

SPECTRAL DENSITY ESTIMATION FOR NORMAL MATRICES

CAMERON MUSCO[†], CHRISTOPHER MUSCO[‡], RIKHAV SHAH^{*}, JOHN URSCHEL^{*}, AND NICHOLAS WEST^{*}

ABSTRACT. The spectral density estimation problem asks for an algorithm that, given an $n \times n$ matrix A , outputs a probability measure that is a good approximation to the uniform distribution on the eigenvalues of A , called the spectral density of A . This paper considers the setting where A is a large normal matrix that is accessible only through matrix-vector product queries. We provide an algorithm that makes just m matrix-vector queries to A and returns, with high probability, a measure within earth mover’s distance $O(1/m + \log m/\sqrt{n})$ of the true spectral density of A . We provide a complementary lower bound that any algorithm producing an ε -approximation to the true spectral density for large matrices must make $\Omega(1/\varepsilon)$ matrix-vector queries. The lower bound holds even for the more restricted case of real symmetric input matrices. In combination with our upper bound, it shows that spectral density estimation is essentially no harder for complex normal matrices than for real symmetric matrices.

1. INTRODUCTION

Given a matrix $A \in \mathbb{C}^{n \times n}$, its *spectral density* μ_A is the uniform distribution on its discrete set of n complex eigenvalues. Spectral density estimation (SDE) is the problem of producing, given access to A , a measure $\hat{\mu}_A$ that approximates μ_A in some metric; we will use the earth mover’s distance (also called the Wasserstein distance, Kantorovich-Rubinstein distance, and optimal transport distance).

Spectral density estimation arises as a natural problem in a variety of contexts: in physics and quantum chemistry (where the spectral density is often referred to as the “density of states”) [JP94, LSY16], in the study of the Hessians of neural nets during optimization [GKX19], and as a subroutine in large-scale eigenvalue computation [LXES19]. In many of these contexts, the matrix A is constrained to be real symmetric or Hermitian, and there are several standard references for spectral density estimation in these cases, e.g. [WWAF06, LSY16, APJ⁺18]. In this paper, we expand the scope of spectral density estimation to *normal* matrices. Normal matrices which are not Hermitian arise in several settings. The most important class is unitary matrices. Other examples include circulant matrices, normal random matrix models in physics [CY92, CZ98], normally regular digraphs [Rys63], and directed Cayley graphs where the set of generators is contained in its own normalizer. Moreover, several matrix operations preserve normality: applying matrix functions, taking Kronecker products, adding commuting normal matrices, and taking the compound matrix (matrix of minors), to name a few.

The spectral density can in principle be computed to a high degree of precision by diagonalizing A , at cost $O(n^\omega)$ (or $O(n^3)$ in practice) [BGVKS22]. For modern problems where n is enormous or A is given implicitly (e.g., as a black box), this is prohibitive. The earth mover’s relaxation—only requiring that μ_A and $\hat{\mu}_A$ are close in earth mover’s distance—opens the door to dramatically faster algorithms.

[†]MANNING COLLEGE OF INFORMATION AND COMPUTER SCIENCES, UNIVERSITY OF MASSACHUSETTS AMHERST, AMHERST, MA 01003

[‡]DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING, NEW YORK UNIVERSITY, NEW YORK, NY 10012

^{*}DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA 02139.

E-mail addresses: cmusco@cs.umass.edu, cmusco@nyu.edu, rdshah@mit.edu, urschel@mit.edu, npwest00@mit.edu.

2020 *Mathematics Subject Classification.* Primary 65F15; Secondary 68Q17.

Key words and phrases. Normal matrices, spectral density estimation, kernel polynomial method, algorithmic lower bounds, Wasserstein distance, randomized algorithms, matrix-vector query complexity.

In particular, our algorithm is a matrix-vector query (matvec) algorithm. These algorithms interact with A only through queries of the form $v \mapsto Av$ (and possibly A^*v). The matvec complexity of a problem is the number of such queries needed to solve it; in the worst-case, each query may cost $O(n^2)$, and so we obtain a substantially faster algorithm than full diagonalization when the query complexity is $o(n^{\omega-2})$ (or just $o(n)$ in practice). Important classes of matrix-vector query algorithms are Krylov subspace methods and randomized linear sketching methods, both of which are frequently used to approximate outlying eigenvalues of matrices and solve related spectral problems [Lan50, Saa11, HMT11, Woo14, MM15, SW23].

For Hermitian matrices, the well-known stochastic Lanczos quadrature (SLQ) and kernel polynomial method (KPM) both show that m matrix-vector queries suffice to achieve an earth mover's approximation error of $O(1/m)$ to the spectral density of any Hermitian matrix when the matrix dimension n is sufficiently large [CTU21, BKM22]. A restricted lower bound has also been shown: if the very small approximation error of $O(1/n)$ is desired, then $\Omega(n)$ queries are required [BJM⁺25, MMRS25].

Both SLQ and KPM use the fact that the spectrum of A lies on the real line, and SLQ additionally uses the property that compressions of Hermitian matrices are themselves Hermitian. Neither of those hold in the general normal setting [FP57]. Because the spectrum of A can be distributed throughout a 2D set in \mathbb{C} , one might expect more matrix-vector queries are needed to resolve it. A heuristic based on packing suggests the problem should be quadratically harder: at error ε , one can fit $\Theta(1/\varepsilon^2)$ well-separated eigenvalue clusters in a bounded region of the plane, versus $\Theta(1/\varepsilon)$ on the line. Indeed, such a loss is inherent for related problems like numerical range estimation [CU24].

Contrary to this intuition, we show that m matvecs suffice to achieve approximation error $O(1/m)$ for large normal matrices A . In particular, normal SDE is no harder than Hermitian SDE in the matvec model. We additionally show a substantially broader lower bound as compared to [BJM⁺25, MMRS25]: for any $\varepsilon > 0$, if an algorithm produces an ε approximation for large matrices, then we show $\Omega(1/\varepsilon)$ queries are required.

Before stating our results more formally, we review some standard notation. We use $\|\cdot\|$ to denote the ℓ_2 (Euclidean) norm of vectors, and the spectral norm of a matrices. $\|\cdot\|_p$ denotes the ℓ_p norm of vectors. For a set S , $\|f\|_S = \sup_{x \in S} |f(x)|$ is the sup-norm of a function f on S . The ℓ_p -Lipschitz constant of $f : \mathbb{C}^n \rightarrow \mathbb{C}$ is $\sup_{x \neq y} \frac{|f(x) - f(y)|}{\|x - y\|_p}$. The Dirac delta distribution $\delta(z)$ is defined by $\int f(z)\delta(z - c)dz = f(c)$ for continuous functions f . \mathcal{W}_1 is the earth mover's distance defined for compactly supported probability measures μ, ν as $\mathcal{W}_1(\mu, \nu) = \sup_f (\int f(z)(d\mu(z) - d\nu(z)))$, where the supremum is over functions f with ℓ_2 -Lipschitz constant bounded by 1. We note that there is an equivalent definition of $\mathcal{W}_1(\cdot, \cdot)$ in terms of the optimal transport distance between μ and ν . The above characterization, which will be easier to work with, is often known as Kantorovich-Rubinstein duality. We denote the spectral density of a matrix $A \in \mathbb{C}^{n \times n}$ as

$$\mu_A(z) \equiv \frac{1}{n} \sum_{j=1}^n \delta(z - \lambda_j), \quad (1.1)$$

where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of A . Our main algorithmic result is the following.

Theorem 1.1. *Let $A \in \mathbb{C}^{n \times n}$ be a normal matrix with spectral density μ_A . Fix $0 < \delta < 1$. Algorithm 3 computes an approximating measure $\hat{\mu}_A$ satisfying, with probability $1 - \delta$,*

$$\frac{1}{\|A\|} \mathcal{W}_1(\mu_A, \hat{\mu}_A) \leq O\left(\frac{1}{m} + \sqrt{\frac{\log(m/\delta) \log(m)}{n}}\right). \quad (1.2)$$

The algorithm uses $O(m)$ matrix-vector products with A and A^ , $O(mn)$ storage, and $O(m^2n)$ additional running time. Moreover, the measure $\hat{\mu}_A$ may be chosen to be supported on just $O(m^2)$ atoms.*




FIGURE 1. Spectral density estimation for two example matrices. The leftmost column is a scatter plot of the eigenvalues of the matrices. The next columns left-to-right are the result of running Algorithm 3 for $m = 16, 32, 64,$ and 128 respectively. The first row is with the $150,000 \times 150,000$ matrix where $100,000$ eigenvalues are uniformly distributed in the disk, and $50,000$ additional eigenvalues are uniformly distributed in the intersection of the disk with the second and fourth quadrants. The second row is with the $27,078 \times 27,078$ matrix which has its eigenvalues trace out a version of the “Anscombosaurus” [Cai16, MF17]. All plots depict the unit square $[-1, 1] + i[-1, 1] \subset \mathbb{C}$. The total time for generating the entire figure (including plotting) on a standard laptop with an i7 core took less than 2 seconds.

Remark 1. Several variants of Algorithm 3 are described in Subsection 2.5 for various purposes and applications, all still achieving (1.2) or the appropriate analog.

Remark 2. In addition to applying to a broader class of matrices, Algorithm 3 improves upon the run-times of the algorithms of [BKM22, MMRS25] for Hermitian SDE. Namely, the m dependence is improved by avoiding an optimization problem used by those algorithms to enforce strict moment-matching not required in our analysis. Additionally, a variant of Algorithm 3 described in Subsection 2.5 allows us to output $\hat{\mu}_A$ supported on just $O(m)$ atoms in the Hermitian case, compared to $O(m^{1.5})$ atoms in the algorithm of [MMRS25]. We also note that the arguments of [MMRS25] for producing atomic measures do extend to the non-Hermitian / 2D setting, but this only gives a construction for $O(m^3)$ atoms compared to $O(m^2)$ of Theorem 1.1.

An alternative way to state Theorem 1.1 is that to achieve a target error of $\varepsilon \|A\|$ in earth mover’s distance for an $n \times n$ matrix with n sufficiently large, Algorithm 3 need only make $O(1/\varepsilon)$ queries. We complement this result with a matching lower bound.

Theorem 1.2. Consider any triplet of parameters ε, δ, n satisfying $\delta \geq \varepsilon$ and $n \gg \max\left(\frac{1}{\varepsilon^4 \delta^2}, \left(\frac{2}{\varepsilon}\right)^{1/\delta}\right)$. If a randomized, adaptive, matrix-vector query algorithm, on input any $n \times n$ real symmetric matrix A , outputs a distribution that is within $\varepsilon \|A\|$ in earth mover’s distance of the spectral density A with probability more than 2δ , then it must make $\Omega(1/\varepsilon)$ matrix-vector queries.

Remark 3. The lower bound of Theorem 1.2 is against algorithms which need only succeed for real symmetric inputs. Thinking of ε and δ as fixed constants, this matches our upper bound in the general

normal case, establishing that there is no gap between the complexities of the normal and real symmetric spectral density estimation problems.

Remark 4. A weaker lower bound changes the order of the quantifiers on the parameters: suppose one wants an algorithm which for fixed n can produce an ε approximation for all $\varepsilon > 0$, including ε going to 0. In this case, there are several existing lower bounds, e.g. [MMM21, MMR25, WZZ22] all prove the following: for each n , there exists an ε such that approximating the spectral density to ε earth mover’s error requires $O(1/\varepsilon)$ queries. In particular, their constructions are for $\varepsilon = \Theta(1/n)$. Stated differently, they show that to achieve $O(1/n)$ error, one must learn a constant fraction of the matrix. However, this is not typically the setting in which SDE is interesting. The goal is to obtain fast algorithms which require very few queries relative to the matrix size and output a relatively coarse approximation.

In the practical regime where ε is a small, fixed constant (or tends to 0 slowly), and n is large, the previous lower bounds do not apply. Theorem 1.2 addresses this gap: it applies whenever $\varepsilon = \Omega(1/n^{1/4})$. This resolves the complexity of spectral density estimation outside the window $1/n \ll \varepsilon \ll 1/n^{1/4}$.

Theorem 1.2 follows from a more general statement, Proposition 3.1, which isolates the information-theoretic bottleneck at the core of spectral density estimation using the ideas of [CdDPL⁺24]. We concretely find that any algorithm making only m matrix vector queries to real symmetric A does not learn much unitarily-invariant information about A beyond the values of $\text{tr}(A^j)$ for $j = 1, \dots, 4m$, which are just the first $4m$ moments of the spectral density of A . These moments are both necessary and sufficient to approximate the spectral density of A .

2. ALGORITHM AND ITS ANALYSIS

In this section, we derive Algorithm 3 for approximating the spectral density of a normal matrix and discuss its theoretical properties, most notably by proving Theorem 1.1. We begin by introducing the algorithm itself: given matrix-vector access to A and A^* , we explain how to compute mixed Chebyshev moments of a random surrogate spectral measure and then use these moments to form an approximation to μ_A . After presenting the algorithm, we turn to its analysis via a multivariate version of the kernel polynomial method (KPM). Although the KPM is well known in the Hermitian setting [WWAF06], rigorous error bounds for approximating probability distributions are more recent [BKM22], and the multivariate version we require here has received comparatively little attention, excepting the appendix of [MMRS25]. For this reason, we include a discussion of Jackson’s theorems, their multivariate extension, and how they are deployed our setting. We conclude the section with some algorithmic variants and numerical examples, the results of which are depicted in Figure 1.

Weighted spectral density. A key object appearing in our analysis is the *weighted* spectral density of a matrix A . If we have the eigendecomposition $A = \sum_{j=1}^n \lambda_j v_j v_j^*$, then the weighted spectral density with respect to a unit vector $b \in \mathbb{C}^n$ is the measure

$$\mu_{A,b}(z) \equiv \sum_{j=1}^n \delta(z - \lambda_j) |b^* v_j|^2. \quad (2.1)$$

When $b = \frac{1}{\sqrt{n}}(v_1 + \dots + v_n)$, this coincides with the unweighted spectral measure $\mu_{A,b} = \mu_A$ defined in (1.1). The analysis of our algorithm consists of two steps: (1) show that we can compute an approximation to $\mu_{A,b}$ for any b of our choice (without ever explicitly forming $\mu_{A,b}$), and (2) show that $\mu_{A,b}$ is an approximation to μ_A with high probability over a Haar random selection of b .

Chebyshev polynomials and moments. Playing a role both steps mentioned above are the Chebyshev polynomials and Chebyshev moments. We let $T_j(x)$ denote the j^{th} Chebyshev polynomial

of the first kind, defined by the three-term recurrence,

$$\begin{aligned} T_0(x) &= 1, & T_1(x) &= x, \\ T_{n+1}(x) &= 2xT_n(x) - T_{n-1}(x), & n &\geq 1. \end{aligned} \tag{2.2}$$

The Chebyshev polynomials are orthogonal on the interval $[-1, 1]$ with respect to the weight function $w(x) = (1 - x^2)^{-1/2}$:

$$\int_{-1}^1 \frac{T_j(x)T_k(x)}{\sqrt{1-x^2}} dx = \begin{cases} 0 & j \neq k \\ \pi & j = k = 0 \\ \pi/2 & j = k \geq 1 \end{cases}. \tag{2.3}$$

We write $\tilde{T}_0(x) = 1/\sqrt{\pi}$ and $\tilde{T}_j(x) = \sqrt{2/\pi} \cdot T_j(x)$ for the *normalized* Chebyshev polynomials, which are an orthonormal family. For a measure μ supported on the square $[-1, 1]^2$, define the (j, k) mixed Chebyshev moment as

$$\Gamma_{jk} = \int_{-1}^1 \int_{-1}^1 \tilde{T}_j(x)\tilde{T}_k(y) d\mu(x, y).$$

That is, Γ_{jk} is obtained by integrating the bivariate polynomial $\tilde{T}_j(x)\tilde{T}_k(y)$ over the unit square with respect to μ . By identifying $z = x + iy \in \mathbb{C}$, this definition applies to measures μ on the unit square $[-1, 1] + i[-1, 1]$ in \mathbb{C} as well.

In our application to spectral density estimation, we assume $\|A\|_2 \leq 1$ so that the real and imaginary parts of the eigenvalues of A each lie in the interval $[-1, 1]$. This is without loss of generality, as one may apply the algorithm to a shifted and rescaled matrix $\tilde{A} = c_1(A - c_2)$ and then shift and rescale the output $\hat{\mu}$ to $\hat{\mu}(z/c_1 + c_2)$.¹

2.1. Chebyshev moments of weighted spectral measure. Our first goal is to compute the mixed Chebyshev moments Γ_{jk} for the weighted spectral measure $\mu_{A,b}$. The key observation is that the Toeplitz decomposition $A = H + iK$, i.e. $H = (A + A^*)/2$ and $K = (A - A^*)/2i$ of a normal matrix allows us to compute these moments without actually forming $\mu_{A,b}$.

Lemma 2.1. *Let $A = H + iK$ be the Toeplitz decomposition of the normal matrix A so that $H = (A + A^*)/2$ and $K = (A - A^*)/2i$. Then for any fixed polynomials p, q it holds that*

$$\int_{-1}^1 \int_{-1}^1 \bar{p}(x)q(y) d\mu_{A,b}(x, y) = b^*p(H)^*q(K)b \tag{2.4}$$

where \bar{p} denotes the polynomial with coefficients that are the conjugates of the coefficients of p .

Proof. Since A is normal, by the spectral theorem we may write its eigendecomposition as $A = U\Lambda U^*$ where U is unitary. From the definitions $H = (A + A^*)/2$ and $K = (A - A^*)/2i$, we have that H and K have the corresponding eigendecompositions

$$H = U\operatorname{Re}(\Lambda)U^*, \quad K = U\operatorname{Im}(\Lambda)U^*.$$

¹The desired values of c_1 and c_2 can be computed efficiently. For example, one may take $c_2 = (x^*Ax)/(x^*x)$ with x sampled randomly; in fact, any element of the numerical range of A will suffice. Then since a constant-factor estimate of $\|A - c_2I\|_2$ suffices one may use a few steps of power iteration in order to estimate $c_1 \leq 1/\|A - c_2I\|_2$.

Then we observe that the quadratic form simplifies considerably:

$$\begin{aligned}
b^* \bar{p}(H) q(K) b &= b^* U \bar{p}(\operatorname{Re}(\Lambda)) U^* U q(\operatorname{Im}(\Lambda)) U^* b \\
&= (p(\operatorname{Re}(\Lambda)) U^* b)^* (q(\operatorname{Im}(\Lambda)) U^* b) \\
&= \sum_{j=1}^n \bar{p}(\operatorname{Re}(\lambda_j)) q(\operatorname{Im}(\lambda_j)) |(U^* b)_j|^2 \\
&= \int_{-1}^1 \int_{-1}^1 \bar{p}(x) q(y) \, d\mu_{A,b}(x, y).
\end{aligned}$$

The first equality follows from the definition of matrix functions that if $B = VDV^{-1}$ then $f(B) = Vf(D)V^{-1}$, the second and third equalities rewrite the result as an inner product, and the final equality rewrites the discrete sum as an integral over the weighted spectral measure $\mu_{A,b}$ as defined in (2.1). \square

If we collect the mixed Chebyshev moments of $\mu_{A,b}$ up to degree m into an $(m+1)$ -by- $(m+1)$ matrix Γ , then Lemma 2.1 provides the factorization

$$\Gamma = X^* Y \tag{2.5}$$

where

$$X = \begin{bmatrix} \tilde{T}_0(H)b & \tilde{T}_1(H)b & \cdots & \tilde{T}_m(H)b \end{bmatrix}$$

and

$$Y = \begin{bmatrix} \tilde{T}_0(K)b & \tilde{T}_1(K)b & \cdots & \tilde{T}_m(K)b \end{bmatrix}.$$

The matrices X, Y can be computed using the three-term recurrence of the Chebyshev polynomials, costing only m matrix-vector products with each of H and K . If A and A^* are only available implicitly, then one can simulate the application of H via $Hx = (Ax + A^*x)/2$, and similarly for K . The resulting procedure is summarized in Algorithm 1, which as a result requires $O(m)$ matrix-vector products with H and K or $O(m)$ matrix-vector products with A and A^* . It also requires $O(mn)$ storage for X and Y and $O(m^2n)$ additional operations for the formation of $\Gamma = X^*Y$.

Algorithm 1 CHEBMOMENTS(A, b, m)

Computes mixed Chebyshev moments of $\mu_{A,b}$ up to maximal degree m using three-term recurrence.

Input: $A \in \mathbb{C}^{n \times n}$, $\|A\| \leq 1$, $b \in \mathbb{C}^n$,

Output: $(\Gamma)_{jk} = \int \tilde{T}_j(x) \tilde{T}_k(y) \, d\mu_{A,b}(x, y)$ are the mixed Chebyshev moments of $\mu_{A,b}$ of orders $0 \leq j, k \leq m$.

- 1: $H \leftarrow (A + A^*)/2$
 - 2: $K \leftarrow (A - A^*)/2i$
 - 3: $x_0 \leftarrow b/\sqrt{\pi}$, $x_1 \leftarrow \sqrt{2}Hx_0$, $x_2 \leftarrow 2Hx_1 - \sqrt{2}x_0$
 - 4: $y_0 \leftarrow b/\sqrt{\pi}$, $y_1 \leftarrow \sqrt{2}Ky_0$, $y_2 \leftarrow 2Ky_1 - \sqrt{2}y_0$
 - 5: **for** $j \leftarrow 2$ to $m-1$ **do**
 - 6: $x_{j+1} \leftarrow 2Hx_j - x_{j-1}$
 - 7: $y_{j+1} \leftarrow 2Ky_j - y_{j-1}$
 - 8: **end for**
 - 9: $X \leftarrow [x_0, x_1, x_2, \dots, x_m]$
 - 10: $Y \leftarrow [y_0, y_1, y_2, \dots, y_m]$
 - 11: $\Gamma \leftarrow X^*Y$
 - 12: **return** Γ
-

2.2. Approximating the weighted spectral measure. The previous subsection gave an algorithm for computing the mixed Chebyshev moments Γ_{jk} of $\mu_{A,b}$. We now show how to use a multivariate version of the kernel polynomial method (KPM) [WWAF06] to convert these moments to an approximation to $\mu_{A,b}$.² Our approximation will be

$$d\mu_{A,b}(x, y) \approx d\mu_{\text{KPM}}(x, y) \equiv \frac{p_m(x, y)}{\sqrt{1-x^2}\sqrt{1-y^2}} dx dy, \quad (2.6)$$

where

$$p_m(x, y) = \sum_{j=0}^m \sum_{k=0}^m \Gamma_{jk} \rho_j \rho_k \tilde{T}_j(x) \tilde{T}_k(y).$$

The \tilde{T}_j are the normalized Chebyshev polynomials and the ρ_j are fixed constants, sometimes called *damping coefficients*, which we specify momentarily. We summarize the resulting procedure in Algorithm 2, which uses Algorithm 1 as a subroutine for computing the Chebyshev moments. The central theoretical result of this subsection is the following theorem, whose proof is found at the end of this subsection:

Theorem 2.2. *Let μ be a probability distribution on $[-1, 1]^2$ and let μ_{KPM} be defined as in (2.6). Then*

$$\mathcal{W}_1(\mu, \mu_{\text{KPM}}) \leq \frac{12}{m}. \quad (2.7)$$

To understand μ_{KPM} , we first present the usual single-variable KPM formulation and discuss its generalization to multiple variables. The central object is the polynomial kernel K_m ,

$$K_m(x, r) \equiv \sum_{j=0}^m \rho_j \tilde{T}_j(x) \tilde{T}_j(r). \quad (2.8)$$

One associates with K_m a linear operator (which we will also denote K_m) mapping bounded functions f supported on $[-1, 1]$ to polynomials of degree at most m . Application of this operator is performed via the formula

$$(K_m f)(x) = \int_{-1}^1 \frac{K_m(x, r) f(r)}{\sqrt{1-r^2}} dr. \quad (2.9)$$

If $g(x)$ is a probability density function on the interval $[-1, 1]$ then the KPM approximation to $g(x)$ is the measure with density

$$\tilde{g}(x) = \frac{1}{\sqrt{1-x^2}} \int_{-1}^1 K_m(x, r) g(r) dr \quad (2.10)$$

which can be interpreted as the composition of three operations on g : pointwise multiplication by $\sqrt{1-x^2}$, application of the kernel K_m , and division by $\sqrt{1-x^2}$. Importantly, this multiplication by the weight function $\sqrt{1-x^2}$ exactly cancels the weight function from the definition of the kernel application (2.9). By expanding the kernel definition in the resulting approximation, one sees that \tilde{g} can be computed from just the knowledge of the Chebyshev moments of g , $\int_{-1}^1 T_j(r) g(r) dr$.

In order to ensure that \tilde{g} approximates g and that \tilde{g} remains a probability density, the damping coefficients ρ_j must be selected carefully. One choice is to pick ρ_j to be the optimum of the following linear program: minimize a quantity called the squared kernel resolution

$$Q(K_m) \equiv \int_{-1}^1 \int_{-1}^1 \frac{K_m(x, r)(x-r)^2}{\sqrt{1-x^2}\sqrt{1-r^2}} dx dr,$$

²This multivariate generalization of the KPM was described briefly but not analyzed in [WWAF06].

which rewards K_m for being close to the identity operation, subject to

$$\int_{-1}^1 \frac{K_m(x, r)}{\sqrt{1-r^2}} dr = 1, \quad K_m(x, r) \geq 0 \text{ for } x, r \in [-1, 1], \quad (2.11)$$

which ensure \tilde{g} is nonnegative and integrates to 1, i.e. is a probability density function. Those optimal values of ρ_j are obtained in, e.g., [WWAF06, RB25] and are given by

$$\rho_j = \frac{(m+2-j) \cos\left(\frac{j\pi}{m+2}\right) + \sin\left(\frac{j\pi}{m+2}\right) \cot\left(\frac{\pi}{m+2}\right)}{m+2}, \quad j = 0, 1, \dots, m. \quad (2.12)$$

A different derivation, which emphasizes trigonometric polynomials rather than algebraic polynomials, for these same coefficients was also obtained by Rivlin [Riv81].

KPM posits that \tilde{g} is close to g in the earth mover's distance. To understand why, we summarize several results from [Riv81, Section 1.1.2], which control how K_m acts on Lipschitz test functions. When reading that reference, one should note that the coefficients in [Riv81, Page 20] are not presented in the form of (2.12), but they are equivalent.³

Lemma 2.3. *Suppose f is an ℓ -Lipschitz function on $[-1, 1]$, and let the damping coefficients $\rho_0, \rho_1, \dots, \rho_m$ be selected such that K_m is positive and normalized. Then*

$$\|f - K_m f\|_{[-1, 1]} \leq \frac{\ell}{m} \left(1 + \frac{m\pi}{\sqrt{2}} \sqrt{1 - \rho_1}\right). \quad (2.13)$$

In particular, with the choice in (2.12), one obtains

$$\|f - K_m f\|_{[-1, 1]} \leq \frac{6\ell}{m}. \quad (2.14)$$

Remark 5. *The constant 6 in (2.14) may be improved, especially for large m ; however, a lower bound on the constant of $\pi/2$ is known [Jac12, Theorem VI].*

Remark 6. *Results similar to Lemma 2.3 of this kind are often called ‘‘Jackson-type theorems,’’ after Dunham Jackson, who showed that algebraic or trigonometric polynomials can be used to approximate Lipschitz functions with a sup-norm error of $O(1/m)$. Although the kernel K_m that we use here arising from the choice (2.12) is sometimes called ‘‘Jackson’s kernel’’ in the literature, we note that Jackson’s original construction as given in [Jac12] is not exactly equivalent. Still, Jackson proves that the kernel he introduced in [Jac12] obtains an analogous result. The kernel of [Jac12] is also analyzed in [BKM22, Appendix C]. For practical purposes, differences between the kernels seem to be negligible.*

Lemma 2.3 directly allows one to control the earth mover’s distance between the measures with densities g and \tilde{g} : fix a 1-Lipschitz test function f and observe

$$\int_{-1}^1 f(x)\tilde{g}(x) dx = \int_{-1}^1 \int_{-1}^1 \frac{f(x)K_m(x, r)g(r)}{\sqrt{1-x^2}} dr dx = \int_{-1}^1 (K_m f)(r)g(r) dr. \quad (2.15)$$

Since Lemma 2.3 implies f and $K_m f$ are close, this means that the integrations of f against g and \tilde{g} are close, which is precisely what is required to show that g and \tilde{g} are close in the earth mover’s distance. From (2.15) we can directly bound:

$$\mathcal{W}_1(g, \tilde{g}) = \sup_{\|f\|_{\text{Lip}} \leq 1} \int_{-1}^1 f(z)(g(z) - \tilde{g}(z)) dx = \sup_{\|f\|_{\text{Lip}} \leq 1} \int_{-1}^1 (f(z) - K_m f(z))g(z) dx \leq \frac{6}{m},$$

where the last inequality follows from (2.14).

³Make the substitutions $x = \cos \theta$ and $r = \cos \phi$ into our definition of K_m and use the property of Chebyshev polynomials that $T_j(\cos \theta) = \cos(j\theta)$.

One final observation of single-variable KPM before generalizing to multiple variables is that K_m is a contraction in the following sense.

Lemma 2.4. *Let K_m be the polynomial kernel defined by (2.8), (2.9), and (2.12). K_m has the following contraction property:*

$$\|K_m f\|_{[-1,1]} \leq \|f\|_{[-1,1]}. \quad (2.16)$$

Proof. Simply expand the definition to see that

$$\begin{aligned} \|K_m f\|_{[-1,1]} &= \max_{x \in [-1,1]} \int \frac{|K_m(x, r)f(r)|}{\sqrt{1-r^2}} dr \\ &\leq \max_{r \in [-1,1]} |f(r)| \max_{x \in [-1,1]} \int_{-1}^1 \frac{K_m(x, r)}{\sqrt{1-r^2}} dr \\ &= \|f\|_{[-1,1]}. \end{aligned}$$

□

The preceding discussion for the univariate KPM generalizes in a straightforward way to multivariate problems in $d \geq 1$ variables on the hypercube $[-1, 1]^d$. In the following we state multivariate analogues of Lemma 2.3 and (2.15), and we consequently bound the KPM approximation error in earth mover's distance in Theorem 2.2.⁴ Define for a multivariate function $f(x_1, x_2, \dots, x_d)$ on $[-1, 1]^d$ the component Lipschitz constants

$$\text{Lip}_j(f) \equiv \sup_{\substack{x_k \in [-1,1] \\ k \neq j}} \sup_{\substack{u, v \in [-1,1] \\ u \neq v}} \frac{|f(x_1, \dots, x_{j-1}, u, x_{j+1}, \dots, x_d) - f(x_1, \dots, x_{j-1}, v, x_{j+1}, \dots, x_d)|}{|u - v|}.$$

$\text{Lip}_j(f)$ is the maximal Lipschitz constant of f when viewed as a function of x_j with all other variables fixed.

The one-dimensional kernel K_m defined by (2.9) extends to a d dimensional analog via the tensor product, i.e., let $\otimes^d K_m = K_m \otimes K_m \otimes \dots \otimes K_m$ be the integral operator defined by applying the one-dimensional Jackson operator separately in each coordinate:

$$(\otimes^d K_m f)(x_1, \dots, x_d) = \int_{[-1,1]^d} f(y_1, \dots, y_d) \prod_{j=1}^d \left(\frac{K_m(x_j, y_j)}{\sqrt{1-y_j^2}} \right) dy_1 \dots dy_d. \quad (2.17)$$

Then we have the following multivariate extension of Lemma 2.3.

Theorem 2.5 (Multivariate Jackson). *Let $f(x_1, x_2, \dots, x_d)$ be a continuous function on $[-1, 1]^d$ and suppose there exist positive constants ℓ_1, \dots, ℓ_d such that $\text{Lip}_j(f) \leq \ell_j$ for $j = 1, 2, \dots, d$. Then*

$$\|f - (\otimes^d K_m)f\|_{[-1,1]^d} \leq \frac{6}{m} \sum_{j=1}^d \ell_j, \quad (2.18)$$

Proof. We prove first $d = 2$ for simplicity, and then describe the general proof afterward. We use $(I \otimes K_m)$ to indicate the linear integral operator that applies Jackson's one-dimensional kernel on the second variable:

$$((I \otimes K_m)f)(x, y) = \int_{-1}^1 f(x, s) \frac{K_m(y, s)}{\sqrt{1-s^2}} ds.$$

⁴While multivariate extensions of Jackson's theorem have been known for a long time (c.f. [Tim63, NS64]), they do not seem to be well-known and we provide a simple proof of Theorem 2.5 using known properties of the one-dimensional K_m .

The operator $(K_m \otimes I)$ is defined similarly as operating on the first variable. Under this notation, we may factor the iterated integrals via $(K_m \otimes K_m) = (I \otimes K_m)(K_m \otimes I) = (K_m \otimes I)(I \otimes K_m)$. Then observe

$$\begin{aligned} \|f - (K_m \otimes K_m)f\|_{[-1,1]^2} &= \|f - (I \otimes K_m)f + (I \otimes K_m)(f - (K_m \otimes I)f)\|_{[-1,1]^2} \\ &\leq \|f - (I \otimes K_m)f\|_{[-1,1]^2} + \|(I \otimes K_m)(f - (K_m \otimes I)f)\|_{[-1,1]^2} \\ &\leq \|f - (I \otimes K_m)f\|_{[-1,1]^2} + \|f - (K_m \otimes I)f\|_{[-1,1]^2} \\ &\leq \frac{6\ell_2}{m} + \frac{6\ell_1}{m}. \end{aligned}$$

The first equality comes from adding and subtracting $(I \otimes K_m)f$. The inequality on the second line is an application of the triangle inequality. For the third line, let x be fixed and define the one-dimensional slice $y \mapsto h_x(y) = (f - (K_m \otimes I)f)(x, y)$, so that we can apply Lemma 2.4 by observing

$$\|(I \otimes K_m)(f - (K_m \otimes I)f)\|_{[-1,1]^2} = \sup_{x \in [-1,1]} \|K_m h_x\| \leq \sup_{x \in [-1,1]} \|h_x\|_{[-1,1]} = \|f - (K_m \otimes I)f\|_{[-1,1]^2}.$$

Finally, the fourth line is a consequence of applying the one-dimensional Lemma 2.3 to the one-dimensional slices analogous to the argument applied to h_x above. More precisely, for fixed x , observe that $y \mapsto f_x(y) \equiv f(x, y)$ is an ℓ_2 -Lipschitz function of y , so that

$$\|f - (I \otimes K_m)f\|_{[-1,1]^2} = \sup_{x \in [-1,1]} \|f_x - K_m f_x\|_{[-1,1]} \leq \frac{6\ell_2}{m}.$$

The same argument may be applied to the analogously defined f_y to control the second term in the third line.

The $d > 2$ case follows from a similar argument. If we let $P_j = (\otimes^j K_m) \otimes (\otimes^{d-j} I)$ be the operator which applies K_m in the first j coordinates, we apply the triangle inequality to

$$f - (\otimes^d K_m)f = \sum_{j=1}^d P_{j-1}f - P_jf,$$

whose d terms lead to the factor of $\sum_{j=1}^d \ell_j$ in the result. \square

Remark 7. If f is ℓ -Lipschitz with respect to any p -norm then $\sum_j \text{Lip}_j(f) \leq d\ell$.

Remark 8. Analogous results can be obtained by using a different kernel K_m corresponding to a different choice of the ρ_j , such as Jackson's original construction [Jac12].

The calculation in (2.15) immediately generalizes to the multivariate setting. For simplicity, we only state it for $d = 2$ variables, which is all we need for complex measures.

Lemma 2.6. Let μ be a probability distribution on $[-1, 1]^2$. Define further the approximation

$$d\mu_{\text{KPM}}(x, y) \equiv \frac{p_m(x, y)}{\sqrt{1-x^2}\sqrt{1-y^2}} dx dy$$

where

$$p_m(x, y) = \sum_{i=0}^m \sum_{j=0}^m \rho_i \rho_j \Gamma_{ij} \tilde{T}_i(x) \tilde{T}_j(y), \quad \Gamma_{ij} = \int_{-1}^1 \int_{-1}^1 \tilde{T}_i(x) \tilde{T}_j(y) d\mu(x, y).$$

Then μ_{KPM} is a probability distribution and

$$\int_{[-1,1]^2} f(x, y) d\mu_{\text{KPM}}(x, y) = \int_{[-1,1]^2} ((K_m \otimes K_m)f)(x, y) d\mu(x, y).$$

Remark 9. We observe that Lemma 2.6 shows that the ansatz (2.6) constructs a measure μ_{KPM} whose action on a test function f reproduces the action of the true measure μ on the test function after polynomial smoothing, $K_m f$. This is independent of the choice of the ρ_j .

We may now prove Theorem 2.2.

Proof of Theorem 2.2. Let \mathcal{F} denote the class of 1-Lipschitz functions with respect to the Euclidean norm on $[-1, 1]^2$. Then applying Lemma 2.6 and Theorem 2.5 we obtain

$$\begin{aligned} \mathcal{W}_1(\mu, \mu_{\text{KPM}}) &= \sup_{f \in \mathcal{F}} \left(\int_{-1}^1 \int_{-1}^1 f(x, y) d\mu(x, y) - \int_{-1}^1 \int_{-1}^1 f(x, y) d\mu_{\text{KPM}}(x, y) \right) \\ &= \sup_{f \in \mathcal{F}} \left(\int_{-1}^1 \int_{-1}^1 (f - (K_m \otimes K_m)f)(x, y) d\mu(x, y) \right) \\ &\leq \frac{12}{m}. \end{aligned}$$

□

To be explicit about the algorithmic importance of Theorem 2.2, it shows that one may compute an approximation to a measure with just knowledge of its Chebyshev moments. The previous section showed how to compute those moments of $\mu_{A,b}$, and so we have an algorithm, Algorithm 2, for computing an approximation $\hat{\mu}_{A,b}$ of it.

Algorithm 2 WSPEC DENS(A, b, m)

Estimates weighted spectral density $\hat{\mu}_{A,b}$ using Algorithm 1 and Jackson's coefficients (2.12).

Input: $A \in \mathbb{C}^{n \times n}$, $\|A\| \leq 1$, $b \in \mathbb{C}^n$, $\|b\| = 1$.

Output: $\hat{\mu}_{A,b}$ approximating $\mu_{A,b}$ with guarantee from Theorem 2.2.

1: $\Gamma \leftarrow \text{CHEB MOMENTS}(A, b, m)$

2: $p_{\text{KPM}}(x, y) \leftarrow \sum_{j=0}^m \sum_{k=0}^m \rho_j \rho_k \Gamma_{jk} \tilde{T}_j(x) \tilde{T}_k(y)$

3: **return** $\frac{p_{\text{KPM}}(x, y)}{\sqrt{1-x^2}\sqrt{1-y^2}} dx dy$.

2.3. Selection of weight vector. For any fixed vector b , the previous two subsections show how to compute an approximation of the weighted spectral density $\mu_{A,b}$. Our goal, however, is to approximate the unweighted spectral density μ_A . Observe that if b is a random vector sampled from an isotropic distribution, then

$$\mathbb{E} \mu_{A,b} = \mu_A.$$

This section argues that $\mu_{A,b}$ exhibits sufficient concentration when b is sampled uniformly at random from the unit sphere. The final algorithm will then be Algorithm 3.

Algorithm 3 SPEC DENS(A, m)

Estimates spectral density $\hat{\mu}_A$ using Algorithm 2.

Input: $A \in \mathbb{C}^{n \times n}$, $\|A\| \leq 1$.

Output: $\hat{\mu}_A$ approximating μ_A with guarantee from Theorem 1.1.

1: Sample $b \sim \text{Unif}(\mathcal{S}^{2n-1})$

2: **return** WSPEC DENS(A, b, m).

Our analysis is in two parts, both appealing to known results. Claim **I** is that the Chebyshev moments of $\mu_{A,b}$ are highly concentrated around the corresponding moments of μ_A . Claim **II** is that this suffices to imply $\mu_{A,b}$ is close to μ_A in earth mover's distance with high probability.

Claim **I** is a consequence of the following concentration result for Lipschitz functions on the sphere. We identify the unit sphere living in \mathbb{C}^n with the unit sphere \mathcal{S}^{2n-1} living in \mathbb{R}^{2n} .

Lemma 2.7 ([BCG23, Theorem 10.3.1]). *Suppose $g : \mathcal{S}^{2n-1} \rightarrow \mathbb{R}$ has ℓ_2 Lipschitz constant bounded by L . Let $b \sim \mathbf{Unif}(\mathcal{S}^{2n-1})$. Then*

$$\Pr(|g(b) - \mathbb{E}[g(b)]| \geq r) \leq 2e^{-(2n-1)r^2/(2L^2)}.$$

We can use this result to bound the fluctuations in the Chebyshev moments of $\mu_{A,b}$ around its expectation μ_A .

Proposition 2.8. *Let Γ_{jk} and Γ'_{jk} be the mixed Chebyshev moments of μ_A and $\mu_{A,b}$, respectively, for $0 \leq j, k \leq m$. Let $b \sim \mathbf{Unif}(\mathcal{S}^{2n-1})$. Then for $0 < \delta < 1$,*

$$\mathbb{P}\left(\max_{(j,k) \neq 0} |\Gamma_{jk} - \Gamma'_{jk}| \geq \sqrt{\frac{8}{2n-1} \log\left(\frac{2(m+1)^2 - 2}{\delta}\right)}\right) \leq \delta.$$

Proof. For any $f : \mathbb{C} \rightarrow \mathbb{R}$, and $b \sim \mathbf{Unif}(\mathcal{S}^{2n-1})$, we have

$$\int f(z) d\mu_A(z) = \frac{1}{n} \operatorname{tr}(f(A)), \quad \int f(z) d\mu_{A,b}(z) = b^* f(A) b.$$

Moreover, $\mathbb{E}bb^* = \frac{1}{n}I$ so

$$\mathbb{E}[b^* f(A) b] = \mathbb{E}[\operatorname{tr}(f(A)bb^*)] = \frac{1}{n} \operatorname{tr}(f(A)).$$

Observe also that $g(b)$ defined by $b \mapsto g(b) = b^* f(A) b \in \mathbb{R}$ is a Lipschitz function with Lipschitz constant bounded by $2\|f(A)\|_2$ since

$$|b^* f(A) b - c^* f(A) c| = |b^* f(A)(b - c) - (c - b)^* f(A) c| \leq 2\|f(A)\| \|b - c\|.$$

Thus, according to Lemma 2.7,

$$\mathbb{P}\left(\left|b^* f(A) b - \frac{1}{n} \operatorname{tr}(f(A))\right| \geq t\right) \leq 2e^{-(2n-1)t^2/(8\|f(A)\|^2)}. \quad (2.19)$$

Let us now specialize to the case $f_{ij}(z) = \tilde{T}_i(x)\tilde{T}_j(y)$ so that

$$\Gamma_{ij} - \Gamma'_{ij} = \frac{1}{n} \operatorname{tr}(f_{ij}(A)) - b^* f_{ij}(A) b,$$

to which we may apply (2.19) with a union bound over the $(m+1)^2 - 1$ polynomials. Note also that $\|f_{ij}(A)\| \leq 1$, so for each $0 < \delta < 1$ we may set $t = \sqrt{\frac{8}{2n-1} \log\left(\frac{2(m+1)^2 - 2}{\delta}\right)}$ for the final result. \square

Claim **II** is proven in [MMRS25]. They show that if two measures have close Chebyshev moments, they must be close in earth mover's distance. Although the result is stated for d -dimensions in [MMRS25], we state it only for 2-dimensions here.

Lemma 2.9 ([MMRS25, Theorem 32]). *Suppose μ, ν are probability distributions on $[-1, 1]^2$ with normalized Chebyshev moments Γ_{kj} and Γ'_{kj} , respectively. If the distributions' normalized Chebyshev moments satisfy the bound*

$$\sum_{\substack{0 \leq k, j \leq m \\ (k, j) \neq (0, 0)}} \frac{1}{k^2 + j^2} (\Gamma_{kj} - \Gamma'_{kj})^2 \leq K^2, \quad (2.20)$$

for some constant $K > 0$, then one may bound the earth mover's distance

$$\mathcal{W}_1(\mu, \nu) \leq \frac{4C}{m} + \pi K,$$

where $\pi/2 \leq C \leq 3$.⁵

In order to combine Proposition 2.8 and Lemma 2.9 to imply a bound on $\mathcal{W}_1(\mu_A, \mu_{A,b})$, we must bound the sum over the terms of the form $(k^2 + j^2)^{-1}$. For this, we introduce the following lemma.

Lemma 2.10. *One may bound*

$$\sum_{\substack{0 \leq k, j \leq m \\ (k, j) \neq (0, 0)}} \frac{1}{k^2 + j^2} \leq 5 + \frac{\pi}{2} \log m.$$

Proof. We have

$$\sum_{\substack{0 \leq k, j \leq m \\ (k, j) \neq (0, 0)}} \frac{1}{k^2 + j^2} = 2 \sum_{k=1}^m \frac{1}{k^2} + \sum_{i=1}^m \sum_{j=1}^m \frac{1}{i^2 + j^2}.$$

The first term is bounded by $2\zeta(2) = \pi^2/3$ and the second term is bounded by

$$\sum_{k=1}^m \sum_{j=1}^m \frac{1}{k^2 + j^2} \leq \sum_{k=1}^m \int_0^\infty \frac{1}{k^2 + x^2} dx = \sum_{k=1}^m \frac{\pi}{2k} \leq \frac{\pi}{2} (1 + \log m).$$

Lastly, $\pi^2/3 + \pi/2 < 5$. □

2.4. Proof of Theorem 1.1. We now collect the results of the previous two subsections into a proof of our main result Theorem 1.1, which bounds the error in Algorithm 3.

Proof of Theorem 1.1. Denote the output of Algorithm 3 by $\hat{\mu}_A$, which is the KPM approximation to the random measure $\mu_{A,b}$. Then observe that with probability $1 - \delta$ we have that

$$\mathcal{W}_1(\mu_A, \hat{\mu}_A) \leq \mathcal{W}_1(\mu_A, \mu_{A,b}) + \mathcal{W}_1(\mu_{A,b}, \hat{\mu}_A) \leq O\left(\frac{1}{m}\right) + O\left(\sqrt{\frac{\log(m) \log(m/\delta)}{n}}\right).$$

The first term $\mathcal{W}_1(\mu_A, \mu_{A,b})$ is bounded via Lemma 2.9, within which the left-hand side of (2.20) is bounded with high probability using Hölder's inequality together with Lemma 2.10 and Proposition 2.8. The second term $\mathcal{W}_1(\mu_{A,b}, \hat{\mu}_A)$ is bounded by Theorem 2.2.

Moreover, the measure $\hat{\mu}_A$ may be (optionally) discretized using Gauss-Chebyshev quadrature on $(m+1)^2$ points. That is, let x_1, \dots, x_{m+1} and w_1, \dots, w_{m+1} be the nodes and weights of the Gauss quadrature associated with the Chebyshev measure. Recall that Gauss quadrature on k nodes is exact for polynomials of degree $2k - 1$. Then consider the discrete measure

$$\hat{\mu}_{A, \text{discrete}} := \sum_{i=1}^{m+1} \sum_{j=1}^{m+1} w_i w_j p_m(x_i, x_j) \delta_{(x_i, x_j)},$$

where p_m is as in (2.6). The exactness of Gauss quadrature ensures

$$\int_{[-1, 1]^2} T_i(x) T_j(y) (d\hat{\mu}_A - d\hat{\mu}_{A, \text{discrete}}) = 0$$

⁵We have modified the result of [MMRS25] by using $\pi/2 \leq C \leq 3$, as derived by Jackson in [Jac12], rather than the number 18 cited in Fact 12 of [MMRS25]. This improves the constant in [MMRS25] in the univariate case from $c \leq 36$ to $c \leq 6$. A similar gain may be obtained by replacing the 18 in Fact 3.2 of [BKM22], since both of these works rely on Jackson's construction in [Jac12].

for all $i, j \leq m$. When combined with Lemma 2.9, this implies that $\mathcal{W}_1(\widehat{\mu}_A, \widehat{\mu}_{A, \text{discrete}}) \leq 4C/m$, which merely alters the constant in the bound. \square

2.5. Algorithmic variants. We include here some comments about algorithmic variants that may be useful.

- (1) It is possible to consider an *adjoint-free* version of Algorithm 3 that avoids the use of the adjoint A^* , which is an important feature in some applications, particularly PDEs [BHOT24]. One way to accomplish this is to use $m - 1$ matrix-vector products to form the explicit Krylov matrix $\mathcal{K}_m = [b, Ab, \dots, A^{m-1}b]$. In analogy to (2.5), the Gram matrix $\mathcal{K}_m^* \mathcal{K}_m$ contains the “polyanalytic” mixed moments corresponding to polynomials in z and \bar{z} . Since there is an invertible linear change of coordinates from $z = x + iy, \bar{z} = x - iy$ into x, y , one could in principle pre-compute the change of basis coefficients from $z^k \bar{z}^l$ into the Chebyshev polynomials $\widetilde{T}_i(x) \widetilde{T}_j(y)$, though this may require high precision. One would likely use the Arnoldi procedure in practice rather than forming the ill-conditioned explicit Krylov matrix \mathcal{K}_m .
- (2) It is possible to consider a low-memory variant of Algorithm 3. We could compute X^*Y in Algorithm 1 one row at a time, which requires $O(m^2)$ matvecs but only $O(n)$ storage.
- (3) Say the spectrum of A is somewhat oblong, e.g., the range of real parts is decently larger than the range of imaginary parts. Then rather than computing m matvecs in both the x and y coordinates, we could decouple the x and y directions into m_1 and m_2 matvecs, so that $X^*Y \in \mathbb{R}^{m_1 \times m_2}$.
- (4) If the $\widetilde{O}(n^{-1/2})$ error from Theorem 1.1 is too high, one could apply Algorithm 3 to the block matrix $A \otimes I_s \in \mathbb{C}^{ns \times ns}$, where I_s is the $s \times s$ identity matrix and \otimes is the Kronecker product. This would decrease the error due to randomness to $O((ns)^{-1/2})$. Note that this is equivalent to replacing $\mu_{A,b}$ with a weighted average of s copies of $\mu_{A,b}$ with independent b . By picking $s = \lceil m^{2/3}/n^{1/3} \rceil$, this improves the overall error to $\widetilde{O}(1/(nm)^{1/3})$ when $m = \Omega(\sqrt{n})$.
- (5) Because of the decoupling of the x - and y -coordinate directions, the resulting measure seems to exhibit axis-aligned artifacts. One can remove these artifacts by “rotating” the input, i.e. applying the algorithm to $e^{i\theta}A$ for a few angles θ , rotating back the returned measures, and averaging. That is, use the estimate $\frac{1}{|\Theta|} \sum_{\theta \in \Theta} \widehat{\mu}_{e^{i\theta}A} \circ (z \mapsto e^{i\theta}z)$ for some set of angles Θ . It may also be possible to work with other polynomial bases, such as disk polynomials [DMN25].

3. LOWER BOUNDS FOR SPECTRAL DENSITY ESTIMATION

In this section, we prove a lower bound for all matrix-vector query algorithms. That is, we consider any algorithm interacts with the input matrix A only by constructing query vectors x and receiving Ax and A^*x . We show that for large matrices, Algorithm 3 makes the asymptotically lowest possible number of queries to achieve its error guarantee stated in Theorem 1.1. In fact, we show the lower bound holds even for algorithms only required to succeed when A is real symmetric.

The main result for this purpose, Proposition 3.1, is a characterization of all adaptive matrix-vector query algorithms which make a small number of queries relative to the matrix size. In particular, closely following ideas of Chewi et al [CdDPL⁺24], we show that any adaptive algorithm which makes at most m matvec queries to large A learns only a small amount of information about A beyond possibly the values of $\text{tr}(A^j)$ for $j = 1, \dots, 4m$. We believe this result may be independently useful for producing related lower bounds.

The second step of our lower bound, Lemma 3.2, is to show the values of $\text{tr}(A^j)$ for $j = 1, \dots, 4m$ is not enough information about A to reconstruct its spectral density to $o(1/m)$ error. In particular, we produce a family of matrices with far spectral densities, but matching values of $\text{tr}(A^j)$.

In order to argue that matvec algorithms do not learn much information beyond $\text{tr}(A^j)$ for $j = 1, \dots, 4m$, we show that the output of any matvec algorithm can be simulated with high probability

by an algorithm which takes as input just the values of $\text{tr}(A^j)$ for $j \leq 4m$. To make this formal, we need the notion of *total variation distance*,

$$\text{dist}_{\text{TV}}(X, Y) = \inf_{(X, Y) \sim \gamma} \Pr(X \neq Y),$$

where the infimum is over all joint distributions γ whose marginals are the laws of X and Y . From this formulation of total variation distance, it is clear for every f that

$$\text{dist}_{\text{TV}}(f(X), f(Y)) \leq \text{dist}_{\text{TV}}(X, Y). \quad (3.1)$$

The notation $X \stackrel{d}{=} Y$ means X and Y are equal in distribution, or equivalently $\text{dist}_{\text{TV}}(X, Y) = 0$. In this section, \otimes denotes the Kronecker product, $\mathbf{1}_p \in \mathbb{R}^p$ the column vector of 1s, I_m the $m \times m$ identity, and e_j the j th elementary basis column vector of the appropriate dimension.

Proposition 3.1. *Let $v_0^{\text{alg}} \in \mathbb{R}^n$ be any fixed vector and let $v_j^{\text{alg}} : (\mathbb{R}^n)^j \rightarrow \mathbb{R}^n$ be measurable functions. Fix any real symmetric matrix $B \in \mathbb{R}^{p \times p}$ and let $A = U(I_{n/p} \otimes B)U^\top$ be the random matrix where U is a Haar-distributed random orthogonal matrix. If*

$$v_0 = v_0^{\text{alg}}, \quad v_j = v_j^{\text{alg}}(Av_0, Av_1, \dots, Av_{j-1}).$$

then for every $f_{\text{alg}} : (\mathbb{R}^n)^m \rightarrow \mathcal{X}$ (for any output space \mathcal{X}), there is a measurable function $f_{\text{sim}} : \mathbb{R}^{4m} \rightarrow \mathcal{X}$ such that

$$\text{dist}_{\text{TV}}(f_{\text{alg}}(Av_0, \dots, Av_{m-1}), f_{\text{sim}}(\text{tr}(A), \text{tr}(A^2), \dots, \text{tr}(A^{4m}))) = O\left(\frac{m^{3/2}p^{1/2}}{n^{1/2}}\right).$$

3.1. Proof of Theorem 1.2 assuming Proposition 3.1. In the following lemma we give a family of matrices with matching traces of powers but separated spectral densities.

Lemma 3.2. *Let $B_m \in \mathbb{R}^{m \times m}$ be the diagonal matrix with $(B_m)_{jj} = \cos\left(\frac{2j-1}{2m}\pi\right)$. Let k, k', m, m', n be such that $km = k'm' = n$. Then $\text{tr}((B_m \otimes I_k)^i) = \text{tr}((B_{m'} \otimes I_{k'})^i)$ for $i \leq \min(m, m')$ but if $m \neq m'$ then*

$$\mathcal{W}_1(\mu_{B_m \otimes I_k}, \mu_{B_{m'} \otimes I_{k'}}) = \Omega\left(\frac{1}{\min(m, m')}\right).$$

Proof. By Chebyshev quadrature, one has $\frac{1}{m} \text{tr}(B_m^i) = \frac{1}{\pi} \int_{-1}^1 \frac{x^i}{\sqrt{1-x^2}} dx$ for $i \leq m$. In particular,

$$\text{tr}(B_{km}^i) = \frac{km}{\pi} \int_{-1}^1 \frac{x^i}{\sqrt{1-x^2}} dx = k \text{tr}(B_m^i) = \text{tr}((B_m \otimes I_k)^i).$$

Since $km = k'm'$, we thus have $\text{tr}((B_m \otimes I_k)^i) = \text{tr}((B_{m'} \otimes I_{k'})^i)$ for $i \leq \min(m, m')$. On the other hand, the earth mover's distance between the spectral measures of $B_m \otimes I_k$ and $B_{m'} \otimes I_{k'}$ is the minimum ℓ_1 cost of a coupling between the diagonal entries of those two matrices. Specifically, denoting $q_r(x) = \frac{\lfloor rx \rfloor + \frac{1}{2}}{r}$ and replacing $x = j/n$, we have

$$\begin{aligned} W &\equiv \mathcal{W}_1(\mu_{B_m \otimes I_k}, \mu_{B_{m'} \otimes I_{k'}}) = \frac{1}{n} \sum_{j=1}^n |(B_m \otimes I_k)_{jj} - (B_{m'} \otimes I_{k'})_{jj}| \\ &= \frac{1}{n} \sum_{j=0}^{n-1} \left| \cos\left(\frac{2\lfloor j/k \rfloor + 1}{2m}\pi\right) - \cos\left(\frac{2\lfloor j/k' \rfloor + 1}{2m'}\pi\right) \right| \\ &= \int_0^1 |\cos(\pi q_m(x)) - \cos(\pi q_{m'}(x))| dx \\ &= \int_0^1 \left| \int_{q_{m'}(x)}^{q_m(x)} \sin(\pi t) dt \right| dx \end{aligned} \quad (3.2)$$

Let S be the set

$$S = \{(x, t) \in [0, 1]^2 : q_m(x) \leq t \leq q_{m'}(x) \text{ or } q_{m'}(x) \leq t \leq q_m(x)\}$$

so that W is simply the integral of $\sin(\pi t)$ over S . Let

$$w(t) = \max\{x : (x, t) \in S\} - \min\{x : (x, t) \in S\}.$$

so that (3.2) reduces to

$$W = \int_0^1 w(t) \sin(\pi t) dt. \quad (3.3)$$

Since $|q_r(x) - x| \leq \frac{1}{2r}$, we have $w(t) \leq \frac{1}{m} + \frac{1}{m'}$. Suppose $\int_0^1 w(t) \geq c(\frac{1}{m} + \frac{1}{m'})$ for some positive constant c . Then (3.3) would be minimized for $w(t) = \frac{1}{m} + \frac{1}{m'}$ for $t \leq c$ and $w(t) = 0$ for $t > c$. In particular,

$$W \geq \left(\frac{1}{m} + \frac{1}{m'}\right) \int_0^c \sin(\pi t) dt \geq \frac{1}{\pi} \left(\frac{1}{m} + \frac{1}{m'}\right) (1 - \cos(\pi c)).$$

It therefore suffices to prove this supposition. To this end, note

$$\begin{aligned} \int_0^1 w(t) dt &= \int_0^1 |q_m(x) - q_{m'}(x)| dx \\ &= \frac{1}{n} \int_0^1 \left| k \lfloor mx \rfloor - k' \lfloor m'x \rfloor + \frac{k - k'}{2} \right| dx \\ &= \frac{1}{n^2} \sum_{j=0}^{n-1} \left| k \lfloor j/k \rfloor - k' \lfloor j/k' \rfloor + \frac{k - k'}{2} \right| dx \\ &= \frac{1}{n^2} \sum_{j=0}^{n-1} \left| (j \bmod k) - (j \bmod k') + \frac{k - k'}{2} \right| dx \end{aligned} \quad (3.4)$$

Let $d = \gcd(k, k')$. Then we group the terms in the sum based on the residue of $j = id + r$ modulo d .

$$\begin{aligned} \int_0^1 w(t) dt &= \frac{1}{n^2} \sum_{r=0}^{d-1} \sum_{i=0}^{n/d-1} \left| (j \bmod k) - (j \bmod k') + \frac{k - k'}{2} \right| \\ &= \frac{d^2}{n^2} \sum_{i=0}^{n/d-1} \left| (i \bmod k/d) - (i \bmod k'/d) + \frac{k - k'}{2d} \right|. \end{aligned} \quad (3.5)$$

Notice that the map $i \mapsto (i \bmod k/d, i \bmod k'/d)$ is a $n/\text{lcm}(k, k')$ -to-1 map from $\{0, \dots, n/d - 1\}$ to $\{0, \dots, k/d - 1\} \times \{0, \dots, k'/d - 1\}$. That is,

$$\int_0^1 w(t) dt = \frac{d^2}{n^2} \frac{n}{\text{lcm}(k, k')} \sum_{x=0}^{k/d-1} \sum_{y=0}^{k'/d-1} \left| x - y + \frac{k - k'}{2d} \right| = \Theta \left(\frac{d^2}{n^2} \cdot \frac{n}{\text{lcm}(k, k')} \cdot \frac{kk'k + k'}{d^2} \right).$$

Which, since $kk'/d = \text{lcm}(k, k')$, simplifies to $\Theta(\frac{1}{m} + \frac{1}{m'})$ as desired. \square

Proof of Theorem 1.2. Let $B_{(\cdot)}$ be as in Lemma 3.2. Let $m_1 < \dots < m_s$ be some divisors of n . Set

$$A_j = B_{m_j} \otimes I_{n/m_j}$$

so that A_j is an $n \times n$ matrix with a spectral measure matching that of B_{m_j} . By Lemma 3.2,

$$\text{tr}(A_j^i) = \text{tr}(A_{j'}^i)$$

for all $i \leq m_1$ and

$$\mathcal{W}_1(\mu_{A_j}, \mu_{A_{j'}}) > \frac{1}{Cm_{s-1}}.$$

Let X_j be the result of some algorithm making at most $m_1/4$ matrix vector queries on input UA_jU^\top . By Proposition 3.1, each X_j are all within $O\left(\frac{m_1^{3/2}m_s^{1/2}}{n^{1/2}}\right)$ in total variation distance to the same random variable (produced by f_{sim}). Thus if the target accuracy ε is less than $\frac{1}{2Cm_{s-1}}$, then with probability $1 - O\left(\frac{m_1^{3/2}m_s^{1/2}}{n^{1/2}}\right)$ over the randomness of U , the algorithm must fail for all but one of the s inputs. The algorithm therefore cannot succeed with overall probability more than

$$\frac{1}{s} + O\left(\frac{m_1^{3/2}m_s^{1/2}}{n^{1/2}}\right).$$

Now, simply set $m_j = \frac{1}{\varepsilon} + j - 1$, $s = 1/\delta$, and $n = \prod_{j=1}^s m_j$. \square

3.2. Proof of Proposition 3.1. We now show how to simulate any matvec algorithm with an algorithm taking as input the traces of powers of A .

Lemma 3.3 ([BN23, Lemma 6.3]). *Let v_j^{alg} , v_j , and A be as in Proposition 3.1. Let $z_j \sim \mathcal{N}(0, I)$. Then there is a function g depending only on v^{alg} such that*

$$g(\{A^i z_j\}_{i+j \leq m}) \stackrel{d}{=} \{Av_j\}_{j < m}.$$

Lemma 3.4. *Let A be as in Proposition 3.1 and $z_j \sim \mathcal{N}(0, I)$. Then there is a function h such that*

$$h(\{z_{j'}^T A^i z_j\}_{(i,j,j') \in H}) \stackrel{d}{=} \{A^i z_j\}_{i+j \leq m}$$

where

$$H = \{(i, j, j') : i + j + j' \leq 2m, j, j' \leq m\}.$$

Proof. Since A is real symmetric, $\{z_j^T A^i z_{j'}\}_{(i,j,j') \in H}$ is equivalent (up to some duplicated entries) to the set of inner products $\{(A^{i'} z_{j'})^\top (A^i z_j)\}_{i+j \leq m, i'+j' \leq m}$ which h may compute a factorization of. For two collections of vectors arranged into the columns of matrices X and Y , if all the inner products agree, i.e. $X^*X = Y^*Y$, then there exists a orthogonal U such that $X = UY$. Finally note by construction of A that $\{UA^i z_j\}_{i+j \leq m} \stackrel{d}{=} \{A^i z_j\}_{i+j \leq m}$, so applying a random orthogonal matrix to the factorization found by h achieves the desired distribution. \square

The following Lemma 3.5 is a non-asymptotic version of a lemma appearing several places, such as [RR18, Theorem 1.2], [BDER16, Theorem 7], and [JL13, Theorem 1]. This non-asymptotic version was first stated without a proof in [CdDPL⁺24, Lemma 5.3], and then again in a concurrent preprint [DEM26]. A non-trivial refinement of the existing asymptotic proofs verifies the non-asymptotic version.⁶

Lemma 3.5 will be used to argue that the bilinear forms $z_{j'}^\top A^i z_j$ appearing in Lemma 3.4 look like Gaussian fluctuations around their means.

Lemma 3.5. *Let $Z \in \mathbb{R}^{t \times m}$ and $X \in \mathbb{R}^{m \times m}$ have i.i.d. $\mathcal{N}(0, 1)$ entries. Then*

$$\text{dist}_{\text{TV}}\left(Z^\top Z, tI + \sqrt{\frac{t}{2}}(X + X^\top)\right) \leq O\left(\frac{m^{3/2}}{t^{1/2}}\right).$$

⁶A full proof was provided in private correspondence with the authors of [DEM26].

The matrix $Z^\top Z$ is called a Wishart matrix.

Lemma 3.6. *Let $Z \in \mathbb{R}^{t \times p \times m}$ have i.i.d. $\mathcal{N}(0, 1)$ entries. Fix any symmetric $B \in \mathbb{R}^{p \times p}$ and let $A = B \otimes I_t$. Then there is a Gaussian matrix G with correlated entries entirely characterized by $\{\text{tr}(A^i)\}_{i \leq \ell}$ and $\{\text{tr}(A^{2i})\}_{i \leq \ell}$ such that*

$$\text{dist}_{\text{TV}}((Z^\top AZ \quad \cdots \quad Z^\top A^\ell Z), G) \leq O\left(\frac{m^{3/2}}{t^{1/2}}\right).$$

Proof. Let $\lambda_1, \dots, \lambda_p$ be the eigenvalues of B and denote the Vandermonde matrix

$$V = \begin{bmatrix} \lambda_1 & \cdots & \lambda_1^\ell \\ \vdots & & \vdots \\ \lambda_p & \cdots & \lambda_p^\ell \end{bmatrix}.$$

Let $B = UDU^\top$ be the diagonalization of B . Set $\tilde{Z} = (U^\top \otimes I_t)Z$.

$$Z^\top A^i Z = Z^\top (B^i \otimes I_t)Z = \tilde{Z}^\top (U^\top \otimes I_t)(B^i \otimes I_t)(U \otimes I_t)\tilde{Z} = \tilde{Z}^\top (D^i \otimes I_t)\tilde{Z}$$

If one breaks \tilde{Z} into a column of p many $t \times m$ blocks, then this becomes simply

$$Z^\top A^i Z = \sum_{j=1}^p \lambda_j^i \tilde{Z}_j^\top \tilde{Z}_j.$$

Rearranging this equality gives

$$(Z^\top AZ \quad \cdots \quad Z^\top A^\ell Z) \stackrel{d}{=} \begin{pmatrix} \tilde{Z}_1^\top \tilde{Z}_1 & \cdots & \tilde{Z}_p^\top \tilde{Z}_p \end{pmatrix} (V \otimes I_m).$$

Notably, by rotational invariance of Z , the matrices $\tilde{Z}_j^\top \tilde{Z}_j$ on the right are all independent Wishart matrices. Set

$$G = \left[(t\mathbf{1}_p^\top \otimes I_m) + \sqrt{t/2} (X_1 + X_1^\top \quad \cdots \quad X_p + X_p^\top) \right] (V \otimes I_m)$$

where X_j has i.i.d. $\mathcal{N}(0, 1)$ entries and note that the entries of G are indeed Gaussian random variables as they are affine functions of the entries of the X_j . By (3.1) and Lemma 3.5, we have

$$\begin{aligned} & \text{dist}_{\text{TV}}((Z^\top AZ \quad \cdots \quad Z^\top A^\ell Z), G) \\ & \leq \text{dist}_{\text{TV}}\left(\begin{pmatrix} \tilde{Z}_1^\top \tilde{Z}_1 & \cdots & \tilde{Z}_p^\top \tilde{Z}_p \end{pmatrix}, (t\mathbf{1}^\top \otimes I) + \sqrt{t/2} (X_1 + X_1^\top \quad \cdots \quad X_p + X_p^\top)\right) \\ & \leq O\left(\frac{m^{3/2}}{t^{1/2}}\right). \end{aligned} \tag{3.6}$$

Let's now compute the first two moments of G . Since $\mathbb{E}(X) = 0$, the first moment is simply

$$\mathbb{E}(G) = (t\mathbf{1}_p^\top \otimes I_m)(V \otimes I_m) = (t\mathbf{1}_p^\top V \otimes I_m) = ((\text{tr}(A) \quad \cdots \quad \text{tr}(A^\ell)) \otimes I_m)$$

so is determined by $\{\text{tr}(A^i)\}_{i \leq \ell}$. The covariance of the entries is given by

$$\begin{aligned} \text{Cov}((1 \otimes e_b^\top)G(e_i \otimes e_a), (1 \otimes e_{b'}^\top)G(e_{i'} \otimes e_{a'})) &= \frac{t}{2} \sum_{j=1}^p \sum_{j'=1}^p \lambda_j^i \lambda_{j'}^{i'} \mathbb{E}(e_b^\top (X_j + X_j^\top) e_a e_{b'}^\top (Y_{j'} + Y_{j'}^\top) e_{a'}) \\ &= \frac{t}{2} \sum_{j=1}^p \lambda_j^{2i} \mathbb{E}(e_b^\top (X_j + X_j^\top) e_a e_{b'}^\top (Y_j + Y_j^\top) e_{a'}) \\ &= \frac{1}{2} \text{tr}(A^{2i}) \mathbb{E}(e_b^\top (X_1 + X_1^\top) e_a e_{b'}^\top (X_1 + X_1^\top) e_{a'}) \end{aligned} \tag{3.7}$$

so is determined by $\{\text{tr}(A^{2i})\}_{i \leq \ell}$. \square

Proof of Proposition 3.1. We describe the operations performed by f_{sim} . First, because it is given $\text{tr}(A^i)$ for $i \leq 4m$, it is able to sample a Gaussian matrix G which, by Lemma 3.6 for $\ell = 2m$ and $t = n/p$ and the definition of total variation distance, can be written as the mixture

$$G = \begin{cases} \begin{pmatrix} Z^\top AZ & \cdots & Z^\top A^{2m}Z \\ \star & & \end{pmatrix} & \text{w.p. } 1 - O\left(\frac{m^{3/2}}{(n/p)^{1/2}}\right) \\ \star & \text{o.w.} \end{cases}$$

where $Z \in \mathbb{R}^{n \times m}$ has i.i.d. $\mathcal{N}(0,1)$ entries and \star is some unspecified random variable. The blocks $Z^\top A^i Z$ contain the inner products $\{z_j^\top A^i z_{j'}\}_{(i,j,j') \in H}$ where H is the index set in Lemma 3.4. By Lemma 3.3 and Lemma 3.4 respectively, there are g and h such that

$$g\left(h\left(\{z_j^\top A^i z_{j'}\}_{(i,j,j') \in H}\right)\right) \stackrel{d}{=} g\left(\{A^i z_j\}_{i+j \leq m}\right) \stackrel{d}{=} \{Av_j\}_{j < m}.$$

From there, we may simply apply f_{alg} . The result follows by (3.1). \square

4. CONCLUSION

In this work, we studied the problem of estimating the spectral density of a large normal matrix using only matrix-vector products with A and A^* . Our main conclusion is that the spectral density of a normal matrix can be estimated with roughly the same complexity as a Hermitian matrix. More specifically, we showed that when $n \gg 1/\varepsilon^2$, for any normal matrix A with $\|A\|_2 \leq 1$, it is sufficient to use $O(1/\varepsilon)$ matrix-vector products in order to return with high probability an $O(\varepsilon)$ -error spectral density estimation in earth mover's distance. Moreover, we proved that this rate is asymptotically optimal, even for Hermitian input matrices, in that any algorithm with such a guarantee must in general make $\Omega(1/\varepsilon)$ matrix-vector queries.

An important direction for future research is to determine whether our arguments extend beyond the normal matrix setting. Our algorithm and analysis relies crucially on the spectral theorem and on the commutativity of A and A^* or the Hermitian and skew-Hermitian parts of A . For non-normal matrices, the matrix will not be unitarily diagonalizable and these commutativity relations break down. Moreover, the eigenvalues of A may not be well-conditioned. It would therefore be interesting either to identify useful classes of nearly normal matrices for which analogues of our results still hold or to prove hardness results showing that the normal case is a genuine boundary of tractability. In our view, understanding which parts of non-Hermitian spectral computation are rendered intractable by non-normality remains one of the central problems of numerical linear algebra.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under grant nos. DMS-2513687, AF-2427362, AF-2427363, AF-2046235, and AF-2045590. Nicholas thanks MathWorks for its support through the MathWorks Fellowship. We would like to thank Rajarshi Bhattacharjee for helpful conversations.

REFERENCES

- [APJ⁺18] Ryan Adams, Jeffrey Pennington, Matthew Johnson, Jamie Smith, Yaniv Ovadia, Brian Patton, and James Saunderson. Estimating the spectral density of large implicit matrices. *arXiv:1802.03451*, 2018.
- [BCG23] Sergey Bobkov, Gennadiy Chistyakov, and Friedrich Götze. *Concentration and Gaussian Approximation for Randomized Sums*. Springer Nature Switzerland, 2023.
- [BDER16] Sébastien Bubeck, Jian Ding, Ronen Eldan, and Miklós Z Rácz. Testing for high-dimensional geometry in random graphs. *Random Structures & Algorithms*, 49(3), 2016.

- [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*, 23(6):1959–2047, 2022.
- [BHOT24] Nicolas Boullé, Diana Halikias, Samuel E Otto, and Alex Townsend. Operator learning without the adjoint. *Journal of Machine Learning Research*, 25(364), 2024.
- [BJM⁺25] Rajarshi Bhattacharjee, Rajesh Jayaram, Cameron Musco, Christopher Musco, and Archan Ray. Improved spectral density estimation via explicit and implicit deflation. In *Proceedings of the 36th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025.
- [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.
- [Cai16] Alberto Cairo. Download the datasaurus: Never trust summary statistics alone; always visualize your data. <https://web.archive.org/web/20160901105257/http://www.thefunctionalart.com/2016/08/download-datasaurus-never-trust-summary.html>, 2016.
- [CdDPL⁺24] Sinho Chewi, Jaume de Dios Pont, Jerry Li, Chen Lu, and Shyam Narayanan. Query lower bounds for log-concave sampling. *Journal of the ACM*, 71(4), 2024.
- [CTU21] Tyler Chen, Thomas Trogdon, and Shashanka Ubaru. Analysis of stochastic Lanczos quadrature for spectrum approximation. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
- [CU24] Cecilia Chen and John Urschel. Estimating the numerical range with a Krylov subspace. *arXiv:2411.19165*, 2024.
- [CY92] Ling-Lie Chau and Yue Yu. Unitary polynomials in normal matrix models and wave functions for the fractional quantum Hall effects. *Physics Letters A*, 167(5-6), 1992.
- [CZ98] Ling-Lie Chau and Oleg Zaboronsky. On the structure of correlation functions in the normal matrix model. *Communications in mathematical physics*, 196(1), 1998.
- [DEM26] Michal Dereziński, Ethan N Epperly, and Raphael A Meyer. The matrix-vector complexity of $Ax = b$. *arXiv:2602.04842*, 2026.
- [DMN25] Francesco Dell’Accio, Francisco Marcellán, and Federico Nudo. An interpolation–regression approach for function approximation on the disk and its application to cubature formulas. *Advances in Computational Mathematics*, 51(6), 2025.
- [FP57] Ky Fan and Gordon Pall. Imbedding conditions for Hermitian and normal matrices. *Canadian Journal of Mathematics*, 9, 1957.
- [GKX19] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via Hessian eigenvalue density. In *International Conference on Machine Learning*, pages 2232–2241. PMLR, 2019.
- [HMT11] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2), 2011.
- [Jac12] Dunham Jackson. On approximation by trigonometric sums and polynomials. *Transactions of the American Mathematical society*, 13(4), 1912.
- [JL13] Tiefeng Jiang and Danning Li. Approximation of rectangular beta-Laguerre ensembles and large deviations. *Journal of Theoretical Probability*, 28(3), 2013.
- [JP94] Janez Jaklic and Peter Prelovsek. Lanczos method for the calculation of finite-temperature quantities in correlated systems. *Physical Review B*, 49(7), 1994.
- [Lan50] Cornelius Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 45(4), 1950.
- [LSY16] Lin Lin, Yousef Saad, and Chao Yang. Approximating spectral densities of large matrices. *SIAM review*, 58(1), 2016.
- [LXES19] Ruipeng Li, Yuanzhe Xi, Lucas Erlandson, and Yousef Saad. The eigenvalues slicing library (EVSL): Algorithms, implementation, and software. *SIAM Journal on Scientific Computing*, 41(4), 2019.
- [MF17] Justin Matejka and George Fitzmaurice. Same stats, different graphs: Generating datasets with varied appearance and identical statistics through simulated annealing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 1290–1294, 2017.
- [MM15] Cameron Musco and Christopher Musco. Randomized block Krylov methods for stronger and faster approximate singular value decomposition. *Advances in Neural Information Processing Systems 28 (NeurIPS)*, 28, 2015.

- [MMM21] Raphael A. Meyer, Cameron Musco, Christopher Musco, and David P. Woodruff. Hutch++: Optimal stochastic trace estimation. In *Proceedings of the 4th Symposium on Simplicity in Algorithms (SOSA)*, 2021.
- [MMRS25] Cameron Musco, Christopher Musco, Lucas Rosenblatt, and Apoorv Vikram Singh. Sharper bounds for Chebyshev moment matching, with applications. *Proceedings of the 38th Annual Conference on Computational Learning Theory (COLT)*, 2025.
- [NS64] D. J. Newman and H.S. Shapiro. Jackson’s theorem in higher dimensions. In *On Approximation Theory / Über Approximationstheorie: Proceedings of the Conference held in the Mathematical Research Institute at Oberwolfach*, pages 208–219, 1964.
- [RB25] I.O. Raikov and Y.M. Beltukov. The kernel polynomial method based on Jacobi polynomials. *Applied Mathematics and Computation*, 490, 2025.
- [Riv81] Theodore J. Rivlin. *An introduction to the approximation of functions*. Courier Corporation, 1981.
- [RR18] Miklós Z. Rácz and Jacob Richey. A smooth transition from Wishart to GOE. *Journal of Theoretical Probability*, 32(2), 2018.
- [Rys63] Herbert Ryser. *Combinatorial Mathematics*. American Mathematical Society, 1963.
- [Saa11] Yousef Saad. *Numerical methods for large eigenvalue problems: revised edition*. SIAM, 2011.
- [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.
- [Tim63] Aleksandr F. Timan. *Theory of Approximation of Functions of a Real Variable*. International Series of Monographs in Pure and Applied Mathematics. Pergamon Press, 1963.
- [Woo14] David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1-2), 2014.
- [WWAF06] Alexander Weiß, Gerhard Wellein, Andreas Alvermann, and Holger Fehske. The kernel polynomial method. *Reviews of modern physics*, 78(1), 2006.
- [WZZ22] David P. Woodruff, Fred Zhang, and Richard Zhang. Optimal query complexities for dynamic trace estimation. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022.