
Stochastic Chebyshev Gradient Descent for Spectral Optimization

Insu Han¹, Haim Avron² and Jinwoo Shin^{1,3}

¹School of Electrical Engineering, Korea Advanced Institute of Science and Technology

²Department of Applied Mathematics, Tel Aviv University

³AITrics

{insu.han, jinwoos}@kaist.ac.kr haimav@post.tau.ac.il

Abstract

A large class of machine learning techniques requires the solution of optimization problems involving spectral functions of parametric matrices, e.g. log-determinant and nuclear norm. Unfortunately, computing the gradient of a spectral function is generally of cubic complexity, as such gradient descent methods are rather expensive for optimizing objectives involving the spectral function. Thus, one naturally turns to stochastic gradient methods in hope that they will provide a way to reduce or altogether avoid the computation of full gradients. However, here a new challenge appears: there is no straightforward way to compute unbiased stochastic gradients for spectral functions. In this paper, we develop unbiased stochastic gradients for spectral-sums, an important subclass of spectral functions. Our unbiased stochastic gradients are based on combining randomized trace estimators with stochastic truncation of the Chebyshev expansions. A careful design of the truncation distribution allows us to offer distributions that are variance-optimal, which is crucial for fast and stable convergence of stochastic gradient methods. We further leverage our proposed stochastic gradients to devise stochastic methods for objective functions involving spectral-sums, and rigorously analyze their convergence rate. The utility of our methods is demonstrated in numerical experiments.

1 Introduction

A large class of machine learning techniques involves *spectral optimization* problems of the form,

$$\min_{\theta \in \mathcal{C}} F(A(\theta)) + g(\theta), \quad (1)$$

where \mathcal{C} is some finite-dimensional parameter space, A is a function that maps a parameter vector θ to a symmetric matrix $A(\theta)$, F is a *spectral function* (i.e., a real-valued function on symmetric matrices that depends only on the eigenvalues of the input matrix), and $g : \mathcal{C} \rightarrow \mathbb{R}$. Examples include hyperparameter learning in Gaussian process regression with $F(X) = \log \det X$ [23], nuclear norm regularization with $F(X) = \text{tr}(X^{1/2})$ [21], phase retrieval with $F(X) = \text{tr}(X)$ [9], and quantum state tomography with $F(X) = \text{tr}(X \log X)$ [16]. In the aforementioned applications, the main difficulty in solving problems of the form (1) is in efficiently addressing the spectral component $F(A(\cdot))$. While explicit formulas for the gradients of spectral functions can be derived [18], it is typically computationally expensive. For example, for $F(X) = \log \det X$ and $A(\theta) \in \mathbb{R}^{d \times d}$, the exact computation of $\nabla_{\theta} F(A(\theta))$ can take as much as $O(d^3 k)$, where k is the number of parameters in θ . Therefore, it is desirable to avoid computing, or at the very least reduce the number of times we compute, the gradient of $F(A(\theta))$ exactly.

It is now well appreciated in the machine learning literature that the use of stochastic gradients is effective in alleviating costs associated with expensive exact gradient computations. Using cheap stochastic gradients, one can avoid computing full gradients altogether by using Stochastic Gradient Descent (SGD). The cost is, naturally, a reduced rate of convergence. Nevertheless, many machine learning applications require only mild suboptimality, in which case cheap iterations often outweigh the reduced convergence rate. When nearly optimal solutions are sought, more recent variance reduced methods (e.g. SVRG [15]) are effective in reducing the number of full gradient computations to $O(1)$. For non-convex objectives, the stochastic methods are even more attractive to use as they allow to avoid a bad local optimum. However, closed-form formulas for computing the full gradients of spectral functions do not lead to efficient stochastic gradients in a straightforward manner.

Contribution. In this paper, we propose stochastic methods for solving (1) when the spectral function F is a *spectral-sum*. Formally, spectral-sums are spectral functions that can be expressed as $F(X) = \text{tr}(f(X))$ where f is a real-valued function that is lifted to the symmetric matrix domain by applying it to the eigenvalues. They constitute an important subclass of spectral functions, e.g., in all of the aforementioned applications of spectral optimization, the spectral function F is a spectral-sum.

Our algorithms are based on recent *biased* estimators for spectral-sums that combine stochastic trace estimation with Chebyshev expansion [12]. The technique used to derive these estimators can also be used to derive stochastic estimators for the gradient of spectral-sums (e.g., see [8]), but the resulting estimator is biased. To address this issue, we propose an *unbiased* estimator for spectral-sums, and use it to derive unbiased stochastic gradients. Our unbiased estimator is based on randomly selecting the truncation degree in the Chebyshev expansion, i.e., the truncated polynomial degree is drawn under some distribution. We remark that similar ideas of sampling unbiased polynomials have been studied in the literature, but for different setups [5, 17, 31, 26], and none of which are suitable for use in our setup.

While deriving unbiased estimators is very useful for ensuring stable convergence of stochastic gradient methods, it is not sufficient: convergence rates of stochastic gradient descent methods depend on the variance of the stochastic gradients, and this can be rather large for naïve choices of degree distributions. Thus, our main contribution is in establishing a provably optimal degree distribution minimizing the estimators’ variances with respect to the Chebyshev series. The proposed distribution gives order-of-magnitude smaller variances compared to other popular ones (Figure 1), which leads to improved convergence of the downstream optimization (Figure 2(c)).

We leverage our proposed unbiased estimators to design two stochastic gradient descent methods, one using the SGD framework and the other using the SVRG one. We rigorously analyze their convergence rates, showing sublinear and linear rate for SGD and SVRG, respectively. It is important to stress that our fast convergence results crucially depend on the proposed optimal degree distributions. Finally, we apply our algorithms to two machine learning tasks that involve spectral optimization: matrix completion and learning Gaussian processes. Our experimental results confirm that the proposed algorithms are significantly faster than other competitors under large-scale real-world instances. In particular, for learning Gaussian process under Szeged humid dataset, our generic method runs up to six times faster than the state-of-art method [8] specialized for the purpose.

2 Preliminaries

We denote the family of real symmetric matrices of dimension d by $\mathcal{S}^{d \times d}$. For $A \in \mathcal{S}^{d \times d}$, we use $\|A\|_{\text{mv}}$ to denote the time-complexity of multiplying A with a vector, i.e., $\|A\|_{\text{mv}} = O(d^2)$. For some structured matrices, e.g. low-rank, sparse or Toeplitz matrices, it is possible to have $\|A\|_{\text{mv}} = o(d^2)$.

2.1 Chebyshev expansion

Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be an analytic function on $[a, b]$ for $a, b \in \mathbb{R}$. Then, the Chebyshev series of f is given by

$$f(x) = \sum_{j=0}^{\infty} b_j T_j \left(\frac{2}{b-a}x - \frac{b+a}{b-a} \right), \quad b_j = \frac{2 - \mathbb{1}_{j=0}}{\pi} \int_{-1}^1 \frac{f\left(\frac{b-a}{2}x + \frac{b+a}{2}\right) T_j(x)}{\sqrt{1-x^2}} dx.$$

In the above, $\mathbb{1}_{j=0} = 1$ if $j = 0$ and 0 otherwise and $T_j(x)$ is the Chebyshev polynomial (of the first kind) of degree j . An important property of the Chebyshev polynomials is the following recursive

formula: $T_{j+1}(x) = 2xT_j(x) - T_{j-1}(x)$, $T_1(x) = x$, $T_0(x) = 1$. The Chebyshev series can be used to approximate $f(x)$ via simply truncating the higher order terms, i.e., $f(x) \approx p_n(x) := \sum_{j=0}^n b_j T_j(\frac{2}{b-a}x - \frac{b+a}{b-a})$. We call $p_n(x)$ the *truncated Chebyshev series* of degree n . For analytic functions, the approximation error (in the uniform norm) is known to decay exponentially [29]. Specifically, if f is analytic with $|f(\frac{b-a}{2}z + \frac{b+a}{2})| \leq U$ for some $U > 0$ in the region bounded by the ellipse with foci $+1, -1$ and sum of major and minor semi-axis lengths equals to $\rho > 1$, then

$$|b_j| \leq \frac{2U}{\rho^j}, \quad \forall j \geq 0, \quad \sup_{x \in [a,b]} |f(x) - p_n(x)| \leq \frac{4U}{(\rho - 1)\rho^n}. \quad (2)$$

2.2 Spectral-sums and their Chebyshev approximations

Given a matrix $A \in \mathcal{S}^{d \times d}$ and a function $f : \mathbb{R} \rightarrow \mathbb{R}$, the *spectral-sum* of A with respect to f is

$$\Sigma_f(A) := \text{tr}(f(A)) = \sum_{i=1}^d f(\lambda_i),$$

where $\text{tr}(\cdot)$ is the matrix trace and $\lambda_1, \lambda_2, \dots, \lambda_d$ are the eigenvalues of A . Spectral-sums constitute an important subclass of spectral functions, and many applications of spectral optimization involve spectral-sums. This is fortunate since spectral-sums can be well approximated using Chebyshev approximations.

For a general f , one needs all eigenvalues to compute $\Sigma_f(A)$, while for some functions, simpler types of decomposition might suffice (e.g., $\log \det A = \Sigma_{\log}(A)$ can be computed using the Cholesky decomposition). Therefore, the general complexity of computing spectral-sums is $O(d^3)$, which is clearly not feasible when d is very large, as is common in many machine learning applications. Hence, it is not surprising that recent literature proposed methods to approximate the large-scale spectral-sums, e.g., [12] recently suggested a fast randomized algorithm for approximating spectral-sums based on Chebyshev series and Monte-Carlo trace estimators (i.e., Hutchinson's method [14]):

$$\Sigma_f(A) = \text{tr}(f(A)) \approx \text{tr}(p_n(A)) = \mathbf{E}_{\mathbf{v}} [\mathbf{v}^\top p_n(A) \mathbf{v}] \approx \frac{1}{M} \sum_{k=1}^M \mathbf{v}^{(k)\top} \left(\sum_{j=0}^n b_j \mathbf{w}_j^{(k)} \right) \quad (3)$$

where $\mathbf{w}_{j+1}^{(k)} = 2 \left(\frac{2}{b-a}A - \frac{b+a}{b-a}I \right) \mathbf{w}_j^{(k)} - \mathbf{w}_{j-1}^{(k)}$, $\mathbf{w}_1^{(k)} = \left(\frac{2}{b-a}A - \frac{b+a}{b-a}I \right) \mathbf{v}$, $\mathbf{w}_0^{(k)} = \mathbf{v}^{(k)}$, and $\{\mathbf{v}^{(k)}\}_{k=1}^M$ are Rademacher random vectors, i.e., each coordinate of $\mathbf{v}^{(k)}$ is an i.i.d. random variable in $\{-1, 1\}$ with equal probability $1/2$ [14, 2, 25]. The approximation (3) can be computed using only matrix-vector multiplications, vector-vector inner-products and vector-vector additions $O(Mn)$ times each. Thus, the time-complexity becomes $O(Mn\|A\|_{\text{mv}} + Mnd) = O(Mn\|A\|_{\text{mv}})$. In particular, when $Mn \ll d$ and $\|A\|_{\text{mv}} = o(d^2)$, the cost can be significantly cheaper than $O(d^3)$ of exact computation. We further note that to apply the approximation (3), a bound on the eigenvalues is necessary. For an upper bound, one can use fast power methods [7]; this does not hurt the total algorithm complexity (see [11]). A lower bound can be enforced by substituting A with $A + \varepsilon I$ for some small $\varepsilon > 0$. We use these techniques in our numerical experiments.

We remark that one may consider other polynomial approximation schemes, e.g. Taylor, but we focus on the Chebyshev approximations since they are nearly optimal in approximation among polynomial series [20]. Another recently suggested powerful technique is *stochastic Lanczos quadrature* [30], however it is not suitable for our needs (our bias removal technique is not applicable for it).

3 Stochastic Chebyshev gradients of spectral-sums

Our main goal is to develop scalable methods for solving the following optimization problem:

$$\min_{\theta \in \mathcal{C} \subseteq \mathbb{R}^{d'}} \Sigma_f(A(\theta)) + g(\theta), \quad (4)$$

where $\mathcal{C} \subseteq \mathbb{R}^{d'}$ is a non-empty, closed and convex domain, $A : \mathbb{R}^{d'} \rightarrow \mathcal{S}^{d \times d}$ is a function of parameter $\theta = [\theta_i] \in \mathbb{R}^{d'}$ and $g : \mathbb{R}^{d'} \rightarrow \mathbb{R}$ is some function whose derivative with respect to

any parameter θ is computationally easy to obtain. Gradient-descent type methods are natural candidates for tackling such problems. However, while it is usually possible to compute the gradient of $\Sigma_f(A(\theta))$, this is typically very expensive. Thus, we turn to stochastic methods, like (projected) SGD [4, 34] and SVRG [15, 33]. In order to apply stochastic methods, one needs unbiased estimators of the gradient. The goal of this section is to propose a computationally efficient method to generate unbiased stochastic gradients of small variance for $\Sigma_f(A(\theta))$.

3.1 Stochastic Chebyshev gradients

Biased stochastic gradients. We begin by observing that if f is a polynomial itself or the Chebyshev approximation is exact, i.e., $f(x) = p_n(x) = \sum_{j=0}^n b_j T_j(\frac{2}{b-a}x - \frac{b+a}{b-a})$, we have

$$\begin{aligned} \frac{\partial}{\partial \theta_i} \Sigma_{p_n}(A) &= \frac{\partial}{\partial \theta_i} \text{tr}(p_n(A)) = \frac{\partial}{\partial \theta_i} \mathbf{E}_{\mathbf{v}}[\mathbf{v}^\top p_n(A) \mathbf{v}] = \mathbf{E}_{\mathbf{v}} \left[\frac{\partial}{\partial \theta_i} \mathbf{v}^\top p_n(A) \mathbf{v} \right] \\ &\approx \frac{1}{M} \sum_{k=1}^M \frac{\partial}{\partial \theta_i} \mathbf{v}^{(k)\top} p_n(A) \mathbf{v}^{(k)} = \frac{1}{M} \sum_{k=1}^M \mathbf{v}^{(k)\top} \left(\sum_{j=0}^n b_j \frac{\partial \mathbf{w}_j^{(k)}}{\partial \theta_i} \right), \end{aligned} \quad (5)$$

where $\{\mathbf{v}^{(k)}\}_{k=1}^M$ are i.i.d. Rademacher random vectors and $\partial \mathbf{w}_j^{(k)} / \partial \theta_i$ are given by the following recursive formula:

$$\frac{\partial \mathbf{w}_{j+1}^{(k)}}{\partial \theta_i} = \frac{4}{b-a} \frac{\partial A}{\partial \theta_i} \mathbf{w}_j^{(k)} + 2\tilde{A} \frac{\partial \mathbf{w}_j^{(k)}}{\partial \theta_i} - \frac{\partial \mathbf{w}_{j-1}^{(k)}}{\partial \theta_i}, \quad \frac{\partial \mathbf{w}_1^{(k)}}{\partial \theta_i} = \frac{2}{b-a} \frac{\partial A}{\partial \theta_i} \mathbf{v}^{(k)}, \quad \frac{\partial \mathbf{w}_0^{(k)}}{\partial \theta_i} = \mathbf{0}, \quad (6)$$

and $\tilde{A} = \frac{2}{b-a}A - \frac{b+a}{b-a}I$. We note that in order to compute (6) only matrix-vector products with A and $\partial A / \partial \theta_i$ are needed. Thus, stochastic gradients of spectral-sums involving polynomials of degree n can be computed in $O(Mn(\|A\|_{\text{mv}} d' + \sum_{i=1}^{d'} \|\frac{\partial A}{\partial \theta_i}\|_{\text{mv}}))$. As we shall see in Section 5, the complexity can be further reduced in certain cases. The above estimator can be leveraged to approximate gradients for spectral-sums of analytic functions via the truncated Chebyshev series: $\nabla_{\theta} \Sigma_f(A(\theta)) \approx \nabla_{\theta} \Sigma_{p_n}(A(\theta))$. Indeed, [8] recently explored this in the context of Gaussian process kernel learning. However, if f is not a polynomial, the truncated Chebyshev series p_n is not equal to f , so the above estimator is biased, i.e. $\nabla_{\theta} \Sigma_f(A) \neq \mathbf{E}[\nabla_{\theta} \mathbf{v}^\top p_n(A) \mathbf{v}]$. The biased stochastic gradients might hurt iterative stochastic optimization as biased errors accumulate over iterations.

Unbiased stochastic gradients. The estimators (3) and (5) are biased since they approximate an analytic function f via a polynomial p_n of fixed degree. Unless f is a polynomial itself, there exists an x_0 (usually uncountably many) for which $f(x_0) \neq p_n(x_0)$, so if A has an eigenvalue at x_0 we have $\Sigma_f(A) \neq \Sigma_{p_n}(A)$. Thus, one cannot hope that the estimator (3), let alone the gradient estimator (5), to be unbiased for *all* matrices A . To avoid deterministic truncation errors, we simply randomize the degree, i.e., design some distribution \mathcal{D} on polynomials such that for every x we have $\mathbf{E}_{p \sim \mathcal{D}}[p(x)] = f(x)$. This guarantees $\mathbf{E}_{p \sim \mathcal{D}}[\text{tr}(p(A))] = \text{tr}(\Sigma_f(A))$ from the linearity of expectation.

We propose to build such a distribution on polynomials by using truncated Chebyshev expansions where the truncation degree is stochastic. Let $\{q_i\}_{i=0}^{\infty} \subseteq [0, 1]$ be a set of numbers such that $\sum_{i=0}^{\infty} q_i = 1$ and $\sum_{i=r}^{\infty} q_i > 0$ for all $r \geq 0$. We now define for $r = 0, 1, \dots$

$$\hat{p}_r(x) := \sum_{j=0}^r \frac{b_j}{1 - \sum_{i=0}^{j-1} q_i} T_j \left(\frac{2}{b-a}x - \frac{b+a}{b-a} \right). \quad (7)$$

Note that $\hat{p}_r(x)$ can be obtained from $p_r(x)$ by re-weighting each coefficient according to $\{q_i\}_{i=0}^{\infty}$. Next, let n be a random variable taking non-negative integer values, and defined according to $\Pr(n=r) = q_r$. Under certain conditions on $\{q_i\}$, $\hat{p}_n(\cdot)$ can be used to derive unbiased estimators of $\Sigma_f(A)$ and $\nabla_{\theta} \Sigma_f(A)$ as stated in the following lemma.

Lemma 1 *Suppose that f is an analytic function and \hat{p}_n is the randomized Chebyshev series of f in (7). Assume that the entries of A are differentiable for $\theta \in \mathcal{C}'$, where \mathcal{C}' is an open set containing \mathcal{C} , and that for $a, b \in \mathbb{R}$ all the eigenvalues of $A(\theta)$ for $\theta \in \mathcal{C}'$ are in $[a, b]$. For any degree*

¹We assume that all partial derivatives $\partial A_{j,k} / \partial \theta_i$ for $j, k = 1, \dots, d, i = 1, \dots, d'$ exist and are continuous.

distribution on non-negative integers $\{q_i \in (0, 1) : \sum_{i=0}^{\infty} q_i = 1, \sum_{r=i}^{\infty} q_r > 0, \forall i \geq 0\}$ satisfying $\lim_{n \rightarrow \infty} \sum_{i=n+1}^{\infty} q_i \hat{p}_n(x) = 0$ for all $x \in [a, b]$, it holds

$$\mathbf{E}_{\mathbf{v}, n} [\mathbf{v}^\top \hat{p}_n(A) \mathbf{v}] = \Sigma_f(A), \quad \mathbf{E}_{\mathbf{v}, n} [\nabla_{\theta} \mathbf{v}^\top \hat{p}_n(A) \mathbf{v}] = \nabla_{\theta} \Sigma_f(A). \quad (8)$$

where the expectations are taken over the joint distribution on random degree n and Rademacher random vector \mathbf{v} (other randomized probing vectors can be used as well).

The proof of Lemma 1 is given in the supplementary material. We emphasize that (8) holds for any distribution $\{q_i\}_{i=0}^{\infty}$ on non-negative integers for which the conditions stated in Lemma 1 hold, e.g., geometric, Poisson or negative binomial distribution.

3.2 Main result: optimal unbiased Chebyshev gradients

It is a well-known fact that stochastic gradient methods converge faster when the gradients have smaller variances. The variance of our proposed unbiased estimators crucially depends on the choice of the degree distribution, i.e., $\{q_i\}_{i=0}^{\infty}$. In this section, we design a degree distribution that is variance-optimal in some formal sense. The variance of our proposed degree distribution decays exponentially with the expected degree, and this is crucial for the convergence analysis (Section 4).

The degrees-of-freedom in choosing $\{q_i\}_{i=0}^{\infty}$ is infinite, which poses a challenge for devising low-variance distributions. Our approach is based on the following simplified analytic approach studying the scalar function f in such a way that one can naturally expect that the resulting distribution $\{q_i\}_{i=0}^{\infty}$ also provides low-variance for the matrix cases of (8). We begin by defining the variance of randomized Chebyshev expansion (7) via the Chebyshev weighted norm as

$$\text{Var}_C(\hat{p}_n) := \mathbf{E}_n [\|\hat{p}_n - f\|_C^2], \quad \text{where } \|g\|_C^2 := \int_{-1}^1 \frac{g(\frac{b-a}{2}x + \frac{b+a}{2})^2}{\sqrt{1-x^2}} dx. \quad (9)$$

The primary reason why we consider the above variance is because by utilizing the orthogonality of Chebyshev polynomials we can derive an analytic expression for it.

Lemma 2 Suppose $\{b_j\}_{j=0}^{\infty}$ are coefficients of the Chebyshev series for analytic function f and \hat{p}_n is its randomized Chebyshev expansion (7). Then, it holds that $\text{Var}_C(\hat{p}_n) = \frac{\pi}{2} \sum_{j=1}^{\infty} b_j^2 \left(\frac{\sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} \right)$.

The proof of Lemma 2 is given in the supplementary material. One can observe from this result that the variance reduces as we assign larger masses to high degrees (due to exponentially decaying property of b_j (2)). However, using large degrees increases the computational complexity of computing the estimators. Hence, we aim to design a good distribution given some target complexity, i.e., the expected polynomial degree N . Namely, the minimization of $\text{Var}_C(\hat{p}_n)$ should be constrained by $\sum_{i=1}^{\infty} i q_i = N$ for some parameter $N \geq 0$.

However, minimizing $\text{Var}_C(\hat{p}_n)$ subject to the aforementioned constraints might be generally intractable as the number of variables $\{q_i\}_{i=0}^{\infty}$ is infinite and the algebraic structure of $\{b_j\}_{i=0}^{\infty}$ is arbitrary. Hence, in order to derive an analytic or closed-form solution, we relax the optimization. In particular, we suggest the following optimization to minimize an upper bound of the variance by utilizing $|b_j| \leq 2U\rho^{-j}$ from (2) as follows:

$$\min_{\{q_i\}_{i=0}^{\infty}} \sum_{j=1}^{\infty} \rho^{-2j} \left(\frac{\sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} \right) \quad \text{subject to} \quad \sum_{i=1}^{\infty} i q_i = N, \sum_{i=0}^{\infty} q_i = 1 \quad \text{and} \quad q_i \geq 0. \quad (10)$$

Figure 1(d) empirically demonstrates that $b_j^2 \approx c\rho^{-2j}$ for constant $c > 0$ under $f(x) = \log x$, in which case the above relaxed optimization (10) is nearly tight. The next theorem establishes that (10) has a closed-form solution, despite having infinite degrees-of-freedom. The theorem is applicable when knowing a $\rho > 1$ and a bound U such that the function f is analytic with $|f(\frac{b-a}{2}z + \frac{b+a}{2})| \leq U$ in the complex region bounded by the ellipse with foci $+1, -1$ and sum of major and minor semi-axis lengths is equal to $\rho > 1$.

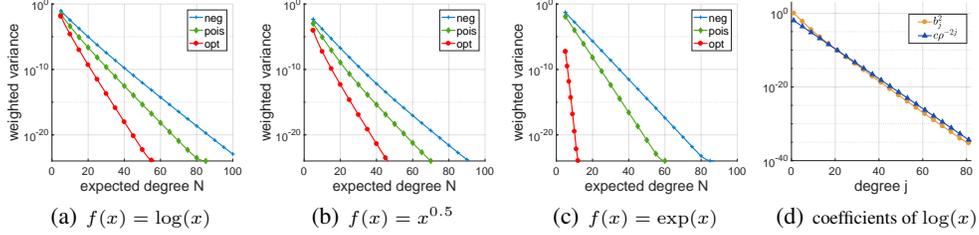


Figure 1: Chebyshev weighted variance for three distinct distributions: negative binomial (neg), Poisson (pois) and the optimal distribution (11) (opt) with the same mean N under (a) $\log x$, (b) \sqrt{x} on $[0.05, 0.95]$ and (c) $\exp(x)$ on $[-1, 1]$, respectively. Observe that “opt” has the smallest variance among all distributions. (d) Comparison between b_j^2 and $c\rho^{-2j}$ for some constant $c > 0$ and $\log x$.

Theorem 3 Let $K = \max\{0, N - \lfloor \frac{\rho}{\rho-1} \rfloor\}$. The optimal solution $\{q_i^*\}_{i=0}^\infty$ of (10) is

$$q_i^* = \begin{cases} 0 & \text{for } i < K \\ 1 - (N - K)(\rho - 1)\rho^{-1} & \text{for } i = K \\ (N - K)(\rho - 1)^2\rho^{-i-1+N-K} & \text{for } i > K, \end{cases} \quad (11)$$

and it satisfies the unbiasedness condition in Lemma 1, i.e., $\lim_{n \rightarrow \infty} \sum_{i=n+1}^\infty q_i^* \hat{p}_n(x) = 0$.

The proof of Theorem 3 is given in the supplementary material. Observe that a degree smaller than K is never sampled under $\{q_i^*\}$, which means that the corresponding unbiased estimator (7) combines deterministic series of degree K with randomized ones of higher degrees. Due to the geometric decay of $\{q_i^*\}$, large degrees will be sampled with exponentially small probability.

The optimality of the proposed distribution (11) (labeled opt) is illustrated by comparing it numerically to other distributions: negative binomial (labeled neg) and Poisson (labeled pois), on three analytic functions: $\log x$, \sqrt{x} and $\exp(x)$. Figures 1(a) to 1(c) show the weighted variance (9) of these distributions where their means are commonly set from $N = 5$ to 100. Observe that the proposed distribution has order-of-magnitude smaller variance compared to other tested distributions.

4 Stochastic Chebyshev gradient descent algorithms

In this section, we leverage unbiased gradient estimators based on (8) in conjunction with our optimal degree distribution (11) to design computationally efficient methods for solving (4). In particular, we propose to randomly sample a degree n from (11) and estimate the gradient via Monte-Carlo method:

$$\frac{\partial}{\partial \theta_i} \Sigma_f(A) = \mathbf{E} \left[\frac{\partial}{\partial \theta_i} \mathbf{v}^\top \hat{p}_n(A) \mathbf{v} \right] \approx \frac{1}{M} \sum_{k=1}^M \mathbf{v}^{(k)\top} \left(\sum_{j=0}^n \frac{b_j}{1 - \sum_{i=0}^{j-1} q_i^*} \frac{\partial \mathbf{w}_j^{(k)}}{\partial \theta_i} \right) \quad (12)$$

where $\partial \mathbf{w}_j^{(k)} / \partial \theta_i$ can be computed using a Rademacher vector $\mathbf{v}^{(k)}$ and the recursive relation (6).

4.1 Stochastic Gradient Descent (SGD)

In this section, we consider the use of projected SGD in conjunction with (12) to numerically solve the optimization (4). In the following, we provide a pseudo-code description of our proposed algorithm.

Algorithm 1 SGD for solving (4)

- 1: **Input:** number of iterations T , number of Rademacher vectors M , expected degree N and $\theta^{(0)}$
 - 2: **for** $t = 0$ to $T - 1$ **do**
 - 3: Draw M Rademacher random vectors $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and a random degree n from (11) given N
 - 4: Compute $\psi^{(t)}$ from (12) at $\theta^{(t)}$ using $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and n
 - 5: Obtain a proper step-size η_t
 - 6: $\theta^{(t+1)} \leftarrow \Pi_{\mathcal{C}} \left(\theta^{(t)} - \eta_t \left(\psi^{(t)} + \nabla g(\theta^{(t)}) \right) \right)$, where $\Pi_{\mathcal{C}}(\cdot)$ is the projection mapping into \mathcal{C}
 - 7: **end for**
-

In order to analyze the convergence rate, we assume that (A0) all eigenvalues of $A(\theta)$ for $\theta \in \mathcal{C}'$ are in the interval $[a, b]$ for some open $\mathcal{C}' \supseteq \mathcal{C}$, (A1) $\Sigma_f(A(\theta)) + g(\theta)$ is continuous and α -strongly convex with respect to θ and (A2) $A(\theta)$ is L_A -Lipschitz for $\|\cdot\|_F$, $g(\theta)$ is L_g -Lipschitz and β_g -smooth. The formal definitions of the assumptions are in the supplementary material. These assumptions hold for many target applications, including the ones explored in Section 5. In particular, we note that assumption (A0) can be often satisfied with a careful choice of \mathcal{C} . It has been studied that (projected) SGD has a sublinear convergence rate for a smooth strongly-convex objective if the variance of gradient estimates is uniformly bounded [24, 22]. Motivated by this, we first derive the following upper bound on the variance of gradient estimators under the optimal degree distribution (11).

Lemma 4 *Suppose that assumptions (A0)-(A2) hold and $A(\theta)$ is L_{nuc} -Lipschitz for $\|\cdot\|_{\text{nuc}}$. Let ψ be the gradient estimator (12) at $\theta \in \mathcal{C}$ using Rademacher vectors $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and degree n drawn from the optimal distribution (11). Then, $\mathbf{E}_{\mathbf{v}, n}[\|\psi\|_2^2] \leq (2L_A^2/M + d' L_{\text{nuc}}^2) (C_1 + C_2 N^4 \rho^{-2N})$ where $C_1, C_2 > 0$ are some constants independent of M, N .*

The above lemma allows us to provide a sublinear convergence rate for Algorithm 1.

Theorem 5 *Suppose that assumptions (A0)-(A2) hold and $A(\theta)$ is L_{nuc} -Lipschitz for $\|\cdot\|_{\text{nuc}}$. If one chooses the step-size $\eta_t = 1/\alpha t$, then it holds that*

$$\mathbf{E}[\|\theta^{(T)} - \theta^*\|_2^2] \leq \frac{4}{\alpha^2 T} \max \left(L_g^2, \left(\frac{2L_A^2}{M} + d' L_{\text{nuc}}^2 \right) \left(C_1 + \frac{C_2 N^4}{\rho^{2N}} \right) \right)$$

where $C_1, C_2 > 0$ are constants independent of M, N , and $\theta^* \in \mathcal{C}$ is the global optimum of (4).

The proofs of Lemma 4 and Theorem 5 are given in the supplementary material. Note that larger M, N provide better convergence but they increase the computational complexity. The convergence is also faster with smaller d' , which is also evident in our experiments (see Section 5).

4.2 Stochastic Variance Reduced Gradient (SVRG)

In this section, we introduce a more advanced stochastic method using a further variance reduction technique, inspired by the stochastic variance reduced gradient method (SVRG) [15]. The full description of the proposed SVRG scheme for solving the optimization (4) is given below.

Algorithm 2 SVRG for solving (4)

- 1: **Input:** number of inner/outer iterations T, S , number of Rademacher vectors M , expected degree N , step-size η and initial parameter $\theta^{(0)} \in \mathcal{C}$
 - 2: $\tilde{\theta}^{(1)} \leftarrow \theta^{(0)}$
 - 3: **for** $s = 1$ to S **do**
 - 4: $\tilde{\mu}^{(s)} \leftarrow \nabla \Sigma_f(A(\tilde{\theta}^{(s)}))$ and $\theta^{(0)} \leftarrow \tilde{\theta}^{(s)}$
 - 5: **for** $t = 0$ to $T - 1$ **do**
 - 6: Draw M Rademacher random vectors $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and a random degree n from (11)
 - 7: Compute $\psi^{(t)}, \tilde{\psi}^{(s)}$ from (12) at $\theta^{(t)}$ and $\tilde{\theta}^{(s)}$, respectively using $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and n
 - 8: $\theta^{(t+1)} \leftarrow \Pi_{\mathcal{C}} \left(\theta^{(t)} - \eta \left(\psi^{(t)} - \tilde{\psi}^{(s)} + \tilde{\mu}^{(s)} + \nabla g(\theta^{(t)}) \right) \right)$
 - 9: **end for**
 - 10: $\tilde{\theta}^{(s+1)} \leftarrow \frac{1}{T} \sum_{t=1}^T \theta^{(t)}$
 - 11: **end for**
-

The main idea of SVRG is to subtract a mean-zero random variable to the original stochastic gradient estimator, where the randomness between them is shared. The SVRG algorithm was originally designed for optimizing finite-sum objectives, i.e., $\sum_i f_i(x)$, whose randomness is from the index i . On the other hand, the randomness in our case is from polynomial degrees and trace probing vectors for optimizing objectives of spectral-sums. This leads us to use the same randomness in $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and n for estimating both $\psi^{(t)}$ and $\tilde{\psi}^{(s)}$ in line 7 of Algorithm 2. We remark that unlike SGD, Algorithm 2 requires the expensive computation of exact gradients every T iterations. The next theorem establishes that if one sets T correctly only $O(1)$ gradient computations are required (for a fixed suboptimality) since we have a linear convergence rate.

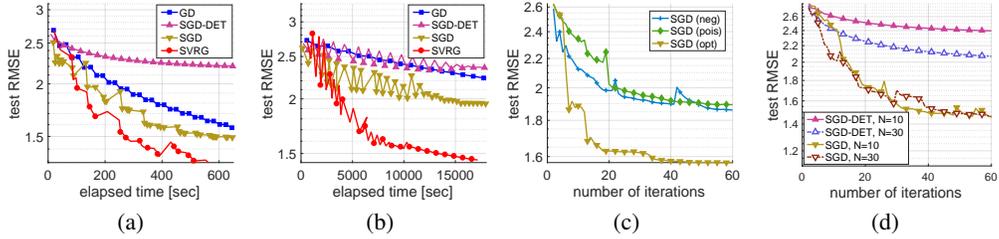


Figure 2: Matrix completion results under (a) MovieLens 1M and (b) MovieLens 10M. (c) Algorithm 1 (SGD) in MovieLens 1M under other distributions such as negative binomial (neg) and Poisson (pois). (d) SGD and SGD-DET under $N = 10, 30$.

Theorem 6 Suppose that assumptions (A0)-(A2) hold and $A(\theta)$ is β_A -smooth for $\|\cdot\|_F$. Let $\beta^2 = 2\beta_g^2 + \left(\frac{L_A^4 + \beta_A^2}{M} + L_A^4\right) \left(D_1 + \frac{D_2 N^8}{\rho^2 N}\right)$ for some constants $D_1, D_2 > 0$ independent of M, N . Choose $\eta = \frac{\alpha}{7\beta^2}$ and $T \geq 25\beta^2/\alpha^2$. Then, it holds that

$$\mathbf{E}[\|\tilde{\theta}^{(S)} - \theta^*\|_2^2] \leq r^S \mathbf{E}[\|\theta^{(0)} - \theta^*\|_2^2],$$

where $0 < r < 1$ is some constant and $\theta^* \in \mathcal{C}$ is the global optimum of (4).

The proof of the above theorem is given in the supplementary material, where we utilize the recent analysis of SVRG for the sum of smooth non-convex objectives [10, 1]. The key additional component in our analysis is to characterize $\beta > 0$ in terms of M, N so that the unbiased gradient estimator (12) is β -smooth in expectation under the optimal degree distribution (11).

5 Applications

In this section, we apply the proposed methods to two machine learning tasks: matrix completion and learning Gaussian processes. These correspond to minimizing spectral-sums Σ_f with $f(x) = x^{1/2}$ and $\log x$, respectively. We evaluate our methods under real-world datasets for both experiments.

5.1 Matrix completion

The goal is to recover a low-rank matrix $\theta \in [0, 5]^{d \times r}$ when a few of its entries are given. Since the rank function is neither differentiable nor convex, its relaxation such as Schatten- p norm has been used in respective optimization formulations. In particular, we consider the smoothed nuclear norm (i.e., Schatten-1 norm) minimization [19, 21] that corresponds to

$$\min_{\theta \in [0, 5]^{d \times r}} \text{tr}(A^{1/2}) + \lambda \sum_{(i,j) \in \Omega} (\theta_{i,j} - R_{i,j})^2$$

where $A = \theta\theta^\top + \varepsilon I$, $R \in [0, 5]^{d \times r}$ is a given matrix with missing entries, Ω indicates the positions of known entries and λ is a weight parameter and $\varepsilon > 0$ is a smoothing parameter. Observe that $\|A\|_{\text{mv}} = \|\theta\|_{\text{mv}} = O(dr)$, and the derivative estimation in this case can be amortized to compute using $O(dM(N^2 + Nr))$ operations. More details on this and our experimental settings are given in the supplementary material.

We use the MovieLens 1M and 10M datasets [13] (they correspond to $d = 3,706$ and $10,677$, respectively) and benchmark the gradient descent (GD), Algorithm 1 (SGD) and Algorithm 2 (SVRG). We also consider a variant of SGD using a deterministic polynomial degree, referred as SGD-DET, where it uses biased gradient estimators. We report the results for MovieLens 1M in Figure 2(a) and 10M in 2(b). For both datasets, SGD-DET performs badly due to its biased gradient estimators. On the other hand, SGD converges much faster and outperforms GD, where SGD for 10M converges much slower than that for 1M due to the larger dimension $d' = dr$ (see Theorem 5). Observe that SVRG is the fastest one, e.g., compared to GD, about 2 times faster to achieve RMSE 1.5 for MovieLens 1M and up to 6 times faster to achieve RMSE 1.8 for MovieLens 10M as shown in Figure 2(b). The gap between SVRG and GD is expected to increase for larger datasets. We also test SGD under other degree distributions: negative binomial (neg) and Poisson (pois) by choosing

parameters so that their means equal to $N = 15$. As reported in Figure 2(c), other distributions have relatively large variances so that they converge slower than the optimal distribution (opt). In Figure 2(d), we compare SGD-DET with SGD of the optimal distribution under the (mean) polynomial degrees $N = 10, 30$. Observe that a larger degree ($N = 30$) reduces the bias error in SGD-DET, while SGD achieves similar error regardless of the degree. The above results confirm that the unbiased gradient estimation and our degree distribution (11) are crucial for SGD.

5.2 Learning for Gaussian process regression

Next, we apply our method to hyperparameter learning for Gaussian process (GP) regression. Given training data $\{\mathbf{x}_i \in \mathbb{R}^\ell\}_{i=1}^d$ with corresponding outputs $\mathbf{y} \in \mathbb{R}^d$, the goal of GP regression is to learn a hyperparameter θ for predicting the output of a new/test input. The hyperparameter θ constructs the kernel matrix $A(\theta) \in \mathcal{S}^{d \times d}$ of the training data $\{\mathbf{x}_i\}_{i=1}^d$ (see [23]). One can find a good hyperparameter by minimizing the negative log-marginal likelihood with respect to θ :

$$\mathcal{L} := -\log p(\mathbf{y}|\{\mathbf{x}_i\}_{i=1}^d) = \frac{1}{2}\mathbf{y}^\top A(\theta)^{-1}\mathbf{y} + \frac{1}{2}\log \det A(\theta) + \frac{n}{2}\log 2\pi.$$

For handling large-scale datasets, [32] proposed the structured kernel interpolation framework assuming $\theta = [\theta_i] \in \mathbb{R}^3$ and

$$A(\theta) = WKW^\top + \theta_1^2 I, \quad K_{i,j} = \theta_2^2 \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / 2\theta_3^2),$$

where $W \in \mathbb{R}^{d \times r}$ is some sparse matrix and $K \in \mathbb{R}^{r \times r}$ is a dense kernel with $r \ll d$. Specifically, in [32], r ‘‘inducing’’ points are selected and entries of W are computed via interpolation with the inducing points. Under the framework, matrix-vector multiplications with A can be performed even faster, requiring $\|A\|_{\text{mv}} = \|W\|_{\text{mv}} + \|K\|_{\text{mv}} = O(d + r^2)$ operations. From $\|A\|_{\text{mv}} = \|\frac{\partial A}{\partial \theta_i}\|_{\text{mv}}$ and $d' = 3$, the complexity for computing gradient estimation (12) becomes $O(MN(d + r^2))$. If we choose $M, N, r = O(1)$, the complexity reduces to $O(d)$. The more detailed problem description and our experimental settings are given in the supplementary material.

We benchmark GP regression under natural sound dataset used in [32] and Szeged humid dataset [6] where they correspond to $d = 35,000$ and $16,930$, respectively. Recently, [8] utilized an approximation to derivatives of log-determinant based on stochastic Lanczos quadrature [30] (LANCZOS). We compare it with Algorithm 1 (SGD) which utilizes with unbiased gradient estimators while SVRG requires the exact gradient computation at least once which is intractable to run in these cases. As reported in Figure 3, SGD converges faster than LANCZOS for both datasets and it runs 2 times faster to achieve RMSE 0.0375 under sound dataset and under humid dataset LANCZOS can be often stuck at a local optimum, while SGD avoids it due to the use of unbiased gradient estimators.

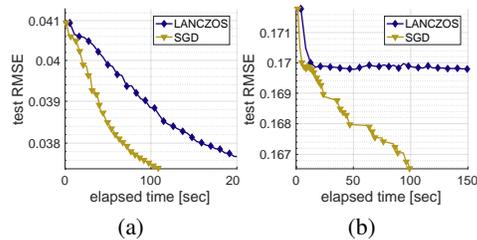


Figure 3: Hyperparameter learning for Gaussian process in modeling (a) sound dataset and (b) Szeged humid dataset comparing SGD to stochastic Lanczos quadrature (LANCZOS).

6 Conclusion

We proposed an optimal variance unbiased estimator for spectral-sums and their gradients. We applied our estimator in the SGD and SVRG frameworks, and analyzed convergence. The proposed optimal degree distribution is a crucial component of the analysis. We believe that the proposed stochastic methods are of broader interest in many machine learning tasks involving spectral-sums.

Acknowledgement

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (2018R1A5A1059921). Haim Avron acknowledges the support of the Israel Science Foundation (grant no. 1272/17).

References

- [1] Allen-Zhu, Zeyuan and Yuan, Yang. Improved SVRG for non-strongly-convex or sum-of-non-convex objectives. In *International Conference on Machine Learning (ICML)*, pp. 1080–1089, 2016.
- [2] Avron, H. and Toledo, S. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *Journal of the ACM*, 58(2):8, 2011.
- [3] Berge, Claude. *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity*. Courier Corporation, 1963.
- [4] Bottou, Léon. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pp. 177–186. Springer, 2010.
- [5] Broniatowski, Michel and Celant, Giorgio. Some overview on unbiased interpolation and extrapolation designs. *arXiv preprint arXiv:1403.5113*, 2014.
- [6] Budincsevity, Norbert. Weather in Szeged 2006-2016. <https://www.kaggle.com/budincsevity/szeged-weather/data>, 2016.
- [7] Davidson, Ernest R. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices. *Journal of Computational Physics*, 17(1): 87–94, 1975.
- [8] Dong, Kun, Eriksson, David, Nickisch, Hannes, Bindel, David, and Wilson, Andrew G. Scalable log determinants for Gaussian process kernel learning. In *Advances in Neural Information Processing Systems*, pp. 6330–6340, 2017.
- [9] Friedlander, Michael P. and Macêdo, Ives. Low-rank spectral optimization via gauge duality. *SIAM Journal on Scientific Computing*, 38(3):A1616–A1638, 2016. doi: 10.1137/15M1034283. URL <https://doi.org/10.1137/15M1034283>.
- [10] Garber, Dan and Hazan, Elad. Fast and simple PCA via convex optimization. *arXiv preprint arXiv:1509.05647*, 2015.
- [11] Han, Insu, Malioutov, Dmitry, and Shin, Jinwoo. Large-scale log-determinant computation through stochastic chebyshev expansions. In *International Conference on Machine Learning*, pp. 908–917, 2015.
- [12] Han, Insu, Malioutov, Dmitry, Avron, Haim, and Shin, Jinwoo. Approximating spectral sums of large-scale matrices using stochastic chebyshev approximations. *SIAM Journal on Scientific Computing*, 39(4):A1558–A1585, 2017.
- [13] Harper, F Maxwell and Konstan, Joseph A. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):19, 2016.
- [14] Hutchinson, M.F. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- [15] Johnson, Rie and Zhang, Tong. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pp. 315–323, 2013.
- [16] Koltchinskii, Vladimir and Xia, Dong. Optimal estimation of low rank density matrices. *J. Mach. Learn. Res.*, 16(1):1757–1792, January 2015. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=2789272.2886806>.
- [17] Lee, Yin Tat, Sidford, Aaron, and Wong, Sam Chiu-wai. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pp. 1049–1065. IEEE, 2015.
- [18] Lewis, A. S. Derivatives of spectral functions. *Mathematics of Operations Research*, 21(3):576–588, 1996. ISSN 0364765X, 15265471. URL <http://www.jstor.org/stable/3690298>.

- [19] Lu, Canyi, Lin, Zhouchen, and Yan, Shuicheng. Smoothed low rank and sparse matrix recovery by iteratively reweighted least squares minimization. *IEEE Transactions on Image Processing*, 24(2):646–654, 2015.
- [20] Mason, John C and Handscomb, David C. *Chebyshev polynomials*. CRC Press, 2002.
- [21] Mohan, Karthik and Fazel, Maryam. Iterative reweighted algorithms for matrix rank minimization. *Journal of Machine Learning Research*, 13(Nov):3441–3473, 2012.
- [22] Nemirovski, Arkadi, Juditsky, Anatoli, Lan, Guanghui, and Shapiro, Alexander. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [23] Rasmussen, Carl Edward. Gaussian processes in machine learning. In *Advanced Lectures on Machine Learning*, pp. 63–71. Springer, 2004.
- [24] Robbins, Herbert and Monro, Sutton. A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400–407, 1951.
- [25] Roosta-Khorasani, Farbod and Ascher, Uri. Improved bounds on sample size for implicit matrix trace estimators. *Foundations of Computational Mathematics*, 15(5):1187–1212, 2015.
- [26] Ryan P Adams, Jeffrey Pennington, Matthew J Johnson Jamie Smith Yaniv Ovadia Brian Patton James Saunderson. Estimating the spectral density of large implicit matrices. *arXiv preprint arXiv:1802.03451*, 2018.
- [27] Saad, Yousef. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [28] Schochetman, Irwin E and Smith, Robert L. Finite dimensional approximation in infinite dimensional mathematical programming. *Mathematical Programming*, 54(1-3):307–333, 1992.
- [29] Trefethen, Lloyd N. *Approximation theory and approximation practice*. SIAM, 2013.
- [30] Ubaru, Shashanka, Chen, Jie, and Saad, Yousef. Fast estimation of $\text{tr}(f(a))$ via stochastic Lanczos quadrature. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1075–1099, 2017.
- [31] Vinck, Martin, Battaglia, Francesco P, Balakirsky, Vladimir B, Vinck, AJ Han, and Pennartz, Cyriel MA. Estimation of the entropy based on its polynomial representation. *Physical Review E*, 85(5):051139, 2012.
- [32] Wilson, Andrew and Nickisch, Hannes. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *International Conference on Machine Learning*, pp. 1775–1784, 2015.
- [33] Xiao, Lin and Zhang, Tong. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- [34] Zinkevich, Martin. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 928–936, 2003.

Stochastic Chebyshev Gradient Descent for Spectral Optimization: Supplementary material

A Details of experiments

A.1 Matrix completion

For matrix completion, the problem can be expressed via the convex smoothed nuclear norm minimization as

$$\min_{\theta \in [0,5]^{d \times r}} \text{tr}(A^{1/2}) + \lambda \sum_{(i,j) \in \Omega} (\theta_{i,j} - R_{i,j})^2, \quad (13)$$

where $A = \theta\theta^\top + \varepsilon I$, $R \in [0, 5]^{d \times r}$ is a given matrix with missing entries, Ω indicates the positions of known entries and λ is a weight parameter and $\varepsilon > 0$ is a smoothing parameter. In this case, the gradient estimator (12) can be amortized as

$$\nabla_{\theta} \text{tr}(A^{p/2}) \approx \frac{2}{M} \sum_{k=1}^M \sum_{i=0}^{n-1} (2 - \mathbf{1}_{i=0}) \mathbf{w}_i^{(k)} \left(\sum_{j=i}^{n-1} \frac{b_{j+1}}{1 - \sum_{\ell=0}^j q_{\ell}^*} \mathbf{y}_{j-i}^{(k)} \right)^\top \theta \quad (14)$$

where

$$\begin{aligned} \mathbf{w}_{j+1}^{(k)} &= 2\mathbf{w}_j^{(k)} - \mathbf{w}_{j-1}^{(k)}, & \mathbf{w}_1^{(k)} &= \tilde{A}\mathbf{v}, & \mathbf{w}_0^{(k)} &= \mathbf{v}^{(k)}, \\ \mathbf{y}_{j+1}^{(k)} &= 2\mathbf{y}_j^{(k)} + \mathbf{y}_{j-1}^{(k)}, & \mathbf{y}_1^{(k)} &= 2\tilde{A}\mathbf{v}^{(k)}, & \mathbf{y}_0^{(k)} &= \mathbf{v}^{(k)} \end{aligned}$$

and $\tilde{A} = \left(\frac{2}{b-a}A - \frac{b+a}{b-a}I \right)$ for the lower/upper bound on A 's eigenvalues $a, b \in \mathbb{R}^+$. This comes from the following lemma, whose proof is in Section B.5.

Lemma 7 *Suppose f is an analytic function and $p_n(x) := \sum_{j=0}^n b_j T_j(x)$ is its truncated Chebyshev series of degree $n \geq 1$ for $x \in [-1, 1]$. Let $A = \theta\theta^\top + \varepsilon I$ for $\theta \in \mathbb{R}^{d \times r}, \varepsilon > 0$ such that all eigenvalues of A are in $[-1, 1]$. Then, for any $\mathbf{v} \in \mathbb{R}^d$, it holds that*

$$\nabla_{\theta} \mathbf{v}^\top p_n(A) \mathbf{v} = 2 \sum_{i=0}^{n-1} (2 - \mathbf{1}_{i=0}) \mathbf{w}_i \left(\sum_{j=i}^{n-1} b_{j+1} \mathbf{y}_{j-i} \right)^\top \theta,$$

where $\mathbf{w}_{j+1} = 2A\mathbf{w}_j - \mathbf{w}_{j-1}$, $\mathbf{w}_1 = A\mathbf{v}$, $\mathbf{w}_0 = \mathbf{v}$ and $\mathbf{y}_{j+1} = 2\mathbf{y}_j + \mathbf{y}_{j-1}$, $\mathbf{y}_1 = 2A\mathbf{v}$, $\mathbf{y}_0 = \mathbf{v}$.

Observe that $\|A\|_{\text{mv}} = \|\theta\|_{\text{mv}} = O(dr)$, and the computation for (14) can be amortized using $O(M(n^2d + ndr))$ operations. For $M, n, r = O(1)$, the complexity reduces to $O(d)$.

After update the parameter θ in a direction of gradient estimator, we project θ onto $[0, 5]^{d \times r}$, that is,

$$\Pi_{\mathcal{C}}(\theta_{i,j}) = \begin{cases} \theta_{i,j}, & \text{if } \theta_{i,j} \in [0, 5], \\ 0, & \text{if } \theta_{i,j} < 0, \\ 5, & \text{otherwise.} \end{cases}$$

In addition, after performing all gradient updates, we finally apply low-rank approximation using truncated SVD with rank 10 once and measure the test root mean square error (RMSE).

Setup. We use matrix R from MovieLens 1M (about 10^6 integer ratings from 1 to 5 from 6,040 users on 3,706 movies) and 10M (about 10^7 ratings from 0.5 to 5 with intervals 0.5 from 10,677

users on 71,567 movies) datasets [13]. We randomly select 90% of each dataset for training and use the rest for testing. We choose the (mean) polynomial degree $N = 15$ and the number of trace random vectors $M = 100$ for SVRG and $M = 200$ for SGD-DET, SGD, respectively, for comparable complexity at each gradient update. Especially, for SVRG, we choose $T = 100$. We decrease step-sizes exponentially with ratio 0.97 over the iterations.

A.2 Gaussian process (GP) regression

Given training data $\{\mathbf{x}_i \in \mathbb{R}^\ell\}_{i=1}^d$ with corresponding outputs $\mathbf{y} \in \mathbb{R}^d$, the goal of GP regression is to learn a hyperparameter θ for predicting the output of a new/test input. GP defines a distribution over functions, which follow multivariate Gaussian distribution with mean function $\mu_\theta : \mathbb{R}^\ell \rightarrow \mathbb{R}$ and covariance (i.e., kernel) function $a_\theta : \mathbb{R}^\ell \times \mathbb{R}^\ell \rightarrow \mathbb{R}$. To this end, we set the kernel matrix $A = A(\theta) \in \mathcal{S}^{d \times d}$ of $\{\mathbf{x}_i\}_{i=1}^d$ such that $A_{i,j} = a_\theta(\mathbf{x}_i, \mathbf{x}_j)$ and the mean function to be zero. One can find a good hyperparameter by minimizing the negative log-marginal likelihood with respect to θ :

$$\mathcal{L} := -\log p(\mathbf{y}|\{\mathbf{x}_i\}_{i=1}^d) = \frac{1}{2}\mathbf{y}^\top A^{-1}\mathbf{y} + \frac{1}{2}\log \det A + \frac{n}{2}\log 2\pi, \quad (15)$$

and predict $y = \mathbf{a}^\top A^{-1}\mathbf{y}$ where $\mathbf{a}_i = a_\theta(\mathbf{x}_i, \mathbf{x})$ (see [23]). Gradient-based methods can be used for optimizing (15) using its partial derivatives:

$$\frac{\partial \mathcal{L}}{\partial \theta_i} = -\frac{1}{2}\left(\mathbf{y}^\top \frac{\partial A}{\partial \theta_i}\right)A^{-1}\left(\frac{\partial A}{\partial \theta_i}\mathbf{y}\right) - \frac{1}{2}\frac{\partial \log \det A}{\partial \theta_i}.$$

Observe that the first term can be computed by an efficient linear solver, e.g., conjugate gradient descents [27], while the second term is computationally expensive for large d . Hence, one can use our proposed gradient estimator (12) for $\Sigma_f(A)$ with $f(x) = \log x$.

For handling large-scale datasets, [32] proposed the structured kernel interpolation framework assuming $\theta = [\theta_i] \in \mathbb{R}^3$ and

$$A(\theta) = WKW^\top + \theta_1^2 I, K_{i,j} = \theta_2^2 \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / 2\theta_3^2),$$

where $W \in \mathbb{R}^{d \times r}$ is some sparse matrix and $K \in \mathbb{R}^{r \times r}$ is a dense kernel with $r \ll d$. Specifically, the authors select r ‘‘inducing’’ points and compute entries of W via interpolation with the inducing points. Under the framework, matrix-vector multiplications with A can be performed even faster, requiring $\|A\|_{\text{mv}} = \|W\|_{\text{mv}} + \|K\|_{\text{mv}} = O(d + r^2)$ operations. From $\|A\|_{\text{mv}} = \|\frac{\partial A}{\partial \theta_i}\|_{\text{mv}}$ and $d' = 3$, the complexity for computing gradient estimation (12) becomes $O(MN(d + r^2))$. If we choose $M, N, r = O(1)$, the complexity reduces to $O(d)$.

Setup. We benchmark GP regression under natural sound dataset used in [32, 8] and Szeged humid data [6]. We randomly choose 35,000 points for training and 691 for testing in sound dataset and choose 16,930 points for training and 614 points for test in Szeged 2015-2016 humid dataset. We set the polynomial degree $N = 15$ and $M = 30$ trace vectors for all algorithms. We also select $r = 3000$ induced points for kernel interpolation. Since GP regression is non-convex problem, the gradient descent methods are sensitive to the initial point. We select a good initial point using random grid search. We observe that our algorithm (SGD) utilizing unbiased gradient estimator performs well for any initial point. On the other hand, since LANCZOS is type of biased gradient descent methods, it is often stuck on a bad local optimum.

B Proof of theorems

B.1 Smoothness and strong convexity of matrix functions

We first provide the formal definitions of the assumptions in Section 4. Let $\mathcal{C} \subseteq \mathbb{R}^{d'}$ be a non-empty, closed convex domain and $h : \mathbb{R}^{d'} \rightarrow \mathbb{R}$ be a continuously differentiable function.

Definition 1 A function h is L -Lipschitz continuous (or L -Lipschitz) on \mathcal{C} if for all $\theta, \theta' \in \mathcal{C}$, there exists a constant $L > 0$ such that

$$|h(\theta) - h(\theta')| \leq L \|\theta - \theta'\|_2.$$

Definition 2 A function h is β -smooth on \mathcal{C} if its gradient is β -Lipschitz such that

$$\|\nabla h(\theta) - \nabla h(\theta')\|_2 \leq \beta \|\theta - \theta'\|_2.$$

Definition 3 A function h is α -strongly convex on \mathcal{C} if for all $\theta, \theta' \in \mathcal{C}$, there exists a constant $\alpha > 0$ such that

$$\langle \nabla h(\theta) - \nabla h(\theta'), \theta - \theta' \rangle \geq \alpha \|\theta - \theta'\|_2^2.$$

The above definition can be extended to functions map into matrix space. For example, suppose $A : \mathbb{R}^{d'} \rightarrow \mathbb{R}^{d \times d}$ is a function of $\theta \in \mathcal{C}$ and assume that all $\partial A_{j,k} / \partial \theta_i$'s exist and are continuous.

Definition 4 A function $A(\theta)$ is L_A -Lipschitz with respect to $\|\cdot\|_F$ if for all $\theta, \theta' \in \mathcal{C}$, there exists a constant $L_A > 0$ such that

$$\|A(\theta) - A(\theta')\|_F \leq L_A \|\theta - \theta'\|_2.$$

Similarly, $A(\theta)$ is L_{nuc} -Lipschitz with respect to $\|\cdot\|_{\text{nuc}}$ (matrix nuclear norm) there exists a constant $L_{\text{nuc}} > 0$ such that

$$\|A(\theta) - A(\theta')\|_{\text{nuc}} \leq L_{\text{nuc}} \|\theta - \theta'\|_2.$$

Definition 5 Let $A : \mathbb{R}^{d'} \rightarrow \mathcal{S}^{d \times d}$ be a continuously differentiable function of $\theta \in \mathcal{C}$. If $A(\theta)$ is β_A -smooth if for all $\theta, \theta' \in \mathcal{C}$, there exists a constant $\beta_A > 0$ such that

$$\left\| \frac{\partial A(\theta)}{\partial \theta} - \frac{\partial A(\theta')}{\partial \theta} \right\|_F \leq \beta_A \|\theta - \theta'\|_2.$$

B.2 Proof of Theorem 3 : optimal degree distribution

By adding $\sum_{j=1}^{\infty} \rho^{-2j} = 1/(\rho^2 - 1)$ in both sides of (10), the optimization (10) is equivalent to

$$\min_{\{q_n\}_{n=0}^{\infty}} \sum_{j=1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} \quad \text{subject to} \quad \sum_{n=1}^{\infty} n q_n = N, \quad \sum_{n=0}^{\infty} q_n = 1 \quad \text{and} \quad q_n \geq 0. \quad (16)$$

Note that the equality conditions can be written as

$$N = \sum_{n=1}^{\infty} n q_n = \sum_{n=1}^{\infty} \sum_{j=1}^n q_n = \sum_{j=1}^{\infty} \sum_{n=j}^{\infty} q_n = \sum_{j=1}^{\infty} \left(1 - \sum_{n=0}^{j-1} q_n \right). \quad (17)$$

By Cauchy-Schwarz inequality for infinite series, we have

$$\begin{aligned} N \sum_{j=1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} &= \left(\sum_{j=1}^{\infty} \left(1 - \sum_{n=0}^{j-1} q_n \right) \right) \left(\sum_{j=1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} \right) \\ &\geq \left(\sum_{j=1}^{\infty} \rho^{-j} \right)^2 = \frac{1}{(\rho - 1)^2} \end{aligned}$$

and the equality holds when $q_0 = 1 - N(\rho - 1)\rho^{-1}$ and $q_n = N(\rho - 1)^2 \rho^{-(n+1)}$ for $n \geq 1$. However, this solution is not feasible when a given integer N is greater than $\frac{\rho}{\rho - 1}$ (due to $q_0 < 0$). The solution of (16) exists since the feasible region is closed and nonempty. For example,

$$q_n^* = \begin{cases} 0 & \text{for } 0 \leq n \leq k, \\ 1 - \frac{(N - k - 1)(\rho - 1)}{\rho} & \text{for } n = k + 1, \\ \frac{(N - k - 1)(\rho - 1)^2}{\rho^{n-k}} & \text{for } k + 2 \leq n \end{cases} \quad (18)$$

with $k := N - 1 - \lfloor \frac{\rho}{\rho-1} \rfloor$ is feasible and achieves the objective function of (16)

$$\frac{1 - \rho^{-2(k+1)}}{\rho^2 - 1} + \frac{1}{(N - k - 1)(\rho - 1)^2 \rho^{2(k+1)}}.^2$$

To figure out that q^* is the optimal solution, one can investigate KKT conditions of (16). However, in general, KKT theorem can not be applied to infinite dimensional problems. Instead, we consider the finite dimensional approximation of (16):

$$\min_{q_0, \dots, q_T} \sum_{j=1}^T \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} \quad \text{subject to} \quad \sum_{n=0}^T n q_n = N, \sum_{n=0}^T q_n = 1 \quad \text{and} \quad q_n \geq 0. \quad (19)$$

As we show in later, one can obtain the optimal solution of (19) for sufficiently large T using KKT conditions, which is

$$q_n = \begin{cases} 0 & \text{for } 0 \leq n \leq k, \\ 1 - \frac{(N - k - 1)(\rho - 1)}{1 - \rho^{-T+k+1}} \rho^{-1} & \text{for } n = k + 1, \\ \frac{(N - k - 1)(\rho - 1)^2}{1 - \rho^{-T+k+1}} \rho^{-n+k} & \text{for } k + 2 \leq n \leq T - 1 \\ \frac{(N - k - 1)(\rho - 1)}{\rho^{T-k-1} - 1} & \text{for } n = T \end{cases} \quad (20)$$

with $k := N - 1 - \lfloor \frac{\rho}{\rho-1} \rfloor$ and achieves the minimum

$$\frac{1 - \rho^{-2(k+1)}}{\rho^2 - 1} + \frac{(1 - \rho^{-T+k+1})^2}{(N - k - 1)(\rho - 1)^2 \rho^{2(k+1)}}. \quad (21)$$

We will show that the minimum of the infinite problem (16) is equivalent to the limit of (21) (a similar approach was introduced in [28]). We first extend q_n to the point with infinite dimension.

Let $q_n^{(T)} = q_n$ for $n \leq T$ and $q_n^{(T)} = 0$ for $n > T$, then $q^{(T)} = (q_0^{(T)}, q_1^{(T)}, \dots)$ is a feasible point of (16). Note that $\lim_{T \rightarrow \infty} q_n^{(T)} = q_n^*$ for all n . Define that

$$f(q, T) = \begin{cases} \sum_{j=1}^T \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} := C(q; T), & T = 1, 2, \dots, \\ \sum_{j=1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} := C(q), & T = \infty \end{cases}$$

for $q = (q_0, q_1, \dots)$. We claim that f is continuous. Suppose $T_i \in \mathbb{N}$ is a nondecreasing infinite sequence such that $T_i > k, T_i \rightarrow \infty$ and $q^{(T_i)} \rightarrow q^*$ as $i \rightarrow \infty$. Consider that

$$\begin{aligned} \left| f(q^*, \infty) - f(q^{(T_i)}, T_i) \right| &= \left| C(q^*) - C(q^{(T_i)}; T_i) \right| \\ &= \left| \sum_{j=1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n^*} - \sum_{j=1}^{T_i} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n^{(T_i)}} \right| \\ &\leq \left| \sum_{j=T_i+1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n^*} \right| + \left| \sum_{j=1}^{T_i} \rho^{-2j} \left(\frac{1}{1 - \sum_{n=0}^{j-1} q_n^*} - \frac{1}{1 - \sum_{n=0}^{j-1} q_n^{(T_i)}} \right) \right| \\ &\leq \frac{\rho^{-T_i-k}}{(N - k - 1)(\rho - 1)} + \frac{\rho^{-T_i-k-1}}{(N - k - 1)(\rho - 1)^2} \end{aligned} \quad (22)$$

²If $k = -1$, it is equivalent to the minimum from Cauchy-Schwarz inequality.

and (22) goes to zero as $i \rightarrow \infty$. In addition, the feasible set of (19) is nondecreasing, i.e., if we define the feasible regions as

$$X(T) := \left\{ q : \sum_{n=0}^T nq_n = N, \sum_{n=0}^T q_n = 1, q_n \geq 0, q_n = 0 \text{ for } n > T \right\},$$

$$X := \left\{ q : \sum_{n=0}^{\infty} nq_n = N, \sum_{n=0}^{\infty} q_n = 1, q_n \geq 0 \right\}$$

then $X(T) \subseteq X(T+1)$ for any T . This leads to $\lim_{T \rightarrow \infty} X(T) = \cup_{T \geq 1} X(T) = X$. Therefore, by the Berge's Maximum Theorem [3], the minimum of the finite dimensional problem (21) converges to that of infinite problem (16), i.e.,

$$\begin{aligned} \min \left\{ \sum_{j=1}^{\infty} \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} : q \in X \right\} &= \lim_{T \rightarrow \infty} \min \left\{ \sum_{j=1}^T \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} : q \in X(T) \right\} \\ &= \lim_{T \rightarrow \infty} \left(\frac{1 - \rho^{-2(k+1)}}{\rho^2 - 1} + \frac{(1 - \rho^{-T+k+1})^2}{(N - k - 1)(\rho - 1)^2 \rho^{2(k+1)}} \right) \\ &= \frac{1 - \rho^{-2(k+1)}}{\rho^2 - 1} + \frac{1}{(N - k - 1)(\rho - 1)^2 \rho^{2(k+1)}}. \end{aligned}$$

Since q^* in (18) achieves the above minimum, it follows that q^* in (18) is the minimizer of (16).

The remaining part is to obtain the solution of the finite dimensional approximation (19) using KKT conditions. Since the objective and all inequality conditions are *convex* functions, any feasible solution that satisfies KKT conditions are optimal. Define the Lagrangian as

$$\mathcal{L}(q, \lambda, \nu) = \sum_{j=1}^T \frac{\rho^{-2j}}{1 - \sum_{n=0}^{j-1} q_n} + \lambda_1 \left(\sum_{n=0}^T nq_n - N \right) + \lambda_2 \left(\sum_{n=0}^T q_n - 1 \right) - \sum_{n=0}^T \nu_n q_n$$

where λ_1, λ_2 and ν_0, \dots, ν_T are the Lagrangian multipliers of equality and inequality condition, respectively. The corresponding KKT conditions are following:

- **Stationary:** For $0 \leq n \leq T$,

$$\frac{\partial \mathcal{L}}{\partial q_n} = \sum_{j=n+1}^T \frac{\rho^{-2j}}{(1 - \sum_{n'=0}^{j-1} q_{n'})^2} + \lambda_1 n + \lambda_2 - \nu_n = 0, \quad (\text{C1})$$

- **Primal feasibility:**

$$\sum_{n=0}^T nq_n = N, \sum_{n=0}^T q_n = 1, q_n \geq 0, \quad (\text{C2})$$

- **Dual feasibility:** For $0 \leq n \leq T$,

$$\nu_n \geq 0, \quad (\text{C3})$$

- **Complementary slackness:** For $0 \leq n \leq T$,

$$\nu_n q_n = 0. \quad (\text{C4})$$

Consider (q, λ, ν) that satisfies the KKT conditions holds that $\nu_n = 0, k+1 \leq n \leq T$ and $\nu_n \neq 0, 0 \leq n \leq k$ for some $k \in [0, T]$. By the complementary slackness (C4), $q_0 = q_1 = \dots = q_k = 0$. Subtracting two consecutive stationary conditions (C1), we obtain for $0 \leq n \leq T-1$

$$\frac{\partial \mathcal{L}}{\partial q_n} - \frac{\partial \mathcal{L}}{\partial q_{n+1}} = \frac{\rho^{-2(n+1)}}{(1 - \sum_{n'=0}^n q_{n'})^2} - \lambda_1 - \nu_n + \nu_{n+1} = 0, \quad (\text{23})$$

which implies that

$$1 - \sum_{n'=0}^n q_{n'} = \frac{\rho^{-(n+1)}}{\sqrt{\lambda_1}} \quad \text{for } k+1 \leq n \leq T-1. \quad (24)$$

Putting them together into the equality condition (17) gives

$$N = \sum_{n=0}^{T-1} \left(1 - \sum_{n'=0}^n q_{n'} \right) = k+1 + \frac{1 - \rho^{-(T-k-1)}}{\sqrt{\lambda_1} \rho^{k+1} (\rho-1)},$$

equivalently, $\sqrt{\lambda_1} = \frac{1 - \rho^{-(T-k-1)}}{(N-k-1)\rho^{k+1}(\rho-1)}$. Therefore, we obtain the solution from (24):

$$q_n = \begin{cases} 0 & \text{for } 0 \leq n \leq k, \\ 1 - \frac{(N-k-1)(\rho-1)}{1 - \rho^{-T+k+1}} \rho^{-1} & \text{for } n = k+1, \\ \frac{(N-k-1)(\rho-1)^2}{1 - \rho^{-T+k+1}} \rho^{-n+k} & \text{for } k+2 \leq n \leq T-1 \\ \frac{(N-k-1)(\rho-1)}{\rho^{T-k-1} - 1} & \text{for } n = T \end{cases}$$

In order to satisfy the primal feasibility (C2), it should hold that

$$\frac{\rho(1 - \rho^{-T+k+1})}{\rho-1} \geq N-k-1 \quad \text{and} \quad k \leq N-1. \quad (25)$$

From (23), the dual variables ν can be written as for $n \leq k$

$$\nu_n - \nu_{n+1} = \rho^{-2(n+1)} - \lambda_1$$

and in order to satisfy the dual feasibility (C3), i.e., $\nu_n > 0$ for $n \leq k$, the sufficient condition is

$$N-k-1 > \frac{(1 - \rho^{-T+k+1})}{\rho-1}. \quad (26)$$

To satisfy both (25) and (26), there exists an integer in the interval $\left[\frac{(1 - \rho^{-T+k+1})}{\rho-1}, \frac{\rho(1 - \rho^{-T+k+1})}{\rho-1} \right]$. We now choose T large enough such that

$$\left\lfloor \frac{\rho}{\rho-1} \right\rfloor \leq \frac{\rho(1 - \rho^{-T+N})}{\rho-1},$$

and it holds that $\left\lfloor \frac{\rho}{\rho-1} \right\rfloor \in \left[\frac{(1 - \rho^{-T+k+1})}{\rho-1}, \frac{\rho(1 - \rho^{-T+k+1})}{\rho-1} \right]$ for some $0 \leq k \leq N-1$. By choosing $k := N-1 - \left\lfloor \frac{\rho}{\rho-1} \right\rfloor$, $\{q_n\}_{n=0}^T$ in (20) satisfies the KKT conditions and achieves the minimum

$$\frac{1 - \rho^{-2(k+1)}}{\rho^2 - 1} + \frac{(1 - \rho^{-T+k+1})^2}{(N-k-1)(\rho-1)^2 \rho^{2(k+1)}}.$$

B.3 Proof of Theorem 5 : convergence analysis of SGD

We recall that $\theta^{(t)} \in \mathcal{C} \subseteq \mathbb{R}^{d'}$ by the parameter in the t -th iteration and $\theta_i^{(t)}$ by its element i -th position for $i = 1, \dots, d'$. For simplicity, we denote that

$$h(\theta) := \Sigma_f(A(\theta)) + g(\theta)$$

and $\theta^* \in \mathcal{C}$ be the optimal of h . Let $\psi^{(t)}$ be our unbiased gradient estimator for $\Sigma_f(A(\theta))$ using $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and n , that is,

$$\mathbf{E}_{n, \mathbf{v}}[\psi^{(t)}] = \frac{\partial}{\partial \theta} \Sigma_f(A(\theta))$$

and $\nabla g^{(t)}$ be the derivative of $g(\theta)$ at $\theta^{(t)}$. Unless stated otherwise, we use $\|\cdot\|$ as the entry-wise L_2 -norm, i.e., L_2 -norm for vectors and Frobenius norm for matrices. Now we are ready to show the convergence guarantee for SGD. The iteration of SGD can be written as

$$\theta^{(t+1)} = \Pi_{\mathcal{C}} \left(\theta^{(t)} - \eta(\psi^{(t)} + \nabla g^{(t)}) \right)$$

where $\Pi_{\mathcal{C}}(\cdot)$ is the projection mapping in \mathcal{C} . The remaining part is similar with standard proof of the projected stochastic gradient descent. First, we write the error between $\theta^{(t)}$ and θ^* as

$$\begin{aligned} \|\theta^{(t+1)} - \theta^*\|^2 &= \|\Pi_{\mathcal{C}}(\theta^{(t)} - \eta(\psi^{(t)} + \nabla g^{(t)})) - \theta^*\|^2 \\ &\leq \|\theta^{(t)} - \eta(\psi^{(t)} + \nabla g^{(t)}) - \theta^*\|^2 \\ &= \|\theta^{(t)} - \theta^*\|^2 - 2\eta \left\langle \psi^{(t)} + \nabla g^{(t)}, \theta^{(t)} - \theta^* \right\rangle + \eta^2 \|\psi^{(t)} + \nabla g^{(t)}\|^2 \\ &\leq \|\theta^{(t)} - \theta^*\|^2 - 2\eta \left\langle \psi^{(t)} + \nabla g^{(t)}, \theta^{(t)} - \theta^* \right\rangle + 2\eta^2 \|\psi^{(t)}\|^2 + 2\eta^2 \|\nabla g^{(t)}\|^2 \\ &\leq \|\theta^{(t)} - \theta^*\|^2 - 2\eta \left\langle \psi^{(t)} + \nabla g^{(t)}, \theta^{(t)} - \theta^* \right\rangle + 2\eta^2 \|\psi^{(t)}\|^2 + 2\eta^2 L_g^2 \end{aligned}$$

where the inequality in the second line holds from the convexity of \mathcal{C} , the inequality in the fourth line follows from that $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$ and the last inequality follows from Lipschitz continuity of g . Taking the expectation with respect to random samples (i.e., random degree and vectors) in t -th iteration, which denoted as $\mathbf{E}_t[\cdot]$, we have

$$\mathbf{E}_t[\|\theta^{(t+1)} - \theta^*\|^2] \leq \|\theta^{(t)} - \theta^*\|^2 - 2\eta \left\langle \nabla h(\theta^{(t)}), \theta^{(t)} - \theta^* \right\rangle + 4\eta^2 B^2 \quad (27)$$

where $B^2 := \max(\mathbf{E}_t[\|\psi^{(t)}\|^2], L_g^2)$. In addition, by α -strong convexity of h , it holds that

$$\alpha \|\theta^{(t)} - \theta^*\|^2 \leq \left\langle \nabla h(\theta^{(t)}), \theta^{(t)} - \theta^* \right\rangle. \quad (28)$$

Combining (27) with (28) and taking the expectation on both sides with respect to all random samples from 1, ..., t iteration, we obtain that

$$\mathbf{E}[\|\theta^{(t+1)} - \theta^*\|^2] \leq (1 - 2\eta\alpha) \mathbf{E}[\|\theta^{(t)} - \theta^*\|^2] + 4\eta^2 B^2$$

Applying $\eta = \frac{1}{\alpha t}$, we have

$$\mathbf{E}[\|\theta^{(t+1)} - \theta^*\|^2] \leq \left(1 - \frac{2}{t}\right) \mathbf{E}[\|\theta^{(t)} - \theta^*\|^2] + \frac{4B^2}{\alpha^2 t^2}.$$

Therefore, if $\mathbf{E}[\|\theta^{(1)} - \theta^*\|^2] \leq 4B^2/\alpha^2$ holds, then the result follows by induction on $t \geq 1$. Under assumption that $\mathbf{E}[\|\theta^{(t)} - \theta^*\|^2] \leq 4B^2/(\alpha^2 t)$, it is straightforward that

$$\mathbf{E}[\|\theta^{(t+1)} - \theta^*\|^2] \leq \left(1 - \frac{2}{t}\right) \frac{4B^2}{\alpha^2 t} + \frac{4B^2}{\alpha^2 t^2} \leq \frac{4B^2}{\alpha^2} \left(\frac{1}{t+1}\right).$$

To show the case of $t = 1$, we recall the strong convexity of h and use Cauchy-Schwartz inequality:

$$\alpha \|\theta^{(1)} - \theta^*\|^2 \leq \left\langle \psi^{(1)} + \nabla g^{(1)}, \theta^{(1)} - \theta^* \right\rangle \leq \|\psi^{(1)} + \nabla g^{(1)}\| \|\theta^{(1)} - \theta^*\|,$$

which leads to that

$$\alpha^2 \mathbf{E}[\|\theta^{(1)} - \theta^*\|^2] \leq \mathbf{E}[\|\psi^{(1)} + \nabla g^{(1)}\|^2] \leq 4B^2.$$

Recall that Lemma 4 implies that for all t

$$\mathbf{E}_t[\|\psi^{(t)}\|^2] \leq (2L_A^2/M + d' L_{\text{nuc}}^2) (C_1 + C_2 N^4 \rho^{-2N}).$$

for some constants $C_1, C_2 > 0$. This completes the proof of Theorem 5.

B.4 Proof of Theorem 6 : convergence analysis of SVRG

Denote the objective as $h(\theta) := \Sigma_f(A(\theta)) + g(\theta)$. Let $\psi^{(t)}, \tilde{\psi}$ be our unbiased gradient estimator for $\Sigma_f(A(\theta))$ at $\theta^{(t)}$ and $\tilde{\theta}^{(s)}$, respectively, and $\tilde{\mu} = \nabla \Sigma_f(A(\tilde{\theta}^{(s)}))$. We use $\nabla g^{(t)}$ by the exact gradient of $g(\theta)$ at $\theta^{(t)}$, which is easy to compute. The iteration of SVRG can be written as

$$\theta^{(t+1)} = \Pi_{\mathcal{C}}(\theta^{(t)} - \eta \xi^{(t)}), \quad \text{where} \quad \xi^{(t)} := \psi^{(t)} - \tilde{\psi} + \tilde{\mu} + \nabla g^{(t)}$$

where $\Pi_{\mathcal{C}}(\cdot)$ is the projection mapping in \mathcal{C} . We first introduce the lemma that implies our unbiased estimator is β -smooth for some $\beta > 0$.

Lemma 8 *Suppose that assumptions (A0)-(A2) hold and assume that $A : \mathcal{C} \rightarrow \mathcal{S}^{d \times d}$ is β_A -smooth function with respect to $\|\cdot\|_F$. Let ψ, ψ' be our unbiased gradient estimator (12) at $\theta, \theta' \in \mathcal{C} \subseteq \mathbb{R}$ using the same $\{\mathbf{v}^{(k)}\}_{k=1}^M$ and n (drawn from (11) with mean N). Then, it holds that*

$$\mathbf{E}_{n, \mathbf{v}} [\|\psi + \nabla g(\theta) - \psi' - \nabla g(\theta')\|_2^2] \leq \left(2\beta_g^2 + \left(\frac{L_A^4 + \beta_A^2}{M} + L_A^4 \right) \left(D_1 + \frac{D_2 N^8}{\rho^{2N}} \right) \right) \|\theta - \theta'\|_2^2,$$

where $D_1, D_2 > 0$ are some constants independent of M, N .

The proof of the above lemma is given in Section B.5. For notational simplicity, we denote

$$\beta^2 := 2\beta_g^2 + \left(\frac{L_A^4 + \beta_A^2}{M} + L_A^4 \right) \left(D_1 + \frac{D_2 N^8}{\rho^{2N}} \right).$$

The remaining part mimics the analysis of [10]. Using the above lemma, the moment of the gradient estimator is bounded as

$$\begin{aligned} \mathbf{E}_t [\|\psi^{(t)} - \tilde{\psi} + \tilde{\mu} + \nabla g^{(t)}\|^2] &\leq 2\mathbf{E}_t [\|\psi^{(t)} + \nabla g^{(t)} - \psi^* - \nabla g^*\|^2] + 2\mathbf{E}_t [\|\tilde{\psi} - \psi^* - \nabla g^* - \tilde{\mu}\|^2] \\ &\leq 2\mathbf{E}_t [\|\psi^{(t)} + \nabla g^{(t)} - \psi^* - \nabla g^*\|^2] + 2\mathbf{E}_t [\|\tilde{\psi} + \nabla \tilde{g} - \psi^* - \nabla g^*\|^2] \\ &\leq 2\beta^2 \left(\|\theta^{(t)} - \theta^*\|^2 + \|\tilde{\theta} - \theta^*\|^2 \right) \end{aligned} \quad (29)$$

where the inequality in the first line holds from $\|a + b\|^2 \leq 2(\|a\|^2 + \|b\|^2)$, the inequality in the second line holds that $\mathbf{E}[\|X - \mathbf{E}[X]\|^2] \leq \mathbf{E}[\|X\|^2]$ for any random variable X and the last inequality holds from Lemma 8.

Now, we use similar procedures of Theorem 5 to obtain

$$\begin{aligned} \|\theta^{(t+1)} - \theta^*\|^2 &= \|\Pi_{\mathcal{C}}(\theta^{(t)} - \eta \xi^{(t)}) - \theta^*\|^2 \\ &\leq \|\theta^{(t)} - \eta \xi^{(t)} - \theta^*\|^2 \\ &= \|\theta^{(t)} - \theta^*\|^2 - 2\eta \langle \theta^{(t)} - \theta^*, \xi_t \rangle + \|\xi_t\|^2. \end{aligned}$$

where the inequality holds from the convexity of \mathcal{C} . Taking the expectation with respect to random samples of t -th iteration, which denoted as $\mathbf{E}_t[\cdot]$, we obtain that

$$\begin{aligned} \mathbf{E}_t [\|\theta^{(t+1)} - \theta^*\|^2] &= \|\theta^{(t)} - \theta^*\|^2 - 2\eta \langle \theta^{(t)} - \theta^*, \nabla h(\theta^{(t)}) \rangle + \eta^2 \mathbf{E}_t [\|\xi_t\|^2] \\ &\leq \|\theta^{(t)} - \theta^*\|^2 - 2\eta \alpha \|\theta^{(t)} - \theta^*\|^2 + \eta^2 \mathbf{E}_t [\|\xi_t\|^2] \\ &\leq \|\theta^{(t)} - \theta^*\|^2 - 2\eta \alpha \|\theta^{(t)} - \theta^*\|^2 + 2\eta^2 \beta^2 \left(\|\theta^{(t)} - \theta^*\|^2 + \|\tilde{\theta} - \theta^*\|^2 \right) \end{aligned}$$

where the inequality in the second line holds from the α -strong convexity of the objective and the last inequality holds from (29). Taking the expectation over the randomness of all iterations, we have

$$\mathbf{E} [\|\theta^{(t+1)} - \theta^*\|^2] - \mathbf{E} [\|\theta^{(t)} - \theta^*\|^2] \leq 2\eta (\eta \beta^2 - \alpha) \mathbf{E} [\|\theta^{(t)} - \theta^*\|^2] + 2\eta^2 \beta^2 \mathbf{E} [\|\tilde{\theta} - \theta^*\|^2]$$

Summing both sides over $t = 1, 2, \dots, T$, it yields that

$$\mathbf{E} [\|\theta^{(T)} - \theta^*\|^2] - \mathbf{E} [\|\theta^{(0)} - \theta^*\|^2] \leq 2\eta (\eta \beta^2 - \alpha) \sum_{t=0}^{T-1} \mathbf{E} [\|\theta^{(t)} - \theta^*\|^2] + 2T\eta^2 \beta^2 \mathbf{E} [\|\tilde{\theta} - \theta^*\|^2]$$

Rearranging and using the facts that $\mathbf{E}[\|\theta^{(T)} - \theta^*\|^2] \geq 0$ and $\tilde{\theta} = \tilde{\theta}^{(s)}$, we get

$$2\eta(\alpha - \eta\beta^2) \sum_{t=0}^{T-1} \mathbf{E}[\|\theta^{(t)} - \theta^*\|^2] \leq (1 + 2T\eta^2\beta^2) \mathbf{E}[\|\theta^{(0)} - \theta^*\|^2].$$

From $\tilde{\theta}^{(s+1)} = \frac{1}{T} \sum_{t=1}^T \theta^{(t)}$ and Jensen's inequality, we have

$$\mathbf{E}[\|\tilde{\theta}^{(s+1)} - \theta^*\|^2] \leq \frac{1}{T} \sum_{t=1}^T \mathbf{E}[\|\theta^{(t)} - \theta^*\|^2] \leq \frac{1 + 2T\eta^2\beta^2}{2\eta T(\alpha - \eta\beta^2)} \mathbf{E}[\|\tilde{\theta}^{(s)} - \theta^*\|^2]$$

Substituting $\eta = \frac{\alpha}{7\beta^2}$ and $T \geq \frac{49\beta^2}{2\alpha^2}$, we have that

$$\mathbf{E}[\|\tilde{\theta}^{(s)} - \theta^*\|^2] \leq r^s \mathbf{E}[\|\tilde{\theta}^{(0)} - \theta^*\|^2]$$

for some $0 < r < 1$.

B.5 Proof of lemmas

B.5.1 Proof of Lemma 1

Without loss of generality, we choose $a = -1, b = 1$. An analytic function f has an (unique) infinite Chebyshev series expansion: $f(x) = \sum_{j=0}^{\infty} b_j T_j(x)$. and recall that our proposed estimator as

$$\hat{p}_n(x) = \sum_{j=0}^n \frac{b_j}{1 - \sum_{i=0}^{j-1} q_i} T_j(x).$$

To prove that $\mathbf{E}_n[\hat{p}_n(x)] = f(x)$, we define two sequences:

$$A_M := \sum_{j=0}^M \sum_{n=j}^M q_n \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i}, \quad B_{M,K} := \sum_{j=0}^M \sum_{n=j}^K q_n \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i}.$$

Then, it is easy to show that

$$\lim_{M \rightarrow \infty} A_M = \sum_{j=0}^{\infty} \sum_{n=j}^{\infty} q_n \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i} = \sum_{n=0}^{\infty} q_n \left(\sum_{j=0}^n \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i} \right) = \sum_{n=0}^{\infty} q_n \hat{p}_n(x) = \mathbf{E}_n[\hat{p}_n(x)],$$

and

$$\lim_{M \rightarrow \infty} \lim_{K \rightarrow \infty} B_{M,K} = \lim_{M \rightarrow \infty} \sum_{j=0}^M \left(\sum_{n=j}^{\infty} q_n \right) \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i} = \lim_{M \rightarrow \infty} \sum_{j=0}^M b_j T_j(x) = f(x).$$

In general, A_M and $B_{M,K}$ might not converge to the same values. Now, consider sufficiently large $K \geq M$. From the condition that $\lim_{n \rightarrow \infty} \sum_{i=n+1}^{\infty} q_i \hat{p}_n(x)$, we have

$$\begin{aligned} \mathbf{E}_n[\hat{p}_n(x)] - f(x) &= \lim_{M \rightarrow \infty} \lim_{K \rightarrow \infty} (A_M - B_{M,K}) = \lim_{M \rightarrow \infty} \lim_{K \rightarrow \infty} \left(\sum_{j=0}^M \sum_{n=M+1}^K q_n \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i} \right) \\ &= \lim_{M \rightarrow \infty} \lim_{K \rightarrow \infty} \left(\sum_{n=M+1}^K q_n \right) \left(\sum_{j=0}^M \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i} \right) \\ &= \lim_{M \rightarrow \infty} \left(\sum_{n=M+1}^{\infty} q_n \right) \left(\sum_{j=0}^M \frac{b_j T_j(x)}{1 - \sum_{i=0}^{j-1} q_i} \right) \\ &= \lim_{M \rightarrow \infty} \left(\sum_{n=M+1}^{\infty} q_n \right) \hat{p}_M(x) = 0. \end{aligned}$$

Therefore, we can conclude that $\widehat{p}_n(x)$ is an unbiased estimator of $f(x)$. In addition, this also holds for the trace of matrices due to its linearity: $\mathbf{E}_n[\text{tr}(\widehat{p}_n(A))] = \text{tr}(f(A))$. By taking expectation over Rademacher random vectors \mathbf{v} and degree n , we establish the unbiased estimator of spectral-sums:

$$\mathbf{E}_{n,\mathbf{v}}[\mathbf{v}^\top \widehat{p}_n(A) \mathbf{v}] = \mathbf{E}_n[\mathbf{E}_{\mathbf{v}}[\mathbf{v}^\top \widehat{p}_n(A) \mathbf{v} | n]] = \mathbf{E}_n[\text{tr}(\widehat{p}_n(A))] = \text{tr}(f(A)),$$

For fixed \mathbf{v} and n , the function $h(\theta) := \mathbf{v}^\top \widehat{p}_n(A(\theta)) \mathbf{v}$ is a linear combination of all entries of A , so the fact that all partial derivatives $\partial A_{j,k} / \partial \theta_i$ exist and are continuous implies that the partial derivatives of h with respect to $\theta_1, \dots, \theta_{d'}$ exist and are continuous. In particular, since expectation over $\mathbf{v} \in [-1, +1]^d$ is a finite sum, it is straightforward that the gradient operator and expectation operator can be interchanged:

$$\nabla_{\theta} \text{tr}(f(A)) = \nabla_{\theta} \mathbf{E}[\mathbf{v}^\top \widehat{p}_n(A) \mathbf{v}] = \mathbf{E}[\nabla_{\theta} \mathbf{v}^\top \widehat{p}_n(A) \mathbf{v}].$$

In the case of trace probing vector \mathbf{v} is a continuous random vector, i.e., Gaussian, we turn to use the Leibniz rule which allows to interchange the gradient operator and expectation operator. Hence, we conclude the same result. This completes the proof of Lemma 1.

B.5.2 Proof of Lemma 2

Without loss of generality, we choose $a = -1, b = 1$. We first introduce the orthogonality of Chebyshev polynomials of the first kind, that is,

$$\int_{-1}^1 \frac{T_i(x)T_j(x)}{\sqrt{1-x^2}} dx = \begin{cases} 0 & i \neq j, \\ \pi & i = j = 0, \\ \frac{\pi}{2} & i = j \neq 0. \end{cases}$$

Given functions f, g defined on $[-1, 1]$, Chebyshev induced inner-product and weighted norm are defined as

$$\langle f, g \rangle_C = \int_{-1}^1 \frac{f(x)g(x)}{\sqrt{1-x^2}} dx, \quad \|f\|_C^2 = \langle f, f \rangle_C.$$

For a fixed n , the square of Chebyshev weighted error can be written as

$$\begin{aligned} \|\widehat{p}_n - f\|_C^2 &= \|\widehat{p}_n - p_n + p_n - f\|_C^2 = \|p_n - f\|_C^2 + 2\langle p_n - f, \widehat{p}_n - p_n \rangle_C + \|\widehat{p}_n - p_n\|_C^2 \\ &\stackrel{(\dagger)}{=} \|p_n - f\|_C^2 + \|\widehat{p}_n - p_n\|_C^2 \\ &= \left\| \sum_{j=n+1}^{\infty} b_j T_j \right\|_C^2 + \left\| \sum_{j=1}^n \frac{\sum_{k=0}^{j-1} q_n}{1 - \sum_{k=0}^{j-1} q_n} b_j T_j \right\|_C^2 \\ &\stackrel{(\ddagger)}{=} \frac{\pi}{2} \sum_{j=n+1}^{\infty} b_j^2 + \frac{\pi}{2} \sum_{j=1}^n \left(\frac{\sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} b_j \right)^2. \end{aligned}$$

Both the second equality (\dagger) and the last equality (\ddagger) come from the orthogonality of Chebyshev polynomials and the following facts:

$$\begin{aligned} p_n - f &: \text{linear combination of } T_{n+1}(x), T_{n+2}(x), \dots, \\ \widehat{p}_n - p_n &: \text{linear combination of } T_0(x), \dots, T_n(x). \end{aligned}$$

The Chebyshev weighted variance can be computed by taking expectation with respect to n :

$$\begin{aligned}
\frac{2}{\pi} \mathbf{E}_n[\|\widehat{p}_n - f\|_C^2] &= \frac{2}{\pi} \sum_{n=0}^{\infty} q_n \|\widehat{p}_n - f\|_C^2 = q_0 \sum_{j=1}^{\infty} b_j^2 + \sum_{n=1}^{\infty} q_n \left(\sum_{j=1}^n \left(\frac{b_j \sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} \right)^2 + \sum_{j=n+1}^{\infty} b_j^2 \right) \\
&= \sum_{j=1}^{\infty} b_j^2 \left(q_0 + \sum_{i=1}^{j-1} q_i + \left(\frac{\sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} \right)^2 \sum_{i=j}^{\infty} q_i \right) \\
&= \sum_{j=1}^{\infty} b_j^2 \left(\sum_{i=0}^{j-1} q_i + \left(\frac{\sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} \right)^2 \left(1 - \sum_{i=0}^{j-1} q_i \right) \right) \\
&= \sum_{j=1}^{\infty} b_j^2 \left(\sum_{i=0}^{j-1} q_i + \frac{\left(\sum_{i=0}^{j-1} q_i \right)^2}{1 - \sum_{i=0}^{j-1} q_i} \right) = \sum_{j=1}^{\infty} b_j^2 \left(\frac{\sum_{i=0}^{j-1} q_i}{1 - \sum_{i=0}^{j-1} q_i} \right).
\end{aligned}$$

This completes the proof of Lemma 2.

B.5.3 Proof of Lemma 4

First, we define the n degree Chebyshev polynomials of the first kind by $T_n(\cdot)$ and the second kind by $U_n(\cdot)$. One important property is that $T'_n(x) := \frac{d}{dx} T_n(x) = nU_{n-1}(x)$ for $n \geq 1$ (see [20]). Consider our unbiased estimator with a single random sample, i.e., a Rademacher vector \mathbf{v} and a degree n drawn from the optimal distribution (11).

From the intermediate result (34) in the proof of Lemma 7, the gradient estimator can be written as following:

$$\psi_i := \frac{\partial}{\partial \theta_i} \mathbf{v}^\top \widehat{p}_n(A(\theta)) \mathbf{v} = \frac{2}{b-a} \mathbf{v}^\top G \mathbf{v} \quad (30)$$

where

$$G = \sum_{j=0}^{n-1} \widehat{b}_{j+1} \left(2 \sum_{r=0}^j {}'T_r(\tilde{A}) \frac{\partial A}{\partial \theta_i} U_{j-r}(\tilde{A}) \right)$$

and $\tilde{A} = \frac{2}{b-a} A(\theta) - \frac{b+a}{b-a} I$ and \sum' implies the summation where the first term is halved. We also note that $\text{tr}(G) = \text{tr} \left(\frac{\partial A}{\partial \theta_i} \widehat{p}'_n(A) \right)$. Here, our goal is to find the upper bound of $\mathbf{E}_{n,\mathbf{v}}[\psi_i^2]$, that is,

$$\frac{(b-a)^2}{4} \mathbf{E}_{n,\mathbf{v}}[\psi_i^2] = \mathbf{E}_{n,\mathbf{v}} \left[(\mathbf{v}^\top G \mathbf{v})^2 \right] = \mathbf{E}_n \left[\mathbf{E}_{\mathbf{v}} \left[(\mathbf{v}^\top G \mathbf{v})^2 \mid n \right] \right].$$

From [14, 2], we have that $\text{Var}_{\mathbf{v}}[\mathbf{v}^\top A \mathbf{v}] = 2(\|A\|_F^2 - \sum_{i=1}^d A_{ii}^2) \leq 2\|A\|_F^2$ and $\mathbf{E}_{\mathbf{v}}[\mathbf{v}^\top G \mathbf{v}] = \text{tr}(G)$ for Rademacher random vector $\mathbf{v} \in [-1, 1]^d$ and $A \in \mathcal{S}^{d \times d}$. Therefore, we have

$$\mathbf{E}_{\mathbf{v}} \left[(\mathbf{v}^\top G \mathbf{v})^2 \mid n \right] = \text{Var}_{\mathbf{v}}[\mathbf{v}^\top G \mathbf{v} \mid n] + \mathbf{E}_{\mathbf{v}}[\mathbf{v}^\top G \mathbf{v} \mid n]^2 \leq 2\|G\|_F^2 + (\text{tr}(G))^2. \quad (31)$$

The first term in (31) is bounded as

$$\begin{aligned}
2\|G\|_F^2 &\leq 2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F^2 \left(\sum_{j=1}^n |\widehat{b}_j| \left(2 \sum_{r=0}^j {}'T_r(\tilde{A}) \|U_{j-r}(\tilde{A})\|_2 \right) \right)^2 \\
&\leq 2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F^2 \left(\sum_{j=1}^n |\widehat{b}_j| \left(2 \sum_{r=0}^j {}'(j-r+1) \right) \right)^2 \\
&= 2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F^2 \left(\sum_{j=1}^n |\widehat{b}_j| j^2 \right)^2
\end{aligned}$$

which the first inequality comes from the triangle inequality of $\|\cdot\|_F$ and the fact that $\|XY\|_F \leq \|X\|_F \|Y\|_2$ for multiplicible matrices X and Y . The inequality in the second line holds from $\|T_i(\tilde{A})\|_2 \leq 1$ and $\|U_i(\tilde{A})\|_2 \leq i + 1$ for $i \geq 0$.

For second term in (31), we use the inequality that $\text{tr}(XY) \leq \|X\|_{\text{nuc}} \|Y\|_2$ for real symmetric matrices X, Y (see Section B.5.8) to obtain

$$\begin{aligned} (\text{tr}(G))^2 &= \left(\text{tr} \left(\frac{\partial A}{\partial \theta_i} \tilde{p}'_n(A) \right) \right)^2 \leq \left\| \frac{\partial A}{\partial \theta_i} \right\|_{\text{nuc}}^2 \|\tilde{p}'_n(A)\|_2^2 \\ &= \left\| \frac{\partial A}{\partial \theta_i} \right\|_{\text{nuc}}^2 \left\| \sum_{j=1}^n \hat{b}_j j U_{j-1}(\tilde{A}) \right\|_2^2 \\ &\leq \left\| \frac{\partial A}{\partial \theta_i} \right\|_{\text{nuc}}^2 \left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \end{aligned}$$

where the equality in the second line uses that $\left(\sum_{j=0}^n \hat{b}_j T_j(x) \right)' = \sum_{j=1}^n \hat{b}_j j U_{j-1}(x)$ and the last inequality holds from $\|U_i(\tilde{A})\|_2 \leq i + 1$. Putting all together into (31) and summing for all $i = 1, \dots, d'$, we obtain that

$$\begin{aligned} \mathbf{E}_{n, \mathbf{v}}[\psi^2] &= \sum_{i=1}^{d'} \mathbf{E}_{n, \mathbf{v}}[\psi_i^2] \leq \frac{4}{(b-a)^2} \sum_{i=1}^{d'} \mathbf{E}_n \left[2 \|G\|_F^2 + (\text{tr}(G))^2 \right] \\ &\leq \frac{4}{(b-a)^2} \sum_{i=1}^{d'} \left(2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F^2 + \left\| \frac{\partial A}{\partial \theta_i} \right\|_{\text{nuc}}^2 \right) \mathbf{E}_n \left[\left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \right] \\ &\leq \frac{4}{(b-a)^2} \left(2 \left\| \frac{\partial A}{\partial \theta} \right\|_F^2 + \sum_{k=1}^{d'} \left\| \frac{\partial A}{\partial \theta_i} \right\|_{\text{nuc}}^2 \right) \mathbf{E}_n \left[\left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \right]. \end{aligned}$$

When we estimate ψ using M Rademacher random vectors $\{\mathbf{v}^{(k)}\}_{k=1}^M$, the variance in (31) is reduced by $1/M$. Hence, we have

$$\begin{aligned} \mathbf{E}_{n, \mathbf{v}}[\psi^2] &\leq \frac{4}{(b-a)^2} \left(\frac{2}{M} \left\| \frac{\partial A}{\partial \theta} \right\|_F^2 + \sum_{k=1}^{d'} \left\| \frac{\partial A}{\partial \theta_i} \right\|_{\text{nuc}}^2 \right) \mathbf{E}_n \left[\left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \right] \\ &\leq \frac{4}{(b-a)^2} \left(\frac{2L_A^2}{M} + d' L_{\text{nuc}}^2 \right) \mathbf{E}_n \left[\left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \right]. \end{aligned}$$

Finally, we introduce the following lemma to bound the right-hand side, where its proof is given in Section B.5.6.

Lemma 9 *Suppose that q_n^* is the optimal degree distribution as defined in (11) and b_j is the Chebyshev coefficients of analytic function f . Define the weighted coefficient \hat{b}_j as $\hat{b}_j = b_j / (1 - \sum_{i=0}^{j-1} q_i^*)$ for $j \geq 0$ and conventionally $q_{-1}^* = 0$. Then, there exists constants $C_1, C_2 > 0$ independent of M, N such that*

$$\sum_{n=1}^{\infty} q_n^* \left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \leq C_1 + \frac{C_2 N^4}{\rho^{2N}}.$$

To sum up, we conclude that

$$\mathbf{E}_{n, \mathbf{v}}[\psi^2] \leq \left(\frac{2L_A^2}{M} + d' L_{\text{nuc}}^2 \right) \left(C_1 + \frac{C_2 N^4}{\rho^{2N}} \right)$$

for some constant $C_1, C_2 > 0$. This completes the proof of Lemma 4.

B.5.4 Proof of Lemma 7

We consider more general case in which $A \in \mathcal{S}^{d \times d}$ is a function of parameter $\theta = [\theta_1, \dots, \theta_d]$, and our goal is to derive a closed form of $\frac{\partial}{\partial \theta_i} \mathbf{v}^\top p_n(A) \mathbf{v}$ with allowing only vector operations. We begin by observing that for any polynomial p_n and symmetric matrix $A \in \mathcal{S}^{d \times d}$, the derivative of $\mathbb{E}_{\mathbf{v}}[\mathbf{v}^\top p_n(A) \mathbf{v}]$ can be expressed by a simple formulation, that is,

$$\frac{\partial}{\partial \theta_i} \mathbb{E}_{\mathbf{v}}[\mathbf{v}^\top p_n(A) \mathbf{v}] = \frac{\partial}{\partial \theta_i} \text{tr}(p_n(A)) = p'_n(A) \frac{\partial A}{\partial \theta_i}.$$

However, it does not holds that

$$\frac{\partial}{\partial \theta_i} \mathbf{v}^\top p_n(A) \mathbf{v} = \frac{\partial}{\partial \theta_i} \text{tr}(p_n(A) \mathbf{v} \mathbf{v}^\top) \neq p'_n(A) \mathbf{v} \mathbf{v}^\top \frac{\partial A}{\partial \theta_i}.$$

for some vector $\mathbf{v} \in \mathbb{R}^d$. This is because of $\frac{\partial}{\partial \theta_i} \text{tr}(A^j \mathbf{v} \mathbf{v}^\top) \neq j A^{j-1} \mathbf{v} \mathbf{v}^\top \frac{\partial A}{\partial \theta_i}$ in general.

If $p_n(x)$ is the truncated Chebyshev series, i.e., $p_n(x) = \sum_{j=0}^n b_j T_j(x)$, we can compute $\frac{\partial}{\partial \theta_i} \mathbf{v}^\top p_n(A) \mathbf{v}$ efficiently using the recursive relation of Chebyshev polynomials, that is,

$$T_{j+1}(x) = 2xT_j(x) - T_{j-1}(x),$$

where $T_j(x)$ is the Chebyshev polynomial of the first-kind with degree j . Let $\mathbf{w}_j := T_j(A) \mathbf{v}$ for $j \geq 0$, and we have that

$$\frac{\partial}{\partial \theta_i} \mathbf{v}^\top p_n(A) \mathbf{v} = \frac{\partial}{\partial \theta_i} \left(\sum_{j=0}^n b_j \mathbf{v}^\top T_j(A) \mathbf{v} \right) = \sum_{j=0}^n b_j \mathbf{v}^\top \left(\frac{\partial}{\partial \theta_i} T_j(A) \mathbf{v} \right) = \sum_{j=0}^n b_j \mathbf{v}^\top \frac{\partial \mathbf{w}_j}{\partial \theta_i}. \quad (32)$$

In the right hand side, $\mathbf{v}^\top \left(\frac{\partial \mathbf{w}_j}{\partial \theta_i} \right)$ can be computed using the recursion $\mathbf{w}_{j+1} = 2A \mathbf{w}_j - \mathbf{w}_{j-1}$:

$$\begin{aligned} \mathbf{v}^\top \frac{\partial \mathbf{w}_{j+1}}{\partial \theta_i} &= \mathbf{v}^\top \frac{\partial}{\partial \theta_i} (2A \mathbf{w}_j - \mathbf{w}_{j-1}) = 2 \mathbf{v}^\top \frac{\partial}{\partial \theta_i} (A \mathbf{w}_j) - \mathbf{v}^\top \frac{\partial \mathbf{w}_{j-1}}{\partial \theta_i} \\ &= 2 \left(\mathbf{v}^\top \frac{\partial A}{\partial \theta_i} \mathbf{w}_j + \mathbf{v}^\top A \frac{\partial \mathbf{w}_j}{\partial \theta_i} \right) - \mathbf{v}^\top \frac{\partial \mathbf{w}_{j-1}}{\partial \theta_i} \end{aligned}$$

where $\mathbf{v}^\top \frac{\partial \mathbf{w}_1}{\partial \theta_i} = \mathbf{v}^\top \frac{\partial A}{\partial \theta_i} \mathbf{v}$ and $\mathbf{v}^\top \frac{\partial \mathbf{w}_0}{\partial \theta_i} = 0$. Applying induction on $j \geq 1$, we can obtain that

$$\mathbf{v}^\top \frac{\partial \mathbf{w}_{j+1}}{\partial \theta_i} = \sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{w}_k^\top \frac{\partial A}{\partial \theta_i} \mathbf{y}_{j-k}, \quad (33)$$

where $\mathbf{y}_{j+1} = 2A \mathbf{y}_j - \mathbf{y}_{j-1} = 2 \mathbf{w}_{j+1} + \mathbf{y}_{j-1}$, $\mathbf{y}_1 = 2A \mathbf{v}$ and $\mathbf{y}_0 = \mathbf{v}$.³ Putting (33) to (32), we get

$$\frac{\partial}{\partial \theta_i} \mathbf{v}^\top p_n(A) \mathbf{v} = \sum_{j=0}^{n-1} b_{j+1} \mathbf{v}^\top \frac{\partial \mathbf{w}_{j+1}}{\partial \theta_i} = \sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{w}_k^\top \frac{\partial A}{\partial \theta_i} \mathbf{y}_{j-k} \right). \quad (34)$$

In case when $A = \theta \theta^\top + \varepsilon I$ and $\theta \in \mathbb{R}^{d \times r}$, it holds that for $\ell = 1, \dots, d$ and $m = 1, \dots, r$,

$$\frac{\partial A}{\partial \theta_{\ell, m}} = \mathbf{e}_\ell \theta_{:, m}^\top + \theta_{:, m} \mathbf{e}_\ell^\top, \quad (35)$$

³ Indeed, $\mathbf{y}_j = U_j(A) \mathbf{v}$ for $j \geq 1$, where $U_j(x)$ is the j -th Chebyshev polynomial of the second-kind.

where $\theta_{:,m} \in \mathbb{R}^d$ is the m -th column of θ and $\mathbf{e}_\ell \in \mathbb{R}^d$ is a unit vector with the index ℓ . Finally, we substitute (35) to (34) to have

$$\begin{aligned}
\left[\frac{\partial}{\partial \theta} \mathbf{v}^\top p_n(A) \mathbf{v} \right]_{\ell,m} &= \frac{\partial}{\partial \theta_{\ell,m}} \mathbf{v}^\top p_n(A) \mathbf{v} = \sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{w}_k^\top \left(\frac{\partial A}{\partial \theta_{\ell,m}} \right) \mathbf{y}_{j-k} \right) \\
&= \sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{w}_k^\top (\mathbf{e}_\ell \theta_{:,m}^\top + \theta_{:,m} \mathbf{e}_\ell^\top) \mathbf{y}_{j-k} \right) \\
&\stackrel{(\dagger)}{=} \sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) (\mathbf{e}_\ell^\top \mathbf{w}_k \mathbf{y}_{j-k}^\top \theta_{:,m} + \mathbf{e}_\ell^\top \mathbf{y}_{j-k} \mathbf{w}_k^\top \theta_{:,m}) \right) \\
&= \mathbf{e}_\ell^\top \left[\sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) (\mathbf{w}_k \mathbf{y}_{j-k}^\top + \mathbf{y}_{j-k} \mathbf{w}_k^\top) \right) \right] \theta_{:,m} \\
&\stackrel{(\ddagger)}{=} \mathbf{e}_\ell^\top \left[\sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) 2 \mathbf{w}_k \mathbf{y}_{j-k}^\top \right) \right] \theta_{:,m} \\
&= \mathbf{e}_\ell^\top \left[2 \sum_{j=0}^{n-1} b_{j+1} \left(\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{w}_k \mathbf{y}_{j-k}^\top \right) \theta \right] \mathbf{e}'_m \\
&= \left[2 \sum_{k=0}^{n-1} (2 - \mathbf{1}_{k=0}) \mathbf{w}_k \left(\sum_{j=k}^{n-1} b_{j+1} \mathbf{y}_{j-k} \right)^\top \theta \right]_{\ell,m}
\end{aligned}$$

where $\mathbf{e}'_m \in \mathbb{R}^r$ is the unit vector with index m satisfying with $\theta_{:,m} = \theta \mathbf{e}'_m$. The equality (†) holds from that $\mathbf{a}^\top \mathbf{b} = \mathbf{b}^\top \mathbf{a}$ for any two vectors \mathbf{a} and \mathbf{b} , and for the equality (‡) it is easy to check that $\sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{w}_k \mathbf{y}_{j-k}^\top = \sum_{k=0}^j (2 - \mathbf{1}_{k=0}) \mathbf{y}_{j-k} \mathbf{w}_k^\top$ using $2\mathbf{w}_j = \mathbf{y}_j - \mathbf{y}_{j-2}$ for $j \geq 2$. Thus,

$$\nabla_{\theta} \mathbf{v}^\top p_n(A) \mathbf{v} = 2 \sum_{k=0}^{n-1} (2 - \mathbf{1}_{k=0}) \mathbf{w}_k \left(\sum_{j=k}^{n-1} b_{j+1} \mathbf{y}_{j-k} \right)^\top \theta$$

This completes the proof of Lemma 7.

B.5.5 Proof of Lemma 8

The proof of Lemma 8 is similar with the proof of Lemma 4. We recall the formulation

$$\psi_i := \frac{\partial}{\partial \theta_i} \mathbf{v}^\top \hat{p}_n(A(\theta)) \mathbf{v} = \frac{2}{b-a} \mathbf{v}^\top G \mathbf{v}$$

where

$$G = \sum_{j=0}^{n-1} \hat{b}_{j+1} \left(2 \sum_{r=0}^j {}'T_r(\tilde{A}) \frac{\partial A}{\partial \theta_i} U_{j-r}(\tilde{A}) \right)$$

and $\tilde{A} = \frac{2}{b-a} A(\theta) - \frac{b+a}{b-a} I$. Define that $\Delta G := G(\theta) - G(\theta')$. Our goal is to find some $\beta \in \mathbb{R}$ such that $\mathbf{E}_{n,\mathbf{v}}[(\mathbf{v}^\top \Delta G \mathbf{v})^2] \leq \beta^2 (\theta_i - \theta'_i)^2$. For notational simplicity, we write that

$$\begin{aligned}
\Delta T_r &:= T_r(\tilde{A}) - T_r(\tilde{A}') = T_r - T'_r, & \Delta U_j &:= U_j(\tilde{A}) - U_j(\tilde{A}') = U_j - U'_j, \\
\Delta A &:= \frac{2}{b-a} (A(\theta) - A(\theta')), & \Delta \frac{\partial A}{\partial \theta} &:= \frac{\partial A(\theta)}{\partial \theta} - \frac{\partial A(\theta')}{\partial \theta}, & \Delta \theta &= \theta - \theta'.
\end{aligned}$$

and ΔG can be expressed as

$$\Delta G = \sum_{j=0}^{n-1} \hat{b}_{j+1} \left(2 \sum_{r=0}^j {}'T_r \frac{\partial A}{\partial \theta_i} U_{j-r} - T'_r \frac{\partial A'}{\partial \theta_i} U'_{j-r} \right).$$

We use similar procedure in the proof of Lemma 4 to obtain

$$\begin{aligned}
\frac{(b-a)^2}{4} \mathbf{E}_{n,\mathbf{v}} [(\psi_i - \psi'_i)^2] &= \mathbf{E}_{n,\mathbf{v}} [(\mathbf{v}^\top \Delta G \mathbf{v})^2] = \mathbf{E}_n [\mathbf{E}_{\mathbf{v}} [(\mathbf{v}^\top \Delta G \mathbf{v})^2 | n]] \\
&= \mathbf{E}_n [\text{Var}_{\mathbf{v}}[\mathbf{v}^\top \Delta G \mathbf{v} | n] + \mathbf{E}_{\mathbf{v}} [\mathbf{v}^\top \Delta G \mathbf{v} | n]^2] \\
&\leq \mathbf{E}_n [2 \|\Delta G\|_F^2 + (\text{tr}(\Delta G))^2]. \tag{36}
\end{aligned}$$

For the first term in (36), we use the triangle inequality to obtain

$$\|\Delta G\|_F \leq \sum_{j=0}^{n-1} |\widehat{b}_{j+1}| \left(\underbrace{2 \sum_{r=0}^j \left\| T_r \frac{\partial A}{\partial \theta_i} U_{j-r} - T'_r \frac{\partial A'}{\partial \theta_i} U'_{j-r} \right\|_F}_{(\ddagger)} \right)$$

and consider that

$$\begin{aligned}
(\ddagger) &\leq \left\| (T_r - T'_r) \frac{\partial A}{\partial \theta_i} U_{j-r} \right\|_F + \left\| T_r \frac{\partial A}{\partial \theta_i} (U_{j-r} - U'_{j-r}) \right\|_F + \left\| T'_r \left(\frac{\partial A}{\partial \theta_i} - \frac{\partial A'}{\partial \theta_i} \right) U'_{j-r} \right\|_F \\
&\leq \|\Delta T_r\|_2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F \|U_{j-r}\|_2 + \|T_r\|_2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F \|\Delta U_{j-r}\|_2 + \|T'_r\|_2 \left\| \Delta \frac{\partial A}{\partial \theta_i} \right\|_F \|U'_{j-r}\|_2 \\
&\leq \|\Delta A\|_2 r^2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F (j-r+1) + \left\| \frac{\partial A}{\partial \theta_i} \right\|_F \frac{(j-r)(j-r+1)(j-r+2)}{3} \|\Delta A\|_2 + \left\| \Delta \frac{\partial A}{\partial \theta_i} \right\|_F (j-r+1)
\end{aligned}$$

where the first inequality is from the triangle inequality of $\|\cdot\|_F$ and the second inequality holds from $\|XY\|_F \leq \|X\|_2 \|Y\|_F$ for multiplicable matrices X, Y and the last is from $\|T_i(\widetilde{A})\|_2 \leq 1$, $\|U_i(\widetilde{A})\|_2 \leq i+1$ for $i \geq 0$ and

$$\|U_i(X+E) - U_i(X)\|_2 \leq \frac{i(i+1)(i+2)}{3} \|E\|_2 \tag{37}$$

for $X, E \in \mathcal{S}^{d \times d}$ satisfying with $\|X+E\|_2, \|X\|_2 \leq 1$ (see Section B.5.8).

Summing (\ddagger) for all $r = 0, 1, \dots, j$, we have

$$\begin{aligned}
\|\Delta G\|_F &\leq \sum_{j=0}^{n-1} |\widehat{b}_{j+1}| \left(\|\Delta A\|_2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F \frac{j(j+1)^2(j+2)}{3} + \left\| \Delta \frac{\partial A}{\partial \theta_i} \right\|_F (j+1)^2 \right) \\
&\leq \max \left(\|\Delta A\|_2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F, \left\| \Delta \frac{\partial A}{\partial \theta_i} \right\|_F \right) \sum_{j=0}^{n-1} |\widehat{b}_{j+1}| \left(\frac{j(j+1)^2(j+2)}{3} + (j+1)^2 \right) \\
&\leq \frac{1}{2} \max \left(\|\Delta A\|_2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F, \left\| \Delta \frac{\partial A}{\partial \theta_i} \right\|_F \right) \sum_{j=0}^{n-1} |\widehat{b}_{j+1}| (j+1)^4.
\end{aligned}$$

If one estimates ψ and ψ' using M Rademacher random vectors, the variance of $\mathbf{v}^\top \Delta G \mathbf{v}$ is reduced by $1/M$ so that we have

$$2 \|\Delta G\|_F^2 \leq \frac{1}{2M} \max \left(\|\Delta A\|_2^2 \left\| \frac{\partial A}{\partial \theta_i} \right\|_F^2, \left\| \Delta \frac{\partial A}{\partial \theta_i} \right\|_F^2 \right) \left(\sum_{j=1}^n |\widehat{b}_j| j^4 \right)^2$$

For the second term in (36), it holds that

$$\begin{aligned}
\mathrm{tr}(\Delta G) &= \mathrm{tr}\left(\frac{\partial A}{\partial \theta_i}(\tilde{p}'_n(A) - \tilde{p}'_n(A'))\right) \leq \left\|\frac{\partial A}{\partial \theta_i}\right\|_F \|\tilde{p}'_n(A) - \tilde{p}'_n(A')\|_F \\
&\leq \left\|\frac{\partial A}{\partial \theta_i}\right\|_F \sum_{j=1}^n |\widehat{b}_j| j \|U_{j-1}(\tilde{A}) - U_{j-1}(\tilde{A}')\|_F \\
&\leq \left\|\frac{\partial A}{\partial \theta_i}\right\|_F \|\Delta A\|_F \sum_{j=1}^n |\widehat{b}_j| \frac{(j^2 - 1)j^2}{3} \\
&\leq \left\|\frac{\partial A}{\partial \theta_i}\right\|_F \frac{\|\Delta A\|_F}{3} \sum_{j=1}^n |\widehat{b}_j| j^4.
\end{aligned}$$

where the inequality in the first line holds from matrix version Cauchy-Schwarz inequality, the inequality in the second line holds from $\tilde{p}'_n(x) = \left(\sum_{j=0}^n \widehat{b}_j T_j(x)\right)' = \sum_{j=1}^n \widehat{b}_j j U_{j-1}(x)$ and inequality in the third line holds from (37).

Putting all together into (36), we obtain that

$$\begin{aligned}
\mathbf{E}_{n,\mathbf{v}}\left[(\psi_i - \psi'_i)^2\right] &= \mathbf{E}_n\left[2\|\Delta G\|_F^2 + (\mathrm{tr}(\Delta G))^2\right] \\
&\leq \left(\frac{1}{2M} \max\left(\|\Delta A\|_2^2 \left\|\frac{\partial A}{\partial \theta_i}\right\|_F^2, \left\|\Delta \frac{\partial A}{\partial \theta_i}\right\|_F^2\right) + \left\|\frac{\partial A}{\partial \theta_i}\right\|_F^2 \frac{\|\Delta A\|_F^2}{9}\right) \mathbf{E}_n\left[\left(\sum_{j=1}^n |\widehat{b}_j| j^4\right)^2\right] \\
&\leq \left(\left(\frac{1}{2M} + \frac{1}{9}\right) \left\|\frac{\partial A}{\partial \theta_i}\right\|_F^2 \|\Delta A\|_F^2 + \frac{1}{2M} \left\|\Delta \frac{\partial A}{\partial \theta_i}\right\|_F^2\right) \mathbf{E}_n\left[\left(\sum_{j=1}^n |\widehat{b}_j| j^4\right)^2\right] \\
&\leq \left(\left(\frac{1}{2M} + \frac{1}{9}\right) \left\|\frac{\partial A}{\partial \theta_i}\right\|_F^2 \frac{4L_A^2 \|\Delta\theta\|_2^2}{(b-a)^2} + \frac{1}{2M} \left\|\Delta \frac{\partial A}{\partial \theta_i}\right\|_F^2\right) \mathbf{E}_n\left[\left(\sum_{j=1}^n |\widehat{b}_j| j^4\right)^2\right]
\end{aligned}$$

where the inequality in the second line holds from $\max(a, b) \leq a + b$ for $a, b \in \mathbb{R}^+$ and the inequality in the third line holds from the Lipschitz continuity on A (assumption $\mathcal{A}(2)$), formally,

$$\|A(\theta) - A(\theta')\|_2 \leq \|A(\theta) - A(\theta')\|_F \leq L_A \|\theta - \theta'\|_2.$$

Summing the above for all $i = 1, 2, \dots, d'$ and using that $\|\partial A / \partial \theta\|_F \leq L_A$ and $\|\Delta(\partial A / \partial \theta)\|_F \leq \beta_A \|\Delta\theta\|_2$, we get

$$\mathbf{E}_{n,\mathbf{v}}\left[\|\psi - \psi'\|_2^2\right] \leq D_0 \left(\frac{L_A^4 + \beta_A^2}{M} + L_A^4\right) \|\Delta\theta\|_2^2 \mathbf{E}_n\left[\left(\sum_{j=1}^n |\widehat{b}_j| j^4\right)^2\right]$$

for some constant $D_0 > 0$.

To bound the right-hand side, we introduce the following lemma, whose proof is in Section B.5.7.

Lemma 10 *Suppose that q_n^* is the optimal degree distribution as defined in (11) and b_j is the Chebyshev coefficients of analytic function f . Define the weighted coefficient \widehat{b}_j as $\widehat{b}_j = b_j / (1 - \sum_{i=0}^{j-1} q_i^*)$ for $j \geq 0$ and conventionally $q_{-1}^* = 0$. Then, there exists constants $D'_1, D'_2 > 0$ independent of M, N such that*

$$\sum_{n=1}^{\infty} q_n^* \left(\sum_{j=1}^n |\widehat{b}_j| j^4\right)^2 \leq D'_1 + \frac{D'_2 N^8}{\rho^{2N}}.$$

Therefore, we obtain the result that

$$\mathbf{E}_{n,\mathbf{v}} [\|\psi - \psi'\|_2^2] \leq \beta^2 \|\theta - \theta'\|_2^2 \quad (38)$$

where

$$\beta^2 := \left(\frac{L_A^4 + \beta_A^2}{M} + L_A^4 \right) \left(D_1 + \frac{D_2 N^8}{\rho^{2N}} \right)$$

Under the assumption that $g(\theta)$ is β_g -smooth function (assumptio $\mathcal{A}(2)$), we have that

$$\|\nabla g(\theta) - \nabla g(\theta')\|_2^2 \leq \beta_g^2 \|\theta - \theta'\|_2^2. \quad (39)$$

Summing both (38) and (39), it yields that

$$\mathbf{E}_{n,\mathbf{v}} \left[\|\psi - \psi'\|_2^2 + \|\nabla g(\theta) - \nabla g(\theta')\|_2^2 \right] \leq (\beta^2 + \beta_g^2) \|\theta - \theta'\|_2^2.$$

Using $\|a + b\| \leq 2(\|a\|^2 + \|b\|^2)$ again, we conclude that

$$\mathbf{E}_{n,\mathbf{v}} \left[\|\psi + \nabla g(\theta) - \psi' - \nabla g(\theta')\|_2^2 \right] \leq 2(\beta^2 + \beta_g^2) \|\theta - \theta'\|_2^2.$$

This completes the proof of Lemma 8.

B.5.6 Proof of Lemma 9

Recall that the optimal degree distribution as

$$q_i^* = \begin{cases} 0 & \text{for } i < K \\ 1 - (N - K)(\rho - 1)\rho^{-1} & \text{for } i = K \\ (N - K)(\rho - 1)^2 \rho^{-i-1+K} & \text{for } i > K. \end{cases}$$

where $K = \max(0, N - \lfloor \frac{\rho}{\rho-1} \rfloor)$. We first use the upper bound on the coefficients from (2), i.e., $|b_j| \leq 2U/\rho^j$ to obtain

$$\sum_{n=1}^{\infty} q_n^* \left(\sum_{j=1}^n |\widehat{b}_j| j^2 \right)^2 = \sum_{n=K}^{\infty} q_n^* \left(\sum_{j=1}^n |\widehat{b}_j| j^2 \right)^2 \leq 4U^2 \sum_{n=K}^{\infty} q_n^* \left(\sum_{j=1}^n \frac{j^2}{(1 - \sum_{i=0}^{j-1} q_i^*) \rho^j} \right)^2 \quad (40)$$

To express (40) more simple, we define that

$$\Lambda := \sum_{j=1}^K \frac{j^2}{(1 - \sum_{i=0}^{j-1} q_i^*) \rho^j} = \sum_{j=1}^K \frac{j^2}{\rho^j} \leq \frac{\rho(\rho+1)}{(\rho-1)^3}$$

which equals to the second term in the summation (40) when $n = K$. For $n \geq K + i, i \geq 1$, we get

$$\sum_{j=1}^{K+i} \frac{j^2}{(1 - \sum_{i=0}^{j-1} q_i^*) \rho^j} = \Lambda + \frac{\sum_{j=1}^i (K+j)^2}{(N-K)(\rho-1)\rho^K}. \quad (41)$$

Putting q_i^* and (41) to the right hand side of (40), we have

$$\begin{aligned} & \left(1 - (N-K) \frac{\rho-1}{\rho} \right) \Lambda^2 + (N-K) \left(\frac{\rho-1}{\rho} \right)^2 \left(\Lambda + \frac{(K+1)^2}{(N-K)(\rho-1)\rho^K} \right)^2 \\ & + (N-K) \left(\frac{\rho-1}{\rho} \right)^2 \frac{1}{\rho} \left(\Lambda + \frac{\sum_{j=1}^2 (K+j)^2}{(N-K)(\rho-1)\rho^K} \right)^2 \\ & + (N-K) \left(\frac{\rho-1}{\rho} \right)^2 \frac{1}{\rho^2} \left(\Lambda + \frac{\sum_{j=1}^3 (K+j)^2}{(N-K)(\rho-1)\rho^K} \right)^2 \\ & + \dots \end{aligned}$$

Rearranging all terms with respect to Λ , we obtain that

$$\Lambda^2 + \frac{2(\rho-1)}{\rho^{K+1}} \left(\sum_{i=1}^{\infty} \frac{\sum_{j=1}^i (K+j)^2}{\rho^i} \right) \Lambda + \frac{1}{(N-K)\rho^{2K+1}} \left(\sum_{i=1}^{\infty} \frac{\left(\sum_{j=1}^i (K+j)^2 \right)^2}{\rho^i} \right).$$

Note that

$$\sum_{i=1}^{\infty} \frac{\sum_{j=1}^i (K+j)^2}{\rho^i} = \frac{K^2\rho(\rho-1)^2 + 2K\rho^2(\rho-1) + \rho^2(\rho+1)}{(\rho-1)^4}$$

and

$$\sum_{i=1}^{\infty} \frac{\left(\sum_{j=1}^i (K+j)^2 \right)^2}{\rho^i} = \text{poly}(K^4).$$

Since $K = O(N)$ and $N - K = O(1)$, we can conclude that

$$\sum_{n=1}^{\infty} q_n^* \left(\sum_{j=1}^n |\hat{b}_j| j^2 \right)^2 \leq C_1 + C_2 \frac{N^4}{\rho^{2N}}$$

for some constants $C_1, C_2 > 0$ not depend on N .

B.5.7 Proof of Lemma 10

The proof of Lemma 10 is straightforward from that of Lemma 9. One can replace j^2 into j^4 in the proof of Lemma 9, which results in N^8 dependence. We omit the details of the proof.

B.5.8 Proof of other lemmas

Lemma 11 *Suppose that $A, A + E \in \mathbb{R}^{d \times d}$ are symmetric matrices and they have eigenvalues in $[-1, 1]$. Let T_i and U_i be the first and the second kind of Chebyshev basis polynomial with degree $i \geq 0$, respectively. Then, it holds that*

$$\|T_i(A + E) - T_i(A)\| \leq i^2 \|E\|, \quad \|U_i(A + E) - U_i(A)\| \leq \frac{i(i+1)(i+2)}{3} \|E\|.$$

where $\|\cdot\|$ can be $\|\cdot\|_2$ (spectral norm) or $\|\cdot\|_F$ (Frobenius norm).

Proof. Denote $R_i := T_i(A + E) - T_i(A)$. From the recursive relation of Chebyshev polynomial, i.e., $T_{j+1}(x) = 2AT_j(x) - T_{j-1}(x)$, R_i has following property:

$$R_{i+1} = 2(A + E)R_i - R_{i-1} + 2ET_i(A)$$

for $i \geq 1$ where $R_1 = E$, $R_0 = \mathbf{0}$. By induction on i , it is easy to show that

$$R_{i+1} = 2 \sum_{j=0}^i U_{i-j}(A + E) E T_j(A)$$

where $U_j(x)$ is the Chebyshev polynomial of the second kind. Therefore, we have

$$\begin{aligned} \|R_{i+1}\|_F &\leq 2 \sum_{j=0}^i \|U_{i-j}(A + E) E T_j(A)\|_F \\ &\leq 2 \sum_{j=0}^i \|U_{i-j}(A + E)\|_2 \|E\|_F \|T_j(A)\|_2 \\ &\leq 2 \sum_{j=0}^i (i+1-j) \|E\|_F = (i+1)^2 \|E\|_F \end{aligned}$$

where the second inequality holds from $\|YX\|_F = \|XY\|_F \leq \|X\|_2 \|Y\|_F$ for matrices X, Y . This also holds for $\|\cdot\|_2$ giving that $\|R_{i+1}\|_2 \leq (i+1)^2 \|E\|_2$. Similarly, we denote $Y_i := U_i(A+E) - U_i(A)$. By induction on i , it is easy to show that

$$Y_{i+1} = 2 \sum_{j=0}^i U_{i-j}(A+E) E U_j(A)$$

Then, we have that for $i \geq 0$

$$\begin{aligned} \|Y_{i+1}\|_F &\leq 2 \sum_{j=0}^i \|U_{i-j}(A+E) E U_j(A)\|_F \\ &\leq 2 \sum_{j=0}^i \|U_{i-j}(A+E)\|_2 \|E\|_F \|U_j(A)\|_2 \\ &\leq 2 \sum_{j=0}^i (i+1-j)(j+1) \|E\|_F \\ &= \frac{(i+1)(i+2)(i+3)}{3} \|E\|_F. \end{aligned}$$

This also holds for $\|\cdot\|_2$ giving that $\|Y_{i+1}\|_2 \leq \frac{(i+1)(i+2)(i+3)}{3} \|E\|_2$. This completes the proof of Lemma 11. ■

Lemma 12 For symmetric matrices $A, B \in \mathcal{S}^{d \times d}$, it holds that $\text{tr}(AB) \leq \|A\|_{\text{nuc}} \|B\|_2$.

Proof. Since A is real symmetric, it can be written as $A = \sum_{i=1}^d \lambda_i \mathbf{u}_i \mathbf{u}_i^\top$ where λ_i and \mathbf{u}_i is i -th eigenvalue and eigenvector, respectively. Then, the result follows that

$$\begin{aligned} \text{tr}(AB) &= \sum_{i=1}^d \lambda_i \text{tr}(\mathbf{u}_i \mathbf{u}_i^\top B) = \sum_{i=1}^d \lambda_i \mathbf{u}_i^\top B \mathbf{u}_i \\ &\leq \sum_{i=1}^d |\lambda_i| \mathbf{u}_i^\top B \mathbf{u}_i \\ &\leq \sum_{i=1}^d |\lambda_i| \|B\|_2 = \|A\|_{\text{nuc}} \|B\|_2. \end{aligned}$$

This completes the proof of Lemma 12. ■