
Optimal Randomized First-Order Methods for Least-Squares Problems

Jonathan Lacotte Mert Pilanci

Abstract

We provide an exact analysis of a class of randomized algorithms for solving overdetermined least-squares problems. We consider first-order methods, where the gradients are pre-conditioned by an approximation of the Hessian, based on a subspace embedding of the data matrix. This class of algorithms encompasses several randomized methods among the fastest solvers for least-squares problems. We focus on two classical embeddings, namely, Gaussian projections and subsampled randomized Hadamard transforms (SRHT). Our key technical innovation is the derivation of the limiting spectral density of SRHT embeddings. Leveraging this novel result, we derive the family of normalized orthogonal polynomials of the SRHT density and we find the optimal pre-conditioned first-order method along with its rate of convergence. Our analysis of Gaussian embeddings proceeds similarly, and leverages classical random matrix theory results. In particular, we show that for a given sketch size, SRHT embeddings exhibits a faster rate of convergence than Gaussian embeddings. Then, we propose a new algorithm by optimizing the computational complexity over the choice of the sketching dimension. To our knowledge, our resulting algorithm yields the best known complexity for solving least-squares problems with no condition number dependence.

1. Introduction

We study the performance of a randomized method, namely, the Hessian sketch [26], in the context of (overdetermined) least-squares problems,

$$x^* := \operatorname{argmin}_{x \in \mathbb{R}^d} \left\{ f(x) := \frac{1}{2} \|Ax - b\|^2 \right\}, \quad (1)$$

where $A \in \mathbb{R}^{n \times d}$ is a given data matrix with $n \geq d$ and $b \in \mathbb{R}^n$ is a vector of observations. For simplicity of notations, we will assume throughout this work that $\operatorname{rank}(A) = d$.

Many works have developed randomized algorithms [4, 27, 9, 25] for solving (1), based on sketching methods. The lat-

ter involve using a random matrix $S \in \mathbb{R}^{m \times n}$ to project the data A and/or b to a lower dimensional space ($m \ll n$), and then approximately solving the least-squares problem using the sketch SA and/or Sb . The most classical sketch is a matrix $S \in \mathbb{R}^{m \times n}$ with independent and identically distributed (i.i.d.) Gaussian entries $\mathcal{N}(0, m^{-1})$, for which forming SA requires in general $\mathcal{O}(mnd)$ basic operations (using classical matrix multiplication). This is larger than the cost $\mathcal{O}(nd^2)$ of solving (1) through standard matrix factorization methods, provided that $m \geq d$. Another well-studied embedding is the (truncated) $m \times n$ Haar matrix S , whose rows are orthonormal and with range uniformly distributed among the subspaces of \mathbb{R}^n with dimension m . However, it requires time $\mathcal{O}(nm^2)$ to be formed, through a Gram-Schmidt procedure, which is also larger than $\mathcal{O}(nd^2)$. An alternative embedding which verifies orthogonality properties is the SRHT [1], which is based on the Walsh-Hadamard transform. Due to the recursive structure of the latter, the sketch SA can be formed in $\mathcal{O}(nd \log m)$ time, so that the SRHT is often viewed as a standard reference point for comparing sketching algorithms.

It has been observed in several contexts that random projections with i.i.d. entries degrade the performance of the approximate solution compared to orthogonal projections [17, 18, 8, 7]. Consequently, along with computational considerations, this suggests to consider the SRHT over Gaussian or Haar projections. On the other hand, in order to pick optimal algorithm's parameters, it is usually necessary to have a tight characterization of the spectrum of $C_S := U^\top S^\top S U$, where U is the matrix of left singular vectors of A , which is the case for Gaussian embeddings.

Using the standard prediction (semi-)norm $\|A(\tilde{x} - x^*)\|^2$ as the evaluation criterion for an approximate solution \tilde{x} , iterative methods (e.g., gradient descent or the conjugate gradient algorithm) have time complexity which usually scales proportionally to the condition number κ of the matrix A – defined as the ratio between the largest and smallest singular values of A –, and this becomes prohibitively large when $\kappa \gg 1$. To address the latter issue, we introduce a pre-conditioning method, namely, the Hessian sketch [26], which approximates the Hessian $H = A^\top A$ of $f(x)$ by $H_S = A^\top S^\top S A$. Then, we consider the following class of pre-conditioned first-order methods (also referred to as a

quasi-Newton method in the optimization literature),

$$x_t \in x_0 + H_S^{-1} \cdot \text{span} \{ \nabla f(x_0), \dots, \nabla f(x_{t-1}) \}, \quad (2)$$

Several versions of (2) have been studied. For instance, it has been recently shown by [22] that the Heavy-ball update

$$x_{t+1} = x_t - \mu_t H_S^{-1} \nabla f(x_t) + \beta_t (x_t - x_{t-1}) \quad (3)$$

yields a sequence of iterates whose convergence rate does not depend on the spectrum of A , but only on the concentration of the matrix C_S around the identity, and they show that this convergence rate is equal to the ratio d/m both for Gaussian and SRHT embeddings. Notably, this rate does not depend on the sample size n . For a Gaussian embedding, this makes intuitively sense since the limiting spectral distribution of C_S is the Marchenko-Pastur law [19] with scale parameter ρ , edge eigenvalues $a = (1 - \sqrt{\rho})^2$ and $b = (1 + \sqrt{\rho})^2$, and density

$$\mu_\rho(x) = \frac{\sqrt{(b-x)_+(x-a)_+}}{2\pi\rho x}, \quad (4)$$

where $y_+ = \max\{y, 0\}$, and it does not depend on the sample size n but only on the limit ratio $\rho := \lim d/m$. However, for a SRHT embedding, it is unclear if the dimension n affects the best achievable convergence rate. In a related vein, the authors of [15] considered the Heavy-ball update (3) where at each iteration the sketching SRHT matrix $S = S_t$ is refreshed (i.e., re-sampled independently of S_0, \dots, S_{t-1}) so that $H_S = H_{S_t}$ is also re-computed. They show that Haar and SRHT embeddings yield the same convergence rate $\rho_h^{\text{ref}} := \rho \cdot \frac{\xi(1-\xi)}{\gamma^2 + \xi - 2\gamma\xi}$, which indeed depends on the three relevant dimensions m, d, n through the aspect ratios $\rho, \gamma := \lim d/n$ and $\xi := \lim m/n$, which is always strictly smaller than ρ , the convergence rate one would obtain with refreshed Gaussian embeddings [13].

In this work, we consider a definition of the SRHT slightly different than its classical version [1], which has been introduced in [7, 16]. For an integer $n = 2^p$ with $p \geq 1$, the Walsh-Hadamard transform is defined recursively as $H_n = \frac{1}{\sqrt{2}} \begin{bmatrix} H_{n/2} & H_{n/2} \\ H_{n/2} & -H_{n/2} \end{bmatrix}$ with $H_1 = 1$. Our transform $A \mapsto SA$ first *randomly permutes* the rows of A , before applying the classical transform. This has negligible cost $\mathcal{O}(n)$ compared to the cost $\mathcal{O}(nd \log m)$ of the matrix multiplication $A \mapsto SA$, and *breaks the non-uniformity* in the data. That is, we define the $n \times n$ subsampled randomized Hadamard matrix as $S = BH_n DP$, where B is an $n \times n$ diagonal sampling matrix of i.i.d. Bernoulli random variables with success probability m/n , H_n is the $n \times n$ Walsh-Hadamard matrix, D is an $n \times n$ diagonal matrix of i.i.d. sign random variables, equal to ± 1 with equal probability, and $P \in \mathbb{R}^{n \times n}$ is a uniformly distributed permutation matrix. At the last step, we discard the zero rows

of S , so that it becomes an $\tilde{m} \times n$ orthogonal matrix with $\tilde{m} \sim \text{Binomial}(m/n, n)$, and the ratio \tilde{m}/n concentrates fast around ξ while $n \rightarrow \infty$. Although the dimension \tilde{m} is random, we refer to S as an $m \times n$ SRHT matrix.

We will focus exclusively on (pre-conditioned) *first-order methods* of the form (2) with a fixed embedding S , and our goal is to answer the following questions. What are the best achievable convergence rates for, respectively, Gaussian and SRHT embeddings? What are the corresponding optimal algorithms? How do these rates compare to each other and to that of state-of-the-art randomized iterative methods for solving (1)?

1.1. Technical background, notations and assumptions

We will assume that $\lim_{n \rightarrow \infty} \frac{d}{n} = \gamma \in (0, 1)$, $\lim_{n \rightarrow \infty} \frac{m}{n} = \xi \in (\gamma, 1)$ and $\rho = \frac{\gamma}{\xi} \in (0, 1)$. We denote $\|z\| \equiv \|z\|_2$ the Euclidean norm of a vector z , $\|M\|_2$ the operator norm of a matrix M , and $\|M\|_F$ its Frobenius norm. Given a sequence of iterates $\{x_t\}$, we denote the error at time t by $\Delta_t = U^\top A(x_t - x^*)$. Note that $\|\Delta_t\|^2 = \|A(x_t - x^*)\|^2$. Our evaluation criterion is the error $\lim_{n \rightarrow \infty} \mathbb{E}[\|\Delta_t\|^2] / \mathbb{E}[\|\Delta_0\|^2]$, and we call its (asymptotic) rate of convergence the quantity $\limsup_{t \rightarrow \infty} (\lim_{n \rightarrow \infty} \mathbb{E}[\|\Delta_t\|^2] / \mathbb{E}[\|\Delta_0\|^2])^{1/t}$.

As we focus on infinite-dimensional regimes, our technical analysis is based on asymptotic random matrix theory, and we refer the reader to [5, 23, 35] for an extensive introduction to this field. For a random Hermitian matrix M_n of size $n \times n$, the empirical spectral distribution (e.s.d.) of M_n is the (cumulative) distribution function of its eigenvalues $\lambda_1, \dots, \lambda_n$, i.e., $F_{M_n}(x) := \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{\lambda_j \leq x\}$ for $x \in \mathbb{R}$, which has density $f_{M_n}(x) = \frac{1}{n} \sum_{j=1}^n \delta_{\lambda_j}(x)$ with δ_λ the Dirac measure at λ . Due to the randomness of the eigenvalues, F_{M_n} is random. The relevant aspect of some classes of large $n \times n$ symmetric random matrices M_n is that, almost surely, the e.s.d. F_{M_n} converges weakly towards a non-random distribution F , as $n \rightarrow \infty$. This function F , if it exists, will be called the *limiting spectral distribution* (l.s.d.) of M_n . Key to our analysis is the notion of orthogonal polynomials, which are fundamental both in optimization [28] and in random matrix theory. We write $\mathbb{R}_t[X]$ the set of real polynomials with degree less than t , and $\mathbb{R}_t^0[X]$ the set of polynomials $P \in \mathbb{R}_t[X]$ such that $P(0) = 1$. For a complex number $z \in \mathbb{C}$, we denote respectively by $\text{Re}(z)$ and $\text{Im}(z)$ its real and imaginary parts, and we use \mathbb{C}_+ for the complex numbers with positive imaginary parts, and \mathbb{R}_+ for the positive real numbers. For two sequences of real positive numbers $\{a_t\}$ and $\{b_t\}$, we write $a_t \asymp b_t$ if $\liminf \frac{a_t}{b_t} > 0$ and $\limsup \frac{a_t}{b_t} < \infty$.

We will assume that the first iterate x_0 is random such that $\mathbb{E}[x_0] = 0$, and, that the condition number of the matrix

$U^\top A \mathbb{E}[x_0 x_0^\top] A^\top U + U^\top b b^\top U$ remains bounded as the dimensions grow. Essentially, this states that the condition number of A does not degenerate to $+\infty$ as the dimensions grow.

1.2. Overview of our results and contributions

We have the following contributions.

1. For Gaussian embeddings, we characterize the algorithm (Algorithm 1) which attains the infimum of the error $\lim_{n \rightarrow \infty} \mathbb{E}[\|\Delta_t^2\|] / \mathbb{E}[\|\Delta_0\|^2]$, and we show that it corresponds to the Heavy-ball method with constant step size $\mu_t = (1 - \rho)^2$ and momentum parameter $\beta_t = \rho$. Further, we show that the infimum of the error is equal to ρ^t .
2. For SRHT embeddings, we perform a similar analysis, and find the optimal first-order method (Algorithm 2). Notably, it is a Heavy-ball update with non-constant step sizes and momentum parameters. Further, we show that its rate of convergence is $\rho_h := \rho \cdot \frac{1-\xi}{1-\gamma}$, which is always strictly smaller than ρ and ρ_h^{ref} , i.e., Algorithm 2 has uniformly better convergence rate than that of Gaussian embeddings or the Heavy-ball method with refreshed SRHT embeddings. Even though our theoretical results hold asymptotically, we verify empirically that our theoretical predictions hold, even for sample sizes $n \gtrsim 1000$, and that Algorithm 2 is faster in practice than the other aforementioned algorithms.
3. We characterize explicitly the density $f_{h,r}$ of the l.s.d of the matrix $\frac{n}{m} C_S$, which is given by

$$f_{h,r}(x) = \frac{\sqrt{(\Lambda_{h,r} - x)_+(x - \lambda_{h,r})_+}}{2\pi\rho x(1 - \xi x)}, \quad (5)$$

where the edge (i.e., extreme) eigenvalues are $\lambda_{h,r} = (\sqrt{1-\gamma} - \sqrt{(1-\xi)\rho})^2$ and $\Lambda_{h,r} = (\sqrt{1-\gamma} + \sqrt{(1-\xi)\rho})^2$. This characterization of the limiting density is of independent interest, as it might have several implications beyond least-squares optimization.

4. Finally, we show that Algorithm 2 has the best known complexity to solve (1) with no condition number dependence.

Except for the time complexity results, all our results regarding the SRHT hold exactly the same with Haar embeddings, since they both yield the same limiting spectral distributions.

1.3. Other related work

The design of optimal first-order methods for quadratic optimization problems has been recently considered in [24]. In contrast, they assume the data matrix to be random and

they require its spectrum to be known beforehand, which is often impractical. On the other hand, our class of first-order methods applies a randomized pre-conditioning, so that only the spectral distribution of the matrix C_S is required, and this is universal, i.e., independent of the spectrum of A . Therefore, by characterizing the l.s.d. of C_S for some classical embeddings, we are able to optimize the *exact* error for any data matrix A . We note that existing methods do not directly minimize the error, but a worst-case upper bound as in Chebyshev iteration and Conjugate Gradient methods [29].

Besides the Hessian sketch, there are many other efficient pre-conditioned iterative methods which aim to address the aforementioned conditioning issue, based on an SRHT sketch of the data (or closely related sketches based on the Fourier transform). Randomized *right* pre-conditioning methods [4, 27] compute first a matrix P – which itself depends on SA – such that the condition number of AP^{-1} is $\mathcal{O}(1)$, and then apply any standard iterative algorithm to the pre-conditioned least-squares objective $\|AP^{-1}y - b\|^2$. SRHT sketches are also used for a wide range of applications across numerical linear algebra, statistics and convex optimization, such as low-rank matrix factorization [10, 33], kernel regression [34], random subspace optimization [14], or, sketch and solve linear regression [7]. Hence, a refined analysis of the SRHT may also lead to better algorithms in these fields.

2. Optimal first-order method for classical embeddings

Let S be an $m \times n$ Gaussian or SRHT embedding. Denote by μ the l.s.d. of C_S . We say that a family of polynomials $\{R_k\}$ is orthogonal with respect to μ if $\int R_k R_\ell d\mu = 0$ for any $k \neq \ell$. The next result establishes the link between polynomials and the pre-conditioned first-order methods (2) we consider, and its proof is deferred to Appendix B.1.

Lemma 2.1. *Let $\{x_t\}$ be generated by some first-order method (2). Then, for any iteration $t \geq 0$, there exists a polynomial $p_t \in \mathbb{R}_t^0[X]$ such that $\Delta_t = p_t(C_S^{-1}) \cdot \Delta_0$. Further, it holds that*

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[\|\Delta_t\|^2]}{\mathbb{E}[\|\Delta_0\|^2]} = \int_{\mathbb{R}} p_t^2(\lambda^{-1}) d\mu(\lambda). \quad (6)$$

Thus, the best achievable error is lower bounded by the infimum of the following variational problem,

$$\mathcal{L}_{\mu,t}^* := \min_{p \in \mathbb{R}_t^0[X]} F_\mu(p), \quad (7)$$

where $F_\mu(p) := \int p^2(\lambda^{-1}) d\mu(\lambda)$. Using the change of variable $x = 1/\lambda$ and setting $d\nu(x) = x^{-1}d\mu(x^{-1})$, we have that $F_\mu(p) = G_\nu(p)$ where $G_\nu(p) := \int p^2(x) \frac{1}{x} d\nu(x)$.

The optimal polynomial can be constructed by leveraging the following result.

Lemma 2.2. *Let ν be some measure with bounded support in $(0, +\infty)$, and suppose that $\{\Pi_t\}$ is a family of orthogonal polynomials with respect to ν such that $\deg(\Pi_t) = t$ and $\Pi_t(0) = 1$. Then, the polynomial Π_t is the unique solution of the optimization problem $\min G_\nu(p)$ over $p \in \mathbb{R}_t^0[X]$.*

Proof. Let $p \in \mathbb{R}_t^0[X]$. Since $\Pi_t(0) = 1$, the polynomial $(p - \Pi_t)$ has a root at 0. Hence, $(p - \Pi_t)(x) = xQ(x)$ with $Q \in \mathbb{R}_{t-1}[X]$. Then,

$$\begin{aligned} G_\nu(p) &= \int p^2(x)x^{-1}d\nu(x) \\ &= \int \Pi_t^2(x)x^{-1}d\nu(x) + 2 \int \Pi_t Q(x)d\nu(x) \\ &\quad + \int xQ^2(x)d\nu(x). \end{aligned}$$

The cross-term is equal to 0 since Q in the span of Π_0, \dots, Π_{t-1} , which are orthogonal to Π_t . The third term is non-negative, and equal to 0 if and only if that $Q = 0$. Therefore, the unique solution to (7) is Π_t . \square

Based on such an orthogonal family $\{\Pi_t\}$, we aim to derive a first-order method which achieves the lower bound $\mathcal{L}_{\mu,t}^*$. We recall a standard result, that is, for such a family of polynomials $\{\Pi_t\}$, there exist sequences $\{a_t\}$ and $\{b_t\}$ such that $\Pi_0(x) = 1$, $\Pi_1(x) = 1 + b_1x$ and for any $t \geq 2$,

$$\Pi_t(x) = (a_t + b_t x)\Pi_{t-1}(x) + (1 - a_t)\Pi_{t-2}(x). \quad (8)$$

Then we can construct an optimal first-order method according to the following result, which is inspired by the work of [24] and whose proof is deferred to Appendix A.1.

Theorem 1. *Given $x_0 \in \mathbb{R}^d$, set $x_1 = x_0 + b_1 H_S^{-1} \nabla f(x_0)$, and for $t \geq 2$,*

$$x_t = x_{t-1} + b_t H_S^{-1} \nabla f(x_{t-1}) + (1 - a_t)(x_{t-2} - x_{t-1}). \quad (9)$$

Then, the sequences of iterates $\{x_t\}$ is asymptotically optimal, i.e.,

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}\|\Delta_t\|^2}{\mathbb{E}\|\Delta_0\|^2} = \mathcal{L}_{\mu,t}^*. \quad (10)$$

Consequently, a strategy to find the optimal first-order method proceeds as follows. First, we characterize the l.s.d. μ of the matrix C_S , and we find the polynomial $\Pi_t \in \mathbb{R}_t^0[X]$ which achieves the lower bound $\mathcal{L}_{\mu,t}^*$. Then, according to Theorem 1, we build from the three-terms recursion (8) of the orthogonal polynomials $\{\Pi_t\}$ a first-order method which yields an asymptotically optimal sequence of iterates $\{x_t\}$. Our analysis of the Gaussian case

is based on standard random matrix theory results, that we recall in details as we leverage them for the analysis of the SRHT case. For the latter, most technicalities actually lie in characterizing the l.s.d. μ of C_S , and in constructing an orthogonal basis of polynomials for the distribution $d\nu(x) = x^{-1}d\mu(x^{-1})$.

2.1. The Gaussian case

Consider an $m \times n$ matrix S with i.i.d. entries $\mathcal{N}(0, m^{-1})$. The l.s.d. of C_S is the Marchenko-Pastur law with density μ_ρ given in (4). Denote by $a = (1 - \sqrt{\rho})^2$ and $b = (1 + \sqrt{\rho})^2$ the edge eigenvalues. Let $\{\Delta_t\}$ be the sequence of error vectors generated by a first-order method as in (2). According to Lemma 2.1, there exists a sequence of polynomials $p_t \in \mathbb{R}_t^0[X]$ such that $\Delta_t = p_t(C_S^{-1}) \Delta_0$, and

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}\|\Delta_t\|^2}{\mathbb{E}\|\Delta_0\|^2} = \int_a^b p_t^2(\lambda^{-1})\mu_\rho(\lambda)d\lambda, \quad (11)$$

Lemma 2.3. *Under the above assumptions and notations, and setting $P_t(x) = p_t\left(\frac{x}{(1-\rho)^2}\right)$, we have*

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}\|\Delta_t\|^2}{\mathbb{E}\|\Delta_0\|^2} = (1 - \rho) \int_a^b P_t^2(x) \frac{1}{x} \mu_\rho(x) dx. \quad (12)$$

Consequently, if $\{\Pi_t\}$ is an orthogonal basis of polynomials with respect to μ_ρ such that $\deg(\Pi_t) = t$ and $\Pi_t(0) = 1$ then $\bar{\Pi}_t(x) := \Pi_t((1 - \rho)^2 x)$ achieves the lower bound $\mathcal{L}_{\mu_\rho,t}^$.*

Proof. Using the change of variable $x = (1 - \rho)^2/\lambda$, a simple calculation yields that $p_t^2(\lambda^{-1})\mu_\rho(\lambda)d\lambda = (1 - \rho)P_t^2(x)\frac{1}{x}\mu_\rho(x)dx$. Applying Lemma 2.2 with $\nu = \mu_\rho$, we get that the optimal polynomial P_t is equal to $\bar{\Pi}_t$, and thus, p_t is exactly $\bar{\Pi}_t(x)$. \square

The Marchenko-Pastur law μ_ρ is well-studied, and such a construction of polynomials is classical. In this section, we provide a definition by recursion, which is enough to state the optimal algorithm. However, for the proof of the next results, we will consider an alternative construction, from which we establish several intermediate properties useful to the analysis. Define $\Pi_0(x) = 1$, $\Pi_1(x) = 1 - x$, and for $t \geq 2$,

$$\Pi_t(x) = (1 + \rho - x)\Pi_{t-1}(x) - \rho\Pi_{t-2}(x). \quad (13)$$

Lemma 2.4. *The family of polynomials $\{\Pi_t\}$ is orthogonal with respect to μ_ρ . Further, we have $\Pi_t(0) = 1$ and $\deg(\Pi_t) = t$ for all $t \geq 0$.*

Proof. We defer the proof to Section B.4. \square

Now, set $\bar{\Pi}_t(x) = \Pi_t((1-\rho)^2x)$. From (13), we obtain that $\bar{\Pi}_0(x) = 1$, $\bar{\Pi}_1(x) = 1 - (1-\rho)^2x$, and for $t \geq 2$,

$$\bar{\Pi}_t(x) = (1+\rho-(1-\rho)^2x)\bar{\Pi}_{t-1}(x) - \rho\bar{\Pi}_{t-2}(x). \quad (14)$$

According to Lemma 2.3, the polynomial $\bar{\Pi}_t$ achieves the lower bound $\mathcal{L}_{\mu,\rho,t}^*$. Further, we identify the recursion formula (14) with the three-terms recursion (8) by setting $b_t = -(1-\rho)^2$ for $t \geq 1$, and $a_t = 1+\rho$ for $t \geq 2$. Using Theorem 1, we immediately have the asymptotically optimal first-order method, which we present in Algorithm 1 in its finite-sample approximation.

Algorithm 1 Optimal First-Order Method for Gaussian embeddings.

Input: Data matrix $A \in \mathbb{R}^{n \times d}$, sketch size $m \geq d+1$, initial point $x_0 \in \mathbb{R}^d$ and (finite-sample) ratio $\rho := d/m$. Sample $S \in \mathbb{R}^{m \times n}$ with i.i.d. entries $\mathcal{N}(0, 1/m)$.

Compute the sketched matrix $S_A = S \cdot A$.

Compute and cache a factorization of $H_S = S_A^\top S_A$.

Set $x_1 = x_0 - (1-\rho)^2 H_S^{-1} \cdot A^\top (Ax_0 - b)$.

for $t = 2$ **to** T **do**

 Compute the gradient $g_{t-1} = A^\top (Ax_{t-1} - b)$.

 Perform the update

$$x_t = x_{t-1} + \rho(x_{t-1} - x_{t-2}) - (1-\rho)^2 \cdot H_S^{-1} g_t. \quad (15)$$

end for

Return the last iterate x_T .

Surprisingly, up to the initialization of the first iterate x_1 , Algorithm 1 corresponds exactly to the Heavy-ball method (3) using the fixed step size $\mu = (1-\rho)^2$ and the fixed momentum parameter $\beta = \rho$, which was obtained in [22, 13] based on edge eigenvalues analysis. Hence, in the Gaussian case, leveraging the whole shape of the limiting distribution, as opposed to using only the edge eigenvalues, yields the same algorithm. We complete the analysis of the Gaussian case by providing the exact asymptotic error $\mathcal{L}_{\mu,\rho,t}^*$.

Theorem 2. *The sequence of iterates $\{x_t\}$ given by Algorithm 1 is asymptotically optimal within the class of first order algorithms as in (2), and the optimal error is given by $\mathcal{L}_{\mu,\rho,t}^* = \rho^t$.*

Proof. We have already argued that $\{x_t\}$ is asymptotically optimal. It remains to show that $\mathcal{L}_{\mu,\rho,t}^* = \rho^t$, whose proof is deferred to Appendix A.2. \square

2.2. The SRHT case

Haar random projections have been shown to have a better performance than Gaussian embeddings in several contexts.

However, they are slow to generate and apply, and we consider instead the SRHT. We recall the definition of the Stieltjes transform m_μ of a distribution μ supported on $[0, +\infty)$, which, for $z \in \mathbb{C} \setminus \mathbb{R}_+$, is given by $m_\mu(z) := \int_{\mathbb{R}} \frac{1}{x-z} d\mu(x)$. It has been recently shown that the SRHT behaves asymptotically as Haar embeddings, as formally stated by the next result.

Lemma 2.5 (Theorem 4.1 in [15]). *Let S be an $m \times n$ SRHT embedding and S_h be an $m \times n$ Haar embedding. Then, the matrices C_S and C_{S_h} have the same limiting spectral distribution F_h , with support included within the interval $(0, 1)$ and whose Stieltjes transform m_h is given by*

$$m_h(z) = \frac{1}{2\gamma} \left(\frac{2\gamma-1}{1-z} + \frac{\xi-\gamma}{z(1-z)} - \frac{R(z)}{z(1-z)} \right), \quad (16)$$

where

$$R(z) = \sqrt{(\gamma + \xi - 2 + z)^2 + 4(z-1)(1-\gamma)(1-\xi)},$$

Remark 1. *Due to the computational benefits of the SRHT over Haar projections, we state all our next results for the former, although all statements also apply to the latter (except for the time complexity results).*

In order to characterize the optimal first-method with SRHT embeddings, we first derive the density of F_h .

Theorem 3. *The distribution F_h admits the following density on \mathbb{R} ,*

$$f_h(x) = \frac{1}{2\gamma\pi} \frac{\sqrt{(\Lambda_h - x)_+(x - \lambda_h)_+}}{x(1-x)}, \quad (17)$$

where

$$\begin{cases} \lambda_h := \left(\sqrt{(1-\gamma)\xi} - \sqrt{(1-\xi)\gamma} \right)^2, \\ \Lambda_h := \left(\sqrt{(1-\gamma)\xi} + \sqrt{(1-\xi)\gamma} \right)^2. \end{cases}$$

Proof. The proof is essentially based on the expression (16) of the Stieltjes transform m_h , and on the inversion formula,

$$f_h(x) = \lim_{y \rightarrow 0^+} \frac{1}{\pi} \text{Im} (m_h(x + iy)), \quad \text{where } y \in \mathbb{R}_+. \quad (18)$$

which holds for any $x \in \mathbb{R}$ provided that the above limit exists [30]. We defer the calculations to Appendix A.3. \square

Using the change of variable $y = x/\xi$, we can also derive the limiting density of the *rescaled* matrix $\frac{n}{m}C_S$ – whose expectation is equal to the identity – which is given by

$$f_{h,r}(y) = \xi f_h(\xi y) = \frac{\sqrt{(\Lambda_{h,r} - y)_+(y - \lambda_{h,r})_+}}{2\rho\pi y(1-\xi y)}, \quad (19)$$

where

$$\begin{cases} \lambda_{h,r} = \lambda_h/\xi = \left(\sqrt{1-\gamma} - \sqrt{(1-\xi)\rho} \right)^2, \\ \Lambda_{h,r} = \Lambda_h/\xi = \left(\sqrt{1-\gamma} + \sqrt{(1-\xi)\rho} \right)^2. \end{cases}$$

The density $f_{h,r}$ resembles the Marchenko-Pastur density μ_ρ , up to the factor $(1-\xi y)$ and corrections in the edge eigenvalues $\lambda_{h,r}$ and $\Lambda_{h,r}$. When $\xi, \gamma \approx 0$, then $\lambda_{h,r} \approx (1-\sqrt{\rho})^2$, $\Lambda_{h,r} \approx (1+\sqrt{\rho})^2$, and $f_{h,r}(x) \approx \mu_\rho(x)$. This is consistent with the fact that provided $m, d = o(n)$ so that $\xi, \gamma = 0$, then the l.s.d. of $\frac{n}{m}C_S$ is the Marchenko-Pastur law with parameter ρ (see [12] for a formal statement). In Figure 1, we compare the empirical spectral density of the matrix $\frac{n}{m}C_S$ with S an $m \times n$ SRHT to $f_{h,r}$, for fixed d and n , and several values of m . We observe that these two densities match very closely, and so does the empirical spectral density using a Haar projection with $f_{h,r}$. Further, as m increases, the limiting density $f_{h,r}$ departs from μ_ρ , and then concentrates more and more around 1. Note in particular that the support of $f_{h,r}$ is always within that of μ_ρ . This can be formally verified by comparing their respective edge eigenvalues.




Figure 1. We use $n = 8192$, $\gamma \approx \frac{d}{n} = 0.2$ and $\xi \approx \frac{m}{n} \in \{0.21, 0.4, 0.6\}$.

2.2.1. ORTHOGONAL POLYNOMIALS AND OPTIMAL FIRST-ORDER METHOD

Given a first-order method as in (2), we know from Lemma 2.1 that for a given iteration t , there exists a polynomial $p \in \mathbb{R}_t^0[X]$ such that $\Delta_t = p(C_S^{-1}) \Delta_0$, and

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} \|\Delta_t\|^2}{\mathbb{E} \|\Delta_0\|^2} = \int_{\lambda_h}^{\Lambda_h} p^2(\lambda^{-1}) f_h(\lambda) d\lambda. \quad (20)$$

Introducing the scaling parameters $\tau = \left(\frac{\sqrt{\Lambda_h} - \sqrt{\lambda_h}}{\sqrt{\Lambda_h} + \sqrt{\lambda_h}} \right)^2$, $c = \frac{4}{(\sqrt{1/\Lambda_h} + \sqrt{1/\lambda_h})^2}$, $\alpha = (1 - \sqrt{\tau})^2$, $\beta = (1 + \sqrt{\tau})^2$, the rescaled polynomial $P(x) = p(x/c)$, and using the change of variable $x = c/\lambda$, we find that

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} \|\Delta_t\|^2}{\mathbb{E} \|\Delta_0\|^2} \quad (21)$$

$$= \frac{c\tau}{(1-\tau)\gamma} \int_{\alpha}^{\beta} P^2(x) \frac{\sqrt{(x-\alpha)(\beta-x)}}{2\pi\tau x(x-c)} dx \quad (22)$$

$$= \frac{c\tau}{(1-\tau)\gamma} \int_{\alpha}^{\beta} P^2(x) \frac{\mu_\tau(x)}{x-c} dx \quad (23)$$

Thus, according to Lemma 2.2, it suffices to find a family of polynomials $\{R_t\}$ orthogonal with respect to the density $\frac{x\mu_\tau(x)}{x-c}$ such that $\deg(R_t) = t$ and $R_t(0) = 1$, in which case the minimizer over $P \in \mathbb{R}_t^0[X]$ of the integral in (23) is equal to R_t , and the minimizer of (20) is then $\bar{R}_t(x) = R_t(cx)$.

Theorem 4. Define the parameters $\omega = \frac{4}{(\sqrt{\beta-c} + \sqrt{\alpha-c})^2}$ and $\kappa = \left(\frac{\sqrt{\beta-c} - \sqrt{\alpha-c}}{\sqrt{\beta-c} + \sqrt{\alpha-c}} \right)^2$. Let $\{\Pi_t\}$ be the orthogonal family of polynomials with respect to μ_κ , that is, $\Pi_0(x) = 1$, $\Pi_1(x) = 1 - x$, and for $t \geq 2$,

$$\Pi_t(x) = (1 + \kappa - x)\Pi_{t-1}(x) - \kappa\Pi_{t-2}(x). \quad (24)$$

Define the polynomials $R_t(x) = \Pi_t(\omega(x-c))/\Pi_t(-\omega c)$. Then, it holds that $R_t(0) = 1$, $\deg(R_t) = t$, and the family $\{R_t\}$ is orthogonal with respect to the density $\frac{x\mu_\tau(x)}{x-c}$.

Proof. For $k \neq \ell$, we have that $\int_{\mathbb{R}} R_k(x)R_\ell(x) \frac{x\mu_\tau(x)}{x-c} dx \propto \int_{\alpha}^{\beta} \Pi_k(\omega(x-c)) \Pi_\ell(\omega(x-c)) \frac{\sqrt{(\beta-x)(x-\alpha)}}{2\pi\rho(x-c)} dx$. Using the change of variable $y = \omega(x-c)$, we find that the latter integral is (up to a constant) equal to $\int_{\mathbb{R}} \Pi_k(y)\Pi_\ell(y)\mu_\kappa(y) dy$, which is itself equal to 0 due to the orthogonality of the Π_t with respect to μ_κ . \square

In order to derive the optimal first-order method, we need to find the three-terms recursion relationship satisfied by the polynomials $\{\bar{R}_t\}$. First, let us compute the normalization factor $u_t := \Pi_t(-\omega c)$. Evaluating (24) at $x = -\omega c$ and denoting $\eta := 1 + \kappa + \omega c$, we find that $u_{t+1} = \eta u_t - \kappa u_{t-1}$, with the initial conditions $u_0 = 1$ and $u_1 = \Pi_1(-\omega c) = 1 + \omega c = \eta - \kappa$. Thus, after solving this second-order linear system, we obtain that

$$u_t = \frac{x_1 - \kappa}{x_1 - x_2} x_1^t + \frac{\kappa - x_2}{x_1 - x_2} x_2^t, \quad (25)$$

where $x_1 = \frac{\eta}{2} + \sqrt{\frac{\eta^2}{4} - \kappa}$ and $x_2 = \frac{\eta}{2} - \sqrt{\frac{\eta^2}{4} - \kappa}$. It is easy to check that $\eta^2/4 > \kappa$, so that x_1 and x_2 are indeed distinct and real. Then, using the change of variable $y = \omega(x-c)$ in (24), we get the following three-terms recurrence relationship, that is, $\bar{R}_0(x) = 1$, $\bar{R}_1(x) = 1 + b_{h,1}x$ and for $k \geq 2$,

$$\bar{R}_t(x) = (a_{h,t} + xb_{h,t})\bar{R}_{t-1}(x) + (1 - a_{h,t})\bar{R}_{t-2}(x), \quad (26)$$

where $a_{h,t} = \frac{\eta u_{t-1}}{u_t}$ for $t \geq 1$, and $b_{h,t} = -\frac{\omega c u_{t-1}}{u_t}$ for $t \geq 2$. Using Theorem 1, we obtain the optimal first-order method, which we present in Algorithm 2 in its finite-sample approximation.

Algorithm 2 Optimal First-Order Method for SRHT (or Haar) embeddings.

Input: Data matrix $A \in \mathbb{R}^{n \times d}$, sketch size $m \geq d + 1$, initial point $x_0 \in \mathbb{R}^d$.

Sample an $m \times n$ SRHT S .

Compute the sketched matrix $S_A = S \cdot A$.

Compute and cache a factorization of $H_S = S_A^\top S_A$.

Set $x_1 = x_0 + b_{h,1} H_S^{-1} A^\top (Ax_0 - b)$.

for $t = 2$ **to** T **do**

 Compute the gradient $g_{t-1} = A^\top (Ax_{t-1} - b)$.

 Perform the update

$$x_t = x_{t-1} + b_{h,t} H_S^{-1} g_t + (1 - a_{h,t})(x_{t-2} - x_{t-1}). \quad (27)$$

 where $a_{h,t}$ and $b_{h,t}$ are as described in Section 2.2.1.

end for

Return the last iterate x_T .

Differently from the Gaussian case, Algorithm 2 does not correspond to the Heavy-ball method (3) using the fixed step size $\mu = (1 - \rho)^2$ and the fixed momentum parameter $\beta = \rho$, which was obtained by [13] based on edge eigenvalues analysis and standard finite-sample concentration bounds on the spectrum of SRHT matrices [31].

Using the new asymptotically exact extreme eigenvalues we derived in Theorem 3 – which are different from the bounds obtained by [31] – and following the same extreme eigenvalues analysis proposed by [13], we can derive an optimal Heavy-ball method for which the step size μ_h and momentum parameter β_h are given by $\mu_h = \frac{4}{\left(\frac{1}{\sqrt{\lambda_h}} + \frac{1}{\sqrt{\lambda_h}}\right)^2}$

$$\text{and } \beta_h = \left(\frac{\sqrt{\lambda_h} - \sqrt{\lambda_h}}{\sqrt{\lambda_h} + \sqrt{\lambda_h}}\right)^2.$$

Hence, leveraging the whole shape of the limiting distribution, as opposed to using only the edge eigenvalues, yields an optimal first-order method which is different, and has non-constant step sizes and momentum parameters. But interestingly, it holds that as the iteration number t grows to $+\infty$, then the update coefficients $a_{h,t}$ and $b_{h,t}$ have respective limits $1 + \beta_h$ and $-\mu_h$, which yields exactly this Heavy-ball method. Thus, we expect the latter and Algorithm 2 to have a similar performance as t grows large.

We complete our analysis of the SRHT case by characterizing the asymptotic error $\mathcal{L}_{f_h,t}^*$.

Theorem 5. *The sequence of iterates $\{x_t\}$ given by Algorithm 2 is asymptotically optimal, and the optimal error satisfies $\mathcal{L}_{f_h,t}^* \asymp \frac{(1-\xi)^t}{(1-\gamma)^t} \rho^t$.*

Proof. We have already argued that $\{x_t\}$ is asymptotically optimal. It remains to show that $\mathcal{L}_{f_h,t}^* \asymp \frac{(1-\xi)^t}{(1-\gamma)^t} \rho^t$, whose

proof is deferred to Appendix A.4. \square

Of natural interest is to compare the rate of convergence $\rho_h := \frac{(1-\xi)}{(1-\gamma)} \rho$ of Algorithm 2 to the rate ρ of Algorithm 1. We have $\frac{\rho_h}{\rho} = \frac{(1-\xi)}{(1-\gamma)}$, which is always smaller than 1 since $\xi > \gamma$. Hence, these rotation matrices yield an optimal first-order method which is uniformly better than that with Gaussian embeddings, by a factor which can be made arbitrarily large by increasing the sketch size m relatively to the other dimensions. Further, if we do not reduce the size of the original matrix, so that $m = n$ and $\xi = 1$, then the algorithm converges in one iteration. This means that we do not lose any information by sketching. In contrast, Gaussian projections introduce more distortions than rotations, even though the rows of a Gaussian matrix are almost orthogonal to each other in the high-dimensional setting.

Further, we compare the rate of Algorithm 2 to the rate of the best Heavy-ball method with refreshed SRHT embeddings which is equal to $\rho_h^{\text{ref}} = \rho \cdot \frac{\xi(1-\xi)}{\gamma^2 + \xi - 2\xi\gamma}$. We have $\rho_h < \rho_h^{\text{ref}}$ if and only if $\frac{1-\xi}{1-\gamma} < \frac{\xi(1-\xi)}{\gamma^2 + \xi - 2\xi\gamma}$, which is equivalent to $\gamma^2 + \xi - 2\gamma\xi < \xi - \gamma\xi$, again equivalent to $\gamma^2 < \gamma\xi$, i.e., $\gamma < \xi$, which holds by assumption. Thus, a fixed embedding yields a first-order method which is uniformly faster than the best Heavy-ball method with refreshed sketches. However, it remains an open problem whether one can find a first-order method with refreshed sketches which yields a rate better than ρ_h^{ref} . We recapitulate the different convergence rates in Table 1.

Table 1. Asymptotic rates of convergence for the best first-order method (2) and the best Heavy-ball method (3), with fixed or refreshed Gaussian or SRHT embeddings. For the best Heavy-ball method rates, we use previously derived results from [22, 13, 15].

Algorithm	Fixed Gaussian	Refreshed Gaussian	Fixed SRHT	Refreshed SRHT
Best first-order method (2)	ρ	unknown	$\frac{1-\xi}{1-\gamma} \rho$	unknown
Best Heavy-ball method (3)	ρ	ρ	ρ	$\frac{\xi(1-\xi)}{\gamma^2 + \xi - 2\xi\gamma}$

In Figure 2, we verify numerically that Algorithm 2 is faster than the best Heavy-ball method with refreshed SRHT sketches (“SRHT (refreshed)”), and than Algorithm 1. Further, we compare Algorithm 2 to the Heavy-ball method with fixed SRHT embedding whose parameters are found based on edge eigenvalues analysis, using either our new density f_h (“SRHT (edge eig.)”) – as described previously in Section 2.2.1 –, or, the previous bounds derived by [31] (“SRHT (baseline)”). As predicted, Algorithm 2 performs very similarly to the former, and better than the latter. Finally, we verify that our predicted convergence rates for Algorithms 1 and 2 are matched empirically, on Figure 3.

We mention that we use small perturbations of the algorithmic parameters derived from our asymptotic analysis. Following the notations introduced in Theorem 1, instead of a_t and b_t , we use $a_t^\delta = (1 + \delta)a_t$ and $b_t^\delta = (1 - \delta)b_t$ with $\delta = 0.01$. These conservative perturbations are necessary in practice due to the finite-sample approximations. We defer a detailed description of the experimental setup to Appendix C.




Figure 2. Error $\mathbb{E}\|\Delta_t\|^2/\mathbb{E}\|\Delta_0\|^2$ versus number of iterations. We use $n = 8192$, $d/n \approx \gamma = 0.2$ and $m/n \approx \xi \in \{0.22, 0.4, 0.7\}$.




Figure 3. Empirical and theoretical convergence rates versus sketch size m . We use $n = 8192$ and $d \in \{500, 1250, 2000\}$.

2.2.2. COMPLEXITY ANALYSIS

We turn to a complexity analysis of Algorithm 2 and compare it to the currently best known algorithmic complexities for solving (1).

Given a fixed (and independent of the dimensions) error $\varepsilon > 0$, we aim to find \tilde{x} such that $\|A(\tilde{x} - x^*)\|^2 \leq \varepsilon$. Among the best complexity algorithms is the pre-conditioned conjugate gradient algorithm [27]. As described in Section 1, it is decomposed into three parts: sketching the data matrix, factoring the pre-conditioned matrix, and then the iterations of the conjugate gradient method. This algorithm prescribes at least the sketch size $m \asymp d \log d$ in order to converge with high-probability guarantees. This theoretical prescription is based on the finite-sample bounds on the extremal eigenvalues of the matrix C_S derived by [31]. Then, the resulting complexity scales as

$$C_{\text{cg}} \asymp nd \log d + d^3 \log d + nd \log(1/\varepsilon), \quad (28)$$

where $nd \log d$ is the sketching cost, $d^3 \log d$ the pre-conditioning cost, and $nd \log(1/\varepsilon)$ is the per-iteration cost nd times the number of iterations $\log 1/\varepsilon$.

Our analysis shows that for $m \asymp d$, Algorithm 2 yields a

complexity no larger than

$$C_{\text{fhs}} \asymp nd \log d + d^3 + nd \log(1/\varepsilon), \quad (29)$$

Note that in the above complexity, we omit the rate of convergence – which would yield an even smaller complexity – to simplify the comparison. Since ε is independent of the dimensions, it follows that

$$\frac{C_{\text{fhs}}}{C_{\text{cg}}} \asymp \frac{1}{\log d}, \quad d \rightarrow \infty. \quad (30)$$

Hence, with a smaller sketch size, the resulting complexity improves by a factor $\log d$ over the current state-of-the-art in randomized preconditioning for dense problems (e.g., see [6, 20]). We also note that the $O(d^3)$ term can be improved to $O(d^\omega)$, where ω is the exponent of matrix multiplication.

It has also been shown by [15] that the Heavy-ball update (3) with refreshed SRHT embeddings yields a complexity C_{ihs} such that $C_{\text{ihs}}/C_{\text{cg}} \asymp \frac{1}{\log d}$, provided that $m \asymp d$. In order to compare more finely Algorithm 2 with this algorithm, we consider an arbitrary sketch size m . Then, the complexity of Algorithm 2 is

$$C_{\text{fhs}} \asymp nd \log m + md^2 + nd \frac{\log(1/\varepsilon)}{\log \rho_h}, \quad (31)$$

whereas the former algorithm yields

$$C_{\text{ihs}} \asymp (nd \log m + md^2 + nd) \frac{\log(1/\varepsilon)}{\log \rho_h^{\text{ref}}}. \quad (32)$$

Since, in particular, ρ_h is uniformly smaller than ρ_h^{ref} , it always holds that

$$C_{\text{fhs}} \leq C_{\text{ihs}}. \quad (33)$$

It should be noted that we translate our asymptotic results to finite-sample versions. Although it is beyond our scope, we believe that our results could be extended to finite-sample versions with high-probability guarantees and with similar rates of convergence.

Acknowledgements

The authors thank Sifan Liu and Edgar Dobriban for helpful discussions. This work was partially supported by the National Science Foundation under grant IIS-1838179.

References

- [1] N. Ailon and B. Chazelle. Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform. *Proceedings of the thirty-eighth annual ACM symposium on Theory of Computing*, 557-563, 2006.

- [2] G. Anderson and B. Farrell. Asymptotically liberating sequences of random unitary matrices. *Advances in Mathematics*, 255:381-413, 2014.
- [3] G. Anderson, A. Guionnet and O. Zeitouni. An introduction to random matrices. *Cambridge University Press*, 2010.
- [4] H. Avron, P. Maymounkov and S. Toledo. Blendepik: Supercharging LAPACK's least-squares solver. *SIAM Journal on Scientific Computing*, 32(3):1217-1236, 2014.
- [5] Z. Bai and J. Silverstein. Spectral analysis of large dimensional random matrices. *Springer series in Statistics*, 2010.
- [6] C. Boutsidis and A. Gittens. Improved matrix algorithms via the subsampled randomized Hadamard transform. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1301-1340, 2013.
- [7] E. Dobriban and Sifan Liu. Asymptotics for sketching in least-squares regression. *Advances in Neural Information Processing Systems*, 3670-3680, 2019.
- [8] P. Drineas and M. Mahoney. RandNLA: randomized numerical linear algebra. *Communications of the ACM*, 59(6):80-90, 2016.
- [9] P. Drineas, M. Mahoney, S. Muthukrishnan and T. Sarlos. Faster least squares approximation. *Journal of Numerical Mathematics*, 1117(2):219-249, 2011.
- [10] N. Halko, P. Martinsson and J. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217-288, 2011.
- [11] F. Hiai and D. Petz. The semicircle law, free random variables and entropy. *Americal Mathematical Society*, 2006.
- [12] T. Jiang. Approximation of Haar distributed matrices and limiting distributions of eigenvalues of Jacobi ensembles. *Probability Theory and related fields*, 144(1-2):221-246, 2009.
- [13] J. Lacotte and M. Pilanci. Faster least-squares optimization. *arXiv:1911.02675*, 2019.
- [14] J. Lacotte, M. Pilanci and M. Pavone. High-dimensional optimization in adaptive random subspaces. *Advances in Neural Information Processing Systems*, 10846-10856, 2019.
- [15] J. Lacotte, S. Liu, E. Dobriban and M. Pilanci. Limiting spectrum of randomized Hadamard transform and optimal iterative sketching methods. *arXiv:2002.00864*, 2020.
- [16] S. Liu and E. Dobriban. Ridge regression: Structure, crossvalidation, and sketching. *arXiv:1910.02373*, 2019.
- [17] M. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2), 2011.
- [18] M. Mahoney and P. Drineas. Structural properties underlying high-quality randomized numerical linear algebra algorithms. *Handbook of Big Data*, 137-154, 2016.
- [19] V.A. Marcenko and L.A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Annals of Probability*, 4(1):457-483, 1967.
- [20] J. Nelson and H. Nguyen. OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings. *IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 117-126. IEEE, 2013.
- [21] A. Nica and R. Speicher. Lectures on the combinatorics of free probability. *Cambridge University Press*, 2006.
- [22] I. Ozaslan, M. Pilanci and O. Arikan. Iterative hessian sketch with momentum. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7470-7474, 2019.
- [23] D. Paul and A. Aue. Random matrix theory in statistics: a review. *Journal of Statistical Planning and Inference*, 150:1-29, 2014.
- [24] F. Pedregosa and D. Scieur. Acceleration through spectral modeling. *NeurIPS workshop Beyond First Order Methods in ML*, 2019.
- [25] M. Pilanci and M. Wainwright. Randomized sketches of convex programs with sharp guarantees. *IEEE Transactions on Information Theory*, 61.9:5096-5115. IEEE, 2015.
- [26] M. Pilanci and M. Wainwright. Iterative Hessian sketch: Fast and accurate solution approximation for constrained least-squares. *IEEE Transactions on Information Theory*, 9(61):5096-5115, 2015.
- [27] V. Rokhlin and M. Tygert. A fast randomized algorithm for overdetermined linear least-squares regression. *Proceedings of the National Academy of Sciences*, 105(36):13212-13217, 2008.
- [28] H. Rutishauser. Theory of gradient methods. *Refined iterative methods for computation of the solution and the eigenvalues of self-adjoint boundary value problems*, pages 24-29, 1959.

- [29] Y. Saad. Iterative methods for sparse linear systems. *SIAM*, volume 82, 2003.
- [30] J. Silverstein and S. Choi. Analysis of the limiting spectral distribution of large dimensional random matrices. *Journal on Multivariate Analysis*, 54(2):295-309, 1995.
- [31] J. Tropp. Improved analysis of the subsampled randomized Hadamard transform. *Advances in Adaptive Data Analysis*, 3:115-126, 2011.
- [32] D. Voiculescu, K. Dykema and A. Nica. Free random variables. *American mathematical society*, 1992.
- [33] R. Witten and E. Candes. Randomized algorithms for low-rank matrix factorizations: sharp performance bounds. *Algorithmica*, 72(1):264-281, 2015.
- [34] Y. Yang, M. Pilanci and M. Wainwright. Randomized sketches for kernels: Fast and optimal nonparametric regression. *The Annals of Statistics*, 45(3):991-1023, 2017.
- [35] J. Yao, Z. Bai and S. Zheng. Large sample covariance matrices and high-dimensional data analysis. *Cambridge University Press*, 2015.

A. Proof of main results

For a polynomial P and a measure (resp. density) μ , we will denote $\mu[P] := \int_{\mathbb{R}} P(x)\mu(x)d\mu(x)$ (resp. $\mu[P] := \int_{\mathbb{R}} P(x)\mu(x)dx$). For a density μ , we stress the fact that $\mu[x]$ and $\mu(x)$ refer to different quantities.

A.1. Proof of Theorem 1

We recall that $\Pi_0(x) = 1$, $\Pi_1(x) = 1 + b_1x$ and for $t \geq 2$,

$$\Pi_t(x) = (a_t + b_t x) \Pi_{t-1}(x) + (1 - a_t) \Pi_{t-2}(x). \quad (34)$$

First, we claim that for any $t \geq 0$, $\Delta_t = \Pi_t(C_S^{-1}) \Delta_0$, and we show it by induction. Since $\Pi_0(x) = 1$, we have that $\Delta_0 = \Pi_0(C_S^{-1}) \cdot \Delta_0$. Since $x_1 = x_0 + b_1 H_S^{-1} \nabla f(x_0)$, subtracting x^* and multiplying by $U^\top A$ the latter equation, we obtain that $\Delta_1 = \Pi_1(C_S^{-1}) \cdot \Delta_0$. Suppose that for some $t \geq 2$, the induction claim holds for $t-1$ and $t-2$. Subtracting x^* and multiplying by $U^\top A$ the update formula (9), we obtain that

$$\begin{aligned} \Delta_t &= \Delta_{t-1} + (1 - a_t)(\Delta_{t-2} - \Delta_{t-1}) + b_t C_S^{-1} \Delta_{t-1} \\ &= (a_t + b_t C_S^{-1}) \Delta_{t-1} + (1 - a_t) \Delta_{t-2} \\ &= ((a_t + b_t C_S^{-1}) \Pi_{t-1}(C_S^{-1}) + (1 - a_t) \Pi_{t-2}(C_S^{-1})) \Delta_0 \\ &= \Pi_t(C_S^{-1}) \Delta_0, \end{aligned}$$

where we used the induction hypothesis for $t-1$ and $t-2$ in the third equality, and the recursion formula (34) in the last equality. Consequently, using Lemma 2.1, we obtain that

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} \|\Delta_t\|^2}{\mathbb{E} \|\Delta_0\|^2} = \int_a^b \Pi_t^2(\lambda^{-1}) d\mu(\lambda) = \mathcal{L}_{\mu, t}^*.$$

A.2. Proof of Theorem 2

We have already argued that $\{x_t\}$ is asymptotically optimal. It remains to prove that $\mathcal{L}_{\mu_\rho, t}^* = \rho^t$.

Set $\lambda_\rho(x) = x^{-1} \mu_\rho(x)$. Let $\{\Pi_t\}$ be an orthogonal basis with respect to μ_ρ such that $\Pi_t(0) = 1$ and $\deg(\Pi_t) = t$. From Lemma 2.3, we have $\mathcal{L}_{\mu_\rho, t}^* = (1 - \rho) \lambda_\rho[\Pi_t^2]$, so that it suffices to show that $\lambda_\rho[\Pi_t^2] = (1 - \rho)^{-1} \rho^t$. On the other hand, in the proof of Lemma 2.4 in Appendix B.4, we establish that there exists a sequence of polynomials $\{T_k\}_{k \geq 1}$ such that for any $t \geq 1$ and $k, \ell \geq 1$,

$$\begin{aligned} \Pi_t(x) &= 1 - \sum_{j=1}^t \lambda_\rho[T_j] T_j(x), \\ \lambda_\rho[T_t] &= (-1)^{t-1} \sqrt{\rho^{t-1}}, \\ \lambda_\rho[T_k T_\ell] &= \delta_{k\ell}, \end{aligned}$$

where $\delta_{k\ell} = 1$ if $k = \ell$, and 0 otherwise. Using the latter properties, it follows that

$$\begin{aligned} \lambda_\rho[\Pi_t^2] &= \lambda_\rho[1] - 2 \sum_{j=1}^t \lambda_\rho[T_j]^2 + \sum_{j=1}^t \lambda_\rho[T_j]^2 \underbrace{\lambda_\rho[T_j^2]}_{=1} \\ &= \lambda_\rho[1] - \sum_{j=1}^t \lambda_\rho[T_j]^2 \\ &= \frac{1}{1 - \rho} - \sum_{j=0}^{t-1} \rho^j \\ &= \frac{\rho^t}{1 - \rho}, \end{aligned}$$

and, in the third equality, we used the standard inverse moment formula $\lambda_\rho[1] = \int x^{-1} \mu_\rho(x) dx = (1 - \rho)^{-1}$. Consequently, we obtain the claimed formula, that is, $\mathcal{L}_{\mu_\rho, t}^* = \rho^t$.

A.3. Proof of Theorem 3

According to Lemma 2.5, the support of F_h is included within the interval $(0, 1)$. Therefore, we fix $x \in (0, 1)$ and we consider the complex number $z = x + iy$, where $y > 0$. Our goal is to compute the quantity

$$\lim_{y \rightarrow 0^+} \frac{1}{\pi} |\operatorname{Im}(m_h(z))|.$$

If the above limit exists, then F_h is differentiable at x and its derivative is equal to this limit [30]. Note that the absolute value is not necessary, since $\operatorname{Im}(m_h(z))$ is positive on \mathbb{C}^+ . But it will avoid to specify explicitly the branch cut of the square-root considered later in this proof, and thus additional technicalities.

From Lemma 2.5, it holds that

$$2\gamma m_h(z) = \frac{2\gamma - 1}{1 - z} + \frac{\xi - \gamma}{z(1 - z)} - \frac{R(z)}{z(1 - z)}. \quad (35)$$

where $R(z) = \sqrt{(\gamma + \xi - 2 + z)^2 + 4(z - 1)(1 - \gamma)(1 - \xi)}$, and the branch cut of the square-root is chosen such that $m_h > 0$ on \mathbb{C}^+ , $m_h < 0$ on \mathbb{C}^- (the complex numbers with negative imaginary parts), and $m_h > 0$ on \mathbb{R}_- (the negative real numbers). Further, we have

$$\frac{1}{z(1 - z)} = \frac{x(1 - x) + y^2 + iy(2x - 1)}{(x(1 - x) + y^2)^2 + y^2(2x - 1)^2}, \quad \frac{1}{1 - z} = \frac{1 - x + iy}{(1 - x)^2 + y^2},$$

from which we deduce that the imaginary parts of the first two terms in the expansion (35) of $2\gamma m_h(z)$ are given by

$$\begin{aligned} \operatorname{Im}\left(\frac{2\gamma - 1}{1 - z}\right) &= \frac{(2\gamma - 1)y}{(1 - x)^2 + y^2}, \\ \operatorname{Im}\left(\frac{\xi - \gamma}{z(1 - z)}\right) &= \frac{(\xi - \gamma)(2x - 1)y}{(x(1 - x) + y^2)^2 + y^2(2x - 1)^2}. \end{aligned}$$

Since $x \in (0, 1)$, the limits $y \rightarrow 0^+$ of the two above quantities exist and are equal to 0. Hence, provided it exists, we have

$$\lim_{y \rightarrow 0^+} 2\gamma |\operatorname{Im}(m_h(z))| = \lim_{y \rightarrow 0^+} \left| \operatorname{Im}\left(\frac{R(z)}{z(1 - z)}\right) \right|. \quad (36)$$

We introduce the function $f(z) = (z - \alpha - \beta)^2 + 4(z - 1)\alpha\beta$ where $\alpha = 1 - \xi$ and $\beta = 1 - \gamma$, so that $R(z) = \sqrt{f(z)}$. We have $f(z) = X + iY$ where

$$\begin{aligned} X &= (x - \alpha - \beta)^2 - y^2 + 4(x - 1)\alpha\beta, \\ Y &= 2(x - \alpha - \beta + 2\alpha\beta)y. \end{aligned}$$

Thus, the absolute values of the real and imaginary parts of $R(z)$ are given by

$$\begin{aligned} |\operatorname{Re}(R(z))| &= \frac{1}{\sqrt{2}} \sqrt{\sqrt{X^2 + Y^2} + X}, \\ |\operatorname{Im}(R(z))| &= \frac{1}{\sqrt{2}} \sqrt{\sqrt{X^2 + Y^2} - X}, \end{aligned}$$

and they have respective limits

$$\begin{aligned} \lim_{y \rightarrow 0^+} |\operatorname{Re}(R(z))| &= \sqrt{|\varphi(x)|} \cdot \mathbf{1}(\varphi(x) > 0), \\ \lim_{y \rightarrow 0^+} |\operatorname{Im}(R(z))| &= \sqrt{|\varphi(x)|} \cdot \mathbf{1}(\varphi(x) < 0), \end{aligned}$$

where $\varphi(x) := (x - \alpha - \beta)^2 + 4(x - 1)\alpha\beta$. Further, we have

$$\operatorname{Im}\left(\frac{R(z)}{z(1 - z)}\right) = \frac{y(2x - 1)\operatorname{Re}(R(z)) + (x(1 - x) + y^2)\operatorname{Im}(R(z))}{g(x, y)},$$

where $g(x, y) = (x(1-x) + y^2)^2 + y^2(2x-1)^2$. Note that $\lim_{y \rightarrow 0^+} g(x, y) = x^2(1-x)^2$, which is non-zero since $x \in (0, 1)$. Then, we obtain

$$\lim_{y \rightarrow 0^+} \left| \operatorname{Im} \left(\frac{R(z)}{z(z-1)} \right) \right| = \frac{\sqrt{|\varphi(x)|} \mathbf{1}(\varphi(x) < 0)}{x(1-x)},$$

Using (36), it follows that for any $x \in (0, 1)$, $\lim_{y \rightarrow 0^+} \frac{1}{\pi} |\operatorname{Im}(m_h(x))|$ exists. This implies that F_h admits a density over $(0, 1)$, given by

$$f_h(x) = \frac{1}{2\gamma\pi} \frac{\sqrt{(\Lambda_h - x)_+(x - \lambda_h)_+}}{x(1-x)},$$

where we used the fact that $\varphi(x) = (x - \Lambda_h)(x - \lambda_h)$, and we recall that the edge eigenvalues Λ_h and λ_h are given by

$$\begin{aligned} \lambda_h &:= \left(\sqrt{(1-\gamma)\xi} - \sqrt{(1-\xi)\gamma} \right)^2 \\ \Lambda_h &:= \left(\sqrt{(1-\gamma)\xi} + \sqrt{(1-\xi)\gamma} \right)^2. \end{aligned}$$

Using the fact that F_h is supported within the interval $(0, 1)$, we have recovered the whole density of the limiting spectral distribution F_h of the matrix $U^\top S^\top S U$.

A.4. Proof of Theorem 5

We have already argued that $\{x_t\}$ is asymptotically optimal. It remains to show that $\mathcal{L}_{f_h, t}^* \asymp \frac{(1-\xi)^t}{(1-\gamma)^t} \rho^t$.

Using (23), we have

$$\begin{aligned} \mathcal{L}_{f_h, t}^* &= \frac{c\tau}{(1-\tau)\gamma} \min_{P \in \mathbb{R}_t^0[X]} \int_\alpha^\beta P^2(t) \frac{\mu_\tau(x)}{x-c} dx \\ &= \frac{c\tau}{(1-\tau)\gamma} \min_{P \in \mathbb{R}_t^0[X]} \int_\alpha^\beta P^2(t) \frac{x}{x-c} \frac{\mu_\tau(x)}{x} dx. \end{aligned}$$

For $x \in [\alpha, \beta]$, it holds that

$$\frac{\beta}{\beta-c} \leq \frac{x}{x-c} \leq \frac{\alpha}{\alpha-c},$$

and consequently, we can lower and upper bound $\mathcal{L}_{f_h, t}^*$ as follows,

$$\frac{c\tau}{(1-\tau)\gamma} \frac{\beta}{\beta-c} \min_{P \in \mathbb{R}_t^0[X]} \int_\alpha^\beta P^2(t) \frac{\mu_\tau(x)}{x} dx \leq \mathcal{L}_{f_h, t}^* \leq \frac{c\tau}{(1-\tau)\gamma} \frac{\alpha}{\alpha-c} \min_{P \in \mathbb{R}_t^0[X]} \int_\alpha^\beta P^2(t) \frac{\mu_\tau(x)}{x} dx.$$

From Lemma 2.3, we know that $\mathcal{L}_{\mu_\tau, t}^* = (1-\tau) \min_{P \in \mathbb{R}_t^0[X]} \int_\alpha^\beta P^2(t) \frac{\mu_\tau(x)}{x} dx$. Thus,

$$\frac{c\tau}{(1-\tau)^2\gamma} \frac{\beta}{\beta-c} \mathcal{L}_{\mu_\tau, t}^* \leq \mathcal{L}_{f_h, t}^* \leq \frac{c\tau}{(1-\tau)^2\gamma} \frac{\alpha}{\alpha-c} \mathcal{L}_{\mu_\tau, t}^*.$$

From Theorem 2, we know that $\mathcal{L}_{\mu_\tau, t}^* = \tau^t$. Thus, we obtain that

$$\mathcal{L}_{f_h, t}^* \asymp \tau^t.$$

A simple calculation gives that $\tau = \frac{1-\xi}{1-\gamma} \rho$, which yields the claimed result. As for the Gaussian case, an exact calculation of $\mathcal{L}_{f_h, t}^*$ is actually possible. But, after investigation, the resulting expression is lengthy and fairly difficult to simplify, whereas we are primarily interested in the scaling in terms of the iteration number t .

B. Proofs of intermediate results

B.1. Proof of Lemma 2.1

Suppose that $\{x_t\}$ is generated by a first-order method (2). Fix $t \geq 1$, then there exists $\alpha_{0,t}, \dots, \alpha_{t-1,t}$ such that

$$x_t = x_{t-1} + \sum_{j=0}^{t-1} \alpha_{j,t} H_S^{-1} A^\top (Ax_j - b). \quad (37)$$

Multiplying both sides of (37) by $U^\top A$, subtracting $U^\top Ax^*$ and using the normal equation $A^\top Ax^* = A^\top b$, we find that

$$\Delta_t = \Delta_{t-1} + \sum_{j=0}^{t-1} \alpha_{j,t} C_S^{-1} \Delta_j. \quad (38)$$

First, we aim to show that there exists a polynomial $p_t \in \mathbb{R}_t^0[X]$ such that $\Delta_t = p_t(C_S^{-1}) \Delta_0$. We proceed by induction over $t \geq 0$. For $t=0$, the claim is true. Suppose that for some $t \geq 1$, it holds that $\Delta_j = p_j(C_S^{-1}) \Delta_0$ with $p_j \in \mathbb{R}_j^0[X]$ for $j=0, \dots, t-1$. Then, we have from (38) that

$$\Delta_t = p_{t-1}(C_S^{-1}) \Delta_0 + \sum_{j=0}^{t-1} \alpha_{j,t} C_S^{-1} p_j(C_S^{-1}) \Delta_0 \quad (39)$$

$$= \left(p_{t-1}(C_S^{-1}) + \sum_{j=0}^{t-1} \alpha_{j,t} C_S^{-1} p_j(C_S^{-1}) \right) \Delta_0. \quad (40)$$

We set $p_t(x) = p_{t-1}(x) + \sum_{j=0}^{t-1} \alpha_{j,t} x p_j(x)$. It holds that $p_t(0) = p_{t-1}(0) + 0 = 1$, and $\deg(p_t) \leq t$ since $\deg(p_{t-1}) \leq t-1$ and $\deg(x p_j(x)) \leq j+1 \leq t$ for $j=0, \dots, t-1$. Then, from (40), we have $\Delta_t = p_t(C_S^{-1}) \Delta_0$, which concludes the induction.

Second, we aim to show that $\lim_{n \rightarrow \infty} \frac{\mathbb{E}[\|\Delta_t\|^2]}{\mathbb{E}[\|\Delta_0\|^2]} = \int_{\mathbb{R}} p^2(\lambda^{-1}) d\mu(\lambda)$, where μ is the l.s.d. of C_S for S an $m \times n$ Gaussian or SRHT embedding. The Gaussian case is straightforward to prove, by using the rotational invariance of the Gaussian distribution. The SRHT case is more involved, and we leverage tools from free probability theory.

B.1.1. THE GAUSSIAN CASE

Let S be an $m \times n$ random matrix with i.i.d. entries $\mathcal{N}(0, 1/m)$. Then, by rotational invariance, SU is an $m \times d$ matrix with i.i.d. entries $\mathcal{N}(0, 1/m)$. Write the eigenvalue decomposition $C_S = V \Sigma V^\top$ where V is a $d \times d$ orthogonal matrix, and Σ a diagonal matrix with positive entries $\lambda_1, \dots, \lambda_d$. A standard result states that V and Σ are independent matrices, and V is Haar-distributed.

Fix $t \geq 0$, and let $p_t \in \mathbb{R}_t^0[X]$ such that $\Delta_t = p_t(C_S^{-1}) \Delta_0$. Taking the squared norm and the expectation, we obtain that

$$\begin{aligned} \mathbb{E}[\|\Delta_t\|^2] &= \mathbb{E}[\Delta_0^\top p_t^2(C_S^{-1}) \Delta_0] \\ &= \mathbb{E}[\Delta_0^\top V p_t^2(\Sigma^{-1}) V^\top \Delta_0]. \end{aligned}$$

Using the independence of Σ , V and Δ_0 and writing $V = [v_1, \dots, v_d]$, we further obtain that

$$\begin{aligned} \mathbb{E}[\|\Delta_t\|^2] &= \mathbb{E}[\Delta_0^\top V \mathbb{E}[p_t^2(\Sigma^{-1})] V^\top \Delta_0] \\ &= \sum_{i=1}^d \mathbb{E}[(v_i^\top \Delta_0)^2] \mathbb{E}[p_t^2(\lambda_i^{-1})]. \end{aligned}$$

Since each v_i is uniformly distributed on the unit sphere, we have that $\mathbb{E}[(v_i^\top \Delta_0)^2] = \frac{1}{d} \mathbb{E}\|\Delta_0\|^2$, so that

$$\begin{aligned} \mathbb{E}[\|\Delta_t\|^2] &= \frac{1}{d} \mathbb{E}\|\Delta_0\|^2 \mathbb{E}\left[\sum_{i=1}^d p_t^2(\lambda_i^{-1})\right] \\ &= \mathbb{E}\|\Delta_0\|^2 \frac{1}{d} \text{trace} \mathbb{E}[p_t^2(C_S^{-1})]. \end{aligned}$$

Dividing both sides of the above equation by $\mathbb{E}\|\Delta_0\|^2$ and taking the limit $d \rightarrow \infty$, we obtain the claimed result,

$$\frac{\mathbb{E}[\|\Delta_t\|^2]}{\mathbb{E}[\|\Delta_0\|^2]} = \int_{\mathbb{R}} p_t^2(\lambda^{-1}) d\mu(\lambda).$$

B.1.2. THE SRHT CASE

The SRHT does not satisfy rotational invariance as the Gaussian distribution (or Haar matrices), and we need to use a different approach for this proof, based on *asymptotically liberating sequences of unitary matrices* [2].

Let S be an $m \times n$ SRHT embedding. We denote by μ the l.s.d. of the matrix C_S . Following the same first steps as for the Gaussian case, we have that

$$\mathbb{E}[\|\Delta_t\|^2] = \mathbb{E} \operatorname{trace} [p_t^2(C_S^{-1}) \Delta_0 \Delta_0^\top] = \mathbb{E} \operatorname{trace} [p_t^2(C_S^{-1}) \Sigma_0], \quad (41)$$

where $\Sigma_0 := \mathbb{E} \Delta_0 \Delta_0^\top$. Writing $p_t^2(x) = \sum_{k=0}^t a_k x^{2k}$, it follows that

$$\mathbb{E}[\|\Delta_t\|^2] = \sum_{k=0}^t a_k \mathbb{E} \operatorname{trace} [C_S^{-2k} \Sigma_0]. \quad (42)$$

Introducing the matrix $\tilde{\Sigma}_0 = \frac{\Sigma_0}{\operatorname{trace} \Sigma_0 / d}$, we further obtain

$$\frac{\mathbb{E}[\|\Delta_t\|^2]}{\mathbb{E}[\|\Delta_0\|^2]} = \sum_{k=0}^t a_k \frac{1}{d} \mathbb{E} \operatorname{trace} [C_S^{-2k} \tilde{\Sigma}_0]. \quad (43)$$

We use the following result, whose proof leverages some notions from free probability theory. We defer the proof to Appendix B.2.

Lemma 1. *It holds that for any $k \geq 0$,*

$$\lim_{n \rightarrow \infty} \frac{1}{d} \mathbb{E} \operatorname{trace} [C_S^{-2k} \tilde{\Sigma}_0] = \lim_{n \rightarrow \infty} \frac{1}{d} \mathbb{E} \operatorname{trace} [C_S^{-2k}] = \int_{\mathbb{R}} \lambda^{-2k} d\mu(\lambda). \quad (44)$$

Combining (43) and the result of Lemma 1, it follows that

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[\|\Delta_t\|^2]}{\mathbb{E}[\|\Delta_0\|^2]} = \sum_{k=0}^t a_k \int_{\mathbb{R}} \lambda^{-2k} d\mu(\lambda) = \int_{\mathbb{R}} p_t^2(\lambda^{-1}) d\mu(\lambda), \quad (45)$$

which is the claimed result.

B.2. Proof of Lemma 1

We introduce a few needed concepts from free probability that will be used in this proof. We refer the reader to [32, 11, 21, 3] for an extensive introduction to this field. Consider the algebra \mathcal{A}_n of $n \times n$ random matrices. For $X_n \in \mathcal{A}_n$, we define the linear functional $\tau_n(X_n) := \frac{1}{n} \mathbb{E}[\operatorname{trace} X_n]$. Then, we say that a family $\{X_{n,1}, \dots, X_{n,I}\}$ of random matrices in \mathcal{A}_n is *asymptotically free* if for every $i \in \{1, \dots, I\}$, $X_{n,i}$ has a limiting spectral distribution, and if $\tau_n \left(\prod_{j=1}^m P_j(X_{n,i_j} - \tau(P_j(X_{n,i_j}))) \right) \rightarrow 0$ almost surely for any positive integer m , any polynomials P_1, \dots, P_m and any indices $i_1, \dots, i_m \in \{1, \dots, I\}$ with $i_1 \neq i_2, \dots, i_{m-1} \neq i_m$.

Let S be an $n \times n$ SRHT embedding (we consider the SRHT before discarding its zero rows). By definition, we can write $S = BHW$, where B is an $n \times n$ matrix with i.i.d. Bernoulli entries on the diagonal, with success probability m/n , $H = H_n$ is the n -th Walsh-Hadamard matrix. The matrix W is an $n \times n$ bi-signed permutation, i.e., $W = DP$, where D is a diagonal matrix with i.i.d. random signs, and P is an $n \times n$ uniformly random permutation matrix.

We aim to show that for any $k \geq 0$,

$$\lim_{d \rightarrow \infty} \tau_d \left(C_S^{-k} \tilde{\Sigma}_0 \right) = \lim_{d \rightarrow \infty} \tau_d \left(C_S^{-k} \right). \quad (46)$$

We reduce the problem of proving (46) to the following, which is more simple to treat. The proof of this reduction is deferred to Appendix D.1.

Lemma 2. *Suppose that for any $k \geq 0$, we have*

$$\lim_{d \rightarrow \infty} \tau_d \left(C_S^k \tilde{\Sigma}_0 \right) = \lim_{d \rightarrow \infty} \tau_d \left(C_S^k \right). \quad (47)$$

Then, the claim (46) is true for any $k \geq 0$.

Thus, we aim to show (47) for all $k \geq 0$.

It holds that

$$\begin{aligned} C_S &= U^\top S^\top S U = (U^\top W^\top H B)(B H W U) \\ &= U^\top W^\top H B^2 H W U \\ &= U^\top W^\top H B H W U, \end{aligned}$$

where we used $B^2 = B$ in the fourth equality. Further, we have the following equality in distribution, whose proof is deferred to Appendix B.3.

Lemma 3. *It holds that*

$$U^\top W^\top H B H W U \stackrel{d}{=} U^\top W^\top H W B W^\top H W U. \quad (48)$$

Consequently,

$$C_S \stackrel{d}{=} U^\top W^\top H W B W^\top H W U. \quad (49)$$

Let $k \geq 0$. We have $W^\top W = I_n$, $U^\top U = I_d$, $B^2 = B$, $H^2 = H$ and $\tau_d(\tilde{\Sigma}_0) = 1$. Using (48), we find

$$\tau_d \left(C_S^k \tilde{\Sigma}_0 \right) = \tau_d \left((U^\top W^\top H W B W^\top H W U)^k \tilde{\Sigma}_0 \right) = \frac{n}{d} \cdot \tau_n \left(X_1 (Y X_2)^k \right), \quad (50)$$

where we introduced the matrices $X_1 := W U \tilde{\Sigma}_0 U^\top W^\top$, $X_2 := W U U^\top W^\top$ and $Y := H W B W^\top H$. These matrices satisfy the following collection of properties, whose proof is deferred to Appendix D.2.

Lemma 4. *It holds that $X_1 X_2 = X_2 X_1 = X_1$, $X_2^2 = X_2$, $Y^2 = Y$,*

$$\lim_{n \rightarrow \infty} \tau_n(X_1) = \lim_{n \rightarrow \infty} \tau_n(X_2), \quad (51)$$

and the sets of matrices $\{X_1, X_2\}$ and $\{Y\}$ are asymptotically free.

Further, for any $k \geq 1$, we have

$$\lim_{n \rightarrow \infty} \tau_n(X_1 (Y X_2)^k) = \lim_{n \rightarrow \infty} \tau_n(X_2 (Y X_2)^k). \quad (52)$$

Now, observe that

$$\begin{aligned} \tau_n(X_2 (Y X_2)^k) &= \tau_n(W U U^\top W^\top (H W B W^\top H W U U^\top W^\top)^k) \\ &= \frac{d}{n} \tau_d((U^\top W^\top H W B W^\top H W U)^k) \\ &= \frac{d}{n} \tau_d(C_S^k), \end{aligned}$$

where we used the commutativity of the trace in the second equality, and the equality in distribution (48) for the third equality. Consequently,

$$\lim_{n \rightarrow \infty} \tau_n(X_1 (Y X_2)^k) = \gamma \lim_{d \rightarrow \infty} \tau_d(C_S^k). \quad (53)$$

Combining the above equality (53) with equality (50), we obtain the claimed result (47).

B.3. Proof of Lemma 3

Note that both B and D are diagonal matrices whose diagonal entries are i.i.d. random variables, and P is a permutation matrix. Define $\tilde{B} = PBP^\top$ and $\tilde{D} = P^\top DP$, then $\tilde{B} \stackrel{d}{=} B$, $\tilde{D} \stackrel{d}{=} D$,

$$DP = P\tilde{D}, \quad P^\top D = \tilde{D}P^\top. \quad (54)$$

It follows that

$$\begin{aligned} U^\top W^\top HWBW^\top HWU &= U^\top P^\top DHDPBP^\top DHDP U \\ &= U^\top P^\top DHP\tilde{D}B\tilde{D}P^\top HDPU \\ &= U^\top P^\top DH_nPB\tilde{D}^2P^\top H_nDPU \\ &= U^\top P^\top DH_nPBP^\top H_nDPU \\ &= U^\top P^\top DH_n\tilde{B}H_nDPU \\ &\stackrel{d}{=} U^\top P^\top DH_nBH_nDPU, \end{aligned}$$

where the second equation follows from (54), the third equation holds because \tilde{D} and B are diagonal so they commute, while the fourth equation holds because $\tilde{D}^2 = I_n$.

B.4. Proof of Lemma 2.4 – Alternative construction of the polynomials $\{\Pi_k\}$

We recall that for a polynomial P and a measure (resp. density) μ , we will denote $\mu[P] := \int_{\mathbb{R}} P(x)\mu(x)d\mu(x)$ (resp. $\mu[P] := \int_{\mathbb{R}} P(x)\mu(x)dx$). Thus, for a density μ , the reader should be aware that $\mu[x]$ and $\mu(x)$ refer to different quantities.

We present an alternative construction of the orthogonal family $\{\Pi_k\}$ with respect to μ_ρ , explicitly based on the Chebyshev polynomials of the second kind. This explicit construction allows us to leverage several properties of the polynomials $\{\Pi_k\}$ which are useful to perform calculations and prove Lemma 2.4, as well as Theorem 2.

We introduce the shifted Chebyshev polynomials of the second kind, which are defined by the recurrence

$$Q_0(x) = 1, \quad Q_1(x) = \frac{x - (1+\rho)}{\sqrt{\rho}}, \quad Q_{k+1}(x) = \frac{x - (1+\rho)}{\sqrt{\rho}} Q_k(x) - Q_{k-1}(x). \quad (55)$$

A standard result states that the polynomials Q_k are orthonormal with respect to the measure $\nu(x)dx := x\mu_\rho(x)dx$. We set $\hat{\Pi}_0(x) = 1$, and for $k \geq 1$,

$$\hat{\Pi}_k(x) := 1 - \sum_{j=1}^k (-1)^{j-1} \sqrt{\rho}^{j-1} x Q_{k-1}(x). \quad (56)$$

For instance, we have $\hat{\Pi}_1(x) = 1 - x$ and $\hat{\Pi}_2(x) = 1 - (2 + \rho)x + x^2$.

We aim to show that $\{\hat{\Pi}_k\}$ is an orthogonal family with respect to μ_ρ and then, that $\hat{\Pi}_k = \Pi_k$.

First, we show that the polynomials $\hat{\Pi}_k$ form an orthonormal family with respect to μ_ρ such that $\deg(\hat{\Pi}_k) = k$ and $\hat{\Pi}_k(0) = 1$. For $k \geq 1$, we define the polynomial $T_k(x) = xQ_{k-1}(x)$ and the measure $\lambda_\rho(x) = x^{-1}\mu_\rho(x)$. We have that $\lambda_\rho[T_k T_\ell] = \nu_\rho[Q_{k-1}Q_{\ell-1}] = \delta_{k\ell}$, so that the T_k are orthonormal with respect to λ_ρ . Since $\deg(Q_{k-1}) = k-1$, we have $\deg(T_k) = k$. We also have $T_k(0) = 0 \cdot Q_{k-1}(0) = 0$.

Second, we show that $\mu_\rho[Q_k] = (-1)^k \sqrt{\rho}^k$, which will immediately imply that

$$\lambda_\rho[T_k] = \lambda_\rho[xQ_{k-1}(x)] = \mu_\rho[Q_{k-1}] = (-1)^{k-1} \sqrt{\rho}^{k-1}. \quad (57)$$

We denote $u_k := \mu_\rho[Q_k]$. The measure μ_ρ is a probability measure, so that $u_0 = 1$. Further, we have

$$u_1 = \mu_\rho[Q_1] = \int_a^b \frac{x - (1+\rho)}{\sqrt{\rho}} \mu_\rho(x) dx = \frac{-1 - \rho + \int_a^b x \mu_\rho(x) dx}{\sqrt{\rho}}.$$

The first moment $\mu_\rho[x]$ is equal to 1, so that $u_1 = -\sqrt{\rho}$. From the recurrence relationship (55), we obtain $u_{k+1} = -\frac{1+\rho}{\sqrt{\rho}}u_k - u_{k-1}$. The characteristic equation $x^2 + \frac{1+\rho}{\sqrt{\rho}}x + 1 = 0$ has roots $-1/\sqrt{\rho}$ and $-\sqrt{\rho}$. Therefore, $u_k = \alpha \frac{(-1)^k}{\sqrt{\rho}^k} + \beta(-1)^k \sqrt{\rho}^k$ for some $\alpha, \beta \in \mathbb{R}$. Using the initial values u_0 and u_1 , we find $\alpha = 0$ and $\beta = 1$. This yields the claimed formula for u_k .

Then, using the definition (56), we have

$$\widehat{\Pi}_k = 1 - \sum_{j=1}^k (-1)^{j-1} \sqrt{\rho}^{j-1} x Q_{k-1}(x) \quad (58)$$

$$= 1 - \sum_{j=1}^k \lambda_\rho[T_k] T_k(x). \quad (59)$$

Hence, recognizing the Gram-Schmidt orthogonalization of the constant polynomial 1 with respect to $\{T_1, \dots, T_k\}$, we deduce that the family $\{\widehat{\Pi}_k, T_1, \dots, T_k\}$ is orthogonal with respect to λ_ρ , and is a basis of $\mathbb{R}_k[X]$. Consider now the variational problem

$$\min_{p \in \mathbb{R}_k^0[X]} \int p^2(x) \lambda_\rho(x) dx. \quad (60)$$

Let $p \in \mathbb{R}_k^0[X]$ and decompose p as $p = \alpha_0 \widehat{\Pi}_k + \sum_{j=1}^k \alpha_j T_j$. Using $p(0) = 1$, $\widehat{\Pi}_k(0) = 1$ and $T_j(0) = 0$, we get that α_0 must be equal to 1. Then,

$$\begin{aligned} \int p^2(x) \lambda_\rho(x) dx &= \int \widehat{\Pi}_k^2(x) \lambda_\rho(x) dx + 2 \sum_{j=1}^k \alpha_j \int \widehat{\Pi}_k(x) T_j(x) \lambda_\rho(x) dx \\ &\quad + \int \left(\sum_{j=1}^k \alpha_j T_j(x) \right)^2 \lambda_\rho(x) dx. \end{aligned}$$

The cross-term is equal to 0 by orthogonality of the family $\{\widehat{\Pi}_k, T_1, \dots, T_k\}$. The third-term is non-negative, and equal to 0 if and only if $p = \widehat{\Pi}_k$. Therefore, the minimizer of the variational problem (60) is exactly $\widehat{\Pi}_k$. On the other hand, applying Lemma 2.2 with $\nu = x \lambda_\rho = \mu_\rho$, we know that the solution of each of the problems (60) (for varying k) is unique, and the solutions form an orthogonal family with respect to $x \lambda_\rho(x) dx = \mu_\rho(x) dx$. Thus, we obtain that the family $\{\widehat{\Pi}_k\}$ is orthogonal with respect to μ_ρ .

Finally, we show that the sequence $\{\widehat{\Pi}_k\}$ satisfies the recurrence relationship (13). Observe that

$$\begin{aligned} x \widehat{\Pi}_k(x) &= x - \sum_{j=1}^k \lambda_\rho[T_j] x T_j(x) = x - \lambda_\rho[T_1] x T_1(x) - \sum_{j=2}^k \lambda_\rho[T_j] x T_j(x) \\ &= x - x^2 - \sum_{j=2}^k \lambda_\rho[T_j] x T_j(x). \end{aligned}$$

Multiplying (55) by x and using the definition $T_k(x) = x Q_{k-1}(x)$, we find that for $k \geq 2$,

$$x T_j(x) = \sqrt{\rho} (T_{j-1}(x) + T_{j+1}(x)) + (1 + \rho) T_j(x).$$

Using the above decomposition of $x T_j(x)$, it obtain $\sum_{j=2}^k \lambda_\rho(T_j) x T_j(x) = s_1 + s_2 + s_3$, where

$$\begin{aligned} s_1 &:= \sqrt{\rho} \sum_{j=2}^k \lambda_\rho(T_j) T_{j+1}(x) = \sum_{j=2}^k (-1)^{j-1} \sqrt{\rho}^j T_{j+1}(x) \\ &= \sum_{j=3}^{k+1} (-1)^j \sqrt{\rho}^{j-1} T_j(x) \\ &= \widehat{\Pi}_{k+1}(x) - 1 + T_1(x) - \sqrt{\rho} T_2(x) \\ &= \widehat{\Pi}_{k+1}(x) - 1 + x - x^2 + (1 + \rho) x, \end{aligned}$$

the second term is $s_2 := \sqrt{\rho} \sum_{j=2}^k \lambda_\rho [T_j] T_{j-1}(x) = \sum_{j=2}^k (-1)^{j-1} \sqrt{\rho^j} T_{j-1}(x) = \rho (\widehat{\Pi}_{k-1}(x) - 1)$ and the third term is $s_3 := (1 + \rho) \sum_{j=2}^k \lambda_\rho [T_j] T_j(x) = -(1 + \rho) (\widehat{\Pi}_k(x) - 1 + x)$. Consequently,

$$\begin{aligned} x\Pi_k(x) &= x - x^2 - s_1 - s_2 - s_3 \\ &= x - x^2 - \widehat{\Pi}_{k+1}(x) + 1 - x + x^2 - (1 + \rho)x - \rho (\widehat{\Pi}_{k-1}(x) - 1) + (1 + \rho) (\widehat{\Pi}_k(x) - 1 + x) \\ &= -\widehat{\Pi}_{k+1}(x) - \rho \widehat{\Pi}_{k-1}(x) + (1 + \rho) \widehat{\Pi}_k(x), \end{aligned}$$

which is the claimed recurrence. We deduce that $\widehat{\Pi}_k = \Pi_k$, and that the family $\{\Pi_k\}$ is orthogonal with respect to μ_ρ .

C. Description of numerical experiments

Numerical simulations are run in *Python* with the numerical linear algebra module *NumPy* and the scientific computation module *SciPy*, on a machine with 256Gb of memory.

To generate an $m \times n$ Haar matrix S_h , we sample an $m \times n$ matrix G with i.i.d. Gaussian entries $\mathcal{N}(0, 1)$, and we set S_h to be its $m \times n$ matrix of right singular vectors. To generate an $m \times n$ SRHT matrix, we follow the description given in Section 1. The plots correspond to one trial for each embedding.

C.1. Figure 1

We set $n = 8192$, $d = 1640$ and $m \in \{1720, 3280, 4915\}$. We generate the plots of μ_ρ and $f_{h,r}$ by discretizing their respective supports with step size $1e-5$.

C.2. Figures 2 and 3

We generate an $n \times d$ Gaussian matrix G with i.i.d. entries, and we compute its left singular matrix U and right singular matrix V . Then, we set $A = U\Sigma V^\top$, where Σ is a $d \times d$ diagonal matrix with entries $\Sigma_j = 0.98^j$ for $j = 1, \dots, d$. We generate a vector b using a planted model $b = Ax_{\text{pl}} + \frac{1}{\sqrt{n}} \mathcal{N}(0, I_n)$, and $x_{\text{pl}} \sim \frac{1}{\sqrt{d}} \mathcal{N}(0, I_d)$. Note that, although the performance of the algorithms do not depend on the data A and b , we choose a standard statistical model to generate the data, and a data matrix with a very large condition number.

Algorithms 1 and 2 are implemented following their pseudo-code description. We use small perturbations of the algorithmic parameters by setting $a_t^\delta = (1 + \delta)a_t$ and $b_t^\delta = (1 - \delta)b_t$ with $\delta = 0.01$ – where a_t and b_t correspond to the parameters as described in Theorem 1. Similarly, for the Heavy-ball method with fixed SRHT embeddings and parameters derived based on our new asymptotic edge eigenvalues (“SRHT (edge eig.)”), we use instead the slightly perturbed edge eigenvalues $\lambda_h^\delta = (1 - \delta)\lambda_h$ and $\Lambda_h^\delta = (1 + \delta)\Lambda_h$, with $\delta = 0.01$. These small perturbations of the parameters are necessary in practice due to the finite-sample approximations. For the Heavy-ball method with fixed SRHT embeddings based on the bounds of [31] (“SRHT (baseline)”), we use the parameters prescribed in [13]. For the Heavy-ball method with refreshed SRHT embeddings (“SRHT (refreshed)”), we use the parameters prescribed in [15]. For each algorithm, results are averaged over 20 independent trials (using the same data A and b).

D. Proofs of auxiliary results

D.1. Proof of Lemma 2

Suppose that (47) holds. Let $k \geq 0$. We have

$$C_S^{-k} = (I_d - (I_d - C_S))^{-k} = \left(\sum_{j=0}^{\infty} (I_d - C_S)^j \right)^k, \quad (61)$$

where the series expansion $(I_d - (I_d - C_S))^{-1} = \sum_{j=0}^{\infty} (I_d - C_S)^j$ holds almost surely, due to the fact that C_S has spectrum within $(0, 1)$ almost surely. There exist coefficients $\{a_\ell\}$ such that $\left(\sum_{j=0}^{\infty} x^j \right)^k = \sum_{\ell=0}^{\infty} a_\ell x^\ell$, and such that the

sum is absolutely convergent, i.e., $\sum_{\ell=0}^{\infty} |a_{\ell}| |x|^{\ell} < +\infty$, for any $x \in (0, 1)$. Consequently,

$$C_S^{-k} = \sum_{\ell=0}^{\infty} a_{\ell} C_S^{\ell} \quad (62)$$

Then, by absolute convergence of $\sum_{\ell} a_{\ell} x^{\ell}$ and using the fact that $\|C_S\|_2 < 1$, we can exchange the operator τ_d and the infinite sum, so that

$$\tau_d(C_S^{-k}) = \tau_d\left(\sum_{\ell=0}^{\infty} a_{\ell} C_S^{\ell}\right) = \sum_{\ell=0}^{\infty} a_{\ell} \tau_d(C_S^{\ell}) . \quad (63)$$

and writing the latter as a series in C_S , we obtain the claimed result. Due to the fact that $\sup_{\ell} \lim_{d \rightarrow \infty} \tau_d(C_S^{\ell}) < 1$, and using again the absolute convergence of $\sum_{\ell} a_{\ell} x^{\ell}$ for $|x| < 1$, it follows that

$$\lim_{d \rightarrow \infty} \tau_d(C_S^{-k}) = \lim_{d \rightarrow \infty} \sum_{\ell=0}^{\infty} a_{\ell} \tau_d(C_S^{\ell}) \quad (64)$$

$$= \sum_{\ell=0}^{\infty} a_{\ell} \lim_{d \rightarrow \infty} \tau_d(C_S^{\ell}) \quad (65)$$

$$= \sum_{\ell=0}^{\infty} a_{\ell} \lim_{d \rightarrow \infty} \tau_d(C_S^{\ell} \tilde{\Sigma}_0) . \quad (66)$$

Using the same arguments, we find that

$$\sum_{\ell=0}^{\infty} a_{\ell} \lim_{d \rightarrow \infty} \tau_d(C_S^{\ell} \tilde{\Sigma}_0) = \lim_{d \rightarrow \infty} \tau_d(C_S^{-k} \tilde{\Sigma}_0) , \quad (67)$$

and we conclude that

$$\tau_d(C_S^{-k} \tilde{\Sigma}_0) = \tau_d(C_S^{-k}) \quad (68)$$

D.2. Proof of Lemma 4

We have

$$X_1 X_2 = W U \tilde{\Sigma}_0 U^{\top} W^{\top} W U U^{\top} W^{\top} = W U \tilde{\Sigma}_0 U^{\top} W^{\top} = X_1$$

where we used in the second equality $U^{\top} W^{\top} W U = I_d$. Similarly, we obtain $X_2 X_1 = X_1$.

We have

$$Y^2 = (H W B W^{\top} H)(H W B W^{\top} H) = H W B W^{\top} H = Y$$

where we used in the second equality $B W^{\top} H H W B = B$.

We have

$$X_2^2 = W U U^{\top} W^{\top} W U U^{\top} W^{\top} = W U U^{\top} W^{\top} = X_2 ,$$

where we used in the second equality $U^{\top} W^{\top} W U = I_d$.

Further, it holds that

$$\lim_{n \rightarrow \infty} \tau_n(X_1) = \gamma \lim_{d \rightarrow \infty} \tau_d(\tilde{\Sigma}_0) = \gamma = \lim_{n \rightarrow \infty} \tau_n(X_2) , .$$

We show asymptotic freeness. Note that the matrices $U U^{\top}$, B and $\tilde{\Sigma}_0$ have l.s.d. compactly supported. For the latter, this directly follows from our initial assumption that the condition number of the matrix $U^{\top} A \mathbb{E}[x_0 x_0^{\top}] A^{\top} U + U^{\top} b b^{\top} U$

remains bounded. Then, applying Corollary 3.2 from [2] with the set of asymptotically liberating matrices $\{W, HW\}$, we immediately obtain asymptotic freeness of $\{X_1, X_2\}$ and $\{Y\}$.

It remains to show that for any $k \geq 0$,

$$\lim_{n \rightarrow \infty} \tau_n(X_1(YX_2)^k) = \lim_{n \rightarrow \infty} \tau_n(X_2(YX_2)^k). \quad (69)$$

For the rest of this proof, we use the more compact notations $a := X_1$, $b := Y$, $c := X_2$ and $\varphi = \lim_{n \rightarrow \infty} \tau_n$. We show (69) by induction over $k \geq 0$. For $k = 0$, the claim is true because $\varphi(a) = \varphi(c)$ as shown above. Fix $k \geq 1$ and suppose that the claim is true for $j = 0, \dots, k-1$. By asymptotic freeness, we have

$$\varphi\left((a - \varphi(a))(b - \varphi(b))(c - \varphi(c))\right)^k = 0. \quad (70)$$

We expand the left-hand side of the above equation as

$$\begin{aligned} & \varphi\left((a - \varphi(a))(b - \varphi(b))(c - \varphi(c))\right)^k \\ = & \varphi(a(bc)^k) + \sum_{\substack{\delta_1, \dots, \delta_{2k} \in \{0,1\} \\ (\delta_1, \dots, \delta_{2k}) \neq (1, \dots, 1)}} \varphi\left(ab^{\delta_1}c^{\delta_2}b^{\delta_3} \dots c^{\delta_{2k}}(-\varphi(b))^{1-\delta_1} \dots (-\varphi(c))^{1-\delta_{2k}}\right) \\ = & \varphi(a(bc)^k) + \sum_{\substack{\delta_1, \dots, \delta_{2k} \in \{0,1\} \\ (\delta_1, \dots, \delta_{2k}) \neq (1, \dots, 1)}} (-\varphi(b))^{1-\delta_1} \dots (-\varphi(c))^{1-\delta_{2k}} \varphi\left(ab^{\delta_1}c^{\delta_2} \dots c^{\delta_{2k}}\right). \end{aligned}$$

For binary exponents $(\delta_1, \dots, \delta_{2k}) \neq (1, \dots, 1)$, the product of non-commutative matrices $b^{\delta_1}c^{\delta_2}b^{\delta_3} \dots c^{\delta_{2k}}$ must have a sub-product of the form bb or cc . Using the fact that $b^2 = b$ and $c^2 = c$, it follows that there exists some integer ℓ such that $0 \leq \ell < k$, and

$$b^{\delta_1}c^{\delta_2}b^{\delta_3} \dots c^{\delta_{2k}} = (bc)^\ell.$$

Using the induction hypothesis, we have

$$\varphi\left(ab^{\delta_1}c^{\delta_2} \dots c^{\delta_{2k}}\right) = \varphi(a(bc)^\ell) = \varphi(c(bc)^\ell) = \varphi\left(cb^{\delta_1}c^{\delta_2} \dots c^{\delta_{2k}}\right).$$

Consequently, we get

$$\begin{aligned} & \varphi\left((a - \varphi(a))(b - \varphi(b))(c - \varphi(c))\right)^k \\ = & \varphi(a(bc)^k) + \sum_{\substack{\delta_1, \dots, \delta_{2k} \in \{0,1\} \\ (\delta_1, \dots, \delta_{2k}) \neq (1, \dots, 1)}} \varphi\left(cb^{\delta_1}c^{\delta_2}b^{\delta_3} \dots c^{\delta_{2k}}(-\varphi(b))^{1-\delta_1} \dots (-\varphi(c))^{1-\delta_{2k}}\right) \end{aligned}$$

On the other hand, using asymptotic freeness again, we have

$$\begin{aligned} 0 &= \varphi\left((c - \varphi(c))(b - \varphi(b))(c - \varphi(c))\right)^k \\ &= \varphi(c(bc)^k) + \sum_{\substack{\delta_1, \dots, \delta_{2k} \in \{0,1\} \\ (\delta_1, \dots, \delta_{2k}) \neq (1, \dots, 1)}} \varphi\left(cb^{\delta_1}c^{\delta_2}b^{\delta_3} \dots c^{\delta_{2k}}(-\varphi(b))^{1-\delta_1} \dots (-\varphi(c))^{1-\delta_{2k}}\right) \end{aligned}$$

Combining the two above sets of equalities, we obtain

$$\varphi(a(bc)^k) = \varphi(c(bc)^k),$$

which concludes the induction, and the proof.