

Improved Spectral Density Estimation via Explicit and Implicit Deflation

Rajarshi Bhattacharjee
UMass Amherst
rbhattacharj@cs.umass.edu

Rajesh Jayaram
Google Research
rkjayaram@google.com

Cameron Musco
UMass Amherst
cmusco@cs.mass.edu

Christopher Musco
New York University
cmusco@nyu.edu

Archan Ray
Sloan Kettering Institute
talk2archan@gmail.com

Abstract

We study algorithms for approximating the spectral density (i.e., the eigenvalue distribution) of a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ that is accessed through matrix-vector product queries. Recent work has analyzed popular Krylov subspace methods for this problem, showing that they output an $\epsilon \cdot \|\mathbf{A}\|_2$ error approximation to the spectral density in the Wasserstein-1 metric using $O(1/\epsilon)$ matrix-vector products. By combining a previously studied Chebyshev polynomial moment matching method with a *deflation* step that approximately projects off the largest magnitude eigendirections of $\mathbf{A}$ before estimating the spectral density, we give an improved error bound of $\epsilon \cdot \sigma_\ell(\mathbf{A})$ using $O(\ell \log n + 1/\epsilon)$ matrix-vector products, where $\sigma_\ell(\mathbf{A})$ is the ℓ^{th} largest singular value of $\mathbf{A}$. In the common case when $\mathbf{A}$ exhibits fast singular value decay and so $\sigma_\ell(\mathbf{A}) \ll \|\mathbf{A}\|_2$, our bound can be much stronger than prior work. We also show that it is nearly tight: any algorithm giving error $\epsilon \cdot \sigma_\ell(\mathbf{A})$ must use $\Omega(\ell + 1/\epsilon)$ matrix-vector products.

We further show that the popular Stochastic Lanczos Quadrature (SLQ) method essentially matches the above bound *for any choice of parameter ℓ* , even though SLQ itself is parameter-free and performs no explicit deflation. Our bound helps to explain the strong practical performance and observed ‘spectrum adaptive’ nature of SLQ, and motivates a simple variant of the method that achieves an even tighter error bound. Technically, our results require a careful analysis of how eigenvalues and eigenvectors are approximated by (block) Krylov subspace methods, which may be of independent interest. Our error bound for SLQ leverages an analysis of the method that views it as an implicit polynomial moment matching method, along with recent results on low-rank approximation with single-vector Krylov methods. We use these results to show that the method can perform ‘implicit deflation’ as part of moment matching.

Contents

1	Introduction	2
1.1	Matrix-Vector Query Algorithms for SDE	2
1.2	Existing Bounds	3
1.3	Our Results	4
1.3.1	Improved SDE via Moment Matching with Explicit Deflation	4
1.3.2	Implicit Deflation Bounds for Stochastic Lanczos Quadrature	5
1.4	Technical Overview	6
1.4.1	Eigenvalue Deflation for SDE	6
1.4.2	Error Analysis of Deflation	7
1.4.3	SLQ and its Existing Analysis	7
1.4.4	Moment Matching-Based Analysis of SLQ	8
1.4.5	Implicit Deflation with SLQ	9
1.4.6	Variance Reduced SLQ	10
1.5	Roadmap	10
2	Notation and Preliminaries	10
2.1	Basic Notation	10
2.2	Linear Algebra Preliminaries	11
2.3	Moment Matching and Wasserstein Distance Preliminaries	12
3	SDE via Moment Matching with Explicit Deflation	13
3.1	Error bounds for Deflation via Block Krylov	14
3.1.1	Eigenvector Alignment	17
3.1.2	Eigenvalue Alignment	20
3.1.3	Bounding the spectral norm after deflation	23
3.2	Error Bounds for moment matching with deflation	25
4	Analysis of Stochastic Lanczos Quadrature	31
4.1	SLQ bounds via Moment Matching	32
4.2	Error Bounds for Lanczos	36
4.3	Improved Error bounds for SLQ via implicit deflation based analysis	41
4.4	Variance reduced SLQ	61
5	Lower Bound	68
6	Empirical Evaluation	69
6.1	Datasets	69
6.2	Implementation Details	70
6.3	Summary of Results	71

1 Introduction

Spectral density estimation (SDE) is a fundamental task in computational linear algebra. Given a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with eigenvalues $\lambda_1(\mathbf{A}), \dots, \lambda_n(\mathbf{A})$, the goal is to approximate $\mathbf{A}$'s eigenvalue distribution (i.e., its spectral density) $s_{\mathbf{A}}$, which is the distribution that places probability mass $1/n$ at each of $\mathbf{A}$'s n eigenvalues. Formally, letting $\delta(\cdot)$ be the Dirac delta function,

$$s_{\mathbf{A}}(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - \lambda_i(\mathbf{A})). \quad (1)$$

Understanding $s_{\mathbf{A}}$ can provide important information about the input matrix $\mathbf{A}$. To what degree does it exhibit low-rank structure (i.e., have a decaying eigenvalues)? How close is the spectrum, or some part of the spectrum, to that of a random matrix? Does $\mathbf{A}$ have many repeated or nearly repeated eigenvalues that may indicate anomalies or other interesting structure? As such, spectral density estimation is applied throughout the sciences [Ski89, SR94, STBB17, SRS20], network science [FDBV01, EG17, DBB19], machine learning and deep learning in particular [RL18, PSG18, MM19, GKX19], numerical linear algebra [DNPS16, LXES19], and beyond.

The spectral density $s_{\mathbf{A}}$ can be computed directly by performing a full eigendecomposition of $\mathbf{A}$, in $O(n^\omega)$ time, for $\omega \approx 2.37$ being the exponent of fast matrix multiplication [Par98, PC99, DDH07]. However, when $\mathbf{A}$ is very large, or in settings where $\mathbf{A}$ can only be accessed through a small number of queries, we often seek an approximation $\tilde{s}_{\mathbf{A}}$, such that $\tilde{s}_{\mathbf{A}}$ and $s_{\mathbf{A}}$ are close in some metric. In this work we will focus on the Wasserstein-1 (i.e., earth mover's) distance, $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}})$, which has been studied in a number of recent works giving formal approximation guarantees for SDE [CSKSV18, CTU21, BKM22, CTU22, JMSS23, JKM⁺24]. When $\tilde{s}_{\mathbf{A}}$ is the uniform distribution over approximate eigenvalues $\tilde{\lambda}_1(\mathbf{A}), \dots, \tilde{\lambda}_n(\mathbf{A})$, and when we order both sets of eigenvalues in decreasing order, $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) = \frac{1}{n} \sum_{i=1}^n |\lambda_i(\mathbf{A}) - \tilde{\lambda}_i(\mathbf{A})|$. I.e., it is the average absolute error of our eigenvalue estimates. More generally, when $\tilde{s}_{\mathbf{A}}$ is any distribution, $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}})$ is the minimum cost of transforming $s_{\mathbf{A}}$ into $\tilde{s}_{\mathbf{A}}$, where moving probability δ from x to y incurs cost $\delta \cdot |x - y|$.

1.1 Matrix-Vector Query Algorithms for SDE

Given its practical importance, efficient algorithms for SDE have been widely studied [BRP92, Wan94, WWAF06, LSY16, CTU21, BKM22]. A large fraction of these methods operate in the *matrix-vector query model*: they only access the input matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ through multiplication on the left or right with a sequence of (possibly adaptively chosen) query vectors $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathbb{R}^n$. The goal is to minimize the number of queries m , which typically dominate the runtime cost.

Matrix-vector query algorithms encompass both linear sketching methods (when queries are chosen non-adaptively) and Krylov subspace methods (when queries are of the form $\mathbf{x}, \mathbf{A}\mathbf{x}, \mathbf{A}^2\mathbf{x}, \dots$, for some starting vector $\mathbf{x}$, or set of starting vectors). Beyond spectral density estimation, they are the dominant algorithms in practice for many linear algebraic problems, including eigenvalue and eigenvector computation [Par98], low-rank approximation [HMT11, MM15], linear system solving [LSY98, Saa03], and beyond. Such methods typically have low-memory overhead, since even when $\mathbf{A}$ is very large, they only need to store the outputs of the matrix-vector products. Further, they can often take advantage of highly optimized software and hardware for matrix-vector multiplication, including parallel hardware like GPUs, and faster matrix-vector multiplication routines when $\mathbf{A}$ is sparse or structured. Moreover, matrix-vector query algorithms are applicable in settings where $\mathbf{A}$

cannot be efficiently materialized, but can be efficiently multiplied by vectors. This is the case e.g., when $\mathbf{A}$ is the Hessian of a neural network [Pea94, GKX19] or a function of some other matrix that can be efficiently applied to vectors using e.g., an iterative method [UCS17].

Recently, matrix-vector query algorithms have received significant attention in theoretical work on numerical linear algebra since in many cases, it is possible to prove (nearly) matching query complexity upper and lower bounds for central problems like trace estimation [MMM21], low-rank approximation [SEAR18, BCW22, BN23], linear regression [BHSW20], structured matrix approximation [HT23, DM23, ACH⁺24], and beyond [SWYZ21, NSW22, SW23].

Most state-of-the-art matrix-vector query algorithms for spectral density estimation are Krylov subspace methods that fall into two general classes.

Moment Matching. The first class of methods approximates $s_{\mathbf{A}}$ by approximating its polynomial moments. I.e., $\mathbb{E}_{s_{\mathbf{A}}}[p(x)] = \frac{1}{n} \sum_{i=1}^n p(\lambda_i(\mathbf{A})) = \frac{1}{n} \text{tr}(p(\mathbf{A}))$, where p is a low-degree polynomial. We can employ stochastic trace estimation methods like Hutchinson’s method [Gir87, Hut90] to approximate this trace using just a small number of matrix-vector products with $p(\mathbf{A})$ and in turn $\mathbf{A}$, since if p has degree k , a single matrix-vector product with $p(\mathbf{A})$ can be performed using k matrix vector products with $\mathbf{A}$. After approximating the moments for a set of low-degree polynomials (e.g., the first k monomials, or the first k Chebyshev polynomials), we can let $\tilde{s}_{\mathbf{A}}$ be a distribution that matches these moments as closely as possible, and thus should closely match $s_{\mathbf{A}}$.

Moment matching methods include the popular Kernel Polynomial Method (KPM) [SR94, Wan94, WWA06] and its variants [CPB10, LSY16, BKM22, Che23]. Several works also use moment matching to give sublinear time SDE methods for graph adjacency matrices [CSKSV18, BKM22, JMSS23], leveraging structure to estimate moments faster than with matrix-vector queries.

Lanczos-Based Methods. The second class of methods computes a small number of approximate eigenvalues of $\mathbf{A}$ using the Lanczos method, and lets $\tilde{s}_{\mathbf{A}}$ be a distribution supported on these eigenvalues, with appropriately chosen probability mass placed at each. The canonical method of this form is Stochastic Lanczos Quadrature (SLQ) [CTU21]. Many other variants have also been studied. Some place probability mass not just at the approximate eigenvalues, but on Gaussian or other simple distributions centered at these eigenvalues [LG82, BRP92, LSY16, HHK72].

1.2 Existing Bounds

While matrix-vector query algorithms for SDE have been studied for decades, theoretical guarantees on their approximation error in terms of the distance between the true spectral density $s_{\mathbf{A}}$ and the approximate density $\tilde{s}_{\mathbf{A}}$ have only recently been formalized. Braverman et al. [BKM22] analyze a Chebyshev Moment Matching method, which can be thought of as a simple variant of KPM, showing that the method can compute $\tilde{s}_{\mathbf{A}}$ satisfying $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \cdot \|\mathbf{A}\|_2$ with probability $\geq 1 - \delta$ using just $O(b/\epsilon)$ matrix vector products, where $b = \max(1, \frac{1}{n\epsilon^2} \log^2 \frac{1}{\epsilon\delta} \log^2 \frac{1}{\epsilon})$. Note that $b = 1$ in the common case when $\epsilon = \tilde{\Omega}(1/\sqrt{n})$. Here $\|\mathbf{A}\|_2$ denotes the spectral norm of $\mathbf{A}$ – i.e., its largest eigenvalue magnitude. They prove a similar guarantee for KPM itself, but with a worse dependence on ϵ . Chen et al. [CTU21, CTU22] prove that the Lanczos-based SLQ method gives essentially the same approximation bound: error $\epsilon \cdot \|\mathbf{A}\|_2$ using $O(1/\epsilon)$ matrix-vector products when $\epsilon = \tilde{\Omega}(1/\sqrt{n})$.¹

¹Work on eigenvalue estimation [AN13, BDD⁺24, SW23] can also give guarantees for SDE. However they are generally weaker than those discussed above. E.g., [SW23] shows how to approximate all eigenvalues of symmetric $\mathbf{A}$ to additive error $\epsilon \|\mathbf{A}\|_F$ using $O(1/\epsilon^2)$ matrix-vector products, which is optimal. Letting $\tilde{s}_{\mathbf{A}}$ be the uniform distribution over their approximate eigenvalues, we obtain the somewhat weak bound of $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \|\mathbf{A}\|_F$.

The above error bounds for KPM, Chebyshev Moment Matching, and SLQ help to justify the effectiveness of these methods in practice. However, in many cases, they can be loose. A bound of $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \cdot \|\mathbf{A}\|_2$ roughly corresponds to estimating each eigenvalue to average error $\epsilon \cdot \|\mathbf{A}\|_2$. Many matrices however exhibit spectral decay: most of their eigenvalues are much smaller in magnitude than their largest (i.e., than $\|\mathbf{A}\|_2$). Thus, this error bound does not guarantee that $\tilde{s}_{\mathbf{A}}$ effectively captures information about $\mathbf{A}$'s smaller magnitude eigenvalues.

1.3 Our Results

Our main contribution is to show that both moment matching and Lanczos based methods for SDE can achieve improved bounds on $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}})$ that depend on $\sigma_{l+1}(\mathbf{A})$, the $(l+1)^{st}$ largest singular value of $\mathbf{A}$ (i.e., the $(l+1)^{st}$ largest eigenvalue magnitude) for some parameter l , instead of $\|\mathbf{A}\|_2$. For matrices that exhibit spectral decay and thus have $\sigma_{l+1}(\mathbf{A}) \ll \sigma_1(\mathbf{A}) = \|\mathbf{A}\|_2$, our bounds can be much stronger than those given in prior work.

1.3.1 Improved SDE via Moment Matching with Explicit Deflation

Our first contribution is a modification of the moment matching method of [BKM22] that first ‘deflates’ off any eigenvalue of $\mathbf{A}$ with magnitude significantly larger than $\sigma_{l+1}(\mathbf{A})$, before estimating the spectral density. Specifically, the method uses a block Krylov subspace method [MM15, Tro18] to first compute highly accurate approximations to the p largest magnitude eigenvalues of $\mathbf{A}$, for some $p \leq l$, along with an orthonormal matrix $\mathbf{Z} \in \mathbb{R}^{n \times p}$ with columns approximating the corresponding eigenvectors. It uses moment matching to estimate the spectral density of $\mathbf{A}$ projected away from these approximate eigendirections $(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)$, achieving error $\epsilon\sigma_{l+1}(\mathbf{A})$ since this matrix has spectral norm bounded by $O(\sigma_{p+1}(\mathbf{A})) = O(\sigma_{l+1}(\mathbf{A}))$ if $\mathbf{Z}$ is sufficiently accurate. It then modifies this approximate density to account for the probability mass at the top p eigenvalues. While block Krylov methods are well understood for the closely related tasks of low-rank approximation and singular value approximation [MM15, MNS⁺18, BN23], our work requires a careful analysis of eigenvalue/eigenvector approximation with these methods that may be of independent interest – see Section 1.4 for details. Overall, the above approach gives the following result:

Theorem 1 (SDE with Explicit Deflation). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric. For any $\epsilon \in (0, 1)$, $l \in [n]$, and any constants $c_1, c_2 > 0$, Algorithm 1 performs $O(l \log n + \frac{b}{\epsilon})$ matrix-vector products with $\mathbf{A}$ where $b = \max(1, \frac{1}{n\epsilon^2} \log^2 \frac{n}{\epsilon} \log^2 \frac{1}{\epsilon})$ and computes a probability density function $\tilde{s}_{\mathbf{A}}$ such that, with probability at least $1 - \frac{1}{n^{c_1}}$,*

$$W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2}}.$$

As compared to the result of [BKM22], Theorem 1 uses $O(l \log n)$ additional matrix-vector products – these are used to compute the approximate top eigenvalues and eigenvectors for deflation. However, the method gives a significantly improved error bound of roughly $\epsilon\sigma_{l+1}(\mathbf{A})$. The additive error $\|\mathbf{A}\|_2/n^{c_2}$ can be thought of as negligible – comparable e.g., to round-off error when directly computing $s_{\mathbf{A}}$ using a full eigendecomposition in finite precision arithmetic [BGVKS22].

We further show that our algorithm is optimal amongst all matrix-vector query algorithms, up to logarithmic factors and the negligible additive error term. Formally:

Theorem 2 (SDE Lower Bound). *Any (possibly randomized) algorithm that given symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ outputs $\tilde{s}_{\mathbf{A}}$ such that, with probability at least $1/2$, $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon\sigma_{l+1}(\mathbf{A})$ for $\epsilon \in (0, 1)$ and $l \in [n]$ must make $\Omega(l + \frac{1}{\epsilon})$ (possibly adaptively chosen) matrix-vector queries to $\mathbf{A}$.*

Theorem 2 leverages an existing lower bound for distinguishing Wishart matrices of different ranks, previously used to give matrix-vector query lower bounds for the closely related problem of eigenvalue estimation [SW23].

Application to Schatten-1 (Nuclear) Norm Estimation. Even for matrices that don't exhibit spectral decay, by balancing the $O(l \log n)$ and $O(b/\epsilon)$ terms we can leverage Theorem 1 in applications that require understanding the small magnitude eigenvalues of $\mathbf{A}$, where previous SDE bounds gave weak results. For example, consider the estimating the Schatten-1 (nuclear) norm $\|\mathbf{A}\|_1 = \sum_{i=1}^n |\lambda_i(\mathbf{A})| = \sum_{i=1}^n \sigma_i(\mathbf{A})$. It is not hard to show that we can estimate $\|\mathbf{A}\|_1$ to relative error $\epsilon \|\mathbf{A}\|_1$ given an approximate spectral density $\tilde{s}_{\mathbf{A}}$ with $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \frac{\epsilon \|\mathbf{A}\|_1}{n}$ (see [CSKSV18], Theorem B.1 in [BKM22]). We can find such a density by applying Theorem 1 with $l = \frac{\sqrt{n}}{\epsilon}$ and $\epsilon' = \frac{1}{\sqrt{n}}$ so that $\epsilon' \cdot \sigma_{l+1}(\mathbf{A}) \leq \frac{\epsilon' \|\mathbf{A}\|_1}{l} \leq \frac{\epsilon \|\mathbf{A}\|_1}{n}$. Doing so yields the following corollary:

Corollary 1 (Schatten-1 Norm Estimation). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric. For any $\epsilon \in (0, 1)$ and any constant $c > 0$, there exists an algorithm that performs $O\left(\frac{\sqrt{n} \log n}{\epsilon} + \sqrt{n} \log^4 n\right)$ matrix vector products with $\mathbf{A}$ and computes M such that, with probability at least $1 - \frac{1}{n^c}$, $|M - \|\mathbf{A}\|_1| \leq \epsilon \|\mathbf{A}\|_1$.*

Prior work on SDE [CTU21, BKM22] could only give error $\epsilon \sqrt{n} \|\mathbf{A}\|_2$ using a comparable number of matrix-vector products, and thus was not able to achieve a relative error guarantee. We note that the $\sqrt{n}$ dependence of Corollary 1 matches the best known matrix-vector product algorithms for Schatten-1 norm estimation [MNS⁺18], while the ϵ dependence improves on prior work.

1.3.2 Implicit Deflation Bounds for Stochastic Lanczos Quadrature

Our second contribution is to show that the popular Stochastic Lanczos Quadrature (SLQ) method for SDE [LSY16, CTU21] nearly matches the improved error bound of Theorem 1 for any choice of l , even though SLQ is ‘parameter-free’ and performs no explicit deflation step. This result helps to justify the strong practical performance of SLQ and the observed ‘spectrum adaptive’ nature of this method as compared to standard moment matching-based methods like KPM [CTU21].

A key idea used to achieve this bound is to view SLQ as an implicit moment matching method as in [CTU21, CTU22], and to analyze it similarly to KPM and other explicit moment matching methods. We combine this analysis approach with recent work on low-rank approximation with single-vector (i.e., non-block) Krylov methods [MMM24] to show that SLQ can perform ‘implicit deflation’ as part of moment matching to achieve the improved error bound. See Section 1.4 for details. Formally, our error bound for SLQ is as follows:

Theorem 3 (SDE with SLQ). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and consider any $l \in [n]$, and $\epsilon, \delta \in (0, 1)$. Let $g_{\min} = \min_{i \in [l]} \frac{\sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A})}{\sigma_i(\mathbf{A})}$ and $\kappa = \frac{\|\mathbf{A}\|_2}{\sigma_{l+1}(\mathbf{A})}$. Algorithm 4 (SLQ) run for $m = O(l \log \frac{1}{g_{\min}} + \frac{1}{\epsilon} \log \frac{n \cdot \kappa}{\delta})$ iterations performs m matrix vector products with $\mathbf{A}$ and outputs a probability density function $\tilde{s}_{\mathbf{A}}$ such that, with probability at least $1 - \delta$, for a fixed constant C ,*

$$W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \frac{C \log(n/\epsilon) \log(1/\epsilon)}{\sqrt{n}} \cdot \sigma_{l+1}(\mathbf{A}) + \frac{Cl \log(1/\epsilon) \sqrt{\log(l/\delta)}}{n} \|\mathbf{A}\|_2.$$

Theorem 3 essentially matches our result for moment matching with explicit deflation (Theorem 1) up to some small caveats, discussed below. First, the number of matrix vector products has a logarithmic dependence on the minimum gap $g_{\min}$ amongst the top l singular values as well as

the condition number $\kappa = \frac{\|\mathbf{A}\|_2}{\sigma_{l+1}(\mathbf{A})}$. The dependence on the minimum gap is inherent, as non-block Krylov methods like SLQ require a dependence on $g_{\min}$ in order to perform deflation/low-rank approximation [MMM24]. We note that by adding a random perturbation to $\mathbf{A}$ with spectral norm bounded by $\frac{\|\mathbf{A}\|_2}{\text{poly}(n)}$, one can ensure that both $g_{\min} \geq \frac{1}{\text{poly}(n)}$ and $\kappa \leq \text{poly}(n)$ with high probability, and thus replace the $O(l \log \frac{1}{g_{\min}})$ term with an $O(l \log n)$ and the $O(\frac{\log(n\kappa)}{\epsilon})$ term with an $O(\frac{\log n}{\epsilon})$ term, matching Theorem 1. See e.g., [MMM24].

Second, Theorem 3 has an additional error term of size $\tilde{O}(\sigma_{l+1}(\mathbf{A})/\sqrt{n})$. This term is lower order whenever $\epsilon = \tilde{\Omega}(1/\sqrt{n})$. Further, we believe that it can be removed by using a variant on SLQ that is popular in practice, where the densities output by multiple independent runs of the method are averaged together to produce $\tilde{s}(\mathbf{A})$. See Section 4 for further discussion.

Finally, Theorem 3 has an additional error term of size $\tilde{O}(\|\mathbf{A}\|_2 \cdot l/n)$. In the natural case when we run for $m \ll n$ iterations and thus $l \ll n$, this term will be small. However, it cannot be avoided: even for a matrix with rank $\leq l$ with well-separated eigenvalues, while the Lanczos method will converge to near-exact approximations to these eigenvalues (with error bounded by $\frac{\|\mathbf{A}\|_2}{n^c}$), the probability distribution output by SLQ will not place mass exactly $1/n$ at these approximate eigenvalues and thus will not achieve SDE error $O(\frac{\|\mathbf{A}\|_2}{n^c})$ – see Section 1.4 for further details.

This limitation motivates us to introduce a simple variant of SLQ, which we call *variance reduced SLQ*, which places mass exactly $1/n$ at any eigenvalue computed by Lanczos that has converged to sufficiently small error. This variant gives the following stronger error bound:

Theorem 4 (SDE with Variance Reduced SLQ). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and consider any $l \in [n]$, and $\epsilon, \delta \in (0, 1)$. Let $g_{\min} = \min_{i \in [l]} \frac{\sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A})}{\sigma_i(\mathbf{A})}$ and $\kappa = \frac{\|\mathbf{A}\|_2}{\sigma_{l+1}(\mathbf{A})}$. Algorithm 5 run for $m = O(l \log \frac{1}{g_{\min}} + \frac{1}{\epsilon} \log \frac{n\kappa}{\delta})$ iterations performs m matrix vector products with $\mathbf{A}$ and outputs a probability density function $\tilde{s}_{\mathbf{A}}$ such that, with probability at least $1 - \delta$, for any constant c and a fixed constant C ,*

$$W_1(s_{\mathbf{A}}, \tilde{s}(\mathbf{A})) \leq \epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \left(\frac{C \log(n/\epsilon) \log(1/\epsilon)}{\sqrt{n}} + \frac{Cl \log(1/\epsilon) \sqrt{\log(l/\delta)}}{n} \right) \cdot \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^c}.$$

1.4 Technical Overview

We next overview the main techniques used to achieve our improved SDE bounds for moment matching with deflation (Theorem 1) and SLQ (Theorems 3 and 4).

1.4.1 Eigenvalue Deflation for SDE

As discussed in Section 1.3.1, the key idea behind Theorem 1 is to apply *eigenvalue deflation*. Assume that we are given $\mathbf{Z} \in \mathbb{R}^{n \times l}$ with columns equal to the eigenvectors of $\mathbf{A}$ corresponding to its l largest magnitude eigenvectors. Then we can write $\mathbf{A} = \mathbf{A}_l + \mathbf{A}_{l,\perp}$, where $\mathbf{A}_l = \mathbf{Z}\mathbf{Z}^T \mathbf{A} \mathbf{Z}\mathbf{Z}^T$ and $\mathbf{A}_{l,\perp} = (\mathbf{I} - \mathbf{Z}\mathbf{Z}^T) \mathbf{A} (\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)$. Referring to (1), since $\mathbf{A}_l$ and $\mathbf{A}_{l,\perp}$ have n eigenvalues in total that are set to 0 due to projecting off $\mathbf{Z}$, we have:

$$s_{\mathbf{A}} = s_{\mathbf{A}_l} + s_{\mathbf{A}_{l,\perp}} - \delta(0), \quad (2)$$

where $\delta(x)$ denotes the Dirac distribution that places probability mass one at x . If we have $\mathbf{Z}$ in hand, we can exactly compute $s_{\mathbf{A}_l}$, whose only non-zero eigenvalues are exactly the eigenvalues

corresponding to the eigenvectors in $\mathbf{Z}$. Further, we can approximate $s_{\mathbf{A}_{l,\perp}}$ using an existing SDE algorithm – e.g., the Chebyshev moment matching method of [BKM22], which will give error $\epsilon \cdot \|\mathbf{A}_{l,\perp}\|_2$ using $O(1/\epsilon)$ matrix vector products when $\epsilon = \tilde{\Omega}(1/\sqrt{n})$.² We have $\|\mathbf{A}_{l,\perp}\|_2 = |\lambda_{l+1}(\mathbf{A})| = \sigma_{l+1}(\mathbf{A})$, where $\lambda_{l+1}(\mathbf{A})$ is the $(l+1)^{st}$ largest magnitude eigenvalue of $\mathbf{A}$. Thus, combining this approximation with (2), we can approximate $s_{\mathbf{A}}$ to error $\epsilon \cdot \sigma_{l+1}(\mathbf{A})$ as desired.

We note that eigenvalue deflation is widely applied throughout numerical linear algebra to problems like linear system solving [BEPW98, FV01, GOSS16, FTU23], trace estimation [GSO17, Lin17, MMMW21], norm estimation [MNS⁺18], and beyond [CS97] – the approach is useful whenever the complexity or approximation error of solving a problem depends on some feature of its eigenspectrum (e.g. the spectral norm, Frobenius norm, or condition number) that can be improved by removing the largest magnitude eigenvalues from the matrix.

1.4.2 Error Analysis of Deflation

Of course, when implementing deflation for SDE, we cannot exactly compute $\mathbf{Z} \in \mathbb{R}^{n \times l}$ spanning the top l eigenvectors of $\mathbf{A}$. Instead we will approximate $\mathbf{Z}$, in our case, using a standard block Krylov subspace method (Algorithm 2). It is well known that when $\mathbf{Z}$ is approximated in this way, using $O(\log n)$ iterations of the block Krylov method, and thus $O(l \log n)$ matrix-vector products in total, we can still ensure that $\|(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\|_2 = O(\sigma_{l+1}(\mathbf{A}))$ [MM15]. This suffices to obtain the error bound of $\epsilon \cdot \sigma_{l+1}(\mathbf{A})$ in approximating the spectral density of $(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)$.

The key challenge however is that when $\mathbf{Z}$ does not exactly span a set of eigenvectors, (2) no longer holds. Further, the spectral density of $\mathbf{Z}\mathbf{Z}^T\mathbf{A}\mathbf{Z}\mathbf{Z}^T$ will no longer exactly match the spectral density of $\mathbf{A}_l$ – it will place mass at the eigenvalues of $\mathbf{Z}\mathbf{Z}^T\mathbf{A}\mathbf{Z}\mathbf{Z}^T$, equal to the eigenvalues of $\mathbf{Z}^T\mathbf{A}\mathbf{Z}$, which approximate, but don't exactly match, the top eigenvalues of $\mathbf{A}$.

Thus, our proof requires showing that $\mathbf{Z}$ is very close to spanning a subspace of eigenvectors, up to the small $\frac{\|\mathbf{A}\|_2}{n^\epsilon}$ additive error of Theorem 1. To do so, in our block Krylov algorithm (Algorithm 2 line 7) we let $\mathbf{Z}$ contain only the approximate eigenvectors that have converged to very small *residual error* – i.e., for which $\mathbf{A}\tilde{\mathbf{z}}_j \approx \lambda_j\tilde{\mathbf{z}}_j$. Via standard backward error analysis bounds for eigenvector approximation (see Theorem 9), this is enough to handle the above two issues. However, we can now no longer be sure that $\|(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\|_2$ is small – what if, e.g., no eigenvectors have converged and thus $\mathbf{Z}$ is empty? Or e.g., if just the top eigenvector has failed to converge.

Our main technical contribution is to argue that that for some $p \leq l$ with $\sigma_{p+1}(\mathbf{A}) = O(\sigma_{l+1}(\mathbf{A}))$, the set of converged eigenvectors (i.e., the columns of $\mathbf{Z}$) will contain approximations to at least the top p magnitude eigenvectors of $\mathbf{A}$, ensuring that $\|(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\|_2 = O(\sigma_{p+1}(\mathbf{A})) = O(\sigma_{l+1}(\mathbf{A}))$ as needed. Theorems 5 and 7 give our convergence bounds for the top p eigenvectors, and Theorem 8 states the ultimate resulting bound on $\|(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\|_2$.

Our work fits into a line of work that focuses on more refined eigenvalue/singular value approximation bounds for block Krylov subspace methods – see e.g. [MM15, MNS⁺18, DIKMI18, Tro22]. We believe that Theorems 5 and 7 may be of independent interest, outside our application to SDE.

1.4.3 SLQ and its Existing Analysis

We next turn our attention to the Stochastic Lanczos Quadrature (SLQ) method, which is detailed in Algorithm 4. This is a popular and extremely simple, parameter-free algorithm for SDE based on the classic Lanczos method (Algorithm 3). SLQ uses Lanczos to compute an orthonormal

²We will focus on the case $\epsilon = \tilde{\Omega}(1/\sqrt{n})$ and so $b = 1$ in Thm. 1 throughout the technical overview for simplicity.

basis $\mathbf{Q}$ for the Krylov subspace $\{\mathbf{g}, \mathbf{A}\mathbf{g}, \mathbf{A}^2\mathbf{g}, \dots, \mathbf{A}^m\mathbf{g}\}$ for random starting vector $\mathbf{g} \in \mathbb{R}^n$ – typically, $\mathbf{g}$ is Gaussian. The algorithm then computes the eigenvalues of $\mathbf{T} = \mathbf{Q}^T \mathbf{A} \mathbf{Q}$ and lets $s_{\tilde{\mathbf{A}}} = \sum_{i=1}^m w_j \cdot \delta(x - \lambda_j(\mathbf{T}))$ where w_j are appropriately chosen weights: $w_j = (\mathbf{g}^T \mathbf{Q} \mathbf{v}_j)^2$ where $\mathbf{v}_j$ is the j^{th} eigenvector of $\mathbf{T}$. I.e., the spectral density of $\mathbf{A}$ is approximated with a reweighted spectral density of $\mathbf{T}$. Lanczos constructs $\mathbf{Q}$ such that $\mathbf{T}$ is tridiagonal – this makes computing its eigenvalues very efficient, but is otherwise not critical in the analysis.

The key idea behind SLQ (and Lanczos-based algorithms in general) is that, for any polynomial with degree $< m$, one can show that $p(\mathbf{A})\mathbf{g} = \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{g}$. Using this fact, if we consider any low-degree polynomial moment for the SLQ approximate spectral density, we have:

$$\begin{aligned} \mathbb{E}_{\tilde{s}_{\mathbf{A}}}[p(x)] &= \sum_{j=1}^m w_j \cdot p(\lambda_j(\mathbf{T})) = \sum_{j=1}^m \mathbf{g}^T \mathbf{Q} \mathbf{v}_j \mathbf{v}_j^T \mathbf{Q}^T \mathbf{g} \cdot p(\lambda_j(\mathbf{T})) \\ &= \mathbf{g}^T \mathbf{Q} \left(\sum_{j=1}^m \mathbf{v}_j \mathbf{v}_j^T p(\lambda_j(\mathbf{T})) \right) \mathbf{Q}^T \mathbf{g} \\ &= \mathbf{g}^T \mathbf{Q} p(\mathbf{T}) \mathbf{Q}^T \mathbf{g} = \mathbf{g}^T p(\mathbf{A}) \mathbf{g}. \end{aligned} \quad (3)$$

We can observe that when $\mathbf{g}$ has i.i.d. Gaussian entries with variance $\frac{1}{n}$, $\mathbb{E}_{\mathbf{g}}[\mathbf{g}^T p(\mathbf{A}) \mathbf{g}] = \frac{1}{n} \text{tr}(p(\mathbf{A})) = \frac{1}{n} \sum_{i=1}^n p(\lambda_i) = \mathbb{E}_{s_{\mathbf{A}}}[p(x)]$. That is, the approximate spectral density output by SLQ matches all low-degree polynomial moments of the true spectral density in expectation.

To formally argue that $\tilde{s}_{\mathbf{A}}$ is close to $s_{\mathbf{A}}$ in Wasserstein distance, Chen, Trogden, and Ubaru [CTU21] argue that any two distributions supported on $[-\|\mathbf{A}\|_2, \|\mathbf{A}\|_2]$ that exactly match on any polynomial moment of degree $O(1/\epsilon)$ have Wasserstein distance at most $\epsilon \cdot \|\mathbf{A}\|_2$. This allows them to argue that $W_1(\bar{s}_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \cdot \|\mathbf{A}\|_2$ where $\bar{s}_{\mathbf{A}} = \sum_{i=1}^n \eta_i \cdot \lambda_i(\mathbf{A})$ with $\eta_i = (\mathbf{g}^T \mathbf{u}_i)^2$, for $\mathbf{u}_i$ being the i^{th} eigenvector of $\mathbf{A}$. Note that $\bar{s}_{\mathbf{A}}$ is constructed so that for any p , $\mathbb{E}_{\bar{s}_{\mathbf{A}}}[p(x)] = \mathbf{g}^T p(\mathbf{A}) \mathbf{g}$ and thus the low-degree moments of $\tilde{s}_{\mathbf{A}}$ and $\bar{s}_{\mathbf{A}}$ match by (3). [CTU21] next uses a concentration bound to argue that the CDFs of $\bar{s}_{\mathbf{A}}$ and $s_{\mathbf{A}}$ are close, and therefore that these two distributions are close in Wasserstein distance. They conclude via triangle inequality that $\tilde{s}_{\mathbf{A}}$ is close to $s_{\mathbf{A}}$.

1.4.4 Moment Matching-Based Analysis of SLQ

Our analysis follows a similar approach to [CTU21], also viewing SLQ as a moment matching method, but analyzing it in essentially an identical manner to the Chebyshev moment matching method of [BKM22]. We consider the first $m = O(1/\epsilon)$ Chebyshev polynomials of the first kind: $T_0, T_1, \dots, T_m$. These polynomials are sine-like functions that are bounded in magnitude by 1 on $[-1, 1]$ (in the SDE setting, we rescale them to be bounded on $[-\|\mathbf{A}\|_2, \|\mathbf{A}\|_2]$). [BKM22] show that any two distributions that are close in their Chebyshev moments are also close in Wasserstein distance. They then approximate the moments of $s_{\mathbf{A}}$ and find a distribution $\tilde{s}_{\mathbf{A}}$ matching these moments (and thus satisfying $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \|\mathbf{A}\|_2$) by solving a constrained optimization problem.

Like many moment matching methods, [BKM22] use Hutchinson’s trace estimator [Gir87, Hut90] to approximate the Chebyshev moments – i.e., they approximate $\frac{1}{n} \text{tr}(T_i(\mathbf{A})) \approx \frac{1}{b} \sum_{j=1}^b \mathbf{g}_j^T T_i(\mathbf{A}) \mathbf{g}_j$ where each $\mathbf{g}_j$ has i.i.d. Gaussian entries with variance $\frac{1}{n}$. They argue that since T_i is bounded, Hutchinson’s method is highly accurate, and thus for $\epsilon = \tilde{\Omega}(1/\sqrt{n})$, they can in fact just set $b = 1$. So their approximate moments are just of the form $\mathbf{g}^T T_i(\mathbf{A}) \mathbf{g}$ for a single random $\mathbf{g}$.

Recall that by (3) these are exactly the moments of $\tilde{s}_{\mathbf{A}}$ output by SLQ! So, SLQ ‘automatically’ finds a distribution matching the approximate Chebyshev moments of [BKM22], and thus satisfies $W_1(\tilde{s}_{\mathbf{A}}, s_{\mathbf{A}}) \leq \epsilon \cdot \|\mathbf{A}\|_2$.³ We prove this bound in Section 4.1, Theorem 10.

1.4.5 Implicit Deflation with SLQ

Armed with the above moment matching-based analysis of SLQ, we next show that the method performs implicit deflation and thus nearly matches the bound of Theorem 1. The key idea is as follows: the Chebyshev moments used in the analysis above are the polynomials $T_1(x/\|\mathbf{A}\|_2), T_2(x/\|\mathbf{A}\|_2), \dots$ – i.e., the Chebyshev polynomials scaled to be bounded on the range $[-\|\mathbf{A}\|_2, \|\mathbf{A}\|_2]$, which contains the eigenvalues of both $\mathbf{A}$ and $\mathbf{T}$. This scaling is what leads to the error term scaling with $\|\mathbf{A}\|_2$.

In our explicit deflation approach, after deflating off the largest magnitude eigenvalues we apply moment matching to $(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)$. We thus approximate the Chebyshev moments of this matrix where the polynomials are scaled to the (possibly much smaller range) $[-\|(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\|_2, \|(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\|_2]$ of width $O(\sigma_{l+1}(\mathbf{A}))$. This leads to error scaling with $\sigma_{l+1}(\mathbf{A})$. We cannot use these moments to approximate $s_{\mathbf{A}}$ directly: these scaled Chebyshev polynomials blow up outside the range on which they are bounded and the moments would be dominated by eigenvalues with magnitudes $\geq \sigma_{l+1}(\mathbf{A})$ and so not informative for approximating $s_{\mathbf{A}}$.

Deflated Polynomial Moments. To handle this issue, we use that the density output by SLQ approximates $s_{\mathbf{A}}$ on *any low-degree polynomial moment*. Instead of the scaled Chebyshev polynomials, we use a set of “deflated polynomials”, denoted $t_1, \dots, t_m$ which are approximately equal to the scaled Chebyshev polynomials on the range $[-O(\sigma_{l+1}(\mathbf{A})), O(\sigma_{l+1}(\mathbf{A}))]$ and have roots placed outside this range so they are equal to zero for any eigenvalue with magnitude significantly larger than $\sigma_{l+1}(\mathbf{A})$. Constructing such polynomials is somewhat delicate – naively placing the roots at large eigenvalues would distort the Chebyshev polynomials on the range of interest. This distortion needs to be canceled using a separate Chebyshev damping polynomial. Our analysis follows ideas from work studying the convergence of single-vector Krylov subspace methods like Lanczos for eigenvector/singular vector approximation [Saa80], and in particular, recent work proving low-rank approximation guarantees for these methods [MMM24].

Once we show that $t_1, \dots, t_m$ can be constructed as described, we know that SLQ matches the moments of $s_{\mathbf{A}}$ with respect to these polynomials. This is enough to argue that the output $\tilde{s}_{\mathbf{A}}$ must closely match the mass of $s_{\mathbf{A}}$ on the eigenvalues with magnitude $\leq \sigma_{l+1}(\mathbf{A})$. But it says nothing about the large eigenvalues, which contribute nothing to these moments. We need to separately argue about this part of the density.

Top Eigenvalue Approximation. To do so, we directly look at the form of $\tilde{s}_{\mathbf{A}}$. Following a similar strategy to our proof for explicit deflation and again building on recent work studying the convergence of single-vector Krylov methods like Lanczos [MMM24], we argue that any eigenvalue $\lambda_j(\mathbf{T})$ with magnitude significantly larger than $\sigma_{l+1}(\mathbf{A})$ is extremely close to $\lambda_j(\mathbf{A})$. This finding is not surprising: the Lanczos method is a popular choice for approximating outlying eigenvalues [Par98]. Thus, on the large magnitude eigenvalues, the error that $\tilde{s}_{\mathbf{A}}$ incurs vs. $s_{\mathbf{A}}$ is roughly:

$$w_j \lambda_j(\mathbf{T}) - \frac{1}{n} \lambda_j(\mathbf{A}) \approx \left(w_j - \frac{1}{n} \right) \lambda_j(\mathbf{A}) \leq \left(w_j - \frac{1}{n} \right) \|\mathbf{A}\|_2.$$

³Even when $\epsilon = \tilde{o}(1/\sqrt{n})$ and so we need $b > 1$ repetitions of Hutchinson’s to estimate the Chebyshev moments in [BKM22], a similar analysis follows from averaging together the densities output by b independent runs of SLQ.

Recall that $w_j = (\mathbf{g}^T \mathbf{Q} \mathbf{v}_j)^2$. Further, since our approximations to the top eigenvectors have converged, $\mathbf{Q} \mathbf{v}_j \approx \mathbf{u}_j$, where $\mathbf{u}_j$ is the eigenvector of $\mathbf{A}$ corresponding to $\lambda_j(\mathbf{A})$. Since $\mathbf{u}_j$ is a fixed unit vector and $\mathbf{g}$ has random Gaussian entries with variance $\frac{1}{n}$, w_j is simply the square of a Gaussian with variance $\frac{1}{n}$. I.e., $\mathbb{E}[w_j] = 1/n$ and we can use a Chi-Squared concentration bound to argue that $|w_j - \frac{1}{n}| = \tilde{O}(\frac{1}{n})$ with high probability. Accounting for the l top eigenvalues, this leads to the $\tilde{O}(l/n \cdot \|\mathbf{A}\|_2)$ error term in Theorem 3 and completes our analysis.

1.4.6 Variance Reduced SLQ

The above argument immediately suggests a simple improvement to SLQ that can avoid the $\tilde{O}(l/n \cdot \|\mathbf{A}\|_2)$ error term due to the large magnitude eigenvalues. If an eigenvector has converged and thus we know that $|\lambda_j(\mathbf{T}) - \lambda_j(\mathbf{A})|$ is small, then in the weighted spectral density $\tilde{s}_{\mathbf{A}}$ output by Lanczos, we know that we should set $w_j = \frac{1}{n}$ rather than setting w_j to be the square of a Gaussian with variance $\frac{1}{n}$. We formalize this approach in Algorithm 5 and show that it obtains the stronger error bound of Theorem 4, with only a negligible additive term depending on $\|\mathbf{A}\|_2$.

A Remark on Numerical Stability. We note that there exist results proving that the Conjugate Gradient (CG) algorithm for solving linear systems, which is closely related to Lanczos, can perform implicit eigenvalue deflation to achieve faster convergence rates, analogous to our analysis for SDE [SW09, AL86]. These bounds are known to suffer from numerical stability issues – when CG is implemented in finite precision arithmetic, it may not be able to apply the necessary polynomials to achieve these convergence rates. It is likely that our bounds also suffer from stability issues. However 1) in finite precision, the method likely works well as long as the number of deflated eigenvalues l is relatively small and 2) our analysis directly applies to SLQ implemented using a stable algorithm to compute a basis for the Krylov subspace – e.g., one that performs full reorthogonalization at each step. Using such an algorithm will not affect the matrix-vector query complexity of the algorithm, and only gives a small runtime overhead when the number of iterations is relatively small.

1.5 Roadmap

The remainder of the paper is organized as follows: In Section 2 we define basic notation and preliminary results used throughout our proofs. In Section 3 we analyze our explicit moment matching method with deflation, culminating in the proof of Theorem 1. In Section 4 we give our analysis of SLQ and prove Theorems 3 and 4. In Section 5 we prove our matching lower bound, Theorem 2, which shows that our SDE bounds are near-optimal. Finally, in Section 6, we report the results of some numerical experiments comparing various moment matching and Lanczos-based algorithms for SDE.

2 Notation and Preliminaries

We first introduce notation and foundational results that we will use throughout.

2.1 Basic Notation

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 $\mathbf{x}$. For a vector $\mathbf{x}$, we let $\|\mathbf{x}\|_2$ denote its Euclidean norm. The Frobenius norm and spectral (i.e., operator) norm of a matrix $\mathbf{A}$ are denoted by $\|\mathbf{A}\|_F$ and $\|\mathbf{A}\|_2$ respectively. We often use that

for any two matrices $\mathbf{A}, \mathbf{B}$ of appropriate dimensions, the spectral norm is sub-multiplicative, i.e., $\|\mathbf{AB}\|_2 \leq \|\mathbf{A}\|_2 \|\mathbf{B}\|_2$. The column span of a matrix $\mathbf{A}$ is denoted by $\text{range}(\mathbf{A})$.

2.2 Linear Algebra Preliminaries

SVD and Eigendecomposition. The singular values of $\mathbf{A} \in \mathbb{R}^{n \times n}$ are denoted by $\sigma_i(\mathbf{A})$ for $i \in [n]$ with $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_n(\mathbf{A})$. We denote the singular value decomposition (SVD) of $\mathbf{A} \in \mathbb{R}^{n \times n}$ by $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ where $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with $\Sigma_{ii} = \sigma_i(\mathbf{A})$ and $\mathbf{U} \in \mathbb{R}^{n \times n}$ and $\mathbf{V} \in \mathbb{R}^{n \times n}$ are orthonormal matrices with columns containing $\mathbf{A}$'s left and right singular vectors respectively.

Our main theorems are concerned with the eigenspectrum of symmetric matrices. We denote the eigenvalues of a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ by $\lambda_i(\mathbf{A})$ for $i \in [n]$ such that $|\lambda_1(\mathbf{A})| \geq \dots \geq |\lambda_n(\mathbf{A})|$. Observe that we have $|\lambda_i(\mathbf{A})| = \sigma_i(\mathbf{A})$ for $i \in [n]$. We denote the eigendecomposition of a symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ by $\mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$ where $\mathbf{\Lambda} \in \mathbb{R}^{n \times n}$ is a diagonal matrix such that $\Lambda_{ii} = \lambda_i(\mathbf{A})$ and $\mathbf{U} \in \mathbb{R}^{n \times n}$ has orthonormal columns equal to the corresponding eigenvectors of $\mathbf{A}$.

For $p \in [n]$, let $\mathbf{U}_p \in \mathbb{R}^{n \times p}$ (or $\mathbf{V}_p \in \mathbb{R}^{n \times p}$) denote the matrix containing the first p columns of $\mathbf{U}$ (or $\mathbf{V}$) – i.e., the singular vectors or the eigenvectors corresponding to the largest p singular values or eigenvalues of $\mathbf{A}$ by magnitude. Similarly, let $\mathbf{U}_{p,\perp} \in \mathbb{R}^{n \times (n-p)}$ (or $\mathbf{V}_{p,\perp} \in \mathbb{R}^{n \times (n-p)}$) denote the matrix which containing the last $n-p$ columns of $\mathbf{U}$ (or $\mathbf{V}$) – i.e., the singular vectors or the eigenvectors corresponding to the smallest $n-p$ singular values or eigenvalues of $\mathbf{A}$ by magnitude. Let $\mathbf{\Sigma}_p \in \mathbb{R}^{p \times p}$ (and $\mathbf{\Sigma}_{p,\perp} \in \mathbb{R}^{(n-p) \times (n-p)}$) denote the matrix containing the top p (or the bottom $n-p$) singular values of $\mathbf{A}$ along its diagonal. For symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ with eigendecomposition $\mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, define $\mathbf{\Lambda}_p \in \mathbb{R}^{p \times p}$ and $\mathbf{\Lambda}_{p,\perp} \in \mathbb{R}^{(n-p) \times (n-p)}$ analogously.

The best p -rank approximation to $\mathbf{A}$ in the spectral or Frobenius norms is denoted by $\mathbf{A}_p$ and is given by $\mathbf{A}_p = \mathbf{U}_p \mathbf{\Sigma}_p \mathbf{V}_p^T$. The pseudoinverse of $\mathbf{A}$ is denoted by $\mathbf{A}^\dagger$ and given by $\mathbf{A}^\dagger = \mathbf{V} \mathbf{\Sigma}^\dagger \mathbf{U}$, where $\Sigma_{ii}^\dagger = 1/\sigma_i(\mathbf{A})$ if $\sigma_i(\mathbf{A}) > 0$ and $\Sigma_{ii}^\dagger = 0$ otherwise.

We will frequently use Weyl's inequality which states that a small perturbation of a symmetric matrix will not significantly change its eigenvalues.

Fact 1 (Weyl's inequality [Wey12]). *For any two symmetric matrices $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{B}^{n \times n}$,*

$$\max_{i \in [n]} |\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})| \leq \|\mathbf{A} - \mathbf{B}\|_2.$$

Matrix functions. Given a function $\phi : \mathbb{R} \rightarrow \mathbb{R}$, we define its application to a matrix $\mathbf{A}$ with SVD $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ as $\phi(\mathbf{A}) = \mathbf{U}\phi(\mathbf{\Sigma})\mathbf{V}^T$, where $\phi(\mathbf{\Sigma})$ is formed by applying ϕ element-wise to the diagonal entries of $\mathbf{\Sigma}$ (i.e., to the singular values of $\mathbf{A}$). We overload notation and, for symmetric $\mathbf{A}$ with eigendecomposition $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, also let $\phi(\mathbf{A}) = \mathbf{U}\phi(\mathbf{\Lambda})\mathbf{U}^T$, where $\phi(\mathbf{\Lambda})$ is formed by applying ϕ element-wise to the diagonal entries of $\mathbf{\Lambda}$ (i.e., to the eigenvalues of $\mathbf{A}$). We will specify which specific form we are using when needed.

Krylov Subspaces. All algorithms analyzed in this work are Krylov subspace methods – we introduce basic notation for Krylov subspaces below.

Definition 2.1 (Block Krylov Subspace). *The Krylov subspace of a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with respect to a starting block $\mathbf{X} \in \mathbb{R}^{n \times l}$ is given by:*

$$\mathcal{K}_q(\mathbf{A}, \mathbf{X}) = [\mathbf{AX}, (\mathbf{AA}^T)\mathbf{AX}, \dots, (\mathbf{AA}^T)^q \mathbf{AX}].$$

Here, l is called the block size and q is the depth or the number of iterations.

Notice that we require $O(ql)$ matrix-vector products with $\mathbf{A}$ to generate a Krylov subspace $\mathcal{K}_q(\mathbf{A}, \mathbf{X})$ with depth q and block size l . We will typically denote an orthonormal basis of the Krylov subspace by $\mathbf{Q} \in \mathbb{R}^{n \times r}$ where r is the dimension to the column span of $\mathcal{K}_q(\mathbf{A}, \mathbf{X})$.

2.3 Moment Matching and Wasserstein Distance Preliminaries

Inner Product between Functions. For any two functions $f : [-L, L] \rightarrow \mathbb{R}$ and $g : [-L, L] \rightarrow \mathbb{R}$, the inner product between f and g is defined as $\langle f, g \rangle = \int_{-L}^L f(x)g(x)dx$.

Chebyshev polynomials. Both our explicit moment matching method and our error analysis for SLQ are based on analyzing approximations to the Chebyshev polynomial moments of the spectral density $s_{\mathbf{A}}$. We now describe some basic properties of Chebyshev polynomials. We will denote the k th Chebyshev polynomial of the first kind by T_k . These polynomials are defined by the recurrence:

$$T_0(x) = 1, T_1(x) = x, T_k(x) = 2x \cdot T_{k-1}(x) - T_{k-2}(x) \text{ for } k \geq 2.$$

We use the well known fact that the Chebyshev polynomials are bounded between $[-1, 1]$ i.e., $\max_{x \in [-1, 1]} |T_k(x)| \leq 1$. These polynomials also have an explicit expression of the form:

$$T_q(x) = \frac{1}{2}[(x + \sqrt{x^2 - 1})^q + (x - \sqrt{x^2 - 1})^q]. \quad (4)$$

Let $w(x) := \frac{1}{\sqrt{1-x^2}}$. It is well known that $\langle T_k, w \cdot T_k \rangle = \int_{-1}^1 \frac{1}{\sqrt{1-x^2}} T_k^2(x) dx = \frac{\pi}{2}$ for $k > 0$, $\langle T_0, w \cdot T_0 \rangle = \pi$ and $\langle T_i, w \cdot T_j \rangle = 0$ for $i \neq j$. We define the k^{th} normalized Chebyshev polynomial to be $\bar{T}_k := T_k / \sqrt{\langle T_k, w \cdot T_k \rangle}$. Note that we have $|\bar{T}_k(x)| \leq \sqrt{\frac{2}{\pi}}$ for $x \in [-1, 1]$. Also note that $T_k(-x) = T_k(x)$ when k is even and $T_k(-x) = -T_k(x)$ when k is odd.

We now describe the Chebyshev Series expansion of a function:

Definition 2.2 (Chebyshev Series Expansion). *The Chebyshev series expansion for a function $f : [-1, 1] \rightarrow \mathbb{R}$ is given by:*

$$\sum_{k=0}^{\infty} \langle f, w \cdot \bar{T}_k \rangle \cdot \bar{T}_k.$$

If f is Lipschitz continuous then this expansion converges absolutely and uniformly to f [Tre19].

Wasserstein Distance. The Wasserstein-1 distance, also known as the earth mover's distance, is a way to measure distance between two distributions. We will use the dual formulation of this distance given by Kantorovich-Rubinstein theorem [KR57]. In particular, for any two probability densities s and q supported on $[-L, L]$ for some $L \in \mathbb{R}^+$ the Wasserstein-1 distance is given by:

$$W_1(s, q) = \sup_{h \in 1\text{-Lip}} \int_{-L}^L h(x)(s(x) - q(x))dx,$$

where $h \in 1\text{-Lip}$ means we are optimizing the integral over all 1-Lipschitz functions.

3 SDE via Moment Matching with Explicit Deflation

In this section, we introduce our deflation-based approach to SDE (Algorithm 1) which combines a block Krylov method for deflation (Algorithm 2) with Chebyshev moment matching (Algorithm 2 of [BKM22]) for SDE. Pseudocode for this algorithm is given below.

Algorithm 1 Spectral Density Estimation with Deflation

Input: Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, error $\epsilon \in (0, 1)$, confidence $\delta \in (0, 1)$, block size l .

- 1: Let $\mathbf{Z} \in \mathbb{R}^{n \times s}$, $\tilde{\mathbf{\Lambda}} \in \mathbb{R}^{s \times s}$ be the outputs of Algorithm 2 (Block Krylov iteration) with inputs $\mathbf{A}$, block size l , and constant β .
 - 2: Let q_1 be the spectral density corresponding to $\tilde{\mathbf{\Lambda}}$: i.e., $q_1(x) = \frac{1}{s} \sum_{i=1}^s \delta(x - (\tilde{\mathbf{\Lambda}})_{ii})$.
 - 3: Let $\mathbf{P} = \mathbf{I} - \mathbf{Z}\mathbf{Z}^T$ and let L be an upper bound on $\|\mathbf{PAP}\|_2$ such that $\|\mathbf{PAP}\|_2 \leq L \leq 2\|\mathbf{PAP}\|_2$. Run Algorithm 2 of [BKM22] (Hutchinson-based Chebyshev moment estimation) with input matrix $\frac{1}{L}\mathbf{PAP}$, number of moments $N = \frac{c_1}{\epsilon}$, and number of repetitions of Hutchinson's method $b = \max(1, \frac{c_2}{n\epsilon^2} \log^2 \frac{1}{\epsilon\delta} \log^2 \frac{1}{\epsilon})$, where $c_1, c_2 > 0$ are sufficiently large constants. Let the approximate moments $\tilde{\tau}_1, \dots, \tilde{\tau}_N$ denote the output of this algorithm.
 - 4: Set $\hat{\tau}_i \rightarrow \frac{1}{n-s}(n \cdot \tilde{\tau}_i - s \cdot \bar{T}_i(0))$ for $i \in [N]$ where $\bar{T}_i(x)$ is the i 'th *normalized* Chebyshev polynomial.
 - 5: Run Algorithm 1 of [BKM22] (Chebyshev moment matching) with the modified approximate moments $\hat{\tau}_1, \dots, \hat{\tau}_N$ as inputs. Let q'_2 be the output of this algorithm.
 - 6: Define q_2 as $q_2(x) = q'_2(x/L)$ if $x \in [-L, L]$.
 - 7: **return** the spectral density $q = \frac{s \cdot q_1 + (n-s) \cdot q_2}{n}$.
-

Description of Algorithm 1. As discussed in Section 1.4.2, Algorithm 1 first uses a randomized block Krylov method (Algorithm 2) to compute a set of approximate eigenvectors and eigenvalues for $\mathbf{A}$, denoted by $\mathbf{Z}$ and $\tilde{\mathbf{\Lambda}}$ respectively. Importantly for our error analysis, Algorithm 2 only returns approximate eigenvectors that satisfy a convergence condition (line 7 of Algorithm 2), which ensures that $\mathbf{A}\mathbf{z}_i \approx \tilde{\lambda}_i \mathbf{z}_i$.

Algorithm 1 computes the spectral density q_1 of these converged approximate eigenvalues. It then *deflates* $\mathbf{A}$ using the corresponding converged eigenvectors in line 3 by computing $\mathbf{PAP}$ where $\mathbf{P} = \mathbf{I} - \mathbf{Z}\mathbf{Z}^T$. The algorithm then approximates the spectral density of $\mathbf{PAP}$ corresponding to the *non-deflated* eigenvalues of $\mathbf{A}$ using a moment matching method. To do so, Algorithm 1 first computes estimates $\tilde{\tau}_1, \dots, \tilde{\tau}_N$ of the normalized Chebyshev moments of $\mathbf{PAP}$ (after normalizing $\mathbf{PAP}$ so that its spectral norm is bounded by 1), using Hutchinson's method (Algorithm 2 of [BKM22]). Since $\mathbf{PAP}$ contains s zero eigenvalues corresponding to the deflated eigenvectors $\mathbf{Z} \in \mathbb{R}^{n \times s}$, the moments need to be adjusted in line 4 so that they do not take into account this mass of zero eigenvalues. These adjusted moments, $\hat{\tau}_1, \dots, \hat{\tau}_N$, are then passed as an input to Algorithm 2 of [BKM22] in line 5. This algorithm computes a spectral density q_2 whose moments are equal to the approximate moments $\hat{\tau}_i$ in line 4. The final output of Algorithm 1 is obtained by combining q_1 and q_2 after appropriately reweighting them.

Outline of Error Analysis. The remainder of this section is dedicating to analyzing Algorithm 1, ultimately culminating in the proof of Theorem 1, which bounds the Wasserstein error of the output spectral density estimate by $\epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^c}$ for any constant c . Our analysis breaks down into the following steps:

- In Section 3.1 we analyze the error of the block Krylov based deflation step (Algorithm 2). The main technical results of this section are Theorems 5 and 7, which argue that, for some $k \leq l$ with $\sigma_{k+1}(\mathbf{A}) = O(\sigma_{l+1})$, the method finds highly accurate approximations to at least the top k eigenvectors and eigenvalues of $\mathbf{A}$. In particular, these approximate eigenpairs meet the convergence condition of the algorithm and thus are returned as columns of $\mathbf{Z}$. With this fact established, we can prove the main export of the section, Theorem 8, which shows that $\|\mathbf{PAP}\|_2 = O(\sigma_{l+1}(\mathbf{A}))$. That is, when we deflate off the converged approximate eigenvectors, we reduce the spectral norm of the matrix to $O(\sigma_{l+1}(\mathbf{A}))$, allowing us to obtain $\epsilon \cdot \sigma_{l+1}(\mathbf{A})$ error when approximating the spectral density of $\mathbf{PAP}$ with Chebyshev moment matching.
- In Section 3.2 we give the error analysis of Algorithm 1 itself. This analysis is fairly straightforward – using existing backward stability analysis for eigenpair approximation, we can argue that, since $\mathbf{P}$ only deflates off highly accurate approximations to $\mathbf{A}$'s eigenvectors, our combined spectral density q has Wasserstein error roughly equal to (up to an additive $\frac{\|\mathbf{A}\|_2}{n^c}$ term) the Wasserstein error of approximating the spectral density of $\mathbf{PAP}$ with moment matching. As discussed above, this error is bounded by $O(\epsilon \cdot \|\mathbf{PAP}\|_2) = \epsilon \cdot \sigma_{l+1}(\mathbf{A})$, allowing us to achieve our final error bound of $\epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^c}$.

Algorithm 2 Block Krylov Iteration for Deflation

Input: Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, block size $l \in [n]$, iterations $q = O(\log n)$, constant $\beta > 0$

- 1: Let $\mathbf{X} \in \mathbb{R}^{n \times l}$ be a starting block with independent $\mathcal{N}(0, 1)$ Gaussian entries.
- 2: Compute $\mathcal{K}_q = [\mathbf{A}\mathbf{X}, (\mathbf{A}\mathbf{A}^T)\mathbf{A}\mathbf{X}, \dots, (\mathbf{A}\mathbf{A}^T)^q\mathbf{A}\mathbf{X}]$.
- 3: Orthonormalize the columns of $\mathcal{K}_q$ to get $\mathbf{Q} \in \mathbb{R}^{n \times r}$ where r is the rank of $\mathcal{K}_q$.
- 4: Compute $\mathbf{T} = \mathbf{Q}^T \mathbf{A} \mathbf{Q}$ and let the eigenvectors of $\mathbf{T}$ be $\mathbf{v}_1, \dots, \mathbf{v}_r$ corresponding to eigenvalues $|\lambda_1(\mathbf{T})| \geq \dots \geq |\lambda_r(\mathbf{T})|$.
- 5: Set $S = \{\}$.
- 6: **for** $j = 1, \dots, r$ **do**
- 7: **if** $\|\mathbf{A}\mathbf{Q}\mathbf{v}_j - \lambda_j(\mathbf{T})\mathbf{Q}\mathbf{v}_j\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^\beta}$ **then**
- 8: $S = S \cup \{j\}$
- 9: **end if**
- 10: **end for**
- 11: **return** $\mathbf{Z}_S = \mathbf{Q}\mathbf{V}_S$ and $\tilde{\mathbf{\Lambda}}_S$ where $\mathbf{V}_S \in \mathbb{R}^{r \times |S|}$ contains all eigenvectors with indices in S and $\tilde{\mathbf{\Lambda}}_S \in \mathbb{R}^{|S| \times |S|}$ is a diagonal matrix containing the corresponding eigenvalues of $\mathbf{T}$.

3.1 Error bounds for Deflation via Block Krylov

In this section, we analyze Algorithm 2 and prove that it outputs highly accurate approximations to the top eigenvalues and eigenvectors of $\mathbf{A}$.

Recall that we let $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ denote the SVD of $\mathbf{A}$ – see Section 2.2. We start by proving Lemma 1, a generalization of Theorem 2.1 of [DIKMI18], which bounds the error $\|(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p\|_F^2$ of projecting the top p left singular vectors of $\mathbf{A}$ onto the span of the block Krylov subspace, for any

⁴Any upper bound on $\|\mathbf{A}\|_2$ off by at most constant multiplicative factors suffices here. We can compute such an upper bound using $O(\log n)$ matrix-vector products via e.g., the power method or a single-vector Krylov method.

$p \leq l^5$. Notice that, for any q and for any polynomial $\phi(x)$ of degree $\leq 2q + 1$ containing only terms with odd powers, $\phi(\mathbf{A})\mathbf{X}$ lies exactly in the span of the Krylov subspace $\mathcal{K}_q$. Lemma 1 bounds $\|(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p\|_F^2$ in terms of the norms of the $\phi(\Sigma_p)$, $\phi(\Sigma_{l,\perp})$ and $\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}$ for any such polynomial $\phi(x)$ which is not zero at any of the top p singular values of $\mathbf{A}$. Later, by choosing such a polynomial $\phi(x)$ of degree at most $2q + 1$ with odd powers so that it is large at the top singular values of $\mathbf{A}$ and small at the rest, we can bound the projection error.

Lemma 1 (Angle between subspaces, generalization of Theorem 2.1 of [DIKMI18]). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ have $\text{rank}(\mathbf{A}) > l$ and SVD given by $\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^T$. Let $\mathbf{X} \in \mathbb{R}^{n \times l}$ be such that $\text{rank}(\mathbf{V}_l^T \mathbf{X}) = l$ and let $\mathbf{Q}$ be an orthonormal basis of the depth q Krylov subspace $\mathcal{K}_q(\mathbf{A}, \mathbf{X})$ (see Def. 2.1). For any $p \leq l$ and any polynomial $\phi(x)$ of degree $2q + 1$ with odd powers only, such that $\phi(\Sigma_p)$ is non-singular,*

$$\|(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p\|_F^2 \leq \|\phi(\Sigma_{l,\perp})\|_2^2 \cdot \|\phi(\Sigma_p)^{-1}\|_2^2 \cdot \|\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\|_2^2.$$

Proof. Let $\Phi = \phi(\mathbf{A})\mathbf{X} = \mathbf{U}\phi(\Sigma)\mathbf{V}^T \mathbf{X}$. As explained above, since $\phi(x)$ consists of only odd powers and has degree at most $2q + 1$, the columns of Φ lie in the span of the Krylov subspace $\mathcal{K}_q(\mathbf{A}, \mathbf{X})$ and thus $\text{range}(\Phi) \subseteq \text{range}(\mathcal{K}_q(\mathbf{A}, \mathbf{X})) = \text{range}(\mathbf{Q})$. $\Phi\Phi^\dagger$ and $\mathbf{Q}\mathbf{Q}^T$ are the orthogonal projectors onto $\text{range}(\Phi)$ and $\text{range}(\mathcal{K}_q(\mathbf{A}, \mathbf{X}))$ respectively. Since $\text{range}(\Phi) \subseteq \text{range}(\mathbf{Q})$ we have:

$$\|(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p\|_F^2 \leq \|(\mathbf{I} - \Phi\Phi^\dagger)\mathbf{U}_p\|_F^2. \quad (5)$$

Thus, to prove the lemma it suffices to upper bound the righthand side of (5). We write $\Phi = \Phi_l + \Phi_{l,\perp}$ where $\Phi_l = \phi(\mathbf{A}_l)\mathbf{X}$ and $\Phi_{l,\perp} = \phi(\mathbf{A}_{l,\perp})\mathbf{X}$. Since $p \leq l$ and ϕ is non-zero on the top p singular values, we have $\Phi_l\Phi_l^\dagger\mathbf{U}_p = \mathbf{U}_p$. We upper bound the righthand side of (5) by:

$$\|(\mathbf{I} - \Phi\Phi^\dagger)\mathbf{U}_p\|_F^2 = \|\mathbf{U}_p - \Phi(\Phi^\dagger\mathbf{U}_p)\|_F^2 = \min_{\Psi \in \mathbb{R}^{\ell \times p}} \|\mathbf{U}_p - \Phi\Psi\|_F^2 \leq \|\mathbf{U}_p - \Phi(\Phi_l^\dagger\mathbf{U}_p)\|_F^2.$$

Then, replacing Φ by $\Phi_l + \Phi_{l,\perp}$ and using that $\Phi_l\Phi_l^\dagger\mathbf{U}_p = \mathbf{U}_p$ we get:

$$\begin{aligned} \|(\mathbf{I} - \Phi\Phi^\dagger)\mathbf{U}_p\|_F^2 &\leq \|\mathbf{U}_p - \Phi(\Phi_l^\dagger\mathbf{U}_p)\|_F^2 \\ &= \|(\mathbf{I} - \Phi_l\Phi_l^\dagger)\mathbf{U}_p - \Phi_{l,\perp}\Phi_l^\dagger\mathbf{U}_p\|_F^2 \\ &= \|\Phi_{l,\perp}\Phi_l^\dagger\mathbf{U}_p\|_F^2 \\ &= \|\mathbf{U}_{l,\perp}\phi(\Sigma_{l,\perp})\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\phi(\Sigma_l)^{-1}\mathbf{U}_l^T \mathbf{U}_p\|_F^2 \\ &= \|\mathbf{U}_{l,\perp}\phi(\Sigma_{l,\perp})\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\phi(\Sigma_p^{-1})\|_F^2 \\ &= \|\phi(\Sigma_{l,\perp})\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\phi(\Sigma_p^{-1})\|_F^2 \\ &\leq \|\phi(\Sigma_{l,\perp})\|_F^2 \cdot \|\phi(\Sigma_p^{-1})\|_F^2 \cdot \|\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\|_F^2. \end{aligned}$$

In line 4 we used that $\mathbf{V}_l^T \mathbf{X} \in \mathbb{R}^{l \times l}$ has rank l by assumption, and is thus invertible. Combined with (5) the above bound completes the proof. $\square$

Using Lemma 1 we can show that for any p with $\sigma_p(\mathbf{A})$ larger than $\sigma_{l+1}(\mathbf{A})$ by a constant multiplicative factor, the Krylov subspace $\mathcal{K}_q(\mathbf{A}, \mathbf{X})$ generated with random width- l starting block $\mathbf{X} \in \mathbb{R}^{n \times l}$ and with depth $q = O(\log n)$ approximately spans the top p singular vectors of $\mathbf{A}$.

⁵In the numerical linear algebra literature, this quantity is often denoted as $\|(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p\|_F^2 = \sin(\mathbf{Q}, \mathbf{U}_p)$

Lemma 2 (Convergence of block Krylov subspace to top singular vectors). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ have $\text{rank}(\mathbf{A}) > l$ and SVD given by $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$. Let $\mathbf{X} \in \mathbb{R}^{n \times l}$ be a matrix with independent $\mathcal{N}(0, 1)$ Gaussian entries and let $\mathbf{Q}$ be an orthonormal basis for the depth $q = O(\log n)$ Krylov subspace $\mathcal{K}_q(\mathbf{A}, \mathbf{X})$ generated by $\mathbf{X}$ (see Def. 2.1). Then, for any $p \leq l$, such that $\sigma_p(\mathbf{A}) \geq \frac{3}{2}\sigma_{l+1}(\mathbf{A})$, for any constants $c, c' > 0$, with probability at least $1 - \frac{1}{n^{c'}}$,*

$$\|\mathbf{U}_p - \mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_F^2 \leq \frac{1}{n^c}.$$

Proof. Since $\mathbf{X} \in \mathbb{R}^{n \times l}$ is a random Gaussian matrix, with probability one, $\text{rank}(\mathbf{V}_l^T \mathbf{X}) = l$. Thus, we can apply Lemma 1 to give, for any degree $2q + 1$ polynomial ϕ consisting of only odd powers where $\phi(\mathbf{\Sigma}_p)$ is non-singular,

$$\|\mathbf{U}_p - \mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_F^2 \leq \|\phi(\mathbf{\Sigma}_{l,\perp})\|_2^2 \cdot \|\phi(\mathbf{\Sigma}_p)^{-1}\|_2^2 \cdot \|\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^\dagger\|_F^2. \quad (6)$$

We will now bound each term on the righthand side of (6). First, we bound $\|\phi(\mathbf{\Sigma}_{l,\perp})\|_2$ and $\|\phi(\mathbf{\Sigma}_p)^{-1}\|_2$ similarly to Lemma 2.4 of [DIKMI18]. We consider a gap amplifying polynomial $\phi(x)$ of degree $2q + 1$ consisting only of odd powers as defined in Lemma 4.5 of [DIKMI18] with parameters $\alpha = \sigma_{l+1}(\mathbf{A})$ and gap $\gamma = \frac{\sigma_p(\mathbf{A})}{\sigma_{l+1}(\mathbf{A})} - 1$. Observe that we have,

$$\|\phi(\mathbf{\Sigma}_p)^{-1}\|_2 = \max_{1 \leq i \leq p} \phi(\sigma_i(\mathbf{A}))^{-1} \leq \max_{1 \leq i \leq p} \sigma_i^{-1}(\mathbf{A}) = \sigma_p^{-1}(\mathbf{A}), \quad (7)$$

where we used the fact that $\sigma_i(\mathbf{A}) > 0$ for $i \leq p \leq l$ as $\text{rank}(\mathbf{A}) > l$ and the fact that for any $i \leq p$, $\phi(\sigma_i(\mathbf{A})) \geq \sigma_i(\mathbf{A})$ as $\sigma_i(\mathbf{A}) \geq \sigma_p(\mathbf{A}) \geq (1 + \gamma)\sigma_{l+1}(\mathbf{A}) = (1 + \gamma)\alpha$ which follows from Lemma 4.5 of [DIKMI18]. Also, as $\sigma_i(\mathbf{A}) \leq \sigma_{l+1}(\mathbf{A})$ for any $i \geq l + 1$, from Lemma 4.5 of [DIKMI18] we have:

$$\|\phi(\mathbf{\Sigma}_{l,\perp})\|_2 = \max_{i \geq l+1} |\phi(\sigma_i(\mathbf{A}))| \leq \frac{4\sigma_{l+1}(\mathbf{A})}{2^{(2q+1)\min(\sqrt{\gamma}, 1)}}. \quad (8)$$

Finally, we bound the middle term of (6), $\|\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\|_F^2 \leq \|\mathbf{V}_{l,\perp}^T \mathbf{X}\|_F^2 \cdot \|(\mathbf{V}_l^T \mathbf{X})^{-1}\|_2^2$. By rotational invariance of the Gaussian distribution, $\mathbf{V}_l^T \mathbf{X}$ and $\mathbf{V}_{l,\perp}^T \mathbf{X}$ are $l \times l$ and $n - l \times l$ random Gaussian matrices respectively. Thus, by Corollary 5.35 of [Ver18], with probability $1 - \frac{1}{n^{c_3}}$ we have $\|\mathbf{V}_{l,\perp}^T \mathbf{X}\|_2^2 \leq c_1 n$ and $\sigma_l^2(\mathbf{V}_l^T \mathbf{X}) \geq c_2 n$ for constants c_1, c_2 and c_3 . So, we have $\|\mathbf{V}_{l,\perp}^T \mathbf{X}\|_F^2 \leq c_1 n^2$ and $\|(\mathbf{V}_l^T \mathbf{X})^{-1}\|_2^2 \leq \sigma_l^{-2}(\mathbf{V}_l^T \mathbf{X}) \leq \frac{1}{c_2 n}$ which overall gives us that

$$\|\mathbf{V}_{l,\perp}^T \mathbf{X}(\mathbf{V}_l^T \mathbf{X})^{-1}\|_F^2 \leq c'' n \quad (9)$$

for some constant c'' . Plugging (9), (8) and (7) back into (6), we get:

$$\|\mathbf{U}_p - \mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_F^2 \leq \frac{c'' \sigma_{l+1}^2(\mathbf{A})}{\sigma_p^2(\mathbf{A})} \frac{n}{2^{(4q+2)\min(\sqrt{\gamma}, 1)}},$$

for some constant c'' . Since $\sigma_p(\mathbf{A}) \geq \frac{3}{2}\sigma_{l+1}(\mathbf{A})$, the gap is given by $\gamma = \frac{\sigma_p(\mathbf{A})}{\sigma_{l+1}(\mathbf{A})} - 1 \geq \frac{1}{2}$. This gives us $\min(\sqrt{\gamma}, 1) \geq \frac{1}{\sqrt{2}}$. Finally, choosing $q = C \log n$ where C is a large enough constant we obtain the final bound. $\square$

Lemma 2 establishes that the Krylov subspace generated by a random starting block $\mathbf{X} \in \mathbb{R}^{n \times l}$ will approximately span any singular vector corresponding to a singular value significantly larger than $\sigma_{l+1}(\mathbf{A})$. Intuitively, projection onto this subspace should thus preserve the largest singular values (and the largest magnitude eigenvalues) of $\mathbf{A}$. Below we prove several Lemmas that allow us to argue this formally. We consider the matrix $\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T$ – the projection of $\mathbf{A}$ onto the Krylov subspace on both the left and right.

3.1.1 Eigenvector Alignment

In this section, we first prove that the eigenvectors corresponding to the large magnitude eigenvalues of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ are also approximate eigenvectors of $\mathbf{A}$. We first prove that if there exists some $k \in [l]$ such that $\sigma_k(\mathbf{A})$ is larger than $\sigma_{l+1}(\mathbf{A})$ by at least a constant multiplicative factor, then there exists some $k \leq p \leq l$ such that the top p singular vectors of $\mathbf{A}$ approximately span the top p singular vectors of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ (i.e. the matrix $\mathbf{A}$ projected on both sides to the Krylov subspace). Note that if no such k exists, then $\|\mathbf{A}\|_2 = O(\sigma_{l+1}(\mathbf{A}))$ and thus there is nothing to gain from deflation.

Lemma 3 (Projection of $\mathbf{Z}_p$ on $\mathbf{U}_p$). *Consider the setting of Lemma 2 where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric. Let $\alpha = \max\left(\frac{\|\mathbf{A}\|_2}{n^{c/2}}, \sigma_{l+1}(\mathbf{A})\right)$ where $c > 0$ is the constant in Lemma 2. Let $k \in [l]$ be such that $\sigma_k(\mathbf{A}) \geq 2\alpha$ and $\sigma_{k+1}(\mathbf{A}) < 2\alpha$. Let $\mathbf{z}_1, \dots, \mathbf{z}_n$ be the eigenvectors of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ corresponding to its eigenvalues $|\lambda_1(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \geq \dots \geq |\lambda_n(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)|$. Let $c_1, c' > 0$ be some large constants. Then, there exists some $k \leq p \leq l$ such that $\sigma_p(\mathbf{A}) \geq \frac{3}{2}\alpha$, $\sigma_p(\mathbf{A}) - \sigma_{p+1}(\mathbf{A}) \geq \frac{\|\mathbf{A}\|_2}{2n^{c/2+1}}$, and, letting $\mathbf{Z}_p \in \mathbb{R}^{n \times p}$ have $\mathbf{z}_1, \dots, \mathbf{z}_p$ as its columns,*

$$\|\mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p - \mathbf{Z}_p\|_2 \leq \frac{1}{n^{c_1}},$$

with probability at least $1 - \frac{1}{n^{c'}}$.

Proof. Let l_1 be the largest index with $k \leq l_1 \leq l$ and $\sigma_{l_1}(\mathbf{A}) \geq \frac{3}{2}\alpha$, i.e., we must have $\sigma_{l_1+1}(\mathbf{A}) < \frac{3}{2}\alpha$. Note that since $\sigma_k(\mathbf{A}) \geq 2\alpha$, such an l_1 must exist. Then, $\sigma_k(\mathbf{A}) - \sigma_{l_1+1}(\mathbf{A}) \geq \alpha/2 \geq \frac{\|\mathbf{A}\|_2}{2n^{c/2}}$ where we use the fact that $\alpha \geq \frac{\|\mathbf{A}\|_2}{n^{c/2}}$. Since there are at most l_1 indices in the range $[k, l_1]$, there must be some $p \in [k, l_1]$ such that:

$$\sigma_p(\mathbf{A}) - \sigma_{p+1}(\mathbf{A}) \geq \frac{\sigma_k(\mathbf{A}) - \sigma_{l_1+1}(\mathbf{A})}{l_1} \geq \frac{\|\mathbf{A}\|_2}{2n^{c/2+1}}, \quad (10)$$

where we loosely bound $l_1 \leq n$. This establishes the first claim of the Lemma. We now prove the second claim.

Since $\sigma_p(\mathbf{A}) \geq \frac{3}{2}\alpha \geq \frac{3}{2}\sigma_{l+1}(\mathbf{A})$, from Lemma 2, we get that $\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p = \mathbf{U}_p + \mathbf{E}$ where $\|\mathbf{E}\|_2 \leq \frac{1}{n^c}$ with probability at least $1 - \frac{1}{n^{c'}}$ for some constant $c' > 0$. Let $\mathbf{Z} \in \mathbb{R}^{n \times n}$ be an orthonormal matrix containing all eigenvectors $\mathbf{z}_1, \dots, \mathbf{z}_n$ of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ as its columns. Since the columns of $\mathbf{Z}$ form an orthonormal basis of $\mathbb{R}^n$, there exists a matrix $\mathbf{C} \in \mathbb{R}^{n \times p}$ such that $\mathbf{U}_p = \mathbf{Z}\mathbf{C}$ and $\mathbf{C}^T\mathbf{C} = \mathbf{I}_p$ where $\mathbf{I}_p$ is the $p \times p$ identity matrix. Then, we have:

$$\mathbf{U}_p = \mathbf{Z}\mathbf{C} = \mathbf{Z}_p\mathbf{C}_1 + \mathbf{Z}_{p,\perp}\mathbf{C}_2, \quad (11)$$

where $\mathbf{C} = [\mathbf{C}_1; \mathbf{C}_2]$ for $\mathbf{C}_1 \in \mathbb{R}^{p \times p}$ and $\mathbf{C}_2 \in \mathbb{R}^{n-p \times p}$. We will now prove that $\|\mathbf{C}_2\|_2$ is very small and $\mathbf{C}_1$ is very close to the identity matrix. This implies that $\mathbf{U}_p$ approximately spans $\mathbf{Z}_p$ which proves our second claim.

First observe that we can write $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p$ as:

$$\begin{aligned} \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p &= \mathbf{Q}\mathbf{Q}^T\mathbf{A}(\mathbf{U}_p + \mathbf{E}) \\ &= \mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\mathbf{A}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E} \\ &= (\mathbf{U}_p + \mathbf{E})\mathbf{A}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E} \\ &= \mathbf{U}_p\mathbf{A}_p + \mathbf{E}\mathbf{A}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}. \end{aligned} \quad (12)$$

Thus, we get:

$$\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p = \mathbf{Z}\mathbf{C}\mathbf{\Lambda}_p + \mathbf{E}\mathbf{\Lambda}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}. \quad (13)$$

Next, observe that, since $\mathbf{Z}$ has columns equal to the eigenvalues of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$, we can also write $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p$ as:

$$\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p = \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{Z}\mathbf{C} = \mathbf{Z}\tilde{\mathbf{\Lambda}}\mathbf{C}, \quad (14)$$

where $\tilde{\mathbf{\Lambda}} \in \mathbb{R}^{n \times n}$ denotes the diagonal matrix containing the eigenvalues of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ on its diagonal. Thus, combining (13) and (14) we get that $\mathbf{Z}\tilde{\mathbf{\Lambda}}\mathbf{C} = \mathbf{Z}\mathbf{C}\mathbf{\Lambda}_p + \mathbf{E}\mathbf{\Lambda}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}$ or

$$\mathbf{Z}\tilde{\mathbf{\Lambda}}\mathbf{C} - \mathbf{Z}\mathbf{C}\mathbf{\Lambda}_p = \mathbf{E}\mathbf{\Lambda}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}.$$

Let $\mathbf{E}' = \mathbf{E}\mathbf{\Lambda}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}$. Using triangle inequality, we get that $\|\mathbf{E}'\|_2 \leq \|\mathbf{E}\mathbf{\Lambda}_p\|_2 + \|\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}\|_2$. Now, since $\|\mathbf{E}\|_2 \leq \frac{1}{n^c}$ and $\|\mathbf{\Lambda}_p\|_2 = \|\mathbf{A}\|_2$, we have $\|\mathbf{E}\mathbf{\Lambda}_p\|_2 \leq \|\mathbf{E}\|_2\|\mathbf{\Lambda}_p\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^c}$. Similarly, since $\mathbf{Q}$ is orthonormal, $\|\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^c}$. So, we get $\|\mathbf{E}'\|_2 \leq \frac{2\|\mathbf{A}\|_2}{n^c}$ which implies that $\|\mathbf{Z}\mathbf{C}\mathbf{\Lambda}_p - \mathbf{Z}\tilde{\mathbf{\Lambda}}\mathbf{C}\|_2 \leq \frac{2\|\mathbf{A}\|_2}{n^c}$ or, equivalently, since $\mathbf{Z}$ has orthonormal columns,

$$\|\mathbf{C}\mathbf{\Lambda}_p - \tilde{\mathbf{\Lambda}}\mathbf{C}\|_2 \leq \frac{2\|\mathbf{A}\|_2}{n^c}.$$

For any $i \in [p]$, the i^{th} column of the $n \times p$ matrix $\mathbf{C}\mathbf{\Lambda}_p - \tilde{\mathbf{\Lambda}}\mathbf{C}$ is given by $\lambda_i(\mathbf{A})\mathbf{C}_{:,i} - \tilde{\mathbf{\Lambda}}\mathbf{C}_{:,i}$ where $\mathbf{C}_{:,i}$ is the i th column of $\mathbf{C}$. So, we have $\|\lambda_i(\mathbf{A})\mathbf{C}_{:,i} - \tilde{\mathbf{\Lambda}}\mathbf{C}_{:,i}\|_2 \leq \frac{2\|\mathbf{A}\|_2}{n^c}$ for all $i \in [p]$. Using the definition of the l_2 norm of the vector, we have $\sqrt{\sum_{j=1}^n (\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T))^2 \mathbf{C}_{ji}^2} \leq \frac{2\|\mathbf{A}\|_2}{n^c}$. This implies that for all $i \in [p]$ and $j \in [n]$,

$$|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| |\mathbf{C}_{ji}| \leq \frac{2\|\mathbf{A}\|_2}{n^c}. \quad (15)$$

Now, by the minimax principle of singular values, $|\lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq |\lambda_{p+1}(\mathbf{A})|$ for all $j \geq p+1$. Thus, for any $i \in [p]$ and $j \in \{p+1, \dots, n\}$, by the triangle inequality, $|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \geq |\lambda_i(\mathbf{A})| - |\lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \geq |\lambda_p(\mathbf{A})| - |\lambda_{p+1}(\mathbf{A})| \geq \frac{\|\mathbf{A}\|_2}{2n^{c/2+1}}$ where the last step follows from the first claim (10). Thus, from (15), we get that for any $i \in [p]$ and $j \in \{p+1, \dots, n\}$,

$$|\mathbf{C}_{ji}| \leq \frac{4}{n^{c/2-1}}.$$

Note that from (11), we have that $\mathbf{C}_2$ contains all $\mathbf{C}_{ji}$ such that $j \in \{p+1, \dots, n\}$ and $i \in [p]$. So, we can bound:

$$\|\mathbf{C}_2\|_2 \leq \|\mathbf{C}_2\|_F \leq \sqrt{\sum_{i=1}^p \sum_{j=p+1}^n \mathbf{C}_{ji}^2} \leq \frac{4}{n^{c/2-2}}. \quad (16)$$

Next, since $\mathbf{C}^T\mathbf{C} = \mathbf{I}_p = \mathbf{C}_1^T\mathbf{C}_1 + \mathbf{C}_2^T\mathbf{C}_2$, we have $\|\mathbf{C}_1^T\mathbf{C}_1 - \mathbf{I}_p\|_2 = \|\mathbf{C}_2^T\mathbf{C}_2\|_2 = \|\mathbf{C}_2\|_2^2 \leq \frac{16}{n^{2(c/2-2)}}$. Thus, we can write $\mathbf{C}_1^T\mathbf{C}_1 = \mathbf{I} + \mathbf{E}_1$ where $\|\mathbf{E}_1\|_2 \leq \frac{16}{n^{2(c/2-2)}}$. Observe that $\mathbf{C}_1\mathbf{C}_1^T$ and $\mathbf{C}_1^T\mathbf{C}_1$ have the same eigenvalues, and thus we also have

$$\mathbf{C}_1\mathbf{C}_1^T = \mathbf{I} + \mathbf{E}'', \quad (17)$$

where $\|\mathbf{E}''\|_2 \leq \frac{16}{n^{2(c/2-2)}}$. Since we have $\|\mathbf{C}_1\|_2 = \sqrt{\|\mathbf{C}_1\mathbf{C}_1^T\|_2} = \sqrt{\|\mathbf{I} + \mathbf{E}_1\|_2} \leq \sqrt{1 + \|\mathbf{E}\|_1} \leq 1 + \frac{4}{n^{c/2-2}}$ (where the second to last step follows using triangle inequality), we also have that:

$$\mathbf{C}_1 = \mathbf{I} + \mathbf{E}_2, \quad (18)$$

where $\|\mathbf{E}_2\|_2 \leq \frac{4}{n^{c/2-2}}$. Then, we have:

$$\begin{aligned} \mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p - \mathbf{Z}_p &= \mathbf{U}_p (\mathbf{C}_1^T \mathbf{Z}_p^T + \mathbf{C}_2^T \mathbf{Z}_{p,\perp}^T) \mathbf{Z}_p - \mathbf{Z}_p \\ &= \mathbf{U}_p \mathbf{C}_1^T - \mathbf{Z}_p \\ &= \mathbf{Z}_p \mathbf{C}_1 \mathbf{C}_1^T + \mathbf{Z}_{p,\perp} \mathbf{C}_2 \mathbf{C}_1^T - \mathbf{Z}_p \\ &= \mathbf{Z}_p (\mathbf{I} + \mathbf{E}'') + \mathbf{Z}_{p,\perp} \mathbf{C}_2 \mathbf{C}_1^T - \mathbf{Z}_p \\ &= \mathbf{Z}_p \mathbf{E}'' + \mathbf{Z}_{p,\perp} \mathbf{C}_2 \mathbf{C}_1^T \\ &= \mathbf{Z}_p \mathbf{E}'' + \mathbf{Z}_{p,\perp} \mathbf{C}_2 (\mathbf{I} + \mathbf{E}_2^T) \\ &= \mathbf{Z}_p \mathbf{E}'' + \mathbf{Z}_{p,\perp} \mathbf{C}_2 + \mathbf{Z}_{p,\perp} \mathbf{C}_2 \mathbf{E}_2^T. \end{aligned}$$

The first and third equality above follows from (11), the fourth equality follows from (17) and the sixth equality follows from (16). Thus, using triangle inequality and spectral submultiplicativity, we get $\|\mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p - \mathbf{Z}_p\|_2 \leq \|\mathbf{Z}_p \mathbf{E}''\|_2 + \|\mathbf{Z}_{p,\perp} \mathbf{C}_2\|_2 + \|\mathbf{Z}_{p,\perp} \mathbf{C}_2 \mathbf{C}_1^T\|_2 \leq \|\mathbf{E}''\|_2 + \|\mathbf{C}_2\|_2 + \|\mathbf{C}_1\|_1 \|\mathbf{C}_2\|_2 \leq \frac{16}{n^{c-4}} + \frac{4}{n^{c/2-2}} + \frac{16}{n^{c-4}} \leq \frac{36}{n^{c/2-2}}$ where the second inequality follows from the fact that $\mathbf{Z}_p$ and $\mathbf{Z}_{p,\perp}$ are orthonormal matrices and the third inequality follows from the error bounds in (16), (17) and (18). Choosing the constant c_1 to be suitably large enough so that $\|\mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p - \mathbf{Z}_p\|_2$ is bounded by $\frac{1}{n^{c_1}}$ completes the proof. $\square$

The next theorem proves the main result of this subsection, i.e., the top eigenvectors of $\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T$ are also approximately the eigenvectors of $\mathbf{A}$, provided there is some $\sigma_k(\mathbf{A})$ that is constant factor larger than $\sigma_{l+1}(\mathbf{A})$, as assumed in Lemma 3.

Theorem 5 (Convergence error of top p eigenvectors). *Consider the setting of Lemma 3. Let $c_1, c' > 0$ be some large constants. Let $\tilde{\Lambda}_p \in \mathbb{R}^{p \times p}$ be a diagonal matrix containing the corresponding eigenvalues $\lambda_1(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T), \dots, \lambda_p(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T)$ on its diagonal. Then, we have*

$$\|\mathbf{A} \mathbf{Z}_p - \mathbf{Z}_p \tilde{\Lambda}_p\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{c_1}}.$$

with probability at least $1 - \frac{1}{n^{c'}}$.

Proof. From Lemma 3, $\mathbf{Z}_p = \mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_1$ where $\|\mathbf{E}_1\|_2 \leq \frac{1}{n^{c_2}}$ with probability at least $1 - \frac{1}{n^{c''}}$ for some constant c_2, c'' . Also, from Lemma 2, we have $\mathbf{U}_p = \mathbf{Q}\mathbf{Q}^T \mathbf{U}_p + \mathbf{E}_2$ where $\|\mathbf{E}_2\|_2 \leq \frac{1}{n^{c_3}}$

with probability at least $1 - \frac{1}{n^{c''}}$ for some constant c_3 . Observe that:

$$\begin{aligned}
\mathbf{AZ}_p &\stackrel{a}{=} \mathbf{A}(\mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_1) \\
&= \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1 \\
&\stackrel{b}{=} \mathbf{U}_p \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1 \\
&\stackrel{c}{=} (\mathbf{QQ}^T \mathbf{U}_p + \mathbf{E}_2) \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1 \\
&= \mathbf{QQ}^T \mathbf{U}_p \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_2 \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1 \\
&\stackrel{d}{=} \mathbf{QQ}^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_2 \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1 \\
&\stackrel{e}{=} \mathbf{QQ}^T \mathbf{A}(\mathbf{QQ}^T \mathbf{U}_p + \mathbf{E}_2) \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_2 \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1 \\
&= \mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{QQ}^T \mathbf{AE}_2 \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_2 \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1. \\
&\stackrel{f}{=} \mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{Z}_p - \mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{E}_1 + \mathbf{QQ}^T \mathbf{AE}_2 \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_2 \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{AE}_1.
\end{aligned}$$

In the above set of equations, (a) follows by replacing $\mathbf{Z}_p$ with $\mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p + \mathbf{E}_1$, (b) and (d) follows from the fact that $\mathbf{AU}_p \mathbf{U}_p^T = \mathbf{U}_p \mathbf{U}_p^T \mathbf{AU}_p \mathbf{U}_p^T$ as $\mathbf{U}_p$ contains the eigenvectors of $\mathbf{A}$, (c) and (e) follows from replacing $\mathbf{U}_p$ in the first term by $\mathbf{QQ}^T \mathbf{U}_p + \mathbf{E}_2$, (f) follows from replacing the $\mathbf{U}_p \mathbf{U}_p^T \mathbf{Z}_p$ in the first term by $\mathbf{Z}_p - \mathbf{E}_1$. Now, $\mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{Z}_p = \mathbf{Z}_p \tilde{\mathbf{\Lambda}}_p$. Also, using spectral submultiplicativity, $\|\mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{E}_1\|_2 \leq \|\mathbf{A}\|_2 \|\mathbf{E}_1\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{c_2}}$. Similarly, each of the last three terms in the final step above can be bounded by $\max\left(\frac{\|\mathbf{A}\|_2}{n^{c_2}}, \frac{\|\mathbf{A}\|_2}{n^{c_3}}\right)$. Thus, using triangle inequality, and after adjusting the constants c_1 and c' appropriately, we have:

$$\|\mathbf{AZ}_p - \mathbf{Z}_p \tilde{\mathbf{\Lambda}}_p\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{c_1}},$$

with probability at least $1 - \frac{1}{n^{c'}}$. □

3.1.2 Eigenvalue Alignment

In this section, we show that as a consequence of Lemma 2, the large magnitude eigenvalues of $\mathbf{A}$ are approximated by those of $\mathbf{QQ}^T \mathbf{AQQ}^T$. We first state a result from [Par98] (Theorem 11.5.1) which states that for a symmetric matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$ and any orthonormal matrix $\mathbf{C} \in \mathbb{R}^{n \times m}$ the eigenvalues of $\mathbf{D}$ can be put into one-to-one correspondence with those of $\mathbf{C}^T \mathbf{D} \mathbf{C}$ such that the error is bounded by the spectral norm of $\mathbf{D} \mathbf{C} - \mathbf{C} \mathbf{C}^T \mathbf{D} \mathbf{C}$.

Theorem 6 (Theorem 11.5.1 of [Par98]). *Let $\mathbf{D} \in \mathbb{R}^{n \times n}$ be a symmetric matrix and $\mathbf{C} \in \mathbb{R}^{n \times m}$ be a matrix with orthonormal columns. Then, there exists m eigenvalues of $\mathbf{D}$, $\{\alpha_i \mid i = 1, \dots, m\}$ such that for $i \in [m]$:*

$$|\alpha_i - \lambda_i(\mathbf{C}^T \mathbf{D} \mathbf{C})| \leq \|\mathbf{D} \mathbf{C} - \mathbf{C} \mathbf{C}^T \mathbf{D} \mathbf{C}\|_2.$$

We now prove a couple of Lemmas which we will use along with Theorem 6 for our main theorem.

Lemma 4. *Consider the setting of Lemma 2 where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric. Let the error bound for Lemma 2 hold for a constant $c > 0$. For some constant $c' > 0$, with probability at least $1 - \frac{1}{n^{c'}}$,*

$$\|\mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{U}_p - \mathbf{U}_p \mathbf{U}_p^T \mathbf{QQ}^T \mathbf{AQQ}^T \mathbf{U}_p\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{c-1}}.$$

Proof. Recall that we denote the eigendecomposition of a symmetric matrix $\mathbf{A}$ by $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$. $\mathbf{A}$'s eigenvectors (i.e., the columns of $\mathbf{U}$) are equal to its singular vectors, and recall that $\mathbf{U}_p \in \mathbb{R}^{n \times p}$ has columns equal to the p eigenvectors corresponding to the p largest magnitude eigenvalues (i.e., the p singular vectors corresponding to the p largest singular values). From Lemma 2, we have $\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p = \mathbf{U}_p + \mathbf{E}$ where $\mathbf{E} \in \mathbb{R}^{n \times p}$ has $\|\mathbf{E}\|_2 \leq \frac{1}{n^c}$ with probability at least $1 - \frac{1}{n^{c'}}.$ Thus, from (12), we get

$$\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p = \mathbf{U}_p\mathbf{\Lambda}_p + \mathbf{E}\mathbf{\Lambda}_p + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{E}. \quad (19)$$

Also, we have:

$$\begin{aligned} \mathbf{U}_p\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p &= \mathbf{U}_p(\mathbf{U}_p^T + \mathbf{E}^T)\mathbf{A}(\mathbf{U}_p + \mathbf{E}) \\ &= \mathbf{U}_p\mathbf{U}_p^T\mathbf{A}\mathbf{U}_p + \mathbf{U}_p\mathbf{U}_p^T\mathbf{A}\mathbf{E} + \mathbf{U}_p\mathbf{E}^T\mathbf{A}\mathbf{U}_p + \mathbf{U}_p\mathbf{E}^T\mathbf{A}\mathbf{E} \\ &= \mathbf{U}_p\mathbf{\Lambda}_p + \mathbf{U}_p\mathbf{U}_p^T\mathbf{A}\mathbf{E} + \mathbf{U}_p\mathbf{E}^T\mathbf{A}\mathbf{U}_p + \mathbf{U}_p\mathbf{E}^T\mathbf{A}\mathbf{E}. \end{aligned}$$

Then, using triangle inequality and spectral submultiplicativity, we have that

$$\|\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p - \mathbf{U}_p\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_2 \leq \|\mathbf{\Lambda}_p\|_2\|\mathbf{E}\|_2 + 3\|\mathbf{A}\|_2\|\mathbf{E}\|_2 + \|\mathbf{A}\|\|\mathbf{E}\|_2^2 \leq \frac{5\|\mathbf{A}\|_2}{n^c}.$$

This completes the proof. $\square$

Next we prove that the eigenvalues of $\mathbf{A}$ are well approximated by those of $\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p$.

Lemma 5. *Consider the setting of Lemma 2 where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric. Let the error bound for Lemma 2 hold for a constant $c > 0$. For some constant $c' > 0$, with probability at least $1 - \frac{1}{n^{c'}}$, we have for all $i \in [p]$,*

$$|\lambda_i(\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p) - \lambda_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\|_2}{n^{c-1}}.$$

Proof. To prove the theorem we will show that $\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p$ is close to $\mathbf{U}_p^T\mathbf{A}\mathbf{U}_p = \mathbf{\Lambda}_p$ in the spectral norm, and thus has nearby eigenvalues. Observe that

$$\begin{aligned} \mathbf{U}_p^T\mathbf{A}\mathbf{U}_p - \mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p &= \mathbf{U}_p^T(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p + \mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p \\ &\quad + \mathbf{U}_p^T(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{A}(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p. \end{aligned} \quad (20)$$

Since all the terms on the righthand side above contain the term $(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{U}_p$, we can use Lemma 2 to show that they are small. For the first term, we have:

$$\|\mathbf{U}_p^T(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_2 \leq \|\mathbf{U}_p^T(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\|_2 \cdot \|\mathbf{A}\|_2 \cdot \|\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^c},$$

where we also use the fact $\mathbf{Q}$ and $\mathbf{U}_p$ are orthonormal matrices and that $\|\mathbf{U}_p^T(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\|_2 \leq \|\mathbf{U}_p^T(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\|_F \leq \frac{1}{n^c}$ by Lemma 2. Similarly we can bound the other two terms in (20) by $\frac{\|\mathbf{A}\|_2}{n^c}$. Applying triangle inequality, we obtain:

$$\|\mathbf{U}_p^T\mathbf{A}\mathbf{U}_p - \mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_2 \leq \frac{3\|\mathbf{A}\|_2}{n^c}.$$

Observe that $\mathbf{U}_p^T\mathbf{A}\mathbf{U}_p = \mathbf{\Lambda}_p$ has eigenvalues $\lambda_1(\mathbf{A}), \dots, \lambda_p(\mathbf{A})$. So, by Weyl's inequality (Fact 1), we have that $|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p)| \leq \|\mathbf{U}_p^T\mathbf{A}\mathbf{U}_p - \mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_2 \leq \frac{3\|\mathbf{A}\|_2}{n^c}$ for $i \in [p]$. This completes the proof. $\square$

The next theorem is the main result of this section which states that the top p eigenvalues of $\mathbf{A}$ are approximated by the top p eigenvalues of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ up to some permutation for some $k \leq p \leq l$ as long as the conditions in Lemma 3 are satisfied. i.e., there exists some $k \in [l]$ with at least some constant multiplicative gap between $\sigma_k(\mathbf{A})$ and $\sigma_{l+1}(\mathbf{A})$.

Theorem 7 (Eigenvalue alignment). *Consider the setting of Theorem 5 and Lemma 3. Let $c_1, c' > 0$ be some constants. Then, there exists a permutation $S : [p] \rightarrow [p]$ such that for every $i \in [p]$, we have*

$$|\lambda_i(\mathbf{A}) - \lambda_{S(i)}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq \frac{\|\mathbf{A}\|_2}{n^{c_1}},$$

with probability at least $1 - \frac{1}{n^{c'}}$.

Proof. We apply Theorem 6 to first prove that each of the top p eigenvalues of $\mathbf{A}$ has an eigenvalue of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ close to it. Let $\mathbf{C} = \mathbf{U}_p$ and $\mathbf{D} = \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$. Let $c > 0$ be the constant in the statements of Lemmas 2, 4 and 5. Then, $\mathbf{C}^T\mathbf{D}\mathbf{C} = \mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p$ and by Theorem 6, there exist p eigenvalues $\alpha_1, \dots, \alpha_p$ of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ such that for $i \in [p]$:

$$\begin{aligned} |\alpha_i - \lambda_i(\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p)| &\leq \|\mathbf{D}\mathbf{C} - \mathbf{C}\mathbf{C}^T\mathbf{D}\mathbf{C}\|_2 \\ &= \|\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p - \mathbf{U}_p\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p\|_2 \\ &\leq \frac{\|\mathbf{A}\|_2}{n^{c-1}}. \end{aligned} \quad (21)$$

The last inequality hold with probability at least $1 - \frac{1}{n^{c'}}$ for some constant $c' > 0$ due to Lemma 4. From Lemma 5, with probability at least $1 - \frac{1}{n^{c'}}$, we get that for all $i \in [p]$, we have:

$$|\lambda_i(\mathbf{U}_p^T\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{U}_p) - \lambda_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\|_2}{n^{c-1}}. \quad (22)$$

Combining (21) and (22), and using triangle inequality we have for every $i \in [p]$ (for some constant $c_4 > 0$),

$$|\alpha_i - \lambda_i(\mathbf{A})| \leq \frac{2\|\mathbf{A}\|_2}{n^{c-1}}. \quad (23)$$

We now prove that the α_i 's (for all $i \in [p]$) are a permutation over $\lambda_1(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T), \dots, \lambda_p(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)$. Suppose that this is not true, i.e., there exists some $j \in [p]$ such that $\alpha_j = \lambda_{p+r}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)$ for some $r \geq 1$. So we have $|\lambda_{p+r}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) - \lambda_j(\mathbf{A})| \leq \frac{2\|\mathbf{A}\|_2}{n^{c-1}}$. Now, by the minimax principle we have

$$|\lambda_{p+r}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq |\lambda_{p+1}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq |\lambda_{p+1}(\mathbf{A})|.$$

Using triangle inequality, we have:

$$|\lambda_p(\mathbf{A})| - |\lambda_{p+1}(\mathbf{A})| \leq |\lambda_j(\mathbf{A})| - |\lambda_{p+r}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq |\lambda_j(\mathbf{A}) - \lambda_{p+r}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq \frac{2\|\mathbf{A}\|_2}{n^{c-1}}.$$

From Lemma 3, we have $\sigma_p(\mathbf{A}) - \sigma_{p+1}(\mathbf{A}) \geq \frac{\|\mathbf{A}\|_2}{2n^{c/2+1}}$ for some constant c . For a large enough c , we have $\frac{2\|\mathbf{A}\|_2}{n^{c-1}} < \frac{\|\mathbf{A}\|_2}{2n^{c/2+1}}$ and thus, we have a contradiction. So we must have $\gamma_j = \lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)$ for some $i \in [p]$. Choosing the constant c_1 to be suitably gives us the bound. $\square$

3.1.3 Bounding the spectral norm after deflation

In this section, we bound the spectral norm of the matrix $\mathbf{A}$ after deflating its top subspace using the converged eigenvectors from the randomized block Krylov algorithm (Algorithm 2). More formally, let $\mathbf{Z}_S = \mathbf{Q}\mathbf{V}_S$ be the output of Algorithm 2. Let $\mathbf{P} = \mathbf{I} - \mathbf{Z}_S\mathbf{Z}_S^T$, i.e., the projection matrix onto the subspace orthogonal to $\mathbf{Z}_S$. Then, we bound $\|\mathbf{P}\mathbf{A}\mathbf{P}\|_2$ based on the fact that $\mathbf{Z}_S$ must contain the top p eigenvectors of $\mathbf{A}$ for some $p \leq l$ as proven in Theorems 5 and 7, provided the assumptions in those theorems hold for $\mathbf{A}$. We first bound the spectral norm after deflating exactly the top p subspace of $\mathbf{A}$.

Lemma 6 (Spectral norm bound after deflation). *Consider the setting of Theorems 5 and 7. Let $c_1, c' > 0$ be some constants. Let $\mathbf{Z}_p \in \mathbb{R}^{n \times p}$ be as defined in Theorem 5. Then,*

$$\|(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\|_2 \leq \sigma_{p+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_1}},$$

with probability at least $1 - \frac{1}{n^{c'}}$.

Proof. Let the bounds from Theorem 5 and 7 hold with some constant $c_2 > 0$. Assume for contradiction, that $\|(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\|_2 > \sigma_{p+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2/2-1}}$. Then, there exists some eigenvector $\mathbf{x}$ of $(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)$ with corresponding eigenvalue λ such that $|\lambda| > \sigma_{p+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2/2-1}}$. Since any eigenvector corresponding to a nonzero eigenvalue of $(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)$ must lie in the column space of $\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T$, they will be orthogonal to $\mathbf{Z}_p$ and thus, we will have $\mathbf{Z}_p^T\mathbf{x} = 0$. Since we have $(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{A}(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{x} = \lambda\mathbf{x}$, we get that $(\mathbf{I} - \mathbf{Z}_p\mathbf{Z}_p^T)\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$. Then, we have

$$\|\mathbf{A}\mathbf{x}\|_2^2 \geq \lambda^2 \geq \sigma_{p+1}^2(\mathbf{A}) + \frac{2\sigma_{p+1}(\mathbf{A})\|\mathbf{A}\|_2}{n^{c_2/2-1}} + \frac{\|\mathbf{A}\|_2^2}{n^{c_2-2}}. \quad (24)$$

Let $\mathbf{Z}' \in \mathbb{R}^{n \times p+1}$ such that the first p columns of $\mathbf{Z}'$ are the columns of $\mathbf{Z}_p$ and $(p+1)$ th column is $\mathbf{x}$. From Theorem 5, we have that $\mathbf{A}\mathbf{Z}_p = \mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p + \mathbf{E}$ where $\|\mathbf{E}\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{c_2}}$. Then, we have $\|\mathbf{E}\|_F \leq \frac{\sqrt{n}\|\mathbf{A}\|_2}{n^{c_2}} \leq \frac{\|\mathbf{A}\|_2}{n^{c_2-0.5}}$ and thus, we get:

$$\begin{aligned} \|\mathbf{A}\mathbf{Z}_p\|_F^2 &\geq (\|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F - \|\mathbf{E}\|_F)^2 \\ &\geq (\|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F - \frac{\|\mathbf{A}\|_2}{n^{c_2-0.5}})^2 \end{aligned} \quad (25)$$

$$\begin{aligned} &\geq \|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F^2 + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - 2\|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F \frac{\|\mathbf{A}\|_2}{n^{c_2-0.5}} \\ &\geq \|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F^2 + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - \frac{2\|\mathbf{A}\|_2^2}{n^{c_2-1}} \\ &= \sum_{i=1}^p \sigma_i^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - \frac{2\|\mathbf{A}\|_2^2}{n^{c_2-1}}, \end{aligned} \quad (26)$$

for some constant $C > 0$. The first inequality above follows from the triangle inequality, the second from the bound $\|\mathbf{E}\|_F \leq \frac{\|\mathbf{A}\|_2}{n^{c_2-0.5}}$ and the third from expanding the quadratic expression. The fourth bound follows from $\|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F = \|\tilde{\mathbf{\Lambda}}_p\|_F \leq \sqrt{n}\|\tilde{\mathbf{\Lambda}}_p\|_2 \leq \sqrt{n}\|\mathbf{A}\|_2$ (where we use the fact that $\|\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\|_2 \leq \|\mathbf{A}\|_2$ and $\mathbf{Z}_p$ is an orthonormal matrix) which gives us $\|\mathbf{Z}_p\tilde{\mathbf{\Lambda}}_p\|_F\|\mathbf{E}\|_F \leq \frac{\|\mathbf{A}\|_2^2}{n^{c_2-1}}$.

The final step follows from the fact that $\mathbf{Z}_p$ and $\tilde{\mathbf{A}}_p$ contain the top p eigenvectors and eigenvalues of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$. By the Pythagorean theorem, we have

$$\begin{aligned}\|\mathbf{AZ}'\|_F^2 &= \|\mathbf{AZ}_p\|_F^2 + \|\mathbf{Ax}\|_2^2 > \sum_{i=1}^p \sigma_i^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - \frac{2\|\mathbf{A}\|_2^2}{n^{c_2-1}} \\ &\quad + \sigma_{p+1}^2(\mathbf{A}) + \frac{2\sigma_{p+1}(\mathbf{A})\|\mathbf{A}\|_2}{n^{c_2/2-1}} + \frac{\|\mathbf{A}\|_2^2}{n^{c_2-2}},\end{aligned}\quad (27)$$

where the last inequality follows from the lower bounds in (24) and (25). From Theorem 7, there exists a permutation $S : [p] \rightarrow [p]$ such that for every $i \in [p]$ (and for some constant $C' > 0$), $|\lambda_i(\mathbf{A}) - \lambda_{S(i)}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq \frac{\|\mathbf{A}\|_2}{n^{c_2}}$ or

$$|\lambda_i^2(\mathbf{A}) - \lambda_{S(i)}^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq \frac{|\lambda_i(\mathbf{A}) + \lambda_{S(i)}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)|\|\mathbf{A}\|_2}{n^{c_2}} \leq \frac{2\|\mathbf{A}\|_2^2}{n^{c_2}},$$

where we upper bounded $\lambda_i(\mathbf{A})$ and $\lambda_{S(i)}(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)$ by $\|\mathbf{A}\|_2$. Thus, we have $\lambda_i^2(\mathbf{A}) \leq \lambda_{S(i)}^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{2\|\mathbf{A}\|_2^2}{n^{c_2}}$ or $\sigma^2(\mathbf{A}) \leq \sigma_{S(i)}^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{2\|\mathbf{A}\|_2^2}{n^{c_2}}$ for every $i \in [p]$. Adding up both sides over $i \in [p]$ we get that $\sum_{i=1}^p \sigma_i^2(\mathbf{A}) \leq \sum_{i=1}^p \sigma_i^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{2p\|\mathbf{A}\|_2^2}{n^{c_2}} \leq \sum_{i=1}^p \sigma_i^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{2\|\mathbf{A}\|_2^2}{n^{c_2-1}}$. So, we get:

$$\begin{aligned}\sum_{i=1}^p \sigma_i^2(\mathbf{A}) + \sigma_{p+1}^2(\mathbf{A}) &\leq \sum_{i=1}^p \sigma_i^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{2\|\mathbf{A}\|_2^2}{n^{c_2-1}} + \sigma_{p+1}^2(\mathbf{A}) \\ &= \sum_{i=1}^p \sigma_i^2(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - \frac{2\|\mathbf{A}\|_2^2}{n^{c_2-1}} + \sigma_{p+1}^2(\mathbf{A}) + \frac{\|\mathbf{A}\|_2^2}{n^{c_2-2}} \\ &\quad + \frac{2\sigma_{p+1}(\mathbf{A})\|\mathbf{A}\|_2}{n^{c_2/2-1}} - \left(\frac{\|\mathbf{A}\|_2^2}{n^{c_2-2}} + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - \frac{4\|\mathbf{A}\|_2^2}{n^{c_2-1}} + \frac{2\sigma_{p+1}(\mathbf{A})\|\mathbf{A}\|_2}{n^{c_2/2-1}} \right).\end{aligned}\quad (28)$$

For large enough n and c_2 , we have $\frac{\|\mathbf{A}\|_2^2}{n^{c_2-2}} + \frac{\|\mathbf{A}\|_2^2}{n^{2c_2-1}} - \frac{4\|\mathbf{A}\|_2^2}{n^{c_2-1}} + \frac{2\sigma_{p+1}(\mathbf{A})\|\mathbf{A}\|_2}{n^{c_2/2-1}} > 0$. Thus from the lower bound (27) on $\|\mathbf{AZ}'\|_F^2$ and from (28), we get $\sum_{i=1}^{p+1} \sigma_i^2(\mathbf{A}) < \|\mathbf{AZ}'\|_F^2$. But, by the minmax principle for singular values, we have $\max_{\mathbf{V} \in \mathbb{R}^{n \times (p+1)}, \mathbf{V}^T \mathbf{V} = \mathbf{I}} \|\mathbf{AV}\|_F^2 = \sum_{i=1}^{p+1} \sigma_i^2(\mathbf{A})$, which results in a contradiction as $\mathbf{Z}'$ is an $n \times (p+1)$ orthonormal matrix. Thus, we must have $\|(\mathbf{I} - \mathbf{Z}_p \mathbf{Z}_p^T) \mathbf{A} (\mathbf{I} - \mathbf{Z}_p \mathbf{Z}_p^T)\|_2 \leq \sigma_{p+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2/2-1}}$. $\square$

We now prove our main result which bounds the spectral norm of the deflated matrix $\mathbf{PAP}$.

Theorem 8. *Consider the setting of Theorem 5 and let the error bound of Theorem 5 hold for some constant $c_3 > 0$. Let Algorithm 2 be run with the inputs $\mathbf{A} \in \mathbb{R}^{n \times n}$, block size $l \in [n]$ and $\beta = c_3/2$. Let S be the set of indices as defined in Algorithm 2 such that for any $i \in S$, the eigenvector $\mathbf{v}_i$ and corresponding eigenvalue $\lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$ of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ satisfies $\|\mathbf{A} \mathbf{Q} \mathbf{v}_i - \lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q}) \mathbf{Q} \mathbf{v}_i\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{c_1}}$ for some constant $c_1 > 0$. Let $\mathbf{P} = \mathbf{I} - \mathbf{Z}_S \mathbf{Z}_S^T$ where $\mathbf{Z}_S = \mathbf{Q} \mathbf{V}_S$ is the output of Algorithm 2. Then, for some constants $c_2, c' > 0$, we have:*

$$\|\mathbf{PAP}\|_2 \leq 2\sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2}},$$

with probability at least $1 - \frac{1}{n^{c'}}$.

Proof. Observe that if $\mathbf{v}$ is an eigenvector of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ (where $\mathbf{Q}$ is an orthonormal basis of the Krylov subspace $\mathcal{K}_q$ in line 3 of Algorithm 2) with corresponding eigenvalue λ , then $\mathbf{z} = \mathbf{Q} \mathbf{v}$ is an eigenvector of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$ with eigenvalue λ as $\mathbf{Q}$ is an orthonormal matrix. Similarly, for any eigenvector $\mathbf{z}$ of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$, there is a corresponding eigenvector $\mathbf{Q}^T \mathbf{z}$ of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ with the same eigenvalue. Thus, we can interchangeably refer to the eigenvalues and eigenvectors of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$ instead of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ in our proof. Let $\alpha = \max\left(\sigma_{l+1}(\mathbf{A}), \frac{\|\mathbf{A}\|_2}{n^{c/2}}\right)$ for some constant c as defined in Lemma 3. We will now prove the Lemma by considering the two cases below:

Case 1: Let there be some $k \in [l]$ such that $\sigma_k(\mathbf{A}) \geq 2\alpha$ and $\sigma_{k+1}(\mathbf{A}) < 2\alpha$. Let $p \in [k, l]$ such that $\sigma_p(\mathbf{A}) \geq \frac{3}{2}\alpha$, $\sigma_p(\mathbf{A}) - \sigma_{p+1}(\mathbf{A}) \geq \frac{\|\mathbf{A}\|_2}{2n^{c/2+1}}$ as defined in Lemma 3. Then, from Theorem 5, for any $i \in [p]$, we have $\|\mathbf{A} \mathbf{z}_i - \lambda_i(\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T) \mathbf{z}_i\|_2 \leq \|\mathbf{A} \mathbf{Z}_p - \mathbf{Z}_p \tilde{\mathbf{\Lambda}}_p\|_F \leq \frac{\|\mathbf{A}\|_2}{n^{c_3}}$. Thus, as $\beta = c_3/2$ in Algorithm 2 we have $\frac{\|\mathbf{A}\|_2}{n^{c_3}} < \frac{\|\mathbf{A}\|_2}{n^\beta}$ and the convergence condition in line 7 of Algorithm 2 is always satisfied for the top p eigenvectors of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$. So, we must have $[p] \subseteq S$, and $\mathbf{Z}_S$ must contain at least the top p eigenvectors of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$. Thus, observe that using Lemma 6, for some constant c_4 , we have $\|\mathbf{P} \mathbf{A} \mathbf{P}\|_2 \leq \|(\mathbf{I} - \mathbf{Z}_p \mathbf{Z}_p^T) \mathbf{A} (\mathbf{I} - \mathbf{Z}_p \mathbf{Z}_p^T)\|_2 \leq \sigma_{p+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_4}} \leq \sigma_{k+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_4}} \leq 2\alpha + \frac{\|\mathbf{A}\|_2}{n^{c_4}} \leq 2(\sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c/2}}) + \frac{\|\mathbf{A}\|_2}{n^{c_4}} \leq 2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2}{n^{c/2}} + \frac{\|\mathbf{A}\|_2}{n^{c_4}}$. So, for some constant $c_5 > 0$ we get:

$$\|\mathbf{P} \mathbf{A} \mathbf{P}\|_2 \leq 2\sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_5}}.$$

Case 2: If no such $k \in [l]$ exists, we will have $\sigma_i(\mathbf{A}) \leq 2\alpha$ for all $i \in [n]$. So, as $\mathbf{P}$ is a projection matrix, we have $\|\mathbf{P} \mathbf{A} \mathbf{P}\|_2 \leq \|\mathbf{A}\|_2 \leq 2\alpha \leq 2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2}{n^{c/2}}$. This completes the proof after choosing the constant carefully. $\square$

3.2 Error Bounds for moment matching with deflation

In this section, we prove the final error bounds for Algorithm 1. We first prove the existence of a matrix $\mathbf{B}$ that is close in spectral norm to $\mathbf{A}$ such that the converged eigenvectors and eigenvalues of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$ as defined in line 7 of Algorithm 2 are a subset of $\mathbf{B}$'s true eigenvectors and eigenvalues. First, we state a simplified version of a backward perturbation bound from [Sun95].

Theorem 9 (Theorem 3.1 of [Sun95]). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Let $\tilde{\mathbf{X}} \in \mathbb{R}^{n \times l}$ and $\tilde{\mathbf{\Lambda}} \in \mathbb{R}^{n \times l}$ (where $\tilde{\mathbf{\Lambda}}$ is diagonal) be such that $\|\mathbf{A} \tilde{\mathbf{X}} - \tilde{\mathbf{X}} \tilde{\mathbf{\Lambda}}\|_2 \leq \Delta$ for some $\Delta > 0$. Let $\hat{\mathbf{R}} = \mathbf{A} \tilde{\mathbf{X}} - \tilde{\mathbf{X}} \tilde{\mathbf{\Lambda}}$. Let $\tilde{\mathbf{X}} = \mathbf{P} \mathbf{H}$ be the polar decomposition of $\tilde{\mathbf{X}}$. Then, there exists a symmetric matrix $\mathbf{H} \in \mathbb{R}^{n \times n}$ such that $(\mathbf{A} + \mathbf{H}) \mathbf{P} = \mathbf{P} \tilde{\mathbf{\Lambda}}$ and $\|\mathbf{H}\|_F \leq \frac{\sqrt{\|\hat{\mathbf{R}}\|_F^2 + \|\mathbf{P}^\perp \hat{\mathbf{R}}\|_F^2}}{\sigma_{\min}(\tilde{\mathbf{X}})}$ where $\mathbf{P}^\perp = \mathbf{I} - \mathbf{P} \mathbf{P}^T$.*

Using Theorem 9, we state a backward error bound for the deflation algorithm which we will use in the final error analysis.

Lemma 7 (Wasserstein Error using Backward Error Bound). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Let $\mathbf{Z}_S$ and $\tilde{\mathbf{\Lambda}}_S$ be the outputs of Algorithm 2 with $\mathbf{A}$ as input. Then there exists a symmetric matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$ such that*

$$\mathbf{B} \mathbf{Z}_S = \mathbf{Z}_S \tilde{\mathbf{\Lambda}}_S \text{ and } \|\mathbf{A} - \mathbf{B}\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}},$$

where $\beta > 0$ is the constant defined in Algorithm 2.

Proof. We will apply the backward perturbation error bound of Theorem 9 to prove the existence of $\mathbf{B}$. Following Theorem 9, we have $\tilde{\mathbf{X}}_1 = \mathbf{Z}_S$ and hence, $\mathbf{P}_1 = \mathbf{Z}_S$ since $\mathbf{Z}_S$ is an orthonormal matrix and its polar decomposition is equal to itself. Next, observe that $\sigma_{\min}(\mathbf{Z}_S) = 1$ and $\|\hat{\mathbf{R}}\|_F^2 = \|\mathbf{A}\mathbf{Z}_S - \mathbf{Z}_S\tilde{\mathbf{A}}_S\|_F^2 = \sum_{i \in S} \|\mathbf{A}\mathbf{z}_i - \lambda_i(\mathbf{Q}^T\mathbf{A}\mathbf{Q})\mathbf{z}_i\|_2^2 \leq \frac{\|\mathbf{A}\|_2^2}{n^{2\beta-1}}$ where the last bound follows from the fact that $\|\mathbf{A}\mathbf{z}_i - \lambda_i(\mathbf{Q}^T\mathbf{A}\mathbf{Q})\mathbf{z}_i\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^\beta}$ as stated in line 7 of Algorithm 2. We also have $\|(\mathbf{Z}_S)_\perp(\mathbf{Z}_S)_\perp^T\hat{\mathbf{R}}\|_F^2 = \|(\mathbf{I} - \mathbf{Z}_S\mathbf{Z}_S^T)(\mathbf{A}\mathbf{Z}_S - \mathbf{Z}_S\tilde{\mathbf{A}}_S)\|_F^2 \leq \|\mathbf{A}\mathbf{Z}_S - \mathbf{Z}_S\tilde{\mathbf{A}}_S\|_F^2 \leq \frac{\|\mathbf{A}\|_2^2}{n^{2\beta-1}}$ where the last step follows as $\mathbf{I} - \mathbf{Z}_S\mathbf{Z}_S^T$ is a projection matrix. Thus, there exists a symmetric matrix $\mathbf{H}$ such that $(\mathbf{A} + \mathbf{H})\mathbf{Z}_S = \mathbf{Z}_S\tilde{\mathbf{A}}_S$ and $\|\mathbf{H}\|_F \leq \frac{\sqrt{\|\hat{\mathbf{R}}\|_F^2 + \|(\mathbf{Z}_S)_\perp\hat{\mathbf{R}}\|_F^2}}{\sigma_{\min}(\mathbf{Z})} \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}$. Setting $\mathbf{B} = \mathbf{A} + \mathbf{H}$ gives us the required matrix. $\square$

We now prove that the output of Algorithm 1 of [BKM22] using the modified moments $\hat{\tau}_i$ as inputs for $i \in N$ (as defined in lines 4-6 of Algorithm 1) must be close in the Wasserstein distance to the spectral density defined by the non-zero eigenvalues of the deflated matrix $\mathbf{PAP}$.

Lemma 8 (Modified SDE bound for deflated matrix). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a rank $n - r$ symmetric matrix for some $r \in [n]$ with spectral density $s_{\mathbf{A}}(x)$ such that $|\lambda_1(\mathbf{A})| \leq 1$. Let $\epsilon \in (0, 1)$, $N = O(\frac{1}{\epsilon})$ and $\tilde{\tau}_1, \dots, \tilde{\tau}_N$ be the estimates of the top N normalized Chebyshev moments of $\mathbf{A}$ estimated using Algorithm 2 of [BKM22]. Define $\hat{\tau}_i = \frac{1}{n-r}(n\tilde{\tau}_i - r\bar{T}_i(0))$ for $i \in [N]$ where $\bar{T}_i(0)$ is the i^{th} normalized Chebyshev polynomial. Let the density function $q(x)$ be the output of Algorithm 1 of [BKM22] with $\hat{\tau}_1, \dots, \hat{\tau}_N$ as the inputs for the moment estimates. Let $s'_{\mathbf{A}}(x) = \frac{1}{n-r} \sum_{j=1}^{n-r} \delta(x - \lambda_j(\mathbf{A}))$ be the probability density defined by the top $n - r$ eigenvalues of $\mathbf{A}$, $\lambda_1(\mathbf{A}), \dots, \lambda_{n-r}(\mathbf{A})$. Then, with probability at least $1 - \delta$, we have:*

$$W_1(s'_{\mathbf{A}}, q) \leq \frac{n\epsilon}{n-r}.$$

Also, estimating the density function q using Algorithm 1 of [BKM22] and also using Algorithm 2 of [BKM22] as a subroutine to estimate the moments, requires $O(\frac{b}{\epsilon})$ matrix-vector products where $b = \max(1, \frac{C'}{n\epsilon^2} \log^2 \frac{1}{\epsilon\delta} \log^2 \frac{1}{\epsilon})$.

Proof. $\frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{A})) = \frac{1}{n} \sum_{j=1}^n \bar{T}_i(\lambda_j(\mathbf{A})) = \frac{1}{n} \sum_{j=1}^{n-r} \bar{T}_i(\lambda_j(\mathbf{A})) + \frac{r}{n} \bar{T}_i(0)$ for $i \in [N]$ are the top N normalized Chebyshev moments of $s_{\mathbf{A}}(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - \lambda_i(\mathbf{A}))$. Then, the i^{th} normalized Chebyshev moment of $s'_{\mathbf{A}}(x)$ is given by $\tau'_i = \frac{1}{n-r} \sum_{j=1}^{n-r} \bar{T}_i(\lambda_j(\mathbf{A})) = \frac{n}{n-r} \cdot \frac{1}{n} \sum_{j=1}^n \bar{T}_i(\lambda_j(\mathbf{A})) - \frac{1}{n-r} \sum_{j=n-r+1}^n \bar{T}_i(\lambda_j(\mathbf{A})) = \frac{n}{n-r} \cdot \frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{A})) - \frac{r}{n-r} \bar{T}_i(0)$.

By setting $\Delta = \frac{1}{N \ln(eN)}$ (where N is the number of Chebyshev moments estimated), using Lemma 4.2 of [BKM22], we get that with $\mathbf{A}$ as the input, the normalized Chebyshev moment estimates $\hat{\tau}_i$ returned by Algorithm 2 of [BKM22] must satisfy $|\hat{\tau}_i - \frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{A}))| \leq \frac{1}{N \ln(eN)}$ for all $i \in [N]$. So, we have:

$$\begin{aligned} |\hat{\tau}_i - \tau'_i| &= \left| \frac{n}{n-r} \tilde{\tau}_i - \frac{r}{n-r} \bar{T}_i(0) - \frac{1}{n-r} \sum_{j=1}^{n-r} \bar{T}_i(\lambda_j(\mathbf{A})) \right| \\ &= \frac{n}{n-r} \left| \tilde{\tau}_i - \frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{A})) \right| \\ &\leq \frac{n}{(n-r)N \ln(eN)}, \end{aligned} \tag{29}$$

where the second step follows from the fact that $\lambda_{n-r+1}(\mathbf{A}) = \dots = \lambda_n(\mathbf{A}) = 0$ by assumption. The density q returned by Algorithm 1 of [BKM22] is defined on a $(d+1)$ -length evenly spaced grid $X_d = [-1, -1 + \frac{2}{d}, \dots, 1 - \frac{2}{d}, 1]$ for $d = \lceil N^3/2 \rceil$. Let $z_{s'} = [\tau'_1/1, \tau'_2/2, \dots, \tau'_N/N]$ and $z_q = \mathcal{T}_N^{d+1} q$ where $(\mathcal{T}_N^{d+1}) \in \mathbb{R}^{N \times (d+1)}$ such that $(\mathcal{T}_N^{d+1})_{ij} = \bar{T}_i(-1 + \frac{2j}{d})/i$ for $i \in [N]$, $j \in \{0, 1, \dots, d\}$. Also, let $z = [\hat{\tau}_1, \hat{\tau}_2/2, \dots, \hat{\tau}_N/N]$. Then, using triangle inequality, we have

$$\|z_q - z_{s'}\|_1 = \|z_q - z\|_1 + \|z - z_{s'}\|_1.$$

Now, $\|z - z_{s'}\|_1 \leq \sum_{i=1}^N (\hat{\tau}_i - \tau'_i)/i \leq \frac{n}{(n-r)N \ln(eN)} \cdot H_n \leq \frac{n}{N(n-r)}$ where we use (29) to bound $\hat{\tau}_i - \tau'_i$ and H_n is the n^{th} harmonic number. Next, consider the following distribution q^* on X_d as defined in Lemma 3.4 of [BKM22]:

$$q^*(x) = \frac{1}{n-r} \sum_{j=1}^{n-r} \delta(x - \arg \min_{p \in X_d} |p - \lambda_j(\mathbf{A})|).$$

q^* is the distribution corresponding to moving the mass from each $\lambda_j(\mathbf{A})$ to its nearest point on the grid X_d . We have $W_1(s'_{\mathbf{A}}, q^*) \leq \frac{1}{d}$ due to the earth mover's distance interpretation of the Wasserstein-1 distance. Let $z_{q^*} = [\langle q^*, \bar{T}_1 \rangle, \dots, \langle q^*, \bar{T}_N \rangle / N]$. Now, from Line 3 of Algorithm 1 of [BKM22], q is defined as the density which minimizes $\|\mathcal{T}_N^{d+1} q - z\|_1$. Thus, we have $\|\mathcal{T}_N^{d+1} q - z\|_1 \leq \|z_{q^*} - z\|_1$. Then, following the proof of Lemma 3.4 of [BKM22] exactly, we can upper bound $\|z_{q^*} - z\|_1$ by $\frac{2}{N}$ using the properties of Chebyshev polynomials which gives us $\|\mathcal{T}_N^{d+1} q - z\|_1 \leq \frac{2}{N}$. Thus, we get $\|z_q - z_{s'}\|_1 \leq \frac{3n}{(n-r)N}$. Finally, following the proof of Lemma 3.4 of [BKM22], Lemma 3.1 of [BKM22] gives us the final bound on $W_1(s'_{\mathbf{A}}, q)$.

Note that the number of matrix-vector products is obtained by setting $\Delta = \frac{1}{N \ln(eN)}$ in Lemma 4.2 of [BKM22]. $\square$

We now prove the final error bound for Algorithm 1. We will first show using Lemma 7 that there exists a matrix $\mathbf{B}$ such that its spectral density is very close to that of $\mathbf{A}$ and $\mathbf{Z}_S$ and $\tilde{\mathbf{A}}_S$ (the output of Algorithm 2) are a subset of the eigenvectors and eigenvalues of $\mathbf{B}$. Then, it is enough to estimate the spectral density of $\mathbf{B}$. We already know the eigenvalues $\tilde{\mathbf{A}}_S$ of $\mathbf{B}$. So we just need to estimate the spectral density of its remaining eigenvalues (which is equal to the part of the spectral density of the deflated matrix $\mathbf{PBP}$ corresponding to the eigenvalues with eigenvectors which lie in the subspace orthogonal to $\mathbf{Z}_S$). To do this, we first observe that the spectral density of $\mathbf{PBP}$ is again close to that of $\mathbf{PAP}$. So, it is enough to estimate the spectral density of the eigenvalues of $\mathbf{PAP}$ corresponding to its non-deflated eigenvectors. We show this is exactly estimated by Algorithm 1 of [BKM22] by appropriately modifying the Chebyshev moments of $\mathbf{PAP}$ to account for the zero eigenvalues corresponding to the deflated eigenvectors.

Theorem 1 (SDE with Explicit Deflation). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix. For any $\epsilon, \delta \in (0, 1)$, $l \in [n]$, and constants $c, c_1 > 0$, Algorithm 1 performs $O(l \log n + \frac{b}{\epsilon})$ matrix-vector products with $\mathbf{A}$ where $b = \max(1, \frac{1}{n\epsilon^2} \log^2 \frac{n}{\epsilon} \log^2 \frac{1}{\epsilon})$ and computes a probability density function q such that*

$$W_1(s_{\mathbf{A}}, q) \leq \epsilon \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^c},$$

with probability at least $1 - \frac{1}{n^{c_1}}$.

Proof. Let $\mathbf{Z}_S \in \mathbb{R}^{n \times |S|}$ and $\tilde{\mathbf{\Lambda}}_S \in \mathbb{R}^{|S| \times |S|}$ be the output of Algorithm 2, in line 2 of Algorithm 1. Recall that $\mathbf{Z}_S = \mathbf{Q}\mathbf{V}_S$ where $\mathbf{Q}$ is an orthonormal basis of the Krylov subspace of $\mathbf{A}$ (Definition 2.1) with a starting block $\mathbf{X} \in \mathbb{R}^{n \times l}$ containing random $\mathcal{N}(0, 1)$ Gaussian entries and the columns of $\mathbf{V}_S$ are all the eigenvectors of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ with the corresponding eigenvalues along the diagonal of $\tilde{\mathbf{\Lambda}}_S \in \mathbb{R}^{|S| \times |S|}$ such that for any $i \in S$, $\|\mathbf{A} \mathbf{Q} \mathbf{v}_i - (\tilde{\mathbf{\Lambda}}_S)_{ii} \mathbf{Q} \mathbf{v}_i\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^\beta}$ for some constant β as defined in line 7 of Algorithm 2 (here $\mathbf{v}_i$ is an eigenvector of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ corresponding to eigenvalue $\lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$). Equivalently, for every $i \in S$, $\mathbf{z}_i = \mathbf{Q} \mathbf{v}_i$ is an eigenvector of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$ corresponding to an eigenvalue $\lambda_i(\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T)$. Now, using Lemma 7, there exists a symmetric matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$ such that $\mathbf{B} \mathbf{Z}_S = \mathbf{Z}_S \tilde{\mathbf{\Lambda}}_S$ and $\|\mathbf{A} - \mathbf{B}\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}$. Using Weyls' inequality (Fact 1), for all $i \in [n]$ we have

$$|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})| \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}. \quad (30)$$

Then, $\frac{1}{n} \sum_{i=1}^n |\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{B})| \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}$. Using the earth mover's interpretation, this implies that:

$$W_1(s_{\mathbf{A}}, s_{\mathbf{B}}) \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}, \quad (31)$$

where $s_{\mathbf{B}}$ is the spectral density of $\mathbf{B}$. Since $\beta = c$ as defined in line 2 of Algorithm 1, it is enough to bound $W_1(s_{\mathbf{B}}, q)$.

Since $\mathbf{B} \mathbf{Z}_S = \mathbf{Z}_S \tilde{\mathbf{\Lambda}}_S$, the eigenvalues $\lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$ such that $i \in S$ are also eigenvalues of $\mathbf{B}$. Let $S_1 \subseteq [n]$ be the set of indices of the eigenvalues of $\mathbf{B}$ such that $|S_1| = |S|$ and for each $i \in S_1$, there exists some $j \in S$ such that $\lambda_i(\mathbf{B}) = \lambda_j(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$. Let $[-L_1, L_1]$ contain the support of both $s_{\mathbf{B}}$ and q (the output of Algorithm 1). Let $\mathbf{P} = \mathbf{I} - \mathbf{Z}_S \mathbf{Z}_S^T$. Note that, from Algorithm 1, we also have $q(x) = \frac{|S|q_1(x) + (n-|S|)q_2(x)}{n}$ where $q_1(x) = \frac{1}{|S|} \sum_{i=1}^{|S|} \delta(x - (\tilde{\mathbf{\Lambda}}_S)_{ii}) = \frac{1}{|S|} \sum_{i \in S} \delta(x - \lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q}))$ (as defined in Line 2 of Algorithm 1) and $q_2(x) = q'_2(x/L)$ for $x \in [-L, L]$ where q'_2 is a density supported on $[-1, 1]$ which is the output of Algorithm 1 of [BKM22] with the modified approximate moments of $\frac{1}{L} \mathbf{P} \mathbf{A} \mathbf{P}$ as defined in lines 4-6 of Algorithm 1 (where $\|\mathbf{P} \mathbf{A} \mathbf{P}\|_2 \leq L \leq 2\|\mathbf{P} \mathbf{A} \mathbf{P}\|_2$). $W_1(s_{\mathbf{B}}, q)$ can be written as:

$$W_1(s_{\mathbf{B}}, q) = \sup_{h \in 1\text{-Lip}} \int_{-L_1}^{L_1} h(x) (s_{\mathbf{B}}(x) - q(x)) dx.$$

Let $h^*(x)$ be the 1-Lipschitz function which maximizes the integral above. Then, using the definitions of $s_{\mathbf{B}}$ and q we have:

$$\begin{aligned} W_1(s_{\mathbf{B}}, q) &= \int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{n} \sum_{i=1}^n \delta(x - \lambda_i(\mathbf{B})) - \frac{1}{n} \sum_{j \in S} \delta(x - \lambda_j(\mathbf{Q}^T \mathbf{A} \mathbf{Q})) - \frac{n-|S|}{n} q_2(x) \right) dx \\ &= \underbrace{\frac{|S|}{n} \int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{|S|} \sum_{i \in S_1} \delta(x - \lambda_i(\mathbf{B})) - \frac{1}{|S|} \sum_{j \in S} \delta(x - \lambda_j(\mathbf{Q}^T \mathbf{A} \mathbf{Q})) \right) dx}_{I_1} \\ &\quad + \underbrace{\frac{n-|S|}{n} \int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{n-|S|} \sum_{i \notin S_1} \delta(x - \lambda_i(\mathbf{B})) - q_2(x) \right) dx}_{I_2}. \end{aligned} \quad (32)$$

We have $I_1 = 0$ by definition of the sets S_1 and S . Thus, we have $W_1(s_{\mathbf{B}}, q) = \frac{n-|S|}{n} I_2$. We now bound I_2 .

Observe that any eigenvalue $\lambda_i(\mathbf{B})$ such that $i \notin S_1$, is also an eigenvalue value of $\mathbf{PBP}$. To see this, let $\mathbf{x}_i$ be the corresponding eigenvector of $\mathbf{B}$. Then, observe that $\mathbf{PBP}\mathbf{x} = \mathbf{PB}\mathbf{x} = \lambda_i(\mathbf{B})\mathbf{x}$ as $\mathbf{Z}_S^T \mathbf{x} = 0$. Also both $\mathbf{PBP}$ and $\mathbf{PAP}$ have $|S|$ (where $|S| = |S_1|$) eigenvalues equal to 0 (with the corresponding eigenvectors being the columns of $\mathbf{Z}$). Let $S_2 \subseteq [n]$ be the set of indices such that $|S_2| = n - |S_1| = n - |S|$ and for every $j \in S_2$, there exists some $i \in [n] \setminus S_1$ such that $\lambda_j(\mathbf{PBP}) = \lambda_i(\mathbf{B})$. So, we can write I_2 as:

$$\begin{aligned}
I_2 &= \int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{n-|S|} \sum_{i \notin S_1} \delta(x - \lambda_i(\mathbf{B})) - q_2(x) \right) dx \\
&= \int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta(x - \lambda_i(\mathbf{PBP})) - q_2(x) \right) dx \\
&= \underbrace{\int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta(x - \lambda_i(\mathbf{PBP})) - \frac{1}{n-|S|} \sum_{i \in S_2} \delta(x - \lambda_i(\mathbf{PAP})) \right) dx}_{t_1} \\
&\quad + \underbrace{\int_{-L_1}^{L_1} h^*(x) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta(x - \lambda_i(\mathbf{PAP})) - q_2(x) \right) dx}_{t_2}. \tag{33}
\end{aligned}$$

Using the fact that $\mathbf{P}$ is a projection matrix and Lemma 6, we again have that $\|\mathbf{PAP} - \mathbf{PBP}\|_2 \leq \|\mathbf{P}(\mathbf{A} - \mathbf{B})\mathbf{P}\|_2 \leq \|\mathbf{A} - \mathbf{B}\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}$. Thus, using Weyl's inequality (Fact 1) we have $|\lambda_i(\mathbf{PBP}) - \lambda_i(\mathbf{PAP})| \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}$ for $i \in [n]$. We can bound t_1 using the earth movers' interpretation of the Wasserstein-1 distance as

$$t_1 \leq \frac{1}{n-|S|} \sum_{i \in S_2} |\lambda_i(\mathbf{PBP}) - \lambda_i(\mathbf{PAP})| \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}. \tag{34}$$

We bound t_2 next. Since all eigenvalues of $\frac{1}{L}\mathbf{PAP}$ are in $[-1, 1]$ and $\mathbf{PAP}$ has at least $n - |S|$ eigenvalues equal to 0 as described previously (corresponding to $\lambda_i(\mathbf{PAP})$ where $i \notin S_2$), according to Lemma 8, the density q'_2 returned by Algorithm 1 of [BKM22] in line 6 of Algorithm 1 satisfies the guarantee

$$W_1(s'_{\mathbf{PAP}}, q'_2) \leq \frac{n\epsilon}{(n-|S|)}, \tag{35}$$

with probability at least $1 - \delta$ where $s'_{\mathbf{PAP}}(x) = \sum_{i \in S_2} \frac{\delta(x - \frac{1}{L}\lambda_i(\mathbf{PAP}))}{n-|S|}$. Since q'_2 is supported on $[-1, 1]$, we set $q_2(x) = q'_2(x/L)$ in line 7 of Algorithm 1 when $x \in [-L, L]$ and $q_2(x) = 0$ otherwise

so that q_2 is now supported on $[-L, L]$. So we get that

$$\begin{aligned}
t_2 &= \int_{-L}^L h^*(x) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta(x - \lambda_i(\mathbf{PAP})) - q_2(x) \right) dx \\
&= \int_{-1}^1 h^*(Lx) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta(Lx - \lambda_i(\mathbf{PAP})) - q_2(Lx) \right) dx \\
&= \int_{-1}^1 h^*(Lx) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta\left(x - \frac{\lambda_i(\mathbf{PAP})}{L}\right) - q'_2(x) \right) dx \\
&\leq \sup_{h_1 \in 1\text{-Lip}} L \int_{-1}^1 h_1(x) \left(\frac{1}{n-|S|} \sum_{i \in S_2} \delta\left(x - \frac{\lambda_i(\mathbf{PAP})}{L}\right) - q'_2(x) \right) dx \\
&= LW_1\left(s'_{\mathbf{PAP}}, q'_2\right) \\
&\leq \frac{n\epsilon}{(n-|S|)} L.
\end{aligned}$$

In the first step above, we use the fact that the densities are supported on $[-L, L]$ instead of $[-L_1, L_1]$. In the second step, we rescale the integrals from $[-L, L]$ to $[-1, 1]$. In the third step, we use the fact that $q_2(Lx) = q'_2(x)$ and $\frac{1}{n-|S|} \sum_{i \in S_2} \delta(Lx - \lambda_i(\mathbf{PAP})) = \frac{1}{n-|S|} \sum_{i \in S_2} \delta\left(x - \frac{\lambda_i(\mathbf{PAP})}{L}\right)$ when $x \in [-1, 1]$. The fourth step follows from the fact that $h^*(Lx)$ is an L -Lipschitz function and the final two steps follow from the definition of the Wasserstein-1 distance and (35).

Since $\|\mathbf{PAP}\|_2 \leq L \leq 2\|\mathbf{PAP}\|_2$ from Theorem 8, for some constant $c_2 > 0$, we have $L \leq 2\|\mathbf{PAP}\|_2 \leq 2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2}{n^{c_2}}$. Thus, we get:

$$t_2 \leq \frac{\epsilon n}{2(n-|S|)} \cdot \left(2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2}{n^{c_2}} \right). \quad (36)$$

From (33), using the bounds on t_1 and t_2 from (34) and (36), we get $I_2 \leq \frac{\|\mathbf{A}\|_2}{n^{\beta-1}} + \frac{\epsilon n}{2(n-|S|)} \cdot \left(2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2}{n^{c_2}} \right)$. Using the bound on I_2 and the fact that $I_1 = 0$ in (32), we get:

$$W_1(s_{\mathbf{B}}, q) \leq \frac{|S|}{n} I_1 + \frac{n-|S|}{n} I_2 \leq \epsilon \sigma_{l+1}(\mathbf{A}) + \frac{\epsilon \|\mathbf{A}\|_2}{n^{c_2}} + \frac{\|\mathbf{A}\|_2}{n^{\beta-1}}.$$

Finally, using triangle inequality and (31), we get that, for some suitable chosen constant $c_1 > 0$:

$$\begin{aligned}
W_1(s, q) &\leq W_1(s_{\mathbf{A}}, s_{\mathbf{B}}) + W_1(s_{\mathbf{B}}, q) \leq \epsilon \sigma_{l+1}(\mathbf{A}) + \frac{\epsilon \|\mathbf{A}\|_2}{n^{c_2}} + \frac{2\|\mathbf{A}\|_2}{n^{\beta-1}} \\
&\leq \epsilon \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_1}}.
\end{aligned}$$

This completes the proof.

Matrix vector products. Line 1 of Algorithm 1 calls Algorithm 2 with $\mathbf{A}$ as input which uses $O(l \log n)$ matrix vector products with $\mathbf{A}$ to form the Krylov subspace in Line 2 of Algorithm 2 (since the Krylov subspace has depth $O(\log n)$ and block size l). By setting the error parameter ϵ in Algorithm 2 of [MM15] to 0.5, we get an output L' in $O(\log n)$ iterations such that $0.5\|\mathbf{PAP}\|_2 \leq L' \leq \|\mathbf{PAP}\|_2$. Then, we can set $L = 2L'$ such that $\|\mathbf{PAP}\|_2 \leq L \leq 2\|\mathbf{PAP}\|_2$ as defined in Line 3 of Algorithm 1. Also, Algorithm 2 of [BKM22] is called with the deflated and scaled matrix $\frac{1}{L}\mathbf{PAP}$ as input in line 4 of Algorithm 1. According to Lemma 8, this uses $O\left(\frac{b}{\epsilon}\right)$ matrix vector products with $\mathbf{PAP}$ where $b = \max\left(1, \frac{1}{n\epsilon^2} \log^2 \frac{n}{\epsilon} \log^2 \frac{1}{\epsilon}\right)$ i.e. $O\left(\frac{b}{\epsilon}\right)$ matrix vector products with $\mathbf{A}$, since $\delta \geq \frac{1}{n^{c'}}$ (for some $c' \geq 0$). Finally, Line 5 of Algorithm 1 calls Algorithm 2 of [BKM22] which uses at most $O(\log n)$ extra matrix vector products. Thus, the total number of matrix vector products is $O(l \log n + \frac{b}{\epsilon})$ where $b = \max(1, \frac{C'}{n\epsilon^2} \log^2 \frac{n}{\epsilon} \log^2 \frac{1}{\epsilon})$. $\square$

We now state a simple corollary which shows that we can get Wasserstein-1 error depending on the Schatten 1-norm of $\mathbf{A}$ by appropriately balancing the errors from deflation and moment matching in Algorithm 1.

Corollary 1. *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric. For any $\epsilon \in (0, 1)$ and some constant $c > 0$, there exists an algorithm that performs $O\left(\frac{\sqrt{n} \log n}{\epsilon} + \sqrt{n} \log^4 n\right)$ matrix vector products with $\mathbf{A}$ and computes M such that, with probability at least $1 - \frac{1}{n^c}$, $|M - \|\mathbf{A}\|_1| \leq \epsilon \|\mathbf{A}\|_1$.*

Proof. Set the block size as $l = \frac{\sqrt{n}}{\epsilon}$ and the error parameter ϵ' to $O(\frac{1}{\sqrt{n}})$ of Algorithm 1. Let q be the output of Algorithm 1 using $O(l \log n + \frac{b}{\epsilon'}) = O\left(\frac{\sqrt{n}}{\epsilon} \log n + \sqrt{n} \log^4 n\right)$ matrix vector products with $\mathbf{A}$. Then, from Theorem 1, we have $W_1(\mathbf{s}_A, q) \leq \frac{\sigma_{l+1}(\mathbf{A})}{\sqrt{n}} + \frac{\|\mathbf{A}\|_2}{n^{c/4}} \leq \frac{\|\mathbf{A}\|_1}{l\sqrt{n}} + \frac{\|\mathbf{A}\|_2}{n^{c/4}} = \frac{\epsilon \|\mathbf{A}\|_1}{n} + \frac{\|\mathbf{A}\|_2}{n^{c/4}} \leq \frac{2\epsilon \|\mathbf{A}\|_1}{n}$. Then, using q , we can construct a list of n values $[\tilde{\lambda}_1, \dots, \tilde{\lambda}_n]$ in time linear in n and $\frac{1}{\epsilon}$ such that $\sum_{i=1}^n |\lambda_i - \tilde{\lambda}_i| \leq \frac{2\epsilon \|\mathbf{A}\|_1 \cdot n}{n} \leq 2\epsilon \|\mathbf{A}\|_1$ (see [CSKSV18], theorem B.1 in [BKM22]). Adjusting ϵ by constant factors gives us the final bound. $\square$

4 Analysis of Stochastic Lanczos Quadrature

In this section, we give our error analysis of Stochastic Lanczos Quadrature (SLQ) (Algorithm 4) by showing that it implicitly performs a deflation of the input matrix. Our analysis shows that for a symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, SLQ achieves an error bound of $\epsilon \sigma_{l+1}(\mathbf{A}) + \tilde{O}(\frac{l \|\mathbf{A}\|_2}{n})$ using $O(l \log \frac{1}{g_{\min}} + \frac{1}{\epsilon} \log \frac{n \cdot \kappa}{\delta})$ matrix vector products with $\mathbf{A}$ for any $l \in [n]$, $\epsilon = \tilde{\Omega}(1/\sqrt{n})$, failure probability $\delta \in (0, 1)$ and where $g_{\min}$ (minimum singular value gap) and κ (condition number) are as stated in Theorem 3. Hence, it almost matches the error bounds of the explicit deflation and moment matching algorithm (Algorithm 1) we described in the previous section (up to the additive $\frac{l \|\mathbf{A}\|_2}{n}$ factor). Roughly, we show that the large magnitude eigenvalues of $\mathbf{A}$ are estimated almost exactly in $O(l \log \frac{1}{g_{\min}} + \frac{1}{\epsilon} \log \frac{n \cdot \kappa}{\delta})$ iterations of Lanczos (Algorithm 3). Estimating the spectral density of the small magnitude eigenvalues requires a further $\tilde{O}(1/\epsilon)$ iterations. Then, we give a simple variant of SLQ, which we call the Variance Reduced SLQ (VRSLQ) which sets the weights of the converged eigenvalues in the final distribution correctly so that we end up getting an error bound of $\epsilon \sigma_{l+1}(\mathbf{A}) + \tilde{O}(\frac{l \sigma_{l+1}(\mathbf{A})}{n})$. We note that the SLQ algorithm described in this paper (Algorithm 4) uses only one random starting vector for simplicity, though in practice, we can get better concentration when the resulting distribution is averaged over multiple random starting vectors.

This section is organized as follows. In Section 4.1, we first derive a loose error bound of $\epsilon\|\mathbf{A}\|_2$ for SLQ using a simple moment matching based analysis. In Section 4.2, we derive error bounds for approximating the top eigenvalue and eigenvector of a matrix using the Lanczos algorithm. In Section 4.3, we give a tighter error bound of $\epsilon\sigma_{l+1}(\mathbf{A}) + \tilde{O}(\frac{\|\mathbf{A}\|_2}{n})$ for SLQ by showing that it implicitly performs deflation and approximates the top eigenvalues (by using the error bounds developed in Section 4.2) followed by moment matching. In Section 4.4 we describe the variance reduced version of SLQ and then give the error bounds.

Notations: Throughout this section, $\mathcal{U}(\mathcal{S}^{n-1})$ denotes a uniform distribution on the unit sphere of dimension n . Also, for two random variables X and Y , $X \stackrel{d}{=} Y$ implies they have the same distribution.

4.1 SLQ bounds via Moment Matching

In this section, we derive an error bound for SLQ (Algorithm 4) using a simple moment matching based argument. We begin by describing the Lanczos algorithm (Algorithm 3) on which SLQ is based. The Lanczos algorithm iteratively constructs an orthonormal basis of the Krylov subspace $[\mathbf{g}, \mathbf{A}\mathbf{g}, \dots, \mathbf{A}^{m-1}\mathbf{g}]$ generated by an input matrix $\mathbf{A}$ and a random starting vector $\mathbf{g}$ of appropriate dimensions such that $\mathbf{T} = \mathbf{Q}^T \mathbf{A} \mathbf{Q}$ where $\mathbf{T}$ is tridiagonal and where $\mathbf{Q}$ is an orthonormal basis of the Krylov subspace that is computed by the Lanczos algorithm.

Algorithm 3 Lanczos algorithm ([Lan52, GM09, CTU21])

Input: Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, starting vector $\mathbf{g} \in \mathbb{R}^n$, number of iterations m .

- 1: Set $\mathbf{q}_1 = \mathbf{g}/\|\mathbf{g}\|_2$, $\alpha_1 = \mathbf{q}_1^T \mathbf{A} \mathbf{q}_1$, $\tilde{\mathbf{q}}_2 = \mathbf{A} \mathbf{q}_1 - \alpha_1 \mathbf{q}_1$ and initialize $\mathbf{T} \in \mathbb{R}^{m \times m}$ with $\mathbf{T}_{11} = \alpha_1$.
 - 2: **for** $i = 2, \dots, m$ **do**
 - 3: Let $\eta_{i-1} = \|\tilde{\mathbf{q}}_{i-1}\|_2$.
 - 4: Compute $\mathbf{q}_i = \tilde{\mathbf{q}}_{i-1}/\eta_{i-1}$ (i.e., normalize $\tilde{\mathbf{q}}_{i-1}$ to obtain Lanczos vectors $\mathbf{q}_i$).
 - 5: Compute $\alpha_i = \mathbf{q}_i^T \mathbf{A} \mathbf{q}_i$.
 - 6: Set $\tilde{\mathbf{q}}_{i+1} = \mathbf{A} \mathbf{q}_i - \alpha_i \mathbf{q}_i - \eta_{i-1} \mathbf{q}_{i-1}$.
 - 7: Set $\mathbf{T}_{ii} = \alpha_i$, and $\mathbf{T}_{i,i-1} = \mathbf{T}_{i-1,i} = \eta_{i-1}$.
 - 8: **end for**
 - 9: **return** $\mathbf{T}$, $\mathbf{Q}$ where $\mathbf{Q} \in \mathbb{R}^{n \times m}$ is a matrix whose i^{th} column is $\mathbf{q}_i$.
-

We have the following well known identity for the Lanczos algorithm (for eg. see [GM09]), which we prove here for completeness.

Lemma 9. ([GM09]) Consider Algorithm 3 run with input $\mathbf{A} \in \mathbb{R}^{n \times n}$, starting vector $\mathbf{g} \in \mathbb{R}^n$, and number of iterations m . Let $\mathbf{T} \in \mathbb{R}^{m \times m}$, $\mathbf{Q} \in \mathbb{R}^{n \times m}$ be the outputs of the algorithm. Then, for any $k \in [m-1]$, we have $\mathbf{A}^k \mathbf{g} = \mathbf{Q} \mathbf{T}^k \mathbf{Q}^T \mathbf{g}$.

Proof. The Krylov subspace generated by Algorithm 3 is $\mathcal{K}_m = [\mathbf{g}, \mathbf{A}\mathbf{g}, \mathbf{A}^2\mathbf{g}, \dots, \mathbf{A}^{m-1}\mathbf{g}]$ after m iterations. Since $\mathbf{Q}$ is an orthonormal basis of $\mathcal{K}_m$, the columns of $\mathcal{K}_m$ are spanned by the columns of $\mathbf{Q}$. Let $\mathbf{x} = \|\mathbf{g}\|_2 \mathbf{e}_1$ (recall $\mathbf{e}_1$ is the first standard basis vector with a 1 in the first position). So, we have $\mathbf{g} = \mathbf{Q} \mathbf{x}$.

We show that for any $p \in [m]$, $\mathbf{A}^p \mathbf{g} = \mathbf{Q} \mathbf{T}^p \mathbf{Q}^T \mathbf{g}$, via induction. Observe that for $p = 1$, $\mathbf{Q} \mathbf{T} \mathbf{Q}^T \mathbf{g} = \mathbf{Q} \mathbf{T} \mathbf{Q}^T \mathbf{Q} \mathbf{x} = \mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{x} = \mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{g} = \mathbf{A} \mathbf{g}$, where in the last equality we use the fact that $\mathbf{A} \mathbf{g}$ is a column of $\mathcal{K}_m$, and so it is spanned by the columns of $\mathbf{Q}$. This shows that the base case

for our induction is true. As the inductive hypothesis, assume $\mathbf{A}^k \mathbf{g} = \mathbf{Q} \mathbf{T}^k \mathbf{Q}^T \mathbf{g}$ for some $k \in [m-2]$. Then observe that $\mathbf{Q} \mathbf{T}^{k+1} \mathbf{Q}^T \mathbf{g} = \mathbf{Q} \mathbf{T} \mathbf{Q}^T \mathbf{Q} \mathbf{T}^k \mathbf{Q}^T \mathbf{g} = \mathbf{Q} \mathbf{T}^k \mathbf{Q}^T \mathbf{A}^k \mathbf{g} = \mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T \mathbf{A}^k \mathbf{g}$, where in the third equality we use our inductive hypothesis. Since $\mathbf{A}^k \mathbf{g}$ is a column of $\mathcal{K}_m$, $\mathbf{A}^k \mathbf{g}$ must be spanned by the columns of $\mathbf{Q}$. So we have, $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T \mathbf{A}^k \mathbf{g} = \mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{A}^k \mathbf{g} = \mathbf{Q} \mathbf{Q}^T \mathbf{A}^{k+1} \mathbf{g} = \mathbf{A}^{k+1} \mathbf{g}$, where in the last step we used the fact that $\mathbf{A}^{k+1} \mathbf{g}$ is also spanned by the columns of $\mathbf{Q}$. This completes our proof. $\square$

We now state the SLQ algorithm for spectral density estimation.

Algorithm 4 Stochastic Lanczos Quadrature (adapted from [CTU21])

Input: Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, number of iterations $m(\leq n)$.

- 1: Sample $\mathbf{g} \sim \mathcal{U}(\mathcal{S}^{n-1})$.
 - 2: Run Lanczos (Algorithm 3) with inputs $\mathbf{A}, \mathbf{g}$ and m to compute symmetric tridiagonal matrix $\mathbf{T} \in \mathbb{R}^{m \times m}$, orthonormal basis $\mathbf{Q} \in \mathbb{R}^{n \times m}$. Let the eigenvectors of $\mathbf{T}$ be $\mathbf{v}_1, \dots, \mathbf{v}_m$.
 - 3: Set $f(x) = \sum_{j=1}^m w_j^2 \delta(x - \lambda_j(\mathbf{T}))$ where $w_j = \mathbf{v}_j^T \mathbf{e}_1$ where $\mathbf{e}_1 \in \mathbb{R}^m$ is the first standard basis unit vector (i.e. a 1 in the first position and a 0 everywhere else)
 - 4: **return** $f(x)$
-

We give an error bound for SLQ by showing that the normalized Chebyshev moments of the output density are approximately equal to the normalized Chebyshev moments of the spectral density of the input matrix $\mathbf{A}$. We start by proving a lemma showing that the j th Chebyshev moment of the output of SLQ f for any $j \leq m-1$ is exactly given by $\mathbf{g}^T \bar{T}_j(\mathbf{A}) \mathbf{g}$ where $\mathbf{g}$ is the random starting vector.

Lemma 10. *Consider Algorithm 4 run with input $\mathbf{A} \in \mathbb{R}^{n \times n}$, number of iterations m , and sampled vector $\mathbf{g} \sim \mathcal{U}(\mathcal{S}^{n-1})$ in line 1. Let $f(x) = \sum_{i=1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}))$ be the output of the algorithm. Then, for any $j \in \{0, 1, 2, \dots, m-1\}$, $\langle \bar{T}_j, f \rangle = \mathbf{g}^T \bar{T}_j(\mathbf{A}) \mathbf{g}$.*

Proof. Let $L = \|\mathbf{T}\|_2$. The j^{th} normalized Chebyshev moment of f is given by $\langle \bar{T}_j, f \rangle = \int_{-L}^L \bar{T}_j(x) f(x) dx = \int_{-L}^L \bar{T}_j(x) \sum_{i=1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T})) dx = \sum_{i=1}^m w_i^2 \bar{T}_j(\lambda_i(\mathbf{T})) = \sum_{i=1}^m \bar{T}_j(\lambda_i(\mathbf{T})) \mathbf{e}_1^T \mathbf{v}_i \mathbf{v}_i^T \mathbf{e}_1 = \mathbf{e}_1^T (\sum_{i=1}^m \bar{T}_j(\lambda_i(\mathbf{T})) \mathbf{v}_i \mathbf{v}_i^T) \mathbf{e}_1 = \mathbf{e}_1^T \bar{T}_j(\mathbf{T}) \mathbf{e}_1$.

Let $\mathbf{Q}$ be the orthonormal basis computed by the Lanczos algorithm (Algorithm 3) in line 2 of Algorithm 4. From Lemma 9 we know that $\mathbf{A}^p \mathbf{g} = \mathbf{Q} \mathbf{T}^p \mathbf{Q}^T \mathbf{g}$ for any $p \in [m-1]$. Thus, $\bar{T}_j(\mathbf{A}) \mathbf{g} = \mathbf{Q} \bar{T}_j(\mathbf{T}) \mathbf{Q}^T \mathbf{g}$ for any $j \leq m-1$. Note that $\mathbf{Q} \mathbf{e}_1 = \mathbf{g}$ since $\mathbf{g}$ is set as the first column of $\mathbf{Q}$ in Algorithm 3 ($\mathbf{g}$ is a random unit vector). Thus, we get $\mathbf{e}_1^T \bar{T}_j(\mathbf{T}) \mathbf{e}_1 = \mathbf{g}^T \mathbf{Q} \bar{T}_j(\mathbf{T}) \mathbf{Q}^T \mathbf{g} = \mathbf{g}^T \bar{T}_j(\mathbf{A}) \mathbf{g}$. $\square$

We next prove that the i^{th} normalized Chebyshev moment of the output of SLQ f is almost equal to the i th normalized Chebyshev moment of the SDE of $\mathbf{A}$ via the error bounds for a modified hutchinson's trace estimator [MMM21] that uses a random vector on the unit sphere as opposed to a gaussian or a random sign vector. We note that the analysis of hutchinson's using a random vector on the unit sphere is different from the usual analysis which assumes the elements of the random vector are independent and identically distributed with zero mean. We also note that we require that SLQ use a random vector on the unit sphere as opposed to a say, a gaussian vector or a random sign vector is because such a vector has the same distribution as a normalized gaussian vector. As we will see, this helps us leverage the rotational invariance of the gaussian distribution

to derive error bounds for the convergence of eigenvectors and eigenvalues while still ensuring that the random starting vector for Lanczos is a unit vector. Hence, we first prove error bounds for hutchinson's trace estimator using a single random vector on the unit sphere in Lemma 11.

Lemma 11 (Hutchinson's with random vectors on the unit sphere). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\delta \in (0, 1/2]$ and $\mathbf{g} \sim \mathcal{U}(\mathcal{S}^{n-1})$. Then, assuming $n \geq \Omega(\log(1/\delta))$ and for some fixed constant $C > 0$, with probability at least $1 - \delta$, we have that:*

$$\left| \frac{1}{n} \text{tr}(\mathbf{A}) - \mathbf{g}^T \mathbf{A} \mathbf{g} \right| \leq \frac{C \log(1/\delta)}{n} \|\mathbf{A}\|_F.$$

Proof. We have $\mathbf{g} \stackrel{d}{=} \frac{\mathbf{y}_j}{\sqrt{\sum_{j=1}^l \mathbf{y}_j^2}}$ where $\mathbf{y} \in \mathbb{R}^n$ is such that $\mathbf{y}_j \sim \mathcal{N}(0, 1)$ for $j \in [n]$ [CTU21]. Now, using triangle inequality, we have that :

$$\left| \frac{1}{n} \text{tr}(\mathbf{A}) - \mathbf{g}^T \mathbf{A} \mathbf{g} \right| \leq \left| \frac{1}{n} \text{tr}(\mathbf{A}) - \frac{1}{n} \mathbf{y}^T \mathbf{A} \mathbf{y} \right| + \left| \mathbf{g}^T \mathbf{A} \mathbf{g} - \frac{1}{n} \mathbf{y}^T \mathbf{A} \mathbf{y} \right|. \quad (37)$$

We will bound the terms individually. The first term is just the error bound for the hutchinson's estimator using a random gaussian vector. The second term can be bounded as the norm of $\mathbf{y}$, which is just a Chi-squared distribution, can be shown to concentrate around n using standard concentration bounds. First observe that, from Lemma 2 of [MMM21], we have

$$\left| \frac{1}{n} \text{tr}(\mathbf{A}) - \frac{1}{n} \mathbf{y}^T \mathbf{A} \mathbf{y} \right| \leq \frac{1}{n} \left| \text{tr}(\mathbf{A}) - \mathbf{y}^T \mathbf{A} \mathbf{y} \right| \leq \frac{\log(1/\delta)}{n} \|\mathbf{A}\|_F, \quad (38)$$

with probability at least $1 - \delta$. This bounds the first term in (37). Now, we bound the second term. From (38), we have:

$$|\mathbf{y}^T \mathbf{A} \mathbf{y}| \leq \log(1/\delta) \|\mathbf{A}\|_F + \text{tr}(\mathbf{A}) \leq \log(1/\delta) \|\mathbf{A}\|_F + \sqrt{n} \|\mathbf{A}\|_F. \quad (39)$$

with probability at least $1 - \delta$. Here, the second inequality follows from the fact that $\text{tr}(\mathbf{A}) \leq \sqrt{n} \|\mathbf{A}\|_F$. Observe that by the concentration properties of the Chi-squared distribution, we have $|\frac{1}{n} \|\mathbf{y}\|_2^2 - 1| \leq \sqrt{\frac{\log(l/\delta)}{n}}$, with probability at least $1 - \delta$, assuming $n \geq \Omega(\log(l/\delta))$ [Wai19a]. Rearranging, we get

$$\left| \frac{n}{\|\mathbf{y}\|_2^2} - 1 \right| \leq 2 \sqrt{\frac{\log(l/\delta)}{n}}. \quad (40)$$

Thus, we have:

$$\begin{aligned} \left| \mathbf{g}^T \mathbf{A} \mathbf{g} - \frac{1}{n} \mathbf{y}^T \mathbf{A} \mathbf{y} \right| &= \left| \frac{\mathbf{y}^T \mathbf{A} \mathbf{y}}{\|\mathbf{y}\|_2^2} - \frac{1}{n} \mathbf{y}^T \mathbf{A} \mathbf{y} \right| \\ &\leq \frac{1}{n} |\mathbf{y}^T \mathbf{A} \mathbf{y}| \left| \frac{n}{\|\mathbf{y}\|_2^2} - 1 \right| \\ &\leq \frac{1}{n} (\log(1/\delta) \|\mathbf{A}\|_F + \sqrt{n} \|\mathbf{A}\|_F) \cdot 2 \sqrt{\frac{\log(1/\delta)}{n}} \\ &\leq \left(\frac{2 \log^{3/2}(1/\delta)}{n \sqrt{n}} + \frac{2 \sqrt{\log(1/\delta)}}{n} \right) \|\mathbf{A}\|_F \\ &\leq \frac{3 \log(1/\delta)}{n} \|\mathbf{A}\|_F. \end{aligned}$$

where in the second step, we used the triangle inequality and in the third step, we used the upper bounds from (39) and (40). In the final step, we used the fact that $n \geq \Omega(\log(1/\delta))$. Thus, using the upper bounds from above and from (38) in (37), we finally get $\left| \frac{1}{n} \text{tr}(\mathbf{A}) - \mathbf{g}^T \mathbf{A} \mathbf{g} \right| \leq \frac{C \log(1/\delta) \|\mathbf{A}\|_F}{n}$ for some constant $C > 0$. $\square$

Lemma 12. *Consider the setting of Lemma 10. For $n \geq \Omega(\log(1/\delta))$, with probability at least $1 - \delta$, $|\langle \bar{T}_i, f \rangle - \langle \bar{T}_i, s_{\mathbf{A}} \rangle| \leq \frac{C \log(m/\delta)}{\sqrt{n}}$ for all $i \in \{0, 1, 2, \dots, m-1\}$ where $C > 0$ is a large constant.*

Proof. Observe that $\langle \bar{T}_i, s_{\mathbf{A}} \rangle = \frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{A}))$. From Lemma 10 we have $\langle \bar{T}_i, f \rangle = \mathbf{g}^T \bar{T}_i(\mathbf{A}) \mathbf{g}$ for $i \in \{0, 1, 2, \dots, m-1\}$. From Lemma (11), since $n \geq \Omega(\log^2(1/\delta))$, for every $i \in \{0, 1, \dots, m-1\}$, with probability at least $1 - \frac{\delta}{m}$ for a constant C , we have:

$$\left| \frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{A})) - \mathbf{g}^T \bar{T}_i(\mathbf{A}) \mathbf{g} \right| \leq \frac{C \log(m/\delta)}{n} \|\bar{T}_i(\mathbf{A})\|_F \leq \frac{C \log(m/\delta)}{\sqrt{n}},$$

where in the second inequality, we used the fact that $\|\bar{T}_i(\mathbf{A})\|_2 \leq 1$ since $\|\mathbf{A}\|_2 \leq 1$. Applying a union bound for all $i \in \{0, 1, \dots, m-1\}$ gives the bound. $\square$

We now state the final result of this section which gives the error bound for SLQ.

Theorem 10. *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Let $f(x)$ be the output of Algorithm 4 with input $\mathbf{A}$ and $m = O(\frac{1}{\epsilon})$. Then, for some constant C and for $\epsilon, \delta \in (0, 1)$, we have*

$$W_1(s_{\mathbf{A}}, f) \leq \epsilon \|\mathbf{A}\|_2 + \frac{C \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \|\mathbf{A}\|_2,$$

with probability at least $1 - \delta$. Also, Algorithm 4 performs $m = O(\frac{1}{\epsilon})$ matrix vector products with $\mathbf{A}$.

Proof. Assume that we run Algorithm 4 with $\mathbf{B} = \frac{1}{\|\mathbf{A}\|_2} \mathbf{A}$ as input. Then, let $s_{\mathbf{B}}$ be the spectral density of $\mathbf{B}$ and let $f_{\mathbf{B}}$ be the output of Algorithm 4 after m iterations. Then, observe that

$$W_1(s_{\mathbf{A}}, f) = \|\mathbf{A}\|_2 \cdot W_1(s_{\mathbf{B}}, f_{\mathbf{B}}).$$

Since $\|\mathbf{B}\|_2 \leq 1$, as stated in Lemma 10, the symmetric tridiagonal matrix $\mathbf{T}$, which is the output of Lanczos with input matrix $\mathbf{B}$ and starting vector $\mathbf{g}$ in Line 2 of Algorithm 4, can be written as $\mathbf{T} = \mathbf{Q}^T \mathbf{B} \mathbf{Q}$ where $\mathbf{Q}$ is an orthonormal matrix. Thus, $\|\mathbf{T}\|_2 = \|\mathbf{Q}^T \mathbf{B} \mathbf{Q}\|_2 \leq \|\mathbf{B}\|_2 \leq 1$. So, the support of density function $f_{\mathbf{B}}$ output by Algorithm 4 is in $[-1, 1]$. Using Lemma 3.1 of [BKM22] for any two distributions $s_{\mathbf{B}}, f_{\mathbf{B}}$ with support in $[-1, 1]$, we have

$$W_1(s_{\mathbf{B}}, f_{\mathbf{B}}) \leq \frac{36}{m-1} + 2 \sum_{i=1}^{m-1} \frac{|\langle \bar{T}_i, s_{\mathbf{A}} \rangle - \langle \bar{T}_i, f \rangle|}{i}.$$

From Lemma 12, $|\langle \bar{T}_i, s_{\mathbf{A}} \rangle - \langle \bar{T}_i, f_{\mathbf{B}} \rangle| \leq \frac{C \log(m/\delta)}{\sqrt{n}}$ with probability $1 - \delta$ for $i \in \{0, 1, 2, \dots, m-1\}$ as long as $\delta \in (\frac{1}{e^n}, 1)$ (due to the assumption $n \geq \Omega(\log(l/n))$). By setting $m = O(1/\epsilon)$, this gives us, for constants C_1 and C_2 , $W_1(s_{\mathbf{B}}, f_{\mathbf{B}}) \leq C_1 \epsilon + \frac{C_2 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}}$ from the equation above. Finally, we get $W_1(s_{\mathbf{A}}, f) \leq \|\mathbf{A}\|_2 W_1(s_{\mathbf{B}}, f_{\mathbf{B}}) \leq C_1 \epsilon \|\mathbf{A}\|_2 + \frac{C_2 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \|\mathbf{A}\|_2$

Since the Lanczos algorithm is run for $m = O(\frac{1}{\epsilon})$ iterations, the number of matrix vector products with $\mathbf{A}$ is $O(\frac{1}{\epsilon})$. $\square$

We also note that the second term $\frac{\log(1/\epsilon\delta)\log(1/\epsilon)}{\sqrt{n}}$ in the error bound $\epsilon\|\mathbf{A}\|_2 + \frac{\log(1/\epsilon\delta)\log(1/\epsilon)}{\sqrt{n}}\|\mathbf{A}\|_2$ can be made smaller ($< \epsilon\|\mathbf{A}\|_2$) by averaging the resulting distributions over multiple random starting vectors in Algorithm 4 instead of a single random starting vector. However, this would complicate the analysis for the improved deflation based bounds for SLQ that we derive in the subsequent section as it would require carefully analyzing the convergence of the ‘average’ distribution for different starting vectors for SLQ. Hence, we do the analysis with a single random starting vector.

4.2 Error Bounds for Lanczos

Since SLQ uses the Lanczos algorithm (Algorithm 3) as a subroutine, we now derive eigenvalue and eigenvector approximation error bounds for the Lanczos algorithm, which is a Krylov method with a single random starting vector. We will use some results and proof techniques developed in [MMM24] to derive our bounds. We start by describing the critical observation in [MMM24] that the span of the Krylov subspace generated with a single vector as the starting block is the same as the span of the Krylov subspace with a large starting block with fewer iterations. Let Krylov subspace generated by Lanczos after q iterations with starting vector $\mathbf{g}$ be

$$\mathcal{K}_q(\mathbf{A}, \mathbf{g}) = [\mathbf{g}, \dots, \mathbf{A}^{q-1}\mathbf{g}].$$

Note that here we overload the notation as in Section 3, we had defined $\mathcal{K}_q(\mathbf{A}, \mathbf{g})$ in terms of $\mathbf{A}\mathbf{A}^T$ and $\mathbf{A}\mathbf{g}$ while here we define it in terms of $\mathbf{A}$ and $\mathbf{g}$. The Lanczos algorithm finds an orthonormal basis $\mathbf{Q}$ of $\mathcal{K}_q(\mathbf{A}, \mathbf{g})$ such that $\mathbf{T} = \mathbf{Q}^T\mathbf{A}\mathbf{Q}$ where $\mathbf{T}$ is a tridiagonal matrix. We are interested in bounding the error between the eigenvalues of $\mathbf{T}$ and the true eigenvalues of $\mathbf{A}$. Since the eigenvalues of $\mathbf{T}$ are the same up to a rotation of the orthonormal basis $\mathbf{Q}$ of the Krylov subspace, the eigenvalues estimated by the Lanczos algorithm depends only on the span of the Krylov subspace $\mathcal{K}$ generated after q iterations and not on the specific $\mathbf{Q}$ found by Algorithm 3. Let $\mathbf{S}_l$ be such that:

$$\mathbf{S}_l = [\mathbf{g}, \mathbf{A}\mathbf{g}, \mathbf{A}^2\mathbf{g}, \dots, \mathbf{A}^{l-1}\mathbf{g}]. \quad (41)$$

Then observe that:

$$\text{span}(\mathcal{K}_q(\mathbf{A}, \mathbf{g})) = \text{span}\left([\mathbf{S}_l, \mathbf{A}\mathbf{S}_l, \mathbf{A}^2\mathbf{S}_l, \dots, \mathbf{A}^{q-l}\mathbf{S}_l]\right). \quad (42)$$

So, the span of the Krylov subspace generated by Lanczos with $\mathbf{g}$ as the starting vector after q iterations matches the span of the Krylov subspace generated after $q - l + 1$ iterations with $\mathbf{S}_l$ as the starting block. So, it is enough to analyze the orthonormal basis of $\mathcal{K}_{q-l+1}(\mathbf{A}, \mathbf{S}_l)$. Let $\mathbf{Q}$ be the orthonormal basis for $\mathcal{K}_{q-l+1}(\mathbf{A}, \mathbf{S}_l)$ found by the Lanczos algorithm (Algorithm 3). Similar to Lemma 2, we first want to show that $\mathbf{Q}$ approximately spans the top subspace of $\mathbf{A}$. However, note that the proof of Lemma 2 relies on Lemma 1 which assumes that for the starting block $\mathbf{X} \in \mathbb{R}^{n \times l}$, we have $\text{rank}(\mathbf{V}_l^T \mathbf{X}) = l$. While this is true for a random Gaussian starting block $\mathbf{X} \in \mathbb{R}^{n \times l}$, the starting block $\mathbf{S}_l$ in (41) is far from being completely random as its columns are highly correlated. We will first prove that $\text{rank}(\mathbf{S}_l^T \mathbf{U}_l) = l$ (where $\mathbf{U}_l \in \mathbb{R}^{n \times l}$ contains the first l columns of $\mathbf{U}$) by following the approach presented in the proof of Theorem 3 in [MMM24]. Then, we will apply Lemma 1 to bound the projection error of the top subspace of $\mathbf{A}$ onto $\mathbf{Q}\mathbf{Q}^T$.

Gap Dependence. We note that the number of iterations of Lanczos will depend logarithmically on the inverse of the minimum relative singular value gap $g_{\min} = \min_{i \in [l]} \frac{\sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A})}{\sigma_i(\mathbf{A})}$ among

the top l singular values of $\mathbf{A}$, similar to the bounds in [MMM24]. We need $g_{\min} > 0$ to prove $\text{rank}(\mathbf{U}_l^T \mathbf{S}_l) = l$ by using the fact that a degree $l - 1$ polynomial cannot be exactly 0 of l distinct points. Note that if $g_{\min} = 0$, then $\mathbf{Q}$ will not converge to the top subspace. For example, when $\mathbf{A}$ is an identity matrix, $\mathbf{Q}$ only spans the starting vector (which is always an eigenvector of the identity matrix) and Lanczos never finds the other eigenvectors. In general, it is reasonable to expect that for most matrices, $g_{\min}$ should be at least a small constant. We also note that the bounds can be made gap independent by a random perturbation analysis as in [MMM24].

Lemma 13. *For a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ such that $\text{rank}(\mathbf{A}) > l$, let $\mathbf{Q}$ be an orthonormal basis of the Krylov subspace $\mathcal{K}_{q-l+1}(\mathbf{A}, \mathbf{S}_l)$ where $\mathbf{S}_l = [\mathbf{g}, \mathbf{A}\mathbf{g}, \mathbf{A}^2\mathbf{g}, \dots, \mathbf{A}^{l-1}\mathbf{g}]$ and $\mathbf{g} \sim \mathcal{U}(\mathcal{S}^{n-1})$.*

Let $g_{\min} = \min_{i \in [l]} \frac{\sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A})}{\sigma_i(\mathbf{A})}$. Let $\alpha = \max\left(\sigma_{l+1}(\mathbf{A}), \frac{\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}}\right)$ for some large constant $c > 0$ and let $k \in [l]$ such that $\sigma_k(\mathbf{A}) \geq 2\alpha$ and $\sigma_{k+1}(\mathbf{A}) \leq 2\alpha$. Then, for any $\epsilon \in (0, 1]$, $\delta \in (0, 1)$ $\kappa = \frac{\|\mathbf{A}\|_2}{2\alpha}$ and $q = O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta}))$, we have

$$\|\mathbf{U}_k - \mathbf{Q}\mathbf{Q}^T \mathbf{U}_k\|_F^2 \leq \frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}},$$

with probability $1 - \delta$.

Proof. Our proof will utilize the results and proof the techniques from Lemma 1, Lemma 2 and Theorem 3 of [MMM24]. We will first prove that $\text{rank}(\mathbf{S}_l^T \mathbf{U}_l) = l$ (where $\mathbf{U}_l \in \mathbb{R}^{n \times l}$ contains the first l columns of $\mathbf{U}$) by following the approach presented in the proof of Theorem 3 in [MMM24]. Note that we can't directly use Theorem 3 from [MMM24] as it is stated for a gaussian starting vector while our starting vector is random on the unit sphere. Observe that for any $\mathbf{x} \in \mathbb{R}^l$, $\mathbf{S}_l \mathbf{x} = \hat{p}(\mathbf{A})\mathbf{g}$ for some degree $l - 1$ polynomial $\hat{p}$ with coefficients determined by the entries in $\mathbf{x}$. Also, by the rotational invariance of the gaussian distribution, we have $(\mathbf{U}_l^T \mathbf{g})_i \stackrel{d}{=} \frac{y_i}{\sqrt{\sum_{j=1}^l y_j^2}}$ where $y_i \sim \mathcal{N}(0, 1)$ for $i \in [l]$. Then,

$$\mathbf{U}_l^T \mathbf{S}_l \mathbf{x} = \mathbf{U}_l^T \hat{p}(\mathbf{A})\mathbf{g} = \mathbf{U}_l^T \mathbf{U} \hat{p}(\mathbf{A}) \mathbf{U}^T \mathbf{g} := \hat{p}(\mathbf{\Lambda}_l) \frac{\mathbf{y}}{\|\mathbf{y}\|_2},$$

where $\mathbf{\Lambda}_l$ contains the top l eigenvalues of $\mathbf{A}$ on its diagonal. By Lemma 4 of [MMM24], we have $\min_{i \in [l]} y_i^2 \geq \frac{2\delta^2}{\pi l^2}$ with probability at least $1 - \delta$. Then, we can bound the numerator above as:

$$\|\hat{p}(\mathbf{\Lambda}_l) \mathbf{y}\|_2^2 = \sum_{i=1}^l (\hat{p}(\lambda_i(\mathbf{A})))^2 y_i^2 \geq \frac{2\delta^2}{\pi l^2} \sum_{i=1}^l (\hat{p}(\lambda_i(\mathbf{A})))^2. \quad (43)$$

Since $\hat{p}$ has degree $l - 1$, and none of the eigenvalues are repeated (as $g_{\min} > 0$), $\|\hat{p}(\mathbf{\Lambda}_l) \mathbf{y}\|_2^2 > 0$. We also have $\|\mathbf{y}\|_2 > 0$ with probability 1. Thus, we get that $\|\mathbf{U}_l^T \mathbf{S}_l \mathbf{x}\|_2^2 > 0$ for any $\mathbf{x}$. So, $\sigma_{\min}(\mathbf{U}_l^T \mathbf{S}_l) > 0$ and hence, $\mathbf{U}_l^T \mathbf{S}_l$ is invertible which means $\text{rank}(\mathbf{S}_l^T \mathbf{U}_l) = l$. Note that we also have $\sigma_k(\mathbf{A}) \geq 2\alpha \geq 2\sigma_{l+1}(\mathbf{A})$ and so $k \in [l]$. So, we can apply Lemma 1 to bound $\|\sin \Theta(\mathbf{Q}, \mathbf{U}_k)\|_F^2$. Without loss of generality, assume that q is odd. Let $\phi(x)$ be a gap amplifying polynomial of degree $q + 1$ consisting of only even powers as defined in Lemma 4.5 of [DIKMI18] with parameters $\alpha = \sigma_{l+1}(\mathbf{A})$ and gap $\gamma = \frac{\sigma_k(\mathbf{A})}{\sigma_{l+1}(\mathbf{A})} - 1 \geq 1$. Note that for even q , we can similarly define an amplifying polynomial of degree q consisting of only even powers. For any $i \in [k]$, $\phi(\sigma_i(\mathbf{A})) \geq \sigma_i(\mathbf{A}) > 0$ as

$\sigma_i(\mathbf{A}) \geq \sigma_k(\mathbf{A}) \geq (1 + \gamma)\sigma_{l+1}(\mathbf{A}) = (1 + \gamma)\alpha$ by our choice of parameters for the gap-amplifying polynomial. Thus, $\phi(\mathbf{A}_p)$ is non-singular for any $p \in [l]$. Let $\Phi = \mathbf{U}\phi(\mathbf{A})\mathbf{U}^T\mathbf{S}_l$. The columns of Φ lie in $\text{span}(\mathcal{K}_{q-1})$ (42) and thus $\text{range}(\Phi) \subseteq \text{range}(\mathcal{K}_{q-1}) = \text{range}(\mathbf{Q})$. Also, $\Phi\Phi^\dagger$ and $\mathbf{Q}\mathbf{Q}^T$ are the orthogonal projectors onto $\text{range}(\Phi)$ and $\text{range}(\mathcal{K}_{q-1})$ respectively. So, following the proof of Lemma 1 we get:

$$\|\mathbf{U}_k - \mathbf{Q}\mathbf{Q}^T\mathbf{U}_k\|_F^2 \leq \|\phi(\mathbf{A}_{l,\perp})\|_2^2 \|\phi(\mathbf{A}_k)^{-1}\|_2^2 \|\mathbf{U}_{l,\perp}^T \mathbf{S}_l (\mathbf{U}_l^T \mathbf{S}_l)^{-1}\|_2^2. \quad (44)$$

Since q is odd, $q + 1$ is even, giving $\phi(\lambda_i(\mathbf{A})) = \phi(-\lambda_i(\mathbf{A})) = \phi(\sigma_i(\mathbf{A}))$ for $i \in [n]$. So, $\phi(\mathbf{A}_{l,\perp}) = \phi(\mathbf{S}_{l,\perp})$ and $\phi(\mathbf{A}_k) = \phi(\mathbf{S}_k)$. Following the proof of Lemma 2 (which in turn uses Lemma 4.5 of [DIKMI18] based on the properties of ϕ), we get the bounds

$$\|\phi(\mathbf{A}_{l,\perp})\|_2 = \|\phi(\mathbf{S}_{l,\perp})\|_2 \leq \frac{4\sigma_{l+1}(\mathbf{A})}{2^{(q+1)\min(\sqrt{\gamma},1)}}. \quad (45)$$

and

$$\|\phi(\mathbf{A}_k)^{-1}\|_2 = \|\phi(\mathbf{S}_k)^{-1}\|_2 \leq \sigma_k(\mathbf{A}). \quad (46)$$

So, the final step is to bound $\|\mathbf{U}_{l,\perp}^T \mathbf{S}_l (\mathbf{U}_l^T \mathbf{S}_l)^{-1}\|_2^2$. We will bound this by following the proof technique presented in Theorem 3 of [MMM24] to bound the same quantity. We just give an outline of the proof and skip the details since the quantity is the same. Observe that we have:

$$\begin{aligned} \|\mathbf{U}_{l,\perp}^T \mathbf{S}_l (\mathbf{U}_l^T \mathbf{S}_l)^{-1}\|_2^2 &= \max_{\mathbf{x}} \frac{\|\mathbf{U}_{l,\perp}^T \mathbf{S}_l (\mathbf{U}_l^T \mathbf{S}_l)^{-1} \mathbf{x}\|_2}{\|\mathbf{x}\|_2} \\ &\leq \max_{\mathbf{x}} \frac{\|\mathbf{U}_{l,\perp}^T \mathbf{S}_l \mathbf{x}\|_2}{\|\mathbf{U}_l^T \mathbf{S}_l \mathbf{x}\|_2} = \max_{\deg(\hat{p}) \leq l-1} \frac{\|\mathbf{U}_{l,\perp}^T \hat{p}(\mathbf{A}) \mathbf{y}\|_2^2}{\|\mathbf{U}_l^T \hat{p}(\mathbf{A})\|_2^2} \\ &= \max_{\deg(\hat{p}) \leq l-1} \frac{\|\hat{p}(\mathbf{A}_{l,\perp}) \mathbf{y}\|_2^2}{\|\hat{p}(\mathbf{A}_l) \mathbf{y}\|_2^2}. \end{aligned} \quad (47)$$

The denominator is already bounded from below in (43). We now bound the numerator following the proof in [MMM24]. Note that $y_i^2 \leq 1 + 4\log(1/\delta)$ for all $i \in [l]$ with probability at least $1 - \delta$ by standard concentration bounds for chi-squared random variables [Wai19b]. Then, by a union bound, $\max_i y_i \leq 5\log(n/\delta)$ for $n > 2$. We thus have:

$$\|\hat{p}(\mathbf{A}_{l,\perp}) \mathbf{y}\|_2^2 \leq 5\log(n/\delta) \sum_{i=1}^n (\hat{p}(\lambda_i(\mathbf{A})))^2 \leq 5n\log(n/\delta) \max_{i \in [n]} (\hat{p}(\lambda_i(\mathbf{A})))^2 \quad (48)$$

Then, combining (48) and (43), we get:

$$\|\mathbf{U}_{l,\perp}^T \mathbf{S}_l (\mathbf{U}_l^T \mathbf{S}_l)^{-1}\|_2^2 \leq \frac{5\pi n l^2 \log(n/\delta)}{2\delta^2} \frac{\max_{i \in [n]} (\hat{p}(\lambda_i(\mathbf{A})))^2}{\sum_{i=1}^l (\hat{p}(\lambda_i(\mathbf{A})))^2}$$

We can then bound $\frac{\max_{i \in [n]} (\hat{p}(\lambda_i(\mathbf{A})))^2}{\sum_{i=1}^l (\hat{p}(\lambda_i(\mathbf{A})))^2} \leq \frac{l}{g_{\min}^{4l}}$ by expanding $\hat{p}$ as a Lagrange interpolating polynomial over $\sigma_1^2(\mathbf{A}), \dots, \sigma_l^2(\mathbf{A})$ in exactly the same way as in the proof of Theorem 3 in [MMM24]. We finally get

$$\|\mathbf{U}_{l,\perp}^T \mathbf{S}_l (\mathbf{U}_l^T \mathbf{S}_l)^{-1}\|_2^2 \leq \frac{5\pi n l^3}{2g_{\min}^{4l} \delta^2} \log\left(\frac{n}{\delta}\right). \quad (49)$$

Then, using the bounds (49), (46) and (45) on the right hand side of (44) and using the fact that $\gamma = \frac{\sigma_k(\mathbf{A})}{\sigma_{l+1}(\mathbf{A})} - 1 \geq 1$ (by assumption) we get that for $q = O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta}))$ (as long as the constant $c > 0$ is large enough):

$$\|\mathbf{U}_k - \mathbf{Q}\mathbf{Q}^T \mathbf{U}_k\|_F^2 \leq O\left(\frac{\sigma_{l+1}^2(\mathbf{A})}{\sigma_k^2(\mathbf{A})} \cdot \frac{nl^3 \log(n/\delta)}{2^{(2q+1) \min(\sqrt{\gamma}, 1)} g_{\min}^{4l} \delta^2}\right) \leq \frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}.$$

□

Based on the results from Lemma 13, we can now generalize the error bounds for the randomized block krylov method in Section 3.1 to the single vector Lanczos Algorithm 3. We will utilize the gap between the eigenvalues ($g_{\min} > 0$) to give stronger convergence guarantees for the top k eigenvalues and eigenvectors of $\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T$. We first state a stronger version of Theorem 7 which gives shows that the large magnitude eigenvalues of $\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T$ converge to those of $\mathbf{A}$ as long as there exists some k such that $\sigma_k(\mathbf{A})$ is larger than $\sigma_{l+1}(\mathbf{A})$ by at least a constant factor. Roughly, we are able to prove the stronger statement as the gaps between the singular values ensure the estimated singular values are also well separated.

Lemma 14. *Consider the setting of Lemma 13. Let $c > 0$ be the constant in the error bound of Lemma 13. Then, for every $i \in [k]$, with probability at least $1 - \delta$, we have:*

$$|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T)| \leq \frac{\|\mathbf{A}\|_2 g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}.$$

Proof. First, we can follow the proofs of Lemmas 4 and 5 along with the stronger bound on $\|\mathbf{U}_k - \mathbf{Q}\mathbf{Q}^T \mathbf{U}_k\|_2$ from Lemma 13 to prove that the following guarantees hold with probability at least $1 - \delta$:

1. $\|\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T \mathbf{U}_k - \mathbf{U}_k \mathbf{U}_k^T \mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T \mathbf{U}_k\|_2 \leq \frac{\|\mathbf{A}\|_2 g_{\min}^{cl}}{n^{c/\epsilon-1} (\kappa)^{c/\epsilon}},$
2. $|\lambda_i(\mathbf{U}_k^T \mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T \mathbf{U}_k) - \lambda_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\|_2 g_{\min}^{cl}}{n^{c/\epsilon-1} (\kappa)^{c/\epsilon}}.$

Next, we can follow the proof of Theorem 7, along with the stronger error bounds above and in Lemma 13, with the stronger bound of $\frac{\|\mathbf{A}\|_2 g_{\min}^{cl}}{n^{c/\epsilon-1} (\kappa)^{c/\epsilon}}$, which (similar to (23)) gives us that for every $i \in [k]$, there exists some $\lambda_j(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T)$ such that:

$$|\lambda_j(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\|_2 g_{\min}^{cl}}{n^{c/\epsilon-1} (\kappa)^{c/\epsilon}}, \quad (50)$$

Using the min-max principle of singular values, $\sigma_i(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T) \leq \sigma_i(\mathbf{A})$ for all $i \in [n]$. We will now prove that $j = i$ for all $i \in [k]$. Assume, for contradiction, there exists some $i \in [k]$ such that (50) is only satisfied by some $\lambda_j(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T)$ such that $j > i$. Then, we get:

$$\begin{aligned} |\lambda_j(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A})| &\geq \sigma_i(\mathbf{A}) - \sigma_j(\mathbf{Q}\mathbf{Q}^T \mathbf{A} \mathbf{Q}\mathbf{Q}^T) \\ &\geq \sigma_i(\mathbf{A}) - \sigma_j(\mathbf{A}) \\ &\geq \sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A}) \\ &\geq g_{\min} \sigma_i(\mathbf{A}) \\ &\geq g_{\min} \sigma_k(\mathbf{A}) \geq 2g_{\min} \alpha \geq \frac{2\|\mathbf{A}\|_2 g_{\min}^{(c/4+1)}}{n^{c/4}} \geq \frac{\|\mathbf{A}\|_2 g_{\min}^{cl}}{n^{c/\epsilon-1} (\kappa)^{c/\epsilon}}. \end{aligned}$$

for a large enough $c > 0$. This contradicts (50) and thus, we must have $j \leq i$. We can similarly rule out the case $j < i$. Thus, we must have $j = i$ for every $i \in [k]$. $\square$

We next prove a stronger version of Theorem 5 below which shows that every eigenvector of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$, corresponding to a large magnitude eigenvalue, converges to the corresponding eigenvector of $\mathbf{A}$ as long as there exists some k such that $\sigma_k(\mathbf{A})$ is larger than $\sigma_{l+1}(\mathbf{A})$ by at least a constant factor.

Lemma 15. *Consider the setting of Lemma 13. Let the eigenvectors of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$ be $\mathbf{z}_1, \dots, \mathbf{z}_n$. Then, for any $i \in [k]$, with probability at least $1 - \delta$, for some constant c_1 , we have*

$$\|\mathbf{z}_i - \mathbf{u}_i\|_2 \leq \frac{g_{\min}^{(c_1 l)}}{(n \cdot \kappa)^{\frac{c_1}{\epsilon}}} \text{ or } \|\mathbf{z}_i + \mathbf{u}_i\|_2 \leq \frac{g_{\min}^{(c_1 l)}}{(n \cdot \kappa)^{\frac{c_1}{\epsilon}}}.$$

Proof. From Lemma 13, we have that for any eigenvector $\mathbf{u}_i$ of $\mathbf{A}$ for $i \in [k]$, $\mathbf{Q}\mathbf{Q}^T\mathbf{u}_i = \mathbf{u}_i + \mathbf{e}_i$ where $\|\mathbf{e}_i\|_2 \leq \frac{\|\mathbf{A}\|2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$ for some constant $c > 0$. So, we have $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{u}_i = \mathbf{Q}\mathbf{Q}^T\mathbf{A}(\mathbf{u}_i + \mathbf{e}_i) = \lambda_i(\mathbf{A})\mathbf{Q}\mathbf{Q}^T\mathbf{u}_i + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{e}_i = \lambda_i(\mathbf{A})(\mathbf{u}_i + \mathbf{e}_i) + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{e}_i = \lambda_i(\mathbf{A})\mathbf{u}_i + \lambda_i(\mathbf{A})\mathbf{e}_i + \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{e}_i$. Thus, we get that

$$\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{u}_i = \lambda_i(\mathbf{A})\mathbf{u}_i + \mathbf{r}_i, \quad (51)$$

where $\|\mathbf{r}_i\|_2 \leq \frac{2\|\mathbf{A}\|2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$.

Since the eigenvectors $\mathbf{z}_1, \dots, \mathbf{z}_n$ form an orthonormal basis for $\mathbb{R}^n$, we can write $\mathbf{u}_i = \sum_{j=1}^n a_j \mathbf{z}_j$ where a_j are constants such that $\sum_{j=1}^n a_j^2 = 1$. Thus,

$$\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{u}_i = \sum_{j=1}^n a_j \lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) \mathbf{z}_j. \quad (52)$$

From (51), we get $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T\mathbf{u}_i = \sum_{j=1}^n \lambda_i(\mathbf{A}) a_j \mathbf{z}_j + \mathbf{r}_i$. Hence, from (52) we have: $\sum_{j=1}^n \lambda_i(\mathbf{A}) a_j \mathbf{z}_j + \mathbf{r}_i = \sum_{j=1}^n a_j \lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) \mathbf{z}_j$. Rearranging, we get: $\sum_{j=1}^n (\lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A})) a_j \mathbf{z}_j = \mathbf{r}_i$. Taking 2-norm on both sides, we get:

$$\left\| \sum_{j=1}^n (\lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A})) a_j \mathbf{z}_j \right\|_2 \leq \frac{2g_{\min}^{cl} \|\mathbf{A}\|_2}{(n \cdot \kappa)^{c/\epsilon}},$$

using the fact that $\|\mathbf{r}_i\|_2 \leq \frac{2g_{\min}^{cl} \|\mathbf{A}\|_2}{(n \cdot \kappa)^{c/\epsilon}}$. Now, $\left\| \sum_{j=1}^n (\lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A})) a_j \mathbf{z}_j \right\|_2 = \sqrt{\sum_{j=1}^n (\lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A}))^2 a_j^2}$. Thus, for all $j \in [n]$ we have

$$|\lambda_j (\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) - \lambda_i(\mathbf{A})| \cdot |a_j| \leq \frac{2\|\mathbf{A}\|2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}. \quad (53)$$

Now, we have $|\lambda_i(\mathbf{A}) - \lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \leq \frac{\|\mathbf{A}\|2g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}$ for $i \in [k]$, from Lemma 14. Also, from our

assumptions on the singular value gaps, for any $i \in [k]$ and $j \in [n]$ such that $i \neq j$,

$$\begin{aligned}
|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| &\geq |\sigma_i(\mathbf{A}) - \sigma_j(\mathbf{A})| \geq g_{\min} \max(\sigma_i(\mathbf{A}), \sigma_j(\mathbf{A})) \\
&\geq g_{\min} \sigma_i(\mathbf{A}) \geq g_{\min} \sigma_k(\mathbf{A}) \\
&\geq 2g_{\min} \alpha \\
&\geq \frac{\|\mathbf{A}\| 2g_{\min}^{3cl/4}}{(n \cdot \kappa)^{3c/4\epsilon}},
\end{aligned}$$

Using the min-max principle of singular values, $\sigma_i(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T) \leq \sigma_i(\mathbf{A})$ for all $i \in [n]$. So, for $j \neq i$ and $i \in [k]$, using triangle inequality we have that $|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \geq |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| - |\lambda_j(\mathbf{A}) - \lambda_j(\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T)| \geq \frac{\|\mathbf{A}\| 2g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} - \frac{\|\mathbf{A}\| 2g_{\min}^{3cl/4}}{(n \cdot \kappa)^{3c/4\epsilon}}$. Thus, from (53), we get that:

$$|a_j| \leq \frac{\frac{2\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}}{\frac{\|\mathbf{A}\| 2g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} - \frac{\|\mathbf{A}\| 2g_{\min}^{3cl/4}}{(n \cdot \kappa)^{3c/4\epsilon}}} \leq \frac{\frac{2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}}{\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} - \frac{g_{\min}^{3cl/4}}{(n \cdot \kappa)^{3c/4\epsilon}}} \leq \frac{g_{\min}^{(c_4 l)}}{(n \cdot \kappa)^{(c_4/\epsilon)}},$$

for some constant $c_4 > 0$. Since $\sum_{j=1}^n a_j^2 = 1$, we get that $|a_i| \geq 1 - \frac{g_{\min}^{(c_4 l)}}{(n \cdot \kappa)^{c_4/\epsilon-1}}$. Let $a_i > 0$. Then, using triangle inequality, we have $\|\mathbf{u}_i - \mathbf{z}_i\|_2 = \|\mathbf{u}_i - a_i \mathbf{z}_i + \mathbf{z}_i(a_i - 1)\|_2 \leq \|\mathbf{u}_i - a_i \mathbf{z}_i\|_2 + \|\mathbf{z}_i(a_i - 1)\|_2 \leq \|\sum_{j \neq i} a_j \mathbf{z}_j\|_2 + |a_i - 1| \leq \frac{C' g_{\min}^{(c_4 l)}}{(n \cdot \kappa)^{c_4/2\epsilon-1}}$ where C' is some constant. Similarly, we can show when $a_i < 0$, $\|\mathbf{u}_i + \mathbf{z}_i\|_2 \leq \frac{C' g_{\min}^{(c_4 l)}}{(n \cdot \kappa)^{c_4/2\epsilon-1}}$. We complete the proof by choosing the constant $c_1 > 0$ suitably. $\square$

4.3 Improved Error bounds for SLQ via implicit deflation based analysis

In this section, we prove the main error bounds for SLQ (Algorithm 4) by showing that it implicitly performs a deflation followed by moment matching. We first show that there exists a polynomial $r(x)$ of degree $O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta}))$ which is almost zero on the large magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ which have magnitude greater than $\sigma_k(\mathbf{A})$ and close to one on the small magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ where $k \in [l]$ is an index as defined in Lemmas 14 and 15 such that $\sigma_k(\mathbf{A})$ is larger than $\sigma_{l+1}(\mathbf{A})$ by at least a constant multiplicative gap. Then, a polynomial of the form $t_i(x) = r(x) \bar{T}_i(x)$ where $\bar{T}_i(x)$ is the i^{th} Chebyshev polynomial for $i \in O(1/\epsilon)$ will have degree at most $O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta}))$ and can be represented in the span of the Krylov subspace (42) generated by the Lanczos algorithm (Algorithm 3) as long as we run Lanczos for at least this many number of iterations. Intuitively, this implies that the polynomial $t_i(x)$ behaves like a Chebyshev polynomial for the small magnitude eigenvalues (with magnitude smaller than $\sigma_k(\mathbf{A})$ and $\sigma_k(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$) of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ while it is almost zero on the large magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$. We then show that the moments of the spectral density of the small magnitude eigenvalues $\mathbf{A}$ and the part of the density output by SLQ which only depends on the small magnitude eigenvalues of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ with respect to the polynomial $t_i(x)$ are very close to each other for $i \in O(1/\epsilon)$. Since $t_i(x)$ behaves like a Chebyshev polynomial on these small magnitude eigenvalues, via an argument similar to the Chebyshev moment matching argument in Section 4.1, we say that the spectral densities of the small magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ are close to each other. On the other hand, via the results in Section 4.2, we know that the large magnitude eigenvalues of $\mathbf{A}$ are approximated by the corresponding large magnitude eigenvalues of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$. Combining the arguments for the large

and small magnitude eigenvalues, we claim the the spectral densities of $\mathbf{A}$ and that defined by SLQ algorithm (Algorithm 4) are close to each other.

The polynomial $r(x)$ is defined as $1 - g(x)L(x)$ where the polynomial $g(x)$ is a Chebyshev minimizing polynomial (Lemma 17) and $L(x)$ can be interpreted as a variant of the Lagrange interpolating polynomial through the points $(\lambda_i(\mathbf{A}), \frac{1}{g(\lambda_i(\mathbf{A}))})$ for $i \in [k]$. We describe this in more detail now. Suppose we have a set of k ‘basis’ polynomials such that the i ’th polynomial in the set is almost 1 at $\lambda_i(\mathbf{A})$ and $\lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$ and almost zero at all other $\lambda_j(\mathbf{A})$ and $\lambda_j(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$. Then, $r(x)$ can be defined as 1 minus the sum of these polynomials i.e. we will have $r(\lambda_i(\mathbf{A}))$ and $r(\lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q}))$ almost to 1 for every $i \in [k]$ and almost 0 at all other eigenvalues. As a first step, in Lemma 16 for each $i \in [k]$ we first define a polynomial $p_i(x)$ which is almost 1 at $\lambda_i(\mathbf{A})$ and $\lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$ (referred to as $\tilde{\lambda}_i(\mathbf{A})$ in the lemmas below) and zero at all other $\lambda_j(\mathbf{A})$ and $\lambda_j(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$ for $j \neq i$ and $j \in [l]$. However, outside the top l eigenvalues, i.e. on the small magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$, $p_i(x)$ can be potentially be very large (up to $(2/g_{\min})^{O(l)}$). To ensure the polynomials $p_i(x)$ do not blow up on the small magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$, we will multiply each $p_i(x)$ with a corresponding minimizing polynomial (defined in Lemma 17) which squishes its values down to almost 0 on the small magnitude eigenvalues (while keeping its values almost same on the large magnitude eigenvalues). The final set of polynomials is each $p_i(x)$ multiplied by its corresponding minimizing polynomials. We can then define $r(x)$ as 1 minus this polynomial.

We first define the polynomials $p_i(x)$ for $i \in [k]$ in the lemma below. In the lemmas below, $\tilde{\lambda}_i(\mathbf{A}) = \lambda_i(\mathbf{Q}^T \mathbf{A} \mathbf{Q})$ for ease of notation.

Lemma 16. *Consider the setting of Lemma 13. Let $\tilde{\lambda}_i(\mathbf{A}) = \lambda_i(\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T)$ for $i \in [n]$. Then, for $i \in [k]$,*

$$p_i(x) = \prod_{j \in [l], j \neq i} \left(\frac{x - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right) \prod_{j \in [l], j \neq i} \left(\frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right).$$

Then, for some constant $c > 0$, with probability at least $1 - \delta$:

1. For any $i \in [k]$ $p_i(\lambda_i(\mathbf{A})) = 1$ and $p_j(\lambda_i(\mathbf{A})) = 0$ for $j \neq i$.
2. For any $i \in \{k+1, \dots, l\}$ and $j \in [k]$, $p_j(\lambda_i(\mathbf{A})) = 0$ and $p_j(\tilde{\lambda}_i(\mathbf{A})) = 0$.
3. $|p_i(x)| \leq \frac{2^{3l}}{g_{\min}^{2l}}$ when $|x| \leq |\lambda_k(\mathbf{A})|$ for $i \in [k]$.
4. For any $i \in [k]$, $|p_i(\tilde{\lambda}_i(\mathbf{A})) - 1| \leq \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon-4}}$ and $p_j(\tilde{\lambda}_i(\mathbf{A})) = 0$ for $j \neq i$.

Proof. First observe that, from Lemma 14, $|\lambda_i(\mathbf{A}) - \tilde{\lambda}_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$ for $i \in [k]$ with probability at least $1 - \delta$ for some constant $c > 0$. The first two claims are straightforward from the definition of $p_i(x)$. So we proceed to prove the third claim.

Claim 3: We now bound $|p_i(x)|$ when $|x| \leq \sigma_k(\mathbf{A})$ to prove the third claim. We have

$$|p_i(x)| = \prod_{j \in [l], j \neq i} \left| \frac{x - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right| \prod_{j \in [l], j \neq i} \left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right|.$$

We bound each term separately.

$\left| \frac{x - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right|$: For any $i \in [k]$, $\left| \frac{x - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right| \leq \frac{2 \max(|x|, \sigma_j(\mathbf{A}))}{|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})|} \leq \frac{2 \max(\sigma_i(\mathbf{A}), \sigma_j(\mathbf{A}))}{|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})|} \leq \frac{2}{g_{\min}}$ where the second inequality follows from the fact that $|x - \lambda_j(\mathbf{A})| \leq 2 \max(|x|, \sigma_j(\mathbf{A}))$, the third inequality from the fact that $|x| \leq \sigma_k(\mathbf{A}) \leq \sigma_i(\mathbf{A})$ for $i \in [k]$ and the final inequality follows from the definition of $g_{\min}$.

$\left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right|$: We consider the cases $j \in [k]$ and $k \leq j \leq l$ separately.

Case 1. ($j \in [k]$): We have $\left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right| \leq \frac{2 \max(|x|, |\tilde{\lambda}_j(\mathbf{A})|)}{|\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})| - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}} \leq \frac{2 \max(|x|, |\lambda_j(\mathbf{A})|)}{|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}}$. Here,

for the numerator we used the fact that $|\tilde{\lambda}_j(\mathbf{A})| \leq |\lambda_j(\mathbf{A})|$ by the minimax principle and for the denominator we used triangle inequality to get $|\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})| \geq |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| - |\lambda_j(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})|$ and the fact that $|\lambda_j(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})| \leq \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$ for $j \in [k]$. Since $|x| \leq |\lambda_i(\mathbf{A})|$ for $i \in [k]$ we have:

$$\begin{aligned} \frac{2 \max(|x|, |\lambda_j(\mathbf{A})|)}{|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}} &\leq \frac{2 \max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|)}{|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}} \\ &\leq \frac{2}{\frac{|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})|}{\max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|)} - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{\max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|)(n \cdot \kappa)^{c/\epsilon}}} \\ &\leq \frac{2}{g_{\min} - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{\max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|)(n \cdot \kappa)^{c/\epsilon}}} \\ &\leq \frac{4}{g_{\min}}. \end{aligned}$$

Here, the second to last step follows from the definition of $g_{\min}$ and the last step follows from the assumptions that $\max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|) \geq \sigma_k(\mathbf{A}) \geq 2\alpha \geq \frac{2\|\mathbf{A}\| 2g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/4}}$ (the first step follows from the fact that $i, j \in [k]$), which gives us $\frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{\max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|)(n \cdot \kappa)^{c/\epsilon}} \leq \frac{\|\mathbf{A}\| 2g_{\min}^{3cl/4} (n \cdot \kappa)^{c/4}}{2(n \cdot \kappa)^{c/\epsilon}} \leq \frac{g_{\min}}{2}$. Thus, we get

$$\left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right| \leq \frac{4}{g_{\min}} \text{ for } j \in [k] \text{ and } |x| \leq \sigma_k(\mathbf{A}).$$

Case 2. ($j \in \{k+1, \dots, l\}$): For $k+1 \leq j \leq p$ and $i \in [k]$, we have $|\tilde{\lambda}_j(\mathbf{A})| \leq |\lambda_j(\mathbf{A})| \leq |\lambda_i(\mathbf{A})|$ where the first inequality follows from our assumption and the second from the fact that $j \geq i$. So, $|\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})| \geq |\lambda_i(\mathbf{A})| - |\tilde{\lambda}_j(\mathbf{A})| \geq |\lambda_i(\mathbf{A})| - |\lambda_j(\mathbf{A})| = \sigma_i(\mathbf{A}) - \sigma_j(\mathbf{A})$. So, we get

$$\left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right| \leq \frac{2 \max(|x|, |\tilde{\lambda}_j(\mathbf{A})|)}{\sigma_i(\mathbf{A}) - \sigma_j(\mathbf{A})} \leq \frac{2 \max(|\lambda_i(\mathbf{A})|, |\lambda_j(\mathbf{A})|)}{\sigma_i(\mathbf{A}) - \sigma_j(\mathbf{A})} \leq \frac{2}{g_{\min}},$$

where in the second inequality we used the fact that $|x| \leq |\lambda_i(\mathbf{A})|$ and $|\tilde{\lambda}_j(\mathbf{A})| \leq |\lambda_j(\mathbf{A})|$ by the minimax principle and the last step follows from the definition of $g_{\min}$. So, we finally have $\left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right| \leq \frac{2}{g_{\min}}$ for $j \in \{k+1, \dots, l\}$.

Combining the two cases above, we get $\left| \frac{x - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right| \leq \frac{4}{g_{\min}}$ when $|x| \leq \sigma_k(\mathbf{A})$. Thus, plugging in the upper bound on each term of $|p_i(x)|$, we get $|p_i(x)| < \frac{2^{3l}}{g_{\min}^{2l}}$ when $|x| \leq \sigma_k(\mathbf{A})$ for $i \in [k]$.

Claim 4: We now prove the fourth claim of our theorem. For any $i, j \in [k]$, $p_j(\tilde{\lambda}_i(\mathbf{A})) = 0$ for $j \neq i$ from the definition of $p_j(x)$. This gives the second part of the claim. We now prove the first part

of the claim. Observe that,

$$\left| p_i(\tilde{\lambda}_i(\mathbf{A})) \right| = \prod_{j \in [l], j \neq i} \left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right| \prod_{j \in [l], j \neq i} \left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right|. \quad (54)$$

We will bound each term in the product individually. First, observe that

$$\left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right| \leq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon} |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})|} \leq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl/2}}}{(n \cdot \kappa)^{c/2\epsilon}},$$

where the second step follows from bounding the numerator using fact that $|\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon}}$ for $i \in [k]$ and last step follows from the fact that for $i, j \in [k+1]$ and $i \neq j$ (and a large enough c):

$$\begin{aligned} |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| &\geq |\sigma_i(\mathbf{A}) - \sigma_j(\mathbf{A})| \geq g_{\min} \max(\sigma_i(\mathbf{A}), \sigma_j(\mathbf{A})) \\ &\geq g_{\min} \sigma_k(\mathbf{A}) \\ &\geq \frac{2\|\mathbf{A}\|_{2g_{\min}^{(c/4)+1}}}{(n \cdot \kappa)^{c/4}} \\ &\geq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl/2}}}{(n \cdot \kappa)^{c/2\epsilon}}. \end{aligned} \quad (55)$$

Thus, we get

$$\begin{aligned} \left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} - 1 \right| &= \left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right| \\ &\leq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon} |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})|} \\ &\leq \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}. \end{aligned} \quad (56)$$

Next, we have

$$\begin{aligned} \left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} - 1 \right| &= \left| \frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_i(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right| \\ &\leq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon} |\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})|} \\ &\leq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon} (|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| - \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon}})} \\ &\leq \frac{1}{n^{c/\epsilon} \frac{g_{\min}^{cl/2}}{g_{\min}^{cl} (n \cdot \kappa)^{c/2\epsilon}} - 1} \leq 2 \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}. \end{aligned} \quad (57)$$

In the second step above, we used the fact that $|\lambda_i(\mathbf{A}) - \tilde{\lambda}_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$ for bounding the numerator, in the third step we used triangle inequality for bounding the denominator and the fourth step follows from the fact that $|\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| \geq \frac{\|\mathbf{A}\| 2g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}$ as described in (55). Thus, from the bounds in (56) and (57) we have

$$\begin{aligned}
\left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right) \left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right) &\leq \left(1 + \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right) \\
&\cdot \left(1 + 2 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right) \\
&\leq 1 + 3 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \\
&+ 2 \left(\frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}} \right) \\
&\leq 1 + 5 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right), \tag{58}
\end{aligned}$$

for a large enough c and also,

$$\begin{aligned}
\left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right) \left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right) &\geq \left(1 - \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right) \\
&\cdot \left(1 - 2 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right) \\
&\geq \left(1 - 3 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right). \tag{59}
\end{aligned}$$

Multiplying the upper bounds in (58) for each $j \in [l]$ together, we get

$$\begin{aligned}
\prod_{j \in [l], j \neq i} \left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right) \left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right) &\leq \left(1 + 5 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right)^l \\
&\leq 1 + \sum_{r=1}^l \binom{l}{r} \left(5 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right)^r \\
&\leq 1 + 5l \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon-2}} \right) \\
&\leq 1 + \frac{5g_{\min}^{cl/2}}{(n \cdot \kappa)^{(c/2\epsilon)-3}}, \tag{60}
\end{aligned}$$

for a large enough c . In the last step above, we bounded l by n . In the third step above, we bounded each term $\binom{l}{r} (5(1/(n \cdot \kappa))^{c/2\epsilon})^r$ (for $r \in [l]$) in the binomial expansion of $(1 + 5(\frac{1}{n \cdot \kappa})^{c/2\epsilon})^l$

as ($a = (n \cdot \kappa)$ below)

$$\left(\frac{el}{r}\right)^r \left(5 \left(\frac{1}{a}\right)^{c/2\epsilon}\right)^r \leq \left(5l^2 \left(\frac{1}{a}\right)^{c/2\epsilon}\right)^r \leq \left(\frac{5}{a^{c/2\epsilon-2}}\right)^r \leq \frac{5}{a^{c/2\epsilon-2}},$$

where in the first step we used the well-known upper bound on the binomial coefficient $\binom{l}{r} \leq \left(\frac{el}{r}\right)^r$ and in the final step, we used the fact that $\frac{5}{(n \cdot \kappa)^{c/2\epsilon-2}} \leq 1$ for large $c > 0$. Similarly, multiplying the lower bounds in (59) for each $j \in [l]$ together, we get

$$\begin{aligned} \prod_{j \in [l], j \neq i} \left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \lambda_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})} \right) \left(\frac{\tilde{\lambda}_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})}{\lambda_i(\mathbf{A}) - \tilde{\lambda}_j(\mathbf{A})} \right) &\geq \left(1 - \frac{3g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right)^l \\ &\geq 1 + \sum_{r=1}^l \binom{l}{r} (-1)^r \left(3 \left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right) \right)^r \\ &\geq 1 - \frac{3lg_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon-2}} \\ &\geq 1 - \frac{3g_{\min}^{cl/2}}{(n \cdot \kappa)^{(c/2\epsilon)-3}}, \end{aligned} \tag{61}$$

for large enough c . Thus, using the bounds in (60), (61) and (54), we get $|p_i(\tilde{\lambda}_i(\mathbf{A})) - 1| \leq \frac{5g_{\min}^{cl/2}}{(n \cdot \kappa)^{(c/2\epsilon)-3}} \leq \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{(c/2\epsilon)-4}}$. This proves the first property in claim 4. $\square$

We now define a Chebyshev Minimizing polynomial similar to the polynomial defined in Lemma 5 of [MM15].

Lemma 17 (Chebyshev Minimizing Polynomial). *For values $\xi > 0$, gap $\gamma > 0$, and some even integer $q \geq 1$, there exists a polynomial $g(x)$ of degree q such that:*

1. $g(x) = (1 + \gamma)\xi$ for $|x| = (1 + \gamma)\xi$.
2. $g(x) \geq |x|$ for all $|x| \geq (1 + \gamma)\xi$.
3. $g(x) \leq \frac{\xi}{2^{q \log(1+\gamma)-1}}$ for $x \in [-\xi, \xi]$.

Proof. We define the polynomial $g(x)$ following the proof of Lemma 5 in [MM15]. Let

$$g(x) = (1 + \gamma)\xi \frac{T_q(x/\xi)}{T_q(1 + \gamma)}, \tag{62}$$

where $T_q(x)$ is a Chebyshev polynomial of degree q where q is even. Since the degree q is even, we have $T_q(x) = T_q(-x)$ and thus, $g(x) = g(-x)$. The proofs of the first two properties follow the proof of the first two properties of Lemma 5 of [MM15] exactly. To prove the third property, first observe that using the property of Chebyshev polynomials that $T_q(x) \leq 1$ for $x \in [-1, 1]$, we have $T_q(x/\xi) \leq 1$. So, it suffices to show that $T_q(1 + \gamma) \geq 2^{q \log(1+\gamma)-1}$. Using equation 15 of [MM15] which gives the expression for $T_q(x)$ for $|x| > 1$, we have $T_q(x) \geq \frac{1}{2}(1 + \gamma)^q = \frac{2^{q \log(1+\gamma)}}{2} = 2^{q \log(1+\gamma)-1}$. This completes the proof of the third property. $\square$

We now define the polynomial $r(x)$ in the lemma below.

Lemma 18. *Consider the setting of Lemmas 16. Let $c_1, c_2 > 0$ be some constants. Let $g(x)$ be a degree $q = O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n\kappa}{\delta}))$ Chebyshev Minimizing polynomial as defined in Lemma 17 for some q , with parameters $\xi = \alpha$ and $\gamma = \frac{\sigma_k(\mathbf{A})}{\sigma_{l+1}(\mathbf{A})} - 1$. Let*

$$r(x) = 1 - g(x) \sum_{i=1}^k \frac{p_i(x)}{g(\lambda_i(\mathbf{A}))}.$$

Then, with probability at least $1 - \delta$, we have:

1. $r(x) = 0$ for any $x \in \{\lambda_1(\mathbf{A}), \dots, \lambda_k(\mathbf{A})\}$.
2. $r(x) = 1$ for any $x \in \{\lambda_{k+1}(\mathbf{A}), \dots, \lambda_l(\mathbf{A})\}$ and any $x \in \{\tilde{\lambda}_{k+1}(\mathbf{A}), \dots, \tilde{\lambda}_l(\mathbf{A})\}$.
3. $|r(x)| \leq \frac{g_{\min}^{c_1 l}}{(n \cdot \kappa)^{c_1/\epsilon}}$ for any $x \in \{\tilde{\lambda}_1(\mathbf{A}), \dots, \tilde{\lambda}_k(\mathbf{A})\}$.
4. $|r(x) - 1| \leq \frac{g_{\min}^{c_2 l}}{(n \cdot \kappa)^{c_2/\epsilon}}$ for $|x| \leq \sigma_{l+1}(\mathbf{A})$.

Proof. First observe that, from Lemma 14, $|\lambda_i(\mathbf{A}) - \tilde{\lambda}_i(\mathbf{A})| \leq \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$ for $i \in [k]$ with probability at least $1 - \delta$. Also note that $\gamma \geq 1$ by the assumption that $\sigma_k(\mathbf{A}) \geq 2\alpha$. We now prove the main claims.

The first property follows directly from property 1 of Lemma 16. The second property follows directly from property 2 of Lemma 16. We prove the third property below.

Claim 3: For the third property, observe that for any $i, j \in [k]$ and $j \neq i$, we have $p_j(\tilde{\lambda}_i(\mathbf{A})) = 0$ from property 4 of Lemma 16, so we have

$$r(\tilde{\lambda}_i(\mathbf{A})) = 1 - g(\tilde{\lambda}_i(\mathbf{A})) \sum_{j=1}^k \frac{p_j(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_j(\mathbf{A}))} = 1 - p_i(\tilde{\lambda}_i(\mathbf{A})) \frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))}. \quad (63)$$

From property 4 of Lemma 16, we get that for any $i \in [k]$:

$$|p_i(\tilde{\lambda}_i(\mathbf{A})) - 1| \leq \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon-4}}. \quad (64)$$

First observe that from definition of $g(x)$ in (62) in Lemma 17, we have $\frac{g(x)}{g(y)} = \frac{T_q(x/\xi)}{T_q(y/\xi)}$ where $T_q(x)$ is the q^{th} Chebyshev polynomial for some even q . Now, for any $i \in [k]$, $|\lambda_i(\mathbf{A})| \geq |\lambda_k(\mathbf{A})| \geq 2\alpha$ and so $\frac{|\lambda_i(\mathbf{A})|}{\alpha} \geq 2$. From Lemma 14, we have that for any $i \in [k]$,

$$\begin{aligned} |\tilde{\lambda}_i(\mathbf{A})| &\geq |\lambda_i(\mathbf{A})| - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}} \\ &\geq |\lambda_k(\mathbf{A})| - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}} \\ &\geq 2\alpha - \frac{\|\mathbf{A}\| 2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}. \end{aligned} \quad (65)$$

Now, $\alpha \geq \frac{\|\mathbf{A}\|_{2g_{\min}^{c/4}}}{(n \cdot \kappa)^{c/4}} \geq \frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{(n \cdot \kappa)^{c/\epsilon}}$. Thus, from (65), we get $|\tilde{\lambda}_i(\mathbf{A})| \geq \alpha$. Thus, we also have $\frac{|\tilde{\lambda}_i|}{\alpha} \geq 1$. Also, since q is even we have $T_q(x) = T_q(|x|)$. Thus, from (4) we have:

$$\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} = \frac{T_q(\tilde{\lambda}_i(\mathbf{A})/\alpha)}{T_q(\lambda_i(\mathbf{A})/\alpha)} = \frac{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q + \left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} - \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q}{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\sigma_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q + \left(\frac{\sigma_i(\mathbf{A})}{\alpha} - \sqrt{(\frac{\sigma_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q}. \quad (66)$$

where $\tilde{\sigma}_i(\mathbf{A}) = |\tilde{\lambda}_i(\mathbf{A})|$. We will first upper bound this ratio. Since $\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} - \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1} \leq 1$, the numerator in (66) can be upper bounded as $\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q + 1$. Now, by the minimax principle, we have $\tilde{\sigma}_i(\mathbf{A}) \leq \sigma_i(\mathbf{A})$. Thus, the denominator in (66) is lower bounded by $\left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\sigma_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q \geq \left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q$. Using the upper bound on the numerator and the lower bound on the denominator in (66), we get:

$$\begin{aligned} \frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} &\leq \frac{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q + 1}{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q} \leq 1 + \frac{1}{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right)^q} \\ &\leq 1 + \frac{1}{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha}\right)^q}. \end{aligned} \quad (67)$$

Observe that since $\alpha \geq \frac{\|\mathbf{A}\|_{2g_{\min}^{c/4}}}{n^{c/4}}$, we have

$$\frac{\|\mathbf{A}\|_{2g_{\min}^{cl}}}{\alpha(n \cdot \kappa)^{c/\epsilon}} \leq \frac{g_{\min}^{cl}(n \cdot \kappa)^{c/4}}{g_{\min}^{c/4}(n \cdot \kappa)^{c/\epsilon}} \leq \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}. \quad (68)$$

Dividing both sides of (65) by α and using the fact $\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \leq \frac{1}{2}$, we get:

$$\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} \geq 2 - \frac{g_{\min}^{cl}}{\alpha(n \cdot \kappa)^{c/\epsilon}} \geq 2 - \frac{1}{2} \geq \frac{3}{2}.$$

Thus, using the bound above in (67), we get that:

$$\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \leq 1 + \frac{1}{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha}\right)^q} \leq 1 + \frac{1}{(3/2)^q} \leq 1 + \frac{g_{\min}^{c_3 l}}{(n \cdot \kappa)^{c_3/\epsilon}}, \quad (69)$$

for some large constant $c_3 > 0$. Here, in the last step, we used the fact that $q = O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n}{\delta}))$. We now lower bound $\frac{g(\tilde{\sigma}_i(\mathbf{A}))}{g(\sigma_i(\mathbf{A}))}$.

First observe that the function $x - \sqrt{x^2 - 1}$ is a decreasing function with respect to x . Thus, we have $\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} - \sqrt{(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha})^2 - 1}\right) \geq \left(\frac{\sigma_i(\mathbf{A})}{\alpha} - \sqrt{(\frac{\sigma_i(\mathbf{A})}{\alpha})^2 - 1}\right)$ as $\frac{\sigma_i(\mathbf{A})}{\alpha} \geq \frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha}$. This implies that

we have

$$\begin{aligned} & \left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} - \sqrt{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right) \left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right) \\ & \geq \left(\frac{\sigma_i(\mathbf{A})}{\alpha} - \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right) \left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right). \end{aligned}$$

So, from (66), we have:

$$\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \geq \frac{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right)^q}{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right)^q}. \quad (70)$$

Now, from Lemma 14 and using triangle inequality, we have $\tilde{\sigma}_i(\mathbf{A}) \geq \sigma_i(\mathbf{A}) - \frac{g_{\min}^{cl} \|\mathbf{A}\|_2}{(n \cdot \kappa)^{c/\epsilon}}$ for $i \in [k]$. Dividing both sides by α , we get $\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha} \geq \frac{\sigma_i(\mathbf{A})}{\alpha} - \frac{g_{\min}^{cl} \|\mathbf{A}\|_2}{\alpha(n \cdot \kappa)^{c/\epsilon}} \geq \frac{\sigma_i(\mathbf{A})}{\alpha} - \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}}$ where in the last step we used (68) to upper bound $\frac{g_{\min}^{cl} \|\mathbf{A}\|_2}{\alpha(n \cdot \kappa)^{c/\epsilon}}$. So, using the lower bound on $\frac{\tilde{\sigma}_i(\mathbf{A})}{\alpha}$ in (70), we get:

$$\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \geq \frac{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} - \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} + \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} - \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right)^2 - 1} \right)^q}{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right)^q}. \quad (71)$$

Observe that we have

$$\begin{aligned} & \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} - \frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} \right)^2 - 1} \\ & = \sqrt{\left(\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1 \right) - \left(\frac{2g_{\min}^{cl/2} \sigma_i(\mathbf{A})}{(n \cdot \kappa)^{c/2\epsilon} \alpha} - \frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}} \right)} \\ & \geq \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} - \sqrt{\frac{2g_{\min}^{cl/2} \sigma_i(\mathbf{A})}{(n \cdot \kappa)^{c/2\epsilon} \alpha} - \frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}}, \end{aligned}$$

from the fact that $\sqrt{a-b} \geq \sqrt{a} - \sqrt{b}$. Thus, using the lower bound on the numerator in (71) we

get:

$$\begin{aligned}
\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} &\geq \left(1 - \frac{\left(\frac{g_{\min}^{cl/2}}{(n \cdot \kappa)^{c/2\epsilon}} + \sqrt{\frac{2g_{\min}^{cl/2}\sigma_i(\mathbf{A})}{(n \cdot \kappa)^{c/2\epsilon}\alpha} - \frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}} \right)}{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right)} \right)^q \\
&\geq \left(1 - \frac{2\sqrt{\frac{2g_{\min}^{cl/2}\sigma_i(\mathbf{A})}{(n \cdot \kappa)^{c/2\epsilon}\alpha}}}{\frac{\sigma_i(\mathbf{A})}{\alpha}} \right)^q \\
&= \left(1 - 2\sqrt{\frac{2g_{\min}^{cl/2}\alpha}{(n \cdot \kappa)^{c/2\epsilon}\sigma_i(\mathbf{A})}} \right)^q \\
&\geq \left(1 - \frac{2g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/4\epsilon}} \right)^q. \tag{72}
\end{aligned}$$

The second step above follow from upper bounding the numerator of $\frac{\left(\frac{1}{(n \cdot \kappa)^{c/2\epsilon}} + \sqrt{\frac{2\sigma_i(\mathbf{A})}{(n \cdot \kappa)^{c/\epsilon}\alpha} - \frac{1}{(n \cdot \kappa)^{c/\epsilon}}} \right)}{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} + \sqrt{\left(\frac{\sigma_i(\mathbf{A})}{\alpha} \right)^2 - 1} \right)}$

by $2\sqrt{\frac{2\sigma_i(\mathbf{A})}{(n \cdot \kappa)^{c/\epsilon}\alpha}}$ and its denominator by $\frac{\sigma_i(\mathbf{A})}{\alpha}$. The final step follows from the fact that $\sigma_i(\mathbf{A}) \geq 2\alpha$ for any $i \in [k]$. Now, observe that we have:

$$\begin{aligned}
\left(1 - \frac{2g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/4\epsilon}} \right)^q &= 1 + \sum_{t=1}^q \binom{q}{t} (-1)^t (2g_{\min}^{cl/4}/(n \cdot \kappa)^{c/4\epsilon})^t \\
&\geq 1 - \sum_{t=1}^q (eq/t)^t (2g_{\min}^{cl/4}/(n \cdot \kappa)^{c/4\epsilon})^t \\
&\geq 1 - \sum_{t=1}^q \left(\frac{2eqg_{\min}^{cl/4}}{(n \cdot \kappa)^{c/4\epsilon}} \right)^t \\
&\geq 1 - \sum_{t=1}^q \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{(c/4\epsilon)-2}} \\
&= 1 - \frac{qg_{\min}^{cl/4}}{(n \cdot \kappa)^{(c/4\epsilon)-2}} \\
&\geq 1 - \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{(c/4\epsilon)-3}}. \tag{73}
\end{aligned}$$

In the second step above, we upper bounded the binomial coefficients $\binom{q}{t}$ by $(eq/t)^t$. In the fourth and final step above, we bounded q by n . In the fourth step, we also used the fact that $\frac{2eq}{(n \cdot \kappa)^{c/2\epsilon}} < 1$ for large enough c . So, using the lower bound from (73) in (72), we get:

$$\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \geq 1 - \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{(c/4\epsilon)-3}}. \tag{74}$$

Now, from the upper bounds on $p_i(\tilde{\lambda}_i(\mathbf{A}))$ and $\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))}$ from (64) and (69) respectively, we get:

$$p_i(\tilde{\lambda}_i(\mathbf{A})) \frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \leq 1 + \frac{g_{\min}^{c_4 l}}{(n \cdot \kappa)^{c_4/\epsilon}},$$

for some large enough c_4 . Similarly, from the lower bounds on $p_i(\tilde{\lambda}_i(\mathbf{A}))$ and $\frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))}$ from (64) and (74) respectively, we get:

$$p_i(\tilde{\lambda}_i(\mathbf{A})) \frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \geq 1 - \frac{g_{\min}^{c_5 l}}{(n \cdot \kappa)^{c_5/\epsilon}},$$

for some c_5 . Finally, using the upper and lower bounds on $p_i(\tilde{\lambda}_i(\mathbf{A})) \frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))}$ in (63), we get:

$$|r(\tilde{\lambda}_i(\mathbf{A}))| = \left| 1 - p_i(\tilde{\lambda}_i(\mathbf{A})) \frac{g(\tilde{\lambda}_i(\mathbf{A}))}{g(\lambda_i(\mathbf{A}))} \right| \leq \frac{g_{\min}^{c_8 l}}{(n \cdot \kappa)^{c_8/\epsilon}},$$

for some constant c_8 . This gives us the third claim.

Claim 4: For the fourth and final property, using property 3 for $|p_i(x)|$ from Lemma 16, we get that for any $|x| \leq \sigma_{l+1}(\mathbf{A})$,

$$\begin{aligned} |r(x) - 1| &\leq \sum_{i=1}^k \frac{|g(x)|}{|g(\lambda_i(\mathbf{A}))|} |p_i(x)| \leq \frac{\alpha}{\alpha 2^{(Cl \log(1/g_{\min}) + \frac{C}{\epsilon} \log(n\kappa/\delta))}} \cdot \frac{2^{3l}}{g_{\min}^{2l}} \\ &\leq \frac{g_{\min}^{c_9 l}}{(n \cdot \kappa)^{c_9/\epsilon}}, \end{aligned}$$

for some large constant $c_9 > 0$. Here, in the second step, we also used property 2 of $g(x)$ from Lemma 17, i.e., $g(\lambda_i(\mathbf{A})) \geq |\lambda_i(\mathbf{A})| \geq 2\alpha$ for $|\lambda_i(\mathbf{A})| \geq |\lambda_k(\mathbf{A})| \geq 2\alpha$ for lower bounding the denominator, and property 3 for $g(x)$ when $|x| \leq \sigma_{l+1}(\mathbf{A}) \leq \alpha$ which gives $g(x) \leq \frac{\alpha}{2^{q \log 2 - 1}}$ from Lemma 17 for upper bounding the numerator. $\square$

Definition 4.1 (Moment Matching Polynomial). Let $t_i(x) = \bar{T}_i(x)r(x)$ where $\bar{T}_i(x)$ is the i^{th} normalized Chebyshev polynomial for $i \in [O(1/\epsilon)]$ and $r(x)$ is a degree $O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n\kappa}{\delta}))$ polynomial as defined in Lemma 18. Thus, $t_i(x)$ has degree $O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n\kappa}{\delta}))$.

We run Lanczos with $\mathbf{A}$ as input for $O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n\kappa}{\delta}))$ iterations. We will now show that the moments of the spectral density of $\mathbf{A}$ as well as the output of Algorithm 4 (SLQ) with respect to $t_i(x)$ is approximately equal to the i^{th} normalized Chebyshev moment of the SDE of $\mathbf{A}$ and the output of SLQ respectively for all $i \in [O(1/\epsilon)]$.

Lemma 19. Consider the setting of Lemma 18. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be such that $\|\mathbf{A}\|_2 \geq 1$ and $\kappa = \|\mathbf{A}\|_2$. Let $k \in [l]$ such that $\sigma_k(\mathbf{A}) \geq 1$ and $\sigma_{k+1}(\mathbf{A}) \leq 1$. Let $t_i(x) = \bar{T}_i(x)r(x)$ for $i \in O(\frac{1}{\epsilon})$ be the polynomials in Definition 4.1. Let $f(x)$ be the final output after running Algorithm 4 with $\mathbf{A}$ as input for $m = O(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n\kappa}{\delta}))$ iterations for some $\epsilon, \delta \in (0, 1)$. Then, with probability at least $1 - \delta$, for $i \in O(\frac{1}{\epsilon})$, we have (for constant $c_1, c_2 > 0$)

$$\left| \langle t_i, s_{\mathbf{A}} \rangle - \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{A})) \right| \leq \frac{g_{\min}^{c_1 l}}{(n \cdot \kappa)^{c_1/\epsilon}}$$

and

$$\left| \langle t_i, f \rangle - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) \right| \leq \frac{g_{\min}^{c_2 l}}{(n \cdot \kappa)^{c_2/\epsilon}}.$$

Proof. We have $\mathbf{T} = \mathbf{Q}^T \mathbf{A} \mathbf{Q}$ where $\mathbf{T} \in \mathbb{R}^{m \times m}$ is the output of Algorithm 4 and $\mathbf{Q}$ is the orthonormal basis of the Krylov subspace generated by the Lanczos algorithm in 3. Note that the nonzero eigenvalues of $\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T$ are the same as those of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$ i.e. $\lambda_i(\mathbf{T}) = \lambda_i(\mathbf{Q} \mathbf{Q}^T \mathbf{A} \mathbf{Q} \mathbf{Q}^T)$ for $i \in [m]$.

Moments of $s_{\mathbf{A}}$: We have:

$$\begin{aligned} \langle t_i, s_{\mathbf{A}} \rangle &= \frac{1}{n} \sum_{j=1}^n t_i(\lambda_j(\mathbf{A})) \\ &= \frac{1}{n} \sum_{j=1}^k \bar{T}_i(\lambda_j(\mathbf{A})) r(\lambda_j(\mathbf{A})) + \frac{1}{n} \sum_{j=k+1}^l \bar{T}_i(\lambda_j(\mathbf{A})) r(\lambda_j(\mathbf{A})) + \frac{1}{n} \sum_{j=l+1}^n \bar{T}_i(\lambda_j(\mathbf{A})) r(\lambda_j(\mathbf{A})) \\ &= 0 + \frac{1}{n} \sum_{j=k+1}^l \bar{T}_i(\lambda_j(\mathbf{A})) + \frac{1}{n} \sum_{j=l+1}^n \bar{T}_i(\lambda_j(\mathbf{A})) r(\lambda_j(\mathbf{A})), \end{aligned}$$

where the second step follows from the definition of $t_i(x)$ and $s_{\mathbf{A}}(x)$ and the last step follows from properties 1 and 2 of $r(x)$ in lemma 18. Finally, by property 4 of $r(x)$ in lemma 18, we have $r(\lambda_j(\mathbf{A})) = 1 + e_j$ such that $|e_j| \leq \frac{g_{\min}^{c_7 l}}{(n \cdot \kappa)^{c_9/\epsilon}}$ for $j \in \{l+1, \dots, n\}$ and for some constant $c_9 > 0$. Also, note that $|\bar{T}_j(\lambda_j(\mathbf{A}))| \leq \sqrt{\frac{2}{\pi}}$ for $j \in \{l+1, \dots, n\}$ as $|\lambda_j(\mathbf{A})| \leq |\lambda_{l+1}(\mathbf{A})| = 1$ by our assumptions. So, using triangle inequality we finally get

$$\left| \langle t_i, s_{\mathbf{A}} \rangle - \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{A})) \right| \leq \frac{1}{n} \sum_{j=l+1}^n |\bar{T}_i(\lambda_j(\mathbf{A}))| |e_j| \leq \frac{g_{\min}^{c_9 l}}{(n \cdot \kappa)^{c_9/\epsilon}},$$

for some constant $c_9 > 0$. This proves the first claim of the lemma. We now prove the second claim.

Moments of f : We have

$$\begin{aligned} \langle t_i, f \rangle &= \sum_{j=1}^m w_j^2 t_i(\lambda_j(\mathbf{T})) \\ &= \sum_{j=1}^k w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T})) + \sum_{j=k+1}^l w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T})) + \sum_{j=l+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T})) \\ &= \sum_{j=1}^k w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T})) + \sum_{j=k+1}^l w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) + \sum_{j=l+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T})), \end{aligned} \quad (75)$$

where the last step follows from the properties of $r(x)$ in Lemma 18.

Suppose $|\lambda_j(\mathbf{T})| \geq 1$ for some $j \in [k]$. So, from (4), we get that $T_i(\lambda_j(\mathbf{T}))$ can be written as

$$T_i(\lambda_j(\mathbf{T})) = \frac{1}{2} \left[\left(|\lambda_j(\mathbf{T})| + \sqrt{|\lambda_j(\mathbf{T})|^2 - 1} \right)^i + \left(|\lambda_j(\mathbf{T})| - \sqrt{|\lambda_j(\mathbf{T})|^2 - 1} \right)^i \right] \leq (2|\lambda_j(\mathbf{T})|)^i.$$

Thus,

$$\bar{T}_i(\lambda_j(\mathbf{T})) \leq \sqrt{\frac{2}{\pi}} (2|\lambda_j(\mathbf{T})|)^i. \quad (76)$$

Thus, for any $j \in [k]$, we have $\bar{T}_i(\lambda_j(\mathbf{T})) \leq \sqrt{\frac{2}{\pi}} \max((2|\lambda_j(\mathbf{T})|)^k, 1)$. We can also bound each w_j^2 as $(w_j)^2 = (\mathbf{v}_i^T \mathbf{e}_1)^2 \leq 1$ for all $j \in [m]$. From property 3 of lemma 18, we have $r(\lambda_j(\mathbf{T})) \leq \frac{g_{\min}^{c_6 l}}{(n \cdot \kappa)^{c_6/\epsilon}}$ for $j \in [k]$ and some constant c_6 . So, using the bounds on w_j , $r(\lambda_j(\mathbf{T}))$ and (76), for any $j \in [k]$ such that $|\lambda_j(\mathbf{T})| \geq 1$, we get:

$$\begin{aligned} |w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T}))| &\leq \sqrt{\frac{2}{\pi}} (2\|\mathbf{A}\|_2)^i \left(\frac{g_{\min}^{c_6 l}}{(n \cdot \kappa)^{c_6/\epsilon}} \right) \\ &\leq \sqrt{\frac{2}{\pi}} (2\|\mathbf{A}\|_2)^{c'/\epsilon} \left(\frac{g_{\min}^{c_6 l}}{(n \cdot \kappa)^{c_6/\epsilon}} \right) \\ &\leq \frac{g_{\min}^{c_8 l}}{(n \cdot \kappa)^{c_8/\epsilon}}, \end{aligned} \quad (77)$$

for a large enough constant c_8 , where we also use the facts that $i \leq O(\frac{1}{\epsilon})$ and $\kappa = \|\mathbf{A}\|_2$ by assumption. Similarly, $|w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T}))| \leq 1$ if $|\lambda_j(\mathbf{T})| \leq 1$ for any $j \in [k]$. From property 2 of Lemma 18, we have $r(\lambda_j(\mathbf{T})) = 1$ for $j \in \{k+1, \dots, l\}$. Also, since $|\lambda_j(\mathbf{T})| \leq |\lambda_j(\mathbf{A})| \leq |\lambda_{l+1}(\mathbf{A})| = 1$ for all $j \in \{l+1, \dots, m\}$, from property 4 of Lemma 18, we again have $r(\lambda_j(\mathbf{T})) = 1 + e_j$ such that $|e_j| \leq \frac{g_{\min}^{c_2 l}}{(n \cdot \kappa)^{c_2/\epsilon}}$ for some constant $c_2 > 0$. Also, note that we again have $|\bar{T}_j(\lambda_j(\mathbf{T}))| \leq 1$ for $j \in \{l+1, \dots, m\}$. So, we have

$$\sum_{j=l+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) r(\lambda_j(\mathbf{T})) = \sum_{j=l+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T})) + \sum_{j=l+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{A})) e_j.$$

Finally, from (75), using (77) and the bound above (where $|e_j| \leq \frac{g_{\min}^{c_2 l}}{(n \cdot \kappa)^{c_2/\epsilon}}$) and using triangle inequality we have that

$$\begin{aligned} |\langle t_i, f \rangle - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}))| &\leq \sum_{j=1}^k w_j^2 |\bar{T}_i(\lambda_j(\mathbf{A}))| |r(\lambda_j(\mathbf{T}))| + \sum_{j=l+1}^m w_j^2 |\bar{T}_i(\lambda_j(\mathbf{A}))| |e_j| \\ &\leq \frac{g_{\min}^{c_{10} l}}{(n \cdot \kappa)^{c_{10}/\epsilon}}, \end{aligned}$$

for some large constant $c_{10} > 0$. This proves the second claim. Note that everything above holds with probability at least $1 - \delta$ after adjusting δ by constant factors (as the bounds in Lemma 18 hold with probability at least $1 - \delta$). $\square$

We now bound the difference between the moments of $s_{\mathbf{A}}$ and f with respect to the polynomials $t_i(x)$ for each $i \in O(1/\epsilon)$.

Lemma 20. Consider the setting of Lemma 19. Then, for $n \geq \Omega(\log(1/\delta))$ and $i \in \{0, 1, \dots, \lceil \frac{C_1}{\epsilon} \rceil\}$ where $C_1 > 0$ is some constant,

$$|\langle t_i, f \rangle - \langle t_i, s_{\mathbf{A}} \rangle| \leq \frac{C \log(1/\epsilon\delta)}{\sqrt{n}},$$

for some constant $C > 0$ with probability at least $1 - \delta$.

Proof. Observe that $\langle t_i, f \rangle = \sum_{j=1}^m w_i^2 t_i(\lambda_j(\mathbf{T})) = \sum_{j=1}^m (\mathbf{e}_1^T \mathbf{v}_j)^2 t_i(\lambda_j(\mathbf{T})) = \sum_{j=1}^m \mathbf{e}_1^T \mathbf{v}_j \mathbf{v}_j^T \mathbf{e}_1 t_i(\lambda_j(\mathbf{T})) = \mathbf{e}_1 (\sum_{j=1}^m t_i(\lambda_j(\mathbf{T})) \mathbf{v}_j \mathbf{v}_j^T) \mathbf{e}_1 = \mathbf{e}_1^T t_i(\lambda_j(\mathbf{T})) \mathbf{e}_1$. From Lemma 19, we have $t_i(x) = \bar{T}_i(x)r(x)$ for $i \in \lceil \frac{C_1}{\epsilon} \rceil$. Then, since $r(x)$ is a polynomial of degree $C_2(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta}))$ for some constant $C_2 > 0$ and $\bar{T}_i(x)$ has degree at most $\frac{C_1}{\epsilon}$, $t_i(x) = \bar{T}_i(x)r(x)$ has degree at most $(C_1 + C_2)(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta}))$. Following the proof of Lemma 10, as long as Algorithm 4 is run for m iterations such that $(C_1 + C_2)(l \log(\frac{1}{g_{\min}}) + \frac{1}{\epsilon} \log(\frac{n \cdot \kappa}{\delta})) \leq m$, we have $\langle t_i, f \rangle = \mathbf{g}^T t_i(\mathbf{A}) \mathbf{g}$.

We also have $\langle t_i, s_{\mathbf{A}} \rangle = \frac{1}{n} \text{tr}(t_i(\mathbf{A}))$. Thus, from Lemma 11, we have that (for a single repetition of the hutchinson's estimator and number), for each $i \in O(1/\epsilon)$, with probability at least $1 - O(\delta/\epsilon)$, for some constant $C > 0$:

$$|\langle t_i, f \rangle - \langle t_i, s_{\mathbf{A}} \rangle| = |\mathbf{g}^T t_i(\mathbf{A}) \mathbf{g} - \frac{1}{n} \text{tr}(t_i(\mathbf{A}))| \leq \frac{C \log(1/\epsilon\delta)}{n} \|t_i(\mathbf{A})\|_F \leq \frac{C \log(1/\epsilon\delta)}{\sqrt{n}} \|t_i(\mathbf{A})\|_2.$$

Now, observe that for $j \in [k]$, from property 1 of Lemma 18, we get that $t_i(\lambda_j(\mathbf{A})) = \bar{T}_i(\lambda_j(\mathbf{A}))r(\lambda_j(\mathbf{A})) = 0$. Also, recall that from the assumptions of Lemma 19, we have that $\sigma_{k+1}(\mathbf{A}) = 1$. Now, for $j \geq k+1$, $|t_i(\lambda_j(\mathbf{A}))| = |\bar{T}_i(\lambda_j(\mathbf{A}))r(\lambda_j(\mathbf{A}))| \leq 2$ where we use property 2 of Lemma 18 and the fact that $|T(x)| \leq 1$ for $|x| \leq 1$. Thus, we have $\|t_i(\mathbf{A})\|_2 \leq \max_{j \in \{k+1, \dots, n\}} |t_i(\lambda_j(\mathbf{A}))| \leq 2$. So, we finally get:

$$|\langle t_i, f \rangle - \langle t_i, s_{\mathbf{A}} \rangle| \leq \frac{C \log(1/\epsilon\delta)}{\sqrt{n}}.$$

Taking a union bound over all $i \in O(1/\epsilon)$ completes the claim. $\square$

Next we bound the weights w_i^2 in the distribution f , the output of SLQ. We will show that w_i^2 for the top k weights is at most $\tilde{O}(1/n)$. This will help us bound the Wasserstein error from the spectral density of the large eigenvalues of $\mathbf{T}$.

Lemma 21. Consider the setting of Lemma 19. Let $w_i = \mathbf{v}_i^T \mathbf{e}_1$ be the weights in the output of distribution of SLQ (Algorithm 4) where $\mathbf{e}_1 \in \mathbb{R}^m$ is the first standard basis vector and $\mathbf{v}_i$ is the i^{th} eigenvector of $\mathbf{Q}^T \mathbf{A} \mathbf{Q}$. Then, for $\delta \in (0, 1)$ such that $n \geq \Omega(\log(1/\delta))$, for all $i \in [k]$, with probability at least $1 - \delta$, we have:

$$w_i^2 \leq \frac{C \sqrt{\log(k/\delta)}}{n},$$

where $C > 0$ is a large constant.

Proof. Let $\mathbf{z}_i = \mathbf{Q}\mathbf{v}_i$ where $\mathbf{Q}$ is the orthonormal basis of the Krylov subspace generated by the Lanczos algorithm (Algorithm 3) after m iterations. Then, $\mathbf{z}_i$ are the eigenvectors of $\mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{Q}\mathbf{Q}^T$. Also, $\mathbf{g} = \mathbf{Q}\mathbf{e}_1$ or $\mathbf{Q}^T\mathbf{g} = \mathbf{e}_1$. From Lemma 15, we have that for every $i \in [k]$, $\mathbf{z}_i = \mathbf{u}_i + \mathbf{b}_i$ where $\mathbf{u}_i$ is the i^{th} eigenvector of $\mathbf{A}$ and

$$\|\mathbf{b}_i\|_2 \leq \frac{g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}},$$

for some constant $c > 0$, with probability at least $1 - \delta$. Also, we have $w_i = \mathbf{v}_i^T \mathbf{e}_1 = (\mathbf{Q}\mathbf{v}_i)^T (\mathbf{Q}\mathbf{e}_1) = \mathbf{z}_i^T \mathbf{g} = (\mathbf{u}_i + \mathbf{b}_i)^T \mathbf{g} = \mathbf{u}_i^T \mathbf{g} + \mathbf{b}_i^T \mathbf{g}$. $w_i^2 = (\mathbf{u}_i^T \mathbf{g})^2 + (\mathbf{b}_i^T \mathbf{g})^2 + 2(\mathbf{u}_i^T \mathbf{g})(\mathbf{b}_i^T \mathbf{g})$

Since $\mathbf{g}$ is sampled from a uniform distribution on the unit sphere, $\mathbf{g} \stackrel{d}{=} \frac{\mathbf{y}}{\sqrt{\sum_{j=1}^n \mathbf{y}_j^2}}$ where $\mathbf{y} \in \mathbb{R}^n$ is a random vector such that $\mathbf{y}_i \sim \mathcal{N}(0, 1)$ for $i \in [n]$. So, we have:

$$\begin{aligned} w_i^2 &= (\mathbf{u}_i^T \mathbf{g})^2 + (\mathbf{b}_i^T \mathbf{g})^2 + 2(\mathbf{u}_i^T \mathbf{g})(\mathbf{b}_i^T \mathbf{g}) \\ &\stackrel{d}{=} \frac{(\sum_{j=1}^n \mathbf{u}_{ij} \mathbf{y}_j)^2}{\sum_{r=1}^n \mathbf{y}_r^2} + \frac{(\sum_{j=1}^n \mathbf{b}_{ij} \mathbf{y}_j)^2}{\sum_{r=1}^n \mathbf{y}_r^2} + 2 \frac{(\sum_{j=1}^n \mathbf{u}_{ij} \mathbf{y}_j)(\sum_{j=1}^n \mathbf{b}_{ij} \mathbf{y}_j)}{\sum_{r=1}^n \mathbf{y}_r^2}. \end{aligned} \quad (78)$$

We will now bound $\sum_{r=1}^n \mathbf{y}_r^2$, $\sum_{j=1}^n \mathbf{u}_{ij} \mathbf{y}_j$ and $\sum_{j=1}^n \mathbf{b}_{ij} \mathbf{y}_j$ individually. Note that $\sum_{r=1}^n \mathbf{y}_r^2$ is a Chi-squared random variable with n degrees of freedom. Then, using well-known tail bounds for Chi-squared variables (see 2.21 [Wai19b]) we have that with probability at least $1 - \delta$, $|\sum_{r=1}^n \mathbf{y}_r^2 - n| \leq 2\sqrt{2n \log \frac{2}{\delta}}$. This gives us $\sum_{r=1}^n \mathbf{y}_r^2 \geq \frac{n}{2}$ (for $\delta \in (\Omega(1/e^n), 1)$). Next, observe that since $\mathbf{u}_i^T \mathbf{y} = \sum_{j=1}^n \mathbf{u}_{ij} \mathbf{y}_j$ is a linear combination of $\mathcal{N}(0, 1)$ random variables and $\|\mathbf{u}_i\|_2 = 1$, $\mathbf{u}_i^T \mathbf{y}$ is another $\mathcal{N}(0, 1)$ random variable. So, $(\mathbf{u}_i^T \mathbf{y})^2$ is a Chi-squared random variable and using the Chi-squared tail bound again, we have $(\mathbf{u}_i^T \mathbf{y})^2 \leq 1 + 2\sqrt{2 \log \frac{2}{\delta}} \leq 3\sqrt{2 \log \frac{2}{\delta}}$ with probability at least $1 - \delta$. So, we have $\frac{(\sum_{j=1}^n \mathbf{u}_{ij} \mathbf{y}_j)^2}{\sum_{r=1}^n \mathbf{y}_r^2} \leq \frac{6\sqrt{2 \log(2/\delta)}}{n}$ with probability at least $1 - 2\delta$. Next, using Cauchy Schwartz inequality, $\sum_{j=1}^n \mathbf{b}_{ij} \mathbf{y}_j \leq \|\mathbf{b}_i\|_2 \|\mathbf{y}\|_2$. So, $\frac{(\sum_{j=1}^n \mathbf{b}_{ij} \mathbf{y}_j)^2}{\sum_{r=1}^n \mathbf{y}_r^2} \leq \|\mathbf{b}_i\|_2^2 \leq \frac{g_{\min}^{2cl}}{(n \cdot \kappa)^{2c/\epsilon}}$. Also, $2 \frac{(\sum_{j=1}^n \mathbf{u}_{ij} \mathbf{y}_j)(\sum_{j=1}^n \mathbf{b}_{ij} \mathbf{y}_j)}{\sum_{r=1}^n \mathbf{y}_r^2} \leq 2\|\mathbf{b}_i\|_2 \leq \frac{2g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}}$. So, taking a union bound over all these events from (78) for all $i \in [k]$, and adjusting δ by a factor of $2k$, we get that with probability at least $1 - \delta$, for every $i \in [k]$:

$$w_i^2 \leq \frac{6\sqrt{2 \log(2k/\delta)}}{n} + \frac{3g_{\min}^{cl}}{(n \cdot \kappa)^{c/\epsilon}} \leq O\left(\frac{\sqrt{\log(k/\delta)}}{n}\right).$$

□

We are now ready to prove our final theorem in this section which bounds the Wasserstein-1 error between $s_{\mathbf{A}}$ and f . Let $\alpha = \max(\sigma_{l+1}(\mathbf{A}), \frac{\|\mathbf{A}\|_{2g_{\min}^{c/4}}}{n^{c/4}})$ for some constant $c > 0$. We will now consider two cases in the proof: when there exists some $k \in [l]$ such that $\sigma_k(\mathbf{A})$ is at least a constant multiplicative factor larger than α and when there isn't any such k . When there is such a k , we will use the bounds from Section 4.2 and Lemma 21 to show that the Wasserstein-1 error of the parts of $s_{\mathbf{A}}$ and f defined by the large magnitude eigenvalues of $\mathbf{A}$ and $\mathbf{T}$ is at most $\tilde{O}(l\|\mathbf{A}\|_2/n)$. Then, we will use Lemma 19 to show that the Chebyshev moments of the spectral density defined by the small magnitude eigenvalues ($\leq \sigma_{k+1}(\mathbf{A})$) of $\mathbf{A}$ and $\mathbf{T}$ are approximately equal. So, the Wasserstein-1 error of the parts of $s_{\mathbf{A}}$ and f defined by the small eigenvalues is bounded by $\epsilon \sigma_{k+1}(\mathbf{A})$. When

there isn't such a k , all singular values of $\mathbf{A}$ are small ($< 2\alpha$), and we can use the bounds using moment matching in Section 4.1 to bound the error by $\epsilon\alpha$ in this case. Before proving the theorem, we state a couple of results which we will use in our proof. We first state a result from [BKM22] on uniform approximation of Lipschitz continuous functions by a Chebyshev series:

Fact 2 (Fact 3.2 of [BKM22]). *Let $f : [-1, 1] \rightarrow \mathbb{R}$ be a Lipschitz continuous function with Lipschitz constant $\lambda > 0$. Then, for every $N \in 4\mathbb{N}^+$, there exists $N + 1$ constants $0 \leq b_N < \dots < b_0 = 1$ such that the polynomial $\bar{f}_N = \sum_{k=0}^N b_k \langle f, w \cdot \bar{T}_k \rangle \bar{T}_k$ has the property that $\max_{x \in [-1, 1]} |f(x) - \bar{f}_N(x)| \leq \frac{18\lambda}{N}$.*

We now state another result from [BKM22] bounding the magnitude of the inner-product of a Lipschitz function f with the k -th normalized Chebyshev polynomial $\bar{T}_k$:

Fact 3 (Fact 3.3 of [BKM22]). *Let $f : [-1, 1] \rightarrow \mathbb{R}$ be a Lipschitz continuous function with Lipschitz constant $\lambda > 0$. Then, for any $k \geq 0$, we have that $|\langle f, w \cdot \bar{T}_k \rangle| = |\int_{-1}^1 f(x) \bar{T}_k(x) w(x) dx| \leq \frac{2\lambda}{k}$.*

Theorem 3. *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and consider any $l \in [n]$ and $\epsilon, \delta \in (0, 1)$. Let $g_{\min} = \min_{i \in [l]} \frac{\sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A})}{\sigma_i(\mathbf{A})}$ and $\kappa = \frac{\|\mathbf{A}\|_2}{2\alpha}$. Let $\alpha = \max \left(\sigma_{l+1}(\mathbf{A}), \frac{\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}} \right)$ for some constant $c > 0$. Algorithm 4 (SLQ) run for $m = O(l \log \frac{1}{g_{\min}} + \frac{1}{\epsilon} \log \frac{n \cdot \kappa}{\delta})$ iterations performs m matrix vector products with $\mathbf{A}$ and outputs a probability density function f such that, with probability at least $1 - \delta$, for a fixed constant C ,*

$$W_1(s_{\mathbf{A}}, f) \leq \epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \frac{C \log(n/\epsilon) \log(1/\epsilon)}{\sqrt{n}} \cdot \sigma_{l+1}(\mathbf{A}) + \frac{Cl \log(1/\epsilon) \sqrt{\log(l/\delta)}}{n} \|\mathbf{A}\|_2.$$

Proof. We will prove the theorem for two complementary cases and analyze them separately below:

Case 1: Let there be some $k \in [l]$ such that $\sigma_k(\mathbf{A}) \geq 2\alpha$ and $\sigma_{k+1}(\mathbf{A}) < 2\alpha$. Assume that we run SLQ (Algorithm 4) with input $\mathbf{B} = \frac{1}{2\alpha} \mathbf{A}$ instead of $\mathbf{A}$ for $m = O(l \log \frac{1}{g_{\min}} + \frac{1}{\epsilon} \log \frac{n \cdot \kappa}{\delta})$ iterations. The output of Lanczos (Algorithm 3) will be the scaled tridiagonal matrix $\mathbf{T}_1 = \frac{1}{2\alpha} \mathbf{T}$ after m iterations. Observe that $\mathbf{B}$ satisfies all the conditions of Lemma 19 since $\sigma_k(\mathbf{B}) = \frac{\sigma_k(\mathbf{A})}{2\alpha} \geq 1$, $\sigma_{k+1}(\mathbf{B}) \leq 1$, $\|\mathbf{B}\|_2 = \frac{\|\mathbf{A}\|_2}{2\alpha} \geq 1$ and $\kappa_{\mathbf{B}} = \|\mathbf{B}\|_2$.

Let the output of Algorithm 4 with $\mathbf{B}$ as input be $f_{\mathbf{B}}(x) = \sum_{j=1}^m w_j^2 \delta(x - \lambda_j(\mathbf{T}_1)) = \sum_{j=1}^m w_j^2 \delta \left(x - \frac{\lambda_j(\mathbf{T})}{\sigma_{l+1}(\mathbf{A})} \right) = f \left(\frac{x}{\sigma_{l+1}(\mathbf{A})} \right)$ where $w_j = \mathbf{v}_j^T \mathbf{e}_1$ (recall that $\mathbf{v}_j$ is the eigenvector of $\mathbf{T}$ corresponding to $\lambda_j(\mathbf{T})$). Also, the spectral density of $\mathbf{B}$ is given by $s_{\mathbf{B}}(x) = \frac{1}{n} \sum_{j=1}^m \delta(x - \lambda_j(\mathbf{T}_1)) = \frac{1}{m} \sum_{j=1}^m \delta \left(x - \frac{\lambda_j(\mathbf{T})}{\sigma_{l+1}(\mathbf{A})} \right) = s_{\mathbf{A}} \left(\frac{x}{\sigma_{l+1}(\mathbf{A})} \right)$. Thus, we must have:

$$W_1(s, f) \leq 2\alpha \cdot W_1(s_{\mathbf{B}}, f_{\mathbf{B}}). \quad (79)$$

Let $L = \|\mathbf{B}\|_2 = \frac{1}{2\alpha} \|\mathbf{A}\|_2$. Then, the spectrum of both $\mathbf{B}$ and $\mathbf{T}_1$ are in $[-L, L]$ (as $|\lambda_i(\mathbf{T}_1)| \leq |\lambda_i(\mathbf{B})|$)

for all $i \in [m]$ by the minimax principle). So, we have:

$$\begin{aligned}
W_1(s_{\mathbf{B}}, f_{\mathbf{B}}) &= \sup_{h \in 1\text{-Lip}} \int_{-L}^L h(x) (s_{\mathbf{B}}(x) - f_{\mathbf{B}}(x)) dx \\
&= \int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^n \delta(x - \lambda_i(\mathbf{B})) - \sum_{i=1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx \\
&= \underbrace{\int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B})) - \sum_{i=1}^k w_i^2 \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx}_{I_1} \\
&\quad + \underbrace{\int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=k+1}^n \delta(x - \lambda_i(\mathbf{B})) - \sum_{i=k+1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx}_{I_2}. \tag{80}
\end{aligned}$$

Here, $h^*(x)$ is a 1-Lipschitz function that maximizes the integral for computing the Wasserstein distance. We will bound I_1 and I_2 separately.

Note that $\sigma_i(\mathbf{B}) \leq \sigma_{k+1}(\mathbf{B}) \leq 1$ for $i \in \{k+1, \dots, n\}$ and $\sigma_i(\mathbf{T}_1) \leq \sigma_i(\mathbf{B}) \leq 1$ for $i \in \{k+1, \dots, m\}$ by the minimax principle. So, the support of $\frac{1}{n} \sum_{i=k+1}^n \delta(x - \lambda_i(\mathbf{B}))$ and $\sum_{i=k+1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))$ is in $[-1, 1]$ and we can write I_2 as follows:

$$I_2 = \int_{-1}^1 h^*(x) \left(\frac{1}{n} \sum_{i=k+1}^n \delta(x - \lambda_i(\mathbf{B})) - \sum_{i=k+1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx = \int_{-1}^1 h^*(x) (r_1(x) - r_2(x)) dx,$$

where $r_1(x) = \sum_{i=k+1}^n \frac{1}{n} \delta(x - \lambda_i(\mathbf{B}))$ and $r_2(x) = \sum_{i=k+1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))$.

Let $N = O(1/\epsilon)$. We will now bound this by following the Chebyshev moment matching proof of Lemma 3.1 in [BKM22] where we match N normalized Chebyshev moments of $r_1(x)$ and $r_2(x)$. Let $\bar{h}_N(x) = \sum_{i=0}^N b_i \langle h^*, w, \bar{T}_i \rangle \bar{T}_i$ (where $w(x) = \frac{1}{\sqrt{1-x^2}}$) be the function from Fact 2 for constants $0 \leq b_N \dots \leq b_0 = 1$ such that $\max_{x \in [-1, 1]} |h^*(x) - \bar{h}_N(x)| \leq \frac{18}{N}$. Then, using triangle inequality, the integral can be upper bounded:

$$I_2 \leq \underbrace{\int_{-1}^1 |h^*(x) - \bar{h}_N(x)| (r_1(x) - r_2(x)) dx}_{t_1} + \underbrace{\int_{-1}^1 \bar{h}_N(x) (r_1(x) - r_2(x)) dx}_{t_2}. \tag{81}$$

Since $\max_{x \in [-1, 1]} |h^*(x) - \bar{h}_N(x)| \leq \frac{18}{N}$, $\int_{-1}^1 r_1(x) = \frac{n-k}{n}$ and $\int_{-1}^1 r_2(x) = \sum_{i=k+1}^m w_i^2$, we have

$$t_1 \leq \frac{18}{N} \left(\frac{n-k}{n} + \sum_{i=k+1}^m w_i^2 \right) \leq \frac{36}{N},$$

where we used the fact that $\sum_{i=k+1}^m w_i^2 \leq 1$. Next, we bound t_2 using the Chebyshev series expansion of $\bar{h}_N(x)$. First note that $r_1(x) - r_2(x) \in [-1, 1]$ and so its Chebyshev series expansion

is $\sum_{i=0}^{\infty} \langle r_1 - r_2, \bar{T}_i \rangle \bar{T}_i$. This gives us:

$$\begin{aligned} t_2 &= \int_{-1}^1 \bar{h}_N(x) w(x) \frac{r_1(x) - r_2(x)}{w(x)} dx = \int_{-1}^1 \bar{h}_N(x) w(x) \cdot \sum_{i=0}^{\infty} \langle r_1 - r_2, \bar{T}_i \rangle \bar{T}_i dx \\ &= \int_{-1}^1 w(x) \left(\sum_{i=0}^N b_i \langle h^*, w \cdot \bar{T}_i \rangle \bar{T}_i \right) \left(\sum_{i=0}^{\infty} \langle r_1 - r_2, \bar{T}_i \rangle \bar{T}_i \right) dx. \end{aligned}$$

By the orthogonality of Chebyshev polynomials under weight $w(x)$ and since $\langle \bar{T}_k, w \bar{T}_k \rangle = 1$ for $k \in [N]$, we can bound t_2 as:

$$\begin{aligned} t_2 &\leq \sum_{i=0}^N \langle h^*, w \cdot \bar{T}_i \rangle \cdot (\langle r_1, \bar{T}_i \rangle - \langle r_2, \bar{T}_i \rangle) \\ &\leq \langle h^*, w \cdot \bar{T}_0 \rangle \cdot (\langle r_1, \bar{T}_0 \rangle - \langle r_2, \bar{T}_0 \rangle) + \sum_{i=1}^N |\langle h^*, w \cdot \bar{T}_i \rangle| \cdot |\langle r_1, \bar{T}_i \rangle - \langle r_2, \bar{T}_i \rangle| \\ &\leq \frac{\langle h^*, w \cdot \bar{T}_0 \rangle}{\sqrt{\pi}} \cdot \left(\frac{n-k}{n} - \sum_{j=k+1}^m w_j^2 \right) + \sum_{i=1}^N \frac{2}{i} \cdot |\langle r_1, \bar{T}_i \rangle - \langle r_2, \bar{T}_i \rangle| \\ &\leq \frac{\langle h^*, w \cdot \bar{T}_0 \rangle}{\sqrt{\pi}} \cdot \left(\frac{n-k}{n} - \sum_{j=k+1}^m w_j^2 \right) + \sum_{i=1}^N \frac{2}{i} \cdot \left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right|, \quad (82) \end{aligned}$$

where for the first step, we used the fact that $\frac{\hat{b}_N(i)}{\hat{b}_N(0)} \leq 1$, and $\int_{-1}^1 \langle r_1 - r_2, \bar{T}_i \rangle = \langle r_1, \bar{T}_i \rangle - \langle r_2, \bar{T}_i \rangle$ for $i \in [N]$. For the second step, we bound the sum from 1 to N by its absolute value. For the third inequality, we used the fact that $\bar{T}_0(x) = \frac{1}{\sqrt{\pi}}$ and so $\langle r_1, \bar{T}_0 \rangle = \frac{n-k}{n\sqrt{\pi}}$ and $\langle r_2, \bar{T}_0 \rangle = \frac{\sum_{j=k+1}^m w_j^2}{\sqrt{\pi}}$ for the first term and use Fact 3 which gives us $|\langle h^*, w \cdot \bar{T}_i \rangle| \leq \frac{2}{i}$ for $i \in [N]$ for the second term. For the final step, we replace $\langle r_1, \bar{T}_i \rangle$ and $\langle r_2, \bar{T}_i \rangle$ by evaluating the integrals. Let $t_i(x)$ be the polynomial defined in Lemma 19. Then, using triangle inequality, for constants $c_9, c_{10} > 0$ and $C > 0$, we get for any $i \in [N]$:

$$\begin{aligned} \left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| &\leq \left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \langle t_i, \mathbf{s}_\mathbf{B} \rangle \right| \\ &\quad + \left| \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) - \langle t_i, \mathbf{f}_\mathbf{B} \rangle \right| + \left| \langle t_i, \mathbf{f}_\mathbf{B} \rangle - \langle t_i, \mathbf{s}_\mathbf{B} \rangle \right| \\ &\leq \frac{g_{\min}^{c_9 l}}{(n \cdot \kappa)^{c_9/\epsilon}} + \frac{g_{\min}^{c_{10} l}}{(n \cdot \kappa)^{c_{10}/\epsilon}} \\ &\quad + \frac{C \log(N/\delta)}{\sqrt{n}}. \quad (83) \end{aligned}$$

To bound $\left| \langle t_i, \mathbf{f}_\mathbf{B} \rangle - \langle t_i, \mathbf{s}_\mathbf{B} \rangle \right|$, the final step uses Lemma 20 and the fact that $\kappa_\mathbf{B} = \|\mathbf{B}\|_2 = \frac{\|\mathbf{A}\|_2}{2\alpha} = \kappa$.

Putting together the bounds on t_1 and t_2 from (82) and (83), and bounding $\sum_{i=1}^N \frac{2}{i} \leq \log N$, we

get from (81) (for some constant c_{11}):

$$I_2 \leq \frac{36}{N} + \frac{\langle h^*, w, \bar{T}_0 \rangle}{\sqrt{\pi}} \cdot \left(\frac{n-k}{n} - \sum_{j=k+1}^m w_j^2 \right) + \frac{2g_{\min}^{c_{11}l} \log N}{(n \cdot \kappa)^{c_{11}/\epsilon}} + \frac{C \log(N/\delta) \log N}{\sqrt{n}}. \quad (84)$$

We now bound I_1 . We have the following:

$$\begin{aligned} I_1 &= \int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B})) - \sum_{i=1}^k w_i^2 \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx \\ &= \frac{k}{n} \int_{-L}^L h^*(x) \left(\frac{1}{k} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B})) - \frac{1}{k} \sum_{i=1}^k \delta(x) \right) dx \\ &\quad + \sum_{j=1}^k w_j^2 \int_{-L}^L h^*(x) \left(\frac{\sum_{i=1}^k w_i^2 \delta(x)}{\sum_{j=1}^k w_j^2} - \frac{\sum_{i=1}^k w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))}{\sum_{j=1}^k w_j^2} \right) dx \\ &\quad + \int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^k \delta(x) - \sum_{i=1}^k w_i^2 \delta(x) \right) dx. \end{aligned} \quad (85)$$

We now bound the three terms above. First observe that $\frac{1}{k} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B}))$ and $\frac{1}{k} \sum_{i=1}^k \delta(x)$ are density functions of distributions (defined by dirac deltas at $\lambda_1(\mathbf{B}), \dots, \lambda_k(\mathbf{B})$ with weights $\frac{1}{n}$ and at 0 with weight 1 respectively). So, $\sup_{h \in 1\text{-Lip}} \int_{-L}^L h^*(x) \left(\frac{1}{k} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B})) - \frac{1}{k} \sum_{i=1}^k \delta(x) \right)$ is the Wasserstein-1 distance between the distributions. Using the Earth mover's distance interpretation of Wasserstein-1 distance, we have:

$$\begin{aligned} \int_{-L}^L h^*(x) \left(\frac{1}{k} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B})) - \frac{1}{k} \sum_{i=1}^k \delta(x) \right) &\leq \sup_{h \in 1\text{-Lip}} \int_{-L}^L h(x) \left(\frac{1}{k} \sum_{i=1}^k \delta(x - \lambda_i(\mathbf{B})) - \frac{1}{k} \sum_{i=1}^k \delta(x) \right) \\ &\leq \frac{\sum_{i=1}^k |\lambda_i(\mathbf{B})|}{k} \leq \|\mathbf{B}\|_2. \end{aligned}$$

Similarly, by interpreting $\frac{\sum_{i=1}^k w_i^2 \delta(x)}{\sum_{j=1}^k w_j^2}$ and $\frac{\sum_{i=1}^k w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))}{\sum_{j=1}^k w_j^2}$ as probability densities and using the earth mover's distance interpretation of Wasserstein-1 distance we have:

$$\int_{-L}^L h^*(x) \left(\frac{\sum_{i=1}^k w_i^2 \delta(x)}{\sum_{j=1}^k w_j^2} - \frac{\sum_{i=1}^k w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))}{\sum_{j=1}^k w_j^2} \right) dx \leq \sum_{i=1}^k \frac{w_i^2 |\lambda_i(\mathbf{T}_1)|}{\sum_{j=1}^k w_j^2} \leq \sum_{i=1}^k \frac{w_i^2 \|\mathbf{T}_1\|_2}{\sum_{j=1}^k w_j^2}.$$

Finally, to bound $\int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^k \delta(x) - \sum_{i=1}^k w_i^2 \delta(x) \right) dx$, we again use the proof technique outlined in Lemma 3.1 of [BKM22] and for bounding I_2 here. Observe that similar to (81) we have:

$$\begin{aligned} &\int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^k \delta(x) - \sum_{i=1}^k w_i^2 \delta(x) \right) dx \\ &\leq \underbrace{\int_{-1}^1 (h^*(x) - \bar{h}_N(x)) \left(\frac{1}{n} \sum_{i=1}^k \delta(x) - \sum_{i=1}^k w_i^2 \delta(x) \right) dx}_{t_1} + \underbrace{\int_{-1}^1 \bar{h}_N(x) \left(\frac{1}{n} \sum_{i=1}^k \delta(x) - \sum_{i=1}^k w_i^2 \delta(x) \right) dx}_{t_2}. \end{aligned}$$

We can bound t_1 and t_2 the same way as in the case of I_2 to get $t_1 \leq \frac{36}{N}$ and (similar to (82))

$$\begin{aligned} t_2 &\leq \frac{\langle h^*, w \cdot \bar{T}_o \rangle}{\sqrt{\pi}} \left(\frac{k}{n} - \sum_{j=1}^k w_j^2 \right) + \sum_{i=1}^N \frac{2}{i} \cdot \left| \frac{1}{n} \sum_{j=1}^k \bar{T}_i(0) - \sum_{j=1}^k w_j^2 \bar{T}_i(0) \right| \\ &\leq \frac{\langle h^*, w \cdot \bar{T}_o \rangle}{\sqrt{\pi}} \left(\frac{k}{n} - \sum_{j=1}^k w_j^2 \right) + \sum_{i=1}^N \frac{2}{i} \left| \frac{k}{n} - \sum_{j=1}^k w_j^2 \right|. \end{aligned}$$

From Lemma 21, we have $\sum_{j=1}^k w_j^2 \leq \frac{Ck\sqrt{\log 2/\delta}}{n}$ for some large $C > 0$. So, using the upper bounds on all the terms on the right hand side of (85) (and using the fact that $\|\mathbf{T}_1\|_2 \leq \|\mathbf{B}\|_2$), we get that:

$$\begin{aligned} I_1 &\leq \frac{k}{n} \|\mathbf{B}\|_2 + \left(\sum_{i=1}^k w_i^2 \right) \|\mathbf{T}_1\|_2 + \frac{\langle h^*, w \cdot \bar{T}_o \rangle}{\sqrt{\pi}} \left(\frac{k}{n} - \sum_{j=1}^k w_j^2 \right) + \sum_{i=1}^N \frac{2}{i} \left| \frac{k}{n} - \sum_{j=1}^k w_j^2 \right| \\ &\leq \frac{k(1 + C\sqrt{\log k/\delta})}{n} \|\mathbf{B}\|_2 + \frac{\langle h^*, w \cdot \bar{T}_o \rangle}{\sqrt{\pi}} \left(\frac{k}{n} - \sum_{j=1}^k w_j^2 \right) + \frac{3Ck \log N \sqrt{\log k/\delta}}{n}, \end{aligned} \quad (86)$$

where in the last step we used triangle inequality to bound the final term. Finally, using the bounds on I_1 and I_2 from (86) and (84) respectively in (80) (and using the fact that $\sum_{j=1}^m w_j^2 = 1$), we get that:

$$\begin{aligned} W_1(s_{\mathbf{B}}, f_{\mathbf{B}}) &\leq \frac{36}{N} + \frac{k(1 + 4C \log N \sqrt{\log k/\delta})}{n} \|\mathbf{B}\|_2 + \frac{2g_{\min}^{c_{11}l} \log N}{(n \cdot \kappa)^{c_{11}/\epsilon}} + \frac{C \log(N/\delta) \log N}{\sqrt{n}} \\ &\leq \frac{36}{N} + \frac{5Ck \log N \sqrt{\log k/\delta}}{n} \|\mathbf{B}\|_2 + \frac{2C \log(N/\delta) \log N}{\sqrt{n}}, \end{aligned}$$

where we also use the fact that the constant c_{11} can be chosen to be large enough so that the third term in the first inequality can be made small enough. Recall that we set $N = O(\frac{1}{\epsilon})$. From (79), we also have $W_1(s, f) \leq 2\alpha W_1(s_{\mathbf{B}}, f_{\mathbf{B}})$. Thus, we finally get (using the fact $\mathbf{B} = \frac{\mathbf{A}}{2\alpha}$):

$$\begin{aligned} W_1(s, f) &\leq 72\epsilon\alpha + \frac{5Ck \log(1/\epsilon) \sqrt{\log k/\delta}}{n} \|\mathbf{A}\|_2 + \frac{4C \log(1/\epsilon\delta) \log(1/\epsilon)\alpha}{\sqrt{n}} \\ &\leq \frac{5Cl \log(1/\epsilon) \sqrt{\log(k/\delta)}}{n} \|\mathbf{A}\|_2 + \left(72\epsilon + \frac{4C \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \right) \sigma_{l+1}(\mathbf{A}) \\ &\leq \frac{5Cl \log(1/\epsilon) \sqrt{\log(l/\delta)}}{n} \|\mathbf{A}\|_2 + \left(72\epsilon + \frac{4C \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \right) \sigma_{l+1}(\mathbf{A}). \end{aligned}$$

where we also used the fact that $2\alpha < 2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}}$ and $k \leq l$. Also note that the term $\left(72\epsilon + \frac{4C \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \right) \frac{2\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}} < \frac{Cl \log(1/\epsilon) \sqrt{\log(l/\delta)}}{n} \|\mathbf{A}\|_2$ for a large c and hence is absorbed into that term by making the constant C larger. We get the final bound by adjusting ϵ by constant factors. This concludes the case when such a k exists.

Case 2: Now, when such a k doesn't exist, i.e. we have $\|\mathbf{A}\|_2 \leq 2\alpha < 2\sigma_{l+1}(\mathbf{A}) + \frac{2\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}}$, from Theorem 10 we directly get error $W_1(s_{\mathbf{A}}, f) \leq 2\epsilon\alpha + \frac{C \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} 2\alpha \leq \epsilon\sigma_{l+1}(\mathbf{A}) + \frac{2C \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \sigma_{l+1}(\mathbf{A}) + \frac{C\|\mathbf{A}\|_2}{n}$ after adjusting ϵ by constants.

Finally, observe that the condition $n \geq \Omega(\log(1/\delta))$ in Lemmas 19 and 12 must always be satisfied if we want a non-trivial ($\leq n$) number of matrix vector products with $\mathbf{A}$. Adjusting δ by some constant factors gives us the final bound. $\square$

4.4 Variance reduced SLQ

We introduce our variant of SLQ with better error guarantees here. The algorithm, which we call the Variance reduced SLQ is described in Algorithm 5. The main difference between this algorithm and Algorithm 4 is that here, we set the weights w_i^2 corresponding to the *converged* large magnitude eigenvalues of $\mathbf{T}$ (the output of Lanczos) to $\frac{1}{n}$ (and adjust the remaining weights so that the square of the weights sum to one). This is described in lines 4-9 of Algorithm 3. From Lemma 14, we know that in the presence of a constant multiplicative gap between $\sigma_k(\mathbf{A})$ and $\sigma_{l+1}(\mathbf{A})$, the top k eigenvalues of $\mathbf{T}$ will approximately be equal to those of $\mathbf{A}$. Hence the *correct* weights corresponding to these eigenvalues must be $\frac{1}{n}$. This lets us avoid the $\tilde{O}(\frac{l}{n}\|\mathbf{A}\|_2)$ error SLQ was incurring on the large magnitude eigenvalues. However, we do not actually know the value of k or even if there exists such a $k \in [l]$ with a constant multiplicative gap. Hence, we check two convergence conditions for each of the top l indices which will always be true for the top $k \leq l$ eigenvectors of $\mathbf{T}$ (the output of Lanczos) and corresponding weights $(\mathbf{v}_j^T \mathbf{e}_1)^2$ if such a $k \in [l]$ exists.

The two convergence conditions to find the indices for which we set the weights to $\frac{1}{n}$ are given in line 5 of Algorithm 5. The first condition ($\|\mathbf{A}\mathbf{Q}\mathbf{v}_j - \lambda_j(\mathbf{T})\mathbf{Q}\mathbf{v}_j\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{\beta/\epsilon}}$) checks if an eigenvector $\mathbf{v}_j$ of $\mathbf{T}$ (and its corresponding eigenvalue) is also approximately an eigenvector and eigenvalue of $\mathbf{A}$. Note that from Lemma 15 this condition will always be true for the top k eigenvalues and eigenvectors of $\mathbf{T}$ as long as there is at least a constant multiplicative factor gap between $\sigma_k(\mathbf{A})$ and $\sigma_{l+1}(\mathbf{A})$. This helps us set the correct weights of $\frac{1}{n}$ for the top k eigenvalues. However, there might be some eigenvectors of $\mathbf{T}$ outside of the top k indices for which this condition is also true. Unfortunately, we can't guarantee that the corresponding eigenvalues of $\mathbf{T}$ outside of the top k have converged to an eigenvalue of $\mathbf{A}$. The second condition ($(\mathbf{v}_j^T \mathbf{e}_1)^2 \leq \frac{\sqrt{\log(l/\delta)}}{n}$) ensures that in case this happens, we don't incur too much error by correcting the weights to $\frac{1}{n}$. Note that from Lemma 21, this condition will also be true for the top k weights of the SLQ distribution as long as there is at least a constant multiplicative factor gap between $\sigma_k(\mathbf{A})$ and $\sigma_{l+1}(\mathbf{A})$. But in case there is spurious convergence of an eigenvector of $\mathbf{T}$ without the corresponding eigenvalue of $\mathbf{T}$ converging to an eigenvalue of $\mathbf{A}$, this condition essentially bound the error we incur by correcting the weights by $\frac{\sqrt{\log(l/\delta)}}{n}$. Also, since the corresponding eigenvalue of $\mathbf{A}$ (and $\mathbf{T}$) is at most $\sigma_{k+1}(\mathbf{A})$, the total error we incur by incorrectly setting weights for these spurious eigenvectors is at most $\tilde{O}(\frac{l\sigma_{k+1}(\mathbf{A})}{n}) \leq \tilde{O}(\frac{l\sigma_{l+1}(\mathbf{A})}{n})$ as opposed to $\tilde{O}(\frac{l\|\mathbf{A}\|_2}{n})$ as the magnitude of these eigenvalues will always be less than $\sigma_{k+1}(\mathbf{A}) \leq O(\sigma_{l+1}(\mathbf{A}))$.

We now analyze the algorithm below. The proof is similar to that of Theorem 3 except now, we use the Backward Perturbation Bound (Lemma 7) to first show that there exists some matrix with the same converged eigenvalues and eigenvectors of $\mathbf{A}$ and which is spectrally close to $\mathbf{A}$ (as we had done in the proof of Theorem 1). This helps us bound the error incurred on the converged eigenvectors. For bounding the Wasserstein error of the spectral density corresponding to the eigenvalues with non-converged eigenvectors (which will also have small magnitude), we use the moment matching analysis of Theorem 3 again.

Algorithm 5 Variance reduced Stochastic Lanczos Quadrature

Input: Symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$, number of iterations $m(\leq n)$, convergence parameters $l \in [n]$, $\beta, C > 0$.⁶

- 1: Sample $\mathbf{g} \sim \mathcal{U}(\mathcal{S}^{n-1})$.
- 2: Run Lanczos (Algorithm 3) on $\mathbf{A}, \mathbf{g}$ for m steps to compute symmetric tridiagonal matrix $\mathbf{T} \in \mathbb{R}^{m \times m}$ and orthonormal basis $\mathbf{Q} \in \mathbb{R}^{n \times m}$ such that $\mathbf{T} = \mathbf{Q}^T \mathbf{A} \mathbf{Q}$ [GM09]. Let the eigenvectors of $\mathbf{T}$ be $\mathbf{v}_1, \dots, \mathbf{v}_m$.
- 3: Set $S = \{\}$
- 4: **for** $j = 1, \dots, l$ **do**
- 5: **if** $\|\mathbf{A} \mathbf{Q} \mathbf{v}_j - \lambda_j(\mathbf{T}) \mathbf{Q} \mathbf{v}_j\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^\beta}$ and $(\mathbf{v}_j^T \mathbf{e}_1)^2 \leq \frac{C \sqrt{\log(l/\delta)}}{n}$ ⁷ **then**
- 6: $S = S \cup \{j\}$
- 7: **end if**
- 8: **end for**
- 9: Set $f(x) = \sum_{j \in S} \frac{1}{n} \delta(x - \lambda_j(\mathbf{T})) + \left(\frac{1 - \frac{|S|}{n}}{\sum_{i \in [m] \setminus S} w_i^2} \right) \sum_{j \in [m] \setminus S} w_j^2 \delta(x - \lambda_j(\mathbf{T}))$, where $w_j = \mathbf{v}_j^T \mathbf{e}_1$.
- 10: **return** $f(x)$.

Theorem 4. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be symmetric and consider any $l \in [n]$ and $\epsilon, \delta \in (0, 1)$. Let $g_{\min} = \min_{i \in [l]} \frac{\sigma_i(\mathbf{A}) - \sigma_{i+1}(\mathbf{A})}{\sigma_i(\mathbf{A})}$ and $\kappa = \frac{\|\mathbf{A}\|_2}{2\alpha}$. Let $\alpha = \max \left(\sigma_{l+1}(\mathbf{A}), \frac{\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}} \right)$ for some constant $c > 0$.

Algorithm 5 run for $m = O(l \log \frac{1}{g_{\min}} + \frac{\log(n \cdot \kappa)}{\epsilon})$ iterations, performs m matrix vector products with $\mathbf{A}$ and outputs a distribution f such that, with probability at least $1 - \delta$, for large fixed constants $C > 0$ and $c_2 > 0$,

$$W_1(s_{\mathbf{A}}, f) \leq \epsilon \cdot \sigma_{l+1}(\mathbf{A}) + \left(\frac{C \log(n/\epsilon) \log(1/\epsilon)}{\sqrt{n}} + \frac{Cl \log(1/\epsilon) \sqrt{\log(l/\delta)}}{n} \right) \cdot \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2}}.$$

Proof. Throughout the proof, for ease of writing, we will abuse notation slightly and assume that the first convergence condition in line 5 of Algorithm 5 is given by $\|\mathbf{A} \mathbf{Q} \mathbf{v}_j - \lambda_j(\mathbf{T}) \mathbf{Q} \mathbf{v}_j\|_2 \leq \frac{\|\mathbf{A}\|_2}{n^{\beta/\epsilon}}$ i.e. the first parameter is β/ϵ instead of just β . Let $k \in [l]$ be such that $\sigma_k(\mathbf{A}) \geq 2\alpha$ and $\sigma_{k+1}(\mathbf{A}) < 2\alpha$. We will again consider two cases as in Theorem 3, i.e. when such a k exists and when it doesn't.

Case 1: We first consider the case when such a k exists. In this case, similar to proof of case 1 of Theorem 10, let $\mathbf{B} = \frac{1}{2\alpha} \mathbf{A}$ and let $\mathbf{T}_1 = \frac{1}{2\alpha} \mathbf{T}$ where $\mathbf{T}_1$ is the output of running Lanczos with $\mathbf{B}$ as input. Then, as in (79), we have:

$$W_1(s_{\mathbf{A}}, f) \leq 2\alpha W_1(s_{\mathbf{B}}, f_{\mathbf{B}}), \quad (87)$$

where $s_{\mathbf{B}}(x) = \sum_{j=1}^m \frac{1}{m} \delta(x - \lambda_j(\mathbf{B}))$ and $f_{\mathbf{B}}(x) = \sum_{j \in S} \frac{1}{n} \delta(x - \lambda_j(\mathbf{T}_1)) + \left(\frac{1 - \frac{|S|}{n}}{\sum_{i \in [m] \setminus S} w_i^2} \right) \sum_{j \in [m] \setminus S} w_j^2 \delta(x - \lambda_j(\mathbf{T}_1))$, where $w_j = \mathbf{v}_j^T \mathbf{e}_1$ for each $j \in [m]$ and S is the set of indices containing the *converged* eigenvectors as defined in Algorithm 5. Let $\mathbf{V}_S \in \mathbb{R}^{m \times |S|}$ contain the set of all eigenvectors $\mathbf{v}_i$ of $\mathbf{T}_1$ such that $i \in S$. Let $\mathbf{Z}_S = \mathbf{Q} \mathbf{V}_S$ where $\mathbf{Q}$ is the

⁶refer to Theorem 4 for instructions on setting these parameters.

⁷ $\delta \in (0, 1)$ is the failure probability.

orthonormal basis of the Krylov subspace generated by the Lanczos algorithm (Algorithm 3). Then, $\|\mathbf{B}\mathbf{Z}_S - \mathbf{Z}_S\tilde{\mathbf{A}}\|_F = \sqrt{\sum_{i \in S} \|\mathbf{B}\mathbf{Q}\mathbf{v}_i - \lambda_i(\mathbf{A})\mathbf{Q}\mathbf{v}_i\|_2^2} \leq \frac{\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}$ where $\tilde{\mathbf{A}}$ is a matrix with the eigenvalues of $\mathbf{T}_1$ corresponding to the eigenvectors of $\mathbf{V}_S$ on its diagonal and zeros everywhere else. Then, using the backwards error bound (Lemma 7) we get that there exists a matrix $\mathbf{C}$ such that $\mathbf{C}\mathbf{Z}_S = \mathbf{Z}_S\tilde{\mathbf{A}}$ and $\|\mathbf{B} - \mathbf{C}\|_2 \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}$. Using Weyls' inequality (Fact 1), we get that for all $i \in [n]$

$$|\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{C})| \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}. \quad (88)$$

Then, $\frac{1}{n} \sum_{i=1}^n |\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{C})| \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}$ which implies that

$$\mathbf{W}_1(s_{\mathbf{B}}, s_{\mathbf{C}}) \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}. \quad (89)$$

where $s_{\mathbf{C}}(x)$ is the spectral density function of $\mathbf{C}$. Let $S_1 \subseteq [n]$ such $|S_1| = |S|$ and for every $j \in S_1$, there exists an $i \in S$ such that $\lambda_j(\mathbf{C}) = \lambda_i(\mathbf{T}_1)$. We know such a set S_1 must exist as $\tilde{\mathbf{A}}$ are eigenvalues of $\mathbf{C}$. Let $L = \max(\|\mathbf{C}\|_2, \|\mathbf{B}\|_2)$. Then we have:

$$\begin{aligned} \mathbf{W}_1(s_{\mathbf{C}}, f_{\mathbf{B}}) &= \sup_{h \in 1\text{-Lip}} \int_{-L}^L h(x)(s_{\mathbf{C}}(x) - f_{\mathbf{B}}(x))dx \\ &= \int_{-L}^L h^*(x) \left(\frac{1}{n} \sum_{i=1}^n \delta(x - \lambda_i(\mathbf{C})) - \sum_{j \in S} \frac{1}{n} \delta(x - \lambda_j(\mathbf{T}_1)) \right. \\ &\quad \left. - \left(\frac{1 - \frac{|S|}{n}}{\sum_{i \in [m] \setminus S} w_i^2} \right) \sum_{j \in [m] \setminus S} w_j^2 \delta(x - \lambda_j(\mathbf{T}_1)) \right) dx \\ &= \underbrace{\frac{|S|}{n} \int_{-L}^L h^*(x) \left(\frac{1}{|S|} \sum_{i \in S_1} \delta(x - \lambda_i(\mathbf{C})) - \frac{1}{|S|} \sum_{i \in S} \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx}_{I_1} \\ &\quad + \underbrace{\frac{n - |S|}{n} \int_{-L}^L h^*(x) \left(\frac{1}{n - |S|} \sum_{i \in [n] \setminus S_1} \delta(x - \lambda_i(\mathbf{C})) - \frac{1}{\sum_{j \in [m] \setminus S} w_j^2} \sum_{i \in [m] \setminus S} w_i^2 \delta(x - \lambda_i(\mathbf{T}_1)) \right) dx}_{I_2}. \end{aligned} \quad (90)$$

Here, in the second step, $h^*(x)$ is the function that maximizes the integral in the first step. By definition of the set S_1 we have $I_1 = 0$. Now we bound I_2 . Let $s_1(x) = \frac{1}{n - |S|} \sum_{i \in [n] \setminus S_1} \delta(x - \lambda_i(\mathbf{C}))$, $s_2(x) = \frac{1}{n - |S|} \sum_{i \in [n] \setminus S_1} \delta(x - \lambda_i(\mathbf{B}))$ and $f_1(x) = \frac{1}{\sum_{j \in [m] \setminus S} w_j^2} \sum_{i \in [m] \setminus S} w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))$ be three density functions. Using triangle inequality, and the fact that $h^*(x)$ is a 1-Lipschitz function, we have:

$$\begin{aligned} I_2 &= \frac{n - |S|}{n} \int_{-L}^L h^*(x)(s_1(x) - f_1(x))dx \\ &\leq \frac{n - |S|}{n} \mathbf{W}_1(s_1, f_1) \leq \frac{n - |S|}{n} (\mathbf{W}_1(s_1, s_2) + \mathbf{W}_1(s_2, f_1)). \end{aligned} \quad (91)$$

Using (88) and the earth movers' interpretation of the 1-Wasserstein distance, we can bound $W_1(s_1, s_2)$ as

$$W_1(s_1, s_2) \leq \frac{1}{n - |S|} \sum_{i \in [n] \setminus S_1} |\lambda_i(\mathbf{C}) - \lambda_i(\mathbf{B})| \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}. \quad (92)$$

We now bound $W_1(s_2, f_1)$. Let $\mathbf{z}_i = \mathbf{Q}\mathbf{v}_i$. Recall that $\mathbf{v}_i$ are eigenvectors of $\mathbf{Q}^T\mathbf{B}\mathbf{Q}$ for eigenvalues $\lambda_i(\mathbf{Q}^T\mathbf{B}\mathbf{Q})$, $\mathbf{z}_i$ are the eigenvectors of $\mathbf{Q}\mathbf{Q}^T\mathbf{B}\mathbf{Q}\mathbf{Q}^T$ corresponding to eigenvalues $\lambda_i(\mathbf{Q}^T\mathbf{B}\mathbf{Q}) = \lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{B}\mathbf{Q}\mathbf{Q}^T)$. From Lemma 15, we have that for every $i \in [k]$, either $\|\mathbf{z}_i - \mathbf{u}_i\|_2 \leq \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/2\epsilon-1}}$ or $\|\mathbf{z}_i + \mathbf{u}_i\|_2 \leq \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/2\epsilon-1}}$. From Lemma 14, we have that for $i \in [k]$,

$$|\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{B}\mathbf{Q}\mathbf{Q}^T)| \leq \frac{g_{\min}^{cl} \|\mathbf{B}\|_2}{(n \cdot \kappa)^{c/\epsilon}}.$$

Let us assume we have $\|\mathbf{z}_i - \mathbf{u}_i\|_2 \leq \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/2\epsilon-1}}$ for some $i \in [k]$. Using triangle inequality and spectral submultiplicativity, we get that for $i \in [k]$:

$$\begin{aligned} \|\mathbf{B}\mathbf{z}_i - \lambda_i(\mathbf{Q}^T\mathbf{B}\mathbf{Q})\mathbf{z}_i\| &\leq \|\mathbf{B}(\mathbf{z}_i - \mathbf{u}_i)\|_2 + \|\mathbf{B}\mathbf{u}_i - \lambda_i(\mathbf{Q}^T\mathbf{B}\mathbf{Q})\mathbf{u}_i\|_2 + \|\lambda_i(\mathbf{Q}^T\mathbf{B}\mathbf{Q})(\mathbf{u}_i - \mathbf{z}_i)\|_2 \\ &\leq \frac{g_{\min}^{cl/4} \|\mathbf{B}\|_2}{(n \cdot \kappa)^{c/2\epsilon-1}} + \|(\lambda_i(\mathbf{B}) - \lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{B}\mathbf{Q}\mathbf{Q}^T))\mathbf{u}_i\|_2 + \frac{g_{\min}^{cl/4} \|\mathbf{B}\|_2}{(n \cdot \kappa)^{c/2\epsilon-1}} \\ &\leq \frac{3g_{\min}^{cl/4} \|\mathbf{B}\|_2}{(n \cdot \kappa)^{c/2\epsilon-3}}. \end{aligned}$$

For the second step, we also used the fact that $\lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{B}\mathbf{Q}\mathbf{Q}^T) = \lambda_i(\mathbf{Q}^T\mathbf{B}\mathbf{Q})$. We can similarly prove that $\|\mathbf{B}\mathbf{z}_i - \lambda_i(\mathbf{Q}\mathbf{Q}^T\mathbf{B}\mathbf{Q}\mathbf{Q}^T)\mathbf{z}_i\| \leq \frac{3g_{\min}^{cl/4} \|\mathbf{B}\|_2}{(n \cdot \kappa)^{c/2\epsilon-3}}$ when $\|\mathbf{z}_i + \mathbf{u}_i\|_2 \leq \frac{g_{\min}^{cl/4}}{(n \cdot \kappa)^{c/2\epsilon-1}}$. Moreover, from Lemma 21, for any $i \in [k]$, we have $w_i^2 = (\mathbf{v}_i^T \mathbf{e}_1)^2 \leq \frac{C\sqrt{\log(k/\delta)}}{n} \leq \frac{C\sqrt{\log(l/\delta)}}{n}$ for some constant $C > 0$. Thus, if we set the constants β, C_1 in Algorithm 5 such that $\beta \leq \frac{\epsilon}{2} - 3\epsilon$, and $C_1 \geq C$, then $\mathbf{v}_1, \dots, \mathbf{v}_k$ must satisfy the conditions in line 5 that $\|\mathbf{B}\mathbf{Q}\mathbf{v}_j - \lambda_j(\mathbf{T}_1)\mathbf{Q}\mathbf{v}_j\|_2 \leq \frac{\|\mathbf{B}\|_2}{n^{\beta/\epsilon}}$ and $(\mathbf{v}_j^T \mathbf{e}_1)^2 \leq \frac{C_1\sqrt{\log(l/\delta)}}{n}$, i.e. $[k] \subseteq S$. Thus, $\max_{i \in [n] \setminus S} |\lambda_i(\mathbf{B})| \leq |\lambda_{k+1}(\mathbf{B})| \leq 1$. Also, by the minimax principle, $\max_{i \in [m] \setminus S} |\lambda_i(\mathbf{T}_1)| \leq |\lambda_{k+1}(\mathbf{T}_1)| \leq |\lambda_{k+1}(\mathbf{B})| \leq 1$. Thus, the support of $s_2(x)$ and $f_1(x)$ is in $[-1, 1]$. To bound $W_1(s_2, f_1)$, we use Lemma 3.1 of [BKM22] to get that for any $N \in 4\mathbb{N}^+$, we have:

$$W_1(s_2, f_1) \leq \frac{36}{N} + 2 \sum_{i=1}^N \frac{|\langle \bar{T}_i, s_2 \rangle - \langle \bar{T}_i, f_1 \rangle|}{i},$$

where $\bar{T}_i$ is the i th normalized Chebyshev polynomial. Then, we get:

$$\frac{n - |S|}{n} W_1(s_2, f_1) \leq \frac{n - |S|}{n} \cdot \frac{36}{N} + 2 \sum_{i=1}^N \frac{1}{i} \left| \frac{1}{n} \sum_{j \in [n] \setminus S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \frac{\frac{n - |S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right|. \quad (93)$$

We can set $N = O(\frac{1}{\epsilon})$ since the total number of iterations of Lanczos is $m = \tilde{O}(l + \frac{1}{\epsilon})$. Then, using triangle inequality, and the fact that $[k] \subseteq S$ and $[k] \subseteq S_1$, we get that for any $i \in [N]$:

$$\begin{aligned}
& \left| \frac{1}{n} \sum_{j \in [n] \setminus S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\
& \leq \left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| + \left| \frac{1}{n} \sum_{j \geq k+1, j \in S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j \geq k+1, j \in S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\
& + \left| \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) - \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right|. \tag{94}
\end{aligned}$$

We will now bound the three terms above individually. Let $f_2(x) = \sum_{i=1}^m w_i^2 \delta(x - \lambda_i(\mathbf{T}_1))$, i.e., the output of SLQ with $\mathbf{B}$ as input. Next, observe that $\mathbf{B}$ satisfies all the conditions of Lemma 19 since $\sigma_k(\mathbf{B}) = \frac{\sigma_k(\mathbf{A})}{2\alpha} \geq 1$ and $\sigma_{k+1}(\mathbf{B}) \leq 1$. Let $t_i(x)$ be the polynomial defined in Lemma 19 for $i \in [N]$. Then, for the first term, using triangle inequality, we have that:

$$\begin{aligned}
\left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| & \leq \left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \langle t_i, s_{\mathbf{B}} \rangle \right| \\
& + \left| \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) - \langle t_i, f_2 \rangle \right| + \left| \langle t_i, f_2 \rangle - \langle t_i, s_{\mathbf{B}} \rangle \right|.
\end{aligned}$$

From Lemma 19, we have that $\left| \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) - \langle t_i, f_2 \rangle \right| \leq \frac{g_{\min}^{c_3 l}}{(n \cdot \kappa)^{c_3/\epsilon}}$ and $\left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \langle t_i, s_{\mathbf{B}} \rangle \right| \leq \frac{g_{\min}^{c_4 l}}{(n \cdot \kappa)^{c_4/\epsilon}}$ for constants $c_3, c_4 > 0$ for all $i \in O(\frac{1}{\epsilon})$. From Lemma 20, we have that $\left| \langle t_i, f_2 \rangle - \langle t_i, s_{\mathbf{B}} \rangle \right| \leq \frac{C_2 \log(N/\delta)}{\sqrt{n}} \leq \frac{C_2 \log(1/\epsilon\delta)}{\sqrt{n}}$ for some constant C_2 and for all $i \in O(1/\epsilon)$ with probability at least $1 - \delta$. Thus, we get that

$$\begin{aligned}
& \left| \frac{1}{n} \sum_{j=k+1}^n \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j=k+1}^m w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\
& \leq \frac{g_{\min}^{c_3 l}}{(n \cdot \kappa)^{c_3/\epsilon}} + \frac{g_{\min}^{c_4 l}}{(n \cdot \kappa)^{c_4/\epsilon}} + \frac{C_2 \log(1/\epsilon\delta)}{\sqrt{n}} \leq \frac{2C_2 \log(1/\epsilon\delta)}{\sqrt{n}},
\end{aligned}$$

for all $i \in O(\frac{1}{\epsilon})$. We now bound the second term in (94). Note that since $|\lambda_j(\mathbf{B})| \leq 1$ for $j \geq k+1$, we have that $\bar{T}_i(\lambda_j(\mathbf{B})) \leq \sqrt{\frac{2}{\pi}}$. Also, $|\lambda_j(\mathbf{T}_1)| \leq |\lambda_j(\mathbf{B})| \leq 1$ for $j \geq k+1$ and $w_j^2 \leq \frac{C_1 \sqrt{\log(l/\delta)}}{n}$ for all $j \in S$. Thus, we get that:

$$\begin{aligned}
& \left| \frac{1}{n} \sum_{j \geq k+1, j \in S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j \geq k+1, j \in S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\
& \leq \frac{1}{n} \sum_{j \geq k+1, j \in S_1} |\bar{T}_i(\lambda_j(\mathbf{B}))| + \sum_{j \geq k+1, j \in S} w_j^2 |\bar{T}_i(\lambda_j(\mathbf{T}_1))| \\
& \leq \sqrt{\frac{2}{\pi}} \frac{|S_1|}{n} + \sqrt{\frac{2}{\pi}} \frac{C_1 |S| \sqrt{\log(l/\delta)}}{n} \leq \frac{C_3 l \sqrt{\log(l/\delta)}}{n}.
\end{aligned}$$

for some constant $C_3 > 0$ and where we used the fact that $|S_1| = |S|$ and $|S| \leq l$ in the last step. We finally bound the last term in (94). Observe that we have (for some constant $C_4 > 0$):

$$\begin{aligned}
& \left| \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) - \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\
& \leq \left| \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \left(\frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} - 1 \right) \right| \\
& \leq \left| \frac{\frac{n-|S|}{n} - \sum_{p \in [m] \setminus S} w_p^2}{\sum_{p \in [m] \setminus S} w_p^2} \right| + \left| \sum_{p \in [m] \setminus S} w_p^2 \sqrt{\frac{2}{\pi}} \right| \\
& \leq \sqrt{\frac{2}{\pi}} \left| \sum_{p \in S} w_p^2 - \frac{|S|}{n} \right| \\
& \leq \sqrt{\frac{2}{\pi}} \frac{C_2 |S| \sqrt{\log(l/\delta)}}{n} + \sqrt{\frac{2}{\pi}} \frac{|S|}{n} \\
& \leq \frac{C_4 l \sqrt{\log(l/\delta)}}{n}.
\end{aligned}$$

In the second step above, we used the fact that since $[k] \subseteq S$, we have $\lambda_j(\mathbf{T}_1) \leq \lambda_k(\mathbf{T}_1) \leq \lambda_k(\mathbf{B}) \leq 1$ for $j \in [m] \setminus S$. So we have $\bar{T}_i(\lambda_j(\mathbf{T}_1)) \leq \sqrt{\frac{2}{\pi}}$ for $j \in [m] \setminus S$. In the third step, we used the fact that $\sum_{i \in [m]} w_i^2 = 1$ and in the fourth step, we bounded $w_p^2 \leq \frac{C_2 \sqrt{\log(l/\delta)}}{n}$. Finally, using the upper bounds on the three terms on the right hand side of (94), and observing that for large enough we get that for $i \in O(\frac{1}{\epsilon})$ (for constants $C_2, C_5 > 0$), with probability at least $1 - \delta$:

$$\left| \frac{1}{n} \sum_{j \in [n] \setminus S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \leq \frac{C_2 \log(1/\epsilon\delta)}{\sqrt{n}} + \frac{C_5 l \sqrt{\log(l/\delta)}}{n}.$$

From (93), we get that (using the fact that $N = O(\frac{1}{\epsilon})$):

$$\begin{aligned}
\frac{n-|S|}{n} W_1(s_2, f_1) & \leq C_6 \epsilon \frac{n-|S|}{n} + \left(\frac{C_6 \log(1/\epsilon\delta)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)}}{n} \right) \sum_{i=1}^N \frac{1}{i} \\
& \leq C_6 \epsilon + \frac{C_6 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n},
\end{aligned}$$

for large constant $C_6 > 0$. Finally, using the bounds on $W_1(s_2, f_1)$ from above, and on $W_1(s_1, s_2)$ from (92), in (91), we get that:

$$I_2 \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}} + C_6 \epsilon + \frac{C_6 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n},$$

where we also used the fact that $\frac{n-|S|}{n} \leq 1$. Thus, from (90), we get that $W_1(s_{\mathbf{C}}, f_{\mathbf{B}}) \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}} + C_6 \epsilon + \frac{C_6 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n}$. Then, from (89) and using triangle inequality, we get

that:

$$\begin{aligned} W_1(s_{\mathbf{B}}, f_{\mathbf{B}}) &\leq W_1(s_{\mathbf{B}}, s_{\mathbf{C}}) + W_1(s_{\mathbf{C}}, f_{\mathbf{B}}) \\ &\leq \frac{4\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}} + C_6\epsilon + \frac{C_6 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n}. \end{aligned}$$

Finally, from (87), we get that:

$$\begin{aligned} W_1(s_{\mathbf{A}}, f) &\leq 2\alpha \left(\frac{4\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}} + C_6\epsilon + \frac{C_6 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n} \right) \\ &\leq \frac{4\|\mathbf{A}\|_2}{n^{(\beta-1)/\epsilon}} + 2\alpha \left(C_6\epsilon + \frac{C_6 \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} + \frac{C_6 l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n} \right), \end{aligned}$$

where the second step follows from the facts that $\mathbf{B} = \frac{\mathbf{A}}{2\alpha}$. Also note that we have $\alpha \leq \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2 g_{\min}^{c/4}}{n^{c/4}}$. Simplifying the above expression we get that for some large constants $C' > 0$ and $c_2 > 0$ such that $\frac{4}{n^{(\beta-1)/\epsilon}} < \frac{1}{n^{c_2}}$, we have:

$$W_1(s_{\mathbf{A}}, f) \leq C'\epsilon\sigma_{l+1}(\mathbf{A}) + \frac{C' \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \sigma_{l+1}(\mathbf{A}) + \frac{C' l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n} \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2}}.$$

Case 2: We now consider the case when such a k doesn't exist. In this case, we have $\|\mathbf{A}\|_2 \leq 2\alpha$. Then, following case 2 of Theorem 3, let $\mathbf{B} = \frac{\mathbf{A}}{2\alpha}$ such that $\|\mathbf{B}\|_2 \leq 1$. Then, following case 1, we can again apply the backwards error to get a matrix $\mathbf{C}$ such that $\|\mathbf{B} - \mathbf{C}\|_2 \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}$. We can again split $W_1(s_{\mathbf{C}}, f_{\mathbf{B}})$ into the integrals I_1 and I_2 as in (90) such that $I_1 = 0$ and I_2 can be bounded as $I_2 \leq \frac{n-|S|}{n} (W_1(s_1, s_2) + W_1(s_2, f_1))$ (as in (91)) such that $W_1(s_1, s_2) \leq \frac{2\|\mathbf{B}\|_2}{n^{(\beta-1)/\epsilon}}$ (as in (92)) and $\frac{n-|S|}{n} W_1(s_2, f_1)$ is bounded using Lemma 3.1 of [BKM22] by the Chebyshev moments of s_2 and f_1 as in (93). Then, using triangle inequality, we can bound the moments of f_1 and s_2 again as in (94) such that we have:

$$\begin{aligned} &\left| \frac{1}{n} \sum_{j \in [n] \setminus S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\ &\leq \left| \frac{1}{n} \sum_{j \in [n]} \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j \in [m]} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| + \left| \frac{1}{n} \sum_{j \in S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j \in S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \\ &+ \left| \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) - \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right|. \end{aligned} \quad (95)$$

Since $w_j^2 \leq \frac{C_2 \sqrt{\log(l/\delta)}}{n}$ for $j \in S$ and $\lambda_j(\mathbf{T}_1) \leq \lambda_j(\mathbf{B}) \leq 1$, the second and third terms on the right hand side above can be bounded by $O\left(\frac{l \sqrt{\log(l/\delta)}}{n}\right)$ as in case 1. To bound the first term, observe that $\frac{1}{n} \sum_{j \in [n]} \bar{T}_i(\lambda_j(\mathbf{B})) = \frac{1}{n} \text{tr}(\bar{T}_i(\mathbf{B}))$ and $\sum_{j \in [m]} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1))$ is the i th Chebyshev moment of the output of SLQ (Algorithm 4) with $\mathbf{B}$ as input. Thus, from Lemma 12 we get

that $\left| \frac{1}{n} \sum_{j \in [n]} \bar{T}_i(\lambda_j(\mathbf{B})) - \sum_{j \in [m]} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \leq O\left(\frac{\log(N/\delta)}{\sqrt{n}}\right) \leq O\left(\frac{\log(1/\epsilon\delta)}{\sqrt{n}}\right)$. Thus, we get $\left| \frac{1}{n} \sum_{j \in [n] \setminus S_1} \bar{T}_i(\lambda_j(\mathbf{B})) - \frac{\frac{n-|S|}{n}}{\sum_{p \in [m] \setminus S} w_p^2} \sum_{j \in [m] \setminus S} w_j^2 \bar{T}_i(\lambda_j(\mathbf{T}_1)) \right| \leq O\left(\frac{\log(1/\epsilon\delta)}{\sqrt{n}} + \frac{l\sqrt{\log(l/\delta)}}{n}\right)$. The rest of the proof follows the proof of case 1 which gives us the final bound of

$$W_1(s, f) \leq \epsilon \sigma_{l+1}(\mathbf{A}) + \frac{C' \log(1/\epsilon\delta) \log(1/\epsilon)}{\sqrt{n}} \sigma_{l+1}(\mathbf{A}) + \frac{C' l \sqrt{\log(l/\delta)} \log(1/\epsilon)}{n} \sigma_{l+1}(\mathbf{A}) + \frac{\|\mathbf{A}\|_2}{n^{c_2}}.$$

Finally, observe that adjusting δ by some constant factors gives us the final bound. $\square$

5 Lower Bound

We now prove the lower bound on the number of matrix vector queries required by any algorithm to estimate the spectral density of any matrix $\mathbf{A}$ upto Wasserstein error $\epsilon \sigma_{l+1}(\mathbf{A})$. Our proof proceeds via a reduction of the spectral density estimation problem to the problem of distinguishing between two Wishart matrices with ranks very close to each other (Theorem 17 of [SW23]). Our lower bound of $O(l + \frac{1}{\epsilon})$ shows that our upper bounds for estimating the SDE via explicit deflation as well as SLQ are nearly tight (upto polylog factors).

Theorem 2. *Any (possibly randomized) algorithm that given symmetric $\mathbf{A} \in \mathbb{R}^{n \times n}$ outputs $\tilde{s}_{\mathbf{A}}$ such that, with probability at least $1/2$, $W_1(s_{\mathbf{A}}, \tilde{s}_{\mathbf{A}}) \leq \epsilon \sigma_{l+1}(\mathbf{A})$ for $\epsilon \in (0, 1)$ and $l \in [n]$ must make $\Omega(l + \frac{1}{\epsilon})$ (possibly adaptively chosen) matrix-vector product queries to $\mathbf{A}$.*

Proof. Let $\mathcal{A}$ be an adaptive algorithm that estimates the spectral density of $\mathbf{A}$ up to error $\epsilon \sigma_{l+1}(\mathbf{A})$ in the Wasserstein-1 norm. Let $W(n, r)$ be the n dimensional Wishart distribution with r degrees of freedom i.e., the distribution of $\mathbf{G}\mathbf{G}^T$ where $\mathbf{G} \in \mathbb{R}^{n \times r}$ has i.i.d. standard normal entries. We will use Theorem 17 of [SW23] which states that at least $\Omega(r)$ (possibly adaptive) matrix vector queries are required by any adaptive algorithm to distinguish between two Wishart matrices $W(n, r)$ and $W(n, r+2)$ with probability at least $\frac{2}{3}$. We prove the lower bound by considering the two cases: $l > \frac{1}{\epsilon}$ and $l \leq \frac{1}{\epsilon}$.

Case 1. ($l > \frac{1}{\epsilon}$): The non-zero eigenvalues of the Wishart ensembles $W(n, l)$ and $W(n, l+2)$ are bounded between $\frac{n}{2}$ and $2n$ with probability at least $5/6$ as long as $n \geq Cl$ for some constant C [Ver18]. Let $n = Cl$. Consider the Wishart ensembles $\mathbf{A}_1 = W(n, l)$ and $\mathbf{A}_2 = W(n, l+2)$. Observe that $\mathbf{A}_1$ and $\mathbf{A}_2$ have ranks of l and $l+2$ respectively. So, $\sigma_{l+3}(\mathbf{A}_1) = \sigma_{l+3}(\mathbf{A}_2) = 0$. Let $\mathcal{A}$ use k matrix-vector products with the input matrix to estimate the spectral density $s_{\mathbf{A}}(x)$ up to error $\epsilon \sigma_{l+3}(\mathbf{A}) = 0$ in both cases with probability at least 0.5 i.e. the spectral density of the input matrix is estimated exactly. So, the rank of $\mathbf{A}$ in both cases is given exactly by $n \int_{n/2}^{2n} s_{\mathbf{A}}(x) dx$ with probability at least 0.5 . Hence, $\mathcal{A}$ can distinguish between $W(n, l)$ and $W(n, l+2)$ with probability at least $\frac{2}{3}$. Thus, from Theorem 17 of [SW23], we must have $k = \Omega(l)$.

Case 2. ($l \leq \frac{1}{\epsilon}$): In this case, let $r = \lfloor \frac{1}{\epsilon} \rfloor$ and let us consider the normalized Wishart ensembles $\mathbf{A}_1 = \frac{1}{2n} W(n, r)$ and $\mathbf{A}_2 = \frac{1}{2n} W(n, r+2)$ where $n = 2Cr$. Let either $\mathbf{A}_1$ or $\mathbf{A}_2$ be the input to $\mathcal{A}$. Since the nonzero eigenvalues of $\mathbf{A}_1$ and $\mathbf{A}_2$ lie in $[0.25, 1]$ with probability at least $5/6$, by setting the error parameter to $\frac{\epsilon}{1000C}$ we can estimate the spectral density of the input to $\mathcal{A}$ up to error $\frac{\epsilon}{1000C} \max(\sigma_{l+1}(\mathbf{A}_1), \sigma_{l+1}(\mathbf{A}_2)) \leq \frac{\epsilon}{1000C}$ in the Wasserstein-1 distance. Let $\lambda_1 \geq \dots \geq \lambda_n$ be the true eigenvalues of the input to $\mathcal{A}$. Then, using the estimated spectral density returned by $\mathcal{A}$, we can

construct a list of n values $[\tilde{\lambda}_1, \dots, \tilde{\lambda}_n]$ in time linear in n and $\frac{1}{\epsilon}$ such that $\sum_{i=1}^n |\lambda_i - \tilde{\lambda}_i| \leq \frac{3\epsilon n}{1000C} \leq \frac{6}{1000}$ (see [CSKSV18], theorem B.1 in [BKM22]). So, we have $|\lambda_{r+1} - \tilde{\lambda}_{r+1}| \leq \sum_{i=1}^n |\lambda_i - \tilde{\lambda}_i| \leq \frac{6}{1000}$. When, $\mathbf{A}_1$ is the input we get $|\tilde{\lambda}_{r+1}| \leq \frac{6}{1000}$ since $\lambda_{r+1} = 0$ in this case. When, $\mathbf{A}_3$ is the input, via triangle inequality, we get $|\tilde{\lambda}_{r+1}| \geq |\lambda_{r+1}| - \frac{6}{1000} \geq 0.25 - \frac{6}{1000}$ since $\lambda_{r+1} > 0.25$ in this case. So, we can distinguish between $\mathbf{A}_1$ and $\mathbf{A}_2$ with probability at least $\frac{2}{3}$. So, using Theorem 17 of [SW23], we again have $k \geq \Omega(\frac{1}{\epsilon})$.

Finally, observe that setting $\delta = \frac{1}{nc'}$ for some constant c' gives us the final bound. $\square$

6 Empirical Evaluation

In this section, we compare the empirical performance of the algorithms studied in this paper and several other standard algorithms in approximating the spectral densities of several synthetic and publicly available matrices. We observe that the SLQ algorithm (Algorithm 4) and its variance reduced variant, VR-SLQ (Algorithm 5) generally out perform explicit moment matching methods like KPM and Chebyshev moment matching in the Wasserstein distance metric. This finding aligns with our results in Theorem 3 and 4 that these methods perform implicit deflation for any deflation parameter ℓ allowing them to adapt to the matrix spectrum to achieve stronger error bounds.

6.1 Datasets

Our comparisons are performed on the following matrices:

- The *Gaussian matrix*, which is a 5000×5000 matrix, constructed by first drawing 5000 eigenvalues from a Gaussian distribution to get $\mathbf{\Lambda} \in \mathcal{N}(0, 1)^{5000}$, normalizing $\mathbf{\Lambda} = \mathbf{\Lambda} / \|\mathbf{\Lambda}\|_\infty$, then generating a random orthonormal matrix $\mathbf{V} \in \mathbb{R}^{5000 \times 5000}$, and finally computing $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$. This matrix is generated from the descriptions of the Gaussian matrix in [BKM22].
- The *Uniform matrix* is constructed analogously to the Gaussian matrix, except its eigenvalues are drawn uniformly and independently from $[-1, 1]$. This matrix is generated from the descriptions of the Uniform matrix in [BKM22].
- The *inverse spectrum matrix* is a 5000×5000 diagonal matrix with entries $1, 1/2, \dots, 1/5000$.
- The *power law spectrum matrix* is a 5000×5000 diagonal matrix with entries $1, 1/2^2, \dots, 1/2^{5000}$.
- The *low-rank matrix* is a 5000×5000 diagonal matrix with 100 entries drawn uniformly and independently from a Gaussian distribution $\mathcal{N}(0, 1)$, and normalized as $\mathbf{\Lambda} = \mathbf{\Lambda} / \|\mathbf{\Lambda}\|_\infty$. The remaining diagonal entries are set at 0.
- Finally, *Erdos992* is the adjacency matrix of the Erdős collaboration network (snapshot of 1992) with 6100 vertices and containing 15030 undirected edges. It is taken from a publicly available sparse matrix collection [DH11]. The adjacency matrix has one eigenvalue at 1, roughly 1000 eigenvalues at 0, and roughly 500 eigenvalues of magnitude greater than 0.2.

We note that all algorithms considered in this paper are ‘rotationally invariant’ in that their performance should not depend on the actual eigenvector basis of $\mathbf{A}$. For this reason, most of our test matrices are simply diagonal matrices. And we indeed see no systematic difference in performance depending on the eigenvector basis.

6.2 Implementation Details

We consider three baseline SDE algorithms: 1) stochastic Lanczos quadrature (SLQ), 2) the Chebyshev moment matching (CMM) algorithm (Algorithm 1 of [BKM22]), and 3) the Jackson damped kernel polynomial method (KPM) (Algorithm 6 of [BKM22]) which is a popular moment matching algorithm that can be thought of as an approximation of CMM. We implement two versions of our explicit deflation algorithm (Algorithm 1) – one of which uses CMM after deflation (the one we analyze), and one that uses KPM. We call these algorithms def-CMM and def-KPM respectively. Along with these algorithms, we also compare the performance of SLQ algorithm (Algorithm 4) and VR-SLQ (Algorithm 5).

Since we test for relatively small n , as the number of iterations/moments performed by each algorithm increases, we may reach small $\epsilon = \tilde{o}(1/\sqrt{n})$ values for which more than one random vectors in Hutchinson’s method in the explicit moment matching methods or more than one independent trial of SLQ are needed. For simplicity, for all moment matching methods we use 15 random vectors for Hutchinson’s method. For SLQ-based methods we perform 15 independent trials of the method (using 15 independent random starting vectors) and average together the output densities to obtain our final spectral density estimate. For block Krylov based deflation we perform 15 iterations to generate the Krylov subspace.

We compare each algorithm based on the Wasserstein-1 error achieved for a fixed number of total matrix-vector queries to the input matrix. Since the iterations in Krylov and number of vectors in Hutchinson’s algorithm are fixed at 15, to ensure that def-CMM or def-KPM uses the same number of matrix-vector queries as other algorithms we then need to split the moment budget of CMM or KPM to accomodate for the matrix-vector queries due to block Krylov. We split this in the ratio 1 : 3, i.e., for every 2 moment computations of CMM or KPM in def-CMM or def-KPM, we use a block size of 3 to compute deflation via block Krylov algorithm. Note, for block size of 3 and ℓ iterations, block Krylov method uses 6ℓ matrix-vector queries. We vary the total matrix-vector query budget for all algorithms, i.e., increase the total moments approximated by CMM-based algorithms, or increase the total number of iterations in SLQ-variants, and report the Wasserstein-1 error in Figure 1 across all matrices.

To compute N moment estimates, the CMM algorithm [BKM22] is evaluated on a discrete and evenly-spaced grid of length $d + 1$ in the interval $[-1, 1]$. Theoretically, the algorithm requires setting $d = \lceil N^3/3 \rceil$. In our experiments, we found setting $d = 20000$ to be sufficient. To solve the ℓ_1 -regression problem in the CMM algorithm, we use HiGHS solvers [HH18] within SciPy [VGO⁺20]. To ensure that the solution vector falls within the probability simplex we use Algorithm 1 of [Con16]. The eigenvalues of any matrix are computed using numpy [Com21].

We repeat each experiment for $t = 10$ independent trials. In the plots in Figure 1, the x -axis denotes the total matrix-vector queries used by each algorithm per trial, and the y -axis represents the corresponding Wasserstein-1 SDE error. The bold lines in the plot represent the mean error across 10 trials. The 10th and the 90th percentile of the observed errors are represented by the faded envelope around the bold lines.

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

6.3 Summary of Results

We observe that across all of the matrices, SLQ (Algorithm 4) and VR-SLQ (Algorithm 5) outperform the explicit moment matching-based algorithms that we test. Among these two algorithms, VR-SLQ more often outperforms SLQ, especially when the spectrum of the input matrix contains only a few large eigenvalues, as in this case, the variance reduction step can have a significant positive effect on the spectral density estimate. We also observe that the variants of Algorithm 1 (def-CMM and def-KPM) more often outperform naive CMM and KPM, in particular for matrices with only a few large eigenvalues, as expected.




Figure 1: **Wasserstein-1 error of spectral density estimation approximation algorithms.** In the figures above, we plot the Wasserstein-1 error of approximating the spectral density of several matrices using the algorithms presented in the paper and some baseline algorithms. We observe that in almost all cases, VR-SLQ algorithm outperforms all other SDE algorithms. We also observe that the variants of our deflation algorithm (Algorithm 1) generally outperform the corresponding baseline of CMM and KPM.

Acknowledgements

Cameron, Rajarshi, and Archan were partially supported by NSF grants 1934846, 2046235, and 2427363. Part of the work was completed while Archan was at the University of Massachusetts Amherst, NY. Christopher was supported by NSF grant 2045590.

References

- [ACH⁺24] Noah Amsel, Tyler Chen, Diana Halikias, Feyza Duman Keles, Cameron Musco, and Christopher Musco. Fixed-sparsity matrix approximation from matrix-vector products. *arXiv:2402.09379*, 2024.
- [AL86] Owe Axelsson and Gunhild Lindskog. On the rate of convergence of the preconditioned conjugate gradient method. *Numerische Mathematik*, 1986.
- [AN13] Alexandr Andoni and Huy L Nguyen. Eigenvalues of a matrix in the streaming model. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2013.
- [BCW22] Ainesh Bakshi, Kenneth L Clarkson, and David P Woodruff. Low-rank approximation with $1/\epsilon^{1/3}$ matrix-vector products. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022.
- [BDD⁺24] Rajarshi Bhattacharjee, Gregory Dexter, Petros Drineas, Cameron Musco, and Archan Ray. Sublinear time eigenvalue approximation via random sampling. *Algorithmica*, 2024.
- [BEPW98] Kevin Burrage, Jocelyne Erhel, Bert Pohl, and Alan Williams. A deflation technique for linear systems of equations. *SIAM Journal on Scientific Computing*, 1998.
- [BGVKS22] Jess Banks, Jorge Garza-Vargas, Archit Kulkarni, and Nikhil Srivastava. Pseudospectral shattering, the sign function, and diagonalization in nearly matrix multiplication time. *Foundations of Computational Mathematics*, 2022.
- [BHSW20] Mark Braverman, Elad Hazan, Max Simchowitz, and Blake Woodworth. The gradient complexity of linear regression. In *Proceedings of the 33rd Annual Conference on Computational Learning Theory (COLT)*, 2020.
- [BKM22] Vladimir Braverman, Aditya Krishnan, and Christopher Musco. Sublinear time spectral density estimation. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022.
- [BN23] Ainesh Bakshi and Shyam Narayanan. Krylov methods are (nearly) optimal for low-rank approximation. In *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2023.
- [BRP92] C Benoit, E Royer, and G Poussiguet. The spectral moments method. *Journal of Physics: Condensed Matter*, 1992.
- [Che23] Tyler Chen. A spectrum adaptive kernel polynomial method. *The Journal of Chemical Physics*, 2023.
- [Com21] The Numpy Community. `numpy.linalg.eigvals`. <https://numpy.org/doc/stable/reference/generated/numpy.linalg.eigvals.html>, 2021.
- [Con16] Laurent Condat. Fast projection onto the simplex and the ℓ_1 ball. *Mathematical Programming*, 2016.

- [CPB10] L Covaci, FM Peeters, and M Berciu. Efficient numerical approach to inhomogeneous superconductivity: The Chebyshev-Bogoliubov–de Gennes method. *Physical review letters*, 2010.
- [CS97] Andrew Chapman and Yousef Saad. Deflated and augmented Krylov subspace techniques. *Numerical linear algebra with applications*, 1997.
- [CSKSV18] David Cohen-Steiner, Weihao Kong, Christian Sohler, and Gregory Valiant. Approximating the spectrum of a graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2018.
- [CTU21] Tyler Chen, Thomas Trogon, and Shashanka Ubaru. Analysis of stochastic Lanczos quadrature for spectrum approximation. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
- [CTU22] Tyler Chen, Thomas Trogon, and Shashanka Ubaru. Randomized matrix-free quadrature for spectrum and spectral sum approximation. [arXiv:2204.01941](https://arxiv.org/abs/2204.01941), 2022.
- [DBB19] Kun Dong, Austin R Benson, and David Bindel. Network density of states. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2019.
- [DDH07] James Demmel, Ioana Dumitriu, and Olga Holtz. Fast linear algebra is stable. *Numerische Mathematik*, 2007.
- [DH11] Timothy A Davis and Yifan Hu. The university of florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)*, 2011.
- [DIKMI18] Petros Drineas, Ilse CF Ipsen, Eugenia-Maria Kontopoulou, and Malik Magdon-Ismail. Structural convergence results for approximation of dominant subspaces from block Krylov spaces. *SIAM Journal on Matrix Analysis and Applications*, 2018.
- [DM23] Prathamesh Dharangutte and Christopher Musco. A tight analysis of hutchinson’s diagonal estimator. *Proceedings of the 6th Symposium on Simplicity in Algorithms (SOSA)*, 2023.
- [DNPS16] Edoardo Di Napoli, Eric Polizzi, and Yousef Saad. Efficient estimation of eigenvalue counts in an interval. *Numerical Linear Algebra with Applications*, 2016.
- [EG17] Nicole Eikmeier and David F Gleich. Revisiting power-law distributions in spectra of real world networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2017.
- [FDBV01] Illés J Farkas, Imre Derényi, Albert-László Barabási, and Tamas Vicsek. Spectra of “real-world” graphs: Beyond the semicircle law. *Physical Review E*, 2001.
- [FTU23] Zachary Frangella, Joel A Tropp, and Madeleine Udell. Randomized Nyström preconditioning. *SIAM Journal on Matrix Analysis and Applications*, 2023.
- [FV01] Jason Frank and Cornelis Vuik. On the construction of deflation-based preconditioners. *SIAM Journal on Scientific Computing*, 2001.

- [Gir87] Didier Girard. Un algorithme simple et rapide pour la validation croisee g  n  ralis  e sur des probl  mes de grande taille. Technical report, 1987.
- [GKX19] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via Hessian eigenvalue density. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- [GM09] Gene H Golub and G  rard Meurant. *Matrices, moments and quadrature with applications*. Princeton University Press, 2009.
- [GOSS16] Alon Gonen, Francesco Orabona, and Shai Shalev-Shwartz. Solving ridge regression using sketched preconditioned SVRG. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, 2016.
- [GSO17] Arjun Singh Gambhir, Andreas Stathopoulos, and Kostas Orginos. Deflation as a method of variance reduction for estimating the trace of a matrix inverse. *SIAM Journal on Scientific Computing*, 2017.
- [HH18] Qi Huangfu and JA Julian Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 2018.
- [HHK72] R Haydock, V Heine, and M J Kelly. Electronic structure based on the local atomic environment for tight-binding bands. *Journal of Physics C: Solid State Physics*, 1972.
- [HMT11] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 2011.
- [HT23] Diana Halikias and Alex Townsend. Structured matrix recovery from matrix-vector products. *Numerical Linear Algebra with Applications*, 2023.
- [Hut90] Michael F. Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 1990.
- [JKM⁺24] Yujia Jin, Ishani Karmarkar, Christopher Musco, Aaron Sidford, and Apoorv Vikram Singh. Faster spectral density estimation and sparsification in the nuclear norm. *Proceedings of the 37th Annual Conference on Computational Learning Theory (COLT)*, 2024.
- [JMSS23] Yujia Jin, Christopher Musco, Aaron Sidford, and Apoorv Vikram Singh. Moments, random walks, and limits for spectrum approximation. In *Proceedings of the 36th Annual Conference on Computational Learning Theory (COLT)*, 2023.
- [KR57] Leonid Vital’evich Kantorovich and Gennadii Shlemovich Rubinshtein. On a functional space and certain extremum problems. In *Doklady Akademii Nauk*. Russian Academy of Sciences, 1957.
- [Lan52] Cornelius Lanczos. Solution of systems of linear equations by minimized iterations. *J. Res. Nat. Bur. Standards*, 1952.

- [LG82] P Lambin and J-P. Gaspard. Continued-fraction technique for tight-binding systems. A generalized-moments method. *Physical Review B*, 1982.
- [Lin17] Lin Lin. Randomized estimation of spectral densities of large matrices made accurate. *Numerische Mathematik*, 2017.
- [LSY98] Richard B Lehoucq, Danny C Sorensen, and Chao Yang. *ARPACK users’ guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods*. SIAM, 1998.
- [LSY16] Lin Lin, Yousef Saad, and Chao Yang. Approximating spectral densities of large matrices. *SIAM Review*, 2016.
- [LXS19] Ruipeng Li, Yuanzhe Xi, Lucas Erlandson, and Yousef Saad. The eigenvalues slicing library (EVSL): Algorithms, implementation, and software. *SIAM Journal on Scientific Computing*, 2019.
- [MM15] Cameron Musco and Christopher Musco. Randomized block Krylov methods for stronger and faster approximate singular value decomposition. In *Advances in Neural Information Processing Systems 28 (NIPS)*, 2015.
- [MM19] Michael Mahoney and Charles Martin. Traditional and heavy tailed self regularization in neural network models. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- [MMM24] Raphael Meyer, Cameron Musco, and Christopher Musco. On the unreasonable effectiveness of single vector Krylov methods for low-rank approximation. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2024.
- [MMM21] Raphael A Meyer, Cameron Musco, Christopher Musco, and David P Woodruff. Hutch++: Optimal stochastic trace estimation. In *Symposium on Simplicity in Algorithms (SOSA)*, 2021.
- [MNS⁺18] Cameron Musco, Praneeth Netrapalli, Aaron Sidford, Shashanka Ubaru, and David P. Woodruff. Spectrum approximation beyond fast matrix multiplication: Algorithms and hardness. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science (ITCS)*, 2018.
- [NSW22] Deanna Needell, William Swartworth, and David P Woodruff. Testing positive semidefiniteness using linear measurements. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2022.
- [Par98] Beresford N Parlett. *The symmetric eigenvalue problem*. SIAM, 1998.
- [PC99] Victor Y Pan and Zhao Q Chen. The complexity of the matrix eigenproblem. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC)*, 1999.
- [Pea94] Barak A. Pearlmutter. Fast exact multiplication by the Hessian. *Neural Computation*, 1994.

- [PSG18] Jeffrey Pennington, Samuel Schoenholz, and Surya Ganguli. The emergence of spectral universality in deep networks. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- [RL18] Aditya Ramesh and Yann LeCun. Backpropagation for implicit spectral densities. [*arXiv:1806.00499*](#), 2018.
- [Saa80] Yousef Saad. On the rates of convergence of the Lanczos and the block-Lanczos methods. *SIAM Journal on Numerical Analysis*, 1980.
- [Saa03] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [SEAR18] Max Simchowitz, Ahmed El Alaoui, and Benjamin Recht. Tight query complexity lower bounds for PCA via finite sample deformed Wigner law. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC)*, 2018.
- [Ski89] John Skilling. The eigenvalues of mega-dimensional matrices. *Maximum Entropy and Bayesian Methods*, 1989.
- [SR94] RN Silver and H Röder. Densities of states of mega-dimensional Hamiltonian matrices. *International Journal of Modern Physics C*, 1994.
- [SRS20] Jürgen Schnack, Johannes Richter, and Robin Steinigeweg. Accuracy of the finite-temperature Lanczos method compared to simple typicality-based estimates. *Physical Review Research*, 2020.
- [STBB17] Björn Sbierski, Maximilian Trescher, Emil J Bergholtz, and Piet W Brouwer. Disordered double Weyl node: Comparison of transport and density of states calculations. *Physical Review B*, 2017.
- [Sun95] Ji-guang Sun. A note on backward perturbations for the hermitian eigenvalue problem. *BIT Numerical Mathematics*, 1995.
- [SW09] Daniel A Spielman and Jaehoh Woo. A note on preconditioning by low-stretch spanning trees. [*arXiv:0903.2816*](#), 2009.
- [SW23] William Swartworth and David P Woodruff. Optimal eigenvalue approximation via sketching. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*, 2023.
- [SWYZ21] Xiaoming Sun, David P Woodruff, Guang Yang, and Jialin Zhang. Querying a matrix through matrix-vector products. *ACM Transactions on Algorithms (TALG)*, 2021.
- [Tre19] Lloyd N Trefethen. *Approximation theory and approximation practice, extended edition*. SIAM, 2019.
- [Tro18] Joel A. Tropp. Analysis of randomized block Krylov methods. Technical report, California Institute of Technology , Pasadena, CA, 2018.
- [Tro22] Joel A Tropp. Randomized block Krylov methods for approximating extreme eigenvalues. *Numerische Mathematik*, 2022.

- [UCS17] Shashanka Ubaru, Jie Chen, and Yousef Saad. Fast estimation of $\text{tr}(f(a))$ via stochastic Lanczos quadrature. *SIAM Journal on Matrix Analysis and Applications*, 2017.
- [Ver18] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*. Cambridge university press, 2018.
- [VGO⁺20] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 2020.
- [Wai19a] Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [Wai19b] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge university press, 2019.
- [Wan94] Lin-Wang Wang. Calculating the density of states and optical-absorption spectra of large quantum systems by the plane-wave moments method. *Physical Review B*, 1994.
- [Wey12] Hermann Weyl. The asymptotic distribution law for the eigenvalues of linear partial differential equations (with applications to the theory of black body radiation). *Math. Ann*, 1912.
- [WWAF06] Alexander Weiße, Gerhard Wellein, Andreas Alvermann, and Holger Fehske. The kernel polynomial method. *Reviews of Modern Physics*, 2006.