

Smoothed Analysis for the Conjugate Gradient Algorithm*

Govind MENON[†] and Thomas TROGDON[‡]

[†] *Division of Applied Mathematics, Brown University,
182 George St., Providence, RI 02912, USA
E-mail: govind_menon@brown.edu*

[‡] *Department of Mathematics, University of California,
Irvine, Rowland Hall, Irvine, CA, 92697-3875, USA
E-mail: ttrogdon@math.uci.edu*

Received May 23, 2016, in final form October 31, 2016; Published online November 06, 2016
<http://dx.doi.org/10.3842/SIGMA.2016.109>

Abstract. The purpose of this paper is to establish bounds on the rate of convergence of the conjugate gradient algorithm when the underlying matrix is a random positive definite perturbation of a deterministic positive definite matrix. We estimate all finite moments of a natural halting time when the random perturbation is drawn from the Laguerre unitary ensemble in a critical scaling regime explored in Deift et al. (2016). These estimates are used to analyze the expected iteration count in the framework of smoothed analysis, introduced by Spielman and Teng (2001). The rigorous results are compared with numerical calculations in several cases of interest.

Key words: conjugate gradient algorithm; Wishart ensemble; Laguerre unitary ensemble; smoothed analysis

2010 Mathematics Subject Classification: 60B20; 65C50; 35Q15

In honor of Percy Deift and Craig Tracy on their 70th birthdays

1 Introduction

It is conventional in numerical analysis to study the worst-case behavior of algorithms, though it is often the case that worst-case behavior is far from typical. A fundamental example of this nature is the behavior of LU factorization with partial pivoting. While the worst-case behavior of the growth factor in LU factorization with partial pivoting is exponential, the algorithm works much better in practice on ‘typical’ problems. The notion of *smoothed analysis* was introduced by Spielman and his co-workers to distinguish between the typical-case and worst-case performance for numerical algorithms in such situations (see [16, 18]). In recent work, also motivated by the distinction between typical and worst case behavior, the authors (along with P. Deift, S. Olver and C. Pfrang) investigated the behavior of several numerical algorithms with random input [2, 14] (see also [15]). These papers differ from smoothed analysis in the sense that ‘typical performance’ was investigated by viewing the algorithms as dynamical systems acting on random input. The main observation in our numerical experiments was an empirical universality of fluctuations for halting times. Rigorous results on universality have now been established in two cases – the conjugate gradient algorithm and certain eigenvalue algorithms [3, 4]. The work of [3] presents a true universality theorem. The work [4] revealed the unexpected emergence of

*This paper is a contribution to the Special Issue on Asymptotics and Universality in Random Matrices, Random Growth Processes, Integrable Systems and Statistical Physics in honor of Percy Deift and Craig Tracy. The full collection is available at <http://www.emis.de/journals/SIGMA/Deift-Tracy.html>

Tracy–Widom fluctuations around the *smallest* eigenvalue of LUE matrices in a regime where universality emerges. Thus, the analysis of a question in probabilistic numerical analysis led to a new discovery in random matrix theory.

Our purpose in this paper is to explore the connections between our work and smoothed analysis. We show that the results of [4] extend to a smoothed analysis of the conjugate gradient algorithm over *strictly* positive definite random perturbations (which constitute a natural class of perturbations for the conjugate gradient algorithm). To the best of our knowledge, this is the first instance of smoothed analysis for the conjugate gradient algorithm. More precisely, the results of [4] are used to establish rigorous bounds on the expected value of a halting time (Theorem 1.1 below). These bounds are combined with numerical experiments that show an interesting improvement of the conjugate gradient algorithm when it is subjected to random perturbations. In what follows, we briefly review the conjugate gradient algorithm, smoothed analysis and the Laguerre unitary ensemble, before stating the main theorem, and illustrating it with numerical experiments.

1.1 The conjugate gradient algorithm

The conjugate gradient algorithm is a Krylov subspace method to solve the linear system $Ax = b$ when $A > 0$ is a positive definite matrix. In this article, we focus on Hermitian positive definite matrices acting on \mathbb{C}^N , though the ideas extend to real, symmetric positive definite matrices.

We use the ℓ^2 inner product on \mathbb{C}^N , $\langle u, v \rangle_{\ell^2} = \sum_{i=1}^N u_i \bar{v}_i$, and $A > 0$ means that A is Hermitian and $\langle u, Au \rangle_{\ell^2} > 0$ for all $u \neq 0$. When $A > 0$, its inverse $A^{-1} > 0$, and we may define the norms

$$\|u\|_w^2 = \langle u, Au \rangle_{\ell^2} \quad \text{and} \quad \|u\|_{w^{-1}}^2 = \langle u, A^{-1}u \rangle_{\ell^2}.$$

In this setting, the simplest formulation of the conjugate gradient algorithm is as follows [8, 10]. In order to solve $Ax = b$, we define the increasing sequence of Krylov subspaces

$$\mathcal{K}_k = \text{span} \{b, Ab, \dots, A^{k-1}b\},$$

and choose the iterates $\{x_k\}_{k=1}^{\infty}$, $x_0 = 0$, to minimize the residual $r_k = b - Ax_k$ in the w^{-1} norm:

$$x_k = \operatorname{argmin}_{x \in \mathcal{K}_k} \|Ax - b\|_{w^{-1}}, \quad \|Ax_k - b\|_{w^{-1}} = \min_{x \in \mathcal{K}_k} \|Ax - b\|_{w^{-1}}.$$

Since $A > 0$, $x \in \mathcal{K}_k$ for some $k \leq N$ (in our random setting it follows that $\mathcal{K}_N = \mathbb{C}^N$ with probability 1), so that the method takes at most N steps in exact arithmetic¹. However, the residual decays exponentially fast, and a useful approximation is obtained in much fewer than N steps. Let λ_{\max} and λ_{\min} denote the largest and smallest eigenvalues of A and $\kappa = \lambda_{\max}/\lambda_{\min}$ the condition number. Then the rate of convergence in the ℓ^2 and w^{-1} norms is [7, Theorem 10.2.6]

$$\|r_k\|_{w^{-1}} \leq 2 \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^k \|r_0\|_{w^{-1}}. \quad (1.1)$$

Since A is positive definite, we have

$$\lambda_{\max}^{-1} \|x\|_{\ell^2}^2 \leq \langle x, A^{-1}x \rangle_{\ell^2} \leq \lambda_{\min}^{-1} \|x\|_{\ell^2}^2.$$

¹In calculations with finite-precision arithmetic the number of steps can be much larger than N and this will be taken into account in the numerical experiments in Section 1.4. The results presented here can be extended to the finite-precision case but only in the limit as the precision tends to ∞ using [8]. Tightening these estimates remains an important open problem.

Applying this estimate to (1.1) we find the rate of convergence of the residual in the ℓ^2 norm

$$\|r_k\|_{\ell^2} \leq 2\sqrt{\kappa} \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^k \|r_0\|_{\ell^2}. \quad (1.2)$$

These rates of convergence provide upper bounds on the following ϵ -dependent run times, which we call *halting times*:

$$\begin{aligned} \tau_\epsilon(A, b) &= \min \left\{ k : \frac{\|r_{k+1}\|_{\ell^2}}{\|r_0\|_{\ell^2}} \leq \epsilon \right\}, \\ \tau_{w,\epsilon}(A, b) &= \min \left\{ k : \frac{\|r_{k+1}\|_{w^{-1}}}{\|r_0\|_{w^{-1}}} \leq \epsilon \right\}. \end{aligned} \quad (1.3)$$

Note that we have set $x_0 = 0$ so that $r_0 = b$ (the estimates above hold for arbitrary x_0). In what follows, we will also assume that $\|b\|_{\ell^2} = 1$, so that the definitions above simplify further.

1.2 Smoothed analysis

Our main results are a theorem (Theorem 1.1) along with numerical evidence to demonstrate that the above worst-case estimates, can be used to obtain bounds on average-case behavior in the sense of smoothed analysis. In order to state the main result, we first review two basic examples of smoothed analysis [18], since these examples clarify the context of our work.

Roughly speaking, the smoothed analysis of a deterministic algorithm proceeds as follows. Given a deterministic problem, we perturb it randomly, compute the expectation of the run-time for the randomly perturbed problem and then take the maximum over all deterministic problems within a fixed class. Subjecting a deterministic problem to random perturbations provides a realistic model of ‘typical performance’, and by taking the maximum over all deterministic problems within a natural class, we retain an important aspect of worst-case analysis. A parameter σ^2 (the variance in our examples) controls the magnitude of the random perturbation. The final estimate of averaged run-time should depend explicitly on σ^2 in way that demonstrates that the average run-time is much better than the worst-case. Let us illustrate this idea with examples.

1.2.1 Smoothed analysis: The simplex algorithm

Assume $\bar{A} = (\bar{a}_1, \bar{a}_2, \dots, \bar{a}_d)$ is a deterministic matrix of size $N \times d$, and \bar{y} and z are deterministic vectors of size N and d , respectively. Let $T(\bar{A}, \bar{y}, z)$ be the number of simplex steps required to solve the linear program

$$\text{maximize } z^T x, \quad \text{subject to } \bar{A}x \leq \bar{y},$$

with the two-phase shadow-vertex simplex algorithm.

We subject the data \bar{A} and \bar{y} to a random perturbation σA , and σy , where A and y have iid normal entries with mean zero and standard deviation $\max_i \|(\bar{y}_i, \bar{a}_i)\|$. It is then shown in [19] that the expected number of simplex steps is controlled by

$$\mathbb{E}[T(\bar{A} + \sigma A, \bar{y} + \sigma y, z)] \leq P(1/\sigma, N, d),$$

where $P(a, b, c)$ is a polynomial. Thus, problems of polynomial complexity occupy a region of high probability.

1.2.2 Smoothed analysis: LU factorization without pivoting

Let \bar{A} be an $N \times N$ non-singular matrix and consider computing its LU factorization, $\bar{A} = LU$, without partial pivoting. The *growth factor* of \bar{A} , defined by

$$\rho(\bar{A}) = \frac{\|U\|_\infty}{\|\bar{A}\|_\infty},$$

may be exponentially large in the size of the matrix, as seen in the following classical example:

$$\bar{A} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 \\ -1 & 1 & 0 & 0 & 1 \\ -1 & -1 & 1 & 0 & 1 \\ -1 & -1 & -1 & 1 & 1 \\ -1 & -1 & -1 & -1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 \\ -1 & -1 & 1 & 0 & 0 \\ -1 & -1 & -1 & 1 & 0 \\ -1 & -1 & -1 & -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 2 \\ 0 & 0 & 1 & 0 & 4 \\ 0 & 0 & 0 & 1 & 8 \\ 0 & 0 & 0 & 0 & 16 \end{bmatrix}.$$

Generalizing this example to all N , we see that $\rho(\bar{A}) = 2^{N-1}/N$. This is close to the worst-case estimate of Wilkinson [22]. Now consider instead $\rho(\bar{A} + \sigma A)$ where the random perturbation A is an $N \times N$ matrix consisting of iid standard normal random variables. One of the results of [16] is

$$\mathbb{P}(\rho(\bar{A} + \sigma A) > 1 + t) \leq \frac{1}{\sqrt{2\pi}} \frac{N(N+1)}{\sigma t}. \quad (1.4)$$

Hence the probability that $\rho(\bar{A} + \sigma A) \geq 2^{N-1}/N$ is exponentially small! The above estimate relies on a tail bound on the condition number

$$\mathbb{P}(\kappa(\bar{A} + \sigma A) > t) \leq \frac{14.1n(1 + \sqrt{2(\log t)/9n})}{t\sigma}.$$

The example above may also be used to demonstrate exponential growth *with partial pivoting*. However, to the best of our knowledge, there are no smoothed analysis bounds analogous to (1.4) that include the effect of pivoting.

1.3 The main result

We now formulate a notion of smoothed analysis for the halting time of the conjugate gradient algorithm. In order to do so, we must choose a matrix ensemble over which to take averages. Since the conjugate gradient algorithm is restricted to positive definite matrices it is natural to choose random perturbations that are also positive definite. The fundamental probability measure on Hermitian positive definite matrices is the Laguerre unitary ensemble (LUE), or Wishart ensemble, defined as follows. Assume N is a positive integer and $\alpha > -N$ is another integer. Let X be an $N \times (N + \alpha)$ matrix of iid standard complex normal random variables². The Hermitian matrix $W = XX^*$ is an LUE matrix with parameter $1 + \alpha/N$.

The parameter α plays an important role in our work. The case $\alpha = 0$ is critical in the following sense. When $-N < \alpha < 0$, the random matrix $W = XX^*$ is positive *semi-definite* and 0 is an eigenvalue of multiplicity $-\alpha$ with probability 1. In particular, the condition number of W is infinite almost surely. On the other hand, when $\alpha \geq 0$, the random matrix W is almost surely strictly positive definite. When $\alpha = 0$, Edelman [5] showed that the condition number of W is heavy-tailed, and does not have a finite mean (see also [16] and the previous examples). On the other hand, if α grows linearly with N , say $\lim_{N \rightarrow \infty} \alpha/N = p > 0$, the leading-order

²A standard complex normal random variable is given by $Z = X + iY$ where X and Y are independent real normal random variables with mean zero and variance $1/2$.

asymptotics of the smallest eigenvalue of W , and thus the condition number, are described by the Marcenko–Pastur distribution with parameter p . In particular, as $N \rightarrow \infty$ the smallest eigenvalue of $W(N)/(N + \alpha)$ remains strictly separated from 0. In recent work with P. Deift, we explored an intermediate regime $\alpha \sim 4cN^{1/2}$, and established Tracy–Widom [21] fluctuations of the smallest eigenvalue and the condition number (see [4, Theorems 1.1 and 1.3]). We further showed numerically that the nontrivial fluctuations of the condition number are reflected in the performance of the conjugate gradient algorithm on Wishart matrices in this regime. In this article, we broaden our exploration of this intermediate regime, choosing

$$\alpha = \alpha(N) = \lfloor \sqrt{4cN^\gamma} \rfloor \quad \text{for some } 0 < \gamma \leq 1/2. \quad (1.5)$$

In order to formulate a notion of smoothed analysis for the conjugate gradient algorithm, we must subject a deterministic positive-definite matrix A with $\|A\| \leq 1$ to a random perturbation of the form $\sigma^2 H$, where $\|H\| = O(1)$, and then take the supremum over all A with $\|A\| \leq 1$. It turns out that the largest eigenvalue of $W = XX^*$ is approximately $\nu = 4N + 2\alpha + 2$. Thus, our implementation of smoothed analysis for the conjugate gradient algorithm involves estimating

$$\sup_{A \geq 0, \|A\| \leq 1} \mathbb{E}[\tau_\epsilon(A + \sigma^2 H, b)], \quad \sup_{A \geq 0, \|A\| \leq 1} \mathbb{E}[\tau_{w,\epsilon}(A + \sigma^2 H, b)], \quad H = W/\nu,$$

with explicit dependence on σ and γ . The factor σ^2 is used here so that σ represents the scaling of the variance of the entries of X . Our main result, proved in Section 3.2, is the following.

Theorem 1.1. *Assume α satisfies (1.5) and $\epsilon > 0$. Let $H = \nu^{-1}XX^*$ where $\nu = 4N + 2\alpha + 2$ and X is an $N \times (N + \alpha)$ matrix of iid standard complex normal random variables. Then with*

$$\rho_\sigma = 2\sqrt{\left(1 + \frac{1}{\sigma^2}\right)\frac{1}{c}}, \quad \sigma > 0,$$

we have the following estimates.

(1) **Halting time with the ℓ^2 norm:**

$$\begin{aligned} & \sup_{\|A\| \leq 1, A \geq 0} \mathbb{E}[\tau_\epsilon(A + \sigma^2 H, b)^j] \\ & \leq \frac{1}{2^j} N^{j(1-\gamma)} \rho_\sigma^j (\log N^{1-\gamma} \rho_\sigma 2\epsilon^{-1})^j (1 + o(1)), \quad \text{as } N \rightarrow \infty. \end{aligned}$$

(2) **Halting time with the weighted norm:**

$$\sup_{\|A\| \leq 1, A \geq 0} \mathbb{E}[\tau_{w,\epsilon}(A + \sigma^2 H, b)^j] \leq \frac{1}{2^j} N^{j(1-\gamma)} \rho_\sigma^j (\log 2\epsilon^{-1})^j (1 + o(1)), \quad \text{as } N \rightarrow \infty.$$

(3) **Successive residuals:** For $r_k = r_k(A + \sigma^2 H, b)$, the k th residual in the solution of $(A + \sigma^2 H)x = b$ with the conjugate gradient algorithm

$$\sup_{\|A\| \leq 1, A \geq 0} \mathbb{E} \left[\frac{\|r_{k+1}\|^j}{\|r_k\|^j} \right] \leq \left(1 - \frac{2}{\rho_\sigma N^{1-\gamma} + 1} (1 + o(1)) \right)^j, \quad \text{as } N \rightarrow \infty,$$

where $\|\cdot\|$ is either $\|\cdot\|_{\ell^2}$ or $\|\cdot\|_{w^{-1}}$.

Remark 1.2. The parameters (γ, σ) control the effect of the random perturbation in very different ways. In Lemma 3.4 we precisely describe how increasing γ leads to better conditioned problems. For all $0 < \gamma \leq 1/2$ (and conjecturally for $\gamma > 1/2$) the asymptotic size of the expectation and the standard deviation of τ_ϵ is $\mathcal{O}(N^{1-\gamma} \log N)$, meaning that the conjugate gradient algorithm will terminate before its maximum of N iterations with high probability. For instance, assume α satisfies (1.5) with $c = 2$. We use Markov’s inequality and Theorem 1.1 for $j > 0$ and sufficiently large³ N to obtain

$$\begin{aligned} \sup_{\|A\| \leq 1, A \geq 0} \mathbb{P}(\tau_\epsilon(A + \sigma^2 H, b) \geq N^{1-\lambda}) \\ \leq \frac{1}{2^{j-2}} \left(1 + \frac{1}{\sigma^2}\right)^{j/2} N^{-j(\gamma-\lambda)} (\log [4N^{1-\gamma} (1 + \sigma^{-2})^{1/2} \epsilon^{-1}])^j. \end{aligned}$$

Hence for $\lambda < \gamma$ and j large, this probability decays rapidly.

Remark 1.3. We only prove Theorem 1.1 for LUE perturbations in the range $0 < \gamma \leq 1/2$. However, we expect Theorem 1.1 to hold for all $0 \leq \gamma \leq 1$, as illustrated in the numerical experiments below. In order to establish Theorem 1.1 in the range $\frac{1}{2} < \gamma \leq 1$ it is only necessary to establish Lemma 3.4 for these values of γ . This will be the focus of future work.

Remark 1.4. Theorem 1.1 provides asymptotic control on the j th moments of halting times for each j . This formally suggests that one may obtain a bound on an exponential generating function of the halting times above. However, we cannot establish this because the condition number κ has only $O(\alpha)$ moments at any finite N .

Remark 1.5. By restricting attention to positive definite perturbations we ensure that the conjugate gradient scheme is always well-defined for the perturbed matrix $A + \sigma^2 H$. This also allows the following simple, but crucial lower bound, on the lowest eigenvalue of the perturbed matrix

$$\lambda_{\min}(A + \sigma^2 H) \geq \lambda_{\min}(A),$$

which then yields an upper bound on the condition number of the perturbed matrix $\kappa(A + \sigma^2 H)$. We have not considered the question of random perturbations of A that are Hermitian, but not necessarily positive definite. Such perturbations are more subtle since they must be scaled according to the smallest eigenvalue of A . Nor have we considered the question of whether such perturbations provide good ‘real-life’ models of a smoothed analysis of the conjugate gradient scheme. Nevertheless, the above framework shares important features with [16] in that the problem is “easier” for large values of σ and the worst case of the supremum over the set $\{A \geq 0, \|A\| \leq 1\}$ can be realized at singular A .

1.4 Numerical simulations and the accuracy of the estimates

In this section we investigate how close our estimates on $\mathbb{E}[\tau_\epsilon(A + \sigma^2 H, b)]$ are to the true value of the expectation. We present numerical evidence that in the “ $\sigma = \infty$ ” (also obtained by $A = 0$ and $\sigma = 1$) case the estimates are better for larger values of γ , and continue to hold beyond the $\gamma = 1/2$ threshold of Theorem 1.1. We also give examples for specific choices of A and demonstrate that, as expected, the actual behavior of the conjugate gradient algorithm is much more complicated for $A \neq 0$.

Because the conjugate gradient algorithm is notoriously affected by round-off error, we adopt the following approach to simulating $\tau_\epsilon(M, b)$, $M = A + \sigma^2 H$ with finite-precision arithmetic:

³Here N should be sufficiently large so as to make the error term in Theorem 1.1 less than unity.

- In exact arithmetic, the conjugate gradient algorithm applied to $Mx = b$, $M = U^*\Lambda U$, with initial guess $x_0 = 0$, has the same residuals as the algorithm applied to $\Lambda y = U^*b$. Indeed, if x_k satisfies $\|Mx_k - b\|_{w^{-1}} = \min_{x \in \mathcal{K}_k} \|Mx - b\|_{w^{-1}}$ then for $y = U^*x$, $y_k = U^*x_k$, $\|Mx_k - b\|_{w^{-1}}^2 = \langle Mx_k - b, x_k - x \rangle = \langle \Lambda y_k - U^*b, y_k - y \rangle$. Thus, defining $\|\cdot\|_{\tilde{w}^{-1}} = \langle \cdot, \Lambda^{-1} \cdot \rangle$ we have

$$\begin{aligned} \|\Lambda y_k - U^*b\|_{\tilde{w}^{-1}} &= \min_{y \in \tilde{\mathcal{K}}_k} \|\Lambda y - U^*b\|_{\tilde{w}^{-1}}, \\ \tilde{\mathcal{K}}_k &= U^*\mathcal{K}_k = \text{span} \{U^*b, \Lambda U^*b, \dots, \Lambda^{k-1}U^*b\}. \end{aligned}$$

This is an exact characterization of the iterates of the conjugate gradient algorithm applied to $\Lambda y = U^*b$.

- Sample a matrix $H = XX^*/\nu$ and compute the spectral decomposition $A + \sigma^2 H = U\Lambda U^*$. Sample a vector b with iid Gaussian entries and normalize⁴ it, so that $\|b\|_{\ell^2} = 1$. Prior to normalization the entries of b are iid Gaussian, thus b is uniformly distributed on the unit sphere in (\mathbb{C}^N, ℓ^2) . Note that if $A = 0$, $M = \sigma^2 H$ is a Wishart matrix, we find that U^*b is also uniformly distributed on the unit sphere in (\mathbb{C}^N, ℓ^2) . That is, b and U^*b have the same law.
- Applying the diagonal matrix Λ to a vector is much less prone to round-off error since it involves only N multiplications, as opposed to N^2 multiplications for the dense matrix H . Thus, to minimize round-off error we compute the iterates of the conjugate gradient algorithm applied to $\Lambda y = b$ with Λ as above, and b uniformly distributed on the unit sphere in (\mathbb{C}^N, ℓ^2) . As noted above when $A = 0$, these iterates have the same law as those of $Hx = \tilde{b}$ when \tilde{b} and b have the same law, and when H is a Wishart matrix. By computing the number of iterations necessary (in high-precision arithmetic) so that $\|y_{k+1}\|_{\ell^2} \leq \epsilon$, we obtain one sample of the halting time $\tau_\epsilon(M, b)$ without significant round-off errors.

1.4.1 The “ $\sigma = \infty$ ” case

Now, we investigate how close our estimates on $\mathbb{E}[\tau_\epsilon(H, b)]$ (which can be obtained from Theorem 1.1 by formally sending $\sigma \rightarrow \infty$). In Figs. 1, 2 and 3 we plot the sample mean $\bar{\tau}_\epsilon$ over 1,000 samples as N increases. Throughout our numerical experiments b is taken to be iid uniform on $[-1, 1]$ and then normalized to be a unit vector. With this consideration, it is clear that the estimate in Theorem 1.1 is good for $\gamma = 2/3$ (despite the fact that we have not proved it holds in this case), fairly tight for $\gamma = 1/2$ and not as good for $\gamma = 1/3$. These calculations demonstrate that the worst case bounds (1.2) and (1.1) provide surprisingly good estimates in a random setting. Further, they appear to be exact in the sense that Theorem 1.1 predicts the correct order of the expectation of τ_ϵ as $N \rightarrow \infty$.

Comparing these numerical results with Theorem 1.1 we conclude that:

- (1) Tail estimates on the condition number derived from tail estimates of the extreme eigenvalues, can be used to obtain near optimal, and in some cases optimal, estimates for the expected moments of the condition number.
- (2) In light of rigorous results and heuristic expectations of universality in random matrix theory, we find it reasonable to expect Lemma 3.4 and Theorem 1.1 to hold for more general real and complex sample covariance matrices, not just LUE.
- (3) The worst-case estimates given in (1.1) and (1.2) produce effective bounds on the moments of the halting time, and predict the correct order of growth of the mean as $N \rightarrow \infty$. The

⁴Choosing b in this way is convenient for these manipulations but it is not necessary. We choose a non-Gaussian vector for our actual experiments.




Figure 1. (a) The sample mean $\bar{\tau}_\epsilon$ as a function of N for $\gamma = 1/2$ and $c = 1$ with $\epsilon = 10^{-4}$. The plot also shows the deterministic maximum of N iterations for the conjugate gradient algorithm (solid line), the upper bound computed in Theorem 1.1 (dashed line) and the curve $7.5N^{1/2}$ (dotted line) to demonstrate that $\bar{\tau}_\epsilon$ grows faster than $N^{1/2}$. (b) A fit of the data points using the function $F(N) = aN^{1/2} \log N + bN^{1/2}$, $a \geq 0$ and $b \geq 0$ plotted against the data. Here we find $a = 0.67$ and $b = 3.51$ indicating that Theorem 1.1 predicts the correct scaling in N for $\bar{\tau}_\epsilon$.




Figure 2. (a) The sample mean $\bar{\tau}_\epsilon$ as a function of N for $\gamma = 2/3$ and $c = 1$ with $\epsilon = 10^{-4}$. The plot includes the linear upper bound $\tau_\epsilon \leq N$ (solid line), the upper bound computed in Theorem 1.1 (dashed line) and the curve $10N^{1/3}$ (dotted line) to demonstrate that $\bar{\tau}_\epsilon$ grows at approximately $N^{1/3}$. (b) A fit of the data points using the function $F(N) = aN^{1/3} \log N + bN^{1/3}$, $a \geq 0$ and $b \geq 0$ plotted against the data. Here we find $a = 0.152$ and $b = 8.66$. Again, Theorem 1.1 predicts the correct scaling for $\bar{\tau}_\epsilon$.

importance of this observation is that these bounds are known to be sub-optimal. Thus, our results show that the matrices for which these estimates are sub-optimal have a small probability of occurrence.

1.4.2 Perturbed discrete Laplacian

The numerical examples of the previous section are dominated by noise. In this subsection and the next, we investigate the effect of small LUE perturbations on structured matrices A . This is a more subtle problem since it is hard to conjecture the growth rate of $\mathbb{E}[\tau_\epsilon(A + \sigma^2 H, b)]$ as σ and γ vary for a given A . We present numerical experiments on random perturbations of two examples that have been studied in the literature on the conjugate gradient algorithm –




Figure 3. (a) The sample mean $\bar{\tau}_\epsilon$ as a function of N for $\gamma = 1/3$ and $c = 1$ with $\epsilon = 10^{-4}$. The plot also shows the upper bound of N iterations (solid line), the upper bound computed in Theorem 1.1 (dashed line) and the curve $6N^{2/3}$ (dashed) line to demonstrate that $\bar{\tau}_\epsilon$ grows faster than $N^{2/3}$. (b) A fit of the data points using the function $F(N) = aN^{2/3} \log N + bN^{2/3}$, $a \geq 0$ and $b \geq 0$ plotted against the data. The parameters $a = 0.916$ and $b = 0$ fit the data very well.

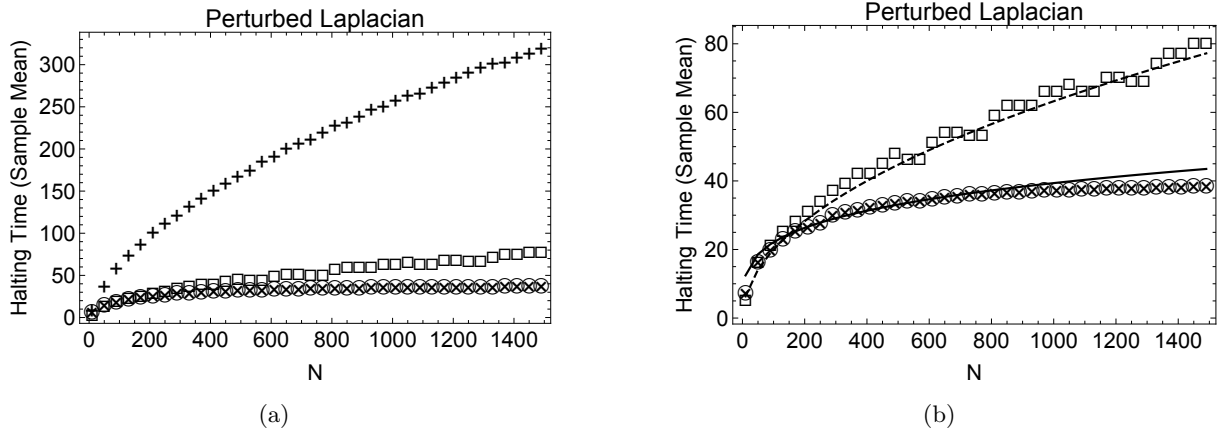


Figure 4. Numerical computations for sample mean of the halting time in the case of a randomly perturbed discrete Laplacian operator $-\Delta_{m,k} + \sigma^2 H$. We let $m = k = \lfloor \sqrt{N} \rfloor$ and make the following choices of parameter values $\sigma = 0$ (\square), $\sigma = 0.1$, $\gamma = 1/2$ (\circ), $\sigma = 0.1$, $\gamma = 1/3$ (\times) and $\sigma = \infty$, $\gamma = 1/2$ ($+$). (a) The sample mean for all four parameter choices. (b) Three parameter choices plotted along with the functions $2\sqrt{N}$ (dashed) and $7N^{1/4}$ (solid).

discrete Laplacians and singular matrices with clusters of eigenvalues. In both these examples, we numerically estimate the growth with N of the halting time $\tau_\epsilon(A + \sigma^2 H, b)$ for a range of σ and γ . These numerical computations are compared with the unperturbed ($\sigma = 0$) and noise dominated ($\sigma = \infty$) cases. Broadly, we observe that finite noise gives faster convergence (smaller halting time) with a different scaling than what is expected with no noise. We also find that when $A \neq 0$, the halting time is not strongly affected by γ . At present, these are numerical observations, not theorems. We hope to investigate the accelerated convergence provided by noise in future work.

In our first example, A is the $mk \times mk$ 2D discrete Laplacian defined by the Kronecker product

$$\Delta_{m,k} = I_m \otimes D_{2,k} + D_{2,m} \otimes I_k,$$

where $D_{2,m}$ is the $m \times m$ symmetric tridiagonal matrix with -2 on the diagonal and 1 on the off-diagonals. We choose $m = k = \lfloor \sqrt{N} \rfloor$ in the computations.

Some results of numerical experiments with this choice of A are shown in Fig. 4. The scaling of the sample mean of the halting time, $\bar{\tau}_\epsilon$, is $O(\sqrt{N})$ in the extreme cases when $\sigma = 0$ or $\sigma = \infty$ (see $+$ and \square in Fig. 4). However, when σ is $O(1)$, we find that $\bar{\tau}_\epsilon \sim N^{1/4}$ (see \circ and \times in Fig. 4). Further, this result is not sensitive to γ . Therefore there is a complicated relationship between the deterministic matrix A , the random perturbation H and the halting time that is not captured by Theorem 1.1.

1.4.3 Perturbed eigenvalue “clusters”

In our second example, we consider random perturbations of a singular matrix with clusters of eigenvalues. This construction is motivated by [11, Section 5.6.5] and [9].

We define A to be the $mk \times mk$ diagonal matrix obtained by sampling the Marchenko–Pastur law as follows⁵. Let $\zeta_{j,k}$, $j = 1, \dots, k$ be defined by

$$\zeta_j = \zeta_{j,k} = \min \left\{ 0 \leq t \leq 4 : \int_{-\infty}^t \frac{1}{2\pi} \sqrt{\frac{4-x}{x}} dx = \frac{j}{k} \right\}.$$

Then we define for $1 \leq j \leq k, 1 \leq \ell \leq m$

$$\lambda_{\ell,j} = \begin{cases} 0, & j = 0, \\ \zeta_j + \left(\frac{\ell - \lfloor m/2 \rfloor}{m} \right) \frac{1}{10k^2}, & j > 0. \end{cases} \quad (1.6)$$

Finally, set $A = M_{m,k} = \text{diag}(\lambda_{\ell,j})$ with any (consistent) ordering.

This produces a $mk \times mk$ diagonal matrix with m zero eigenvalues, and $k(m-1)$ eigenvalues that are each clustered at quantiles of the Marchenko–Pastur law. Note that

$$\frac{1}{k} = \int_{-\infty}^{\zeta_{1,k}} \frac{1}{2\pi} \sqrt{\frac{4-x}{x}} dx \sim c\zeta_{1,k}^{1/2}, \quad k \rightarrow \infty,$$

and so $\zeta_{1,k} = \mathcal{O}(k^{-2})$. We divide by k^2 in (1.6) to ensure we maintain a positive semi-definite matrix. As in the previous section we set $m = k = \lfloor \sqrt{N} \rfloor$ in the computations and plot similar sample mean results in Fig. 5. Despite the fact that $M_{m,k}$ is a singular matrix, the conjugate gradient algorithm converges rapidly for the perturbed matrix $M_{m,k} + \sigma^2 H$. In particular, Fig. 5(b) shows a rate of growth that is only $O(\sqrt{\log N})$.

Finally, in the construction of $M_{m,k}$ we imposed the condition that m eigenvalues are zero. If we considered a case where more and more eigenvalues are set to zero as $m, k \rightarrow \infty$ we would expect a transition to the $\sigma = \infty$ case.

2 Estimating the halting time

2.1 Outline of the proof

In this section, we explain the main steps in the proof of Theorem 1.1. We also abstract the properties that are known to hold for LUE perturbations (as established in Section 3), stating these estimates as a general condition on the tails of the smallest and largest eigenvalues that suffice to prove Theorem 1.1.

⁵We use the Marchenko–Pastur law because it gives the asymptotic density of eigenvalues of the LUE we are considering [12].




Figure 5. Numerical computations for sample mean of the halting time in the case of a randomly perturbed eigenvalue clusters: $M_{m,k} + \sigma^2 H$. Again, we let $m = k = \lfloor \sqrt{N} \rfloor$ and make the following choices of parameter values $\sigma = 0.1, \gamma = 1/2$ (\circ), $\sigma = 0.1, \gamma = 1/3$ (\times) and $\sigma = \infty, \gamma = 1/2$ ($+$). (a) The sample mean for all four parameter choices. (b) Three parameter choices plotted along with the function $6 \log^{1/2}(1+N) + 28.5$ to demonstrate how slow the halting time grows for $0 < \sigma < \infty$.

In order to explain the main idea, we focus on controlling the halting time $\tau_{w,\epsilon}$ using estimate (1.1). For brevity, let us define the parameter

$$\theta(\kappa) = \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}.$$

Since $\kappa \geq 1$, the parameter $\theta(\kappa) \in [0, 1)$. Let us also define the positive real number

$$K_\epsilon(\kappa) = \inf_{k>0} \{ \theta(\kappa)^k \leq \epsilon/2 \} = \frac{\log \epsilon/2}{\log \theta(\kappa)}.$$

It follows immediately from (1.3) and the normalization $\|r_0\|_{\ell^2} = 1$ that $\tau_{w,\epsilon} \leq K_\epsilon$, so that for every $a \geq 0$,

$$\mathbb{P}(\tau_{w,\epsilon} > a) \leq \mathbb{P}(K_\epsilon > a) = \mathbb{P}(\log \epsilon/2 < a \log \theta(\kappa)).$$

Note that $\log \theta(\kappa) < 0$ and that as $\kappa \rightarrow \infty$,

$$\log \theta(\kappa) \sim -\frac{2}{\sqrt{\kappa}}.$$

Thus, basic convergence properties of the conjugate gradient algorithm may be obtained from tail bounds on the condition number. Finally, the condition number is estimated as follows. Since $\kappa = \lambda_{\max}/\lambda_{\min}$ it is clear that upper bounds of the form $\mathbb{P}(\lambda_{\max} > t)$ and $\mathbb{P}(\lambda_{\min}^{-1} > t)$ for arbitrary $t \in (0, \infty)$ may be combined to yield an upper bound on $\mathbb{P}(\kappa > a)$ by suitably choosing $t(a)$. As noted in the first three lines of the proof of Theorem 1.1 below, estimates of upper and lower eigenvalues for Wishart ensembles established in [4] immediately extend to estimates for matrices of the form $A + \sigma^2 H$.

2.2 A general sufficient condition

The abstract property we use to establish Theorem 1.1 is the following.

Condition 2.1. *Given a random positive-definite matrix H , assume there exist positive constants c_1 and δ , constants C_1, C_2 , and a that are greater than 1, and a positive function $f: (0, \infty) \rightarrow (0, \infty)$ such that*

$$\begin{aligned} T_{\max}(t) &:= \mathbb{P}(\lambda_{\max}(H) > t) \leq \tilde{T}_{\max}(t) := C_1 e^{-c_1 N(t-a)}, & t \geq 1, \\ T_{\min}(t) &:= \mathbb{P}(\lambda_{\min}^{-1}(H) > t) \leq \tilde{T}_{\min}(t) := C_2 [t/f(N)]^{-\alpha/2}, & t \geq (1 + \delta)f(N). \end{aligned} \quad (2.1)$$

Assume further that T_{\max}/\min are strictly monotone functions of t and $\lim_{N \rightarrow \infty} f(N) = \infty$.

While the conditions above seem arbitrary at first sight, we will show how they emerge naturally for the LUE ensemble in the next section. In particular, we show that these conditions are satisfied by a class of LUE matrices in Lemmas 3.3 and 3.4.

Lemma 2.2. *Assume Condition 2.1 and that α grows with N as in (1.5). Then there exists a constant $C > 0$ such that*

$$\mathbb{P}(\kappa(H) > t) \leq C [a^{-1}t/f(N)]^{-\alpha/2 + e_N}, \quad e_N = \frac{1}{2} \frac{\alpha^2}{2ac_1N + \alpha},$$

when $t \geq a(1 + \delta_N)f(N)$ where

$$\delta_N = (1 + \delta)(1 + C(N^{-1} + \delta N^{\gamma-1})) - 1.$$

Proof. First, if $xy > ab$ and all these numbers are positive, then either $x > a$ or $y > b$. Thus for $0 \leq s \leq 1$, the tails bounds of Condition 2.1 imply

$$\mathbb{P}\left(\frac{\lambda_{\max}(H)}{\lambda_{\min}(H)} > T_{\max}^{-1}(s)T_{\min}^{-1}(s)\right) \leq 2s.$$

This bound may be ‘inverted’ in the following way. If we define $T^{-1}(s) = T_{\max}^{-1}(s)T_{\min}^{-1}(s)$, then

$$\mathbb{P}\left(\frac{\lambda_{\max}(H)}{\lambda_{\min}(H)} > t\right) \leq 2T(t).$$

Our goal is to obtain an upper bound on $T(t)$ using the upper bounds \tilde{T}_{\max} and \tilde{T}_{\min} . If $T_{\max}(t_1) = s$ and $\tilde{T}_{\max}(t_2) = s$ then $t_2 \geq t_1$. Therefore,

$$\begin{aligned} T_{\max}^{-1}(s) &\leq a + \frac{\log C_1/s}{c_1 N}, & s \leq T_{\max}(a), \\ T_{\min}^{-1}(s) &\leq f(N) \left[\frac{s}{C_2}\right]^{-2/\alpha}, & s \leq T_{\min}((1 + \delta)f(N)). \end{aligned}$$

Since $(1 + \log t/n)^n \leq t$ when $t > 1$ and $n \geq 1$, we can estimate

$$a + \frac{\log C_1/s}{c_1 N} = a \left(1 + \frac{\log C_1/s}{ac_1 N}\right) \leq a \left(\frac{C_1}{s}\right)^{1/(ac_1 N)}.$$

Hence

$$T^{-1}(s) \leq a C_1^{1/(ac_1 N)} f(N) \left[\frac{s}{C_2}\right]^{-2/\alpha} s^{-1/(ac_1 N)}, \quad s \leq \min\{T_{\max}(a), T_{\min}((1 + \delta)f(N))\}.$$

Inverting this expression, we find

$$\begin{aligned} T(t) &\leq \left[a^{-1}t/f(N)C_1^{-1/(ac_1 N)}C_2^{-2/\alpha}\right]^{-\alpha/2 + e_N}, & e_N = \frac{1}{2} \frac{\alpha^2}{2ac_1 N + \alpha}, \\ t &\geq \max\{T^{-1}(T_{\max}(a)), T^{-1}(T_{\min}((1 + \delta)f(N)))\}. \end{aligned}$$

Let us examine this lower bound more carefully. We increase C_1 so that $C_1 \leq C_2$, if necessary, so that

$$\begin{aligned} T^{-1}(T_{\max}(a)) &= aT_{\min}^{-1}(T_{\max}(a)) \leq af(N) \left[\frac{C_2}{C_1} \right]^{2/\alpha} \leq af(N), \\ T^{-1}(T_{\min}((1+\delta)f(N))) &= (1+\delta)f(N)T_{\max}^{-1}(T_{\min}((1+\delta)f(N))) \\ &\leq a(1+\delta)f(N) \left(1 + \frac{\log C_1/C_2 + \alpha/2 \log(1+\delta)}{ac_1N} \right) \\ &\leq a(1+\delta)(1 + CN^{-1}(1+\alpha\delta))f(N), \end{aligned}$$

where C is a suitable constant. Then using the assumption (1.5) and $C_2 \geq 1$, $C_1^{\alpha/(2ac_1N)} C_2^{1-2e_N/\alpha}$ is bounded by a constant, say, $C/2 > 0$. This establishes the lemma. \blacksquare

Let Σ_N denote the set of $N \times N$ strictly positive definite complex matrices and recall that the constant δ_N is defined in Lemma 2.2. The following lemma is applied to control the halting time in terms of the condition number, and the reader may turn to the lemmas that follow to see instances of functions g .

Lemma 2.3. *Let $g: [1, \infty) \rightarrow \mathbb{R}$ be continuous and differentiable on $(1, \infty)$. Assume g satisfies $g(1) = 0$ and $g'(x) \leq Cx^\ell$ for $C > 0$, $\ell \in \mathbb{R}$ and x sufficiently large. Assume a function $M: \Sigma_N \rightarrow \mathbb{R}$ satisfies $M(H) \leq g(\kappa(H))$. Then if H satisfies Condition 2.1 and $\alpha/2 - e_N - \ell > 1$ there exist constants $C, K > 0$ such that*

$$\mathbb{E}[M(H)] \leq g(b_N) + C(1 + \delta_N)^{-\alpha/K} \frac{b_N^{1+\ell}}{\alpha/2 - e_N - \ell - 1},$$

for $b_N = af(N)(1 + \delta_N)$.

Proof. First,

$$\mathbb{E}[M(H)] \leq \mathbb{E}[g(\kappa(H))],$$

and by integration by parts

$$\mathbb{E}[g(\kappa(H))] = \int_1^\infty g(s) d\mathbb{P}(\kappa(H) \leq s) = \int_1^\infty g'(s) \mathbb{P}(\kappa(H) > s) ds.$$

Using Lemma 2.2 for $b_N = af(N)(1 + \delta_N)$ and N sufficiently large

$$\begin{aligned} \int_1^\infty g'(s) \mathbb{P}(\kappa(H) > s) ds &\leq \int_1^{b_N} g'(s) ds + C \left(\frac{b_N}{1 + \delta_N} \right)^{-\alpha/2 + e_N} \int_{b_N}^\infty s^{-\alpha/2 + e_N + \ell} ds \\ &\leq g(b_N) + C(1 + \delta_N)^{-\alpha/K} \frac{b_N^{1+\ell}}{\alpha/2 - e_N - \ell - 1}. \end{aligned}$$

This last inequality follows from the scaling (1.5):

$$e_N = \frac{\alpha}{2} \frac{1}{2ac_2\alpha^{-1}N + 1} = \mathcal{O}(\alpha N^{\gamma-1}),$$

and hence $\alpha/2 - e_N = O(\alpha)$. For any fixed ℓ there exists a constant $K = K_\ell$

$$\frac{1}{\alpha/2 - e_N - \ell - 1} \leq K/\alpha. \tag{2.2}$$

\blacksquare

We apply this lemma to the following functions.

Lemma 2.4. *Let b be a fixed vector then for any $j > 0$*

(1) **Halting time with the ℓ^2 norm:** $\tau_\epsilon(H, b)^j \leq g(\kappa(H))^j$ where

$$g(s) = \frac{\log \sqrt{s} 2\epsilon^{-1}}{\log \left(\frac{\sqrt{s+1}}{\sqrt{s-1}} \right)}.$$

Further, for every $\eta > 0$ and $\epsilon > 0$ there exists a constant $C_{\epsilon, \eta}$ such that for $s \in [1, \infty)$

$$g(s) \leq \frac{1}{2} \sqrt{s} \log \sqrt{s} 2\epsilon^{-1} \leq C_{\epsilon, \eta} s^{1/2+\eta}, \quad g'(s) \leq \frac{1}{2} \frac{\log \sqrt{s} 2\epsilon^{-1}}{\sqrt{s}} \leq C_{\epsilon, \eta} s^{-1/2+\eta}.$$

(2) **Halting time with the weighted norm:** $\tau_{w, \epsilon}(H, b)^j \leq g(\kappa(H))^j$ where

$$g(s) = \frac{\log 2\epsilon^{-1}}{\log \left(\frac{\sqrt{s+1}}{\sqrt{s-1}} \right)}.$$

This function g satisfies the following estimates

$$g(s) \leq \frac{1}{2} \sqrt{s} \log 2\epsilon^{-1}, \quad g'(s) \leq \frac{1}{2} \frac{\log 2\epsilon^{-1}}{\sqrt{s}}.$$

(3) **Successive residuals:** For $r_k = r_k(A, b)$, $\left(\frac{\|r_{k+1}\|}{\|r_k\|} \right)^j \leq g(\kappa(H))^j$ where

$$g(s) = \left(\frac{\sqrt{s} - 1}{\sqrt{s} + 1} \right) \leq 1, \quad g'(s) \leq s^{-3/2},$$

and $\|\cdot\|$ stands for either $\|\cdot\|_{\ell^2}$ or $\|\cdot\|_A$.

Proof. All the bounds follow from (1.1) and (1.2) as explained in the introduction to this section. The estimates on the functions $g(s)$, each of which satisfies $g(1) = 0$, may be obtained by elementary manipulations. \blacksquare

We can now prove our generalized result.

Theorem 2.5. *Assume a random matrix H satisfies Condition 2.1 and α satisfies (1.5). Then for $b_N = af(N)(1 + \delta_N)$, $\delta_N > 0$ and, any vector $b = 1$, the following estimates hold:*

(1) **Halting time with the ℓ^2 norm:**

$$\mathbb{E}[\tau_\epsilon(H, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log b_N^{1/2} \epsilon^{-1})^j (1 + o(b_N^{-1/2} (1 + \delta_N)^{-\alpha/K})), \quad \text{as } N \rightarrow \infty.$$

(2) **Halting time with the weighted norm:**

$$\mathbb{E}[\tau_{w, \epsilon}(H, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log \epsilon^{-1})^j (1 + O(b_N^{-1/2} N^{-\gamma} (1 + \delta_N)^{-\alpha/K})), \quad \text{as } N \rightarrow \infty.$$

(3) **Successive residuals:** For $r_k = r_k(H, b)$

$$\mathbb{E} \left[\frac{\|r_{k+1}\|^j}{\|r_k\|^j} \right] \leq \left(1 - \frac{2}{\sqrt{b_N} + 1} \right)^j + O(b_N^{-1/2} N^{-\gamma} (1 + \delta_N)^{-\alpha/K}), \quad \text{as } N \rightarrow \infty,$$

where $\|\cdot\|$ stands for either $\|\cdot\|_{\ell^2}$ or $\|\cdot\|_{w^{-1}}$.

Proof. Before we begin, we recall (2.2).

(1) **Halting time with the ℓ^2 norm:** As H satisfies Condition 2.1 we can apply Lemma 2.3 with the estimates in Lemma 2.4(1) for $j > 0$ and $\eta > 0$. We use that in this case

$$\frac{d}{ds}g(s)^j = jg'(s)g(s)^{j-1} \leq jC_{\epsilon,\eta}^j s^{(j-3)/2+j\eta},$$

and hence $\ell = (j-3)/2 + j\eta$ in Lemma 2.3. Therefore,

$$\mathbb{E}[\tau_\epsilon(H, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log b_N^{1/2} \epsilon^{-1})^j + jC_\ell C_{\epsilon,\eta}^j (1 + \delta_N)^{-\alpha/K} b_N^{j/2} b_N^{-1/2+j\eta} N^{-\gamma}.$$

If $b_N = \mathcal{O}(N^\zeta)$ for some $\zeta > 0$, we choose $\eta > 0$ such that $j\zeta\eta < \gamma$. Thus,

$$\mathbb{E}[\tau_\epsilon(H, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log b_N^{1/2} \epsilon^{-1})^j (1 + o(b_N^{-1/2} (1 + \delta_N)^{-\alpha/K})), \quad \text{as } N \rightarrow \infty.$$

(2) **Halting time with the weighted norm:** We follow the same calculations as (1) with the estimates in Lemma 2.4(2) for $j > 0$. Here

$$\frac{d}{ds}g(s)^j = jg'(s)g(s)^{j-1} \leq \frac{j}{2^{j+1}} s^{(j-3)/2} (\log \epsilon^{-1})^j,$$

and hence $\ell = (j-3)/2$ in Lemma 2.3. Therefore,

$$\mathbb{E}[\tau_{w,\epsilon}(H, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log \epsilon^{-1})^j + jC_\ell \frac{1}{2^{j+1}} (1 + \delta_N)^{-\alpha/K} b_N^{j/2} b_N^{-1/2} N^{-\gamma}.$$

Finally,

$$\mathbb{E}[\tau_{w,\epsilon}(H, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log \epsilon^{-1})^j (1 + \mathcal{O}(b_N^{-1/2} N^{-\gamma} (1 + \delta_N)^{-\alpha/K})), \quad \text{as } N \rightarrow \infty.$$

(3) **Successive residuals:** We follow the same calculations as (1), (2) with the estimates in Lemma 2.4(3) for $j > 0$. Here

$$\frac{d}{ds}g(s)^j = jg'(s)g(s)^{j-1} \leq js^{-3/2}$$

and hence $\ell = -3/2$ in Lemma 2.3. Then with $R_k = \|r_{k+1}\|/\|r_k\|$

$$\mathbb{E}[R_k^j] \leq \left(1 - \frac{2}{\sqrt{b_N} + 1}\right)^j + jC_\ell b_N^{-1/2} N^{-\gamma} (1 + \delta_N)^{-\alpha/K}.$$

Therefore

$$\mathbb{E}[R_k^j] \leq \left(1 - \frac{2}{\sqrt{b_N} + 1}\right)^j + \mathcal{O}(b_N^{-1/2} N^{-\gamma} (1 + \delta_N)^{-\alpha/K}), \quad \text{as } N \rightarrow \infty. \quad \blacksquare$$

The constant $\delta_N > 0$ is used in the above theorem to make precise the fact that if we integrate the tail of the condition number distribution just beyond $af(N)$ the error term is exponentially small as $\alpha \rightarrow \infty$.

3 The Laguerre unitary ensemble

The following is well-known and may be found in [6, Section 2], for example. This discussion is modified from [4, Section 2]. Let $W = XX^*$ where X is an $N \times (N + \alpha)$ matrix of iid standard complex Gaussian random variables. Recall that the (matrix-valued) random variable W is the Laguerre unitary ensemble (LUE). Then it is known that the eigenvalues $0 \leq \lambda_{\min} = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N = \lambda_{\max}$ of W have the joint probability density

$$p_N(\lambda_1, \dots, \lambda_N) = \frac{1}{C_N^{(\alpha)}} \prod_{j=1}^N \lambda_j^\alpha e^{-\lambda_j} \prod_{1 \leq j < k \leq N} |\lambda_j - \lambda_k|^2.$$

Recall that the Laguerre polynomials, $\{L_j^{(\alpha)}(x)\}_{j=0}^\infty$, are a family of orthogonal polynomials on $[0, \infty)$, orthogonal with respect to the weight $e^{-x}x^\alpha$. We normalize them as follows [13]

$$L_j^{(\alpha)}(x) = k_j x^j + \mathcal{O}(x^{j-1}), \quad k_j = \frac{(-1)^j}{j!},$$

$$\int_0^\infty L_i^{(\alpha)}(x) L_j^{(\alpha)}(x) e^{-x} x^\alpha dx = \delta_{ij} \frac{\Gamma(j + \alpha + 1)}{j!}.$$

Then the following are orthonormal with respect to Lebesgue measure on $[0, \infty)$,

$$\psi_j(x) := \left(\frac{j!}{\Gamma(j + \alpha + 1)} \right)^{1/2} e^{-x/2} x^{\alpha/2} L_j^{(\alpha)}(x), \quad \int_0^\infty \psi_j(x) \psi_i(x) dx = \delta_{ij}.$$

Define the correlation kernel

$$\mathcal{K}_N(x, y) = \sum_{j=0}^{N-1} \psi_j(x) \psi_j(y), \quad 0 < x, y < \infty.$$

The kernel \mathcal{K}_N defines a positive, finite-rank and hence trace-class operator on $L^2([a, b])$. To see that \mathcal{K}_N is positive, consider $f \in C^\infty((s, t))$ with compact support and note that

$$\begin{aligned} \int_s^t \int_s^t \mathcal{K}_N(x, y) f(x) f^*(y) dx dy &= \int_s^t \int_s^t \sum_{j=0}^{N-1} \psi_j(x) \psi_j(y) f(x) f^*(y) dx dy \\ &= \sum_{j=0}^{N-1} \left(\int_s^t \psi_j(x) f(x) dx \right) \left(\int_s^t \psi_j(x) f^*(x) dx \right) = \sum_{j=0}^{N-1} \left| \int_s^t \psi_j(x) f(x) dx \right|^2. \end{aligned}$$

The eigenvalues $\lambda_1 \leq \dots \leq \lambda_N$ may be described in terms of Fredholm determinants of the kernel \mathcal{K}_N [1, 6]. In particular, the statistics of the extreme eigenvalues are recovered from the determinantal formula

$$\mathbb{P}(\text{no eigenvalues in } [a, b]) = \det(I - \mathcal{K}_N|_{L^2([a, b])}).$$

By the Christoffel–Darboux formula [20], we may also write

$$\begin{aligned} \mathcal{K}_N(x, y) &= \frac{N!}{\Gamma(N + \alpha)} \left(\frac{\Gamma(N + \alpha + 1) \Gamma(N + \alpha)}{N! (N - 1)!} \right)^{1/2} \left(\frac{\psi_{N-1}(x) \psi_N(y) - \psi_N(x) \psi_{N-1}(y)}{x - y} \right) \\ &= \frac{N!}{\Gamma(N + \alpha)} e^{-(x+y)/2} x^{\alpha/2} y^{\alpha/2} \frac{L_N^{(\alpha)}(y) L_{N-1}^{(\alpha)}(x) - L_N^{(\alpha)}(x) L_{N-1}^{(\alpha)}(y)}{x - y}. \end{aligned}$$

Thus, questions about the asymptotic behavior of $\mathcal{K}_N(x, y)$ as $N \rightarrow \infty$ reduce to the study of the large N asymptotics of $L_N^{(\alpha)}$ and $L_{N-1}^{(\alpha)}$.

3.1 Kernel estimates

We use Fredholm determinants to show that Condition 2.1 holds with appropriate constants when $W = XX^*$ is distributed according to LUE. The main reference for these ideas is [17]. Let $A: L^2([t, \infty)) \rightarrow L^2([t, \infty))$ be a positive trace-class operator with kernel $\mathcal{K}(x, y)$. Assume

$$\mathbb{P}(X \leq t) = \det(I - A_t),$$

then

$$\begin{aligned} \mathbb{P}(X > t) &= |1 - \det(I - A_t)| = |\det(I) - \det(I - A_t)| \\ &\leq \left(\int_t^\infty |\mathcal{K}(x, x)| dx \right) \exp \left(1 + \int_t^\infty |\mathcal{K}(x, x)| dx \right). \end{aligned}$$

In this way we can get estimates on the tail directly from the large x behavior of $\mathcal{K}(x, x)$. Similar considerations follow if, say, $A_t: L^2([0, t]) \rightarrow L^2([0, t])$.

As was done in [4], we consider the scaled kernel

$$\begin{aligned} \mathcal{K}_N^s(x, y) &= \sum_{j=0}^{N-1} \mathfrak{z}_j \pi_j(x) \pi_j(y) x^{\alpha/2} y^{\alpha/2} e^{-\nu(x+y)/2}, & \mathfrak{z}_j^{-1} &= \int_0^\infty \pi_j^2(x) x^\alpha y^{\alpha/2} e^{-\nu x} dx, \\ \pi_j(x) &= \frac{L_j^{(\alpha)}(\nu x)}{\nu^j k_j} = x^j + \mathcal{O}(x^{j-1}), \end{aligned}$$

so that

$$\begin{aligned} \mathbb{P}(\lambda_{\max}(W)/\nu \leq t) &= \det(1 - \mathcal{K}_N^s|_{L^2(t, \infty)}), \\ \mathbb{P}(\lambda_{\min}(W)/\nu > t) &= \det(1 - \mathcal{K}_N^s|_{L^2(0, t)}). \end{aligned}$$

Next, we pull results from [4] to estimate the kernel $\mathcal{K}_N^s(x, y)$ for LUE near the largest and smallest eigenvalue of W . We first look for the asymptotics of $\mathcal{K}_N^s(x, y)$ for $(x, y) \approx (1, 1)$, called the soft edge. Let $\check{x} = 1 + x/(2^{2/3}M^{2/3})$ and define

$$\check{\mathcal{K}}_N(x, y) = \mathcal{K}_N^s(\check{x}, \check{y}) \frac{1}{2^{2/3}M^{2/3}}, \quad M = N + \frac{1}{2}(\alpha + 1).$$

Then from [4, Proposition 2]:

Proposition 3.1. *As $N \rightarrow \infty$ the rescaled kernels converge pointwise,*

$$\check{\mathcal{K}}_N(x, y) \rightarrow \frac{\text{Ai}(x) \text{Ai}'(y) - \text{Ai}'(x) \text{Ai}(y)}{x - y}, \quad (x, y) \in \mathbb{R}^2,$$

and the convergence is uniform for (x, y) in a compact subset of $[L, \infty)^2$ for any $L \in \mathbb{R}$. If $x = y$ then the limit is determined by continuity. Further, there exists a positive, piecewise-continuous function $\bar{G}: (L, \infty)^2 \rightarrow (0, \infty)$, such that

$$|\check{\mathcal{K}}_N(x, y)| \leq \bar{G}(x, y), \quad \int_L^\infty \int_L^\infty \bar{G}(x, y) dx dy < \infty, \quad \int_L^\infty \bar{G}(x, x) dx < \infty.$$

Furthermore, it suffices to take for a constant $C = C(L) > 0$

$$\bar{G}(x, y) = C(\bar{g}(x)^2 \chi_{|x-y|<1}(x, y) + \bar{g}(x)\bar{g}(y)), \quad \bar{g}(x) = \begin{cases} 1, & \text{if } x < 0, \\ e^{-\frac{1}{3}x^{3/2}}, & \text{if } x \geq 0. \end{cases}$$

For $\mathcal{K}_N^s(x, y)$ near $(0, 0)$ as the scaling of the kernel depends critically on γ . So, we define

$$\hat{x} = \frac{\alpha^2}{\nu^2} \left(1 + x \left(\frac{2}{\alpha} \right)^{2/3} \right)$$

and

$$\hat{\mathcal{K}}_N(x, y) = \mathcal{K}_N^s(\hat{x}, \hat{y}) \frac{2^{2/3} \alpha^{4/3}}{\nu^2}.$$

The next proposition essentially follows directly from [4, Proposition 1] and is in fact a little simpler with the scaling chosen here.

Proposition 3.2. *As $N \rightarrow \infty$ the rescaled kernels converge pointwise,*

$$\hat{\mathcal{K}}_N(x, y) \rightarrow \frac{\text{Ai}(-y) \text{Ai}'(-x) - \text{Ai}'(-y) \text{Ai}(-x)}{x - y}, \quad (x, y) \in \mathbb{R}^2,$$

and the convergence is uniform for (x, y) in any compact subset of $(-\infty, L]^2$ for any $L \in \mathbb{R}$. If $x = y$ then the limit is determined by continuity. Further, there exists a positive, piecewise-continuous function $G: (-\infty, L)^2 \rightarrow (0, \infty)$, such that

$$|\hat{\mathcal{K}}_N(x, y)| \leq G(x, y), \quad \int_{-\infty}^L \int_{-\infty}^L G(x, y) dx dy < \infty, \quad \int_{-\infty}^L G(x, x) dx < \infty.$$

Furthermore, it suffices to take for a constant $C = C(L) > 0$

$$G(x, y) = C(\check{g}(x)\check{h}(x)\chi_{|x-y|<1} + \check{g}(x)\check{g}(y)),$$

where

$$\check{g}(x) = \begin{cases} 0, & \text{if } -\infty < x \leq -(\alpha/2)^{2/3}, \\ \left[\left(1 + x \left(\frac{2}{\alpha} \right)^{2/3} \right) \right]^{d\alpha/4}, & \text{if } -(\alpha/2)^{2/3} \leq x \leq (\ell(d)^2 - 1)(\alpha/2)^{2/3}, \\ e^{-\frac{1}{6}|x+1|^{3/2}}, & \text{if } -b(\alpha/2)^{2/3} < x \leq -1, \\ 1, & \text{otherwise,} \end{cases}$$

$$\check{h}(x) = \begin{cases} 0, & \text{if } -\infty < x \leq -(\alpha/2)^{2/3}, \\ \left[\left(1 + x \left(\frac{2}{\alpha} \right)^{2/3} \right) \right]^{d\alpha/4-1}, & \text{if } -(\alpha/2)^{2/3} \leq x \leq (\ell(d)^2 - 1)(\alpha/2)^{2/3}, \\ \frac{1}{1 - \ell(d)^2} e^{-\frac{1}{6}|x+1|^{3/2}}, & \text{if } -b(\alpha/2)^{2/3} < x \leq -1, \\ 1, & \text{otherwise.} \end{cases}$$

Here $\ell(d) < 1$ satisfies $\ell(d) = 1 + \mathcal{O}(d^{2/3})$ as $d \rightarrow 0$.

We now briefly describe how the estimates in terms of \check{g} and \check{h} arise. The asymptotics of the kernel \mathcal{K}_N^s is given in terms of Bessel functions, after a change of variables. In the regime

$\alpha \rightarrow \infty$, the Bessel functions asymptote to Airy functions, as follows [13]

$$\begin{aligned} J_\alpha(\alpha t) &= \left(\frac{4\zeta}{1-t^2} \right)^{1/4} \alpha^{-1/3} (\text{Ai}(\alpha^{2/3}\zeta) + \text{Ai}'(\alpha^{2/3}\zeta)\mathcal{O}(\alpha^{-4/3})), \quad t > 0, \\ J_\alpha'(\alpha t) &= -\frac{2}{t} \left(\frac{4\zeta}{1-t^2} \right)^{-1/4} \alpha^{-2/3} (\text{Ai}'(\alpha^{2/3}\zeta) + \text{Ai}(\alpha^{2/3}\zeta)\mathcal{O}(\alpha^{-2/3})), \quad t > 0, \\ \frac{2}{3}\zeta^{3/2} &= \int_t^1 \frac{\sqrt{1-s^2}}{s} ds, \quad 0 < t \leq 1, \\ \frac{2}{3}(-\zeta)^{3/2} &= \int_1^t \frac{\sqrt{s^2-1}}{s} ds, \quad t > 1. \end{aligned}$$

This expansion is uniform for $t \in (0, \infty)$. Assume $z \in (0, \delta')$ where $\delta' < \delta < 1/2$ and t is given in terms of z by $t = -i\frac{\nu}{4\alpha}\phi_\rightarrow^+(z)$. The following are from [4]

$$\begin{aligned} -i\phi_\rightarrow^+(z) &= 2 \int_0^z \sqrt{\frac{1-s}{s}} ds \leq 4\sqrt{z}, \quad 0 \leq z \leq \delta', \\ -i\phi_\rightarrow^+(z) &\geq 4\sqrt{z}|1-\delta|^{1/2} \geq 2\sqrt{z}, \quad 0 \leq z \leq \delta', \\ \frac{2}{3}\zeta^{3/2} &= \int_t^1 \sqrt{1-s} \frac{\sqrt{1+s}}{s} ds \geq \frac{4}{3}(1-t)^{3/2}, \quad 0 \leq t \leq 1, \quad \zeta \geq 2^{1/3}(1-t), \quad 0 \leq t \leq 1, \\ \frac{2}{3}\zeta^{3/2} &\geq \int_t^1 \frac{\sqrt{1-s}}{s} ds = -\log 2t + \log 2 + 2 \log(1 + \sqrt{1-t}) - 2\sqrt{1-t}, \quad 0 \leq t \leq 1, \\ \zeta &\geq \left(\frac{3}{2} \right)^{2/3} |\log 2t|^{2/3}, \quad 0 \leq t \leq 1/2. \end{aligned}$$

A subtle issue is the validity of the last bound. We see that $-i\phi_\rightarrow^+(z) \leq 4\sqrt{z}$, and so $t(z) \leq \frac{\nu}{\alpha}\sqrt{z}$ and $t(\tilde{z}) \leq (1 + z(2/\alpha)^{2/3})^{1/2}$. Then considering Lemma 2.3, we see that the dominant contribution arises from the interval over which the estimate $\mathbb{P}(\kappa(H) > s) \leq 1$ is used. Thus, we try to extend the validity of a lower bound on ζ to $t \in [0, 1]$. It follows that

$$\zeta \geq d^{2/3} |\log t|^{2/3}, \quad 0 \leq t \leq \ell(d) < 1.$$

Note that $\ell(d) \neq 1$ as $\zeta'(t)$ has a bounded derivative at $t = 0$ and the right-hand side does not. But as $d \rightarrow 0$, $\ell(d) \rightarrow 1$. A quick calculation, using an expansion near $t = 1$ gives

$$\zeta(t) = a(1 - \ell(d))(1 + \mathcal{O}(1 - \ell(d))) = d^{2/3}(-\log \ell(d))^{2/3},$$

and this implies:

$$\begin{aligned} [1 - \ell(d)]/d^{2/3} &\rightarrow 0, \quad \ell(d) = 1 + o(d^{2/3}), \\ [1 - \ell(d)]/d^{2/3+\epsilon} &\rightarrow \infty, \quad \epsilon > 0. \end{aligned} \tag{3.1}$$

Then, following [4, equations (C.3) and (C.4)],

$$\begin{aligned} e^{-\alpha \frac{1}{3} |\zeta|^{3/2}} &\leq t^{d\alpha/2}, \quad 0 \leq t \leq \ell(d), \\ e^{-\alpha \frac{1}{3} |\zeta|^{3/2}} &\leq \left(\frac{\nu}{\alpha} \sqrt{z} \right)^{d\alpha/2}, \quad 0 \leq z \leq \ell(d)^2 \frac{\alpha^2}{\nu^2}. \end{aligned}$$

This last inequality implies that $t \leq \ell(d)$. Then

$$\begin{aligned} 0 &\leq \tilde{z} \leq \ell(d)^2 \frac{\alpha^2}{\nu^2}, \\ 0 &\leq 1 + z(2/\alpha)^{2/3} \leq \ell(d)^2, \\ -(\alpha/2)^{2/3} &\leq z \leq (\ell(d)^2 - 1)(\alpha/2)^{2/3}. \end{aligned}$$

These estimates can then be plugged into [4, Lemma C.2] to get the estimates in Lemma 3.2.

3.2 Tail bounds

It follows that

$$\mathbb{P}(\lambda_{\max}(W)/\nu > t) \leq \left(\int_t^\infty |\mathcal{K}_N^s(x, x)| dx \right) \exp \left(1 + \int_t^\infty |\mathcal{K}_N^s(x, x)| dx \right).$$

So, we estimate for $t \geq 1$ and $C \geq 1$

$$\begin{aligned} \int_t^\infty |\mathcal{K}_N^s(x, x)| dx &\leq 2^{5/3} M^{2/3} C \int_t^\infty \bar{g}(2^{2/3} M^{2/3} (x-1))^2 dx \\ &= 2C \int_{2^{2/3} M^{2/3} (t-1)}^\infty \bar{g}^2(s) ds \leq 2C e^{-\frac{4}{3} M (t-1)}. \end{aligned}$$

So, for a new constant C

$$T_{\max}(t) = \mathbb{P}(\lambda_{\max}(W)/\nu > t) \leq C e^{-\frac{4}{3} M (t-1)}.$$

The more delicate estimate is to consider T_{\min} :

$$\begin{aligned} T_{\min}(t) &= \mathbb{P}(\nu \lambda_{\min}^{-1}(W) > t) = \mathbb{P}(\nu^{-1} \lambda_{\min}(W) < t^{-1}) = 1 - \mathbb{P}(\nu^{-1} \lambda_{\min}(W) \geq t^{-1}) \\ &\leq \left(\int_0^{t^{-1}} |\mathcal{K}_N^s(x, x)| dx \right) \exp \left(1 + \int_0^{t^{-1}} |\mathcal{K}_N^s(x, x)| dx \right). \end{aligned}$$

We use Proposition 3.2 and invert the scaling \hat{x} . If \hat{x} lies in $[0, \alpha^2/\nu^2]$ then $x \in (-\infty, 1]$. We note that $\check{g}(x)^2 \leq \check{g}(x)\check{h}(x)$ so we only need to estimate

$$\check{g} \left(\left(\frac{\alpha}{2} \right)^{2/3} \left(\frac{\nu^2}{\alpha^2} x - 1 \right) \right) \check{h} \left(\left(\frac{\alpha}{2} \right)^{2/3} \left(\frac{\nu^2}{\alpha^2} x - 1 \right) \right) \leq \left[\frac{\nu^2}{\alpha^2} x \right]^{d\alpha/2} x^{-1},$$

for $0 \leq x \leq \ell(d)^2 \frac{\alpha^2}{\nu^2}$. For $t \leq \ell(d)^{-2} \frac{\nu^2}{\alpha^2}$ we just use $T_{\min}(t) \leq 1$. For $t \geq \ell(d)^{-2} \frac{\nu^2}{\alpha^2}$ and $\alpha > 0$

$$\begin{aligned} \int_0^{t^{-1}} |\mathcal{K}_N^s(x, x)| dx &\leq 2C \frac{\nu^2}{\alpha^2} \alpha^{2/3} \int_0^{t^{-1}} \left[\frac{\nu^2}{\alpha^2} x \right]^{d\alpha/2} x^{-1} dx \\ &= 2C \frac{\alpha^{2/3}}{d\alpha/2} \left[\frac{\nu^2}{\alpha^2} \right]^{d\alpha/2+1} t^{-d\alpha/2} \leq C d^{-1} \alpha^{-1/3} \left[\frac{\nu^2}{\alpha^2} \right]^{d\alpha/2+1} t^{-d\alpha/2}. \end{aligned}$$

It then follows that for a constant $C > 0$

$$T_{\min}(t) \leq C d^{-1} \alpha^{-1/3} \left[\frac{\nu^2}{\alpha^2} \right]^{d\alpha/2+1} t^{-\alpha/2}, \quad t \geq \ell(d)^{-2} \frac{\nu^2}{\alpha^2}.$$

We arrive at the following.

Lemma 3.3. *If $W = XX^*$ where X is an $N \times (N + \alpha)$ matrix of iid standard complex normal random variables, $\alpha = \lfloor \sqrt{4cN} \rfloor$ and $\nu = 4N + 2\alpha + 2$ then*

$$\begin{aligned} \mathbb{P}(\lambda_{\max}(\nu^{-1}W) > t) &\leq C e^{4/3 M (t-1)}, \quad \mathbb{P}(\lambda_{\min}^{-1}(\nu^{-1}W) > t) \leq C [f(N)/t]^{d\alpha/2}, \\ f(N) &= [d^{-1} \alpha^{-1/3}]^{2/(d\alpha)} \left[\frac{\nu^2}{\alpha^2} \right]^{1+2/(d\alpha)}, \quad t \geq \ell(d)^{-2} \frac{\nu^2}{\alpha^2}. \end{aligned}$$

The following lemma is a generalization of this result, it essentially follows from the analysis in [4], by allowing $c \rightarrow 0$ in (1.5) at some rate in N as the estimates there are uniform for c bounded. We do not present a proof here as this will be included in a forthcoming work.

Lemma 3.4. *If $W = XX^*$ where X is an $N \times (N + \alpha)$ matrix of iid standard complex normal random variables, $\alpha = \lfloor \sqrt{4cN^\gamma} \rfloor$, $0 < \gamma \leq 1/2$ and $\nu = 4N + 2\alpha + 2$ then*

$$\begin{aligned} \mathbb{P}(\lambda_{\max}(\nu^{-1}W) > t) &\leq Ce^{-4/3M(t-1)}, & \mathbb{P}(\lambda_{\min}^{-1}(\nu^{-1}W) > t) &\leq C[f(N)/t]^{d\alpha/2}, \\ f(N) &= [d^{-1}\alpha^{-1/3}]^{2/(d\alpha)} \left[\frac{\nu^2}{\alpha^2} \right]^{1+2/(d\alpha)}, & t &\geq \ell(d)^{-2} \frac{\nu^2}{\alpha^2}. \end{aligned}$$

It is conjectured that these same estimates hold for $1/2 < \gamma \leq 1$ also but this does not follow immediately from the work in [4].

Proof of Theorem 1.1. It follows that

$$\kappa(A + \sigma^2 H) = \kappa(\sigma^{-2}A + H).$$

Then

$$\begin{aligned} \mathbb{P}(\lambda_{\max}(\sigma^{-2}A + H) > t) &\leq \mathbb{P}(\lambda_{\max}(H) > t - \sigma^{-2}), \\ \mathbb{P}(\lambda_{\min}^{-1}(\sigma^{-2}A + H) > t) &\leq \mathbb{P}(\lambda_{\min}^{-1}(H) > t). \end{aligned}$$

Since we may choose d as needed, we assume that $\alpha d \sim N^\lambda \rightarrow \infty$. We then have from Lemma 3.4, with a possibly new constant C ,

$$\begin{aligned} \mathbb{P}(\lambda_{\max}(H) > t) &\leq Ce^{-4/3M(x-1)}, \\ \mathbb{P}(\lambda_{\min}^{-1}(H) > t) &\leq C[f(N)/t]^{d\alpha/2}, & f(N) &= \left[\frac{\nu^2}{\alpha^2} \right]. \end{aligned}$$

This follows from the fact that $(N^q)^{N^{-\lambda}} \rightarrow 1$ for any value of q . We define δ by $1 + \delta = \ell(d)^{-2}$ and then the matrix $A + \sigma^2 H$ satisfies Condition 2.1 with $c_1 = 4/3$, $a = 1 + \sigma^{-2}$, $\alpha \rightarrow d\alpha$ and f and δ as defined here. Different values of λ can be used to create different estimates. But for simplicity, we take $\lambda = \gamma/2$ or $d = N^{-\gamma/2}$. Then by (3.1), for ϵ small, $\delta = 1 - \ell(d)^{-2} \geq N^{-\gamma/3-\epsilon}$ if N is sufficiently large and

$$\log(1 + \delta)^{N^{\gamma/2}} \rightarrow \infty,$$

with some power of N . Therefore $(1 + \delta)^{-d\alpha/K}$ tends to zero faster than any power of N if $K > 0$ is fixed. We now establish each estimate by appealing to Theorem 2.5.

(1) **Halting time with the ℓ^2 norm:** Using $b_N = (1 + \sigma^{-2})f(N)(1 + \delta_N)$ it follows directly that

$$\mathbb{E}[\tau_\epsilon(A + \sigma^2 X, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log b_N^{1/2} 2\epsilon^{-1})^j + \mathcal{O}(N^{-k}), \quad \text{for all } k > 0.$$

(2) **Halting time with the weighted norm:** Again, using $b_N = (1 + \sigma^{-2})f(N)(1 + \delta_N)$ it follows directly that

$$\mathbb{E}[\tau_{w,\epsilon}(A + \sigma^2 X, b)^j] \leq \frac{1}{2^j} b_N^{j/2} (\log 2\epsilon^{-1})^j + \mathcal{O}(N^{-k}), \quad \text{for all } k > 0,$$

(3) **Successive residuals:** Similarly, it follows directly that

$$\mathbb{E} \left[\frac{\|r_{k+1}\|^j}{\|r_k\|^j} \right] \leq \left(1 - \frac{2}{\sqrt{b_N} + 1} \right)^j + \mathcal{O}(N^{-k}), \quad \text{for all } k > 0.$$

By equation (1.5), $b_N = c^{-1}(1 + \sigma^{-2})4N^{2-2\gamma}(1 + \mathcal{O}(N^{-\gamma/3}))$, and the result follows. \blacksquare

Acknowledgments

This work was supported in part by grants NSF-DMS-1411278 (GM) and NSF-DMS-1303018 (TT). The authors thank Anne Greenbaum and Zdeněk Strakoš for useful conversations, Folkmar Bornemann for suggesting that we consider the framework of smoothed analysis and the anonymous referees for suggesting additional numerical experiments.

References

- [1] Deift P.A., Orthogonal polynomials and random matrices: a Riemann–Hilbert approach, *Courant Lecture Notes in Mathematics*, Vol. 3, New York University, Courant Institute of Mathematical Sciences, New York, Amer. Math. Soc., Providence, RI, 1999.
- [2] Deift P.A., Menon G., Olver S., Trogdon T., Universality in numerical computations with random data, *Proc. Natl. Acad. Sci. USA* **111** (2014), 14973–14978, [arXiv:1407.3829](https://arxiv.org/abs/1407.3829).
- [3] Deift P.A., Trogdon T., Universality for the Toda algorithm to compute the eigenvalues of a random matrix, [arXiv:1604.07384](https://arxiv.org/abs/1604.07384).
- [4] Deift P.A., Trogdon T., Menon G., On the condition number of the critically-scaled Laguerre unitary ensemble, *Discrete Contin. Dyn. Syst.* **36** (2016), 4287–4347, [arXiv:1507.00750](https://arxiv.org/abs/1507.00750).
- [5] Edelman A., Eigenvalues and condition numbers of random matrices, *SIAM J. Matrix Anal. Appl.* **9** (1988), 543–560.
- [6] Forrester P.J., The spectrum edge of random matrix ensembles, *Nuclear Phys. B* **402** (1993), 709–728.
- [7] Golub G.H., Van Loan C.F., Matrix computations, 4th ed., *Johns Hopkins Studies in the Mathematical Sciences*, Johns Hopkins University Press, Baltimore, MD, 2013.
- [8] Greenbaum A., Behavior of slightly perturbed Lanczos and conjugate-gradient recurrences, *Linear Algebra Appl.* **113** (1989), 7–63.
- [9] Greenbaum A., Strakoš Z., Predicting the behavior of finite precision Lanczos and conjugate gradient computations, *SIAM J. Matrix Anal. Appl.* **13** (1992), 121–137.
- [10] Hestenes M.R., Stiefel E., Methods of conjugate gradients for solving linear systems, *J. Research Nat. Bur. Standards* **49** (1952), 409–436.
- [11] Liesen J., Strakoš Z., Krylov subspace methods. Principles and analysis, *Numerical Mathematics and Scientific Computation*, Oxford University Press, Oxford, 2013.
- [12] Marchenko V.A., Pastur L.A., Distribution of eigenvalues for some sets of random matrices, *Math. USSR Sb.* **1** (1967), 457–483.
- [13] Olver F.W.J., Lozier D.W., Boisvert R.F., Clark C.W. (Editors), NIST handbook of mathematical functions, U.S. Department of Commerce, National Institute of Standards and Technology, Washington, DC, Cambridge University Press, Cambridge, 2010, available at <http://dlmf.nist.gov/>.
- [14] Pfrang C.W., Deift P., Menon G., How long does it take to compute the eigenvalues of a random symmetric matrix?, in Random Matrix Theory, Interacting Particle Systems, and Integrable Systems, *Math. Sci. Res. Inst. Publ.*, Vol. 65, Cambridge University Press, New York, 2014, 411–442, [arXiv:1203.4635](https://arxiv.org/abs/1203.4635).
- [15] Sagun L., Trogdon T., LeCun Y., Universality in halting time and its applications in optimization, [arXiv:1511.06444](https://arxiv.org/abs/1511.06444).
- [16] Sankar A., Spielman D.A., Teng S.H., Smoothed analysis of the condition numbers and growth factors of matrices, *SIAM J. Matrix Anal. Appl.* **28** (2006), 446–476, [cs.NA/0310022](https://arxiv.org/abs/cs.NA/0310022).
- [17] Simon B., Trace ideals and their applications, *Mathematical Surveys and Monographs*, Vol. 120, 2nd ed., Amer. Math. Soc., Providence, RI, 2005.
- [18] Spielman D., Teng S.H., Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time, in Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing, *ACM*, New York, 2001, 296–305.
- [19] Spielman D.A., Teng S.H., Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time, *J. ACM* **51** (2004), 385–463, [cs.DS/0111050](https://arxiv.org/abs/cs.DS/0111050).
- [20] Szegő G., Orthogonal polynomials, *American Mathematical Society, Colloquium Publications*, Vol. 23, 4th ed., Amer. Math. Soc., Providence, RI, 1975.
- [21] Tracy C.A., Widom H., Level-spacing distributions and the Airy kernel, *Comm. Math. Phys.* **159** (1994), 151–174, [hep-th/9211141](https://arxiv.org/abs/hep-th/9211141).
- [22] Wilkinson J.H., Error analysis of direct methods of matrix inversion, *J. ACM* **8** (1961), 281–330.