A Randomized Homotopy for the Hermitian Eigenpair Problem

Diego Armentano
Centro de Matemática
Universidad de la República
Montevideo, URUGUAY
e-mail: diego@cmat.edu.uy

Felipe Cucker*
Department of Mathematics
City University of Hong Kong
HONG KONG
e-mail: macucker@cityu.edu.hk

Abstract. We describe and analyze a randomized homotopy algorithm for the Hermitian eigenvalue problem. Given an $n \times n$ Hermitian matrix A the algorithm returns, almost surely, a pair (λ, v) which approximates, in a very strong sense, an eigenpair of A. We prove that the expected cost of this algorithm, where the expectation is both over the random choices of the algorithm and a probability distribution on the input matrix A, is $\mathcal{O}(n^6)$, that is, cubic on the input size. Our result relies on a cost assumption for some pseudo-random number generators whose rationale is argued by us.

AMS classification numbers: 65F15.

Keywords: Eigenpair computation, homotopy methods.

Dedicated to Mike Shub on his 70th birthday, for years of friendship.

Communicated by Teresa Krick and James Renegar.

1 Introduction

Numerical linear algebra is, arguably, the best developed part of numerical analysis. This is not surprising since all numerical computations ultimately reduce to linear algebra and this character of linear algebra has directed substantial efforts towards its study. The above notwithstanding there are a few problems where the gap between algorithmic practice and theoretical understanding is still large. Notably among them is the eigenvalue problem.

*Partially funded by a GRF grant from the Research Grants Council of the Hong Kong SAR (project number CityU 100810).

In a nutshell, we have come up with algorithms for which either the complexity is well understood but appear to be numerically unstable in practice or they exhibit a stable behavior but defy a complexity analysis. The best known example of the former is the procedure consisting in computing the characteristic polynomial χ_A of the $n \times n$ input matrix A and then approximating the zeros $\lambda_1, \ldots, \lambda_n$ of χ_A. Both ingredients are known to have low complexity but the whole procedure often fails to work due to numerical instability. Today's algorithm of choice for eigenvalue computations (and a perfect example of the latter) is the QR iteration with implicit shifts. The lights and shadows of this procedure are described by Jim Demmel [11, p. 139] as follows:

> It is interesting that after more than 30 years of dependable service, convergence failures of this algorithm have quite recently been observed, analyzed, and patched [. . .]. But there is still no global convergence proof, even though the current algorithm is considered quite reliable.

Demmel then continues with the following words, unambiguosly stating the agenda:

> So the problem of devising an algorithm that is numerically stable and globally (and quickly!) convergent remains open.

One might believe that in the special case of symmetric (or Hermitian) matrices the pleasant properties of these matrices would help to carry out this agenda. And indeed, already in 1980, a book by Beresford Parlett [19] displayed a toolbox of procedures for the symmetric eigenvalue problem. Nonetheless, twenty eight years later, Percy Deift collected a list of open problems in random matrix theory [10] the twelfth of which reads as follows:

> how long does it take on average to diagonalize a random symmetric $n \times n$ matrix M?

After which, Deift added

> For definitiveness, let us assume that M is chosen from GOE and that we use the standard QR eigenvalue algorithm. The fundamental question then takes the following more concrete form: given $\epsilon > 0$, how many QR steps does it take on average to compute the eigenvalues of a GOE random matrix to order ε?

Here GOE refers to the *Gaussian orthogonal ensemble*, a class of random real symmetric matrices that naturally extends (to the symmetric situation) the class of Gaussian matrices. An in-depth empirical study on the behavior of the QR and Toda's algorithms for computing eigenvalues was recently carried out by Pfrang, Deift and Menon in [20].

In spite of the lack of a complete understanding of the QR algorithm on symmetric matrices, one may argue that Demmel's agenda has satisfactory solutions in

this case. For algorithms with a $\mathcal{O}(n^3)$ running time exist which are shown to be numerically stable (see [11, Chapter 5]).

The purpose of this paper is to provide a different answer in the Hermitian case, by designing and analyzing a homotopy method. We do not expect our algorithm to be competitive with state-of-the-art practice but rather to provide a form of analysis that has not been performed for the eigenvalue problem.

Homotopy methods have been used for decades and we are not the first to propose their use for eigenvalue problems (see, e.g., [15]). But previous works on this usage did not provide any form of complexity analysis. Actually, the impetus on complexity analysis for homotopy methods comes from a different problem namely, that of approximating zeros of complex polynomial systems, a task that Steve Smale listed as the 17th in his list of 18 problems for the mathematicians of the 21st century [24]. A major development for this problem, carried out by Carlos Beltrán and Luis Miguel Pardo [3, 4], introduced the idea of drawing the initial data for the homotopy from a random distribution and then averaging both with respect to these random choices (this is the usual notion of cost for Las Vegas type randomized algorithms) and with respect to the input data (the standard average analysis of algorithms). The bound on the cost thus obtained is polynomially bounded on the size of the input system. Our analysis rests on these ideas and closely follows their subsequent development in [7]. The way has been paved in [2] where a geometric framework, and the main properties of the condition number, for eigenvalue computations were established. Loosely speaking our main result is the following.

Main Theorem. *We describe a randomized algorithm which, given an $n \times n$ Hermitian matrix A, returns a pair (λ, v) which is a good approximation of an eigenpair of A. The average cost of this algorithm (over the random choices made by the algorithm and assuming A drawn from a $\mathrm{GUE}(n)$ distribution) is bounded by $\mathcal{O}(n^6)$. A simple modification of the algorithm computes approximate eigenpairs for all the eigenpairs of A with an average cost of $\mathcal{O}(n^7)$.*

In this statement the word *cost* refers to number of arithmetic operations. A couple of other notions (what is a good approximation, what is the $\mathrm{GUE}(n)$ distribution) will be properly defined in the next section. Once this is done, we will close that section by providing a formal statement of our main result (see Theorem 4), and comparing this result with the current algorithmic solutions for the same problem.

Proofs are given in Section 3.

Our algorithm relies on the existence of pseudo-random number generators for a certain family of densities. Since random number generation is not our main concern, we stopped short from analyzing algorithms doing this task. Instead, we assumed these algorithms at hand along with a bound on their complexity. A brief description of a common tecnique in the subject with a rationale for the complexity assumption are given in §2.8.

Remark 1 (i) We do not carry out any accuracy analysis in this paper. That is,

we assume all computations are performed with infinite precision. We finish this introduction, however, with two observations regarding finite precision. Firstly, the fact that homotopy methods are widely used in practice and are considered to be numerically stable. This is to be expected since the core of the method is to follow a curve by a sequence of points which are near the curve. Since only approximations are involved, the use of finite-precision arithmetic should not harm. Secondly, for the problem of approximating zeros of complex polynomial systems mentioned above, the numerical stability of linear homotopies has been shown in [5]. A similar analysis can be carried out in our context.

(ii) The complexity bound $\mathcal{O}(n^6)$ looks (and is) outrageous when compared with the $\mathcal{O}(n^3)$ mentioned above. Yet, on closer examination, it turns out that such comparison is between birds of different feather. We deal with this issue in §2.9.

Acknowledgments. We are grateful to Mike Shub —who raised our awareness on the gaps in the theoretical understanding of eigenvalue computations— to Carlos Andradas and Teresa Krick for useful discussions on algebraic geometry, and to Luc Devroye, Gabriel Illanes, and Gabor Lugosi, for the same on random number generation. We want to express our gratitute to Laurent Mertz, through whom we became aware of [16] and [18]. Finally, we are thankful to the anonymous referees for the enlightening feedback in their reports.

2 Preliminaries, Basic Ideas, and Main Result

2.1 Canonical Metric Structures

Let $\mathcal{H}_n$ be the set of $n \times n$ Hermitian matrices. That is, the set of matrices $H \in \mathbb{C}^{n\times n}$ such that the conjugate transpose H^* of H, coincides with H. This is a real linear subspace of $\mathbb{C}^{n\times n}$ with dimension n^2.

The space $\mathcal{H}_n$ is endowed with the restriction of the real part of the *Frobenius Hermitian product* $\langle\ ,\ \rangle_F$ on $\mathbb{C}^{n\times n}$ given by

$$\langle A, B\rangle_F := \text{trace } (B^*A) = \sum_{i,j=1}^{n} a_{ij}\,\overline{b_{ij}},$$

where $A = (a_{ij})$ and $B = (a_{ij})$. The *Frobenius norm* $\|\ \|_F$ on $\mathbb{C}^{n\times n}$ (and its restriction on $\mathcal{H}_n$) is the norm induced by $\langle\ ,\ \rangle_F$.

On the real vector space $\mathcal{H}_n \times \mathbb{R}$ we introduce the canonical inner product structure and its associated norm structure.

Abusing notation, let us denote by the symbol $d_{\mathbb{S}}$ the associated Riemannian distance on the unit spheres of both $\mathcal{H}_n$ and $\mathcal{H}_n \times \mathbb{R}$. This Riemannian distance —the *angular* distance— amounts to the angle between the arguments.

The space $\mathbb{C}^n$ is equipped with the canonical Hermitian inner product $\langle\ ,\ \rangle$.

Let $\mathbb{P}(\mathbb{C}^n)$ be the projective space associated to $\mathbb{C}^n$. This projective space is a smooth manifold which carries a natural Riemannian metric, namely, the real part of the *Fubini-Study metric* on $\mathbb{P}(\mathbb{C}^n)$. The Fubini-Study metric is the Hermitian structure on $\mathbb{P}(\mathbb{C}^n)$ given in the following way: for $x \in \mathbb{C}^n$,

$$\langle w, w' \rangle_x := \frac{\langle w, w' \rangle}{\|x\|^2},$$

for all w, w' in the Hermitian complement $x^\perp$ of x in $\mathbb{C}^n$. We denote by $d_{\mathbb{P}}$ the associated Riemannian distance.

The space $\mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$ is endowed with the Riemannian product structure. Besides the induced Riemannian distance, we consider on this space the product of the angle distances on $\mathcal{H}_n \times \mathbb{R}$ and $\mathbb{P}(\mathbb{C}^n)$. We denote this product distance by $d_{\mathbb{S}\times\mathbb{P}}$.

Remark 2 A reasonable assumption is that the nature of our problem is invariant under the scaling of the matrix and its eigenvalues. In this way, one may consider the real projective space, or the real sphere, associated to the vector space $\mathcal{H}_n \times \mathbb{R}$ (see [2]). Here, for the sake of simplicity in the exposition, we carry out our computations on the affine vector space $\mathcal{H}_n \times \mathbb{R}$, however the reader must keep in mind that all quantities involved will be invariant under scaling and therefore the main distance is the angular distance $d_{\mathbb{S}}$.

2.2 The Varieties $\mathcal{V}$, $\mathcal{W}$, Σ' and Σ

Motivated by [2], we define the *solution variety* for the Hermitian eigenpair problem as

$$\mathcal{V} := \{(A, \lambda, v) \in \mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n) :\ (\lambda \mathsf{Id} - A)v = 0\}.$$

Proposition 1 *The solution variety $\mathcal{V}$ is a smooth submanifold of $\mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$, of the same dimension as $\mathcal{H}_n$.*

Proof. See §3.5. □

This set inherits the Riemannian structure of the ambient space.

Associated to $\mathcal{V}$ there is a natural projection, namely, the canonical projection on the first component $\pi : \mathcal{V} \to \mathcal{H}_n$. Its derivative $D\pi$ at (A, λ, v) is a linear operator between spaces of equal dimension.

The *subvariety $\mathcal{W}$ of well-posed triples* is the subset of triples $(A, \lambda, v) \in \mathcal{V}$ for which $D\pi(A, \lambda, v)$ is an isomorphism. In particular, when $(A, \lambda, v) \in \mathcal{W}$, the projection π has a branch of its inverse (locally defined) taking $A \in \mathcal{H}_n$ to $(A, \lambda, v) \in \mathcal{V}$. This branch of π^{-1} is called the *solution map*, and is the map that eigenpair algorithms attempt to compute.

Given $(A, \lambda, v) \in \mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$, let $A_{\lambda,v}$ be the linear operator on the Hermitian complement $v^\perp$ of v given by

$$A_{\lambda,v} := \Pi_{v^\perp}(\lambda \mathsf{Id} - A)|_{v^\perp},$$

where $\Pi_{v^\perp} : \mathbb{C}^n \to v^\perp$ is the orthogonal projection. Then, one can prove that the set of well-posed triples is given by

$$\mathcal{W} = \{(A, \lambda, v) \in \mathcal{V} : \ A_{\lambda,v} \ \text{is invertible}\}, \tag{1}$$

(see Lemma 2.7 in [2]).

Let $\Sigma' := \mathcal{V} \setminus \mathcal{W}$ be the variety of *ill-posed problems*, and $\Sigma = \pi(\Sigma') \subset \mathcal{H}_n$ be the *discriminant variety*, i.e., the subset of ill-posed inputs.

Remark 3 From (1) it is clear that the subset Σ' is the set of triples $(A, \lambda, v) \in \mathcal{V}$ such that λ is an eigenvalue of A of algebraic multiplicity at least 2. It follows that Σ is the set of matrices $A \in \mathcal{H}_n$ with multiple eigenvalues. In particular, when $A \in \mathcal{H}_n \setminus \Sigma$ then, $\pi^{-1}(A)$ is the set of triples $(A, \lambda_1, v_1), \ldots, (A, \lambda_n, v_n)$, where (λ_i, v_i) are the n-eigenpairs of A.

Proposition 2 *The discriminant variety $\Sigma \subset \mathcal{H}_n$ is a real algebraic variety, and for all $n \geq 2$ we have* $\dim \Sigma = n^2 - 3$.

Proof. See §3.4. □

2.3 Unitary invariance

Let $\mathcal{U}(n)$ be the group of $n \times n$ unitary matrices. The group $\mathcal{U}(n)$ naturally acts on $\mathbb{P}(\mathbb{C}^n)$. In addition, $\mathcal{U}(n)$ acts on $\mathcal{H}_n$ by conjugation (i.e., $U \cdot A := UAU^{-1}$), and on $\mathcal{H}_n \times \mathbb{R}$ by $U \cdot (A, \lambda) := (UAU^{-1}, \lambda)$. These actions define an action on the product space $\mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$, namely,

$$U \cdot (A, \lambda, v) := (UAU^{-1}, \lambda, Uv). \tag{2}$$

Remark 4 The varieties $\mathcal{V}$, $\mathcal{W}$, Σ', and Σ, are invariant under the action of $\mathcal{U}(n)$ (see [2] for details).

2.4 Condition of a triple

We define the *condition number* of $(A, \lambda, v) \in \mathcal{W}$ as

$$\mu(A, \lambda, v) := \max\left\{1, \|A\|_F \left\|{A_{\lambda,v}}^{-1}\right\|\right\},$$

where $\| \ \|$ is the operator norm. This condition number fusions two condition numbers measuring, respectively, first-order variations for eigenvalues and eigenvectors (see Section 3 in [2] for details). It admits the following characterization.

Lemma 1 (Lemma 3.12 in [2]) *Let $A \in \mathcal{H}_n \setminus \Sigma$, and let $(\lambda_1, v_1), \ldots, (\lambda_n, v_n)$ be its eigenpairs. Then*

$$\mu(A, \lambda_1, v_1) = \max\left\{1, \frac{\|A\|_F}{\min\limits_{i=2,\ldots,n} |\lambda_i - \lambda_1|}\right\}. \qquad \square$$

Remark 5 The condition number μ is invariant under the action of the unitary group $\mathcal{U}(n)$, i.e., $\mu(UAU^{-1}, \lambda, Uv) = \mu(A, \lambda, v)$ for all $U \in \mathcal{U}(n)$. Also, it is clear from Lemma 1, μ is scale invariant on the first two components. That is, $\mu(sA, s\lambda, v) = \mu(A, \lambda, v)$ for all nonzero real s.

The following is an easy consequence of the proof of Proposition 3.14(ii) in [2].

Corollary 1 *Let $\Gamma : [0, 1] \to \mathcal{V}$, $\Gamma(t) = (A_t, \lambda_t, v_t)$ be a smooth curve such that A_t lies in the unit sphere of $\mathcal{H}_n$, for all t. Then $\left\|\frac{d}{dt}\Gamma(t)\right\| \leq 2\mu(\Gamma(t))\left\|\frac{d}{dt}A_t\right\|$.* $\square$

Remark 6 Since the property of $A_{\lambda,v}$ being invertible is open on $\mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$, the condition number μ can be naturally extended to an open neighborhood of $\mathcal{W}$ in $\mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$. We will denote this extension also by μ. In addition, when $A_{\lambda,v}$ is non-invertible we will let $\mu(A, \lambda, v) := \infty$.

Proposition 3 *Given $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that, if $(A, \lambda, v), (A', \lambda', v')$ are arbitrary triples such that*

$$d_{\mathbb{S}\times\mathbb{P}}((A, \lambda, v), (A', \lambda', v')) \leq \frac{C_\varepsilon}{\mu(A, \lambda, v)}$$

then

$$\frac{\mu(A, \lambda, v)}{1 + \varepsilon} \leq \mu(A', \lambda', v') \leq (1 + \varepsilon)\mu(A, \lambda, v).$$

One may choose $C_\varepsilon = \frac{\arctan\left(\frac{\varepsilon}{\sqrt{2}+\varsigma(1+\varepsilon)}\right)}{(1+\varepsilon)}$ where $\varsigma := (1 + \sqrt{5})2\sqrt{2}$. $\square$

This proposition, is essentially, Proposition 3.22 in [2], where it was shown only for pairs of triples in the solution variety. The extension to arbitrary triples is not problematic and may, at most, change the values of the constants.

2.5 Newton's method and approximate eigenpairs

Following [2], for a nonzero matrix $A \in \mathcal{H}_n$, we define the *Newton map* associated to A,

$$N_A : \mathbb{R} \times (\mathbb{C}^n \setminus \{0\}) \to \mathbb{R} \times (\mathbb{C}^n \setminus \{0\}),$$

by $N_A(\lambda, v) = (\lambda - \dot{\lambda}, v - \dot{v})$ where

$$\dot{v} = {A_{\lambda,v}}^{-1}\, \Pi_{v^\perp}(\lambda\, \mathsf{Id} - A)v, \qquad \dot{\lambda} = \frac{\langle (\lambda \mathsf{Id} - A)(v - \dot{v}), v\rangle}{\langle v, v\rangle}.$$

This map is defined for every $(\lambda, v) \in \mathbb{R} \times (\mathbb{C}^n \setminus \{0\})$ such that $A_{\lambda,v} := \Pi_{v^\perp}(\lambda \mathsf{Id} - A)|_{v^\perp}$ is invertible. Indeed, in this case, and since $A \in \mathcal{H}_n$, it is easily seen that $\langle (\lambda \mathsf{Id} - A)\dot{v}, v\rangle \in \mathbb{R}$, and therefore that $\dot{\lambda} \in \mathbb{R}$ as well. This map has been introduced in [2] as the Newton operator associated to the evaluation map $(\lambda, v) \mapsto (\lambda \mathsf{Id} - A)v$ for a fixed A. See Section 4 of [2] for more details.

Definition 1 Given $(A, \lambda, v) \in \mathcal{W}$ we say that $(\nu, w) \in \mathbb{R} \times (\mathbb{C}^n \setminus \{0\})$ is an *approximate eigenpair* of A with associated eigenpair (λ, v) when for all $k \geq 1$ the kth iterate of the Newton map at (ν, w) is well defined and satisfies

$$d_{\mathbb{S}\times\mathbb{P}}\big((A, N_A^k(\nu, w)), (A, \lambda, v)\big) \leq \left(\frac{1}{2}\right)^{2^k - 1} d_{\mathbb{S}\times\mathbb{P}}\big((A, \nu, w), (A, \lambda, v)\big).$$

Remark 7 Note that, if $N_A(\nu, w) = (\nu', w')$ then $N_{sA}(s\nu, \beta w) = (s\nu', \beta w')$, for every $s \in \mathbb{R}\setminus\{0\}$ and $\beta \in \mathbb{C}\setminus\{0\}$. Hence, the notion of approximate eigenpair scales correctly; cf. Remark 2.

Remark 8 The notion of approximate solution as a point where Newton's method converges to a true solution immediately and quadratically fast was introduced by Steve Smale [23]. It allows to elegantly talk about polynomial time without dependencies on pre-established accuracies. In addition, these approximate solutions are "good approximations" (as mentioned in the statement of the Main Theorem) in a very strong sense. The distance to the exact solution dramatically decreases with a single iteration of Newton's method.

The following theorem is a special case of Theorem 2 in [2]. It estimates, in terms of the condition of an eigenpair, the radius of a ball of approximate eigenpairs associated to it.

Theorem 1 *There is a universal constant $c_0 > 0$ with the following property. Let $A \in \mathcal{H}_n$ be nonzero and let $(\lambda, v), (\lambda_0, v_0) \in \mathbb{R} \times (\mathbb{C}^n \setminus \{0\})$. If (λ, v) is a well-posed eigenpair of A and*

$$d_{\mathbb{S}\times\mathbb{P}}((A, \lambda, v), (A, \lambda_0, v_0)) < \frac{c_0}{\mu(A, \lambda, v)}$$

then (λ_0, v_0) is an approximate eigenpair of A with associated eigenpair (λ, v). One may choose $c_0 = 0.0739$. □

Remark 9 We note that $N_A(\nu, w)$ can be computed from the matrix A and the pair (ν, w) in $\mathcal{O}(n^3)$ operations, since the cost of this computation is dominated by that of inverting a matrix.

2.6 Gaussian Measure on $\mathcal{H}_n$

We next introduce a natural probability measure on $\mathcal{H}_n$ which is invariant under the action of $\mathcal{U}(n)$. This measure is the so called *Gaussian Unitary Ensemble*, (GUE(n)), and has a density $\rho_{\mathcal{H}_n} : \mathcal{H}_n \to \mathbb{R}$ given by

$$\rho_{\mathcal{H}_n}(H) := \frac{1}{2^{n/2}} \frac{1}{\pi^{n^2/2}} \exp\left(-\frac{\|H\|_F^2}{2}\right),$$

(see [13, 17]). An equivalent description of GUE(n), based on random variables, will be useful to us. We give it here.

We say that the complex random variable $Z = X + \sqrt{-1}Y$ has distribution $\mathcal{N}_{\mathbb{C}}(0,1)$ when the real part X and the imaginary part Y are independent and identically distributed (i.i.d.) drawn from $\mathcal{N}(0, 1/2)$, i.e., they are Gaussian centered random variables with variance $1/2$.

(We will write $v \sim \mathcal{N}_{\mathbb{C}}(0, \mathbf{1}_n)$ to indicate that the vector $v \in \mathbb{C}^n$ is random with i.i.d. coordinates drawn from $\mathcal{N}_{\mathbb{C}}(0,1)$.)

If $Z \sim \mathcal{N}_{\mathbb{C}}(0,1)$ then its density $p_Z : \mathbb{C} \to \mathbb{R}$ with respect to the Lebesgue measure is given by

$$p_Z(z) := \frac{1}{\pi} e^{-|z|^2}.$$

Thus, it is clear from the definition of $\rho_{\mathcal{H}_n}$ that the Gaussian Unitary Ensemble corresponds to the distribution law of a random matrix on $\mathcal{H}_n$ whose diagonal entries are independent $\mathcal{N}(0,1)$ and those above (or below) the diagonal are independent $\mathcal{N}_{\mathbb{C}}(0,1)$.

A classical result in random matrix theory exhibits the joint density of the eigenvalues of a GUE(n) matrix (cf. [17, §6.2]).

Theorem 2 *Let $H \in$ GUE(n). Then, the joint density ζ_n of the (non ordered) eigenvalues of H is given by*

$$\zeta_n(\lambda_1, \ldots, \lambda_n) = \frac{1}{Z_n} \prod_{i<j} (\lambda_i - \lambda_j)^2 \exp\left(-\sum_{k=1}^{n} \frac{\lambda_k^2}{2}\right),$$

where

$$Z_n = (2\pi)^{n/2} \prod_{j=1}^{n} j!. \qquad \square$$

In the sequel we will use the notation

$$\Delta_n(\lambda_1, \ldots, \lambda_n) := \prod_{i<j} (\lambda_i - \lambda_j).$$

We will also denote by W_n the distribution on $\mathbb{R}^n$ with the density in Theorem 2.

The following is an immediate consequence of this theorem.

Corollary 2 *Let $f : \mathcal{H}_n \to \mathbb{R}$ be a measurable function such that there exist $g : \mathbb{R}^n \to \mathbb{R}$ satisfying $f(H) = g(\lambda_1, \ldots, \lambda_n)$ where $\lambda_1, \ldots, \lambda_n$ are the eigenvalues of H. Then*

$$\mathop{\mathbb{E}}_{H \sim \mathrm{GUE}(n)} f(H) = \mathop{\mathbb{E}}_{x \sim W_n} g(x).$$

Here $\mathbb{E}_{\eta \sim L}\, \varphi(\eta)$ stands for the expected value of the real-valued function φ at η drawn from the distribution L. □

The following lemma shows that some known properties of the Gaussian distribution are shared by the Gaussian Unitary Ensemble.

Lemma 2 **(i)** *Let G and R be independent* GUE(n) *random matrices. Then, for $t \in [0, 1]$,*

$$H_t = (1 - t)G + tR$$

has distribution $((1 - t)^2 + t^2)^{1/2} H$ where H is GUE(n).

(ii) *If H is* GUE(n) *then $\|H\|_F^2$ has the distribution of a $\chi^2_{n^2}$ random variable.*

Proof. **(i)** Write $G = (g_{ij})$ and $R = (r_{ij})$. Then, the (i, j)-entry of H_t is given by $(1 - t)g_{ij} + tr_{ij}$. Since linear combination of Gaussians is Gaussian, the proof follows by computing the variance of each entry.

(ii) If $H = (h_{ij})$ with $h_{ij} = a_{ij} + \sqrt{-1}\, b_{ij}$. Then

$$\|H\|_F^2 = 2 \sum_{i<j\leq n} \left(a_{ij}^2 + b_{ij}^2\right) + \sum_{i \leq n} a_{ii}^2 = \sum_{i<j\leq n} \left((\sqrt{2}a_{ij})^2 + (\sqrt{2}b_{ij})^2\right) + \sum_{i \leq n} a_{ii}^2$$

which shows the statement since $\{\sqrt{2}a_{ij}, \sqrt{2}b_{ij}, a_{ii}\}$ are n^2 independent random variables with distribution $\mathcal{N}(0, 1)$. □

Crucial in our development is the following bound on average condition for GUE(n) matrices. We will prove it in §3.3.

Theorem 3 *Let Q be a random matrix drawn from* GUE(n)*, and let $(\lambda_1, v_1), \ldots, (\lambda_n, v_n)$ be its random eigenpairs. Then,*

$$\mathop{\mathbb{E}}_{Q \sim \mathrm{GUE}(n)} \sum_{k=1}^{n} \frac{\mu^2(Q, \lambda_k, v_k)}{\|Q\|_F^2} \leq n(n - 1).$$

Remark 10 From the joint density of the eigenvalues of a GUE(n) it is clear that, with probability one, matrices have all its eigenvalues different. Therefore the expected value in Theorem 3 is well-defined.

2.7 The eigenpair continuation algorithm

Suppose that we are given an input matrix $A \in \mathcal{H}_n$ and an *initial triple* (M, λ, v) in the solution variety $\mathcal{V} \subseteq \mathcal{H}_n \times \mathbb{R} \times (\mathbb{C}^n \setminus \{0\})$ such that A and M are $\mathbb{R}$-linearly independent. Let $\alpha := d_{\mathbb{S}}(M, A) \in (0, \pi)$ denote the *angle* between the rays $\mathbb{R}_+A$ and $\mathbb{R}_+M$. Consider the line segment $[M, A]$ in $\mathcal{H}_n$ with endpoints M and A.

We parameterize this segment by writing

$$[M, A] = \{Q_\tau \in \mathcal{H}_n \mid \tau \in [0, 1]\}$$

with Q_τ being the only point in $[M, A]$ such that $d_{\mathbb{S}}(M, Q_\tau) = \tau\alpha$ (see Figure 1).

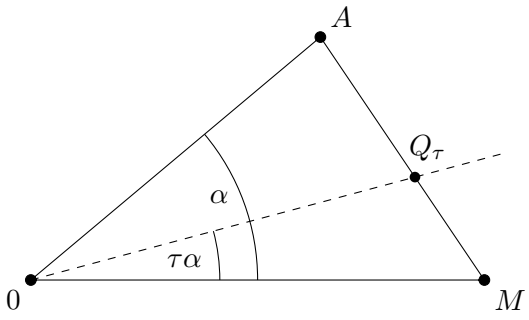

Figure 1: The family Q_τ, $\tau \in [0, 1]$.

If the line segment $[M, A]$ does not intersect the discriminant variety Σ, then starting at the eigenpair (λ, v) of M, the map $[0, 1] \to \mathcal{H}_n, \tau \mapsto Q_\tau$, can be uniquely extended to a continuous map

$$[0, 1] \to \mathcal{V}, \quad \tau \mapsto (Q_\tau, \lambda_\tau, v_\tau), \tag{3}$$

such that $(\lambda_0, v_0) = (\lambda, v)$. We call this map the *lifting* of $[M, A]$ with origin (M, λ, v). We shall also call $\tau \mapsto (Q_\tau, \lambda_\tau, v_\tau)$ the *solution path* in $\mathcal{V}$ corresponding to the input matrix A and initial triple (M, λ, v). Central to our algorithm is the trivial fact that the pair (λ_1, v_1), corresponding to $\tau = 1$, is an eigenpair of A.

In order to find an approximation of the eigenpair (λ_1, v_1) of A we may start with the eigenpair $(\lambda, v) = (\lambda_0, v_0)$ of $M = Q_0$ and numerically follow the path $(Q_\tau, \lambda_\tau, v_\tau)$ by subdividing the interval $[0, 1]$ into subintervals with extremities at $0 = \tau_0 < \tau_1 < \cdots < \tau_K = 1$ and by successively computing approximations (ν_i, w_i) of $(\lambda_{\tau_i}, v_{\tau_i})$ by Newton's method. We want to do so ensuring that for all i, (ν_i, w_i) is an approximate eigenpair of $Q_{\tau_{i+1}}$. Figure 2 attempts to convey the general idea.

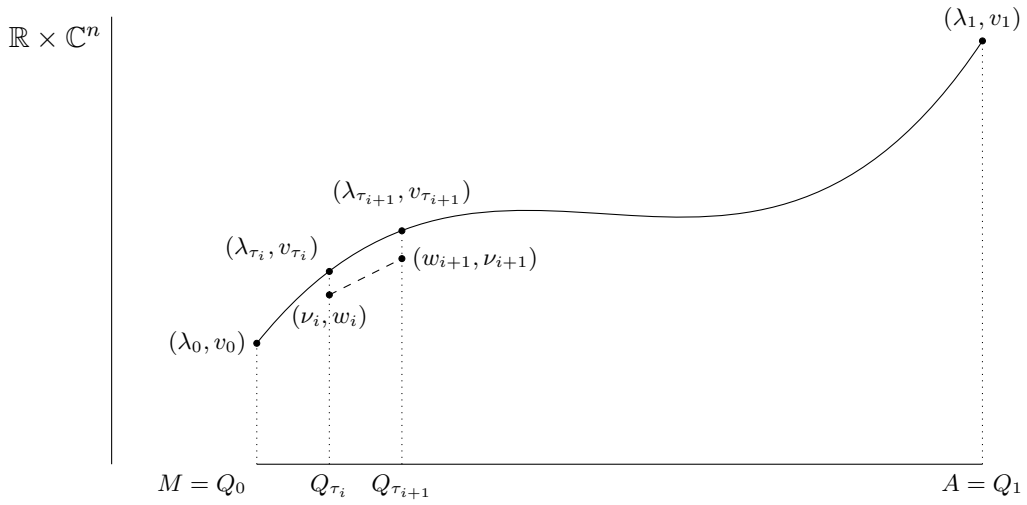

Figure 2: The continuation of the solution path.

The following algorithm gives a precise description of how this is done. The letter ξ denotes a constant, namely $\xi = 0.008535284254$.

Algorithm 1 EC

Input: $M, A \in \mathcal{H}_n$ and $(\lambda, v) \in \mathbb{R} \times \mathbb{C}^n$

Preconditions: $(M - \lambda\,\mathsf{Id})v = 0$, $M \notin \mathbb{R}A$, $v \neq 0$

$\alpha := d_{\mathbb{S}}(M, A)$, $r := \|A\|_F$, $s := \|M\|_F$
$\tau := 0$, $Q := M$, $(\nu, w) := (\lambda, v)$
`repeat`
  $\Delta\tau := \frac{\xi}{\alpha\mu^2(Q,\nu,w)}$
  $\tau := \min\{1, \tau + \Delta\tau\}$
  $t := \frac{s}{r \sin\alpha \cot(\tau\alpha) - r\cos\alpha + s}$
  $Q := tA + (1-t)M$
  $(\nu, w) := N_Q(\nu, w)$
`until` $\tau = 1$
`return` (ν, w)

Output: $(\nu, w) \in \mathbb{R} \times \mathbb{C}^n$

Postconditions: The algorithm halts if the lifting of $[M, A]$ at (λ, v) does not cut Σ'. In this case, (ν, w) is an approximate eigenpair of A.

The following result estimates the number of iterations performed by algorithm EC.

Proposition 4 *Suppose that $[M, A]$ does not intersect the discriminant variety Σ. Then the algorithm* EC *stops after at most $K := K(A, M, \lambda)$ steps with*

$$K \;\leq\; 822\; d_{\mathbb{S}}(M, A) \int_0^1 \mu^2(Q_\tau, \lambda_\tau, v_\tau)\, d\tau.$$

The returned pair (ν, w) is an approximate eigenpair of A with associated eigenpair (λ_1, v_1). Furthermore, the bound above is optimal up to a constant: we have

$$K \;\geq\; 383\; d_{\mathbb{S}}(M, A) \int_0^1 \mu^2(Q_\tau, \lambda_\tau, v_\tau)\, d\tau.$$

Proof. See §3.1. □

2.8 Randomization and a Las Vegas procedure

Our algorithm for computing eigenpairs is a randomized Las Vegas algorithm, as are the main procedures in [4, 7]. This means that we assume we have at hand a routine returning a random vector from a given distribution. In both [4] and [7] this distribution was a multivariate standard Gaussian (that is, in both cases and with input $n \in \mathbb{N}$, the routine returned a vector $v \sim \mathcal{N}(0, \mathbf{1}_n)$). Our setting demands to draw vectors from a different distribution and because of this, a few words on random number generation are in order. We will be brief, noting that a wealth of information on random number generators can be found in [12, 14].

Random number generators are standard in scientific computing and many programming languages offer implementations for the most common distributions. The basic routine draws a number from a uniform distribution in a set of k elements. To draw from other distributions one relies on this one as well as reasonable approximations in the case of continuous distributions. Thus, for instance, elements from the uniform distribution on the interval $(0, 1)$ may be drawn by chosing an element in $\{\frac{1}{k}, \ldots, \frac{k-1}{k}\}$ with k large enough. And elements from the standard Gaussian $\mathcal{N}(0, 1)$ by first drawing $u \in (0, 1)$ from the uniform distribution and then return $x = \Phi^{-1}(u)$ where Φ is the cumulative distribution function of the Gaussian.

The results in [4, 7] assumed a routine drawing real numbers from $\mathcal{N}(0, 1)$ and disregarded the actual implementation of this routine along with the fact that such implementation can only compute approximations of the theoretical generator. Such assumption is coherent with the overall assumption of infinite precision in these papers. Furthermore, they assumed that the cost of drawing from $\mathcal{N}(0, 1)$ was $\mathcal{O}(1)$. A rationale for this assumption on cost (none was offered in [4, 7]) could be the following. Given a number $u \in (0, 1)$ we may draw x from $\mathcal{N}(0, 1)$ by computing $\Phi^{-1}(u)$ (note though, that other ways of doing so, such as the Box-Müller transform,

are possible). This is done up to a prespecified precision by performing a finite number of arithmetic operations and evaluations $\varphi(y)$ for various values of $y \in \mathbb{R}$ (here φ is the density of $\mathcal{N}(0,1)$). Assuming these evaluations are done in constant time as well, the overall time for computing $\Phi^{-1}(u)$ (i.e., for drawing x from $\mathcal{N}(0,1)$) is *independent of the input size* of whichever execution is calling for the random draw. Hence, with respect to this input size, this time is $\mathcal{O}(1)$. Of course, this implies that the cost of producing a v from $\mathcal{N}(0, \mathbf{1}_n)$ is $\mathcal{O}(n)$.

The densities ρ we will draw real numbers from are a combination of polynomial and exponential functions (their exact shape is described below). They are evaluated at a point $x \in \mathbb{R}$ by a sequence of arithmetic operations and calls to the exponential function and —in contrast with the picture above for the drawing of points from $\mathcal{N}(0,1)$— for the various densities we will have to draw from, the cost of such an evaluation (i.e., the number of operations and calls to the exponential) depends on n. In line with the rationale above, we will make the following assumption.

Randomization assumption: Let $\Xi_\rho(n)$ denote the evaluation cost for a density ρ. Then we can draw a real number from ρ with cost $\mathcal{O}(\Xi_\rho(n))$.

We can now proceed to describe our randomization routine.

For $k = 1, \ldots, n-1$, let

$$R_n^{(k)}(x_1, \ldots, x_k) := \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \zeta_n(x_1, \ldots, x_k, x_{k+1}, \ldots, x_n)\, dx_{k+1} \ldots dx_n,$$

where ζ_n is given in Theorem 2. Also, let $R_n^{(n)} := \zeta_n$. Then, for $k = 1, \ldots, n$, the function $R_n^{(k)}$ is a joint density function on $\mathbb{R}^k$.

These densities are easily computable, because of the following well-known result (see, e.g., [17, §6.2]).

Lemma 3 *For $k = 1, \ldots, n$,*

$$R_n^{(k)}(x_1, \ldots, x_k) = \det\left(K_n(x_i, x_j)\right)_{1 \le i,j \le k},$$

where K_n is the reproducing kernel given by

$$K_n(x,y) := \frac{1}{\sqrt{2\pi}} \sum_{j=1}^{n} \frac{H_j(x)\, H_j(y)}{j!} \cdot \exp\left(\frac{-(x^2+y^2)}{4}\right).$$

Here $H_j(x) := (-1)^j \exp(x^2/2) \frac{d^j}{dx^j} \exp(-x^2/2)$, is the jth Hermite's polynomial. □

Our randomization routine is the following.

Algorithm 2 random_triple

Input: $n \in \mathbb{N}$

`draw` $\lambda_1 \in \mathbb{R}$ `from the density` $R_n^{(1)}(*)$
`from` $k = 2, \ldots, n$:
    `draw` λ_k `from the density` $\frac{R_n^{(k)}(\lambda_1,\ldots,\lambda_{k-1},*)}{R_n^{(k-1)}(\lambda_1,\ldots,\lambda_{k-1})}$
`draw` $U \in \mathcal{U}(n)$ `from the uniform distribution`
`let` $v := Ue_1$
`let` $M := U\,\mathsf{diag}(\lambda_1, \ldots, \lambda_n)\,U^*$; $\lambda := \lambda_1$

Output: $(M, \lambda, v) \in \mathcal{H}_n \times \mathbb{R} \times \mathbb{C}^n$
Postconditions: $M \sim \mathrm{GUE}(n)$ and $(M - \lambda\,\mathsf{Id})v = 0$

The first property we want to show about algorithm random_triple is its correctness.

Proposition 5 *The output* (M, λ, v) *of* random_triple *satisfies* $(M - \lambda\,\mathsf{Id})v = 0$. *Furthermore, the matrix* M *follows the* $\mathrm{GUE}(n)$ *distribution, the (conditional to be an eigenpair of* M*) probability distribution of* $(\lambda, v) \in \mathbb{R} \times \mathbb{P}(\mathbb{C}^n)$ *is the discrete uniform on the* n *eigenpairs and we have, for all function* $F : \mathcal{H}_n \times \mathbb{R} \times \mathbb{P}(\mathbb{C}^n) \to \mathbb{R}$, *that*

$$\mathop{\mathbb{E}}_{(M,\lambda,v)\sim\text{random\_triple}} F(M, \lambda, v) = \mathop{\mathbb{E}}_{M\sim\mathrm{GUE}(n)} \frac{1}{n} \sum_{i=1}^{n} F(M, \lambda_i, v_i)$$

where $(\lambda_1, v_1), \ldots, (\lambda_n, v_n)$ *are the* n *eigenpairs of* M.

Proof. By construction, the distribution of M is invariant under the action of the unitary group. Hence, the probability density of M is a function of its eigenvalues $\lambda_1, \ldots, \lambda_n$. Let $\rho(\lambda_1, \ldots, \lambda_n)$ denote the joint probability density of $(\lambda_1, \ldots, \lambda_n) \in \mathbb{R}^n$ generated in random_triple. Then, by Theorem 2, it is enough to prove that $\rho(x_1, \ldots, x_n)$ coincides with $\zeta_n(x_1, \ldots, x_n)$ for all $(x_1, \ldots, x_n) \in \mathbb{R}^n$.

Again by construction, one has that $\rho(x_1, \ldots, x_n)$ is the product

$$\rho_{\lambda_1}(x_1)\,\rho_{\lambda_2|\lambda_1=x_1}(x_2) \cdots \rho_{\lambda_n|(\lambda_1,\ldots,\lambda_{n-1})=(x_1,\ldots,x_{n-1})}(x_n),$$

where $\rho_{\lambda_1}(x_1)$ is the probability density that λ_1 takes the value x_1, and where $\rho_{\lambda_k|(\lambda_1,\ldots,\lambda_{k-1})=(x_1,\ldots,x_{k-1})}(x_k)$ is the conditional probability density of λ_k at x_k given that the random vector $(\lambda_1, \ldots, \lambda_{k-1})$ in $\mathbb{R}^{k-1}$ takes the value $(x_1, \ldots, x_{k-1})$. That is,

$$\rho_{\lambda_1}(x_1) = R_n^{(1)}(x_1); \quad \rho_{\lambda_k|(\lambda_1,\ldots,\lambda_{k-1})=(x_1,\ldots,x_{k-1})}(x_k) = \frac{R_n^{(k)}(x_1, \ldots, x_{k-1}, x_k)}{R_n^{(k-1)}(x_1, \ldots, x_{k-1})}.$$

Hence,

$$\begin{aligned}\rho(x_1,\ldots,x_n) &= R_n^{(1)}(x_1)\,\frac{R_n^{(2)}(x_1,x_2)}{R_n^{(1)}(x_1)}\cdots\frac{R_n^{(n)}(x_1,\ldots,x_{n-1},x_n)}{R_n^{(n-1)}(x_1,\ldots,x_{n-1})}\\ &= R_n^{(n)}(x_1,\ldots,x_n) \;=\; \zeta_n(x_1,\ldots,x_n),\end{aligned}$$

the last equality by the definition of $R_n^{(n)}$. □

The second property we want to show of random_triple is a bound on its cost.

Proposition 6 *Under the randomization assumption, the cost of* random_triple *with input n is $\mathcal{O}(n^4)$.*

Proof. The densities we want to draw from are of the form

$$\rho_k(x) := \frac{R_n^{(k)}(\lambda_1,\ldots,\lambda_{k-1},x)}{R_n^{(k-1)}(\lambda_1,\ldots,\lambda_{k-1})}$$

where the numbers $\lambda_1,\ldots,\lambda_{k-1}$ have been previously computed and are therefore given. To evaluate these quotients at a given $x\in\mathbb{R}$ we will rely on the explicit expression for the functions $R_n^{(j)}$ given in Lemma 3.

Let $k>1$. We first note that the quantities $K_n(\lambda_i,\lambda_j)$ for $i,j\le k-1$ have already been computed when dealing with ρ_ℓ when $\ell<k$. We still need to compute the values $K_n(\lambda_i,x)$.

Now, for $x\in\mathbb{R}$ the cost of computing $H_j(x)$ is $\mathcal{O}(j)$. We can compute all of them, for $j=1,\ldots,n$, with cost $\mathcal{O}(n^2)$. With these quantities at hand, each $K_n(\lambda_i,x)$ is computed with cost $\mathcal{O}(n)$ and, again, the $k-1$ of them are obtained with cost $\mathcal{O}(kn)$. With $\mathcal{O}(k^3)$ further operations to compute a determinant we obtain $R_n^{(k)}(\lambda_1,\ldots,\lambda_{k-1},x)$. The denominator $R_n^{(k-1)}(\lambda_1,\ldots,\lambda_{k-1})$ is computed with a similar cost.

All in all, $\rho_k(x)$ is computed with cost $\mathcal{O}(nk^2)$ (it is immediate to check that this holds for $k=1$ as well). It follows that the total cost of the n drawings (from $\rho_1,\ldots,\rho_n$) is

$$\sum_{k=1}^{n}\mathcal{O}(nk^2) \;=\; \mathcal{O}(n^4).$$

Finally, we note that the cost of drawing U and compute M is dominated by that of these drawings (and, actually, this step is not necessary; see Remark 13 below) which finishes the proof. □

Remark 11 (i) Two recent papers related with the need of the randomization assumption are [18] and [16]. In the former, an algorithm is presented and it is claimed it performs precisely what we assume. Unfortunately, though,

there are no proofs arguing for this claim, which is sustained only by numerical experiments. The latter presents a different procedure to draw a triple as in random_triple, also with cost $\mathcal{O}(n^4)$. It is based on Dyson Brownian Motion. Unfortunately, again, there are no proofs in this paper either.

(ii) Because $\exp(\frac{-(\lambda_i^2+x^2)}{4}) = \exp(\frac{-\lambda_i^2}{4})\exp(\frac{-x^2}{4})$ it is immediate to check that the densities $\rho_k(x)$ are of the form $P_{\lambda_1,\dots,\lambda_k}(x)\exp(\frac{-x^2}{4})$ where $P_{\lambda_1,\dots,\lambda_k}(x)$ is a polynomial of degree $2n$ in x whose coefficients depend on the tuple $(\lambda_1,\dots,\lambda_k)$. These densities are therefore, in a sense, close to Gaussian densities.

The following code puts together the randomization and path-following routines.

Algorithm 3 Eigenpairs

Input: $A \in \mathcal{H}_n$

$(M, \lambda, v) := \mathsf{random\_triple}(n)$
$(\nu, w) := \mathsf{EC}(A, M, \lambda, v)$

Output: $(\nu, w) \in \mathbb{R} \times \mathbb{C}^n$
Postconditions: The algorithm halts if $[M, A] \cap \Sigma = \emptyset$. In this case, the pair (ν, w) is an approximate eigenpair of A.

Remark 12 The fact that the codimension of Σ in $\mathcal{H}_n$ is 3 (shown in Proposition 2) ensures that, almost surely, the segment $[M, A]$ does not intersect Σ and therefore, that Algorithm Eigenpairs halts.

Given a matrix $A \in \mathcal{H}_n$ the cost of algorithm Eigenpairs with input A depends on the triple (M, λ, v) which is random. We therefore consider the *randomized cost* of this algorithm on input A. Because of Proposition 5, the expected number of iterations of algorithm EC with input A is given by

$$\mathsf{Num\_Iter}(A) := \mathop{\mathbb{E}}_{M\sim \mathrm{GUE}(n)} \frac{1}{n}\sum_{i=1}^{n} K(A, M, \lambda_i, v_i).$$

Since we are interested on the *average complexity* of Eigenpairs we will further take the expectation of $\mathsf{Num\_Iter}(A)$ when A is drawn from $\mathrm{GUE}(n)$. We therefore obtain

$$\mathsf{Num\_Iter}(n) := \mathop{\mathbb{E}}_{A\sim \mathrm{GUE}(n)} \mathop{\mathbb{E}}_{M\sim \mathrm{GUE}(n)} \frac{1}{n}\sum_{i=1}^{n} K(A, D, \lambda_i, v_i).$$

Multiplying this expression by the cost $\mathcal{O}(n^3)$ of each iteration (and adding an $\mathcal{O}(n^4)$ term for the execution of $\mathsf{random\_triple}(n)$) we obtain the *average cost* $\mathsf{Cost}(n)$ of Eigenpairs.

Theorem 4 *Algorithm* Eigenpairs *returns (almost surely) an approximate eigenpair of its input $A \in \mathcal{H}_n$. Under the randomization assumption, its average cost satisfies*

$$\mathsf{Cost}(n) = \mathcal{O}(n^6).$$

A simple modification of Eigenpairs *that follows all the eigenpairs of M returns all the eigenpairs of A with average cost $\mathcal{O}(n^7)$.*

Proof. See §3.2. □

Remark 13 We note here that, in practice, we do not need to draw U from the uniform distribution in $\mathcal{U}(n)$ and then return $M := U\,\mathsf{diag}(\lambda_1, \ldots, \lambda_n)\,U^*$ in random_triple. It is enough to return (M_0, λ_1, e_1) where $M_0 := \mathsf{diag}(\lambda_1, \ldots, \lambda_n)$. In doing so, because of unitary invariance, we obtain the same value of $\mathsf{Cost}(n)$.

2.9 A brief comparison with current results

It will be useful to begin noting that the word *approximation* —say $\widetilde{\zeta}$, of a solution ζ of a problem with datum A— is used in the literature with at least three different meanings (cf. [8, §O.2 and §15.2]):

Backward approximation. The element $\widetilde{\zeta}$ is the solution of a datum $\widetilde{A}$ close to A. In floating-point computations, this is usually understood as a multiple of the machine epsilon.

Forward approximation. The quantity $\|\widetilde{\zeta} - \zeta\|$ (or some component-wise version of it) is small. Again, this is usually understood as a multiple of the machine epsilon.

Approximation à la Smale. An appropriate version of Newton's iteration, starting at $\widetilde{\zeta}$, converges immediately, quadratically fast, to ζ.

Without deepening into technical details, what is relevant to our discussion is that we listed these notions in an order that shows the more demanding last. Indeed, to compare the first two, one notes that for a backward approximation, the forward error is obtained, generally speaking, multiplying the backward error by the condition number of the data at hand (in its usual meaning as the largest possible magnification of the function value for small perturbations of the argument). Hence, for poorly conditioned data, good backward approximations can lead to large forward errors (i.e., they do not correspond to forward approximations). Also, it is clear that an approximation à la Smale is stronger than a forward approximation.

To obtain the latter with an error ε it is enough to perform $\mathcal{O}(\log|\log\varepsilon|)$ Newton's steps.

The $\mathcal{O}(n^3)$ algorithms mentioned in the introduction are stable in the sense that they produce eigenpairs which are backward approximations of true eigenpairs. But for poorly conditioned matrices will not produce forward approximations of such pairs. These algorithms are called *direct* (as opposed to iterative) and consequently their complexity bound $\mathcal{O}(n^3)$ is considered a worst-case bound. But Demmel [11, p. 139] warns

> Note that "direct" methods must still iterate, since finding eigenvalues is mathematically equivalent to finding zeros of polynomials, for which no noniterative methods can exist. We call a method *direct* if experience shows that it (nearly) never fails to converge in a fixed number of iterations.

That is, if we want to ensure that a forward approximation of a given quality (say, forward error is at most ε) is returned for a poorly conditioned input, such algorithms will have to iterate and the number of iterations will depend on the condition of the input and on ε. If, in addition, we want to compute an approximation à la Smale —what we do in this paper—, then a supplementary computational effort is expected and there will be no dependence in ε in the associated cost.

Assuming an $\mathcal{O}(n^3)$ cost for each iteration, the fact that we no longer perform "a fixed number of iterations" entails an increase in the computation cost which, as pointed above, will depend on the condition of the input. And since this condition is infinity in the worst case (for ill-posed data), it is common to consider average complexity bounds that eliminate condition numbers from them.

The absence of this kind of analysis for the direct methods mentioned in the introduction, as well as the fact that they do not compute the same kind of approximation, makes their comparison with our results inappropriate: we are comparing apples and oranges.

Once said that, we believe that in practice these methods will certainly outperform the algorithm described in this paper in the sense that for most inputs (those reasonably well conditioned) they return a good forward approximation within a very satisfying time bound. But of course, the pursuit of the contrary —outperforming the current practical algorithms— was never the goal of this paper.

3 Proofs

3.1 Proof of Proposition 4

We associate with the solution path (3) in $\mathcal{V}$ the following curve in $\mathbb{S}(\mathcal{H}_n)\times\mathbb{R}\times\mathbb{C}^n$:

$$[0,1]\to V,\quad \tau\mapsto(P_\tau,\overline{\lambda_\tau},v_\tau):=\Big(\frac{Q_\tau}{\|Q_\tau\|_F},\frac{\lambda_\tau}{\|Q_\tau\|_F},v_\tau\Big),\tag{4}$$

where $\mathbb{S}(\mathcal{H}_n) := \{Q \in \mathcal{H}_n \mid \|Q\|_F = 1\}$. Recall that $\alpha = d_{\mathbb{S}}(M, A)$. The meaning of the parameterization by τ is that $\alpha\tau$ is the parameterization of $\tau \mapsto P_\tau$ by arc length, which means that $\left\| \frac{dP_\tau}{d\tau} \right\| = \alpha$.

Let now $[0,1] \to [0,1]$, $\tau \mapsto t(\tau)$, be any smooth bijective map such that $dt/d\tau > 0$. Then we have

$$\left\| \frac{dP_{\tau(t)}}{dt} \right\| = \left\| \frac{dP_\tau}{d\tau} \right\| \frac{d\tau}{dt} = \alpha \frac{d\tau}{dt},$$

and hence, by variable transformation,

$$\alpha \int_0^1 \mu^2(P_\tau, \lambda_\tau, v_\tau)\, d\tau \;=\; \int_0^1 \mu^2\big(P_{\tau(t)}, \lambda_{\tau(t)}, v_{\tau(t)}\big) \left\| \frac{dP_{\tau(t)}}{dt} \right\| dt. \tag{5}$$

In fact, for the probabilistic analysis later on, it will be essential to consider a different parameterization of $[M, A]$.

Proposition 7 *For all $\tau \in [0,1]$ we have $Q_\tau = tA + (1-t)M$, where $t = t(\tau)$ is given by*

$$t(\tau) = \frac{\|M\|_F}{\|A\|_F \sin\alpha \cot(\tau\alpha) - \|A\|_F \cos\alpha + \|M\|_F}.$$

Proof. We use some elementary geometry. For this, we introduce Cartesian coordinates (x, y) in the plane spanned by A and M (see Figure 1) and assume that M has the coordinates $(s, 0)$ and A has the coordinates $(r\cos\alpha, r\sin\alpha)$ so that $r = \|A\|_F$ and $s = \|M\|_F$.

Then, the lines determining Q_τ have the equations

$$x = y \frac{\cos(\tau\alpha)}{\sin(\tau\alpha)} \qquad \text{and} \qquad x = y \frac{r\cos\alpha - s}{r\sin\alpha} + s,$$

from which it follows that the coordinate y of Q_τ is

$$y = \frac{rs\sin\alpha\sin(\tau\alpha)}{r\sin\alpha\cos(\tau\alpha) - r\cos\alpha\sin(\tau\alpha) + s\sin(\tau\alpha)}.$$

Since $t(\tau) = \frac{y}{r\sin\alpha}$, we conclude that

$$t(\tau) = \frac{s}{r\sin\alpha\cot(\tau\alpha) - r\cos\alpha + s}. \qquad \square$$

Proof of Proposition 4. Set $\varepsilon := 0.1$ and hence $C_\varepsilon \approx 0.0087$. Furthermore, let $\xi := \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)^4} \approx 0.00178$.

We will carry out the proof on the curve (4) in the sphere $\mathbb{S}(\mathcal{H}_n)$. We do so to simplify the exposition and without implying that algorithm EC should be modified to normalize matrices. Indeed, all the quantities involved in our proof depending

on triples in $\mathcal{H}_n \times \mathbb{R} \times \mathbb{C}^n$ —distances $d_{\mathbb{S}\times\mathbb{P}}$, condition numbers μ, and Newton's method— are scale invariant on the first two components. Furthermore, to avoid burdening the notation, we will write λ instead of $\overline{\lambda}$. This should introduce no confusion.

Let $0 = \tau_0 < \tau_1 < \cdots < \tau_K = 1$ and $(\lambda_0, v_0) = (\nu_0, w_0), (\nu_1, w_1), \ldots, (\nu_K, w_K)$ be the sequences of τ-values and pairs in $\mathbb{R} \times \mathbb{C}^n$ generated by the algorithm EC. To simplify notation we write P_i instead of P_{τ_i} and (λ_i, v_i) instead of $(\lambda_{\tau_i}, v_{\tau_i})$.

We claim that for $i = 0, \ldots, K-1$, the following statements are true:

(a) $d_{\mathbb{S}\times\mathbb{P}}((P_i, \nu_i, w_i), (P_i, \lambda_i, v_i)) \leq \dfrac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)}$.

(b) $\dfrac{\mu(P_i, \nu_i, w_i)}{1+\varepsilon} \leq \mu(P_i, \lambda_i, v_i) \leq (1+\varepsilon)\mu(P_i, \nu_i, w_i)$.

(c) $d_{\mathbb{S}\times\mathbb{P}}((P_i, \lambda_i, v_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})) \leq \dfrac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \dfrac{2(1-\varepsilon)}{3(1+\varepsilon)}$.

(d) $d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_i, w_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})) \leq \dfrac{2C_\varepsilon}{(1+\varepsilon)\mu(P_i, \lambda_i, v_i)}$.

(e) (ν_i, w_i) is an approximate eigenpair of P_{i+1} with associated eigenpair (λ_{i+1}, v_{i+1}).

We proceed by induction, showing that

$$(\mathbf{a}, i) \Rightarrow (\mathbf{b}, i) \Rightarrow (\mathbf{c}, i) \Rightarrow (\mathbf{d}, i) \Rightarrow \big((\mathbf{e}, i) \text{ and } (\mathbf{a}, i+1)\big).$$

Inequality (a) for $i = 0$ is trivial.

Assume now that (a) holds for some $i \leq K-1$. Then, Proposition 3 (with $A = A' = P_i$) implies

$$\frac{\mu(P_i, \nu_i, w_i)}{1+\varepsilon} \leq \mu(P_i, \lambda_i, v_i) \leq (1+\varepsilon)\mu(P_i, \nu_i, w_i)$$

and thus (b). We now prove (c). To do so, let $\tau_* > \tau_i$ be such that

$$\int_{\tau_i}^{\tau_*} \left\| \frac{d(P_\tau, \lambda_\tau, v_\tau)}{d\tau} \right\| d\tau = \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \frac{2(1-\varepsilon)}{3(1+\varepsilon)}$$

or $\tau_* = 1$, whichever is smaller. Then, for all $t \in [\tau_i, \tau_*]$,

$$\begin{aligned} d_{\mathbb{S}\times\mathbb{P}}((P_i, \lambda_i, v_i), (P_t, \lambda_t, v_t)) &\leq \int_{\tau_i}^{t} \left\| \frac{d(P_\tau, \lambda_\tau, v_\tau)}{d\tau} \right\| d\tau \\ &\leq \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \frac{2(1-\varepsilon)}{3(1+\varepsilon)}. \end{aligned} \tag{6}$$

It is therefore enough to show that $\tau_{i+1} \leq \tau_*$. This is trivial if $\tau_* = 1$. We therefore assume $\tau_* < 1$. The bound above allows us to apply Proposition 3 and to deduce, for all $\tau \in [\tau_i, \tau_*]$,

$$\mu(P_\tau, \lambda_\tau, v_\tau) \leq (1+\varepsilon)\mu(P_i, \lambda_i, v_i).$$

Corollary 1 implies that

$$\left\| \frac{d}{d\tau}(P_\tau, \lambda_\tau, v_\tau) \right\| \leq 2\mu(P_\tau, \lambda_\tau, v_\tau) \left\| \frac{d}{d\tau} P_\tau \right\|$$

We now deduce that

$$\begin{aligned} \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \frac{2(1-\varepsilon)}{3(1+\varepsilon)} &= \int_{\tau_i}^{\tau_*} \left\| \frac{d(P_\tau, \lambda_\tau, v_\tau)}{d\tau} \right\| d\tau \leq \int_{\tau_i}^{\tau_*} 2\mu(P_\tau, \lambda_\tau, v_\tau) \left\| \frac{d}{d\tau} P_\tau \right\| d\tau \\ &\leq 2(1+\varepsilon)\, \mu(P_i, \lambda_i, v_i) \int_{\tau_i}^{\tau_*} \left\| \frac{d}{d\tau} P_\tau \right\| d\tau = 2(1+\varepsilon)\, \mu(P_i, \lambda_i, v_i)\, d_{\mathbb{S}}(P_i, P_{\tau_*}). \end{aligned}$$

Consequently, using (b), we obtain

$$d_{\mathbb{S}}(P_i, P_{\tau_*}) \geq \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)^2 \mu^2(P_i, \lambda_i, v_i)} \geq \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)^4 \mu^2(P_i, \nu_i, w_i)}.$$

Recall that the parameter ξ in EC was chosen as $\xi = \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)^4}$. By the definition of $\tau_{i+1} - \tau_i$ in EC we have $\alpha(\tau_{i+1} - \tau_i) = \frac{\xi}{\mu^2(P_i, \nu_i, w_i)}$. So we obtain

$$d_{\mathbb{S}}(P_i, P_{\tau_*}) \geq \alpha(\tau_{i+1} - \tau_i) = d_{\mathbb{S}}(P_i, P_{i+1}).$$

This implies $\tau_{i+1} \leq \tau_*$ as claimed, and hence inequality (c) follows from (6) with $t = \tau_{i+1}$. With it, we may apply Proposition 3 once more to deduce, for all $\tau \in [\tau_i, \tau_{i+1}]$,

$$\frac{\mu(P_i, \lambda_i, v_i)}{1+\varepsilon} \leq \mu(P_\tau, \lambda_\tau, v_\tau) \leq (1+\varepsilon)\mu(P_i, \lambda_i, v_i). \tag{7}$$

We now observe that

$$\begin{aligned} d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_i, w_i), (P_i, \nu_i, w_i)) &\leq d_{\mathbb{S}}(P_i, P_{i+1}) = \alpha(\tau_{i+1} - \tau_i) = \frac{\xi}{\mu^2(P_i, \nu_i, w_i)} \\ &\leq \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)\mu(P_i, \nu_i, w_i)} \end{aligned}$$

and use this bound, together with the triangle inequality, (a), and (c) to obtain

$$\begin{aligned} d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_i, w_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})) &\leq d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_i, w_i), (P_i, \nu_i, w_i)) \\ &\quad + d_{\mathbb{S}\times\mathbb{P}}((P_i, \nu_i, w_i), (P_i, \lambda_i, v_i)) \\ &\quad + d_{\mathbb{S}\times\mathbb{P}}((P_i, \lambda_i, v_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})) \\ &\leq \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)\mu(P_i, \nu_i, w_i)} + \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} + \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \frac{2}{3} \frac{1-\varepsilon}{1+\varepsilon} \\ &= \frac{2C_\varepsilon}{(1+\varepsilon)\mu(P_i, \lambda_i, v_i)}, \end{aligned} \tag{8}$$

which proves (d). We then use that $\frac{2C_\varepsilon}{1+\varepsilon} \approx 0.0158 < c_0$ and apply Theorem 1 to deduce that (ν_i, w_i) is an approximate eigenpair of P_{i+1} associated with its eigenpair (λ_{i+1}, v_{i+1}), and hence (e) holds.

It follows from (e) that $(\nu_{i+1}, w_{i+1}) = N_{P_{i+1}}(\nu_i, w_i)$ satisfies

$$d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_{i+1}, w_{i+1}), (P_{i+1}, \lambda_{i+1}, v_{i+1})) \leq \frac{1}{2}\, d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_i, w_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})).$$

Using this bound, (d) and the right-hand inequality in (7) with $\tau = \tau_{i+1}$, we obtain

$$d_{\mathbb{S}\times\mathbb{P}}((P_{i+1}, \nu_{i+1}, w_{i+1}), (P_{i+1}\lambda_{i+1}, v_{i+1})) \leq \frac{C_\varepsilon}{(1+\varepsilon)\mu(P_i, \lambda_i, v_i)} \leq \frac{C_\varepsilon}{\mu(P_{i+1}, \lambda_{i+1}, v_{i+1})},$$

which proves (a) for $i+1$. The claim is thus proved.

Note that (e) for $K-1$ shows that (ν_{K-1}, w_{K-1}) is an approximate eigenpair of $Q_K = A$ with associated eigenpair (λ_1, v_1) and consequently, so is the returned point $(\nu_K, w_K) = N_A(\nu_{K-1}, w_{K-1})$.

Consider now any $i \in \{0, \ldots, K-1\}$. Using (7), (b), and by the choice of the step size $\Delta\tau$ in Algorithm 1, we obtain

$$\begin{aligned}
\int_{\tau_i}^{\tau_{i+1}} \mu^2(P_\tau, \lambda_\tau, v_\tau) d\tau &\geq \int_{\tau_i}^{\tau_{i+1}} \frac{\mu^2(P_i, \lambda_i, v_i)}{(1+\varepsilon)^2} d\tau = \frac{\mu^2(P_i, \lambda_i, v_i)}{(1+\varepsilon)^2}(\tau_{i+1} - \tau_i) \\
&\geq \frac{\mu^2(P_i, \nu_i, w_i)}{(1+\varepsilon)^4}(\tau_{i+1} - \tau_i) \\
&= \frac{\mu^2(P_i, \nu_i, w_i)}{(1+\varepsilon)^4} \frac{\xi}{\alpha\mu^2(P_i, \nu_i, w_i)} \\
&= \frac{\xi}{(1+\varepsilon)^4\alpha} = \frac{C_\varepsilon(1-\varepsilon)}{3(1+\varepsilon)^8}\frac{1}{\alpha} \\
&\geq \frac{1}{822\,\alpha}.
\end{aligned}$$

This implies

$$\int_0^1 \mu^2(P_\tau, \lambda_\tau, v_\tau) d\tau \geq \frac{K}{822\,\alpha},$$

which proves the stated upper bound on K. The lower bound follows from

$$\begin{aligned}
\int_{\tau_i}^{\tau_{i+1}} \mu^2(P_\tau, \lambda_\tau, v_\tau) d\tau &\leq \int_{\tau_i}^{\tau_{i+1}} \mu^2(P_i, \lambda_i, v_i)(1+\varepsilon)^2 d\tau \\
&= \mu^2(P_i, \lambda_i, v_i)(1+\varepsilon)^2(\tau_{i+1} - \tau_i) \\
&\leq \mu^2(P_i, \nu_i, w_i)(1+\varepsilon)^4(\tau_{i+1} - \tau_i) \\
&= \frac{\xi(1+\varepsilon)^4}{\alpha} = \frac{C_\varepsilon(1-\varepsilon)}{3\alpha} \leq \frac{1}{383\,\alpha}.
\end{aligned}$$

□

3.2 Proof of Theorem 4

In order to apply Proposition 4, it will be central in our development to calculate the integral (5) of the squared condition number with respect to the parameterization t of $[M, A]$ introduced in Proposition 7. Abusing notation, we shall write $Q_t = (1 - t)M + tA$. For this parameterization we have the following bound on the norm of the speed of the spherical curve $t \mapsto P_t := \frac{Q_t}{\|Q_t\|_F}$.

Lemma 4 *We have*

$$\left\| \frac{dP_t}{dt} \right\| \leq \frac{\|A\|_F \, \|M\|_F}{\|Q_t\|_F^2}.$$

PROOF. Note that $\frac{dQ_t}{dt} = A - M$. Hence, if Pr denotes the orthogonal projection of $\mathcal{H}_n$ onto the tangent space $T_{Q_t}\mathbb{S}(\mathcal{H}_n)$, we have (we are using [8, Lemma 14.10]),

$$\frac{dP_t}{dt} \;=\; \frac{1}{\|Q_t\|_F} \, \mathsf{Pr}(A - M).$$

We show now by some elementary geometry that $\|\mathsf{Pr}(A - M)\|_F \leq \frac{\|A\|_F \, \|M\|_F}{\|Q_t\|_F}$. For this, as for Proposition 7, we introduce Cartesian coordinates in the plane spanned by A and M and assume that M has the coordinates $(s, 0)$ and A has the coordinates $(r\cos\alpha, r\sin\alpha)$; see Figure 3. We write $Q := Q_t$ and $L := \|A - M\|_F$. Then

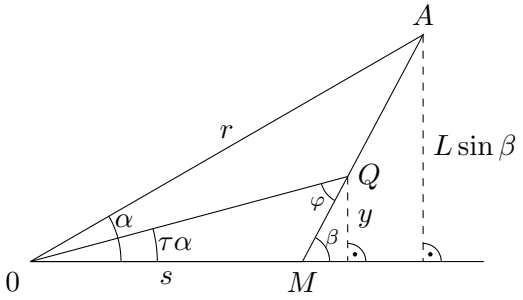

Figure 3: An elementary geometric argument.

$\|Q - M\|_F = tL$, and trigonometry tells us that

$$\frac{\sin\varphi}{\sin(\tau\alpha)} \;=\; \frac{s}{tL}.$$

Hence

$$\|\mathsf{Pr}(A - M)\|_F = L\sin\varphi = \frac{s}{t}\,\sin(\tau\alpha) = \frac{s}{t}\,\frac{y}{\|Q\|_F}.$$

We have

$$\frac{y}{t} \;=\; L\sin\beta \;\leq\; r,$$

and therefore

$$\|\mathsf{Pr}(A-M)\|_F \;\le\; \frac{r\,s}{\|Q\|_F} \;=\; \frac{\|A\|_F\,\|M\|_F}{\|Q_t\|_F}$$

as claimed. □

Recall that our goal is to bound

$$\mathsf{Num\_Iter}(n) = \mathop{\mathbb{E}}_{A\sim \mathrm{GUE}(n)} \mathop{\mathbb{E}}_{M\sim \mathrm{GUE}(n)} \frac{1}{n}\sum_{i=1}^{n} K(A,M,\lambda_i,v_i) \tag{9}$$

where $\lambda_1,\ldots,\lambda_n$ are the eigenvalues of M. The next proposition is a step in this direction.

Proposition 8 *We have*

$$\mathsf{Num\_Iter}(n) \le 1644\, n\pi \mathop{\mathbb{E}}_{Q\sim \mathrm{GUE}(n)} \sum_{i=1}^{n} \frac{\mu^2(Q,\lambda_i,v_i)}{\|Q\|_F^2}.$$

Proof. The starting point is the bound for $K(A,M,\lambda_i,v_i)$ in Proposition 4.

We consider the *truncated* GUE(n) on $\mathcal{H}_n$ (which we will denote by $\mathrm{GUE}_T(n)$) given by the density

$$\rho_T(A) = \begin{cases} \frac{\rho_{\mathcal{H}_n}(A)}{P_n} & \text{if } \|A\|_F \le n, \\ 0 & \text{otherwise,} \end{cases} \tag{10}$$

where $P_n := \mathsf{Prob}_{A\sim \mathrm{GUE}(n)}\left\{\|A\|_F \le n\right\}$, and $\rho_{\mathcal{H}_n}$ is the density of GUE(n). We now observe that, since the random variable $\|A\|_F^2$ is chi-square distributed with n^2 degrees of freedom (by Lemma 2(ii)), the expectation of such random variable is n, and since the median of a chi-square is bounded by its expectation (see [9, Corollary 6]), we have $P_n \ge \frac{1}{2}$.

Writing $\lambda_\tau^{(i)}$ to denote the continuation of the eigenvalue λ_i of M (and similarly for $v_\tau^{(i)}$) we obtain

$$\begin{aligned}
\mathsf{Num_Iter}(n) &\le 822 \mathop{\mathbb{E}}_{A\sim \mathrm{GUE}(n)} \mathop{\mathbb{E}}_{M\sim \mathrm{GUE}(n)} \frac{1}{n}\sum_{i=1}^{n} \int_0^1 \mu^2(Q_\tau,\lambda_\tau^{(i)},v_\tau^{(i)})\,d\tau \\
&= 822 \mathop{\mathbb{E}}_{A\sim \mathrm{GUE}_T(n)} \mathop{\mathbb{E}}_{M\sim \mathrm{GUE}_T(n)} \frac{1}{n}\sum_{i=1}^{n} \int_0^1 \mu^2(Q_\tau,\lambda_\tau^{(i)},v_\tau^{(i)})\,d\tau \\
&\le 822 \mathop{\mathbb{E}}_{A\sim \mathrm{GUE}_T(n)} \mathop{\mathbb{E}}_{M\sim \mathrm{GUE}_T(n)} \frac{1}{n}\sum_{i=1}^{n} \int_0^1 \frac{\|A\|_F\|M\|_F}{\|Q_t\|_F^2}\,\mu^2(Q_t,\lambda_t^{(i)},v_t^{(i)})\,dt.
\end{aligned}$$

The first inequality is a trivial consequence of Proposition 4. The equality after it follows from the fact that the expressions $\int_0^1 \mu^2(Q_\tau,\lambda_\tau^{(i)},v_\tau^{(i)})\,d\tau$ are scale invariant

with respect to both A and M. Hence, we can replace the GUE(n) distribution by its truncation. For the bottom inequality we used (5) and Lemma 4. We can next use that $\|A\|_F^2\|M\|_F^2 \leq n^2$ and replace the expectations by integrals against densities to bound the last expression as follows:

$$\begin{aligned}
822\,\frac{n^2}{P_n^2}\int_{\|A\|_F\leq n}\int_{\|M\|_F\leq n}\frac{1}{n}\sum_{i=1}^{n}\int_0^1 \frac{\mu^2(Q_t,\lambda_t^{(i)},v_t^{(i)})}{\|Q_t\|_F^2}\,dt\;\rho_{\mathcal{H}_n}(M)\,\rho_{\mathcal{H}_n}(A)\,dM\,dA \\
&\leq\; 3288\,n^2 \mathop{\mathbb{E}}_{A\sim\mathrm{GUE}(n)}\mathop{\mathbb{E}}_{M\sim\mathrm{GUE}(n)}\frac{1}{n}\sum_{i=1}^{n}\int_0^1\frac{\mu^2(Q_t,\lambda_t^{(i)},v_t^{(i)})}{\|Q_t\|_F^2}\,dt \\
&=\; 3288\,n^2\frac{1}{n}\mathop{\mathbb{E}}_{Q\sim\mathrm{GUE}(n)}\sum_{i=1}^{n}\left(\frac{\mu^2(Q,\lambda_i,v_i)}{\|Q\|_F^2}\right)\int_0^1\frac{1}{(1-t)^2+t^2}dt \\
&=\; 1644\,n\pi\mathop{\mathbb{E}}_{Q\sim\mathrm{GUE}(n)}\sum_{i=1}^{n}\frac{\mu^2(Q,\lambda_i,v_i)}{\|Q\|_F^2}
\end{aligned}$$

where the equality before the last follows from Lemma 2(i) and the last from the fact that the integral evaluates to $\frac{\pi}{2}$. □

The proof of Theorem 4 now trivially follows by replacing in the right-hand side of Proposition 8, the expectation $\mathbb{E}_{Q\sim\mathrm{GUE}(n)}\sum_{i=1}^{n}\frac{\mu^2(Q,\lambda_i,v_i)}{\|Q\|_F^2}$ by $n(n-1)$ (we use Theorem 3) and noting that

$$\mathsf{Cost}(n) = \mathcal{O}(n^4) + \mathcal{O}(n^3)\,\mathsf{Num\_Iter}(n) = \mathcal{O}(n^6)$$

where the first term in the addition comes from Proposition 6.

3.3 Proof of Theorem 3

Given $(A,\lambda,v)\in\mathcal{W}$, we define the *Frobenius condition number* μ_F as

$$\mu_F(A,\lambda,v) := \|A\|_F\,\|{A_{\lambda,v}}^{-1}\|_F.$$

Since the operator norm of matrices is bounded from above by the Frobenius norm, and from the fact that the expression $\|A\|_F\,\|{A_{\lambda,v}}^{-1}\|$ is bounded from below by $1/\sqrt{2}$ (cf. Lemma 3.8 in [2]), we have that

$$\mu(A,\lambda,v) \leq \sqrt{2}\|A\|_F\,\|{A_{\lambda,v}}^{-1}\| \leq \sqrt{2}\mu_F(A,\lambda,v), \tag{11}$$

for every $(A,\lambda,v)\in\mathcal{W}$.

The proof of Theorem 3 relies on several lemmas the first of which is a crucial linear algebra result.

Lemma 5 *Let $A\in\mathbb{C}^{n\times n}$ and $v\in\mathbb{C}^n$ such that $\|v\|=1$. Let $\widehat{A}: v^\perp\to v^\perp$ be given by $\widehat{A} := \Pi_{v^\perp}A|_{v^\perp}$, and $w := \Pi_{v^\perp}(Av)\in v^\perp$. If A and $\widehat{A}$ are invertible, then,*

$$\big|\det(A)\big|^2\cdot\big\|A^{-1}v\big\|^2 = \big|\det(\widehat{A})\big|^2 + \big|\det(\widehat{A})\big|^2\cdot\big\|\widehat{A}^{-1}w\big\|^2.$$

PROOF. Fix an orthonormal basis of $\mathbb{C}^n$, with first element $e_1 = v$. Then (slightly abusing notation) we may write $A = \begin{pmatrix} a & * \\ w & \widehat{A} \end{pmatrix}$. By Cramer's rule we have,

$$\left\|A^{-1}e_1\right\|^2 \cdot \left|\det(A)\right|^2 = \sum_{i=1}^{n} \left|\det([A_i; e_1])\right|^2, \tag{12}$$

where $[A_i; e_1]$ is the matrix formed by replacing the ith column of A by the column vector e_1. That is,

$$\left\|A^{-1}e_1\right\|^2 \cdot \left|\det(A)\right|^2 = \left|\det(\widehat{A})\right|^2 + \sum_{i=2}^{n} \left|\det([A_i; e_1])\right|^2.$$

Furthermore, it is immediate to check that $|\det([A_i; e_1])| = |\det([(\widehat{A})_{i-1}; w])|$, for $i = 2, \ldots, n$. The proof follows replacing in (12), A and e_1 by $\widehat{A}$ and w, respectively. □

Lemma 6 *Let $A \in \mathbb{C}^{m\times m}$. If η is a standard Gaussian vector in $\mathbb{C}^m$, then*

$$\mathop{\mathbb{E}}_{\eta\sim\mathcal{N}_{\mathbb{C}}(0,\mathbf{1}_m)} \|A\eta\|^2 = \|A\|_F^2.$$

PROOF. The proof is similar to that of [1, Proposition 2]. We give it nonetheless, for the sake of completeness.

Let $A = UDV$ be a singular value decomposition of A, where U and V are unitary matrices and $D = \mathsf{diag}(\sigma_1, \ldots, \sigma_m)$ is a diagonal matrix with real positive entries. By the unitary invariance of the Gaussian distribution $\mathcal{N}_{\mathbb{C}}(0, \mathbf{1}_m)$, $V\eta$ has the same distribution as η. Furthermore, by the unitary invariance of the norm in $\mathbb{C}^m$ we obtain

$$\mathop{\mathbb{E}}_{\eta\sim\mathcal{N}_{\mathbb{C}}(0,\mathbf{1}_m)} \|A\eta\|^2 = \mathop{\mathbb{E}}_{\eta\sim\mathcal{N}_{\mathbb{C}}(0,\mathbf{1}_m)} \|D\eta\|^2 = \sum_{i=1}^{m} \sigma_i^2 \mathop{\mathbb{E}}_{\eta_i\sim\mathcal{N}_{\mathbb{C}}(0,1)} |\eta_i|^2.$$

Since $\mathbb{E}_{z\sim\mathcal{N}_{\mathbb{C}}(0,1)} |z|^2 = 1$, the statement follows. □

In all what follows, for shorteness, we will write

$$\mathbb{E}_n := \mathop{\mathbb{E}}_{Q\sim\mathrm{GUE}(n)} \sum_{k=1}^{n} \frac{\mu_F^2(Q, \lambda_k, v_k)}{\|Q\|_F^2},$$

where (λ_k, v_k), for $k = 1, \ldots, n$, denote the eigenpairs of Q.

Lemma 7

$$\mathbb{E}_n = \frac{1}{(n-1)!} \mathop{\mathbb{E}}_{\lambda\sim\mathcal{N}(0,1)} \mathop{\mathbb{E}}_{B\sim\mathrm{GUE}(n-1)} \left(\left\|(B - \lambda\,\mathsf{Id}_{n-1})^{-1}\right\|_F^2 \cdot \left|\det(B - \lambda\,\mathsf{Id}_{n-1})\right|^2\right)$$

where $\mathcal{N}(0,1)$ denotes the standard Gaussian distribution on $\mathbb{R}$.

PROOF. Let $Q \in \mathcal{H}_n \setminus \Sigma$, and let $(\lambda_1, v_1), \dots, (\lambda_n, v_n)$ be its eigenpairs. Then, for $k = 1, \dots, n$,

$$\frac{\mu_F(Q, \lambda_k, v_k)^2}{\|Q\|_F^2} = \sum_{i=1, i\neq k}^{n} \frac{1}{|\lambda_i - \lambda_k|^2}.$$

By Corollary 2, and the invariance under permutations of the joint density of the eigenvalues, we get

$$\begin{aligned}\mathbb{E}_n &= \mathop{\mathbb{E}}_{(\lambda_1,\dots,\lambda_n)\sim W_n} \sum_{k=1}^{n} \left(\sum_{i=1,\, i\neq k}^{n} \frac{1}{|\lambda_i - \lambda_k|^2} \right) = n \mathop{\mathbb{E}}_{(\lambda_1,\dots,\lambda_n)\sim W_n} \left(\sum_{i=1}^{n-1} \frac{1}{|\lambda_i - \lambda_n|^2} \right) \\ &= \frac{n}{Z_n} \int_{\mathbb{R}^n} \left(\sum_{i=1}^{n-1} \frac{1}{|\lambda_i - \lambda_n|^2} \right) \Delta_n^2(\lambda_1, \dots, \lambda_n) \exp\left(-\sum_{i=1}^{n} \frac{\lambda_i^2}{2} \right) d\lambda_1 \dots d\lambda_n.\end{aligned}$$

Furthermore, applying Fubini's theorem and writing λ in the place of λ_n, we have

$$\begin{aligned}\mathbb{E}_n = \frac{nZ_{n-1}}{Z_n} \int_{\mathbb{R}} \exp\left(-\frac{\lambda^2}{2}\right) \Bigg\{ \frac{1}{Z_{n-1}} \int_{\mathbb{R}^{n-1}} \left(\sum_{i=1}^{n-1} \frac{1}{|\lambda_i - \lambda|^2} \right) \prod_{i=1}^{n-1} (\lambda_i - \lambda)^2 \times \\ \Delta_{n-1}^2(\lambda_1, \dots, \lambda_{n-1}) \exp\left(-\sum_{i=1}^{n-1} \frac{\lambda_i^2}{2} \right) d\lambda_1 \dots d\lambda_{n-1} \Bigg\} d\lambda,\end{aligned}$$

that is,

$$\mathbb{E}_n = \frac{nZ_{n-1}}{Z_n} \int_{\mathbb{R}} \exp\left(-\frac{\lambda^2}{2}\right) \left[\mathop{\mathbb{E}}_{(\lambda_1,\dots,\lambda_{n-1})\sim W_{n-1}} \left(\sum_{i=1}^{n-1} \frac{1}{|\lambda_i - \lambda|^2} \right) \prod_{i=1}^{n-1} (\lambda_i - \lambda)^2 \right] d\lambda.$$

Using Corollary 2 we obtain

$$\mathbb{E}_n = \frac{nZ_{n-1}}{Z_n} \sqrt{2\pi} \mathop{\mathbb{E}}_{\lambda\sim\mathcal{N}(0,1)} \mathop{\mathbb{E}}_{B\sim \mathrm{GUE}(n-1)} \left(\left\| (B - \lambda\, \mathsf{Id}_{n-1})^{-1} \right\|_F^2 \cdot \left| \det(B - \lambda\, \mathsf{Id}_{n-1}) \right|^2 \right).$$

The statement now follows from the definition of the normalization constant Z_n given in Theorem 2. □

For fixed $\lambda \in \mathbb{R}$, and $m \geq 1$, let

$$\begin{aligned} c_m(\lambda) &:= \mathop{\mathbb{E}}_{B\sim \mathrm{GUE}(m)} \left(\left\| (B - \lambda\, \mathsf{Id}_m)^{-1} \right\|_F^2 \cdot \left| \det(B - \lambda\, \mathsf{Id}_m) \right|^2 \right); \\ d_m(\lambda) &:= \mathop{\mathbb{E}}_{B\sim \mathrm{GUE}(m)} \left(\left| \det(B - \lambda\, \mathsf{Id}_m) \right|^2 \right). \end{aligned}$$

Lemma 8 *For every $\lambda \in \mathbb{R}$ and $m \geq 2$, one has*

$$c_m(\lambda) = m! \left(1 + \sum_{k=1}^{m-1} \frac{d_k(\lambda)}{k!} \right).$$

PROOF. Let $\eta \sim \mathcal{N}_{\mathbb{C}}(0, \mathbf{1}_m)$ be independent of $B \sim \mathrm{GUE}(m)$. Then, by Lemma 6 and Fubini's theorem, one has,

$$c_m(\lambda) = \mathop{\mathbb{E}}_{\eta \sim \mathcal{N}_{\mathbb{C}}(0,\mathbf{1}_m)} \mathop{\mathbb{E}}_{B \sim \mathrm{GUE}(m)} \left(\left\| (B - \lambda \, \mathsf{Id}_m)^{-1} \eta \right\|^2 \cdot \left| \det(B - \lambda \, \mathsf{Id}_m) \right|^2 \right).$$

Since η is a standard Gaussian in $\mathbb{C}^m$, integrating in polar coordinates, it is easily seen that $u := \eta / \|\eta\|$ and $\|\eta\|$ are independent random variables. Furthermore, from Lemma 6 for $A = \mathsf{Id}_m$, it follows that $\mathbb{E}_{\eta \sim \mathcal{N}_{\mathbb{C}}(0,\mathbf{1}_m)}(\|\eta\|^2) = m$. Then

$$c_m(\lambda) = m \mathop{\mathbb{E}}_{u \sim \mathcal{U}(\mathbb{S}^{2m-1})} \mathop{\mathbb{E}}_{B \sim \mathrm{GUE}(m)} \left(\left\| (B - \lambda \, \mathsf{Id}_m)^{-1} u \right\|^2 \cdot \left| \det(B - \lambda \, \mathsf{Id}_m) \right|^2 \right),$$

where $\mathcal{U}(\mathbb{S}^{2m-1})$ denotes the uniform measure on the real unit sphere of $\mathbb{C}^m$.

By the unitary invariance of the $\mathrm{GUE}(m)$ distribution and the fact that $\lambda \, \mathsf{Id}_m$ is fixed under the conjugation action, we have that the distribution of $B - \lambda \, \mathsf{Id}_m$ is invariant under conjugations by unitary matrices. Furthermore, from the fact that the Euclidean norm is unitarily invariant, we conclude

$$c_m(\lambda) = m \mathop{\mathbb{E}}_{B \sim \mathrm{GUE}(m)} \left(\left\| (B - \lambda \, \mathsf{Id}_m)^{-1} e_1 \right\|^2 \cdot \left| \det(B - \lambda \, \mathsf{Id}_m) \right|^2 \right),$$

where e_1 is the first element of the canonical basis of $\mathbb{C}^m$. Then, by Lemma 5,

$$c_m(\lambda) = m \mathop{\mathbb{E}}_{B \sim \mathrm{GUE}(m)} \left(\left| \det(B_{\lambda,e_1}) \right|^2 + \left\| (B_{\lambda,e_1})^{-1} b \right\|^2 \left| \det(B_{\lambda,e_1}) \right|^2 \right),$$

where $B_{\lambda,e_1} := \Pi_{e_1^\perp} (B - \lambda \mathsf{Id}_m)|_{e_1^\perp}$, and $b := \Pi_{e_1^\perp}(B - \lambda \, \mathsf{Id}_m) e_1 = \Pi_{e_1^\perp}(B e_1)$. Since $B \sim \mathrm{GUE}(m)$, we have that $\Pi_{e_1^\perp} B|_{e_1^\perp}$ and b are independent $\mathrm{GUE}(m-1)$ and $\mathcal{N}_{\mathbb{C}}(0, \mathbf{1}_{m-1})$, respectively. Hence, again by Lemma 6 and the definition of $d_m(\lambda)$, we conclude that

$$c_m(\lambda) = m(d_{m-1}(\lambda) + c_{m-1}(\lambda)).$$

Working by induction and using the fact that $c_1(\lambda) = 1$, the result follows. □

Lemma 9 *For every $k = 1, 2, \ldots$, one has*

$$\mathop{\mathbb{E}}_{\lambda \sim \mathcal{N}(0,1)} d_k(\lambda) = (k+1)!.$$

PROOF. By the definition of the normalizing constant Z_n we have

$$1 = \frac{1}{Z_{k+1}} \int_{\mathbb{R}^{k+1}} \Delta_{k+1}^2(\lambda_1, \ldots, \lambda_{k+1}) \exp\left(-\sum_{i=1}^{k+1} \frac{\lambda_i^2}{2} \right) d\lambda_1, \ldots, d\lambda_{k+1}.$$

Then, by arguments similar to those in the proof of Lemma 7, we have

$$\begin{aligned} 1 &= \frac{Z_k}{Z_{k+1}}\sqrt{2\pi} \mathop{\mathbb{E}}_{\lambda\sim\mathcal{N}(0,1)} \mathop{\mathbb{E}}_{B\sim\mathrm{GUE}(k)} \left(\left|\det(B-\lambda\,\mathsf{Id}_k)\right|^2\right) \\ &= \frac{Z_k}{Z_{k+1}}\sqrt{2\pi} \mathop{\mathbb{E}}_{\lambda\sim\mathcal{N}(0,1)} d_k(\lambda) \;=\; \frac{1}{(k+1)!} \mathop{\mathbb{E}}_{\lambda\sim\mathcal{N}(0,1)} d_k(\lambda) \end{aligned}$$

from where the statement follows. □

Proof of Theorem 3. From Lemmas 7 and 8 one gets

$$\mathbb{E}_n = \frac{1}{(n-1)!} \mathop{\mathbb{E}}_{\lambda\sim\mathcal{N}(0,1)} (c_{n-1}(\lambda)) = 1 + \sum_{k=1}^{n-2} \frac{\mathbb{E}_{\lambda\sim\mathcal{N}(0,1)}(d_k(\lambda))}{k!}.$$

Then, Lemma 9 yields

$$\mathbb{E}_n = 1 + \sum_{k=1}^{n-2}(k+1) = \frac{n(n-1)}{2},$$

and the bound in (11) finishes the proof. □

3.4 Proof of Proposition 2

Let $k \leq n$ and $V_{n,k} := \{A \in \mathcal{H}_n \mid \mathsf{rank}(A) = k\}$. The set $V_{n,k}$ is a locally closed subset of $\mathcal{H}_n$ in the Zariski topology.

Lemma 10 (Proposition 1.1 in [6]) *For $0 \leq k \leq n$ we have $\dim V_{n,k} = k(2n-k)$.* □

Proof of Proposition 2. Let $W_2 := \{A \in \mathcal{H}_n \mid \mathsf{rank}(A) \leq n-2\}$. Then W_2 is the disjoint union $W_2 = \bigsqcup_{k\leq n-2} V_{n,k}$ and we have

$$\dim W_2 = \max_{k\leq n-2} \dim V_{n,k} = \dim V_{n,n-2} = (n-2)(2n-n+2) = n^2-4.$$

Now consider the function

$$\begin{aligned} \Phi : \mathcal{H}_n \times \mathbb{R} &\to \mathcal{H}_n \\ (A,\lambda) &\mapsto A - \lambda\mathsf{Id}. \end{aligned}$$

It is immediate to see that Φ is a linear map (and hence smooth). Also, since matrices in $\mathcal{H}_n$ diagonalize, that $\Phi(W_2 \times \mathbb{R}) = \Sigma$ (this equality would be false for arbitrary matrices in $\mathbb{C}^{n\times n}$). This shows that $\dim\Sigma \leq \dim W_2 + 1 = n^2 - 3$ and the fact that for all $A \in \Sigma$ the fiber $\Phi^{-1}_{|W_2\times\mathbb{R}}(A) \subset W_2 \times \mathbb{R}$ is finite (it is actually the number of different multiple eigenvalues of A) shows that this inequality is indeed an equality. □

3.5 Proof of Proposition 1

Proof of Proposition 1. Since the system of equations defining the solution variety is homogeneous in the variable v, it is enough to prove that the set

$$\hat{\mathcal{V}} := \{(A, \lambda, v) \in \mathcal{H}_n \times \mathbb{R} \times (\mathbb{C}^n \setminus \{0\}) : \ (\lambda \mathsf{Id} - A)v = 0\},$$

is a smooth manifold of real dimension $n^2 + 2$.

Let $\mathrm{Skew}(n) := \{B \in \mathbb{C}^{n\times n} : \ B^* = -B\}$ be the set of $n \times n$ skew-symmetric matrices. Note that the real dimension of $\mathrm{Skew}(n)$ is n^2.

Let $F_1 : \mathbb{C}^{n\times n} \times \mathbb{C} \times (\mathbb{C}^n \setminus \{0\}) \to \mathbb{C}^n$ given by

$$F_1(A, \lambda, v) = (\lambda \mathsf{Id} - A)v,$$

and let $F_2 : \mathbb{C}^{n\times n} \times \mathbb{C} \times (\mathbb{C}^n \setminus \{0\}) \to \mathrm{Skew}(n)$ given by

$$F_2(A, \lambda, v) = A - A^*.$$

Note that $\hat{\mathcal{V}} = F^{-1}(0, 0)$, where $F = (F_1, F_2)$. We will prove that $(0, 0)$ is a regular value of F. If this is the case, it is easily seen that the dimension of $\hat{\mathcal{V}}$ is $2n^2 + 2 + 2n - 2n - n^2 = n^2 + 2$.

The derivative $DF_2(A, \lambda, v)$ is a surjective map onto the space $\mathrm{Skew}(n)$. Therefore, it is enough for us to show that $\mathbb{C}^n \times \{0\} \subset \mathtt{range} DF(A, \lambda, v)$.

Claim: For all $v \in \mathbb{C}^n \setminus \{0\}$ one has $\{\Pi_{v^\perp} Bv : \ B \in \mathcal{H}_n\} = v^\perp$.

To prove this claim we first note that the space $\mathcal{H}_n$ is invariant under unitary conjugations. Hence, we may assume that v is the first vector of the canonical basis in $\mathbb{C}^n$ (up to a scaling by a positive number). Now the claim follows immediately from the matrix entries' structure of an Hermitian matrix.

Since at any triple (A, λ, v) the derivative of F_1 applied to $(\dot{A}, \dot{\lambda}, 0)$ is given by $\dot{\lambda}v - \dot{A}v$, we conclude from the Claim (and the fact that $\mathcal{H}_n \times \mathbb{C} \times \{0\}$ is in the kernel of $DF_2(A, \lambda, v)$) that $DF(A, \lambda, v)\,(\mathcal{H}_n \times \mathbb{C} \times \{0\}) = \mathbb{C}^n \times \{0\}$. □

References

[1] D. Armentano. Stochastic perturbations and smooth condition numbers. *J. Complexity*, 26(2):161–171, 2010.

[2] D. Armentano. Complexity of Path-Following Methods for the Eigenvalue Problem. *Found. Comput. Math.*, 14(2):185–236, 2014.

[3] C. Beltrán and L.M. Pardo. Smale's 17th problem: average polynomial time to compute affine and projective solutions. *J. Amer. Math. Soc.*, 22(2):363–385, 2009.

[4] C. Beltrán and L.M. Pardo. Fast linear homotopy to find approximate zeros of polynomial systems. *Found. Comput. Math.*, 11(1):95–129, 2011.

[5] I. Briquel, F. Cucker, J. Peña, and V. Roshchina. Fast computation of zeros of polynomial systems with bounded degree under finite-precision. To appear in *Math. of Computation*, 2013.

[6] W. Bruns and U. Vetter. *Determinantal rings*, volume 1327 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1988.

[7] P. Bürgisser and F. Cucker. On a problem posed by Steve Smale. *Annals of Mathematics*, 174:1785–1836, 2011.

[8] P. Bürgisser and F. Cucker. *Condition*, volume 349 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2013.

[9] K.P. Choi. On the medians of gamma distributions and an equation of Ramanujan. *Proc. Amer. Math. Soc.*, 121(1):245–251, 1994.

[10] P. Deift. Some open problems in random matrix theory and the theory of integrable systems. In *Integrable systems and random matrices*, volume 458 of *Contemp. Math.*, pages 419–430. Amer. Math. Soc., Providence, RI, 2008.

[11] J.W. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.

[12] L. Devroye. *Nonuniform random variate generation*. Springer-Verlag, New York, 1986.

[13] Alan Edelman and N. Raj Rao. Random matrix theory. *Acta Numer.*, 14:233–297, 2005.

[14] W. Hrmann, J. Leydold, and G. Derflinger. *Automatic nonuniform random variate generation*. Springer-Verlag, New York, 2004.

[15] T.-Y. Li and T. Sauer. Homotopy method for generalized eigenvalue problems $Ax = \lambda Bx$. *Linear Algebra Appl.*, 91:65–74, 1987.

[16] X.H. Li and G. Menon. Numerical solution of Dyson Brownian motion and a sampling scheme for invariant matrix ensembles. *J. Stat. Phys.*, 153(5):801–812, 2013.

[17] M.L. Mehta. *Random matrices*, volume 142 of *Pure and Applied Mathematics (Amsterdam)*. Elsevier/Academic Press, Amsterdam, third edition, 2004.

[18] S. Olver and A. Townsend. Fast inverse transform sampling in one and two dimensions. Available at `http://arxiv.org/pdf/1307.1223.pdf`, 2013.

[19] B.N. Parlett. *The symmetric eigenvalue problem*, volume 20 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998. Corrected reprint of the 1980 original.

[20] C.W. Pfrang, P. Deift, and G. Menon. How long does it take to compute the eigenvalues of a random symmetric matrix? Available at `http://arxiv.org/pdf/1203.4635.pdf`, 2013.

[21] M. Shub. Complexity of Bézout's Theorem VI geodesics in the condition (number) metric. *Found. Comput. Math.*, 9:171–178, 2009.

[22] M. Shub and S. Smale. Complexity of Bézout's Theorem I: geometric aspects. *Journal of the Amer. Math. Soc.*, 6:459–501, 1993.

[23] S. Smale. Newton's method estimates from data at one point. In R. Ewing, K. Gross, and C. Martin, editors, *The Merging of Disciplines: New Directions in Pure, Applied, and Computational Mathematics*. Springer-Verlag, 1986.

[24] S. Smale. Mathematical problems for the next century. In V. Arnold, M. Atiyah, P. Lax, and B. Mazur, editors, *Mathematics: Frontiers and Perspectives*, pages 271–294. AMS, 2000.