_2^2 \leq \|\mathbf{A}\|_i \|\mathbf{V}_o\|_2^2 \leq n.$$

Therefore, by (2), we have:

$$\sum_{j=1}^r \lambda_j^2 \cdot \mathbf{V}_{o,i,j}^2 \leq n. \quad (3)$$

Since by definition of $\mathbf{\Lambda}_o$, $|\lambda_j| \geq \epsilon\sqrt{\delta}n$ for all j , we finally have

$$\|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i}\|_2^2 = \sum_{j=1}^r \lambda_j \cdot \mathbf{V}_{o,i,j}^2 \leq \frac{n}{\epsilon\sqrt{\delta}n} = \frac{1}{\epsilon\sqrt{\delta}}$$

and

$$\|\mathbf{V}_{o,i}\|_2^2 = \sum_{j=1}^r \mathbf{V}_{o,i,j}^2 \leq \frac{n}{\epsilon^2 \delta n^2} = \frac{1}{\epsilon^2 \delta n}.$$

□

Let $\bar{\mathbf{S}} \in \mathbb{R}^{n \times |S|}$ be the scaled sampling matrix satisfying $\bar{\mathbf{S}}^T \mathbf{A} \bar{\mathbf{S}} = \frac{n}{s} \cdot \mathbf{A}_S$. We next apply Lemma 2 in conjunction with a matrix Bernstein bound to show that $\mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2}$ concentrates around its expectation, $\mathbf{\Lambda}_o$. Since by Fact 1, this matrix has identical eigenvalues to $\frac{n}{s} \cdot \mathbf{A}_{o,S} = \bar{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o \mathbf{V}_o^T \bar{\mathbf{S}}$, this allows us to argue that the eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,S}$ approximate those of $\mathbf{\Lambda}_o$.

Lemma 3 (Concentration of outlying eigenvalues) *Let $S \subseteq [n]$ be sampled as in Algorithm 1 for $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^3 \sqrt{\delta}}$ where c is a sufficiently large constant. Let $\bar{\mathbf{S}} \in \mathbb{R}^{n \times |S|}$ be the scaled sampling matrix satisfying $\bar{\mathbf{S}}^T \mathbf{A} \bar{\mathbf{S}} = \frac{n}{s} \cdot \mathbf{A}_S$. Letting $\mathbf{\Lambda}_o, \mathbf{V}_o$ be as in Definition 1.1, with probability at least $1 - \delta$,*

$$\|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2} - \mathbf{\Lambda}_o\|_2 \leq \epsilon n.$$

Proof Define $\mathbf{E} = \mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2} - \mathbf{\Lambda}_o$. For all $i \in [n]$, let $\mathbf{V}_{o,i}$ be the i^{th} row of $\mathbf{V}_o$ and define the matrix valued random variable

$$\mathbf{Y}_i = \begin{cases} \frac{n}{s} \mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}_o^{1/2}, & \text{with probability } s/n \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Define $\mathbf{Q}_i = \mathbf{Y}_i - \mathbb{E}[\mathbf{Y}_i]$. Observe that $\mathbf{Q}_1, \dots, \mathbf{Q}_n$ are independent random variables and that $\sum_{i=1}^n \mathbf{Q}_i = \mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2} - \mathbf{\Lambda}_o = \mathbf{E}$. Further, observe that $\|\mathbf{Q}_i\|_2 \leq \max(1, \frac{n}{s} - 1) \cdot \|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}_o^{1/2}\|_2 \leq \max(1, \frac{n}{s} - 1) \cdot \|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i}\|_2^2$. Now, $\|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i}\|_2^2 \leq \frac{1}{\epsilon \sqrt{\delta}}$ by Lemma 2. Thus, $\|\mathbf{Q}_i\|_2 \leq \frac{n}{\epsilon \sqrt{\delta}}$. The variance $\mathbf{Var}(\mathbf{E}) \stackrel{\text{def}}{=} \mathbb{E}(\mathbf{E} \mathbf{E}^T) = \mathbb{E}(\mathbf{E}^T \mathbf{E}) = \sum_{i=1}^n \mathbb{E}[\mathbf{Q}_i^2]$ can be bounded as:

$$\begin{aligned} \sum_{i=1}^n \mathbb{E}[\mathbf{Q}_i^2] &= \sum_{i=1}^n \left[\frac{s}{n} \cdot \left(\frac{n}{s} - 1 \right)^2 + \left(1 - \frac{s}{n} \right) \right] \cdot (\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}_o \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}_o^{1/2}) \\ &\leq \sum_{i=1}^n \frac{n}{s} \cdot \|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i}\|_2^2 \cdot (\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}_o^{1/2}). \end{aligned} \quad (5)$$

Again by Lemma 2, $\|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i}\|_2^2 \leq \frac{1}{\epsilon \sqrt{\delta}}$. Plugging back into (5) we can bound,

$$\sum_{i=1}^n \mathbb{E}[\mathbf{Q}_i^2] \leq \sum_{i=1}^n \frac{n}{s} \cdot \frac{1}{\epsilon \sqrt{\delta}} \cdot (\mathbf{\Lambda}_o^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}_o^{1/2}) = \frac{n}{s \epsilon \sqrt{\delta}} \mathbf{\Lambda}_o \leq \frac{n^2}{s \epsilon \sqrt{\delta}} \cdot \mathbf{I}.$$

Since $\mathbf{Q}_i^2$ is PSD, this establishes that $\|\mathbf{Var}(\mathbf{E})\|_2 \leq \frac{n^2}{s \epsilon \sqrt{\delta}}$. We then apply Theorem 6 (the matrix Bernstein inequality) with $L = \frac{n}{s \epsilon \sqrt{\delta}}$, $v = \frac{n^2}{s \epsilon \sqrt{\delta}}$, and $d \leq \frac{1}{\epsilon^2 \delta}$ since there are at most $\frac{\|\mathbf{A}\|_F^2}{\delta \epsilon^2 n^2} \leq \frac{1}{\epsilon^2 \delta}$ outlying eigenvalues with magnitude $\geq \sqrt{\delta} \epsilon n$ in $\mathbf{\Lambda}_o$. This gives:

$$\begin{aligned}
\mathbb{P}(\|\mathbf{E}\|_2 \geq \epsilon n) &\leq \frac{2}{\epsilon^2 \delta} \cdot \exp\left(\frac{-\epsilon^2 n^2/2}{v + L\epsilon n/3}\right) \\
&\leq \frac{2}{\epsilon^2 \delta} \cdot \exp\left(\frac{-\epsilon^2 n^2/2}{\frac{n^2}{s\epsilon\sqrt{\delta}} + \frac{\epsilon n^2}{3s\epsilon\sqrt{\delta}}}\right) \\
&\leq \frac{2}{\epsilon^2 \delta} \cdot \exp\left(\frac{-s\epsilon^3\sqrt{\delta}}{4}\right).
\end{aligned}$$

Thus, if we set $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^3\sqrt{\delta}}$ for large enough c , then the probability is bounded above by δ , completing the proof. $\square$

We cannot prove an analogous leverage score bound to Lemma 2 for the interior eigenvectors of $\mathbf{A}$ appearing in $\mathbf{V}_m$. Thus we cannot apply a matrix Bernstein bound as in Lemma 3. However, we can use Theorem 4 to show that the spectral norm of the random principal submatrix $\mathbf{A}_{m,S}$ is not too large, and thus that the eigenvalues of $\mathbf{A}_S = \mathbf{A}_{o,S} + \mathbf{A}_{m,S}$ are close to those of $\mathbf{A}_{o,S}$.

Lemma 4 (Spectral norm bound—sampled middle eigenvalues) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$. Let $\mathbf{A}_m$ be as in Definition 1.1. Let S be sampled as in Algorithm 1. If $s \geq \frac{c \log n}{\epsilon^2 \delta}$ for some sufficiently large constant c , then with probability at least $1 - \delta$, $\|\mathbf{A}_{m,S}\|_2 \leq \epsilon s$.*

Proof Let $\mathbf{A}_m = \mathbf{D}_m + \mathbf{H}_m$ where $\mathbf{D}_m$ is the matrix of diagonal elements and $\mathbf{H}_m$ the matrix of off-diagonal elements. Let $\mathbf{S} \in \mathbb{R}^{n \times |S|}$ be the binary sampling matrix with $\mathbf{A}_{m,S} = \mathbf{S}^T \mathbf{A}_m \mathbf{S}$. From Theorem 4, we have for some constant C ,

$$\begin{aligned}
\mathbb{E}_2[\|\mathbf{A}_{m,S}\|_2] &\leq C \left[\log n \cdot \mathbb{E}_2[\|\mathbf{S}^T \mathbf{H}_m \mathbf{S}\|_\infty] \right. \\
&\quad \left. + \sqrt{\frac{s \log n}{n}} \mathbb{E}_2[\|\mathbf{H}_m \mathbf{S}\|_{1 \rightarrow 2}] + \frac{s}{n} \|\mathbf{H}_m\|_2 \right] + \mathbb{E}_2[\|\mathbf{S}^T \mathbf{D}_m \mathbf{S}\|_2]. \quad (6)
\end{aligned}$$

Considering the various terms in (6), we have $\|\mathbf{S}^T \mathbf{H}_m \mathbf{S}\|_\infty \leq \|\mathbf{A}_m\|_\infty$ and $\|\mathbf{S}^T \mathbf{D}_m \mathbf{S}\|_2 = \|\mathbf{S}^T \mathbf{D}_m \mathbf{S}\|_\infty \leq \|\mathbf{A}_m\|_\infty$. We also have

$$\|\mathbf{H}_m\|_2 \leq \|\mathbf{A}_m\|_2 + \|\mathbf{D}_m\|_2 \leq \|\mathbf{A}_m\|_2 + \|\mathbf{A}_m\|_\infty \leq \epsilon \delta^{1/2} n + \|\mathbf{A}_m\|_\infty$$

and

$$\|\mathbf{H}_m \mathbf{S}\|_{1 \rightarrow 2} \leq \|\mathbf{A}_m \mathbf{S}\|_{1 \rightarrow 2} \leq \|\mathbf{A}_m\|_{1 \rightarrow 2} \leq \sqrt{n}.$$

The final bound follows since $\mathbf{A}_m = \mathbf{V}_m \mathbf{V}_m^T \mathbf{A}$, where $\mathbf{V}_m \mathbf{V}_m^T$ is an orthogonal projection matrix. Thus, $\|\mathbf{A}_m\|_{1 \rightarrow 2} \leq \|\mathbf{A}\|_{1 \rightarrow 2} \leq \sqrt{n}$ by our assumption that $\|\mathbf{A}\|_\infty \leq 1$. Plugging all these bounds into (6) we have, for some constant C ,

$$\mathbb{E}_2[\|\mathbf{A}_{m,S}\|_2] \leq C \left[\log n \cdot \|\mathbf{A}_m\|_\infty + \sqrt{\log n \cdot s} + s \cdot \epsilon \delta^{1/2} \right]. \quad (7)$$

It remains to bound $\|\mathbf{A}_m\|_\infty$. We have $\mathbf{A} = \mathbf{A}_m + \mathbf{A}_o$ and thus by triangle inequality,

$$\|\mathbf{A}_m\|_\infty \leq \|\mathbf{A}\|_\infty + \|\mathbf{A}_o\|_\infty = 1 + \|\mathbf{A}_o\|_\infty. \quad (8)$$

Writing $\mathbf{A}_o = \mathbf{V}_o \mathbf{\Lambda}_o \mathbf{V}_o^T$ (see Definition 1.1), and letting $\mathbf{V}_{o,i}$ denote the i^{th} row of $\mathbf{V}_o$, the $(i, j)^{\text{th}}$ element of $\mathbf{A}_o$ has magnitude

$$|\mathbf{A}_{o,i,j}| = |\mathbf{V}_{o,i} \mathbf{\Lambda}_o \mathbf{V}_{o,j}^T| \leq \|\mathbf{V}_{o,i}\|_2 \cdot \|\mathbf{\Lambda}_o \mathbf{V}_{o,j}^T\|_2,$$

by Cauchy-Schwarz. From Lemma 2, we have $\|\mathbf{V}_{o,i}\|_2 \leq \frac{1}{\epsilon \delta^{1/2} \sqrt{n}}$. Also, from (2), $\|\mathbf{\Lambda}_o \mathbf{V}_{o,j}^T\|_2 = \|[\mathbf{A} \mathbf{V}_o]_j\|_2 \leq \sqrt{n}$. Overall, for all i, j we have $\mathbf{A}_{o,i,j} \leq \frac{1}{\epsilon \delta^{1/2} \sqrt{n}} \cdot \sqrt{n} = \frac{1}{\epsilon \delta^{1/2}}$, giving $\|\mathbf{A}_o\|_\infty \leq \frac{1}{\epsilon \delta^{1/2}}$. Plugging back into (8) and in turn (7), we have for some constant C ,

$$\mathbb{E}_2[\|\mathbf{A}_{m,s}\|_2] \leq C \left[\frac{\log n}{\epsilon \delta^{1/2}} + \sqrt{s \log n} + s \epsilon \delta^{1/2} \right].$$

Setting $s \geq \frac{c \log n}{\epsilon^2 \delta}$ for sufficiently large c , all terms in the right hand side of the above equation are bounded by $\epsilon \sqrt{\delta} s$ and so

$$\mathbb{E}_2[\|\mathbf{A}_{m,s}\|_2] \leq 3\epsilon \sqrt{\delta} s$$

Thus, by Markov's inequality, with probability at least $1 - \delta$, we have $\|\mathbf{A}_{m,s}\|_2 \leq 3\epsilon s$. We can adjust ϵ by a constant to obtain the required bound. $\square$

4.2 Main Accuracy Bounds

We now restate our main result, and give its proof via Lemmas 3 and 4.

Theorem 1 (Sublinear Time Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including each index independently with probability s/n as in Algorithm 1. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, with eigenvalues $\lambda_1(\mathbf{A}_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}_S)$.*

For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all other $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$. If $s \geq \frac{c \log(1/(\epsilon \delta)) \cdot \log^3 n}{\epsilon^3 \delta}$, for large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$,

$$\lambda_i(\mathbf{A}) - \epsilon n \leq \tilde{\lambda}_i(\mathbf{A}) \leq \lambda_i(\mathbf{A}) + \epsilon n.$$

Proof Let $\mathbf{S} \in \mathbb{R}^{n \times |S|}$ be the binary sampling matrix with a single one in each column such that $\mathbf{S}^T \mathbf{A} \mathbf{S} = \mathbf{A}_S$. Let $\tilde{\mathbf{S}} = \sqrt{n/s} \cdot \mathbf{S}$. Following Definition 1.1, we write $\mathbf{A} = \mathbf{A}_o + \mathbf{A}_m$. By Fact 1 we have that the nonzero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,S} = \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o \mathbf{V}_o^T \tilde{\mathbf{S}}$

are identical to those of $\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}$ where $\Lambda_o^{1/2}$ is the square root matrix of Λ_o such that $\Lambda_o^{1/2} \Lambda_o^{1/2} = \Lambda_o$.

Note that Λ_o is Hermitian. However $\Lambda_o^{1/2}$ may be complex, and hence $\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}$ is *not necessarily Hermitian*, although it does have real eigenvalues. Thus, we can apply the perturbation bound of Fact 4 to Λ_o and $\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}$ to claim for all $i \in [n]$, and some constant C ,

$$|\lambda_i(\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}) - \lambda_i(\Lambda_o)| \leq C \log n \|\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2} - \Lambda_o\|_2.$$

By Lemma 3 applied with error $\frac{\epsilon}{2C \log n}$, with probability at least $1 - \delta$, for any $s \geq \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}$ (for a large enough constant c) we have $\|\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2} - \Lambda_o\|_2 \leq \frac{\epsilon n}{2C \log n}$. Thus, for all i ,

$$\left| \lambda_i(\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}) - \lambda_i(\Lambda_o) \right| < \frac{\epsilon n}{2}. \quad (9)$$

We note that the conceptual part of the proof is essentially complete: the nonzero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,s}$ are identical to those of $\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}$, which we have shown well approximate those of Λ_o and in turn $\mathbf{A}_o$. i.e., the non-zero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,s}$ approximate all outlying eigenvalues of $\mathbf{A}$. It remains to carefully argue how these approximations should be ‘lined up’ given the presence of zero eigenvalues in the spectrum of these matrices. We also must account for the impact of the interior eigenvalues in $\mathbf{A}_{m,s}$, which is limited by the spectral norm bound of Lemma 4.

Eigenvalue alignment and effect of interior eigenvalues. First recall that $\mathbf{A}_S = \mathbf{A}_{o,s} + \mathbf{A}_{m,s}$. By Lemma 4 applied with error $\epsilon/2$, we have $\|\mathbf{A}_{m,s}\|_2 \leq \epsilon/2 \cdot s$ with probability at least $1 - \delta$ when $s \geq \frac{c \log n}{\epsilon^2 \delta}$. By Weyl’s inequality (Fact 3), for all $i \in [|S|]$ we thus have

$$\left| \frac{n}{s} \lambda_i(\mathbf{A}_S) - \frac{n}{s} \lambda_i(\mathbf{A}_{o,s}) \right| \leq \frac{n}{s} \cdot \frac{\epsilon s}{2} = \frac{\epsilon n}{2}. \quad (10)$$

Consider $i \in [|S|]$ with $\lambda_i(\mathbf{A}_{o,s}) > 0$. Since the nonzero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,s}$ are identical to those of $\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2}$, $\frac{n}{s} \cdot \lambda_i(\mathbf{A}_{o,s}) = \lambda_i(\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2})$, and so by (9),

$$\left| \frac{n}{s} \cdot \lambda_i(\mathbf{A}_{o,s}) - \lambda_i(\Lambda_o) \right| < \frac{\epsilon n}{2}. \quad (11)$$

Analogously, consider $i \in [|S|]$ such that $\lambda_i(\mathbf{A}_{o,s}) < 0$. We have $\frac{n}{s} \cdot \lambda_i(\mathbf{A}_{o,s}) = \lambda_{r-(|S|-i)}(\Lambda_o^{1/2} \mathbf{V}_o^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o \Lambda_o^{1/2})$, where $r = \text{rank}(\mathbf{A}_{o,s})$ is the dimension of Λ_o —i.e., the number of outlying eigenvalues in $\mathbf{A}$. Again by (9) we have

$$\left| \frac{n}{s} \cdot \lambda_i(\mathbf{A}_{o,s}) - \lambda_{r-(|S|-i)}(\Lambda_o) \right| < \frac{\epsilon n}{2}. \quad (12)$$

Now the nonzero eigenvalues of $\mathbf{A}_o$ are identical to those of $\mathbf{A}_o$. Consider $i \in [|S|]$ such that $\lambda_i(\mathbf{A}_S) \geq \epsilon s$. In this case, by (10), (11), and the triangle inequality, we have $\lambda_i(\mathbf{A}_o) > 0$ and thus we have $\lambda_i(\mathbf{A}_o) = \lambda_i(\mathbf{A}_o)$. In turn, again applying (10), (11), and the triangle inequality, we have

$$\left| \frac{n}{s} \lambda_i(\mathbf{A}_S) - \lambda_i(\mathbf{A}_o) \right| \leq \left| \frac{n}{s} \lambda_i(\mathbf{A}_{o,S}) - \lambda_i(\mathbf{A}_o) \right| + \left| \frac{n}{s} \lambda_i(\mathbf{A}_S) - \lambda_i(\mathbf{A}_{o,S}) \right| \leq \epsilon n.$$

Analogously, for $i \in [|S|]$ such that $\lambda_i(\mathbf{A}_S) \leq -\epsilon s$, we have by (10) and (12) that $\lambda_{r-(|S|-i)}(\mathbf{A}_o) < 0$. Thus $\lambda_{r-(|S|-i)}(\mathbf{A}_o) = \lambda_{n-(|S|-i)}(\mathbf{A}_o)$. Again by (10), (12), and triangle inequality this gives

$$\left| \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S) - \lambda_{n-(|S|-i)}(\mathbf{A}_o) \right| \leq \epsilon n.$$

Now, consider all $i \in [n]$ such that $\lambda_i(\mathbf{A}_o)$ is not well approximated by one of the outlying eigenvalues of $\mathbf{A}_S$ as argued above. By (10), (11), and (12), all such eigenvalues must have $|\lambda_i(\mathbf{A}_o)| \leq 2\epsilon n$. Thus, if we approximate them in any way either by the remaining eigenvalues of $\mathbf{A}_S$ with magnitude $\leq \epsilon s$, or else by 0, we will approximate all to error at most $3\epsilon n$. Thus, if (as in Algorithm 1) for $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, we let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$ and for $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$, and let $\tilde{\lambda}_i(\mathbf{A}) = 0$ for all other i , we will have for all i ,

$$\left| \tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A}_o) \right| \leq 3\epsilon n.$$

Finally by definition, for all i , $|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}_o)| \leq \epsilon\sqrt{\delta}n \leq \epsilon n$ and thus, via triangle inequality, $\left| \tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A}) \right| \leq 4\epsilon n$. This gives our final error bound after adjusting constants on ϵ .

Recall that we require $s \geq \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}$ for the outer eigenvalue bound of (9) to hold with probability $1 - \delta$. We require $s \geq \frac{c \log n}{\epsilon^2 \delta}$ for $\|\mathbf{A}_{m,S}\|_2 \leq \epsilon/2 \cdot s$ to hold with probability $1 - \delta$ by Lemma 4. Thus, for both conditions to hold simultaneously with probability $1 - 2\delta$ by a union bound, it suffices to set $s = \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \delta} \geq \max \left(\frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}, \frac{c \log n}{\epsilon^2 \delta} \right)$, where we use that $\log(1/(\epsilon\delta)) \leq O(\log n)$, as otherwise our algorithm can take $\mathbf{A}_S$ to be the full matrix $\mathbf{A}$. Adjusting δ to $\delta/2$ completes the theorem. $\square$

Remark: The proof of Lemma 3 and consequently, Theorem 1 can be modified to give better bounds for the case when the eigenvalues of $\mathbf{A}_o$ lie in a bounded range—between $\epsilon^a \sqrt{\delta}n$ and $\epsilon^b n$ where $0 \leq b \leq a \leq 1$. See Theorem 9 in “Appendix C” for details. For example, if all the top eigenvalues are equal, one can show that $s = \tilde{O} \left(\frac{\log^2 n}{\epsilon^2} \right)$ suffices to give $\pm \epsilon n$ error, nearly matching the lower bound of [11]. This seems to indicate that improving Theorem 1 in general requires tackling the case when the outlying eigenvalues in $\mathbf{A}_o$ have a wide range.

5 Improved Bounds via Sparsity-Based Sampling

We now prove the $\pm\epsilon\sqrt{\text{nnz}(\mathbf{A})}$ approximation bound of Theorem 2, assuming the ability to sample each row with probability proportional to $\frac{\text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}$. Pseudocode for our algorithm is given in Algorithm 2. Unlike in the uniform sampling case (Algorithm 1), we cannot simply sample a principal submatrix of $\mathbf{A}$ and compute its eigenvalues. We must carefully zero out entries lying at the intersection of sparse rows and columns to ensure accuracy of our estimates. A similar approach is taken for the norm-based sampling result of Theorem 3. We defer that proof to “Appendix E”.

5.1 Preliminary Lemmas

Our first step is to argue that zeroing out entries in sparse rows/columns in step 5 of Algorithm 2 does not introduce significant error. We define $\mathbf{A}' \in \mathbb{R}^{n \times n}$ to be the extension of $\mathbf{A}'_S$ to the original matrix—i.e., $\mathbf{A}'_{ij} = 0$ whenever $i = j$ or $\text{nnz}(\mathbf{A}_i) \text{nnz}(\mathbf{A}_j) < \frac{\epsilon^2 \text{nnz}(\mathbf{A})}{c_2 \log^2 n}$. Otherwise $\mathbf{A}'_{ij} = \mathbf{A}_{ij}$. We argue via a strengthening of Girshgorin’s theorem that $|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}')| \leq \epsilon\sqrt{\text{nnz}(\mathbf{A})}$ for all i .

After this step is complete, our proof follows the same general outline as that of Theorem 1 in Sect. 3. We split $\mathbf{A}' = \mathbf{A}'_o + \mathbf{A}'_m$, arguing that (1) after sampling $\|\mathbf{A}'_{m,S}\|_2 \leq \epsilon\sqrt{\text{nnz}(\mathbf{A})}$ and (2) that the eigenvalues of $\mathbf{A}'_{o,S}$ are $\pm\epsilon\sqrt{\text{nnz}(\mathbf{A})}$ approximations to those of $\mathbf{A}'_o$. In both cases, we critically use that the rescaling factors introduced in line 4 of Algorithm 2 do not introduce too much variance, due to the zeroing out of entries in $\mathbf{A}'$.

Algorithm 2 Eigenvalue estimator using sparsity-based sampling

- 1: **Input:** Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$, Accuracy $\epsilon \in (0, 1)$, failure prob. $\delta \in (0, 1)$. $\text{nnz}(\mathbf{A}_i)$ for all $i \in [n]$ and $\text{nnz}(\mathbf{A})$.
- 2: Fix $s = \frac{c_1 \log^8 n}{\epsilon^8 \delta^4}$ where c_1 is a sufficiently large constant.
- 3: Add each $i \in [n]$ to sample set S independently, with probability $p_i = \min\left(1, \frac{s \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}\right)$. Let the principal submatrix of $\mathbf{A}$ corresponding to S be $\mathbf{A}_S$.
- 4: Let $\mathbf{A}_S = \mathbf{D}\mathbf{A}_S\mathbf{D}$ where $\mathbf{D} \in \mathbb{R}^{|S| \times |S|}$ is diagonal with $\mathbf{D}_{i,i} = \frac{1}{\sqrt{p_j}}$ if the i^{th} element of S is j .
- 5: Construct $\mathbf{A}'_S \in \mathbb{R}^{|S| \times |S|}$ from $\mathbf{A}_S$ as follows:

$$[\mathbf{A}'_S]_{i,j} = \begin{cases} 0 & \text{if } i = j \text{ or } \text{nnz}(\mathbf{A}_i) \text{nnz}(\mathbf{A}_j) < \frac{\epsilon^2 \text{nnz}(\mathbf{A})}{c_2 \log^2 n} \text{ for sufficient large constant } c_2 \\ [\mathbf{A}_S]_{i,j} & \text{otherwise.} \end{cases}$$

- 6: Compute the eigenvalues of $\mathbf{A}'_S$: $\lambda_1(\mathbf{A}'_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}'_S)$.
 - 7: For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}'_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \lambda_i(\mathbf{A}'_S)$. For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}'_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \lambda_i(\mathbf{A}'_S)$. For all remaining $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$.
 - 8: **Return:** Eigenvalue estimates $\tilde{\lambda}_1(\mathbf{A}) \geq \dots \geq \tilde{\lambda}_n(\mathbf{A})$.
-

Remark Throughout, we will assume that $\mathbf{A}$ does not have any rows/columns that are all 0, as such rows will never be sampled and will have no effect on the output of

Algorithm 2. Additionally, we will assume that $\text{nnz}(\mathbf{A}) \geq \frac{c_1 \log^8 n}{\epsilon^8 \delta^4}$, as otherwise, $\mathbf{A}$ has at most $s = \frac{c_1 \log^8 n}{\epsilon^8 \delta^4}$ non-zero rows. Thus, rather than running Algorithm 2, we can directly compute the eigenvalues of $\mathbf{A}$.

Lemma 5 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and $\text{nnz}(\mathbf{A}) \geq 2/\epsilon^2$. Let $\mathbf{A}' \in \mathbb{R}^{n \times n}$ have $\mathbf{A}'_{ij} = 0$ if $i = j$ or $\text{nnz}(\mathbf{A}_i) \cdot \text{nnz}(\mathbf{A}_j) < \frac{\epsilon^2 \text{nnz}(\mathbf{A})}{c_2 \log^2 n}$ for a sufficiently large constant c_2 and $\mathbf{A}'_{ij} = \mathbf{A}_{ij}$ otherwise. Then, for all $i \in [n]$,*

$$|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}')| \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}.$$

Proof We consider the matrix $\mathbf{A}''$, which is defined identically to $\mathbf{A}'$ except we only set $\mathbf{A}''_{ij} = 0$ if $\text{nnz}(\mathbf{A}_i) \cdot \text{nnz}(\mathbf{A}_j) < \frac{\epsilon^2 \text{nnz}(\mathbf{A})}{c_2 \log^2 n}$. That is, we do not have the condition requiring setting the diagonal to 0. We will show that $|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}'')| \leq \epsilon/2 \cdot \sqrt{\text{nnz}(\mathbf{A})}$. By Weyl's inequality, and the assumption that $\text{nnz}(\mathbf{A}) \geq 2/\epsilon^2$, we then have $|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}')| \leq \epsilon/2 \cdot \sqrt{\text{nnz}(\mathbf{A})} + 1 \leq \epsilon \cdot \sqrt{\text{nnz}(\mathbf{A})}$ as required.

Let $\mathcal{I}_k \subset [n]$ be the set of rows/columns with $\text{nnz}(\mathbf{A}_i) \in \left[\frac{\text{nnz}(\mathbf{A})}{2^k}, \frac{\text{nnz}(\mathbf{A})}{2^{k-1}}\right)$ and $\mathbf{A}_{kl} = \mathbf{A}(\mathcal{I}_k, \mathcal{I}_l)$ be the submatrix of $\mathbf{A}$ formed with rows in $\mathcal{I}_k$ and columns in $\mathcal{I}_l$. Define $\mathbf{A}''_{kl}$ in the same way and observe that $\mathbf{A}''_{kl} = \mathbf{A}_{kl}$ whenever $2^{k+l} \leq \frac{c_2 \text{nnz}(\mathbf{A}) \log^2 n}{\epsilon^2}$.

When $2^{k+l} > \frac{c_2 \text{nnz}(\mathbf{A}) \log^2 n}{\epsilon^2}$, we may zero out some entries of $\mathbf{A}_{kl}$ to produce $\mathbf{A}''_{kl}$. Let $\widehat{\mathbf{A}}_{kl}$ be equal to $\mathbf{A}_{kl}$ on this set of zeroed out entries, and 0 everywhere else. Observe that $(\widehat{\mathbf{A}}_{kl} \widehat{\mathbf{A}}_{kl}^T)_{m,:} = (\widehat{\mathbf{A}}_{kl})_{m,:} \widehat{\mathbf{A}}_{kl}^T$. Next observe that $(\widehat{\mathbf{A}}_{kl})_{m,:}$ has at most $\text{nnz}(\mathbf{A}_m) \leq \frac{\text{nnz}(\mathbf{A})}{2^{k-1}}$ non-zero entries. Similarly, each row of $\widehat{\mathbf{A}}_{kl}^T$ has at most $\frac{\text{nnz}(\mathbf{A})}{2^{l-1}}$ non-zero elements. Thus, for all $m \in |\mathcal{I}_k|$, using that $\|\mathbf{A}\|_\infty \leq 1$,

$$\|(\widehat{\mathbf{A}}_{kl} \widehat{\mathbf{A}}_{kl}^T)_{m,:}\|_1 \leq \frac{\text{nnz}(\mathbf{A})^2}{2^{k+l-2}} = \frac{4 \text{nnz}(\mathbf{A})^2}{2^{k+l}}.$$

Applying Girshgorin's circle theorem (Theorem 2) we thus have:

$$\|\widehat{\mathbf{A}}_{kl}\|_2^2 = \|\widehat{\mathbf{A}}_{kl} \widehat{\mathbf{A}}_{kl}^T\|_2 \leq \max_m \|(\widehat{\mathbf{A}}_{kl} \widehat{\mathbf{A}}_{kl}^T)_{m,:}\|_1 \leq \frac{4 \text{nnz}(\mathbf{A})^2}{2^{k+l}}. \quad (13)$$

Let $\bar{\mathbf{A}}_{kl} \in \mathbb{R}^{n \times n}$ be a symmetric matrix such that $\bar{\mathbf{A}}_{kl}(\mathcal{I}_k, \mathcal{I}_l) = \widehat{\mathbf{A}}_{kl}$, $\bar{\mathbf{A}}_{kl}(\mathcal{I}_l, \mathcal{I}_k) = \widehat{\mathbf{A}}_{lk}$, and $\bar{\mathbf{A}}_{kl}$ is zero everywhere else. By triangle inequality and the bound of (13),

$$\|\bar{\mathbf{A}}_{kl}\|_2 \leq \|\widehat{\mathbf{A}}_{kl}\|_2 + \|\widehat{\mathbf{A}}_{lk}\|_2 \leq \frac{4 \text{nnz}(\mathbf{A})}{2^{(k+l)/2}}.$$

Observe that, since we assume all rows have at least one non-zero entry, $\text{nnz}(\mathbf{A}_i) \geq 1$ and $\text{nnz}(\mathbf{A}) \leq n^2$. Therefore, k, l can range from 1 to $\log(n^2) = 2 \log n$. By triangle inequality,

$$\|\mathbf{A} - \mathbf{A}''\|_2 \leq \sum_{(k,l): 2^{k+l} > \frac{c_2 \text{nnz}(\mathbf{A}) \log^2 n}{\epsilon^2}} \|\bar{\mathbf{A}}_{kl}\|$$

$$\begin{aligned}
&\leq \sum_{k=1}^{2 \log n} \frac{4\epsilon \sqrt{\text{nnz}(\mathbf{A})}}{\sqrt{c_2} \cdot \log n} \cdot \sum_{i=1}^{2 \log n} \frac{1}{2^{i-1}} \\
&\leq \frac{16\epsilon \sqrt{\text{nnz}(\mathbf{A})}}{\sqrt{c_2}}.
\end{aligned}$$

Finally, setting c_2 large enough and using Weyl's inequality (Fact 4) we have the required bound:

$$|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}'')| \leq \epsilon/2\sqrt{\text{nnz}(\mathbf{A})}.$$

□

We next give a bound on the coherence of the outlying eigenvectors of $\mathbf{A}'$. This bound is analogous to Lemma 2, but is more refined, taking into account the sparsity of each row.

Lemma 6 (Incoherence of outlying eigenvectors in terms of sparsity) *Let $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{n \times n}$ be as in Lemma 5. Let $\mathbf{A}'_o = \mathbf{V}'_o \mathbf{\Lambda}'_o \mathbf{V}'_o{}^T$ where $\mathbf{\Lambda}'_o$ is diagonal, with the eigenvalues of $\mathbf{A}'$ with magnitude $\geq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$ on its diagonal, and $\mathbf{V}'_o$ has columns equal to the corresponding eigenvectors. Let $\mathbf{V}'_{o,i}$ denote the i^{th} row of $\mathbf{V}'_o$. Then,*

$$\|\mathbf{A}'^{1/2} \mathbf{V}'_{o,i}\|_2^2 \leq \frac{\text{nnz}(\mathbf{A}_i)}{\epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}} \quad \text{and} \quad \|\mathbf{V}'_{o,i}\|_2^2 \leq \frac{\text{nnz}(\mathbf{A}_i)}{\epsilon^2\delta \text{nnz}(\mathbf{A})}.$$

Proof The proof is nearly identical to that of Lemma 2. Observe that $\mathbf{A}'\mathbf{V}'_o = \mathbf{V}'_o\mathbf{\Lambda}'_o$. Letting $[\mathbf{A}'\mathbf{V}'_o]_i$ denote the i^{th} row of the $\mathbf{A}'\mathbf{V}'_o$, we have

$$\|[\mathbf{A}'\mathbf{V}'_o]_i\|_2^2 = \|[\mathbf{V}'_o\mathbf{\Lambda}'_o]_i\|_2^2 = \sum_{j=1}^r \lambda_j^2 \cdot \mathbf{V}'_{o,i,j}{}^2, \quad (14)$$

where $r = \text{rank}(\mathbf{A}'_o)$, $\mathbf{V}'_{o,i,j}$ is the $(i, j)^{\text{th}}$ element of $\mathbf{V}'_o$ and $\lambda_j = \mathbf{\Lambda}'_o(j, j)$. Since $\mathbf{V}'_o$ has orthonormal columns, we thus have $\|[\mathbf{A}'\mathbf{V}'_o]_i\|_2^2 \leq \|\mathbf{A}'_i\|_2^2 \leq \|\mathbf{A}_i\|_2^2 \leq \text{nnz}(\mathbf{A}_i)$. Therefore, by (14),

$$\sum_{j=1}^r \lambda_j^2 \cdot \mathbf{V}'_{o,i,j}{}^2 \leq \text{nnz}(\mathbf{A}_i). \quad (15)$$

Since by definition $|\lambda_j| \geq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$ for all j , we can conclude that $\|\mathbf{A}'^{1/2} \mathbf{V}'_{o,i}\|_2^2 = \sum_{j=1}^r \lambda_j \cdot \mathbf{V}'_{o,i,j}{}^2 \leq \frac{\text{nnz}(\mathbf{A}_i)}{\epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}}$ and $\|\mathbf{V}'_{o,i}\|_2^2 = \sum_{j=1}^r \mathbf{V}'_{o,i,j}{}^2 \leq \frac{\text{nnz}(\mathbf{A}_i)}{\epsilon^2\delta \text{nnz}(\mathbf{A})}$, which completes the lemma. □

5.2 Outer and Middle Eigenvalue Bounds

Using Lemma 6, we next argue that the eigenvalues of $\mathbf{A}'_{o,s}$ will approximate those of $\mathbf{A}'$, and in turn those of $\mathbf{A}$. The proof is very similar to Lemma 3 in the uniform sampling case.

Lemma 7 (Concentration of outlying eigenvalues with sparsity-based sampling) *Let $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{n \times n}$ be as in Lemmas 5 and 6. Let $\mathbf{A}' = \mathbf{A}'_m + \mathbf{A}'_o$, where $\mathbf{A}'_m = \mathbf{V}'_m \mathbf{\Lambda}'_m \mathbf{V}'_m{}^T$, and $\mathbf{A}'_o = \mathbf{V}'_o \mathbf{\Lambda}'_o \mathbf{V}'_o{}^T$ are projections onto the eigenspaces with magnitude $< \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}$ and $\geq \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}$ respectively (analogous to Definition 1.1) As in Algorithm 2, for all $i \in [n]$ let $p_i = \min \left(1, \frac{s \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})} \right)$ and let $\tilde{\mathbf{S}}$ be a scaled diagonal sampling matrix such that the $\tilde{\mathbf{S}}_{ii} = \frac{1}{\sqrt{p_i}}$ with probability p_i and $\tilde{\mathbf{S}}_{ii} = 0$ otherwise. If $s \geq \frac{c \log(1/(\epsilon \delta))}{\epsilon^3 \sqrt{\delta}}$ for a large enough constant c , then with probability at least $1 - \delta$, $\|\mathbf{A}'^{1/2} \mathbf{V}'_o{}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}'_o \mathbf{\Lambda}'_o{}^{1/2} - \mathbf{\Lambda}'_o\|_2 \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$.*

Proof Define $\mathbf{E} = \mathbf{A}'^{1/2} \mathbf{V}'_o{}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}'_o \mathbf{\Lambda}'_o{}^{1/2} - \mathbf{\Lambda}'_o$. For all $i \in [n]$, let $\mathbf{V}_{o,i}$ be the i^{th} row of $\mathbf{V}'_o$ and define the matrix valued random variable

$$\mathbf{Y}_i = \begin{cases} \frac{1}{p_i} \mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}'_{o,i} \mathbf{V}'_{o,i}{}^T \mathbf{\Lambda}'_o{}^{1/2}, & \text{with probability } p_i \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

Define $\mathbf{Q}_i = \mathbf{Y}_i - \mathbb{E}[\mathbf{Y}_i]$. We can observe that $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_n$ are independent random variables and that $\sum_{i=1}^n \mathbf{Q}_i = \mathbf{A}'^{1/2} \mathbf{V}'_o{}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}'_o \mathbf{\Lambda}'_o{}^{1/2} - \mathbf{\Lambda}'_o = \mathbf{E}$. Let $P = \{i \in [n] : p_i < 1\}$. Then, observe that $\sum_{i \in [n] \setminus P} \mathbf{Q}_i = 0$. So, $\mathbf{E} = \sum_{i \in P} \mathbf{Q}_i$. Then, similar to the proof of Lemma 3, we need to bound $\|\mathbf{Q}_i\|_2$ for all $i \in P$ and $\mathbf{Var}(\mathbf{E}) \stackrel{\text{def}}{=} \mathbb{E}(\mathbf{E} \mathbf{E}^T) = \mathbb{E}(\mathbf{E}^T \mathbf{E}) = \sum_{i \in P} \mathbb{E}[\mathbf{Q}_i^2]$ using the improved row norm bounds of Lemma 5. In particular, we have

$$\begin{aligned} \sum_{i \in P} \mathbb{E}[\mathbf{Q}_i^2] &= \sum_{i \in P} \left[p_i \cdot \left(\frac{1}{p_i} - 1 \right)^2 + (1 - p_i) \right] \cdot (\mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}'_o \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}'_o{}^{1/2}) \\ &\leq \sum_{i \in P} \frac{1}{p_i} \cdot \|\mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}_{o,i}\|_2^2 \cdot (\mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}'_o{}^{1/2}). \end{aligned} \quad (17)$$

By Lemma 6, $\|\mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}_{o,i}\|_2^2 \leq \frac{\text{nnz}(\mathbf{A}_i)}{\epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}}$. Plugging back into (17),

$$\begin{aligned} \sum_{i \in P} \mathbb{E}[\mathbf{Q}_i^2] &\leq \sum_{i \in P} \frac{1}{p_i} \cdot \frac{\text{nnz}(\mathbf{A}_i)}{\epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}} \cdot (\mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}'_o{}^{1/2}) \\ &\leq \frac{\sqrt{\text{nnz}(\mathbf{A})}}{s \epsilon \sqrt{\delta}} \left(\sum_{i \in P} \mathbf{\Lambda}'_o{}^{1/2} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T \mathbf{\Lambda}'_o{}^{1/2} \right) \\ &= \frac{\sqrt{\text{nnz}(\mathbf{A})}}{s \epsilon \sqrt{\delta}} \mathbf{\Lambda}'_o \leq \frac{\text{nnz}(\mathbf{A})}{s \epsilon \sqrt{\delta}} \cdot \mathbf{I}. \end{aligned}$$

Since $\mathbf{Q}_i^2$ is PSD this establishes that $v \leq \|\mathbf{Var}(\mathbf{E})\|_2 \leq \frac{\text{nnz}(\mathbf{A})}{s\epsilon\sqrt{\delta}}$. Since there are at most $\frac{\text{nnz}(\mathbf{A})}{\delta\epsilon^2\text{nnz}(\mathbf{A})} = \frac{1}{\epsilon^2\delta}$ eigenvalues with absolute value $\geq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$, we can apply the matrix Bernstein inequality exactly as in the proof of Lemma 3 with $d = \frac{1}{\epsilon^2\delta}$ to show that when $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^3\sqrt{\delta}}$ for large enough c , with probability at least $1 - \delta$, $\|\mathbf{E}\|_2 \leq \epsilon\sqrt{\text{nnz}(\mathbf{A})}$. $\square$

We next bound the spectral norm of $\mathbf{A}'_{m,S}$. This is the most challenging part of the proof—the rows of this matrix are sampled non-uniformly and scaled proportional to their inverse sampling probabilities, so we cannot apply existing bounds on the spectral norms of uniformly sampled random submatrices [9]. We extend these bounds to the non-uniform case, critically using that entries which would be scaled up significantly after sampling (i.e. those lying in sparse rows/columns), have already been set to 0 in $\mathbf{A}'_{m,S}$, and thus do not contribute to the spectral norm.

Lemma 8 (Concentration of middle eigenvalues with sparsity-based sampling) *Let $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{n \times n}$ be as in Lemmas 5 and 6. Let $\mathbf{A}' = \mathbf{A}'_m + \mathbf{A}'_o$, where $\mathbf{A}'_m = \mathbf{V}'_m \mathbf{A}'_m \mathbf{V}'_m{}^T$, and $\mathbf{A}'_o = \mathbf{V}'_o \mathbf{A}'_o \mathbf{V}'_o{}^T$ are projections onto the eigenspaces with magnitude $< \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$ and $\geq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$ respectively (analogous to Definition 1.1). As in Algorithm 2, for all $i \in [n]$ let $p_i = \min\left(1, \frac{s \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}\right)$ and let $\tilde{\mathbf{S}}$ be a scaled diagonal sampling matrix such that the $\tilde{\mathbf{S}}_{ii} = \frac{1}{\sqrt{p_i}}$ with probability p_i and $\tilde{\mathbf{S}}_{ii} = 0$ otherwise.*

If $s \geq \frac{c \log^8 n}{\epsilon^8 \delta^4}$ for a large enough constant c , then with probability at least $1 - \delta$,

$$\|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq \epsilon\sqrt{\text{nnz}(\mathbf{A})}.$$

Proof The initial part of the proof follows the outline of proof of the spectral norm bound for uniformly random submatrices (Theorem 4) of [10]. From Lemma 6, we have $\|\mathbf{V}'_{o,i}\|_2 \leq \frac{\sqrt{\text{nnz}(\mathbf{A}_i)}}{\epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}}$. Also, following the proof of Lemma 6, we have $\|\mathbf{A}'_o \mathbf{V}'_{o,j}\|_2 = \|[\mathbf{A}' \mathbf{V}'_o]_j\|_2 \leq \sqrt{\text{nnz}(\mathbf{A}_j)}$. Thus, for all $i, j \in [n]$, using Cauchy Schwarz's inequality, we have

$$|\mathbf{A}'_{o,i,j}| = |\mathbf{V}'_{o,i} \mathbf{A}'_o \mathbf{V}'_{o,j}{}^T| \leq \|\mathbf{V}'_{o,i}\|_2 \cdot \|\mathbf{A}'_o \mathbf{V}'_{o,j}{}^T\|_2 \leq \frac{\sqrt{\text{nnz}(\mathbf{A}_i)}}{\epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}} \cdot \sqrt{\text{nnz}(\mathbf{A}_j)}. \quad (18)$$

Let $\mathbf{A}'_m = \mathbf{H}_m + \mathbf{D}_m$ where $\mathbf{H}_m$ and $\mathbf{D}_m$ contain the off-diagonal and diagonal elements of $\mathbf{A}'_m$ respectively. Note that while $\mathbf{A}'$ is zero on the diagonal, $\mathbf{A}'_m$ may not be. We have:

$$\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq \mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \tilde{\mathbf{S}}\|_2 + \mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{D}_m \tilde{\mathbf{S}}\|_2.$$

Using Lemma 1 (decoupling) on $\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \tilde{\mathbf{S}}\|_2$, we get

$$\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq 2\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \hat{\mathbf{S}}\|_2 + \mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{D}_m \tilde{\mathbf{S}}\|_2, \quad (19)$$

where $\hat{\mathbf{S}}$ is an independent copy of $\tilde{\mathbf{S}}$. Upper bounding the rank of $\mathbf{H}_m$ as n and applying Theorem 5 twice to $\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m\tilde{\mathbf{S}}\|_2$, once for each operator, we get

$$\begin{aligned} \mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_2 &\leq 5\sqrt{\log n} \mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} + \mathbb{E}_2 \|\hat{\mathbf{S}}\mathbf{H}_m\|_2 \\ &\leq 5\sqrt{\log n} \mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} + 5\sqrt{\log n} \mathbb{E}_2 \|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} + \|\mathbf{H}_m\|_2. \end{aligned} \quad (20)$$

Plugging (20) into (19), we have:

$$\begin{aligned} \mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{A}'_m\tilde{\mathbf{S}}\|_2 &\leq 10\sqrt{\log n} \left(\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} + \mathbb{E}_2 \|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} \right) \\ &\quad + 2\|\mathbf{H}_m\|_2 + \mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{D}_m\tilde{\mathbf{S}}\|_2 \end{aligned} \quad (21)$$

We now proceed to bound each of the terms on the right hand side of (21). We start with $\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{D}_m\tilde{\mathbf{S}}\|_2$. First, observe that $\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{D}_m\tilde{\mathbf{S}}\|_2 \leq \max_i \frac{1}{p_i} |(\mathbf{D}_m)_{ii}|$. We consider two cases.

Case 1 $p_i < 1$. Then, $p_i = \frac{s \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}$ and $|(\mathbf{D}_m)_{ii}| = |(\mathbf{A}'_m)_{ii}| = |(\mathbf{A}'_o)_{ii}|$ (since $\mathbf{A}'_{ii} = 0$). Then by (18), we have $\frac{1}{p_i} |(\mathbf{D}_m)_{ii}| \leq \frac{\sqrt{\text{nnz}(\mathbf{A})}}{s\epsilon\sqrt{\delta}}$.

Case 2 $p_i = 1$. Then we have $\frac{1}{p_i} |(\mathbf{D}_m)_{ii}| = |(\mathbf{D}_m)_{ii}| \leq \max_j |(\mathbf{D}_m)_{jj}| \leq \|\mathbf{A}'_m\|_2 \leq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$.

From the two cases above, for $s \geq \frac{1}{\epsilon^2\delta}$, we have:

$$\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{D}_m\tilde{\mathbf{S}}\|_2 \leq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}. \quad (22)$$

We can bound $\|\mathbf{H}_m\|_2$ similarly. Since $\mathbf{H}_m = \mathbf{A}'_m - \mathbf{D}_m$ and $\|\mathbf{A}'_m\|_2 \leq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$,

$$\begin{aligned} \|\mathbf{H}_m\|_2 &\leq \|\mathbf{A}'_m\|_2 + \|\mathbf{D}_m\|_2 \\ &\leq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})} + \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})} \\ &= 2\epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})} \end{aligned} \quad (23)$$

where the second step follows from the fact that $\|\mathbf{D}_m\|_2 \leq \max_i |(\mathbf{D}_m)_{ii}| \leq \|\mathbf{A}'_m\|_2$.

We next bound the term $\mathbb{E}_2 \|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2}$. Observe that $\mathbb{E}_2 \|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \frac{\max_i \|\mathbf{A}'_{m,i}\|_2}{\sqrt{p_i}}$, where $\mathbf{A}'_{m,i}$ is the i^{th} column/row of $\mathbf{A}'_m$. We again consider the two cases when $p_i = 1$ and $p_i < 1$:

Case 1 $p_i = 1$. Then $\|\mathbf{A}'_{m,i}\|_2 \leq \|\mathbf{A}'_m\|_2 \leq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}$.

Case 2 $p_i < 1$. Then $\|\mathbf{A}'_{m,i}\|_2 \leq \|\mathbf{A}'_i\|_2 \leq \sqrt{\text{nnz}(\mathbf{A}_i)}$. Thus, setting $s \geq \frac{1}{\epsilon^2\delta}$ we have:

$$\frac{\|\mathbf{A}'_{m,i}\|_2}{\sqrt{p_i}} \leq \sqrt{\frac{\text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_i)}} \cdot \|\mathbf{A}'_i\|_2 \leq \sqrt{\frac{\text{nnz}(\mathbf{A})}{s}} \leq \epsilon\sqrt{\delta}\sqrt{\text{nnz}(\mathbf{A})}.$$

Thus, from the two cases above, for all $i \in [n]$, adjusting ϵ by a $\frac{1}{\sqrt{\log n}}$ factor, we have for $s \geq \frac{\log n}{\epsilon^2 \delta}$:

$$\mathbb{E}_2 \|\mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \frac{\epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}}{\sqrt{\log n}}. \quad (24)$$

Overall, plugging (22), (23), and (24) back into (21), we have:

$$\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq 10 \sqrt{\log n} \cdot \mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} + 15 \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}. \quad (25)$$

It remains to bound $\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2}$, which is the most complex part of the proof. Since $\hat{\mathbf{S}}$ is an independent copy of $\tilde{\mathbf{S}}$, we upper bound the norm of the i^{th} column of $\tilde{\mathbf{S}} \mathbf{H}_m \hat{\mathbf{S}}$ by $\frac{\|(\tilde{\mathbf{S}} \mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}}$. Then $\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \mathbb{E}_2 \left(\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}} \mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \right)$. We will argue that $\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}} \mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}}$ is bounded by $\epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}$ with probability $1 - 1/\text{poly}(n)$. Since our sampling probabilities are all at least $1/n^2$ and since $\|\mathbf{H}_m\|_F \leq \|\mathbf{A}\|_F \leq n$, this value is also deterministically bounded by n^2 . Thus, our high probability bound implies the needed bound on $\mathbb{E}_2 \left(\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}} \mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \right)$.

We begin by observing that since $\mathbf{A}'_m = \mathbf{H}_m + \mathbf{D}_m$, $\|(\tilde{\mathbf{S}} \mathbf{A}'_m)_{:,i}\|_2 \geq \|(\tilde{\mathbf{S}} \mathbf{H}_m)_{:,i}\|_2$, and so to bound $\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}} \mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}}$, it suffices to bound $\frac{\|(\tilde{\mathbf{S}} \mathbf{A}'_m)_{:,i}\|_2}{\sqrt{p_i}}$ for all $i \in [n]$. Towards this end, for a fixed i and any $j \in [n]$, define

$$z_j = \begin{cases} \frac{1}{p_j} |\mathbf{A}'_{m,i,j}|^2 & \text{with probability } p_j \\ 0 & \text{otherwise.} \end{cases}$$

Then $\sum_{j=1}^n z_j = \|(\tilde{\mathbf{S}} \mathbf{A}'_m)_{:,i}\|_2^2$ and $\mathbb{E}[\sum_{j=1}^n z_j] = \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_i\|_2^2 \leq \text{nnz}(\mathbf{A}_i)$. Since $\sum_{j=1}^n z_j = \|(\tilde{\mathbf{S}} \mathbf{A}'_m)_{:,i}\|_2^2$ is a sum of independent random variables, we can bound this quantity by applying Bernstein's inequality. To do this, we must bound $|z_j|$ for all $j \in [n]$ and $\text{Var} \left(\sum_{j=1}^n z_j \right)$. We will again consider the cases of $p_i < 1$ and $p_i = 1$ separately.

Case 1: $p_i < 1$. Then, we have $p_i = s \text{nnz}(\mathbf{A}_i) / \text{nnz}(\mathbf{A})$. If $\mathbf{A}'_{i,j} \neq 0$ then

$$\begin{aligned} |z_j| &\leq \frac{1}{p_j} |\mathbf{A}'_{m,i,j}|^2 \leq \max \left(1, \frac{\text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} \right) |\mathbf{A}'_{m,i,j}|^2 \\ &\leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2 \text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} \left(|\mathbf{A}'_{i,j}|^2 + |\mathbf{A}'_{o,i,j}|^2 \right) \\ &\leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2 \text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} \left(|\mathbf{A}'_{i,j}|^2 + \frac{\text{nnz}(\mathbf{A}_i) \text{nnz}(\mathbf{A}_j)}{\epsilon^2 \delta \text{nnz}(\mathbf{A})} \right) \\ &\leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2 \text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} |\mathbf{A}'_{i,j}|^2 + \frac{2 \text{nnz}(\mathbf{A}_i)}{\epsilon^2 \delta s}, \end{aligned}$$

where the fourth inequality uses (18). By the thresholding procedure which defines $\mathbf{A}'$, if $\mathbf{A}'_{ij} \neq 0$,

$$\text{nnz}(\mathbf{A}_i) \cdot \text{nnz}(\mathbf{A}_j) \geq \frac{\epsilon^2 \text{nnz}(\mathbf{A})}{c_2 \log^2 n} \Rightarrow \text{nnz}(\mathbf{A}_j) \geq \frac{\epsilon^2 \text{nnz}(\mathbf{A})}{c_2 \log^2 n \text{nnz}(\mathbf{A}_i)}, \quad (26)$$

and thus for $p_i < 1$ and $\mathbf{A}'_{ij} \neq 0$ we have

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2c_2 \log^2 n \text{nnz}(\mathbf{A}_i)}{s\epsilon^2} + \frac{2 \text{nnz}(\mathbf{A}_i)}{\epsilon^2 \delta s}.$$

If $\mathbf{A}'_{i,j} = 0$ then we simply have

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{\text{nnz}(\mathbf{A}_i)}{s\epsilon^2 \delta}.$$

Overall for all $j \in [n]$,

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2 \text{nnz}(\mathbf{A}_i)}{s\epsilon^2 \delta} + \frac{2c_2 \log^2 n \text{nnz}(\mathbf{A}_i)}{s\epsilon^2}, \quad (27)$$

and since $|\mathbf{A}'_{m,i,j}|^2 \leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^2 = \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_i\|_2^2 \leq \text{nnz}(\mathbf{A}_i)$,

$$|z_j| \leq \text{nnz}(\mathbf{A}_i) + \frac{2 \text{nnz}(\mathbf{A}_i)}{s\epsilon^2 \delta} + \frac{2c_2 \log^2 n \text{nnz}(\mathbf{A}_i)}{s\epsilon^2}. \quad (28)$$

For $s \geq c \left(\frac{\log^2 n}{\epsilon^2} + \frac{1}{\epsilon^2 \delta} \right)$ and large enough c , we thus have $|z_j| \leq 2 \text{nnz}(\mathbf{A})$.

We next bound the variance by:

$$\begin{aligned} \text{Var} \left(\sum_{j=1}^n z_j \right) &\leq \sum_{j=1}^n \mathbb{E}[z_j^2] \leq \sum_{j=1}^n p_j \frac{1}{p_j^2} |\mathbf{A}'_{m,i,j}|^4 \\ &= \sum_{j=1}^n \max \left(1, \frac{\text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} \right) |\mathbf{A}'_{m,i,j}|^4 \\ &\leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^4 + \sum_{j=1}^n \frac{12 \text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} \left(|\mathbf{A}'_{i,j}|^4 + |\mathbf{A}'_{o,i,j}|^4 \right) \\ &\leq \|\mathbf{A}'_{m,i}\|_2^4 + \sum_{j=1}^n \frac{12 \text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_j)} \left(|\mathbf{A}'_{i,j}|^4 + \frac{\text{nnz}(\mathbf{A}_i)^2 \text{nnz}(\mathbf{A}_j)^2}{\epsilon^4 \delta^2 \text{nnz}(\mathbf{A})^2} \right), \end{aligned}$$

where the last inequality uses (18). Now since $\mathbf{A}'_{ii} = 0$ for all i and $\|\mathbf{A}'\|_\infty \leq 1$ we have

$$\mathbf{Var} \left(\sum_{j=1}^n z_j \right) \leq \|\mathbf{A}'_{m,i}\|_2^4 + \sum_{j:\mathbf{A}'_{i,j} \neq 0} \frac{12 \operatorname{nnz}(\mathbf{A})}{s \operatorname{nnz}(\mathbf{A}_j)} + \sum_{j=1}^n \frac{12 \operatorname{nnz}(\mathbf{A}_i)^2 \operatorname{nnz}(\mathbf{A}_j)}{s \epsilon^4 \delta^2 \operatorname{nnz}(\mathbf{A})}. \quad (29)$$

Combining (26) with the second term to the right of (29) we have

$$\begin{aligned} \mathbf{Var} \left(\sum_{j=1}^n z_j \right) &\leq \|\mathbf{A}'_{m,i}\|_2^4 + \sum_{j:\mathbf{A}'_{i,j} \neq 0} \frac{12c_2 \log^2 n \cdot \operatorname{nnz}(\mathbf{A}_i)}{s \epsilon^2} \\ &\quad + \sum_{j=1}^n \frac{12 \operatorname{nnz}(\mathbf{A}_i)^2 \operatorname{nnz}(\mathbf{A}_j)}{s \epsilon^4 \delta^2 \operatorname{nnz}(\mathbf{A})}, \end{aligned}$$

and since $|\{j : \mathbf{A}'_{i,j} \neq 0\}| = \operatorname{nnz}(\mathbf{A}_i)$, we have

$$\mathbf{Var} \left(\sum_{j=1}^n z_j \right) \leq \|\mathbf{A}'_{m,i}\|_2^4 + \frac{12c_2 \log^2 n \cdot \operatorname{nnz}(\mathbf{A}_i)^2}{s \epsilon^2} + \sum_{j=1}^n \frac{12 \operatorname{nnz}(\mathbf{A}_i)^2 \operatorname{nnz}(\mathbf{A}_j)}{s \epsilon^4 \delta^2 \operatorname{nnz}(\mathbf{A})}. \quad (30)$$

Finally since $\sum_{j=1}^n \operatorname{nnz}(\mathbf{A}_j) = \operatorname{nnz}(\mathbf{A})$ and $\|\mathbf{A}'_{m,i}\|_2^4 \leq \|\mathbf{A}'_i\|_2^4 \leq \operatorname{nnz}(\mathbf{A}_i)^2$ we have

$$\mathbf{Var} \left(\sum_{j=1}^n z_j \right) \leq \operatorname{nnz}(\mathbf{A}_i)^2 + \frac{12c_2 \log^2 n \cdot \operatorname{nnz}(\mathbf{A}_i)^2}{s \epsilon^2} + \frac{12 \operatorname{nnz}(\mathbf{A}_i)^2}{s \epsilon^4 \delta^2}. \quad (31)$$

For $s \geq c \left(\frac{\log^2 n}{\epsilon^2} + \frac{1}{\epsilon^4 \delta^2} \right)$ for large enough c , we have $\mathbf{Var} \left(\sum_{j=1}^n z_j \right) \leq 2 \operatorname{nnz}(\mathbf{A}_i)^2$. Therefore, using (28) and (31) with $s \geq c \left(\frac{\log^2 n}{\epsilon^2} + \frac{1}{\epsilon^4 \delta^2} \right)$, we can apply Bernstein inequality (Theorem 7) (for some constant c) to get

$$\begin{aligned} \mathbb{P} \left(\left\| (\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i} \right\|_2^2 \geq \mathbb{E} \left\| (\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i} \right\|_2^2 + t \right) &\leq \mathbb{P} \left(\sum_{j=1}^n z_j \geq \operatorname{nnz}(\mathbf{A}_i) + t \right) \\ &\leq \exp \left(\frac{-t^2/2}{c \operatorname{nnz}(\mathbf{A}_i)^2 + ct \operatorname{nnz}(\mathbf{A}_i)/3} \right). \end{aligned}$$

If we set $t = \log n \cdot \operatorname{nnz}(\mathbf{A}_i)$, for some constant c' we have

$$\mathbb{P} \left(\left\| (\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i} \right\|_2^2 \geq \mathbb{E} \left\| (\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i} \right\|_2^2 + \log n \cdot \operatorname{nnz}(\mathbf{A}_i) \right)$$

$$\leq \exp\left(\frac{-(\log n)^2/2}{c + c(\log n)/3}\right) \leq \exp(-c' \log n) \leq 1/n^{c'}.$$

Since $\mathbf{A}'_m = \mathbf{H}_m + \mathbf{D}_m$, we have $\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2 \geq \|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2$. Then with probability at least $1 - 1/n^{c'} \geq 1 - \delta$, for any row i with $p_i < 1$, we have

$$\frac{1}{p_i} \cdot \|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 \leq \frac{\text{nnz}(\mathbf{A})}{s \text{nnz}(\mathbf{A}_i)} \cdot c(\log n) \text{nnz}(\mathbf{A}_i) \leq \frac{\epsilon^2 \delta \text{nnz}(\mathbf{A})}{\log n},$$

for $s \geq c \left(\frac{\log^2 n}{\epsilon^2} + \frac{1}{\epsilon^4 \delta^2} \right)$ for large enough c . Observe that, as in Lemma 3 w.l.o.g. we have assumed $1 - n^{c'} \geq 1 - \delta$, since otherwise, our algorithm would read all n^2 entries of the matrix.

Case 2: $p_i = 1$. Then, we have $\text{nnz}(\mathbf{A}_i) \geq \text{nnz}(\mathbf{A})/s$. As in the $p_i < 1$ case, we have from (27):

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2 \text{nnz}(\mathbf{A}_i)}{s \epsilon^2 \delta} + \frac{2c_2 \log^2 n \text{nnz}(\mathbf{A}_i)}{s \epsilon^2}.$$

Now, we observe that $|\mathbf{A}'_{m,i,j}|^2 \leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^2 \leq \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_m\|_2^2 \leq \epsilon^2 \delta \text{nnz}(\mathbf{A})$, which gives us

$$|z_j| \leq \epsilon^2 \delta \text{nnz}(\mathbf{A}) + \frac{2 \text{nnz}(\mathbf{A}_i)}{s \epsilon^2 \delta} + \frac{2c_2 \log^2 n \text{nnz}(\mathbf{A}_i)}{s \epsilon^2}. \quad (32)$$

Thus, for $s \geq c \left(\frac{\log^2 n}{\epsilon^4 \delta} + \frac{1}{\epsilon^4 \delta^2} \right)$ for a large enough constant c and adjusting for other constants we have $|z_j| \leq 2\epsilon^2 \delta \text{nnz}(\mathbf{A})$. Also observe that the expectation of $\sum z_j$ can be bounded by:

$$\mathbb{E} \left(\sum_{j=1}^n z_j \right) = \mathbb{E} \|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 = \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_m\|_2^2 \leq \epsilon^2 \delta \text{nnz}(\mathbf{A}).$$

Next, the variance of the sum of the random variables $\{z_j\}$ can again be bounded by following the analysis presented in and prior to (30) and (31) we have

$$\begin{aligned} \text{Var} \left(\sum_{j=1}^n z_j \right) &\leq \|\mathbf{A}'_{m,i}\|_2^4 + \frac{12c_2 \log^2 n \cdot \text{nnz}(\mathbf{A}_i)^2}{s \epsilon^2} + \frac{12 \text{nnz}(\mathbf{A}_i)^2}{s \epsilon^4 \delta^2} \\ &\leq \epsilon^4 \delta^2 \text{nnz}(\mathbf{A})^2 + \frac{12c_2 \log^2 n \cdot \text{nnz}(\mathbf{A}_i)^2}{s \epsilon^2} + \frac{12 \text{nnz}(\mathbf{A}_i)^2}{s \epsilon^4 \delta^2}, \end{aligned} \quad (33)$$

where we again bound $\|\mathbf{A}'_{m,i,j}\|_2^4$ using

$$|\mathbf{A}'_{m,i,j}|^2 \leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^2 \leq \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'\|_2^2 \leq \epsilon^2 \delta \text{nnz}(\mathbf{A}).$$

Then for $s \geq c(\frac{\log^2 n}{\epsilon^6 \delta^2} + \frac{1}{\epsilon^8 \delta^4})$, we have $\text{Var}\left(\sum_{j=1}^n z_j\right) \leq 2\epsilon^4 \delta^2 \text{nnz}(\mathbf{A})^2$ for large enough constant c .

Using (32) and (33) and noting that $\sum_{j=1}^n \mathbb{E}\left(z_j^2\right) \geq \text{Var}\left(\sum_{j=1}^n z_j\right) - \mathbb{E}^2\left(\sum_{j=1}^n z_j\right)$ we can apply the Bernstein inequality (Theorem 7):

$$\begin{aligned} \mathbb{P}\left(\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \geq \mathbb{E}\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 + t\right) &\leq \mathbb{P}\left(\sum_{j=1}^n z_j \geq \epsilon^2 \delta \text{nnz}(\mathbf{A}_i) + t\right) \\ &\leq \exp\left(\frac{-t^2/2}{c\epsilon^4 \delta^2 \text{nnz}(\mathbf{A})^2 + c\epsilon^2 \delta \text{nnz}(\mathbf{A})t/3}\right). \end{aligned}$$

If we set $t = (\log n)\epsilon^2 \delta \text{nnz}(\mathbf{A})$, then for some constant c' we have

$$\mathbb{P}\left(\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \geq \mathbb{E}\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 + t\right) \leq \exp(-c' \log n) \leq 1/n^{c'}.$$

This, since $\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 \leq \|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2$, when $p_i = 1$, setting $s \geq c(\frac{\log^2 n}{\epsilon^6 \delta^2} + \frac{1}{\epsilon^8 \delta^4})$ for large enough c , we have with probability $\geq 1 - 1/n^{c'}$ $\frac{1}{p_i} \|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 = \|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 \leq \|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \leq (\log n)\epsilon^2 \delta \text{nnz}(\mathbf{A})$.

We thus have, that with probability $\geq 1 - 1/n^{c'}$, for both cases when $p_i < 1$ and $p_i = 1$, $\frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2}{p_i} \leq (\log n)\epsilon^2 \delta \text{nnz}(\mathbf{A})$. Taking a union bound over all $i \in [n]$, with probability at least $1 - 1/n^{c'-1}$, $\max_i \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \leq \sqrt{\log n} \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}$ for $s \geq c(\frac{\log^2 n}{\epsilon^6 \delta^2} + \frac{1}{\epsilon^8 \delta^4})$. As stated before, since $p_i \geq \frac{1}{n^2}$ for all $i \in [n]$, and since $\|\mathbf{H}_m\|_F \leq \|\mathbf{A}\|_F \leq n$, we also have $\max_i \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \leq n^2$. Thus,

$$\begin{aligned} \mathbb{E}_2 \left(\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \right) &\leq \sqrt{\log n} \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})} \left(1 - \frac{1}{n^{c'-1}}\right) \\ &+ \frac{1}{n^{c'-3}} \leq \sqrt{\log n} \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}. \end{aligned}$$

after adjusting ϵ by at most some constants. Overall, we finally get

$$\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \mathbb{E}_2 \left(\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \right) \leq \epsilon \sqrt{\log n} \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}.$$

Plugging this bound into (25), we have for $s \geq c(\frac{\log^2 n}{\epsilon^6 \delta^2} + \frac{1}{\epsilon^8 \delta^4})$,

$$\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq (\log n) \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}.$$

Finally after adjusting ϵ by a $\frac{1}{\log n}$ factor, we have for $s \geq c(\frac{\log^8 n}{\epsilon^6 \delta^2} + \frac{\log^8 n}{\epsilon^8 \delta^4})$ or $s \geq \frac{c \log^8 n}{\epsilon^8 \delta^4}$,

$$\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}.$$

The final bound then follows via Markov's inequality on $\|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2$. $\square$

5.3 Main Accuracy Bound

We are finally ready to prove our main result for sparsity-based sampling, which we restate below.

Theorem 2 (Sparse Matrix Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including the i^{th} index independently with probability $p_i = \min\left(1, \frac{s \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}\right)$ as in Algorithm 2. Here $\text{nnz}(\mathbf{A}_i)$ is the number of non-zero entries in the i^{th} row of $\mathbf{A}$. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, and let $\tilde{\lambda}_i(\mathbf{A})$ be the estimate of $\lambda_i(\mathbf{A})$ computed from $\mathbf{A}_S$ as in Algorithm 2. If $s \geq \frac{c \log^8 n}{\epsilon^8 \delta^4}$, for large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$, $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$.*

Proof With Lemmas 7 and 8 in place, the proof is nearly identical to that of Theorem 1, with the additional need to apply Lemma 5 to show that the eigenvalues of $\mathbf{A}'$ are close to those of $\mathbf{A}$.

For all $i \in [n]$ let $p_i = \min\left(1, \frac{s \text{nnz}(\mathbf{A}_i)}{\text{nnz}(\mathbf{A})}\right)$ and let $\tilde{\mathbf{S}}$ be a scaled diagonal sampling matrix such that the $\tilde{\mathbf{S}}_{ii} = \frac{1}{\sqrt{p_i}}$ with probability p_i and $\tilde{\mathbf{S}}_{ii} = 0$ otherwise. Let $\mathbf{A}'$ be the matrix constructed from $\mathbf{A}$ by zeroing out its elements as described in Lemma 5. Then, note that $\tilde{\mathbf{S}} \mathbf{A}' \tilde{\mathbf{S}} = \mathbf{A}'_S$ where $\mathbf{A}'_S$ is the submatrix constructed as in Algorithm 2. We first show that the eigenvalues of $\mathbf{A}'_S$ approximate those of $\mathbf{A}'$ up to error $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$. The steps are almost identical to those in the proof of Theorem 1. We provide a brief outline of the steps but skip the details.

We split $\mathbf{A}'$ as $\mathbf{A}' = \mathbf{A}'_o + \mathbf{A}'_m$ where $\mathbf{A}'_o$ and $\mathbf{A}'_m$ contain eigenvalues of $\mathbf{A}'$ of magnitudes $< \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}$ and $\geq \epsilon \sqrt{\delta} \sqrt{\text{nnz}(\mathbf{A})}$. This implies $\mathbf{A}'_S = \mathbf{A}'_{o,S} + \mathbf{A}'_{m,S}$ where $\mathbf{A}'_{o,S} = \tilde{\mathbf{S}} \mathbf{A}'_o \tilde{\mathbf{S}}$ and $\mathbf{A}'_{m,S} = \tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}$. By Fact 1 we have that the nonzero eigenvalues of $\mathbf{A}'_{o,S} = \tilde{\mathbf{S}} \mathbf{V}'_o \mathbf{A}'_o \mathbf{V}'_o{}^T \tilde{\mathbf{S}}$ are identical to those of $\mathbf{A}'_o{}^{1/2} \mathbf{V}'_o{}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}} \mathbf{V}'_o \mathbf{A}'_o{}^{1/2}$. Thus, applying the perturbation bound of Fact 4, we have:

$$\left| \lambda_i(\mathbf{A}'_o{}^{1/2} \mathbf{V}'_o{}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}} \mathbf{V}'_o \mathbf{A}'_o{}^{1/2}) - \lambda_i(\mathbf{A}'_o) \right| \leq C \log n \|\mathbf{A}'_o{}^{1/2} \mathbf{V}'_o{}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}} \mathbf{V}'_o \mathbf{A}'_o{}^{1/2} - \mathbf{A}'_o\|_2.$$

From Lemma 7, we get $\|\mathbf{A}'^{1/2} \mathbf{V}_o'^T \bar{\mathbf{S}} \bar{\mathbf{S}} \mathbf{V}_o' \mathbf{A}'^{1/2} - \mathbf{A}'_o\|_2 \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ for $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^3 \sqrt{\delta}}$ with probability at least $1 - \delta$. Thus, setting the error parameter to $\frac{\epsilon}{\log n}$ in Lemma 7, for $s \geq \frac{c \log(1/(\epsilon\delta)) \log^3 n}{\epsilon^3 \sqrt{\delta}}$, with probability at least $1 - \delta$ we have:

$$\left| \lambda_i(\mathbf{A}'^{1/2} \mathbf{V}_o'^T \bar{\mathbf{S}} \bar{\mathbf{S}} \mathbf{V}_o' \mathbf{A}'^{1/2}) - \lambda_i(\mathbf{A}') \right| < \epsilon \sqrt{\text{nnz}(\mathbf{A})}. \quad (34)$$

We have thus shown that the non-zero eigenvalues of $\mathbf{A}'_{o,S}$ approximate all outlying eigenvalues of $\mathbf{A}'$. Note that by Lemma 8, we also have $\|\mathbf{A}'_{m,S}\|_2 \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ with probability at least $1 - \delta$ for $s \geq \frac{c \log^8 n}{\epsilon^8 \delta^4}$. Then, similarly to the section on eigenvalue alignment of Theorem 1, we can argue how these approximations ‘line up’ in the presence of zero eigenvalues in the spectrum of these matrices, concluding that, for all $i \in [n]$,

$$\left| \tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A}') \right| \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}.$$

Finally, by Lemma 5, we have $|\lambda_i(\mathbf{A}') - \lambda_i(\mathbf{A})| \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ for all $i \in [n]$. Thus, via triangle inequality, $\left| \tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A}) \right| \leq 2\epsilon \sqrt{\text{nnz}(\mathbf{A})}$, which gives the required bound after adjusting ϵ to $\epsilon/2$.

Recall that we require $s \geq \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}$ for (34) to hold with probability $1 - \delta$.

We also require $s \geq \frac{c \log^8 n}{\epsilon^8 \delta^4}$ for $\|\mathbf{A}'_{m,S}\|_2 \leq \epsilon \sqrt{\text{nnz}(\mathbf{A})}$ to hold with probability $1 - \delta$ by Lemma 8. Thus, for both conditions to hold simultaneously with probability $1 - 2\delta$ by a union bound, it suffices to set $s = \frac{c \log^8 n}{\epsilon^8 \delta^4} \geq \max\left(\frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}, \frac{c \log^8 n}{\epsilon^8 \delta^4}\right)$, where we use that $\log(1/(\epsilon\delta)) \leq O(\log n)$, as otherwise our algorithm can take $\mathbf{A}_S$ to be the full matrix $\mathbf{A}$. Adjusting δ to $\delta/2$ completes the theorem. $\square$

6 Empirical Evaluation

We complement our theoretical results by evaluating Algorithms 1 (uniform sampling) and Algorithm 2 (sparsity-based sampling) in approximating the eigenvalues of several symmetric matrices. We defer an evaluation of Algorithm 3 (norm-based sampling) to later work. Algorithm 1 and Algorithm 2 perform very well. They seem to have error dependence roughly $1/\epsilon^2$ in practice, as compared to the $1/\epsilon^3$ dependence proven in Theorem 1 and $1/\epsilon^8$ dependence in Theorem 2. Closing the gap between the theory and observed results would be very interesting.

6.1 Datasets

We test Algorithm 1 (uniform sampler) on three dense matrices. We also compare the relative performance of Algorithm 1 and Algorithm 2 (sparsity sampler) on three other synthetic and real world matrices.




Fig. 2 Kong dataset. The image on the left is the original synthetic binary image and the image on the right shows the 5000 sampled points from the outline used as dataset in our experiments

The first two dense matrices, following [49], are created by sampling 5000 points from a binary image. We then normalize all the points in the range $[0, 1]$ in both axes. The original image and resulting set of points are shown in Fig. 2. We then compute a similarity matrix for the points using two common similarity functions used in machine learning and computer graphics: $\delta(\mathbf{x}, \mathbf{y}) = \tanh\left(\frac{\langle \mathbf{x}, \mathbf{y} \rangle}{2}\right)$, the hyperbolic tangent; and $\delta(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^2 \cdot \log(\|\mathbf{x} - \mathbf{y}\|_2^2)$, the thin plane spline. These measures lead to symmetric, indefinite, and entrywise bounded similarity matrices.

Our next dense matrix (called the block matrix) is based on the construction of the hard instance for the lower bound in [11] which shows that we need $\Omega(1/\epsilon^2) \times \Omega(1/\epsilon^2)$ samples to compute ϵn approximations to the eigenvalues of a bounded entry matrix. It is a 5000×5000 matrix containing a 2500×2500 principal submatrix of all 1s, with the rest of the entries set to 0. It has $\lambda_1(\mathbf{A}) = 2500$ and all other eigenvalues equal to 0.

We now describe the three matrices used to compare Algorithm 1 and Algorithm 2. All three are graph adjacency matrices, which are symmetric, indefinite, entrywise bounded and sparse. Spectral density estimation for graph structured matrices is an important primitive in network analysis [19]. The first is a dense Erdős-Rényi graph with 5000 nodes and connection probability 0.1. The second two are real world graphs, taken from SNAP [50]; namely Facebook [51] and Arxiv COND-MAT [52]. The Facebook graph contains 4039 nodes and 88234 directed edges. We symmetrize the adjacency matrix. Arxiv COND-MAT is a collaboration network between authors of Condensed Matter papers published on arXiv, containing 23,133 nodes and 93497 undirected edges. Both these graphs are very sparse—the number of edges is $\leq 1\%$ of the total edges in a complete graph with same number of nodes. The true eigenvalue spectra of each of these matrices are plotted in Fig. 3.

6.2 Implementation Details

Apart from uniform random sampling (Algorithm 1), we also apply the sparsity-based sampling technique in Algorithm 2 and a modification to Algorithm 2, where we do not zero out the elements of the sampled submatrix $\mathbf{A}_S$ (we call this *simple sparsity sampler*). In practice, to apply Algorithm 2, we zero out element $[\mathbf{A}_S]_{i,j}$ (line 5 of




Fig. 3 Eigenspectrum plots. Plots of the eigenspectra of the matrices we studied in Sect. 6. The eigenvalues are ranked in decreasing ordered

Algorithm 2) if $i = j$ or $\text{nnz}(\mathbf{A}_i) \text{nnz}(\mathbf{A}_j) < \frac{\text{nnz}(\mathbf{A})}{c_2 s}$, where c_2 is a constant and s is the size of the sample. We set $c_2 = 0.1$ experimentally as this results in consistent behavior across datasets.

6.3 Experimental Setup

We subsample each matrix and compute its eigenvalues using numpy [53]. We then use our approximation algorithms to estimate the eigenvalues of $\mathbf{A}$ by scaling the eigenvalues of the sampled submatrix. For t trials, we report the logarithm of the average absolute scaled error, $\log \left(\frac{1}{t} \sum \frac{|\tilde{\lambda}_{i,t}(\mathbf{A}) - \lambda_i(\mathbf{A})|}{\sqrt{\text{nnz}(\mathbf{A})}} \right)$, where $\tilde{\lambda}_{i,t}(\mathbf{A})$ is the estimated eigenvalue in the t^{th} trial, $\lambda_i(\mathbf{A})$ is the true eigenvalue and $\text{nnz}(\mathbf{A})$ is the number of non-zero elements in $\mathbf{A}$. Recall that $\sqrt{\text{nnz}(\mathbf{A})} \geq \|\mathbf{A}\|_F$ is an upper bound on all eigenvalue magnitudes. Also note that for the fully dense matrices, $\sqrt{\text{nnz}(\mathbf{A})} \approx n$.

We repeat our experiments for $t = 50$ trials at different sampling rates and aggregate the results. The resultant errors of estimation for dense matrices are plotted in Fig. 4 and for the graph matrices are plotted in Fig. 5. The x-axis is the log proportion of the number of random samples chosen from the matrix. If we sample 1% of the rows/columns, then the log comes to around -4.5 . In these log-log plots, if the sample size has polynomial dependence on ϵ , e.g., ϵn or $\epsilon \sqrt{\text{nnz}(\mathbf{A})}$ error is achieved with sample size proportional to $1/\epsilon^p$, we expect to see error falling off linearly, with slope equal to $-1/p$ where p is the exponent on ϵ .

As a baseline we also show the error if we approximate all eigenvalues with 0 which results in an error of $\frac{\lambda_i}{\sqrt{\text{nnz}(\mathbf{A})}}$. This helps us to observe how the approximation

algorithms perform for both large and small order eigenvalues, as opposed to just approximating everything by 0.

Code. All codes are written in Python and available at https://github.com/archanray/eigenvalue_estimation.

6.4 Summary of Results

Our results are plotted in Figs. 4 and 5. We observe relatively small error in approximating all eigenvalues, with the error decreasing as the number of samples increases. What is more interesting is that the relationship between sample size and error ϵn seems to be generally on the order of $1/\epsilon^2$, our expected lower bound for approximating eigenvalues by randomly sampling a principal submatrix. This can be seen by observing the slope of approximately $-1/2$ on the log-log error plots. In some cases, we do better in approximating small eigenvalues of $\mathbf{A}$ —if the eigenvalue lies well within the range of middle eigenvalues, i.e. $\{-\epsilon n, \epsilon n\}$, we may achieve a very good absolute error estimate simply by approximating it to 0.

As expected, on the graph adjacency matrices (in Fig. 5), sparsity-based sampling techniques generally achieve better error than uniform sampling. For the Erdős-Rényi graph, we expect the node degrees (and hence row sparsities) to be similar. Thus the sampling probability for each row will be roughly uniform, which leads to similar performance of sparsity-based techniques and uniform sampling. For the real world graphs, which have power law degree distributions, sparsity-based sampling techniques has a significant effect. As a result Algorithm 2, and the simple sparsity sampler variant significantly outperform uniform sampling.

Algorithm 2 almost always dominates simple sparsity sampler. In some cases simple sparsity sampler performs better or equivalent to Algorithm 2. This may happen because for two reasons: (1) if Algorithm 2 zeroes out almost all of the sampled submatrix $\mathbf{A}_S$ for small samples, the algorithm will underestimate the corresponding eigenvalue, and (2) the cut-off threshold for the term $\text{nnz}(\mathbf{A}_i) \text{nnz}(\mathbf{A}_j)$ may be too high leading to no difference between simple sparsity sampler and Algorithm 2.

We also observe that approximating all eigenvalues with 0 results in very good approximation for small eigenvalues of the Erdős-Rényi graph. We believe this is because the smaller eigenvalues are significantly less than the largest eigenvalue (of the order of 3500). We see similar trends of approximating eigenvalues with zero for the real world graphs too. But since eigenvalues at the extreme spectrum are of a larger order, we see reasonably good approximation for the sampling algorithms. Algorithm 2 outperforms approximation by 0 in all of these cases.

In the dense matrices uniform sampling almost always outperforms approximation by 0 when estimating any reasonably large eigenvalues. Additionally, note that the block matrix is rank-1 with true eigenvalues $\{2500, 0, \dots, 0\}$. Any sampled principal submatrix will also have rank at most 1. Thus, outside the top eigenvalue, the submatrix will have all zero eigenvalues, which are perfectly approximated by our algorithm. The only non-trivial approximation for this matrix is for the top eigenvalue. This approximation seems to have error dependency around $1/\epsilon^2$, as expected.

7 Conclusion

We present efficient algorithms for estimating all eigenvalues of a symmetric matrix with bounded entries up to additive error ϵn , by reading just a $\text{poly}(\log n, 1/\epsilon) \times \text{poly}(\log n, 1/\epsilon)$ random principal submatrix. We give improved error bounds of $\epsilon\sqrt{\text{nnz}(\mathbf{A})}$ and $\epsilon\|\mathbf{A}\|_F$ when the rows/columns are sampled with probabilities proportional to their sparsities or squared ℓ_2 norms, respectively.

As discussed, our work leaves several open questions. In particular, it is open if our query complexity for $\pm\epsilon n$ approximation can be improved, possibly to $\tilde{O}(\log^c n/\epsilon^4)$ total entries using principal submatrix queries or $\tilde{O}(\log^c/\epsilon^2)$ entries using general queries. The later bound is open even when $\mathbf{A}$ is PSD, a setting where we know that sampling a $O(1/\epsilon^2) \times O(1/\epsilon^2)$ principal submatrix (with $O(1/\epsilon^4)$ total entries) does suffice. Additionally, it is open if we can achieve sample complexity independent of n , by removing all $\log n$ factors, as have been done for the easier problem of testing positive semidefiniteness [11]. See Sect. 1.4 for more details.

It would also be interesting to extend our results to give improved approximation bounds for other properties of the matrix spectrum, such as various Schatten- p norms and spectral summaries. For many of these problems large gaps in understanding exist—e.g., for $\pm n^{3/2}$ approximation to the Schatten-1 norm, which requires $\Omega(n)$ queries, but for which no $o(n^2)$ query algorithm is known. Applying our techniques to improve sublinear time PSD testing algorithms under an ℓ_2 rather than ℓ_∞ approximation requirement [11] would also be interesting. Finally, it would be interesting to identify additional assumptions on $\mathbf{A}$ or on the sampling model where stronger approximation guarantees (e.g., relative error) can be achieved in sublinear time.

Acknowledgements We thank Ainesh Bakshi, Rajesh Jayaram, Anil Damle, and Christopher Musco for helpful conversations about this work. RB, CM and AR was partially supported by an Adobe Research grant, along with NSF Grants 2046235 and 1763618. PD and GD were partially supported by NSF AF 1814041, NSF FRG 1760353, and DOE-SC0022085.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

A Eigenvalue Approximation for PSD Matrices

Here we give a simple proof that shows if Algorithm 1 is used to approximate the eigenvalues of positive semidefinite (PSD) matrices (i.e., with all non-negative eigenvalues) using a $O(1/\epsilon^2) \times O(1/\epsilon^2)$ random submatrix, then the ℓ_2 norm of the error of eigenvalue approximations is bounded by ϵn . This much stronger result immediately implies that each eigenvalue of a PSD matrix can be approximated to $\pm\epsilon n$ additive error using just a $O(1/\epsilon^2) \times O(1/\epsilon^2)$ random submatrix. The proof follows from a bound in [40] which bounds the ℓ_2 norm of the difference vector of eigenvalues of a Hermitian matrix and any other random matrix by the Frobenius norm of the difference of the two matrices. This improves on the bound of Theorem 1 for general




Fig. 4 Approximation error of eigenvalues of dense matrices. Log scale absolute error versus log sampling rate for Algorithm 1 and approximation by 0, as described in Sect. 6.3, for approximating the largest, smallest and fourth largest of three of the example matrices. The corresponding true eigenvalues for each matrix in-order are: (hyperbolic tangent) $\{4.52\text{e}+03, -7.85\text{e}+00, 3.18\text{e}-01\}$, (thin plane spline) $\{3.54\text{e}+02, -1.22\text{e}+03, 1.28\text{e}+02\}$ and (block matrix) $\{2.50\text{e}+03, 0, 0\}$. For the block matrix, all eigenvalues except the largest magnitude one are 0, and so our algorithm perfectly approximates these eigenvalues. Therefore, we only plot the error of approximation for the largest eigenvalue of the block matrix

entrywise bounded matrices by a $1/\epsilon^2$ factor, and matches the $O(1/\epsilon^4)$ lower bound for principal submatrix queries in [11]. Note that the hard instance used to prove the lower bound in [11] can in fact be negated to be PSD, thus showing that our upper bound here is tight.

We first state the result from [40] which we will be using in our proof.

Fact 5 (ℓ_2 -norm bound on eigenvalues [40]) *Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ be Hermitian and $\mathbf{B} \in \mathbb{C}^{n \times n}$ be any matrix whose eigenvalues are $\lambda_1(\mathbf{B}), \dots, \lambda_n(\mathbf{B})$ such that $\text{Re}(\lambda_1(\mathbf{B})) \geq \dots \geq \text{Re}(\lambda_n(\mathbf{B}))$ (where $\text{Re}(\lambda_i(\mathbf{B}))$ denotes the real part of $\lambda_i(\mathbf{B})$). Let $\mathbf{A} - \mathbf{B} = \mathbf{E}$. Then*

$$\left(\sum_{i=1}^n |\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})|^2 \right)^{1/2} \leq \sqrt{2} \|\mathbf{E}\|_F.$$




Fig. 5 Approximation error of eigenvalues of sparse matrices. Log scale absolute error vs. log sampling rate for Algorithms 1, 2, simple sparsity sampler and approximation by 0, as described in Sect. 6.3, for approximating the largest, smallest, and fourth largest of remaining three example matrices. The corresponding true eigenvalues for each matrix in-order are: (Erdős-Rényi) {500.57, −42.52, 42.02}, (Facebook) {162.37, −23.75, 73.28} and (arXiv) {37.95, −15.58, 26.92}

Our result is based on the following Lemma, we prove at the end of the section.

Lemma 9 Consider a PSD matrix $\mathbf{A} = \mathbf{B}\mathbf{B}^T$ with $\|\mathbf{A}\|_\infty \leq 1$. Let S be sampled as in Algorithm 1 for $s \geq \frac{1}{\epsilon^2\delta}$. Let $\tilde{\mathbf{S}} \in \mathbb{R}^{n \times |S|}$ be the scaled sampling matrix satisfying $\tilde{\mathbf{S}}^T \mathbf{A} \tilde{\mathbf{S}} = \frac{n}{s} \cdot \mathbf{A}_S$. Then with probability at least $1 - \delta$,

$$\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F \leq \epsilon n.$$

From the above Lemma we have:

Corollary 2 (Spectral norm bound—PSD matrices) Consider a PSD matrix $\mathbf{A}$ with $\|\mathbf{A}\|_\infty \leq 1$. Let S be a subset of indices formed by including each index in $[n]$ independently with probability s/n as in Algorithm 1. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, with eigenvalues $\lambda_1(\mathbf{A}_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}_S)$.

For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all other $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$. Then if $s \geq \frac{2}{\epsilon^2 \delta}$, with probability at least $1 - \delta$,

$$\left(\sum_{i=1}^n \left| \tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A}) \right|^2 \right)^{1/2} \leq \epsilon n,$$

which implies that for all $i \in [n]$,

$$\lambda_i(\mathbf{A}) - \epsilon n \leq \tilde{\lambda}_i(\mathbf{A}) \leq \lambda_i(\mathbf{A}) + \epsilon n.$$

Proof Let S be sampled as in Algorithm 1 and let $\tilde{\mathbf{S}} \in \mathbb{R}^{n \times |S|}$ be the scaled sampling matrix satisfying $\tilde{\mathbf{S}}^T \mathbf{A} \tilde{\mathbf{S}} = \frac{n}{s} \cdot \mathbf{A}_S$. Since $\mathbf{A}$ is PSD, we can write $\mathbf{A} = \mathbf{B} \mathbf{B}^T$ for some matrix $\mathbf{B} \in \mathbb{R}^{n \times \text{rank}(\mathbf{A})}$. From Lemma 9, for $s \geq \frac{1}{\epsilon^2 \delta}$, we have with probability at least $1 - \delta$:

$$\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F \leq \epsilon n$$

Using Fact 5, we have,

$$\left(\sum_{i=1}^{\text{rank}(\mathbf{A})} \left| \lambda_i(\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B}) - \lambda_i(\mathbf{B}^T \mathbf{B}) \right|^2 \right)^{1/2} \leq \sqrt{2} \|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F \leq \sqrt{2} \epsilon n. \quad (35)$$

Also from Fact 1, we have $\lambda_i(\mathbf{B}^T \mathbf{B}) = \lambda_i(\mathbf{B} \mathbf{B}^T) = \lambda_i(\mathbf{A})$ for all $i \leq \text{rank}(\mathbf{A})$. Thus,

$$\left(\sum_{i=1}^{\text{rank}(\mathbf{A})} \left| \lambda_i(\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B}) - \lambda_i(\mathbf{A}) \right|^2 \right)^{1/2} \leq \sqrt{2} \epsilon n$$

Also by Fact 1, all non-zero eigenvalues of $\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B}$ are equal to those of $\tilde{\mathbf{S}}^T \mathbf{B} \mathbf{B}^T \tilde{\mathbf{S}} = \frac{n}{s} \cdot \mathbf{A}_S$. All other eigenvalue estimates are set to 0. Further, for all $i > \text{rank}(\mathbf{A})$, $\lambda_i(\mathbf{A}) = 0$. Thus,

$$\left(\sum_{i=1}^n \left| \tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A}) \right|^2 \right)^{1/2} \leq \sqrt{2} \epsilon n.$$

Adjusting ϵ to $\epsilon/\sqrt{2}$ then gives us the bound. $\square$

We now prove Lemma 9, using a standard approach for sampling based approximate matrix multiplication—see e.g. [44].

Proof of Lemma 9 For $k = 1, \dots, n$ let $\mathbf{Y}_k = \frac{n}{s} - 1$ with probability $\frac{s}{n}$ and $\mathbf{Y}_k = -1$ with probability $1 - \frac{s}{n}$. Thus $\mathbb{E}[\mathbf{Y}_k] = 0$ and

$$\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F^2 = \sum_{i=1}^n \sum_{j=1}^n \left(\sum_{k=1}^n \mathbf{Y}_k \cdot \mathbf{B}_{ik} \mathbf{B}_{jk} \right)^2.$$

Fixing i, j , each the $\mathbf{Y}_k \cdot \mathbf{B}_{ik} \mathbf{B}_{jk}$ are 0 mean independent random variables. Thus we have:

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F^2 \right] &= \sum_{i=1}^n \sum_{j=1}^n \mathbb{E} \left[\left(\sum_{k=1}^n \mathbf{Y}_k \cdot \mathbf{B}_{ik} \mathbf{B}_{jk} \right)^2 \right] \\ &= \sum_{i=1}^n \sum_{j=1}^n \text{Var} \left[\sum_{k=1}^n \mathbf{Y}_k \cdot \mathbf{B}_{ik} \mathbf{B}_{jk} \right] \\ &= \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \text{Var} [\mathbf{Y}_k \cdot \mathbf{B}_{ik} \mathbf{B}_{jk}] \\ &\leq \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \frac{n}{s} \cdot \mathbf{B}_{ik}^2 \mathbf{B}_{jk}^2. \end{aligned}$$

since $\text{Var}[\mathbf{Y}_k] = \left(\frac{n}{s} - 1\right)^2 \cdot \frac{s}{n} + \left(1 - \frac{s}{n}\right) = \frac{n}{s} - 2 + \frac{s}{n} + 1 - \frac{s}{n} = \frac{n}{s} - 1$. Rearranging the sums we have:

$$\mathbb{E}[\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F^2] \leq \frac{n}{s} \sum_{k=1}^n \sum_{i=1}^n \mathbf{B}_{ik}^2 \sum_{j=1}^n \mathbf{B}_{jk}^2.$$

Observe that $\sum_{j=1}^n \mathbf{B}_{jk}^2 = \mathbf{A}_{kk} \leq \|\mathbf{A}\|_\infty \leq 1$, thus overall we have:

$$\mathbb{E}[\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F^2] \leq \frac{n^2}{s} \leq \epsilon^2 \delta n^2.$$

So by Markov's inequality, with probability $\geq 1 - \delta$, $\|\mathbf{B}^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{B} - \mathbf{B}^T \mathbf{B}\|_F^2 \leq \epsilon^2 n^2$. This completes the theorem after taking a square root. $\square$

Remark: The proof of Lemma 9 can be easily modified to show that the i^{th} row of $\mathbf{A}$ can be sampled with probability proportional to $\frac{|\mathbf{A}_{ii}|}{\text{tr}(\mathbf{A})}$ to approximate the eigenvalues of any PSD $\mathbf{A}$ up to $\pm \epsilon \cdot \text{tr}(\mathbf{A})$ error ($\text{tr}(\mathbf{A})$ is the trace of $\mathbf{A}$). When sampling with probability proportional to $\frac{|\mathbf{A}_{ii}|}{\text{tr}(\mathbf{A})}$, we do not require a bounded entry assumption on $\mathbf{A}$.

B Alternate Bound for Uniform Sampling

In this section we provide an alternate bound for approximating eigenvalues with uniform sampling. The sample complexity is worse by a factor of $1/\epsilon$ for this approach, but better by a factor $\log^2 n$ as compared to Theorem 1. We start with an analog to Lemma 3, showing that the outlying eigenspace remains nearly orthogonal after sampling. In particular, we show concentration of the Hermitian matrix $\mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o$ about its expectation $\mathbf{V}_o^T \mathbf{V}_o = \mathbf{I}$ rather than the non-Hermitian $\mathbf{A}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{A}_o^{1/2}$ as in Lemma 3. This allows us to use Weyl's inequality in our final analysis, rather than the non-Hermitian eigenvalue perturbation bound of Fact 4, saving a $\log^2 n$ factor in the sample complexity.

Lemma 10 (Near orthonormality—sampled outlying eigenvalues) *Let S be sampled as in Algorithm 1 for $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^4 \delta}$ where c is a sufficiently large constant. Let $\tilde{\mathbf{S}} \in \mathbb{R}^{n \times |S|}$ be the scaled sampling matrix satisfying $\tilde{\mathbf{S}}^T \mathbf{A} \tilde{\mathbf{S}} = \frac{n}{s} \cdot \mathbf{A}_S$. Then with probability at least $1 - \delta$, $\|\mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o - \mathbf{I}\|_2 \leq \epsilon$.*

Proof The result is standard in randomized numerical linear algebra—see e.g., [55]. For completeness, we give a proof here. Define $\mathbf{E} = \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o - \mathbf{I}$. For all $i \in [n]$, let $\mathbf{V}_{o,i}$ be the i^{th} row of $\mathbf{V}_o$ and define the matrix valued random variable

$$\mathbf{Y}_i = \begin{cases} \frac{n}{s} \mathbf{V}_{o,i} \mathbf{V}_{o,i}^T, & \text{with probability } s/n \\ 0 & \text{otherwise.} \end{cases}$$

Then, similar to the proof of Lemma 3, define $\mathbf{Q}_i = \mathbf{Y}_i - \mathbb{E}[\mathbf{Y}_i]$. Since $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_n$ are independent random variables and $\sum_{i=1}^n \mathbf{Q}_i = \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o - \mathbf{I} = \mathbf{E}$, we need to bound $\|\mathbf{Q}_i\|_2$ for all $i \in [n]$ and $\text{Var}(\mathbf{E}) \stackrel{\text{def}}{=} \mathbb{E}(\mathbf{E} \mathbf{E}^T) = \mathbb{E}(\mathbf{E}^T \mathbf{E}) = \sum_{i=1}^n \mathbb{E}[\mathbf{Q}_i^2]$. Observe $\|\mathbf{Q}_i\|_2 \leq \max(1, \frac{n}{s} - 1) \|\mathbf{V}_{o,i} \mathbf{V}_{o,i}^T\|_2 = \max(1, \frac{n}{s} - 1) \|\mathbf{V}_{o,i}\|_2^2 \leq \frac{1}{\epsilon^2 \delta s}$, by row norm bounds of Lemma 2. Again, using Lemma 2 we have

$$\begin{aligned} \sum_{i=1}^n \mathbb{E}[\mathbf{Q}_i^2] &= \sum_{i=1}^n \frac{s}{n} \cdot \left(\frac{n}{s} - 1\right)^2 (\mathbf{V}_{o,i} \mathbf{V}_{o,i}^T)^2 + \left(1 - \frac{s}{n}\right) (\mathbf{V}_{o,i} \mathbf{V}_{o,i}^T)^2 \\ &\leq \sum_{i=1}^n \frac{n}{s} \|\mathbf{V}_{o,i}\|_2^2 (\mathbf{V}_{o,i} \mathbf{V}_{o,i}^T) \\ &\leq \sum_{i=1}^n \frac{n}{s} \frac{1}{\epsilon^2 \delta n} (\mathbf{V}_{o,i} \mathbf{V}_{o,i}^T) \\ &\leq \frac{1}{s \epsilon^2 \delta} \cdot \mathbf{I} \end{aligned}$$

where $\mathbf{I}$ is the identity matrix of appropriate dimension. By setting $d = \frac{1}{\epsilon^2 \delta}$, we can finally bound the probability of the event $\|\mathbf{E}\|_2 \geq \epsilon n$ using Theorem 6 (the matrix Bernstein inequality) with δ if $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^4 \delta}$. Since these steps follow Lemma 3 nearly exactly, we omit them here. $\square$

With Lemma 10 in place, we can now give our alternate sample complexity bound.

Theorem 8 (Sublinear Time Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including each index independently with probability s/n as in Algorithm 1. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, with eigenvalues $\lambda_1(\mathbf{A}_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}_S)$.*

For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all $i \in \{1, \dots, |S|\}$ with $\lambda_i(\mathbf{A}_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all other $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$. If $s \geq \frac{c \log n}{\epsilon^4 \delta}$, for a large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$,

$$\lambda_i(\mathbf{A}) - \epsilon n \leq \tilde{\lambda}_i(\mathbf{A}) \leq \lambda_i(\mathbf{A}) + \epsilon n.$$

Proof Let $\mathbf{S} \in \mathbb{R}^{n \times |S|}$ be the binary sampling matrix with a single one in each column such that $\mathbf{S}^T \mathbf{A} \mathbf{S} = \mathbf{A}_S$. Let $\tilde{\mathbf{S}} = \sqrt{n/s} \cdot \mathbf{S}$. Following Definition 1.1, we write $\mathbf{A} = \mathbf{A}_o + \mathbf{A}_m$. By Fact 1 we have that the nonzero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,S} = \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{A}_o \mathbf{V}_o^T \tilde{\mathbf{S}}$ are identical to those of $\mathbf{A}_o \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o$.

Note that $\mathbf{H} = \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o$ is positive semidefinite. Writing its eigendecomposition $\mathbf{H} = \mathbf{U} \mathbf{W} \mathbf{U}^T$ we can define the matrix squareroot $\mathbf{H}^{1/2} = \mathbf{U} \mathbf{W}^{1/2} \mathbf{U}^T$ with $\mathbf{H}^{1/2} \mathbf{H}^{1/2} = \mathbf{H}$. By Lemma 10 applied with error $\epsilon/6$, with probability at least $1 - \delta$, all eigenvalues of $\mathbf{H}$ lie in the range $[1 - \epsilon/6, 1 + \epsilon/6]$. In turn, all eigenvalues of $\mathbf{H}^{1/2}$ also lie in this range. Again using Fact 1, we have that the nonzero eigenvalues of $\mathbf{A}_o \mathbf{H}$, and in turn those of $\frac{n}{s} \cdot \mathbf{A}_{o,S}$, are identical to those of $\mathbf{H}^{1/2} \mathbf{A}_o \mathbf{H}^{1/2}$.

Let $\mathbf{E} = \mathbf{H}^{1/2} - \mathbf{I} = \mathbf{U} \mathbf{W}^{1/2} \mathbf{U}^T - \mathbf{U} \mathbf{U}^T = \mathbf{U} (\mathbf{W}^{1/2} - \mathbf{I}) \mathbf{U}^T$. Since the diagonal entries of $\mathbf{W}^{1/2}$ lie in $[1 - \epsilon/6, 1 + \epsilon/6]$, those of $\mathbf{W}^{1/2} - \mathbf{I}$ lie in $[-\epsilon/6, \epsilon/6]$. Thus, $\|\mathbf{E}\|_2 \leq \epsilon/6$. We can write

$$\lambda_i(\mathbf{H}^{1/2} \mathbf{A}_o \mathbf{H}^{1/2}) = \lambda_i((\mathbf{I} + \mathbf{E}) \mathbf{A}_o (\mathbf{I} + \mathbf{E})) = \lambda_i(\mathbf{A}_o + \mathbf{E} \mathbf{A}_o + \mathbf{A}_o \mathbf{E} + \mathbf{E} \mathbf{A}_o \mathbf{E}).$$

We can then bound

$$\begin{aligned} \|\mathbf{E} \mathbf{A}_o + \mathbf{A}_o \mathbf{E} + \mathbf{E} \mathbf{A}_o \mathbf{E}\|_2 &\leq \|\mathbf{E} \mathbf{A}_o\|_2 + \|\mathbf{A}_o \mathbf{E}\|_2 + \|\mathbf{E} \mathbf{A}_o \mathbf{E}\|_2 \\ &\leq \|\mathbf{E}\|_2 \|\mathbf{A}_o\|_2 + \|\mathbf{A}_o\|_2 \|\mathbf{E}\|_2 + \|\mathbf{E}\|_2 \|\mathbf{A}_o\|_2 \|\mathbf{E}\|_2 \\ &\leq \epsilon n/6 + n\epsilon/6 + \epsilon^2 n/36 \\ &\leq \epsilon/2 \cdot n. \end{aligned}$$

Applying Weyl's eigenvalue perturbation theorem (Fact 3), we thus have for all i ,

$$\left| \lambda_i(\mathbf{H}^{1/2} \mathbf{A}_o \mathbf{H}^{1/2}) - \lambda_i(\mathbf{A}_o) \right| < \epsilon/2 \cdot n. \quad (36)$$

Note that we have shown that the nonzero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,S}$ are identical to those of $\mathbf{H}^{1/2} \mathbf{A}_o \mathbf{H}^{1/2}$, which we have shown well approximate those of $\mathbf{A}_o$ and in turn $\mathbf{A}_o$ i.e., the non-zero eigenvalues of $\frac{n}{s} \cdot \mathbf{A}_{o,S}$ approximate all outlying eigenvalues of $\mathbf{A}$.

We can also bound the middle eigenvalues using Lemma 4 as in Theorem 1. Now the only thing left is to argue how these approximations ‘line up’ in the presence of zero eigenvalues in the spectrum of these matrices. This part of the proof again proceeds similarly to that of Theorem 1 in Sect. 4.2.

Analogous to Theorem 1, from Lemma 10 equation (36) holds with probability $1 - \delta$ if $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^4 \delta}$. We also require $s \geq \frac{c \log n}{\epsilon^2 \delta}$ for $\|\mathbf{A}_{m,S}\|_2 \leq \epsilon n$ to hold with probability $1 - \delta$ by Lemma 4. Thus, for both conditions to hold simultaneously with probability $1 - 2\delta$ by a union bound, it suffices to set $s = \frac{c \log n}{\epsilon^4 \delta} \geq \max\left(\frac{c \log(1/(\epsilon\delta))}{\epsilon^4 \delta}, \frac{c \log n}{\epsilon^2 \delta}\right)$, where we use that $\log(1/(\epsilon\delta)) = O(\log n)$, as otherwise our algorithm can take $\mathbf{A}_S$ to be all of $\mathbf{A}$. Adjusting δ to $\delta/2$ completes the theorem. $\square$

C Refined Bounds

In this section, we show how it is possible to get better query complexity or tighter approximation factors by modifying the proof of Theorem 1 and Lemmas 3 and 2 under some assumptions. We give an extension to Theorem 1 in Theorem 9 for the case when the eigenvalues of $\mathbf{A}_o$ lie in a bounded range—between $\epsilon^a \sqrt{\delta} n$ and $\epsilon^b n$ where $0 \leq b \leq a \leq 1$.

Theorem 9 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $\mathbf{A}_o$ be as in Definition 1.1 such that for all eigenvalues $\lambda_i(\mathbf{A}_o)$ we have either $\epsilon^a \sqrt{\delta} n \leq |\lambda_i(\mathbf{A}_o)| \leq \epsilon^b n$ or $\lambda_i(\mathbf{A}_o) = 0$ where $0 \leq b \leq a \leq 1$. Let $S \subseteq [n]$ be formed by including each index independently with probability s/n as in Algorithm 1. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, with eigenvalues $\lambda_1(\mathbf{A}_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}_S)$.*

For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \frac{n}{s} \cdot \lambda_i(\mathbf{A}_S)$. For all other $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$. If $s \geq \frac{c \log(1/(\epsilon\delta)) \log^{2+a-b} n}{\epsilon^{2+a-b} \delta}$, for large enough c , then with probability at least $1 - \delta$, for all $i \in [n]$,

$$\lambda_i(\mathbf{A}) - \epsilon n \leq \tilde{\lambda}_i(\mathbf{A}) \leq \lambda_i(\mathbf{A}) + \epsilon n.$$

Proof The proof follows by modifying the proofs of Theorem 1, Lemmas 2 and 3 to account for the tighter intervals. First observe that since $|\lambda_i(\mathbf{A}_o)| \leq \epsilon^a \sqrt{\delta} n$ for all i , we can give a tighter row norm bound for $\mathbf{V}_o$ from the proof of Lemma 2. In particular, from equation (3) we get:

$$\|\mathbf{A}_o^{1/2} \mathbf{V}_{o,i}\|_2^2 \leq \frac{1}{\epsilon^a \sqrt{\delta}} \quad \text{and} \quad \|\mathbf{V}_{o,i}\|_2^2 \leq \frac{n}{\epsilon^{2a} \delta n^2} = \frac{1}{\epsilon^{2a} \delta n}.$$

We can then bound the number of samples we need to take such that for $\mathbf{A}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{A}_o^{1/2}$ (as defined in Theorem 8) we have $\|\mathbf{A}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{A}_o^{1/2} - \mathbf{A}_o\|_2 \leq \epsilon n$ with probability at least $1 - \delta$ via a matrix Bernstein bound. By appropriately modifying the proof of Lemma 3 to incorporate the stronger row norm bound

for $\mathbf{V}_o$, we can show that sampling with probability s/n for $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^{2+a-b}\sqrt{\delta}}$ for large enough c suffices. Specifically, we get $L \leq \frac{n}{\epsilon^a \sqrt{\delta s}}$, $v \leq \frac{n^2}{\epsilon^{a-b} \sqrt{\delta s}}$ and $d \leq \log(1/(\epsilon^2 \delta))$ for the Bernstein bound in Lemma 3 which enables us to get the tighter bound. Thus, we have $\|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2} - \mathbf{\Lambda}_o\|_2 \leq \epsilon n$ with probability $1 - \delta$ for $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^{2+a-b} \sqrt{\delta}}$ following Lemma 3. We also require $s \geq \frac{c \log n}{\epsilon^2 \delta}$ for $\|\mathbf{A}_{m,s}\|_2 \leq \epsilon n$ to hold with probability $1 - \delta$ by Lemma 4. Then, following the proof of Theorem 1, by Fact 4, for all $i \in [n]$, and some constant C , we have:

$$|\lambda_i(\mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2}) - \lambda_i(\mathbf{\Lambda}_o)| \leq C \log n \|\mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2} - \mathbf{\Lambda}_o\|_2.$$

As in the proof of Theorem 1, adjusting ϵ by a $\frac{1}{C \log n}$ factor, we get $|\lambda_i(\mathbf{\Lambda}_o^{1/2} \mathbf{V}_o^T \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T \mathbf{V}_o \mathbf{\Lambda}_o^{1/2}) - \lambda_i(\mathbf{\Lambda}_o)| \leq \epsilon n$ with probability $1 - \delta$ for $s \geq \frac{c \log(1/(\epsilon\delta)) \log^{2+a-b} n}{\epsilon^{2+a-b} \sqrt{\delta}}$. Then we follow the proof of Theorem 1 to align the eigenvalues completing the proof. $\square$

D Spectral Norm Bounds for Non-uniform Random Submatrices

In this section, we prove an extension of Theorem 4 to the setting where rows are sampled non-uniformly. This allows us to show that the middle eigenvalues do not grow too large when sampling and rescaling with the non-uniform sampling probabilities used by Algorithms 2 and 3.

Theorem 5 (Non-uniform column sampling—spectral norm bound) *Let $\mathbf{A}$ be an $m \times n$ matrix with rank r . Let δ_j be a sequence of independent random variables such that $\delta_j = \frac{1}{\sqrt{p_j}}$ with probability p_j and 0 otherwise. Let $\mathbf{S}$ be a square diagonal sampling matrix with j^{th} diagonal entry set to δ_j .*

$$\mathbb{E}_2 \|\mathbf{AS}\|_2 \leq 5\sqrt{\log r} \cdot \mathbb{E}_2 \|\mathbf{AS}\|_{1 \rightarrow 2} + \|\mathbf{A}\|_2$$

Proof The proof follows from [46]. We begin by first defining the following term

$$E := \mathbb{E}_2 \|\mathbf{AS}\|_2.$$

Now we have

$$E^2 = \mathbb{E} \|\mathbf{AS}\|_2^2 = \mathbb{E} \|\mathbf{ASSA}^*\|_2 = \mathbb{E} \left\| \sum_{j=1}^n \delta_j^2 \mathbf{A}_j \mathbf{A}_j^* \right\|_2,$$

where δ_j is the sequence of independent random variables such that $\delta_j = \frac{1}{\sqrt{p_j}}$ with probability p_j and 0 otherwise, and $\mathbf{A}_j$ is the j^{th} column of $\mathbf{A}$. Then, $\mu_j = \mathbb{E}[(\delta_j)^2] = 1$. Let $\{\delta'_j\}$ be an independent copy of the sequence $\{\delta_j\}$. Subtracting the mean and applying triangle inequality we have

$$E^2 \leq \mathbb{E} \left\| \sum_{j=1}^n (\delta_j^2 - \mathbb{E}[(\hat{\delta})^2]) \mathbf{A}_j \mathbf{A}_j^* \right\|_2 + \left\| \sum_{j=1}^n \mathbf{A}_j \mathbf{A}_j^* \right\|_2.$$

Using Jensen's inequality we have

$$E^2 \leq \mathbb{E} \left\| \sum_{j=1}^n (\delta_j^2 - (\delta'_j)^2) \mathbf{A}_j \mathbf{A}_j^* \right\|_2 + \|\mathbf{A} \mathbf{A}^*\|_2.$$

The random variables $(\delta_j^2 - (\delta'_j)^2)$ are symmetric and independent. Let $\{\epsilon_j\}$ be i.i.d Rademacher random variables for all $j \in [n]$. Then applying the standard symmetrization argument followed by triangle inequality, we have:

$$E^2 \leq 2\mathbb{E} \left\| \sum_{j=1}^n \epsilon_j \delta_j^2 \mathbf{A}_j \mathbf{A}_j^* \right\|_2 + \|\mathbf{A} \mathbf{A}^*\|_2.$$

Let $\Omega = \{j : \delta_j = \frac{1}{\sqrt{p_j}}\}$. Let $\mathbb{E}_\epsilon$ be the partial expectation with respect to $\{\epsilon_j\}$ while keeping the other random variables fixed, and let $\mathbb{E}_\Omega$ be the partial expectation with respect to Ω . Then, we get:

$$E^2 \leq 2\mathbb{E}_\Omega \left[\mathbb{E}_\epsilon \left\| \sum_{\Omega} \epsilon_j \delta_j^2 \mathbf{A}_j \mathbf{A}_j^T \right\|_2 \right] + \|\mathbf{A}\|_2^2.$$

Using Rudelson's Lemma 11 of [46] for any matrix $\mathbf{X}$ with columns $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ and any $q = 2 \log n$ we have

$$\left(\mathbb{E} \left\| \sum_{j=1}^n \epsilon_j \mathbf{x}_j \mathbf{x}_j^* \right\|_2^q \right)^{1/q} \leq 1.5\sqrt{q} \|\mathbf{X}\|_{1 \rightarrow 2} \|\mathbf{X}\|_2.$$

Since $(\cdot)^{1/q}$ is concave for $q \geq 1$, using Jensen's inequality we get:

$$\mathbb{E} \left\| \sum_{j=1}^n \epsilon_j \mathbf{x}_j \mathbf{x}_j^* \right\|_2 \leq 1.5\sqrt{q} \|\mathbf{X}\|_{1 \rightarrow 2} \|\mathbf{X}\|_2$$

Applying the above result to the matrix $\mathbf{A}\mathbf{S}$, we get:

$$E^2 \leq 3\sqrt{q} [\mathbb{E}(\|\mathbf{A}\mathbf{S}\|_{1 \rightarrow 2} \|\mathbf{A}\mathbf{S}\|_2)] + \|\mathbf{A}\|_2^2.$$

Applying Cauchy Schwartz we get:

$$E^2 \leq 3\sqrt{q} (\mathbb{E} \|\mathbf{A}\mathbf{S}\|_{1 \rightarrow 2}^2)^{1/2} (\mathbb{E} \|\mathbf{A}\mathbf{S}\|_2^2)^{1/2} + \|\mathbf{A}\|_2^2.$$

The above equation is of the form $E^2 \leq bE + c$. Thus, the values of E from which the above equation is true is given by $E \leq \frac{b + \sqrt{b^2 + 4c}}{2} \leq b + \sqrt{c}$. Thus, we get:

$$\mathbb{E}_2 \|\mathbf{AS}\|_2 \leq 3\sqrt{q} \mathbb{E}_2 \|\mathbf{AS}\|_{1 \rightarrow 2} + \|\mathbf{A}\|_2.$$

This gives us the final bound. $\square$

E Improved Bounds via Row-Norm-Based Sampling

Building on the sparsity-based sampling results presented in Sect. 5, we now show how to obtain improved approximation error of $\pm\epsilon \|\mathbf{A}\|_F$ assuming we can sample the rows of $\mathbf{A}$ with probabilities proportional to their squared ℓ_2 norms. The ability to sample by norms also allows us to remove the assumption that $\mathbf{A}$ has bounded entries—our results apply to any symmetric matrix.

For technical reasons, we mix row norm sampling with uniform sampling, forming a random principal submatrix by sampling each index $i \in [n]$ independently with probability $p_i = \min\left(1, \frac{s\|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{n^2}\right)$ and rescaling each sampled row/column by $1/\sqrt{p_i}$. As in the sparsity-based sampling setting, we must carefully zero out entries of the sampled submatrix to ensure concentration of the sampled eigenvalues. Pseudocode for the full algorithm is given in Algorithm 3.

E.1 Preliminary Lemmas

Our proof closely follows that of Theorem 2 in Sect. 5. We start by defining $\mathbf{A}' \in \mathbb{R}^{n \times n}$ obtained by zeroing out entries of $\mathbf{A}$ as described in Algorithm 3. We have $\mathbf{A}'_{ij} = 0$ whenever 1) $i = j$ and $\|\mathbf{A}_i\|_2^2 < \frac{\epsilon^2}{4} \|\mathbf{A}\|_F^2$ or 2) $i \neq j$ and $\|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 < \frac{\epsilon^2 \|\mathbf{A}\|_F^2 \cdot |\mathbf{A}_{ij}|^2}{c_2 \log^4 n}$. Otherwise $\mathbf{A}'_{ij} = \mathbf{A}_{ij}$. Similar to the sparsity sampling case, we argue that the eigenvalues of $\mathbf{A}'$ are close to $\mathbf{A}$ i.e., zeroing out entries of $\mathbf{A}$ according to the given condition doesn't change its eigenvalues by too much (Lemma 11). Then, we again split $\mathbf{A}' = \mathbf{A}'_o + \mathbf{A}'_m$ such that $\|\mathbf{A}'_m\|_2 \leq \epsilon\sqrt{\delta} \|\mathbf{A}\|_F$. We argue that after sampling, we have $\|\mathbf{A}'_{m,S}\|_2 \leq \epsilon \|\mathbf{A}\|_F$ and the eigenvalues of $\mathbf{A}'_{o,S}$ approximate those of $\mathbf{A}'_o$ up to $\pm\epsilon \|\mathbf{A}\|_F$ error.

Lemma 11 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric. Let $\mathbf{A}' \in \mathbb{R}^{n \times n}$ have $\mathbf{A}'_{ij} = 0$ if either 1) $i = j$ and $\|\mathbf{A}_i\|_2^2 < \frac{\epsilon^2}{4} \|\mathbf{A}\|_F^2$ or 2) $i \neq j$ and $\|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 < \frac{\epsilon^2 \|\mathbf{A}\|_F^2 \cdot |\mathbf{A}_{ij}|^2}{c_2 \log^4 n}$ for a sufficiently large constant c_2 . Otherwise, $\mathbf{A}'_{ij} = \mathbf{A}_{ij}$. Then, for all $i \in [n]$,*

$$|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{A}')| \leq \epsilon \|\mathbf{A}\|_F.$$

Proof Consider the matrix $\mathbf{A}''$, which is defined identically to $\mathbf{A}'$ except we only set $\mathbf{A}''_{ij} = 0$ if $i \neq j$ and $\|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 < \frac{\epsilon^2 \|\mathbf{A}\|_F^2 \cdot |\mathbf{A}_{ij}|^2}{c_2 \log^4 n}$. That is, we do not zero out

Algorithm 3 Eigenvalue estimator using ℓ_2 norm-based sampling

- 1: **Input:** Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, Accuracy $\epsilon \in (0, 1)$, failure prob. $\delta \in (0, 1)$. $\|\mathbf{A}_i\|_2$ for all $i \in [n]$.
- 2: Fix $s = \frac{c_1 \log^{10} n}{\epsilon^8 \delta^4}$ where c_1 is a sufficiently large constant.
- 3: Add each $i \in [n]$ to sample set S independently, with probability $p_i = \min \left(1, \frac{s \|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{n^2} \right)$. Let the principal submatrix of $\mathbf{A}$ corresponding to S be $\mathbf{A}_S$.
- 4: Let $\mathbf{A}_S = \mathbf{D} \mathbf{A}_S \mathbf{D}$ where $\mathbf{D} \in \mathbb{R}^{|S| \times |S|}$ is diagonal with $\mathbf{D}_{i,i} = \frac{1}{\sqrt{p_j}}$ if the i^{th} element of S is j .
- 5: Construct $\mathbf{A}'_S \in \mathbb{R}^{|S| \times |S|}$ from $\mathbf{A}_S$ as follows:

$$[\mathbf{A}'_S]_{i,j} = \begin{cases} 0 & \text{if } i = j \text{ and } \|\mathbf{A}_i\|_2^2 < \frac{\epsilon^2}{4} \|\mathbf{A}\|_F^2 \\ 0 & \text{if } i \neq j \text{ and } \|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 < \frac{\epsilon^2 \|\mathbf{A}\|_F^2 \cdot |\mathbf{A}_{ij}|^2}{c_2 \log^4 n} \text{ for sufficient large constant } c_2 \\ [\mathbf{A}_S]_{i,j} & \text{otherwise.} \end{cases}$$

- 6: Compute the eigenvalues of $\mathbf{A}'_S$: $\lambda_1(\mathbf{A}'_S) \geq \dots \geq \lambda_{|S|}(\mathbf{A}'_S)$.
- 7: For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}'_S) \geq 0$, let $\tilde{\lambda}_i(\mathbf{A}) = \lambda_i(\mathbf{A}'_S)$. For all $i \in [|S|]$ with $\lambda_i(\mathbf{A}'_S) < 0$, let $\tilde{\lambda}_{n-(|S|-i)}(\mathbf{A}) = \lambda_i(\mathbf{A}'_S)$. For all remaining $i \in [n]$, let $\tilde{\lambda}_i(\mathbf{A}) = 0$.
- 8: **Return:** Eigenvalue estimates $\tilde{\lambda}_1(\mathbf{A}) \geq \dots \geq \tilde{\lambda}_n(\mathbf{A})$.

any entries on the diagonal as in $\mathbf{A}'$. We will show that $\|\mathbf{A} - \mathbf{A}''\|_2 \leq \frac{\epsilon}{2} \|\mathbf{A}\|_F$. If $\mathbf{A}_{ii}$ is zeroed out in $\mathbf{A}'$ this implies that $\|\mathbf{A}_i\|_2^2 < \frac{\epsilon^2}{4} \|\mathbf{A}\|_F^2$. Thus, $|\mathbf{A}_{ii}| \leq \|\mathbf{A}_i\|_2 \leq \frac{\epsilon}{2} \|\mathbf{A}\|_F$ and so $\|\mathbf{A}'' - \mathbf{A}'\|_2 \leq \frac{\epsilon}{2} \|\mathbf{A}\|_F$. So, by triangle inequality, we will then have $\|\mathbf{A} - \mathbf{A}'\|_2 \leq \epsilon \cdot \|\mathbf{A}\|_F$. The lemma then follows from Weyl's inequality

To show that $\|\mathbf{A} - \mathbf{A}''\|_2 \leq \frac{\epsilon}{2} \|\mathbf{A}\|_F$, we use a variant of Girshgorin's theorem, as in the proof of Lemma 5. First, we split the entries of $\mathbf{A}$ into level sets, according to their magnitudes. Let $\mathbf{A} = \sum_{k=0}^{\log \frac{n}{\epsilon}} \mathbf{A}_k$ where $(\mathbf{A}_0)_{ij} = \mathbf{A}_{ij}$ if $|\mathbf{A}_{ij}| \in [0, \frac{\epsilon}{n} \|\mathbf{A}\|_F]$ and $(\mathbf{A}_0)_{ij} = 0$ otherwise. For $1 \leq k \leq \log \frac{n}{\epsilon}$, $(\mathbf{A}_k)_{ij} = \mathbf{A}_{ij}$ if $|\mathbf{A}_{ij}| \in [\frac{\|\mathbf{A}\|_F}{2^k}, \frac{\|\mathbf{A}\|_F}{2^{k-1}}]$ and $(\mathbf{A}_k)_{ij} = 0$ otherwise. We can also define $\mathbf{A}'' = \sum_{k=0}^{\log \frac{n}{\epsilon}} \mathbf{A}''_k$ where each $\mathbf{A}''_k$ are defined similarly. By triangle inequality, $\|\mathbf{A} - \mathbf{A}''\|_2 \leq \sum_{k=0}^{\log n/\epsilon} \|\mathbf{A}_k - \mathbf{A}''_k\|_2$. First observe that $\|\mathbf{A}_0 - \mathbf{A}''_0\|_2 \leq \|\mathbf{A}_0 - \mathbf{A}''_0\|_F \leq n \cdot \|\mathbf{A}_0\|_\infty \leq \epsilon \|\mathbf{A}\|_F$. Further, we can assume without loss of generality that $\epsilon > 1/n$ and so $\log(n/\epsilon) \leq 2 \log n$, as otherwise our algorithm can afford to read all of $\mathbf{A}$. So, it suffices to show that for all $k \geq 1$,

$$\|\mathbf{A}_k - \mathbf{A}''_k\|_2 \leq \frac{\epsilon}{\log n} \cdot \|\mathbf{A}\|_F. \quad (37)$$

This will give $\|\mathbf{A} - \mathbf{A}''\|_2 \leq \epsilon \cdot \|\mathbf{A}\|_F + \sum_{k=1}^{\log n/\epsilon} \frac{\epsilon}{\log n} \cdot \|\mathbf{A}\|_F \leq 3\epsilon \cdot \|\mathbf{A}\|_F$, which gives the lemma after adjusting ϵ by a constant factor.

We now prove (37) for each $k \geq 1$. For $p \in \{0, 1, \dots, \log(n^2)\}$, let $\mathcal{I}_p \subset [n]$ be the set of rows/columns in $\mathbf{A}_k$ with $\text{nnz}((\mathbf{A}_k)_i) \in \left[\frac{\text{nnz}(\mathbf{A}_k)}{2^p}, \frac{\text{nnz}(\mathbf{A}_k)}{2^{p-1}} \right)$ and let $\mathbf{A}_{kpq} = \mathbf{A}_k(\mathcal{I}_p, \mathcal{I}_q)$ be the submatrix of $\mathbf{A}_k$ formed with rows in $\mathcal{I}_p$ and columns in $\mathcal{I}_q$. Define the submatrix $\mathbf{A}''_{kpq}$ of $\mathbf{A}''_k$ in the same way. Let $\hat{\mathbf{A}}_{kpq} = \mathbf{A}_{kpq} - \mathbf{A}''_{kpq}$ and finally,

let $\bar{\mathbf{A}}_{kpq} \in \mathbb{R}^{n \times n}$ be the symmetric error matrix such that $\bar{\mathbf{A}}_{kpq}(\mathcal{I}_p, \mathcal{I}_q) = \hat{\mathbf{A}}_{kpq}$ and $\bar{\mathbf{A}}_{kpq}(\mathcal{I}_q, \mathcal{I}_p) = \hat{\mathbf{A}}_{kpq}^T$.

Note that all rows from which we zero out entries must have at least one non-zero entry $\text{nnz}((\mathbf{A}_k)_i) \geq 1$ (otherwise all entries in that row/column are already zero), thus all such rows have $\text{nnz}((\mathbf{A}_k)_i) \geq \frac{\text{nnz}(\mathbf{A}_k)}{n^2}$ and so are covered by the submatrices $\mathbf{A}_{kpq}$. Thus, by triangle inequality, we can bound

$$\|\mathbf{A}_k - \mathbf{A}_k''\|_2 \leq \sum_{p=0}^{\log(n^2)} \sum_{q=0}^{\log(n^2)} \|\mathbf{N}_{kpq}\|_2. \quad (38)$$

To prove (37), we need to bound $\|\mathbf{A}_{kpq} - \mathbf{A}_{kpq}''\|_2$ for all $k \geq 1$ and p, q . We use a case analysis.

Case 1 $\frac{4 \text{nnz}(\mathbf{A}_k)^2 \cdot c_2 \log^4 n}{\epsilon^2 \cdot 2^{2k}} > 2^{p+q}$. In this case, first observe that since the nonzero entries of $\mathbf{A}_k$ lie in $\left[\frac{\|\mathbf{A}\|_F}{2^k}, \frac{\|\mathbf{A}\|_F}{2^{k-1}}\right)$, for any $i \in \mathcal{I}_p, j \in \mathcal{I}_j$,

$$\begin{aligned} \|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 &\geq \|(\mathbf{A}_k)_i\|_2^2 \cdot \|(\mathbf{A}_k)_j\|_2^2 \\ &\geq \frac{\|\mathbf{A}\|_F^4}{2^{4k}} \cdot \text{nnz}((\mathbf{A}_k)_i) \cdot \text{nnz}((\mathbf{A}_k)_j) \\ &\geq \frac{\|\mathbf{A}\|_F^4}{2^{4k} \cdot 2^{p+q}} \cdot \text{nnz}(\mathbf{A}_k)^2. \end{aligned}$$

Thus, by the assumed bound on 2^{p+q} , we have for any i, j where $(\mathbf{A}_k)_{ij}$ is nonzero,

$$\|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 \geq \frac{\epsilon^2 \|\mathbf{A}\|_F^4}{4 \cdot 2^{2k} c_2 \log^4 n} \geq \frac{\epsilon^2 \|\mathbf{A}\|_F^2 \cdot |\mathbf{A}_{ij}|^2}{c_2 \log^4 n},$$

where the second inequality follows again from the fact that the nonzero entries of $\mathbf{A}_k$ lie in $\left[\frac{\|\mathbf{A}\|_F}{2^k}, \frac{\|\mathbf{A}\|_F}{2^{k-1}}\right)$. Thus, any i, j with $(\mathbf{A}_{kpq})_{ij}$ nonzero is *not zeroed out* in line 5 of Algorithm 3. So $\mathbf{N}_{kpq} = \mathbf{0}$. Plugging into (38), we thus have:

$$\|\mathbf{A}_k - \mathbf{A}_k''\|_2 \leq \sum_{p=0}^{\log(n^2)} \sum_{q: 2^{p+q} \geq \frac{16 \text{nnz}(\mathbf{A}_k)^2 \cdot c_2 \log^4 n}{\epsilon^2 \cdot 2^{2k}}} \|\mathbf{N}_{kpq}\|_2. \quad (39)$$

Case 2 $\frac{16 \text{nnz}(\mathbf{A}_k)^2 \cdot c_2 \log^4 n}{\epsilon^2 \cdot 2^{2k}} \leq 2^{p+q}$. In this case, observe that $(\hat{\mathbf{A}}_{kpq} \hat{\mathbf{A}}_{kpq}^T)_m = (\hat{\mathbf{A}}_{kpq})_m \hat{\mathbf{A}}_{kpq}^T$. We can see that $(\hat{\mathbf{A}}_{kpq})_m$ has at most $\text{nnz}((\mathbf{A}_k)_m) \leq \frac{\text{nnz}(\mathbf{A}_k)}{2^{p-1}}$ non-zero entries. Similarly, each row of $\hat{\mathbf{A}}_{kpq}^T$ has at most $\frac{\text{nnz}(\mathbf{A}_k)}{2^{q-1}}$ non-zero elements. Thus, for all $m \in |\mathcal{I}_p|$, using the fact that all non-zero entries of $\mathbf{A}_{kpq}$ are bounded by $\frac{\|\mathbf{A}\|_F}{2^{k-1}}$, we

have:

$$\|(\widehat{\mathbf{A}}_{kpq} \widehat{\mathbf{A}}_{kpq}^T)_m\|_1 \leq \frac{\text{nnz}(\mathbf{A}_k)^2}{2^{p+q-2}} \cdot \frac{\|\mathbf{A}\|_F^2}{2^{2k-2}}.$$

Applying Girshgorin's circle theorem (Theorem 2) we thus have:

$$\|\widehat{\mathbf{A}}_{kpq}\|_2^2 = \|\widehat{\mathbf{A}}_{kpq} \widehat{\mathbf{A}}_{kpq}^T\|_2 \leq \frac{\text{nnz}(\mathbf{A}_k)^2}{2^{p+q-2}} \cdot \frac{\|\mathbf{A}\|_F^2}{2^{2k-2}}$$

and so

$$\|\bar{\mathbf{A}}_{kpq}\|_2 \leq 2\|\widehat{\mathbf{A}}_{kpq}\|_2 \leq \frac{8 \cdot \|\mathbf{A}\|_F \cdot \text{nnz}(\mathbf{A}_k)}{2^k 2^{\frac{(p+q)}{2}}}.$$

Plugging to (39), we thus have:

$$\begin{aligned} \|\mathbf{A}_k - \mathbf{A}_k''\|_2 &\leq \sum_{p=0}^{\log(n^2)} \sum_{q: 2^{p+q} \geq \frac{16 \text{nnz}(\mathbf{A}_k)^2 \cdot c_2 \log^4 n}{\epsilon^2 \cdot 2^{2k}}} \frac{8 \cdot \|\mathbf{A}\|_F \cdot \text{nnz}(\mathbf{A}_k)}{2^k 2^{\frac{(p+q)}{2}}} \\ &\leq \sum_{p=0}^{\log(n^2)} \frac{2\epsilon \cdot \|\mathbf{A}\|_F}{\sqrt{c_2} \log^2 n} \cdot \sum_{i=0}^{\infty} \frac{1}{\sqrt{2}} \leq \frac{8\epsilon \|\mathbf{A}\|_F}{\sqrt{c_2}}. \end{aligned}$$

Setting $c_2 \geq 64$, we thus have (37), and in turn the lemma. $\square$

We next give a bound on the incoherence of the outlying eigenvectors of $\mathbf{A}'$. This bound is again similar to Lemmas 2 and 6.

Lemma 12 (Incoherence of outlying eigenvectors in terms of ℓ_2 norms) *Let $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{n \times n}$ be as in Lemma 11. Let $\mathbf{A}'_o = \mathbf{V}'_o \mathbf{\Lambda}'_o \mathbf{V}'_o{}^T$ where $\mathbf{\Lambda}'_o$ is diagonal, with the eigenvalues of $\mathbf{A}'$ with magnitude $\geq \epsilon \sqrt{\delta} \|\mathbf{A}\|_F$ on its diagonal, and $\mathbf{V}'_o$ has columns equal to the corresponding eigenvectors. Let $\mathbf{V}'_{o,i}$ denote the i^{th} row of $\mathbf{V}'_o$. Then,*

$$\|\mathbf{A}'^{1/2} \mathbf{V}'_{o,i}\|_2^2 \leq \frac{\|\mathbf{A}_i\|_2^2}{\epsilon \sqrt{\delta} \|\mathbf{A}\|_F} \quad \text{and} \quad \|\mathbf{V}'_{o,i}\|_2^2 \leq \frac{\|\mathbf{A}_i\|_2^2}{\epsilon^2 \delta \|\mathbf{A}\|_F^2}.$$

Proof The proof is again nearly identical to that of Lemma 2. Observe that $\mathbf{A}' \mathbf{V}'_o = \mathbf{V}'_o \mathbf{\Lambda}'_o$. Letting $[\mathbf{A}' \mathbf{V}'_o]_i$ denote the i^{th} row of the $\mathbf{A}' \mathbf{V}'_o$, we have

$$\|[\mathbf{A}' \mathbf{V}'_o]_i\|_2^2 = \|[\mathbf{V}'_o \mathbf{\Lambda}'_o]_i\|_2^2 = \sum_{j=1}^r \lambda_j^2 \cdot \mathbf{V}'_{o,i,j}^2, \quad (40)$$

where $r = \text{rank}(\mathbf{A}'_o)$, $\mathbf{V}'_{o,i,j}$ is the $(i, j)^{\text{th}}$ element of $\mathbf{V}'_o$ and $\lambda_j = \mathbf{\Lambda}'_o(j, j)$. Since $\mathbf{V}'_o$ has orthonormal columns, we have $\|[\mathbf{A}' \mathbf{V}'_o]_i\|_2^2 \leq \|\mathbf{A}'_i\|_2^2 \leq \|\mathbf{A}_i\|_2^2$. Therefore, by

(40),

$$\sum_{j=1}^r \lambda_j^2 \cdot \mathbf{V}_{o,i,j}'^2 \leq \|\mathbf{A}_i\|_2^2. \quad (41)$$

Since by definition $|\lambda_j| \geq \epsilon\sqrt{\delta}\|\mathbf{A}\|_F$ for all j , we can conclude that $\|\mathbf{A}_o'^{1/2}\mathbf{V}_{o,i}'\|_2^2 = \sum_{j=1}^r \lambda_j \cdot \mathbf{V}_{o,i,j}'^2 \leq \frac{\|\mathbf{A}_i\|_2^2}{\epsilon\sqrt{\delta}\|\mathbf{A}\|_F}$ and $\|\mathbf{V}_{o,i}'\|_2^2 = \sum_{j=1}^r \mathbf{V}_{o,i,j}'^2 \leq \frac{\|\mathbf{A}_i\|_2^2}{\epsilon^2\delta\|\mathbf{A}\|_F^2}$, which completes the lemma. $\square$

E.2 Outer and Middle Eigenvalue Bounds

Using Lemma 12, we next argue that the eigenvalues of $\mathbf{A}'_{o,S}$ will approximate those of $\mathbf{A}'$, and in turn those of $\mathbf{A}$. The proof is very similar to Lemmas 3 and 7.

Lemma 13 (Concentration of outlying eigenvalues with ℓ_2 norm based sampling) *Let $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{n \times n}$ be as in algorithm 3. Let $\mathbf{A}' = \mathbf{A}'_m + \mathbf{A}'_o$, where $\mathbf{A}'_m = \mathbf{V}'_m \mathbf{\Lambda}'_m \mathbf{V}'_m{}^T$, and $\mathbf{A}'_o = \mathbf{V}'_o \mathbf{\Lambda}'_o \mathbf{V}'_o{}^T$ are projections onto the eigenspaces with magnitude $< \epsilon\sqrt{\delta}\|\mathbf{A}\|_F$ and $\geq \epsilon\sqrt{\delta}\|\mathbf{A}\|_F$ respectively. For all $i \in [n]$ let $p_i = \min\left(1, \frac{s\|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{n^2}\right)$ and let $\bar{\mathbf{S}}$ be a scaled diagonal sampling matrix such that the $\bar{\mathbf{S}}_{ii} = \frac{1}{\sqrt{p_i}}$ with probability p_i and $\bar{\mathbf{S}}_{ii} = 0$ otherwise. If $s \geq \frac{c \log(1/(\epsilon\delta))}{\epsilon^3\sqrt{\delta}}$ for a large enough constant c , then with probability at least $1 - \delta$, $\|\mathbf{A}_o'^{1/2}\mathbf{V}_o'^T \bar{\mathbf{S}} \bar{\mathbf{S}}^T \mathbf{V}_o' \mathbf{\Lambda}_o'^{1/2} - \mathbf{A}_o'\|_2 \leq \epsilon\|\mathbf{A}\|_F$.*

Proof We define the random variables $\mathbf{Q}_1, \dots, \mathbf{Q}_n$ and the set $P = \{i \in [n] : p_i < 1\}$ exactly as in the proof of Lemma 7. Then, as explained in the proof of Lemma 7 it is sufficient to bound $\sum_{i \in P} \mathbb{E}[\mathbf{Q}_i^2]$. From 17 we have $\sum_{i \in P} \mathbb{E}[\mathbf{Q}_i^2] \leq \sum_{i \in P} \frac{1}{p_i} \cdot \|\mathbf{A}_o'^{1/2}\mathbf{V}_{o,i}\|_2^2 \cdot (\mathbf{\Lambda}_o'^{1/2}\mathbf{V}_{o,i}\mathbf{V}_{o,i}^T\mathbf{\Lambda}_o'^{1/2})$. Also from Lemma 11, we have $\|\mathbf{A}_o'^{1/2}\mathbf{V}_{o,i}\|_2^2 \leq \frac{\|\mathbf{A}_i\|_2^2}{\epsilon\sqrt{\delta}\|\mathbf{A}\|_F}$ and for all $i \in P$, $\frac{1}{p_i} \leq \frac{\|\mathbf{A}\|_F^2}{s\|\mathbf{A}_i\|_2^2}$. We thus get,

$$\begin{aligned} \sum_{i \in P} \mathbb{E}[\mathbf{Q}_i^2] &\leq \sum_{i \in P} \frac{1}{p_i} \cdot \frac{\|\mathbf{A}_i\|_2^2}{\epsilon\sqrt{\delta}\|\mathbf{A}\|_F} \cdot (\mathbf{\Lambda}_o'^{1/2}\mathbf{V}_{o,i}\mathbf{V}_{o,i}^T\mathbf{\Lambda}_o'^{1/2}) \\ &\leq \frac{\|\mathbf{A}\|_F}{s\epsilon\sqrt{\delta}} \left(\sum_{i \in P} \mathbf{\Lambda}_o'^{1/2}\mathbf{V}_{o,i}\mathbf{V}_{o,i}^T\mathbf{\Lambda}_o'^{1/2} \right) \\ &= \frac{\|\mathbf{A}\|_F}{s\epsilon\sqrt{\delta}} \mathbf{\Lambda}_o' \leq \frac{\|\mathbf{A}\|_F^2}{s\epsilon\sqrt{\delta}} \cdot \mathbf{I}. \end{aligned}$$

Since $\mathbf{Q}_i^2$ is PSD this establishes that $v \leq \|\mathbf{Var}(\mathbf{E})\|_2 \leq \frac{\|\mathbf{A}\|_F^2}{s\epsilon\sqrt{\delta}}$. We can then apply the matrix Bernstein inequality exactly as in the proof of Lemma 3 to show that when $s \geq \frac{c}{\epsilon^3\sqrt{\delta}}$ for large enough c , with probability at least $1 - \delta$, $\|\mathbf{E}\|_2 \leq \epsilon\|\mathbf{A}\|_F$. $\square$

We now bound the middle eigenvalues.

Lemma 14 (Concentration of middle eigenvalues with ℓ_2 - norm based sampling) *Let $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{n \times n}$ be as in Lemma 12. Let $\mathbf{A}' = \mathbf{A}'_m + \mathbf{A}'_o$, where $\mathbf{A}'_m = \mathbf{V}'_m \mathbf{A}'_m \mathbf{V}'_m{}^T$, and $\mathbf{A}'_o = \mathbf{V}'_o \mathbf{A}'_o \mathbf{V}'_o{}^T$ are projections onto the eigenspaces with magnitude $< \epsilon \sqrt{\delta} \|\mathbf{A}\|_F$ and $\geq \epsilon \sqrt{\delta} \|\mathbf{A}\|_F$ respectively (analogous to Definition 1.1). As in Algorithm 2, for all $i \in [n]$ let $p_i = \min \left(1, \frac{s \|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{n^2} \right)$ and let $\tilde{\mathbf{S}}$ be a scaled diagonal sampling matrix such that the $\tilde{\mathbf{S}}_{ii} = \frac{1}{\sqrt{p_i}}$ with probability p_i and $\tilde{\mathbf{S}}_{ii} = 0$ otherwise. If $s \geq \frac{c \log^{10} n}{\epsilon^8 \delta^4}$ for a large enough constant c , then with probability at least $1 - \delta$,*

$$\|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq \epsilon \|\mathbf{A}\|_F.$$

Proof First observe that since $s \geq \frac{4}{\epsilon^2}$ (for large enough c), the results of Lemmas 11 and 12 still hold. The proof follows the same structure as the proof of bounding the middle eigenvalues for sparsity sampling in Lemma 8. From Lemma 12, we have $\|\mathbf{V}'_{o,i}\|_2 \leq \frac{\|\mathbf{A}_i\|_2}{\epsilon \sqrt{\delta} \|\mathbf{A}\|_F}$. Also, following the proof of Lemma 12, we have $\|\mathbf{A}'_o \mathbf{V}'_{o,j}\|_2 = \|[\mathbf{A}' \mathbf{V}'_o]_j\|_2 \leq \|\mathbf{A}_j\|_2$. Thus, for all $i, j \in [n]$, using Cauchy Schwarz's inequality, we have

$$|\mathbf{A}'_{o,i,j}| = |\mathbf{V}'_{o,i} \mathbf{A}'_o \mathbf{V}'_{o,j}{}^T| \leq \|\mathbf{V}'_{o,i}\|_2 \cdot \|\mathbf{A}'_o \mathbf{V}'_{o,j}{}^T\|_2 \leq \frac{\|\mathbf{A}_i\|_2}{\epsilon \sqrt{\delta} \|\mathbf{A}\|_F} \cdot \|\mathbf{A}_j\|_2. \quad (42)$$

Let $\mathbf{A}'_m = \mathbf{H}_m + \mathbf{D}_m$ where $\mathbf{H}_m$ and $\mathbf{D}_m$ contain the off-diagonal and diagonal elements of $\mathbf{A}'_m$ respectively. Then following the proof of Lemma 8, we get:

$$\begin{aligned} \mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{A}'_m \tilde{\mathbf{S}}\|_2 &\leq 10 \sqrt{\log n} \left(\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} + \mathbb{E}_2 \|\mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} \right) \\ &\quad + 2 \|\mathbf{H}_m\|_2 + \mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{D}_m \tilde{\mathbf{S}}\|_2 \end{aligned} \quad (43)$$

We now proceed to bound each of the terms on the right hand side of (43). We start with $\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{D}_m \tilde{\mathbf{S}}\|_2$. First, observe that $\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{D}_m \tilde{\mathbf{S}}\|_2 \leq \max_i \frac{1}{p_i} |(\mathbf{D}_m)_{ii}|$. We consider two cases.

Case 1: $p_i < 1$. Then, as $p_i \geq \frac{s \|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2}$ we have $\|\mathbf{A}\|_F^2 \leq \frac{1}{s} \|\mathbf{A}_i\|_2^2$ since $\frac{1}{s} < \frac{\epsilon^2}{4}$. So we must have that have $|(\mathbf{D}_m)_{ii}| = |(\mathbf{A}'_m)_{ii}| = |(\mathbf{A}'_o)_{ii}|$ (since $\mathbf{A}'_{ii} = 0$). Then by (42), we have $\frac{1}{p_i} |(\mathbf{D}_m)_{ii}| \leq \frac{\|\mathbf{A}\|_F}{s \epsilon \sqrt{\delta}}$.

Case 2: $p_i = 1$. Then we have $\frac{1}{p_i} |(\mathbf{D}_m)_{ii}| = |(\mathbf{D}_m)_{ii}| \leq \max_j |(\mathbf{D}_m)_{jj}| \leq \|\mathbf{A}'_m\|_2 \leq \epsilon \sqrt{\delta} \|\mathbf{A}\|_F$.

From the two cases above, for $s \geq \frac{1}{\epsilon^2 \delta}$, we have:

$$\mathbb{E}_2 \|\tilde{\mathbf{S}} \mathbf{D}_m \tilde{\mathbf{S}}\|_2 \leq \epsilon \sqrt{\delta} \|\mathbf{A}\|_F. \quad (44)$$

We can bound $\|\mathbf{H}_m\|_2$ similarly. Since $\mathbf{H}_m = \mathbf{A}'_m - \mathbf{D}_m$ and $\|\mathbf{A}'_m\|_2 \leq \epsilon\sqrt{\delta}\|\mathbf{A}\|_F$,

$$\begin{aligned}\|\mathbf{H}_m\|_2 &\leq \|\mathbf{A}'_m\|_2 + \|\mathbf{D}_m\|_2 \\ &\leq \epsilon\sqrt{\delta}\|\mathbf{A}\|_F + \epsilon\sqrt{\delta}\|\mathbf{A}\|_F. \\ &= 2\epsilon\sqrt{\delta}\|\mathbf{A}\|_F.\end{aligned}\quad (45)$$

where the second step follows from the fact that $\|\mathbf{D}_m\|_2 \leq \max_i |(\mathbf{D}_m)_{ii}| \leq \|\mathbf{A}'_m\|_2$.

We next bound the term $\mathbb{E}_2\|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2}$. Observe that $\mathbb{E}_2\|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \frac{\max_i \|\mathbf{A}'_{m,i}\|_2}{\sqrt{p_i}}$, where $\mathbf{A}'_{m,i}$ is the i^{th} column/row of $\mathbf{A}'_m$. We again consider the two cases when $p_i = 1$ and $p_i < 1$:

Case 1: $p_i = 1$. Then $\|\mathbf{A}'_{m,i}\|_2 \leq \|\mathbf{A}'_m\|_2 \leq \epsilon\sqrt{\delta}\|\mathbf{A}\|_F$.

Case 2: $p_i < 1$. Then $\|\mathbf{A}'_{m,i}\|_2 \leq \|\mathbf{A}'_i\|_2 \leq \|\mathbf{A}\|_F$. Thus, setting $s \geq \frac{1}{\epsilon^2\delta}$ we have:

$$\begin{aligned}\frac{\|\mathbf{A}'_{m,i}\|_2}{\sqrt{p_i}} &\leq \frac{\|\mathbf{A}\|_F}{\sqrt{s}\|\mathbf{A}_i\|_2} \cdot \|\mathbf{A}'_i\|_2 \\ &\leq \frac{\|\mathbf{A}\|_F}{\sqrt{s}} \leq \epsilon\sqrt{\delta}\|\mathbf{A}\|_F.\end{aligned}$$

Thus, from the two cases above, for all $i \in [n]$, adjusting ϵ by a $\frac{1}{\sqrt{\log n}}$ factor, we have for $s \geq \frac{\log n}{\epsilon^2\delta}$:

$$\mathbb{E}_2\|\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \frac{\epsilon\sqrt{\delta}\|\mathbf{A}\|_F}{\sqrt{\log n}}. \quad (46)$$

Overall, plugging (44), (45), and (46) back into (43), we have:

$$\mathbb{E}_2\|\bar{\mathbf{S}}\mathbf{A}'_m\bar{\mathbf{S}}\|_2 \leq 10\sqrt{\log n} \cdot \mathbb{E}_2\|\bar{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} + 15\epsilon\sqrt{\delta}\|\mathbf{A}\|_F. \quad (47)$$

Finally we bound $\mathbb{E}_2\|\bar{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2}$. As in the proof of Lemma 8, we have $\mathbb{E}_2\|\bar{\mathbf{S}}\mathbf{H}_m\hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \mathbb{E}_2\left(\max_{i:i \in [n]} \frac{\|(\bar{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}}\right)$ and we will argue that $\max_{i:i \in [n]} \frac{\|(\bar{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}}$ is bounded by $\epsilon\sqrt{\delta}\|\mathbf{A}\|_F$ with probability $1 - 1/\text{poly}(n)$. Also as argued in the proof of Lemma 8, since $p_i \geq \frac{1}{n^2}$, it suffices to bound $\frac{\|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2}{\sqrt{p_i}}$ for all $i \in [n]$ with high probability. Again, for a fixed i and any $j \in [n]$, define the random variables z_j as:

$$z_j = \begin{cases} \frac{1}{p_j} |\mathbf{A}'_{m,i,j}|^2 & \text{with probability } p_j \\ 0 & \text{otherwise.} \end{cases}$$

Then $\sum_{j=1}^n z_j = \|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2$ and $\mathbb{E}[\sum_{j=1}^n z_j] = \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_i\|_2^2 \leq \|\mathbf{A}\|_F^2$. We will again use Bernstein's inequality to bound $\sum_{j=1}^n z_j = \|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2$ by bounding

bound $|z_j|$ for all $j \in [n]$ and $\mathbf{Var} \left(\sum_{j=1}^n z_j \right)$. We consider the cases of $p_i < 1$ and $p_i = 1$ separately.

Case 1: $p_i < 1$. Then, we have $p_i \geq s \|\mathbf{A}_i\|_2^2 / \|\mathbf{A}\|_F^2$. If $\mathbf{A}'_{i,j} \neq 0$ then

$$\begin{aligned} |z_j| &\leq \frac{1}{p_j} |\mathbf{A}'_{m,i,j}|^2 \leq \max \left(1, \frac{\|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} \right) |\mathbf{A}'_{m,i,j}|^2 \\ &\leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2\|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} \left(|\mathbf{A}'_{i,j}|^2 + |\mathbf{A}'_{o,i,j}|^2 \right) \\ &\leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2\|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} \left(|\mathbf{A}'_{i,j}|^2 + \frac{\|\mathbf{A}_i\|_2^2 \|\mathbf{A}_j\|_2^2}{\epsilon^2 \delta \|\mathbf{A}\|_F^2} \right) \\ &\leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2\|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} |\mathbf{A}'_{i,j}|^2 + \frac{2\|\mathbf{A}_i\|_2^2}{\epsilon^2 \delta s}, \end{aligned}$$

where the fourth inequality uses (42). By the thresholding procedure which defines $\mathbf{A}'$, if $i \neq j$ and $\mathbf{A}'_{ij} \neq 0$,

$$\|\mathbf{A}_i\|_2^2 \cdot \|\mathbf{A}_j\|_2^2 \geq \frac{\epsilon^2 \|\mathbf{A}\|_F^2 |\mathbf{A}'_{ij}|^2}{c_2 \log^4 n} \Rightarrow \frac{\|\mathbf{A}_j\|_2^2}{|\mathbf{A}'_{i,j}|^2} \geq \frac{\epsilon^2 \|\mathbf{A}\|_F^2}{c_2 \cdot \log^4 n \cdot \|\mathbf{A}_i\|_2^2}, \quad (48)$$

and thus for $p_i < 1$ and $\mathbf{A}'_{ij} \neq 0$ we have

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^2}{s \epsilon^2} + \frac{2\|\mathbf{A}_i\|_2^2}{\epsilon^2 \delta s}.$$

Also $\mathbf{A}'_{ii} = 0$ since we must have $\|\mathbf{A}_i\|_2^2 < \frac{\epsilon^2}{4} \|\mathbf{A}\|_F^2$ as $p_i < 1$. If $\mathbf{A}'_{i,j} = 0$ or $i = j$, then we simply have

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2\|\mathbf{A}_i\|_2^2}{s \epsilon^2 \delta}.$$

Overall for all $j \in [n]$,

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2\|\mathbf{A}_i\|_2^2}{s \epsilon^2 \delta} + \frac{2c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^2}{s \epsilon^2}, \quad (49)$$

and since $|\mathbf{A}'_{m,i,j}|^2 \leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^2 = \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_i\|_2^2 \leq \|\mathbf{A}_i\|_2^2$,

$$|z_j| \leq \|\mathbf{A}_i\|_2^2 + \frac{2\|\mathbf{A}_i\|_2^2}{s \epsilon^2 \delta} + \frac{2c_2 \cdot \log^4 n \cdot \|\mathbf{A}_i\|_2^2}{s \epsilon^2}. \quad (50)$$

For $s \geq c \left(\frac{\log^4 n}{\epsilon^2} + \frac{1}{\epsilon^2 \delta} \right)$ and large enough c , we thus have $|z_j| \leq 2\|\mathbf{A}_i\|_2^2$.

We next bound the variance by:

$$\begin{aligned}
 \text{Var} \left(\sum_{j=1}^n z_j \right) &\leq \sum_{j=1}^n \mathbb{E}[z_j^2] \leq \sum_{j=1}^n p_j \frac{1}{p_j^2} |\mathbf{A}'_{m,i,j}|^4 \\
 &= \sum_{j=1}^n \max \left(1, \frac{\|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} \right) |\mathbf{A}'_{m,i,j}|^4 \\
 &\leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^4 + \sum_{j=1}^n \frac{12 \|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} \left(|\mathbf{A}'_{i,j}|^4 + |\mathbf{A}'_{o,i,j}|^4 \right) \\
 &\leq \|\mathbf{A}'_{m,i}\|_2^4 + \sum_{j=1}^n \frac{12 \|\mathbf{A}\|_F^2}{s \|\mathbf{A}_j\|_2^2} \left(|\mathbf{A}'_{i,j}|^4 + \frac{\|\mathbf{A}_i\|_2^4 \|\mathbf{A}_j\|_2^4}{\epsilon^4 \delta^2 \|\mathbf{A}\|_F^4} \right),
 \end{aligned}$$

where the last inequality uses (42). We thus get:

$$\text{Var} \left(\sum_{j=1}^n z_j \right) \leq \|\mathbf{A}'_{m,i}\|_2^4 + \sum_{j: \mathbf{A}'_{i,j} \neq 0} \frac{12 \|\mathbf{A}\|_F^2 |\mathbf{A}'_{i,j}|^4}{s \|\mathbf{A}_j\|_2^2} + \sum_{j=1}^n \frac{12 \|\mathbf{A}_i\|_2^4 \|\mathbf{A}_j\|_2^2}{s \epsilon^4 \delta^2 \|\mathbf{A}\|_F^2}. \quad (51)$$

Now $\mathbf{A}'_{ii} = 0$ as $p_i < 1$ (and thus, $\|\mathbf{A}\|_i^2 < \frac{\epsilon^2}{4} \|\mathbf{A}\|_F^2$). Combining (48) with the second term to the right of (51) we have

$$\begin{aligned}
 \text{Var} \left(\sum_{j=1}^n z_j \right) &\leq \|\mathbf{A}'_{m,i}\|_2^4 + \sum_{j: \mathbf{A}'_{i,j} \neq 0} \frac{12 c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^2 \cdot |\mathbf{A}'_{i,j}|^2}{s \epsilon^2} \\
 &\quad + \sum_{j=1}^n \frac{12 \|\mathbf{A}_i\|_2^4 \|\mathbf{A}_j\|_2^2}{s \epsilon^4 \delta^2 \|\mathbf{A}\|_F^2},
 \end{aligned}$$

and since $\sum_j |\mathbf{A}'_{i,j}|^2 = \|\mathbf{A}_i\|_2^2$, we have

$$\begin{aligned}
 \text{Var} \left(\sum_{j=1}^n z_j \right) &\leq \|\mathbf{A}'_{m,i}\|_2^4 + \frac{12 c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^4}{s \epsilon^2} \\
 &\quad + \sum_{j=1}^n \frac{12 \|\mathbf{A}_i\|_2^4 \|\mathbf{A}_j\|_2^2}{s \epsilon^4 \delta^2 \|\mathbf{A}\|_F^2}. \quad (52)
 \end{aligned}$$

Finally since $\sum_{j=1}^n \|\mathbf{A}_j\|_2^2 = \|\mathbf{A}\|_F^2$ and $\|\mathbf{A}'_{m,i}\|_2^4 \leq \|\mathbf{A}'_i\|_2^4 \leq \|\mathbf{A}_i\|_2^4$ we have

$$\text{Var} \left(\sum_{j=1}^n z_j \right) \leq \|\mathbf{A}_i\|_2^4 + \frac{12 c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^4}{s \epsilon^2} + \frac{12 \|\mathbf{A}_i\|_2^4}{s \epsilon^4 \delta^2}. \quad (53)$$

For $s \geq c \left(\frac{\log^4 n}{\epsilon^2} + \frac{1}{\epsilon^4 \delta^2} \right)$ for large enough c , we have $\text{Var} \left(\sum_{j=1}^n z_j \right) \leq 2 \|\mathbf{A}_i\|_2^4$.

Therefore, using (50) and (53) with $s \geq c \left(\frac{\log^4 n}{\epsilon^2} + \frac{1}{\epsilon^4 \delta^2} \right)$, we can apply Bernstein inequality (Theorem 7) (for some constant c) to get

$$\begin{aligned} \mathbb{P} \left(\|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \geq \mathbb{E}\|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 + t \right) &\leq \mathbb{P} \left(\sum_{j=1}^n z_j \geq \|\mathbf{A}_i\|_2^2 + t \right) \\ &\leq \exp \left(\frac{-t^2/2}{c\|\mathbf{A}_i\|_2^4 + ct\|\mathbf{A}_i\|_2^2/3} \right). \end{aligned}$$

If we set $t = \log n \cdot \|\mathbf{A}_i\|_2^2$, for some constant c' we have

$$\begin{aligned} \mathbb{P} \left(\|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \geq \mathbb{E}\|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 + \log n \cdot \|\mathbf{A}_i\|_2^2 \right) \\ \leq \exp \left(\frac{-(\log n)^2/2}{c + c(\log n)/3} \right) \leq \exp(-c' \log n) \leq 1/n^{c'}. \end{aligned}$$

Since $\mathbf{A}'_m = \mathbf{H}_m + \mathbf{D}_m$, we have $\|(\bar{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2 \geq \|(\bar{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2$. Then with probability at least $1 - 1/n^{c'} \geq 1 - \delta$, for any row i with $p_i < 1$, we have

$$\frac{1}{p_i} \cdot \|(\bar{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 \leq \frac{\|\mathbf{A}\|_F^2}{s\|\mathbf{A}_i\|_2^2} \cdot c(\log n)\|\mathbf{A}_i\|_2^2 \leq \frac{\epsilon^2 \delta \|\mathbf{A}\|_F^2}{\log n},$$

for $s \geq c \left(\frac{\log^4 n}{\epsilon^2} + \frac{1}{\epsilon^4 \delta^2} \right)$ for large enough c . Observe that, as in Lemma 3 w.l.o.g. we have assumed $1 - \frac{1}{n^{c'}} \geq 1 - \delta$, since otherwise, our algorithm would read all n^2 entries of the matrix.

Case 2: $p_i = 1$. Then, we have $\|\mathbf{A}_i\|_2^2 \geq \|\mathbf{A}\|_F^2/s$. As in the $p_i < 1$ case, when $\mathbf{A}_{ii} = 0$, (and this $\mathbf{A}'_{ii} = \mathbf{A}_{ii} = 0$) we have from (49):

$$|z_j| \leq |\mathbf{A}'_{m,i,j}|^2 + \frac{2\|\mathbf{A}_i\|_2^2}{s\epsilon^2\delta} + \frac{2c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^2}{s\epsilon^2}.$$

Now, we observe that $|\mathbf{A}'_{m,i,j}|^2 \leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^2 \leq \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_m\|_2^2 \leq \epsilon^2 \delta \|\mathbf{A}\|_F^2$, which gives us

$$|z_j| \leq \epsilon^2 \delta \|\mathbf{A}\|_F^2 + \frac{2\|\mathbf{A}_i\|_2^2}{s\epsilon^2\delta} + \frac{2c_2 \log^4 n \cdot \|\mathbf{A}_i\|_2^2}{s\epsilon^2}. \quad (54)$$

Note that if $\mathbf{A}_{ii} \neq 0$, the second term in (49) is bounded as $\frac{2\|\mathbf{A}\|_F^2}{s\|\mathbf{A}_i\|_2^2} \cdot |\mathbf{A}'_{ii}|^2 \leq \frac{2\|\mathbf{A}\|_F^2}{s} \leq 2\epsilon^2 \delta \|\mathbf{A}\|_F^2$ for $s \geq O(\frac{1}{\epsilon^2 \delta})$. Thus, for $s \geq c \left(\frac{\log^4 n}{\epsilon^4 \delta} + \frac{1}{\epsilon^4 \delta^2} \right)$ for a large enough constant

c and adjusting for other constants we have $|z_j| \leq 2\epsilon^2\delta\|\mathbf{A}\|_F^2$. Also observe that the expectation of $\sum z_j$ can be bounded by:

$$\mathbb{E}\left(\sum_{j=1}^n z_j\right) = \mathbb{E}\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 = \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_m\|_2^2 \leq \epsilon^2\delta\|\mathbf{A}\|_F^2.$$

Next, the variance of the sum of the random variables $\{z_j\}$ can again be bounded by following the analysis presented in and prior to (52) and (53) we have

$$\begin{aligned} \mathbf{Var}\left(\sum_{j=1}^n z_j\right) &\leq \|\mathbf{A}'_{m,i,j}\|_2^4 + \frac{12c_2 \log^2 n \cdot \|\mathbf{A}_i\|_2^4}{s\epsilon^2} + \frac{12\|\mathbf{A}_i\|_2^4}{s\epsilon^4\delta^2} \\ &\leq \epsilon^4\delta^2\|\mathbf{A}\|_F^4 + \frac{12c_2 \log^2 n \cdot \|\mathbf{A}_i\|_2^4}{s\epsilon^2} + \frac{12\|\mathbf{A}_i\|_2^4}{s\epsilon^4\delta^2}, \end{aligned} \quad (55)$$

where we again bound $\|\mathbf{A}'_{m,i,j}\|_2^4$ using

$$|\mathbf{A}'_{m,i,j}|^2 \leq \sum_{j=1}^n |\mathbf{A}'_{m,i,j}|^2 \leq \|\mathbf{A}'_{m,i}\|_2^2 \leq \|\mathbf{A}'_m\|_2^2 \leq \epsilon^2\delta\|\mathbf{A}\|_F^2.$$

Then for $s \geq c(\frac{\log^4 n}{\epsilon^6\delta^2} + \frac{1}{\epsilon^8\delta^4})$, we have $\mathbf{Var}\left(\sum_{j=1}^n z_j\right) \leq 2\epsilon^4\delta^2\|\mathbf{A}\|_F^4$ for large enough constant c .

Using (54) and (55) and noting that $\sum_{j=1}^n \mathbb{E}(z_j^2) \geq \mathbf{Var}\left(\sum_{j=1}^n z_j\right) - \mathbb{E}^2\left(\sum_{j=1}^n z_j\right)$ we can apply the Bernstein inequality (Theorem 7):

$$\begin{aligned} \mathbb{P}\left(\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \geq \mathbb{E}\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 + t\right) &\leq \mathbb{P}\left(\sum_{j=1}^n z_j \geq \epsilon^2\delta\|\mathbf{A}_i\|_2^2 + t\right) \\ &\leq \exp\left(\frac{-t^2/2}{c\epsilon^4\delta^2\|\mathbf{A}\|_F^4 + c\epsilon^2\delta\|\mathbf{A}\|_F^2 t/3}\right). \end{aligned}$$

If we set $t = (\log n)\epsilon^2\delta\|\mathbf{A}\|_F^2$, then for some constant c' we have

$$\mathbb{P}\left(\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \geq \mathbb{E}\|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 + t\right) \leq \exp(-c' \log n) \leq 1/n^{c'}.$$

This, since $\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 \leq \|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2$, when $p_i = 1$, setting $s \geq c(\frac{\log^4 n}{\epsilon^6\delta^2} + \frac{1}{\epsilon^8\delta^4})$ for large enough c , we have with probability $\geq 1 - 1/n^{c'}$ $\frac{1}{p_i}\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 = \|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2 \leq \|(\tilde{\mathbf{S}}\mathbf{A}'_m)_{:,i}\|_2^2 \leq (\log n)\epsilon^2\delta \text{nnz}(\mathbf{A})$.

We have proven that with probability $\geq 1 - 1/n^{c'}$, for both cases when $p_i < 1$ and $p_i = 1$, $\frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2^2}{p_i} \leq (\log n)\epsilon^2\delta\|\mathbf{A}\|_F^2$. Taking a union bound over all $i \in [n]$, with

probability at least $1 - 1/n^{c'-1}$, $\max_i \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \leq \sqrt{\log n} \epsilon \sqrt{\delta} \|\mathbf{A}\|_F$ for $s \geq c(\frac{\log^4 n}{\epsilon^6 \delta^2} + \frac{1}{\epsilon^8 \delta^4})$. Also, since $p_i \geq \frac{1}{n^2}$ for all $i \in [n]$, $\frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \leq \sqrt{\sum_{j=1}^n \frac{\mathbf{A}_{m,i,j}^2}{p_i \cdot p_j}} \leq \frac{n \cdot \|\mathbf{A}\|_F}{\sqrt{s}}$. Thus, $\max_i \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \leq n \|\mathbf{A}\|_F$ and we get,

$$\mathbb{E}_2 \left(\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \right) \leq \sqrt{\log n} \epsilon \sqrt{\delta} \|\mathbf{A}\|_F + \frac{1}{n^{c'-3}} \leq \sqrt{\log n} \epsilon \sqrt{\delta} \|\mathbf{A}\|_F.$$

after adjusting ϵ by at most some constants. Overall, we finally get

$$\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{H}_m \hat{\mathbf{S}}\|_{1 \rightarrow 2} \leq \mathbb{E}_2 \left(\max_{i:i \in [n]} \frac{\|(\tilde{\mathbf{S}}\mathbf{H}_m)_{:,i}\|_2}{\sqrt{p_i}} \right) \leq \epsilon \sqrt{\log n} \sqrt{\delta} \|\mathbf{A}\|_F.$$

Plugging this bound into (47), we have for $s \geq c(\frac{\log^4 n}{\epsilon^6 \delta^2} + \frac{1}{\epsilon^8 \delta^4})$,

$$\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq (\log n) \epsilon \sqrt{\delta} \|\mathbf{A}\|_F.$$

Finally after adjusting ϵ by a $\frac{1}{\log n}$ factor, we have for $s \geq c(\frac{\log^{10} n}{\epsilon^6 \delta^2} + \frac{\log^8 n}{\epsilon^8 \delta^4})$ or $s \geq \frac{c \log^{10} n}{\epsilon^8 \delta^4}$,

$$\mathbb{E}_2 \|\tilde{\mathbf{S}}\mathbf{A}'_m \tilde{\mathbf{S}}\|_2 \leq \epsilon \sqrt{\delta} \|\mathbf{A}\|_F.$$

The final bound then follows via Markov's inequality on $\|\tilde{\mathbf{S}}\mathbf{A}'_m \tilde{\mathbf{S}}\|_2$. $\square$

E.3 Main Accuracy Bound

We are finally ready to state our main result for ℓ_2 norm based sampling.

Theorem 3 (Row Norm Based Matrix Eigenvalue Approximation) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Let $S \subseteq [n]$ be formed by including the i^{th} index independently with probability $p_i = \min \left(1, \frac{s \|\mathbf{A}_i\|_2^2}{\|\mathbf{A}\|_F^2} + \frac{1}{n^2} \right)$ as in Algorithm 3. Here $\|\mathbf{A}_i\|_2$ is the ℓ_2 norm of the i^{th} row of $\mathbf{A}$. Let $\mathbf{A}_S$ be the corresponding principal submatrix of $\mathbf{A}$, and let $\tilde{\lambda}_i(\mathbf{A})$ be the estimate of $\lambda_i(\mathbf{A})$ computed from $\mathbf{A}_S$ as in Algorithm 3. If $s \geq \frac{c \log^{10} n}{\epsilon^8 \delta^4}$, for large enough constant c , then with probability $\geq 1 - \delta$, for all $i \in [n]$, $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \epsilon \|\mathbf{A}\|_F$.*

Proof The proof follows exactly the same structure as the proofs of Theorems 1 and 2 for uniform and sparsity based sampling respectively. We use the results of Lemmas 14 and 13 on the concentration of the middle and large eigenvalues for ℓ_2 norm based sampling.

Analogous to Theorem 2, from Lemma 13 with error parameter $\frac{\epsilon}{\log n}$ the eigenvalues of $\mathbf{A}'_{o,s}$ approximate those of $\mathbf{A}'_o$ up to error $\epsilon \|\mathbf{A}\|_F$ with probability $1 - \delta$ if $s \geq \frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}$. We also require $s \geq \frac{c \log^{10} n}{\epsilon^8 \delta^4}$ for $\|\mathbf{A}'_{m,s}\|_2 \leq \epsilon \|\mathbf{A}\|_F$ to hold with probability $1 - \delta$ by Lemma 14. Thus, for both conditions to hold simultaneously with probability $1 - 2\delta$ by a union bound, it suffices to set $s = \max\left(\frac{c \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \sqrt{\delta}}, \frac{c \log^{10} n}{\epsilon^8 \delta^4}\right) = \frac{c \log^{10} n}{\epsilon^8 \delta^4}$, where we use that $\log(1/(\epsilon\delta)) \leq \log n$, as otherwise our algorithm can take $\mathbf{A}_S$ to be the full matrix $\mathbf{A}$. Adjusting δ to $\delta/2$ completes the theorem. $\square$

F Eigenvalue Approximation via Entrywise Sampling

In this section we show that sampling $\tilde{O}(n/\epsilon^2)$ entries from a bounded entry matrix preserves its eigenvalues up to error $\pm\epsilon n$. We use this result to improve the sample complexity of Theorem 1 from $\tilde{O}\left(\frac{\log^6 n}{\epsilon^6}\right)$ to $\tilde{O}\left(\frac{\log^3 n}{\epsilon^5}\right)$ by applying entrywise sampling to further sparsify the submatrix $\mathbf{A}_S$ that is sampled in Algorithm 1. Entrywise sampling results similar to what we show are well-known in the literature. See for example [8] and [56]. For completeness, we give a proof here using standard matrix concentration bounds.

Theorem 10 (Entrywise sampling—spectral norm bound) *Consider $\mathbf{A} \in \mathbb{R}^{n \times n}$ with $\|\mathbf{A}\|_\infty \leq 1$. Let $\mathbf{C} \in \mathbb{R}^{n \times n}$ be constructed by setting $\mathbf{C}_{i,i} = \mathbf{A}_{i,i}$ for all $i \in [n]$ and*

$$\mathbf{C}_{j,i} = \mathbf{C}_{i,j} = \begin{cases} \frac{1}{p} \cdot \mathbf{A}_{i,j} & \text{with probability } p \\ 0 & \text{otherwise.} \end{cases}$$

For any $\epsilon, \delta \in (0, 1)$, if $p \geq \frac{c \log(n/\delta)}{n\epsilon^2}$ for a large enough constant c , then with probability at least $1 - \delta$, $\|\mathbf{A} - \mathbf{C}\|_2 \leq \epsilon n$.

Note that by Weyl's inequality (Fact 3), Theorem 10 immediately implies that the eigenvalues of $\mathbf{C}$ approximate those of $\mathbf{A}$ up to $\pm\epsilon n$ error with good probability.

Proof For any $i < j$, define the symmetric random matrix $\mathbf{E}^{(ij)}$ with

$$\mathbf{E}_{i,j}^{(ij)} = \mathbf{E}_{j,i}^{(ij)} = \begin{cases} (\frac{1}{p} - 1) \cdot \mathbf{A}_{i,j} & \text{with probability } p \\ -\mathbf{A}_{i,j} & \text{otherwise.} \end{cases}$$

Observe that $\mathbf{C} - \mathbf{A} = \sum_{i,j \in [n], i < j} \mathbf{E}^{(ij)}$. Further, each $\mathbf{E}^{(ij)}$ has just two non-zero values in different rows and columns. So

$$\|\mathbf{E}^{(ij)}\|_2 = |\mathbf{C}_{i,j} - \mathbf{A}_{i,j}| \leq \left(\frac{1}{p} - 1\right) \cdot |\mathbf{A}_{i,j}| \leq \frac{1}{p},$$

where the last inequality uses that $\|\mathbf{A}\|_\infty \leq 1$. Additionally, $\mathbf{E}^{(ij)}\mathbf{E}^{(ij)T}$ is diagonal with two diagonal entries at (i, i) or (j, j) equal to $(\mathbf{C}_{i,j} - \mathbf{A}_{i,j})^2$. Thus,

$\mathbf{V} = \sum_{i,j \in [n], i < j} \mathbb{E}[\mathbf{E}^{(ij)} \mathbf{E}^{(ij)T}]$ is also diagonal. We have

$$\begin{aligned} \mathbf{V}_{i,i} &= \sum_{j \neq i} \mathbb{E}[(\mathbf{C}_{i,j} - \mathbf{A}_{i,j})^2] = \sum_{j \neq i} \mathbf{A}_{i,j}^2 \cdot \left(p \cdot \left(\frac{1}{p} - 1 \right)^2 + (1-p) \cdot (-1)^2 \right) \\ &= \sum_{j \neq i} \mathbf{A}_{i,j}^2 \cdot \left(\frac{1}{p} - 1 \right) \leq \frac{n}{p}, \end{aligned}$$

where in the final inequality we use that $\|\mathbf{A}\|_\infty \leq 1$. Thus, since $\mathbf{V}$ is diagonal, $\|\mathbf{V}\|_2 \leq \frac{n}{p}$. Putting the above together using Theorem 6 we get,

$$\mathbb{P}(\|\mathbf{A} - \mathbf{C}\|_2 \geq \epsilon n) = \mathbb{P}\left(\left\| \sum_{i,j \in [n], i < j} \mathbf{E}^{(ij)} \right\|_2 \geq \epsilon n\right) \leq 2n \cdot \exp\left(\frac{-\epsilon^2 n^2 / 2}{\frac{n}{p} + \frac{\epsilon n}{3p}}\right).$$

Thus, for $p \geq \frac{c \log(n/\delta)}{n\epsilon^2}$ for large enough c , with probability at least $1 - \delta$ we have $\|\mathbf{A} - \mathbf{C}\|_2 \leq \epsilon n$. $\square$

F.1 Improved Sample Complexity via Entrywise Sampling

We can combine Theorem 10 directly with Theorem 1 to give an improved sample complexity for eigenvalue estimation. we have:

Corollary 1 (Improved Sample Complexity via Entrywise Sampling) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric with $\|\mathbf{A}\|_\infty \leq 1$ and eigenvalues $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. For any $\epsilon, \delta \in (0, 1)$, there is an algorithm that reads $\tilde{O}\left(\frac{\log^3 n}{\epsilon^3 \delta}\right)$ entries of $\mathbf{A}$ and returns, with probability at least $1 - \delta$, $\tilde{\lambda}_i(\mathbf{A})$ for each $i \in [n]$ satisfying $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \epsilon n$.*

Proof Letting $s = \frac{c_1 \log(1/(\epsilon\delta)) \cdot \log^3 n}{\epsilon^3 \delta}$ for large enough constant c_1 , by Theorem 1, for a random principal submatrix $\mathbf{A}_S$ formed by sampling each index with probability s/n , the eigenvalues of $\mathbf{A}_S$, after scaling up by a factor of n/s approximate those of $\mathbf{A}$ to error $\pm \epsilon n$ with probability at least $1 - \delta$. By Theorem 10, if we sample off-diagonal entries of $\mathbf{A}_S$ with probability $p \geq \frac{c_2 \log(|S|/\delta)}{|S| \cdot \epsilon^2}$ to produce $\mathbf{C}$, then we preserve its eigenvalues to error $\pm \epsilon |S|$. Thus, after scaling by $\frac{n}{s}$, the eigenvalues of $\mathbf{C}$ approximate those of $\mathbf{A}$ to error $\pm (\epsilon n + \frac{n}{s} \cdot \epsilon |S|)$. Finally, observe that by a standard Chernoff bound, $|S| \leq 2s$ with probability at least $1 - \delta$. So adjusting ϵ by a constant, the scaled eigenvalues of $\mathbf{C}$ give $\pm \epsilon n$ approximations to $\mathbf{A}$'s eigenvalues. The expected number of entries read is $|S| + p \cdot |S|^2 = \tilde{O}\left(\frac{s \cdot \log(1/\delta)}{\epsilon^2}\right) = \tilde{O}\left(\frac{\log^3 n}{\epsilon^3 \delta}\right)$. Additionally, by a standard Chernoff bound at most $\tilde{O}\left(\frac{\log^3 n}{\epsilon^3 \delta}\right)$ are read with probability at least $1 - \delta$. $\square$

G Singular Value Approximation via Sampling

We now show how to estimate the singular values of a bounded-entry matrix via random subsampling. Unlike in eigenvalue estimation, instead of sampling a random principal submatrix, here we sample a random submatrix with independent rows and columns. This allows us to apply known interior eigenvalue matrix Chernoff bounds to bound the perturbation in the singular values [11, 12]. We first state a simplified version of Theorem 4.1 from [12] (also stated as Theorem 4.6 in [11]), simplified using standard upper bounds on the Chernoff bounds in [57].

Theorem 11 (Interior Eigenvalue Matrix Chernoff bounds—Theorem 4.1 of [12]) *Let $\{\mathbf{X}_j\}$ be a finite sequence of independent, random, positive-semidefinite matrices with dimension n , and assume that $\|\mathbf{X}_j\|_2 \leq L$ for some value L almost surely. Given an integer $k \leq n$, define*

$$\mu_k = \lambda_k \left(\sum_j \mathbb{E}[\mathbf{X}_j] \right).$$

Then we have the tail inequalities:

$$\begin{cases} \mathbb{P} \left(\lambda_k \left(\sum_j \mathbf{X}_j \right) \geq (1 + \Delta) \mu_k \right) \leq (n - k + 1) \cdot e^{-\frac{\Delta \mu_k}{3L}}, & \text{for } \Delta \geq 1 \\ \mathbb{P} \left(\lambda_k \left(\sum_j \mathbf{X}_j \right) \geq (1 + \Delta) \mu_k \right) \leq (n - k + 1) \cdot e^{-\frac{\Delta^2 \mu_k}{3L}}, & \text{for } \Delta \in [0, 1) \\ \mathbb{P} \left(\lambda_k \left(\sum_j \mathbf{X}_j \right) \leq (1 - \Delta) \mu_k \right) \leq k \cdot e^{-\frac{\Delta^2 \mu_k}{2L}}, & \text{for } \Delta \in [0, 1) \end{cases}$$

We are now ready to state and prove the main theorem.

Theorem 12 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a matrix with $\|\mathbf{A}\|_\infty \leq 1$ and singular values $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_n(\mathbf{A})$. Let $\tilde{\mathbf{S}} \in \mathbb{R}^{n \times n}$ be a scaled diagonal sampling matrix such that $\tilde{\mathbf{S}}_{ii} = \sqrt{\frac{n}{s}}$ with probability $\frac{s}{n}$ and $\tilde{\mathbf{S}}_{ii} = 0$ otherwise. Let $\tilde{\mathbf{T}} \in \mathbb{R}^{n \times n}$ be an independent and identically distributed random sampling matrix. Let $\mathbf{Z} = \tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}$ be the sampled submatrix from $\mathbf{A}$ with singular values $\sigma_1(\mathbf{Z}) \geq \dots \geq \sigma_n(\mathbf{Z})$. Then, if $s \geq \frac{c \log(n/\delta)}{\epsilon^2}$ for some constant c , with probability at least $1 - \delta$, for all $i \in [n]$,*

$$\sigma_i(\mathbf{A}) - \epsilon n \leq \sigma_i(\mathbf{Z}) \leq \sigma_i(\mathbf{A}) + \epsilon n.$$

Proof We first prove that singular values of $\tilde{\mathbf{S}}\mathbf{A}$ are close to those of $\mathbf{A}$. Let $\mathbf{X}_i \in \mathbb{R}^{n \times n}$ be matrix valued r.v.'s for $i \in [n]$ such that:

$$\mathbf{X}_i = \begin{cases} \frac{n}{s} \mathbf{A}_i \mathbf{A}_i^T, & \text{with probability } s/n \\ 0 & \text{otherwise} \end{cases}$$

where $\mathbf{A}_i$ is the i^{th} row of $\mathbf{A}$ written as a column vector. Then, $\sum_i \mathbf{X}_i = (\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})$ and $\mathbb{E}[\sum_i \mathbf{X}_i] = \mathbf{A}^T \mathbf{A}$. We have $\|\mathbf{X}_i\|_2 \leq \max_j \frac{n}{s} \|\mathbf{A}_j\|_2^2 \leq \frac{n^2}{s}$ and $\lambda_k(\mathbb{E}[\sum_i \mathbf{X}_i]) = \lambda_k(\mathbf{A}^T \mathbf{A}) = \sigma_k^2(\mathbf{A})$ for $k \in [n]$.

Case 1 We will first prove that $\sigma_k(\mathbf{A}) - \epsilon n \leq \sigma_k(\bar{\mathbf{S}}\mathbf{A})$ for all $k \in [n]$. Note that when $\sigma_k(\mathbf{A}) \leq \epsilon n$, $\sigma_k(\mathbf{A}) - \epsilon n \leq \sigma_k(\bar{\mathbf{S}}\mathbf{A})$ is trivially true. We now consider all $k \in [n]$ such that $\sigma_k(\mathbf{A}) > \epsilon n$. Setting $\mu_k = \lambda_k(\mathbf{A}^T \mathbf{A})$, $L = \frac{n^2}{s}$ and $\Delta = \frac{\epsilon n}{\sigma_k(\mathbf{A})}$ (note that $\Delta < 1$) in Theorem 11, we get:

$$\mathbb{P}\left(\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) \leq (1 - \Delta)\lambda_k(\mathbf{A}^T \mathbf{A})\right) \leq k \cdot e^{-c \frac{\Delta^2 \lambda_k(\mathbf{A}^T \mathbf{A})}{L}} \leq k \cdot e^{-c \frac{\epsilon^2 n^2}{\lambda_k(\mathbf{A}^T \mathbf{A})} \cdot \frac{\lambda_k(\mathbf{A}^T \mathbf{A})}{(n^2/s)}}$$

where c is constant. So, for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$ for any k , we have $\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) = \sigma_k^2(\bar{\mathbf{S}}\mathbf{A}) \geq (1 - \Delta)\sigma_k^2(\mathbf{A})$ with probability at least $1 - \frac{\delta}{n}$. Taking a square root on both sides we get $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \geq \sqrt{1 - \Delta}\sigma_k(\mathbf{A}) \geq (1 - \Delta)\sigma_k(\mathbf{A}) = \sigma_k(\mathbf{A}) - \epsilon n$. Taking a union bound over all k with $\sigma_k(\mathbf{A}) > \epsilon n$, $\sigma_k(\mathbf{A}) - \epsilon n \leq \sigma_k(\bar{\mathbf{S}}\mathbf{A})$ holds for all such k with probability at least $1 - \delta$.

Case 2 We now prove that $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \leq \sigma_k(\mathbf{A}) + \epsilon n$ for all $k \in [n]$. We first consider the case when $\sigma_k(\mathbf{A}) \leq \epsilon n$. Setting $\mu_k = \lambda_k(\mathbf{A}^T \mathbf{A})$, $L = \frac{n^2}{s}$ and $\Delta = \frac{\epsilon^2 n^2}{\sigma_k^2(\mathbf{A})}$ (note that $\Delta \geq 1$) in Theorem 11, we get (for some constant c):

$$\begin{aligned} \mathbb{P}\left(\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) \geq (1 + \Delta)\lambda_k(\mathbf{A}^T \mathbf{A})\right) &\leq n \cdot e^{-\frac{c\Delta\lambda_k(\mathbf{A}^T \mathbf{A})}{L}} \\ &\leq n \cdot e^{-\frac{c\epsilon^2 n^2}{\lambda_k(\mathbf{A}^T \mathbf{A})} \cdot \frac{\lambda_k(\mathbf{A}^T \mathbf{A})}{(n^2/s)}} \end{aligned}$$

Thus, if $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$, we have $\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) \leq (1 + \Delta)\lambda_k(\mathbf{A}^T \mathbf{A}) \leq \lambda_k(\mathbf{A}^T \mathbf{A}) + \epsilon^2 n^2$ for all $k \in [n]$ such that $\sigma_k(\mathbf{A}) \leq \epsilon n$ with probability at least $1 - \delta$ via a union bound. Taking square root on both sides and using the facts that $\lambda_k(\mathbf{A}^T \mathbf{A}) = \sigma_k^2(\mathbf{A})$, $\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) = \sigma_k^2(\bar{\mathbf{S}}\mathbf{A})$ and $\sqrt{a+b} < \sqrt{a} + \sqrt{b}$, we get $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \leq \sigma_k(\mathbf{A}) + \epsilon n$.

We now consider the case $\sigma_k(\mathbf{A}) > \epsilon n$. Setting $\mu_k = \lambda_k(\mathbf{A}^T \mathbf{A})$, $L = \frac{n^2}{s}$ and $\Delta = \frac{\epsilon n}{\sigma_k(\mathbf{A})}$ (note that $\Delta < 1$) in Theorem 11, we get (for some constant c):

$$\begin{aligned} \mathbb{P}\left(\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) \geq (1 + \Delta)\lambda_k(\mathbf{A}^T \mathbf{A})\right) &\leq n \cdot e^{-\frac{c\Delta^2 \lambda_k(\mathbf{A}^T \mathbf{A})}{L}} \\ &\leq n \cdot e^{-\frac{c\epsilon^2 n^2}{\lambda_k(\mathbf{A}^T \mathbf{A})} \cdot \frac{\lambda_k(\mathbf{A}^T \mathbf{A})}{(n^2/s)}}. \end{aligned}$$

Thus, if $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$, we have $\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) \leq (1 + \Delta)\lambda_k(\mathbf{A}^T \mathbf{A})$ for all $k \in [n]$ such that $\sigma_k(\mathbf{A}) > \epsilon n$ with probability at least $1 - \delta$ via a union bound. Taking square root on both sides and using the fact that $\lambda_k(\mathbf{A}^T \mathbf{A}) = \sigma_k^2(\mathbf{A})$, $\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T (\bar{\mathbf{S}}\mathbf{A})) = \sigma_k^2(\bar{\mathbf{S}}\mathbf{A})$ and $\sqrt{a} < a$ for any $a > 1$, we get $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \leq (1 + \Delta)\sigma_k(\mathbf{A}) \leq \sigma_k(\mathbf{A}) + \epsilon n$.

Thus, via a union bound over all $k \in [n]$, we have $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \leq \sigma_k(\mathbf{A}) + \epsilon n$ with probability $1 - 2\delta$.

Thus, via a union bound over the two cases above, for all $k \in [n]$ with probability at least $1 - 3\delta$ for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$ we have, for all $k \in [n]$,

$$|\sigma_k(\bar{\mathbf{S}}\mathbf{A}) - \sigma_k(\mathbf{A})| \leq \epsilon n. \quad (56)$$

Next we prove that the singular values of $\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}}$ are close to those of $\bar{\mathbf{S}}\mathbf{A}$, using essentially the same approach as above. Let $\mathbf{Y}_i$ be a matrix valued random variable for $i \in [n]$ such that:

$$\mathbf{Y}_i = \begin{cases} \frac{n}{s}(\bar{\mathbf{S}}\mathbf{A})_i(\bar{\mathbf{S}}\mathbf{A})_i^T, & \text{with probability } s/n \\ 0 & \text{otherwise} \end{cases}$$

where $(\bar{\mathbf{S}}\mathbf{A})_i$ is the i^{th} column of $\bar{\mathbf{S}}\mathbf{A}$. Then, $\sum_i \mathbf{Y}_i = (\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})^T(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})$. Also, we have $\lambda_k(\mathbb{E}[\sum_i \mathbf{Y}_i]) = \lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A})) = \sigma_k^2(\bar{\mathbf{S}}\mathbf{A})$. First, using a standard Chernoff bound, we can claim that $\bar{\mathbf{S}}$ will sample at most $2s$ rows from $\mathbf{A}$ with probability at least $1 - \delta$ for any $s \geq O(\log(1/\delta))$. Thus, we have $\|\mathbf{Y}_i\|_2 = \frac{n}{s}\|\bar{\mathbf{S}}\mathbf{A}\|_2^2 \leq \frac{n}{s} \cdot \frac{n}{s} \cdot 2s \leq \frac{2n^2}{s}$ with probability $1 - \delta$. Let this event be called E_2 . We now consider two cases conditioned on the event E_2 .

Case 1 We first prove that $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) - \epsilon n \leq \sigma_k(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})$ for all $k \in [n]$. Again note that when $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \leq \epsilon n$ this is trivially true. So we consider all $k \in [n]$ such that $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) > \epsilon n$. Setting $\mu_k = \lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))$, $L = \frac{2n^2}{s}$ (as we have conditioned on E_2) and $\Delta = \frac{\epsilon n}{\sigma_k(\bar{\mathbf{S}}\mathbf{A})}$ (note that $\Delta < 1$) in Theorem 11, we get:

$$\begin{aligned} \mathbb{P}\left(\lambda_k((\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})^T(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})) \leq (1 - \Delta)\lambda_k(\mathbf{A}^T\mathbf{A})\right) \\ \leq k \cdot e^{-c \frac{\Delta^2 \lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))}{L}} \leq k \cdot e^{-c \frac{\epsilon^2 n^2}{\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))} \cdot \frac{\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))}{(n^2/s)}} \end{aligned}$$

where c is some constant. So, for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$ for any k , we have $\lambda_k((\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})^T(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})) = \sigma_k^2(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}}) \geq (1 - \Delta)\sigma_k^2(\bar{\mathbf{S}}\mathbf{A})$ with probability at least $1 - \frac{\delta}{n}$. Taking a square root on both sides we get $\sigma_k(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}}) \geq \sqrt{1 - \Delta}\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \geq (1 - \Delta)\sigma_k(\bar{\mathbf{S}}\mathbf{A}) = \sigma_k(\bar{\mathbf{S}}\mathbf{A}) - \epsilon n$. Taking a union bound over all k with $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) > \epsilon n$, $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) - \epsilon n \leq \sigma_k(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})$ holds for all such k with probability at least $1 - \delta$.

Case 2 We now prove $\sigma_k(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}}) \leq \sigma_k(\bar{\mathbf{S}}\mathbf{A}) + \epsilon n$ for all $k \in [n]$. We again first consider the case $\sigma_k(\bar{\mathbf{S}}\mathbf{A}) \leq \epsilon n$. Setting $\mu_k = \lambda_k(\mathbf{A}^T\mathbf{A})$, $L = \frac{n^2}{s}$ and $\Delta = \frac{\epsilon^2 n^2}{\sigma_k^2(\bar{\mathbf{S}}\mathbf{A})}$ (note that $\Delta \geq 1$) in Theorem 11:

$$\begin{aligned} \mathbb{P}\left(\lambda_k((\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})^T(\bar{\mathbf{S}}\mathbf{A}\bar{\mathbf{T}})) \geq (1 + \Delta)\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))\right) &\leq n \cdot e^{-\frac{c\Delta\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))}{L}} \\ &\leq n \cdot e^{-\frac{c\epsilon^2 n^2}{\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))} \cdot \frac{\lambda_k((\bar{\mathbf{S}}\mathbf{A})^T(\bar{\mathbf{S}}\mathbf{A}))}{(n^2/s)}} \end{aligned}$$

Then, similar to the case $\sigma_k(\mathbf{A}) \leq \epsilon n$ in the previous case 2, taking square root of both sides and via a union bound, we get $\sigma_k(\tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}) \leq \sigma_k(\tilde{\mathbf{S}}\mathbf{A}) + \epsilon n$ for all $k \in [n]$ such that $\sigma_k(\tilde{\mathbf{S}}\mathbf{A}) \leq \epsilon n$ with probability at least $1 - \delta$ for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$. The case $\sigma_k(\tilde{\mathbf{S}}\mathbf{A}) > \epsilon n$ will again be similar as $\sigma_k(\mathbf{A}) > \epsilon n$ in the previous case 2. We set $\Delta = \frac{\epsilon n}{\sigma_k(\tilde{\mathbf{S}}\mathbf{A})}$ and apply Theorem 11 and take the square root on both sides to get $\sigma_k(\tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}) \leq \sigma_k(\tilde{\mathbf{S}}\mathbf{A}) + \epsilon n$ with probability $1 - \delta$ for all $k \in [n]$ for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$. Thus, with probability $1 - 2\delta$, conditioned on the event E_2 , we have $\sigma_k(\tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}) \leq \sigma_k(\tilde{\mathbf{S}}\mathbf{A}) + \epsilon n$ for all $k \in [n]$. Finally, via a union bound over the two cases above, and conditioned on E_2 , for all $k \in [n]$ with probability at least $1 - 2\delta$ for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$ we get

$$|\sigma_k(\tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}) - \sigma_k(\tilde{\mathbf{S}}\mathbf{A})| \leq \epsilon n. \quad (57)$$

Thus, taking a union bound over all the cases above (including E_2), from equation (56) and (57) and via a triangle inequality, we get: $|\sigma_k(\tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}) - \sigma_k(\mathbf{A})| \leq 2\epsilon n$ with probability at least $1 - c\delta$ (where c is a small constant) for $s \geq O(\frac{\log(n/\delta)}{\epsilon^2})$. Adjusting ϵ and δ by constant factors gives us the final bound. $\square$

Remark on Rectangular Matrices Though we have considered $\mathbf{A}$ to be a square matrix for simplicity, notice that Theorem 12 also holds for any arbitrary (non-square) matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, with n replaced by $\max(n, m)$ in the sample complexity bound.

Remark on Non-Uniform Sampling As discussed in Sect. 1.3.1, simple non-uniform random submatrix sampling via row/column sparsities or norms does not suffice to estimate the singular values up to improved error bounds of $\epsilon\sqrt{\text{nnz}(\mathbf{A})}$ or $\epsilon\|\mathbf{A}\|_F$. A more complex strategy, such as the zeroing out used in Theorems 2 and 3 must be used. It is worth noting that following the same proof as Theorem 12, it is easy to show that if $\tilde{\mathbf{S}}$ is sampled according to row norms or sparsities and appropriately weighted, then the singular values of $\tilde{\mathbf{S}}\mathbf{A}$ do approximate those of $\mathbf{A}$ up to these improved error bounds. The proof breaks down when analyzing $\tilde{\mathbf{S}}\mathbf{A}\tilde{\mathbf{T}}$. $\tilde{\mathbf{T}}$ would have to be sampled according to the row norms/sparsities of $\tilde{\mathbf{S}}\mathbf{A}$, not $\mathbf{A}$, for the proof to go through. However, in general, these sampling probabilities can differ significantly between $\tilde{\mathbf{S}}\mathbf{A}$ and $\mathbf{A}$.

References

- Demmel, J., Dumitriu, I., Holtz, O., Kleinberg, R.: Fast matrix multiplication is stable. *Numer. Math.* **106**(2), 199–224 (2007)
- Alman, J., Williams, V.V.: A refined laser method and faster matrix multiplication. In: *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, vol. 32, pp. 522–539 (2021)
- Saad, Y.: *Numerical Methods for Large Eigenvalue Problems*, Revised Edition. SIAM (2011)
- Gu, M., Eisenstat, S.C.: A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. *SIAM J. Matrix Anal. Appl.* **16**(1), 172–191 (1995)
- Greengard, L., Strain, J.: The fast gauss transform. *SIAM J. Sci. Stat. Comput.* **12**(1), 79–94 (1991)
- Hardt, M., Price, E.: The noisy power method: a meta algorithm with applications. *Adv. Neural Inf. Process. Syst.* **27** (2014)
- Backurs, A., Indyk, P., Musco, C., Wagner, T.: Faster kernel matrix algebra via density estimation. In: *Proceedings of the 38th International Conference on Machine Learning (ICML)*, vol. 38, pp. 500–510 (2021)

8. Achlioptas, D., McSherry, F.: Fast computation of low-rank matrix approximations. *J. ACM (JACM)* **54**, 9–27 (2007)
9. Rudelson, M., Vershynin, R.: Sampling from large matrices: an approach through geometric functional analysis. *J. ACM (JACM)* **54**(4), 21 (2007)
10. Tropp, J.A.: Norms of random submatrices and sparse approximation. *C.R. Math.* **346**(23–24), 1271–1274 (2008)
11. Bakshi, A., Chepurko, N., Jayaram, R.: Testing positive semi-definiteness via random submatrices. In: *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, vol. 61, pp. 1191–1202 (2020)
12. Gittens, A., Tropp, J.A.: Tail bounds for all eigenvalues of a sum of random matrices. [arXiv:1104.4513](https://arxiv.org/abs/1104.4513) (2011)
13. Goldreich, O., Ron, D.: Property testing in bounded degree graphs. In: *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*, vol. 29, pp. 406–415 (1997)
14. Frieze, A., Kannan, R., Vempala, S.: Fast Monte-Carlo algorithms for finding low-rank approximations. *J. ACM (JACM)* **51**, 1025–1041 (2004)
15. Tang, E.: Quantum-inspired classical algorithms for principal component analysis and supervised clustering. [arXiv:1811.00414](https://arxiv.org/abs/1811.00414) (2018)
16. Benjamini, I., Schramm, O., Shapira, A.: Every minor-closed property of sparse graphs is testable. *Adv. Math.* **40**, 393–402 (2010)
17. Weiße, A., Wellein, G., Alvermann, A., Fehske, H.: The kernel polynomial method. *Rev. Mod. Phys.* **78**(1), 275 (2006)
18. Lin, L., Saad, Y., Yang, C.: Approximating spectral densities of large matrices. *SIAM Rev.* **58**(1), 34–65 (2016)
19. Dong, K., Benson, A.R., Bindel, D.: Network density of states. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, vol. 25, pp. 1152–1161 (2019)
20. Sagun, L., Bottou, L., LeCun, Y.: Eigenvalues of the hessian in deep learning: singularity and beyond. [arXiv:1611.07476](https://arxiv.org/abs/1611.07476) (2016)
21. Yao, Z., Gholami, A., Lei, Q., Keutzer, K., Mahoney, M.W.: Hessian-based analysis of large batch training and robustness to adversaries. [arXiv:1802.08241](https://arxiv.org/abs/1802.08241) (2018)
22. Ghorbani, B., Krishnan, S., Xiao, Y.: An investigation into neural net optimization via hessian eigenvalue density. In: *Proceedings of the 36th International Conference on Machine Learning (ICML)*, vol. 36, pp. 2232–2241 (2019)
23. Li, R., Xi, Y., Erlandson, L., Saad, Y.: The eigenvalues slicing library (EVSL): algorithms, implementation, and software. *SIAM J. Sci. Comput.* **41**, C393–C415 (2019)
24. Wang, L.-W.: Calculating the density of states and optical-absorption spectra of large quantum systems by the plane-wave moments method. *Phys. Rev. B* **49**(15), 10154 (1994)
25. Silver, R., Röder, H.: Densities of states of mega-dimensional Hamiltonian matrices. *Int. J. Mod. Phys. C* **5**(04), 735–753 (1994)
26. Helsen, J., Battistel, F., Terhal, B.M.: Spectral quantum tomography. *Quantum Inf.* **5**(1), 74 (2019)
27. Braverman, V., Krishnan, A., Musco, C.: Linear and sublinear time spectral density estimation. In: *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)* (2022)
28. Cohen-Steiner, D., Kong, W., Sohler, C., Valiant, G.: Approximating the spectrum of a graph. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, vol. 24, pp. 1263–1271 (2018)
29. Balcan, M.-F., Li, Y., Woodruff, D.P., Zhang, H.: Testing matrix rank, optimally. In: *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, vol. 30, pp. 727–746 (2019)
30. Krauthgamer, R., Sasson, O.: Property testing of data dimensionality. In: *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, vol. 3, pp. 18–27 (2003)
31. Li, Y., Wang, Z., Woodruff, D.P.: Improved testing of low rank matrices. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, vol. 20, pp. 691–700 (2014)
32. Needell, D., Swartworth, W., Woodruff, D.P.: Testing positive semidefiniteness using linear measurements. In: *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, vol. 63, pp. 87–97 (2022)

33. Andoni, A., Nguyễn, H.L.: Eigenvalues of a matrix in the streaming model. In: Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), vol. 24, pp. 1729–1737 (2013)
34. Li, Y., Nguyen, H.L., Woodruff, D.P.: On sketching matrix norms and the top singular vector. In: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), vol. 25, pp. 1562–1581 (2014)
35. Li, Y., Woodruff, D.P.: On approximating functions of the singular values in a stream. In: Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC), vol. 48, pp. 726–739 (2016)
36. Weyl, H.: The asymptotic distribution law of the eigenvalues of linear partial differential equations (with an application to the theory of cavity radiation). *Math. Ann.* **71**(1), 441–479 (1912)
37. Stewart, G.W., Sun, J.: *Matrix Perturbation Theory*. Academic Press, New York (1990)
38. Horn, R.A., Johnson, C.R.: *Matrix Analysis*. Cambridge University Press, New York (2012)
39. Kahan, W.: Spectra of nearly Hermitian matrices. *Proc. Am. Math. Soc.* **48**(1), 11–17 (1975)
40. Bhatia, R.: *Matrix Analysis*. Springer, New York (2013)
41. Goldreich, O., Ron, D.: Approximating average parameters of graphs. *Rand. Struct. Algorithms* **32**, 473–493 (2008)
42. Eden, T., Rosenbaum, W.: On sampling edges almost uniformly. In: *SIAM Symposium on Simplicity in Algorithms (SOSA)* (2018)
43. Chepurko, N., Clarkson, K.L., Horeh, L., Lin, H., Woodruff, D.P.: Quantum-inspired algorithms from randomized numerical linear algebra. [arXiv:2011.04125](https://arxiv.org/abs/2011.04125) (2020)
44. Drineas, P., Kannan, R.: Fast Monte-Carlo algorithms for approximate matrix multiplication. In: Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), vol. 42, pp. 452–459 (2001)
45. Gershgorin, S.A.: Über die abgrenzung der eigenwerte einer matrix. *Izvestiya Rossiyskoy akademii nauk. Seriya matematicheskaya* 749–754 (1931)
46. Tropp, J.A.: The random paving property for uniformly bounded matrices. *Stud. Math.* **185**, 67–82 (2008)
47. Tropp, J.A.: An introduction to matrix concentration inequalities. [arXiv:1501.01571](https://arxiv.org/abs/1501.01571) (2015)
48. Bernstein, S.: Sur l’extension du théorème limite du calcul des probabilités aux sommes de quantités dépendantes. *Math. Ann.* **97**, 1–59 (1927)
49. Cai, D., Nagy, J., Xi, Y.: Fast and stable deterministic approximation of general symmetric kernel matrices in high dimensions. [arXiv:2102.05215](https://arxiv.org/abs/2102.05215) (2021)
50. Leskovec, J., Krevl, A.: SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data> (2014)
51. McAuley, J.J., Leskovec, J.: Learning to discover social circles in ego networks. *Adv. Neural Inf. Process. Syst.* **25** (2012)
52. Leskovec, J., Kleinberg, J., Faloutsos, C.: Graph evolution: densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data TKDD* **1**(1), 2 (2007)
53. Community, T.N.: `numpy.linalg.eigvals`. <https://numpy.org/doc/stable/reference/generated/numpy.linalg.eigvals.html> (2021)
54. Erdős, P., Rényi, A.: On random graphs I. *Publ. Math. Debr.* **6**(290–297), 18 (1959)
55. Cohen, M.B., Lee, Y.T., Musco, C., Musco, C., Peng, R., Sidford, A.: Uniform sampling for matrix approximation. In: Proceedings of the 6th Conference on Innovations in Theoretical Computer Science (ITCS), vol. 6, pp. 181–190 (2015)
56. Braverman, V., Krauthgamer, R., Krishnan, A.R., Sapir, S.: Near-optimal entrywise sampling of numerically sparse matrices. In: Proceedings of the 34th Annual Conference on Computational Learning Theory (COLT), vol. 34, pp. 759–773 (2021)
57. Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, Cambridge (2017)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.