

SAMPLING UNITARY INVARIANT ENSEMBLES

SHEEHAN OLVER*, RAJ RAO NADAKUDITI†, AND THOMAS TROGDON‡

Abstract.

We develop an algorithm for sampling from the unitary invariant random matrix ensembles. The algorithm is based on the representation of their eigenvalues as a determinantal point process whose kernel is given in terms of orthogonal polynomials. Using this algorithm, statistics beyond those known through analysis are calculable through Monte Carlo simulation. Unexpected phenomena are observed in the simulations.

1. Introduction. The *unitary invariant ensembles* (UIE) are an important class of random matrices which are invariant under conjugation by unitary matrices. This corresponds to the physical situation where the frame of reference does not affect the underlying statistics. UIEs are defined as $n \times n$ random Hermitian matrices

$$M = \begin{pmatrix} M_{11} & M_{12}^R + iM_{12}^I & \cdots & M_{1n}^R + iM_{1n}^I \\ M_{12}^R - iM_{12}^I & M_{22} & \cdots & M_{2n}^R + iM_{2n}^I \\ \vdots & \ddots & \ddots & \vdots \\ M_{1n}^R - iM_{1n}^I & \cdots & M_{(n-1)n}^R - iM_{(n-1)n}^I & M_{nn} \end{pmatrix}$$

whose entries are distributed according to a given *potential* Q , by the rule

$$\frac{1}{Z_n} e^{-\text{Tr } Q(M)} dM,$$

where Z_n is the normalization constant and

$$dM = \prod_{i=1}^n dM_{ii} \prod_{i < j} (dM_{ij}^R dM_{ij}^I).$$

Associated with the UIE potential Q is the weight $w(x) = e^{-Q(x)}$. We include weights of the form $x^\alpha e^{-Q(x)}$, which corresponds to the UIE

$$e^{-\text{Tr } [Q(M) - \alpha \log M]} = (e^{\text{Tr } \log M})^\alpha e^{-\text{Tr } Q(M)} = (\det M)^\alpha e^{-\text{Tr } Q(M)},$$

using $\det e^M = e^{\text{Tr } M}$. We also allow Q to depend on n , particularly $Q(x) = nV(x)$.

The contribution of this paper is the development of an efficient algorithm for sampling from invariant ensembles. Sampling the entries directly is prohibitively expensive: it is difficult to sample many random variables that depend on each other in a complicated manner. Instead we exploit the fact that the eigenvalues of invariant ensembles are described by determinantal point processes whose kernel is written in terms of the associated orthogonal polynomials. Thus, the task of sampling invariant ensembles is reduced to the following:

1. Construct the orthogonal polynomials associated to the weight w .
2. Sample a determinantal point process defined through this sequence of orthogonal polynomials.

*The University of Sydney, Australia. (Sheehan.Olver@sydney.edu.au)

†University of Michigan, USA. (rajnrao@umich.edu)

‡Courant Institute, New York University, USA. (trogdon@cims.nyu.edu)

The first task can be accomplished via either Stieljes procedure [16], or using Riemann–Hilbert techniques [32]. For the second task, we adapt a recently developed algorithm [17, 28].

Very recently, an alternative approach for sampling invariant ensembles based on simulating Dyson Brownian motion was developed by Li and Menon [18]. Both approaches have the same complexity of $\mathcal{O}(n^3)$ operations to sample an $n \times n$ matrix, although the method of Li and Melon is likely to be significantly faster. However, our approach samples the correct distribution to essentially machine precision accuracy, a task that is computationally impractical using numerical simulation of stochastic differential equations. See Section 3.5 and Figure 3.6 for a demonstration and discussion of this.

Our motivation for sampling invariant ensembles is to extend the class of matrices for which simulations can be performed, to facilitate a deeper understanding of the rich phenomena that random matrices exhibit. Currently, the ensembles that can easily be sampled are restricted to those that are generated from independent entries. These include the following classical ensembles:

- *Wigner ensembles*: Hermitian ensembles with independent entries.
- *Wishart ensembles*: XX^* where X is rectangular with independent complex entries.
- *Manova ensembles*: $A(A + B)^{-1}$ for $A = a^*a$, $B = b^*b$ and a, b rectangular with independent complex entries.

We note that three of the most studied ensembles lie right at the intersection of these ensembles and invariant ensembles:

- *Gaussian Unitary Ensemble (GUE)*: Wigner ensemble whose entries are independent complex Gaussian with twice the variance on the diagonal, is also a UIE with potential $Q(x) = x^2$.
- *Laquerre Unitary Ensemble (LUE)*: Wishart ensemble where X is rectangular with iid complex Gaussian entries, is also an UIE with weight $x^\alpha e^{-x}$ on the half line.
- *Jacobi Unitary Ensemble (JUE)*: Manova ensemble where a, b are rectangular with iid complex Gaussian entries, is also a UIE with weight $(1 - x)^\alpha (1 + x)^\beta$ [34, pp. 1440].

While ensembles generated from independent entries are easily sampled — particularly GUE, LUE and JUE which have banded representations that enable faster computation [12, 14] — they are by no means general: under broad conditions [33, 26], Wigner ensembles follow the semicircle law — i.e., the global distribution of eigenvalues tends to a semicircle — while GUE is the only invariant ensemble that follows the semicircle law. Similarly, Wishart ensembles follow the Marčenko–Pastur law [20]. On the other hand, whereas simulation has been primarily limited to ensembles generated from independent entries, analysis of invariant ensembles is quite developed (cf. [9, 10]).

A particular phenomena studied in depth is *universality*, and universality laws for many canonical families of random matrices, in particular Wigner, Wishart and invariant ensembles, have been proved in the bulk [9, 19, 15] and edge [30, 10]. The techniques of these proofs differ greatly, hinting that universality is applicable to a much wider class of matrices, beyond the reach of current analytical approaches. Using the results of this paper, it is now possible to guide this conjecture through simulation. For example, we include in Section 3 some experiments that suggest algebraic manipulations of invariant ensembles also follow universality laws.

Notably, a very special subset of invariant ensembles have been shown to violate the Tracy–Widom universality law; rather, they follow higher-order analogues of the Tracy–Widom law [5]. While this was discovered analytically, using the sampling algorithm we are able to verify it through simulation. We further show that adding a sufficient number of these degenerate ensembles (precisely, four) causes the standard Tracy–Widom law to reemerge. This demonstrates why broadening the possible simulations is important: if one is restricted to ensembles generated from independent entries, such phenomena may have gone undiscovered.

The paper is organized as follows:

Section 2: We give a brief overview of invariant ensembles and the distribution of their eigenvalues.

Section 3: We present experimental results that are achieved using the sampling algorithm, verifying that special classes of ensembles do indeed follow alternative universality results. We also compare the sampling accuracy of our approach to that of Li and Menon [18].

Section 4: We give background on orthogonal polynomials and recurrence relationships, and describe their approximation by Chebyshev series for use in the algorithm.

Section 5: We close our discussion with a brief overview of determinantal point processes, and present the algorithm as specialized for the orthogonal polynomial setting. We prove that the algorithm does indeed sample unitary invariant ensembles.

2. Unitary Invariant Ensembles. We begin with a brief overview of how the eigenvalues of a UIE are reduced to a determinantal point process. The eigenvectors of any UIE are simply Haar distributed unitary matrices¹. Upon integrating out the eigenvectors, the distribution of the eigenvalues is determined to be [9, Section 5.4]

$$\frac{1}{\hat{Z}_n} \prod_{i < j} (\lambda_i - \lambda_j)^2 w(\lambda_1) \cdots w(\lambda_n) d\boldsymbol{\lambda} \quad (2.1)$$

where \hat{Z}_n is a normalization constant.

This distribution can be rewritten in terms of the determinant of orthogonal polynomials. Let p_0, p_1, \dots be the polynomials orthonormal with respect to $w(x)dx$, and define

$$\phi_k(x) := p_k(x) \sqrt{w(x)},$$

which are orthonormal in $L^2(\mathbb{R})$. Define the kernel

$$\mathcal{K}_n(x, y) = \sum_{k=0}^{n-1} \phi_k(x) \phi_k(y). \quad (2.2)$$

then we have [9, (5.30)]

$$(2.1) = \frac{1}{n!} \det \begin{pmatrix} \mathcal{K}_n(\lambda_1, \lambda_1) & \cdots & \mathcal{K}_n(\lambda_n, \lambda_1) \\ \vdots & \ddots & \vdots \\ \mathcal{K}_n(\lambda_1, \lambda_n) & \cdots & \mathcal{K}_n(\lambda_n, \lambda_n) \end{pmatrix} d\boldsymbol{\lambda}.$$

¹This follows from GUE having Haar distributed eigenvectors [2, Corollary 2.5.4] and the distribution of the eigenvectors of all UIEs being independent of V [10, pp. 25].

This form of distribution defines a *determinantal point process* (see, e.g., [1, Chapter 11]).

2.1. Spectral densities and equilibrium measures. The spectral density of a random matrix is the distribution of any single eigenvalue, without ordering. In other words, we integrate out all but one of the eigenvalues in (2.1). The so-called integrating out lemma [9, pg. 103] tells us that the spectral density is precisely

$$\frac{\mathcal{K}_n(x, x)}{n} dx.$$

As $n \rightarrow \infty$, the spectral density approaches the *limiting spectral density*. In the case of UIE with $Q(x) = nV(x)$, the scaling by n causes the limiting spectral density to have compact support. The limiting spectral density is then precisely the equilibrium measure of V [9, Section 6.4]:

DEFINITION 2.1. *The equilibrium measure μ is the unique minimizer of*

$$\iint \log \frac{1}{|x - y|} d\mu(x) d\mu(y) + \int V(x) d\mu(x)$$

among Borel probability measures on \mathbb{R} .

Some ensembles of interest do not have the scaled form $nV(x)$. From classical probability, however, we see that multiplication of the random matrix by a constant $\alpha(n)$ induces a scaling of Q : if M is an invariant ensemble with potential $Q(x)$, then $\alpha(n)^{-1}M$ is also an invariant ensemble with potential $Q(\alpha(n)x)$. We can choose $\alpha(n)$ so that $Q(\alpha(n)x)$ has the desired form.

Example The unscaled GUE has potential $Q(x) = x^2$. Multiplication by $\frac{1}{\sqrt{n}}$ gives the new potential $Q(\sqrt{n}x) = nx^2 = nV(x)$, which is of the desired form.

Example Consider a general polynomial potential $Q(x) = q_m x^m + \dots + q_0$. Multiplying the ensemble by $\frac{1}{n^{1/m}}$ gives the new potential

$$Q(n^{1/m}x) = n(q_m x^m + n^{-\frac{1}{m}} q_{m-1} x^{m-1} + \dots + n^{-1} q_0) = nV_n(x).$$

While V_n now depends on n , it tends to the n independent monomial $q_m x^m$, and the limiting spectral density is the equilibrium measure² of $q_m x^m$. If we can sample M from the invariant ensemble with potential $Q(x) = nV_n(x)$, then $n^{1/m}M$ gives a sample from the unscaled $Q(x)$.

Remark If we restrict our attention to weights supported on compact sets (say, $[-1, 1]$), then it is not necessary to induce a scaling. When Q is entire but not polynomial, there is no explicit scaling, however, a choice of $\alpha(n)$ that causes the equilibrium measure to have compact support for all n can be computed numerically [32].

For the potentials V we consider, the equilibrium measure is supported on a single interval $[a_V, b_V]$. The measure can be readily calculated using [22], returning an approximation to the representation in terms of Chebyshev U series:

$$d\mu(x) = \psi(x) dx \quad \text{for} \quad \psi\left(\frac{b_V + a_V}{2} + \frac{b_V - a_V}{2}x\right) = \frac{\sqrt{1-x^2}}{2\pi} \sum_{k=1}^{\infty} V_k U_{k-1}(x) dx,$$

²This follows from the leading order asymptotics of the equilibrium measure of V_n only depending on the leading order of the polynomial [11, (5.26)], and the expression of the finite spectral density as $\frac{\mathcal{K}_n(x, x)}{n} dx$.

where V_k are the Chebyshev coefficients of V' :

$$V' \left(\frac{b_V + a_V}{2} + \frac{b_V - a_V}{2} x \right) = \sum_{k=0}^{\infty} V_k T_k(x).$$

Associated with the equilibrium measure is a scaling constant for the edge universality law. In the non-degenerate case where the equilibrium measure has a precisely square root singularity at b_V , this is [25]

$$c_V = (b_V - a_V)^{-1/3} \left(2\pi \sum_{k=1}^{\infty} V_k \right)^{2/3}. \quad (2.3)$$

2.2. Gap statistics and universality. The gap statistics are local statistics, measuring the probability that there are no eigenvalues in an interval Ω . These statistics can be written in terms of a Fredholm determinant [1, Proposition 4.6.2]:

$$\det(I - \mathcal{K}_n|_{L^2(\Omega)}).$$

The notation $\mathcal{K}_n|_{L^2(\Omega)}$ refers to the integral operator on $L^2(\Omega)$ with kernel \mathcal{K}_n . A special case of gap statistics is the edge statistic, which is the probability that the largest eigenvalue is greater than s ; i.e., $\Omega = [s, \infty)$.

2.3. Universality. There are two types of universality laws we discuss: *bulk universality* and *edge universality*. The bulk universality law states that the gap statistic of a scaled neighbourhood of x in the support of the limiting spectral measure of a Hermitian random matrix tends (in an appropriate sense) to the sine-kernel law:

$$\det(I - \mathcal{K}_n|_{L^2[x + \frac{(-s, s)}{\psi(x)n}]}) \rightarrow \rho^{\sin}(s),$$

where ψ is the density of the equilibrium measure and

$$\rho^{\sin}(s) = \det(I - \mathcal{S}|_{L^2(-s, s)}) \quad \text{for} \quad \mathcal{S} = \frac{\sin(x-y)}{x-y}.$$

This universality law has been proved under broad conditions for both Wigner ensembles [31, 15] and UIE [9, 19].

Soft edge universality states that the distribution of a scaled largest eigenvalue tends to the Tracy–Widom law:

$$\det \left(I - \mathcal{K}_n|_{L^2\left[b_V + \frac{s}{c_V n^{2/3}}, \infty\right]} \right) \rightarrow \rho^{\text{Ai}}(s)$$

where b_V is again the right endpoint of the equilibrium measure, c_V is the edge scaling constant (2.3) and

$$\rho^{\text{Ai}}(s) = \det(I - \mathcal{A}|_{L^2(s, \infty)}) \quad \text{for} \quad \mathcal{A}(x, y) = \frac{\text{Ai}(x)\text{Ai}'(y) - \text{Ai}'(x)\text{Ai}(y)}{x-y}.$$

This law only applies if the behaviour at the right endpoint of the limiting spectral density has precisely square root decay. Since Wigner ensembles satisfy the semicircle law (having precisely square root decay), they broadly satisfy the Tracy–Widom law [30]. When the limiting spectral density has a singularity, as can be the case at the




Fig. 3.1: Monte Carlo simulation with $n = 5, 15$ and 25 of the quartic ensemble (left) and the higher-order decay ensemble (right) compared to the standard and higher-order Tracy–Widom distributions, respectively (dashed line).

left endpoint of a Wishart ensemble, hard edge universality applies, which is described in terms of the Bessel kernel [1, pp. 107].

However, invariant ensembles can have limiting spectral densities that exhibit faster decay, in which case the largest eigenvalue follows the higher-order Tracy–Widom law [7, 5]. Below, Section 3.2, is the first time the higher-order Tracy–Widom law have been observed through Monte Carlo simulations.

3. Experimental results. In this section, we present preliminary experimental results that are now possible using the algorithm for sampling invariant ensembles.

3.1. Quartic ensemble. The quartic ensemble has the potential $V(x) = x^4$. In the left-hand side of Figure 3.1, we see the convergence of the empirical CDF of the shifted and scaled largest eigenvalue $c_V(\lambda_{\max} - b_V)$ to the Tracy–Widom law. In Figure 3.2 we plot the difference between the empirical CDF and the Tracy–Widom law, and also compare the empirical complementary CDF³ of

$$n\psi(0) |\lambda_{\min}|$$

to the sine kernel law, where λ_{\min} is the eigenvalue with smallest absolute value and ψ is the density of the equilibrium measure.

Remark Monte Carlo simulation for the bulk statistics in Figure 3.2 and Figure 3.3 had not fully converged to the true distribution after a million samples, hence these should be taken as a rough estimate. For invariant ensembles, high accuracy and fast calculation of the true statistic is possible via the methodology of [25].

3.2. Higher-order decay ensemble. A special ensemble has the potential

$$V(x) = \frac{x^4}{20} - \frac{4}{15}x^3 + \frac{x^2}{5} + \frac{8}{5}x,$$

and the equilibrium measure [5]

$$\psi(x) dx = \frac{1}{10\pi}(x+2)^{1/2}(2-x)^{5/2} dx \quad \text{for } -2 < x < 2,$$

³The empirical complementary CDF is the probability that a random variable is larger than a value x , as measured through Monte Carlo simulation.




Fig. 3.2: Convergence of Monte Carlo simulation to universality for $w(x) = e^{-x^4}$, for $n = 25$ (dot-dashed), 50 (dotted), 75 (dashed) and 100 (plain), with a million samples. Left: difference between the empirical bulk statistic and the sine kernel law ($n = 100$ not pictured as the distribution was not sufficiently resolved). Right: convergence of the edge statistic to the Tracy–Widom law.

which is the dashed line in the left-hand side of Figure 3.5. The higher-order decay of the equilibrium causes the large eigenvalue statistics to be different: rather than following the Tracy–Widom distribution, they follow the higher-order Tracy–Widom distributions [5].

In Figure 3.1, we compare Monte Carlo simulations with the finite n distributions to the predicted higher-order Tracy–Widom distribution as calculated in [6]. We plot the edge statistic in the neighbourhood $[2 + \frac{s}{c_V^{\text{HO}} n^{2/7}}, \infty)$, for the scaling constant⁴ $c_V^{\text{HO}} = 5^{-2/7}$.

Remark While the Monte Carlo simulations roughly follow the predicted universality law, they differ by a substantial amount due to n not being sufficiently high. We have verified the accuracy of Monte Carlo simulation by comparing to a high accuracy approximation of the finite n statistic using a numerical Riemann–Hilbert approach [25] (not pictured). Using the Riemann–Hilbert approach, we see that high accuracy agreement of the finite n statistic with the universality law does not occur until $n \approx 100,000$, beyond the current reach of simulation.

3.3. Non-varying cosh ensemble.

Consider the weight

$$w(x) = e^{-\cosh x}$$

which does not scale with n , hence the existing theory on universality breaks down. However, we can still perform sampling as we can calculate the recurrence relationship for the associated orthogonal polynomials, using the Riemann–Hilbert approach [32].

With sampling in hand, we verify that the ensemble does indeed follow both bulk and edge universality in Figure 3.3. Here, to get convergence to the universality laws, we choose the scaling constants associated with the equilibrium measure of $\frac{1}{n} \cosh x$. It is important to note that in this case c_V varies with n . The convergence to the Tracy–Widom law of the edge statistic is surprising: the equilibrium measure should not have a nice limit at the right endpoint.

⁴The constant is stated in [7] as $6^{1/7}$. However, the normalization used in [5, 6] differ by a factor of 30, giving $c_V^{\text{HO}} = (6/30)^{1/7} = 5^{-2/7}$.




Fig. 3.3: Convergence of Monte Carlo simulation to universality for $w(x) = e^{-\cosh x}$, for $n = 25$ (dot-dashed), 50 (dotted), 75 (dashed) and 100 (plain). Left: convergence of the bulk statistic to the sine kernel law. Right: convergence of the edge statistic to the Tracy–Widom law.

3.4. Addition of invariant ensembles. While the statistics of invariant ensembles themselves are already known in detail, both asymptotically [9] and numerically [25], the theory and numerics break down as soon as further manipulations (*e.g.* matrix addition) of invariant ensembles are performed. The ability to sample invariant ensembles directly means that we can now perform Monte Carlo simulations to help understand such manipulations. In this section, we investigate the global distribution of the eigenvalues for the addition of invariant ensembles.

Our first experiment consists of calculating the spectral density of quartic ensemble added to a GUE. The $n = 1$ is equivalent to classical convolution, while as n increases it approaches the free probability convolution, which has been calculated symbolically in [27] and numerically [23]. In Figure 3.4, we see that the simulated statistics agree in the $n = 1$ case, as expected. Surprisingly, close agreement is seen for $n = 10$ between simulation and the free probability convolution. This relationship is only known to hold in the large n limit.

In the second example, we sample H_1, \dots, H_k independent higher-order decay ensembles, as defined in Section 3.2, and investigate the statistics of $H_1 + \dots + H_k$. Free probability tells us the limiting spectral density, which we show in the left-hand side of Figure 3.5, calculated numerically via [23] and verified via Monte Carlo simulation (not pictured). On the right, we plot the empirical edge statistic: the CDF of $.3n^{2/3}(\lambda_{\max} - b_k)$, where b_k is the right endpoint of the limiting spectral density. (The choice of scaling by $.3$ is arbitrary.) In the limiting spectral density we see the emergence of a precisely square root singularity at $k = 4$. This coincides with the edge statistic appearing to follow the standard Tracy–Widom law, rather than the higher-order analogues.

3.5. The Kolmogorov–Smirnov statistic. In following the error analysis of [18], we consider Kolmogorov–Smirnov (KS) statistics. The empirical distribution function for m samples $\{\lambda^m\}$, each consisting of n eigenvalues, is

$$F_{n,m}(x) = \frac{1}{nm} \sum_{j=1}^m \sum_{k=1}^n \mathbf{1}_{\lambda_k^m < x}.$$




Fig. 3.4: Spectral density of the addition of an $n \times n$ Quartic ensemble and GUE for $n = 1, 2, 5$ and 10 . The dotted line in the first figure is a numerical classical additive convolution, and the dotted line in the last figure is the numerically calculated free additive convolution. Close agreement is seen for $n = 10$ between simulation and the free probability convolution ($n = \infty$).




Fig. 3.5: Statistics of the addition of k higher-order decay ensembles, for $k = 1$ (dashed), 2 (dotted), 3 (dot-dashed), 4 (long-dashed) and 5 (plain). Left: the limiting spectral densities. Right: Monte Carlo simulation of the edge statistics.

This should be compared with both the spectral distribution function

$$F_n(x) = \frac{1}{n} \int_{-\infty}^x \mathcal{K}_n(y, y) dy,$$




Fig. 3.6: The estimated KS statistics for the Quartic ensemble. +: $E_{n,1000}^{\infty}$ plotted versus n , O: $E_{n,1000}$ plotted versus n , x: $E_{n,5000}^{\infty}$ plotted versus n , □: $E_{n,5000}$ plotted versus n . We see an essentially monotonic decay in the statistic as the matrix size, n , increases. As expected, lower errors are seen when comparing with \mathcal{K}_n .

and the limiting distribution function $F(x) = \mu((-\infty, x])$ where μ is the equilibrium measure. The two relevant KS statistics are

$$E_{n,m} = \sup_{x \in \mathbb{R}} |F_{n,m}(x) - F_n(x)|, \quad E_{n,m}^{\infty} = \sup_{x \in \mathbb{R}} |F_{n,m}(x) - F(x)|.$$

We show these statistics for the Quartic ensemble in Figure 3.6. It should be noted that the data point for $E_{100,1000}^{\infty}$ (see the + at $n = 100$ in Figure 3.6) appears to lie below the comparable statistic in [18, Figure 5(b)]. Furthermore, our errors do not saturate in the same way. This can be explained by the fact that while truncation errors are a concern for the numerical solution of stochastic differential equations, they are much less significant (at least to $n = 1000$, $m = 5000$) for the method presented here.

4. Chebyshev expansion of weighted orthogonal polynomials. The algorithm below will depend on calculating $\mathcal{K}_n(x, y)$, which requires calculating the orthogonal polynomials. Recall that if $p_k(x)$ are orthonormal with respect to $w(x) dx$, then they satisfy a symmetric three-term recurrence relationship

$$\beta_{k-1} p_{k-1}(x) + \alpha_k p_k(x) + \beta_k p_{k+1}(x) = x p_k(x)$$

with recurrence coefficients α_k and β_k . We use these recurrence coefficients, along with the constant $\int w(x) dx$, to calculate orthogonal polynomials pointwise.

Computation of the recurrence coefficients α_k and β_k is required. These are known in closed form for classical orthogonal polynomials: Hermite, Jacobi and Laguerre polynomials. Otherwise, the standard approach is the Stieltjes procedure [16], which is essentially the modified Gram–Schmidt method applied to a discretized inner product using a quadrature rule. Alternatively, one can calculate these via the

numerical solution of Riemann–Hilbert problems [32]. While the latter approach is more complicated, it is an $\mathcal{O}(n)$ algorithm, which is significantly better complexity than Stieljes procedure for large n , see the discussion in [32].

Pointwise evaluation of the orthogonal polynomials is not sufficient on its own: in the algorithm, we need to perform the following further operations on $\phi(x) = (\phi_0(x), \dots, \phi_{n-1}(x))^\top$ where $\phi_k(x) = \sqrt{w(x)}p_k(x)$:

1. Evaluation of $\phi(r)$.
2. Pointwise evaluation of the indefinite integral of $\phi(x)A\phi(x)$ where A is a constant matrix.

A convenient method of accomplishing these tasks is to initially expand ϕ_k into a Chebyshev expansion on an interval $[a, b]$ chosen so that ϕ_k is negligible off the interval:

$$\phi_k\left(\frac{a+b}{2} + \frac{b-a}{2}x\right) \approx \sum_{j=0}^{m-1} c_{kj}T_j(x).$$

We determine this expansion by adaptively doubling \tilde{m} until the last eight coefficients of

$$\phi_0k\left(\frac{a+b}{2} + \frac{b-a}{2}x\right) \approx \sum_{j=0}^{\tilde{m}-1} c_{0j}T_k(x)$$

are negligible, *a lá* the CHEBFUN package [3]. At each test value of \tilde{m} , c_{kj} are determined by applying the DCT to $\phi_0\left(\frac{a+b}{2} + \frac{b-a}{2}\mathbf{x}_m\right)$ for m Chebyshev points of the first kind:

$$\mathbf{x}_m = \left(-1, \cos \pi \left[1 - \frac{1}{m-1}\right], \dots, \cos \pi \frac{1}{m-1}, 1\right)^\top.$$

The negligibility of the last coefficients of ϕ_0 means that it is effectively a polynomial of degree less than \tilde{m} , hence $\phi_1, \dots, \phi_{n-1}$ are effectively polynomials of degree less than $\tilde{m} + n$, and $\phi_k(x)\phi_j(x)$ are effectively polynomials of degree less than $m = 2(\tilde{m} + n)$. From the values of $\phi_0(\mathbf{x}_m)$, we can successfully calculate the values of $\phi_k(\mathbf{x}_m)$ using the recurrence relationship.

Using the values of ϕ_k at the Chebyshev points \mathbf{x}_m , $\phi(r)$ can be approximated in $\mathcal{O}(m)$ operations using the barycentric formula [4]. We can efficiently multiply Chebyshev series at each Chebyshev point, and then the Chebyshev expansion of $\phi(x)A\phi(x)$ can be calculated using the DCT. Indefinite integration of the resulting Chebyshev expansion is possible in $\mathcal{O}(m)$ operations [21, pp. 32–33], and then the resulting Chebyshev expansion can be evaluated pointwise in $\mathcal{O}(m)$ operations using Clenshaw’s method [8].

5. Sampling determinantal point processes. Here we present an algorithm for sampling the determinantal point process associated with the eigenvalues of UIEs. Determinantal point processes are point processes whose distribution has a determinantal representation

$$\det [\mathcal{K}_n(x_i, x_j)]_{i,j=1}^n$$

in terms of a kernel $\mathcal{K}_n(x, y)$, see [1, Chapter 11] for a general exposition. In our case, the kernel \mathcal{K}_n is given in terms of orthogonal polynomials, recall definition (2.2).

In [17, 28] an approach for sampling determinantal point processes was introduced. In our concrete setting, the algorithm takes the following form.

DEFINITION 5.1. For $A \in \mathbb{R}^{m \times n}$ with row rank r and $m < n$, the function $\text{null}(A)$ is defined to return an orthogonal matrix $Q \in \mathbb{R}^{n \times n-r}$ whose columns span the kernel of A , so that $AQ = 0$. It will be clear that our results do not depend on the choice of matrix Q .

Algorithm 1 Sample determinantal processes, adapted from [17, 28]

Input: Chebyshev interpolations of $\boldsymbol{\phi} = (\phi_0, \dots, \phi_{n-1})^\top$
Output: n UIE eigenvalues $\mathbf{r} = (r_1, \dots, r_n)$
Initialize $\mathbf{q}_n(x) = \boldsymbol{\phi}(x)$
for $k = n, \dots, 2$ **do**
 Obtain r_k by sampling the PDF $\frac{\mathbf{q}_k(x)^\top \mathbf{q}_k(x)}{k}$
 Let $\mathbf{f}_k = \mathbf{q}_k(r_k) \in \mathbb{R}^k$
 Let $Q_k = \text{null}(\mathbf{f}_k^\top)$ (so that $Q_k^\top \mathbf{f}_k = 0$, and $Q_k \in \mathbb{R}^{k \times k-1}$)
 Let $\mathbf{q}_{k-1}(x) = Q_k^\top \mathbf{q}_k(x)$
end for
Obtain r_1 by sampling the PDF $\mathbf{q}_1^\top \mathbf{q}_1$

We now use this algorithm, combined with calculation of orthogonal polynomials, to sample unitary invariant ensembles.

Algorithm 2 Sample unitary invariant ensembles

Input: Chebyshev interpolations of $\boldsymbol{\phi} = (\phi_0, \dots, \phi_{n-1})^\top$
Output: $n \times n$ UIE matrix
Obtain n eigenvalues \mathbf{r} by calling Algorithm 1
Sample unitary matrix V from the Haar distribution. E.g., orthogonalize a random $n \times n$ complex matrix whose entries are iid complex Gaussians [29].
Return $V \text{diag}\{\mathbf{r}\} V^*$

We must ensure that we are sampling the correct distribution with this algorithm. Algorithm 1 can be interpreted as a construction of mathematical random variables, however, on a computer these are in fact a deterministic sequence of pseudo-random numbers. In particular, we sample $\frac{\mathbf{q}_k^\top(x) \mathbf{q}_k(x)}{k}$ using the numerical inverse transform sampling approach of [24]: calculate the CDF

$$F(x) = \int_a^x \frac{\mathbf{q}_k(x)^\top \mathbf{q}_k(x)}{k} dx$$

using indefinite integration of the Chebyshev series representation, sample pseudo-random variable Y uniformly in $[0, 1]$, and find X satisfying $F(X) = Y$ using the bisection method, which only requires evaluation of F .

We use the following definition in order to be precise on what is meant by sampling a distribution. Note that this definition only encodes convergence to the correct distribution; pseudo-randomness is a subtle issue beyond the scope of this paper.

DEFINITION 5.2. A method is said to sample a Borel probability measure μ on

\mathbb{R}^n if it produces a (non-random) sequence of points $\{\mathbf{x}_1, \mathbf{x}_2, \dots\} \subset \mathbb{R}^n$ such that

$$\lim_{n \rightarrow \infty} \frac{\#\{j : \mathbf{x}_j \in B\}}{n} = \lim_{n \rightarrow \infty} \int_B \left(\frac{1}{n} \sum_{j=1}^n \delta_{\mathbf{x}_j}(\mathbf{x}) \right) d\mathbf{x} = \mu(B), \quad (5.1)$$

for all rectangles $B = [a_1, b_1] \times \dots \times [a_n, b_n] \subset \mathbb{R}^n$.

Remark From a stochastic process point of view, we want the sequence $\{\mathbf{x}_1, \mathbf{x}_2, \dots\}$ to behave like a generic sample path of $\{\mathbf{X}_1(\omega), \mathbf{X}_2(\omega), \dots\}$ where \mathbf{X}_i are iid vector-valued random variables with joint distribution μ . From the Glivenko–Cantelli theorem [13, pg. 76], (5.1) is a necessary consequence of this. Furthermore, if we can sample with a given method, then we can, in principle, completely determine the distribution μ , even if we do not have an expression for μ .

If Y samples the uniform distribution on $[0, 1]$, then X will sample a distribution that will be on the order of machine precision ($\approx 2.22 \times 10^{-16}$) when compared with the true distribution. In what follows, we ignore this small error and treat our sampling procedure as exact.

We aim to show that r_1, \dots, r_n sample the probability measure

$$P_n(x_1, \dots, x_n) d\mathbf{x} = \frac{1}{Z_n} e^{-\sum_{i=1}^n Q(x_i)} \prod_{i < j} (x_i - x_j)^2 d\mathbf{x},$$

and we use the representation

$$P_n(x_1, \dots, x_n) = \frac{1}{n!} \det(\mathcal{K}_n(x_i, x_j))_{1 \leq i, j \leq n}$$

Instead of sampling an n -dimensional distribution, as mentioned above, we reduce the complexity by exploiting the determinantal representation. Indeed, the sampling is reduced to sampling a sequence of n one-dimensional distributions.

We now show that Algorithm 1 samples $P_n(x_1, \dots, x_n) d\mathbf{x}$. As a first step, we show that the above algorithm samples

$$p(x_1, \dots, x_n) d\mathbf{x} = \prod_{j=1}^n \frac{\mathbf{q}_j(x_j)^\top \mathbf{q}_j(x_j)}{j} d\mathbf{x}.$$

This follows from the following lemma.

LEMMA 5.3. *Assume a probability measure μ on $U \subset \mathbb{R}^n$ has the form*

$$\begin{aligned} d\mu(x_1, \dots, x_n) &= F(x_1, \dots, x_n) d\mathbf{x}, \\ &= f_1(x_1) f_2(x_1, x_2) \cdots f_n(x_1, \dots, x_n) d\mathbf{x}, \quad F > 0, \end{aligned}$$

where $\int_{\mathbb{R}^{m-1}} f_m(x_1, \dots, x_m) dx_1 \cdots dx_{m-1}$ defines a probability measure on \mathbb{R} (for the x_m variable). If the sequence $\mathbf{X}_i = (X_1^i, \dots, X_n^i)$, $i = 1, 2, \dots$ samples the standard uniform distribution on $[0, 1]^n$ then the sequence, $\mathbf{x}_i = (x_1^i, \dots, x_n^i)$, $i = 1, 2, \dots$

uniquely defined by the relation

$$\begin{aligned}
X_1^i &= \int_{-\infty}^{x_1^i} f_1(x_1') dx_1', \\
X_2^i &= \int_{-\infty}^{x_2^i} f_2(x_1^i, x_2') dx_2', \\
&\vdots \\
X_n^i &= \int_{-\infty}^{x_n^i} f_n(x_1^i, \dots, x_{n-1}^i, x_n') dx_n',
\end{aligned} \tag{5.2}$$

samples μ .

Proof. It is clear that the determinant of the Jacobian of the transformation is just $F(x_1^i, \dots, x_n^i)$ and therefore the mapping $\mathbf{x}_i \mapsto \mathbf{X}_i$ is invertible. We know that for any Borel set $S \subset [0, 1]^n$

$$\lim_{n \rightarrow \infty} \frac{\#\{j : j = 1, \dots, n \text{ and } \mathbf{X}^j \in S\}}{n} = \int_S dX_1 \cdots dX_n.$$

A simple calculation using the determinant of the Jacobian shows that

$$\lim_{n \rightarrow \infty} \frac{\#\{j : j = 1, \dots, n \text{ and } \mathbf{x}^j \in T\}}{n} = \int_T F(x_1, \dots, x_n) dx_1 \cdots dx_n,$$

where T is the inverse image of S under the transformation (5.2). Therefore, for any Borel set T let S be the image of T under (5.2) and this shows that the method samples μ .

□

Next, we show that $p(x_1, \dots, x_n) = P_n(x_1, \dots, x_n)$. Define

$$\psi_j^{(k)}(x) = \mathbf{q}_k(x_j)^\top \mathbf{q}_k(x),$$

so that $\psi_k^{(k)}(x) = \mathbf{f}_k^\top \mathbf{q}_k(x)$. By orthogonality we have

$$\langle \psi_j^{(k)}, \psi_\ell^{(k)} \rangle = \mathbf{q}_k(x_j)^\top \mathbf{q}_k(x_\ell).$$

therefore

$$\begin{aligned}
P_n(x_1, \dots, x_n) &= \frac{1}{n!} A \quad \text{for} \\
A &= \det \begin{bmatrix} \langle \psi_1^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_1^{(1)}, \psi_2^{(1)} \rangle & \langle \psi_1^{(1)}, \psi_3^{(1)} \rangle & \cdots & \langle \psi_1^{(1)}, \psi_n^{(1)} \rangle \\ \langle \psi_2^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_2^{(1)}, \psi_2^{(1)} \rangle & \cdots & \cdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \langle \psi_n^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_n^{(1)}, \psi_2^{(1)} \rangle & \cdots & \cdots & \langle \psi_n^{(1)}, \psi_n^{(1)} \rangle \end{bmatrix}.
\end{aligned}$$

We wish to rewrite $\psi_j^{(1)}$ in terms of $\psi_j^{(k)}(x) = \mathbf{q}_k(x_j)^\top \mathbf{q}_k(x)$. First observe that prepending the vector $\mathbf{f}_k / \|\mathbf{f}_k\|$ to Q_k we obtain a $k \times k$ orthogonal matrix

$$\tilde{Q}_k = \begin{bmatrix} \mathbf{f}_k \\ \|\mathbf{f}_k\|, Q_k \end{bmatrix}.$$

Thus we have (observing that $\mathbf{f}_1 = \phi(x_1)$)

$$\begin{aligned}\psi_j^{(1)}(x) &= \phi(x_j)^\top \tilde{Q}_1 \tilde{Q}_1^\top \mathbf{q}_1(x) = \frac{\phi(x_j)^\top \mathbf{f}_1 \mathbf{f}_1^\top \mathbf{q}_1(x)}{\|\mathbf{f}_1\|^2} + \phi(x_j)^\top Q_1 Q_1^\top \mathbf{q}_1(x) \\ &= c_j \psi_1^{(1)}(x) + \mathbf{q}_2(x_j)^\top \mathbf{q}_2(x) = c_j \psi_1^{(1)}(x) + \psi_j^{(2)}(x)\end{aligned}$$

By column reductions we get

$$A = \det \begin{bmatrix} \langle \psi_1^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_1^{(1)}, \psi_2^{(2)} \rangle & \langle \psi_1^{(1)}, \psi_3^{(2)} \rangle & \cdots & \langle \psi_1^{(1)}, \psi_n^{(2)} \rangle \\ \langle \psi_2^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_2^{(1)}, \psi_2^{(2)} \rangle & \cdots & \cdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \langle \psi_n^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_n^{(1)}, \psi_2^{(2)} \rangle & \cdots & \cdots & \langle \psi_n^{(1)}, \psi_n^{(2)} \rangle \end{bmatrix}$$

Similarly, we have

$$\psi_j^{(k)}(x) = c_j \psi_k^{(k)}(x) + \psi_j^{(k+1)}(x),$$

and thus we can further reduce to

$$A = \det \begin{bmatrix} \langle \psi_1^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_1^{(1)}, \psi_2^{(2)} \rangle & \langle \psi_1^{(1)}, \psi_3^{(3)} \rangle & \cdots & \langle \psi_1^{(1)}, \psi_n^{(n)} \rangle \\ \langle \psi_2^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_2^{(1)}, \psi_2^{(2)} \rangle & \cdots & \cdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \langle \psi_n^{(1)}, \psi_1^{(1)} \rangle & \langle \psi_n^{(1)}, \psi_2^{(2)} \rangle & \cdots & \cdots & \langle \psi_n^{(1)}, \psi_n^{(n)} \rangle \end{bmatrix}$$

Noting that for any vectors \mathbf{u} and \mathbf{v}

$$\langle \mathbf{u}^\top \phi, \mathbf{v}^\top \phi \rangle = \mathbf{u}^\top \mathbf{v},$$

we have for $j < k$

$$\begin{aligned}\langle \psi_j^{(1)}, \psi_k^{(k)} \rangle &= \langle \phi(x_j)^\top \phi, \mathbf{f}_k^\top \mathbf{q}_k \rangle = \langle \phi(x_j)^\top \phi, \mathbf{f}_k^\top Q_{k-1}^\top \cdots Q_1^\top \phi \rangle \\ &= \phi(x_j)^\top Q_1 \cdots Q_{k-1} \mathbf{f}_k = \mathbf{q}_j(x_j)^\top Q_{k-1}^\top \cdots Q_1^\top Q_1 \cdots Q_{k-1} \mathbf{f}_k \\ &= \mathbf{f}_j^\top Q_j \cdots Q_{k-1} \mathbf{f}_k = 0.\end{aligned}$$

Using this, we subsequently obtain

$$\langle \psi_k^{(1)}, \psi_k^{(k)} \rangle = \langle \psi_k^{(2)}, \psi_k^{(k)} \rangle = \cdots = \langle \psi_k^{(k)}, \psi_k^{(k)} \rangle$$

Thus we have an upper triangular representation:

$$\begin{aligned}A &= \det \begin{bmatrix} \langle \psi_1^{(1)}, \psi_1^{(1)} \rangle & & & & \\ \vdots & \langle \psi_2^{(2)}, \psi_2^{(2)} \rangle & & & \\ \vdots & \vdots & \ddots & & \\ \vdots & \cdots & \cdots & \langle \psi_n^{(n)}, \psi_n^{(n)} \rangle & \end{bmatrix} \\ &= \prod_{k=1}^n \mathbf{q}_k(x_k)^\top \mathbf{q}_k(x_k).\end{aligned}$$

In other words,

$$P_n(x_1, \dots, x_n) = \frac{A}{n!} = p(x_1, \dots, x_n).$$

Acknowledgements. We would like to thank Paul Cheung, Percy Deift and Govind Menon for discussions concerning this method. We acknowledge the generous support of the National Science Foundation through grant NSF-DMS-130318 (TT) and the Australian Research Council through the Discovery Early Career Research Award (SO). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding sources.

REFERENCES

- [1] G. Akemann, J. Baik, and P. Di Francesco. *The Oxford Handbook of Random Matrix Theory*. Oxford University Press, 2011.
- [2] G. Anderson, A. Guionnet, and O. Zeitouni. *Introduction to Random Matrices (Studies in advanced mathematics, N 118)*. Cambridge University Press, Cambridge, 2009.
- [3] Z. Battles and L. N. Trefethen. An extension of MATLAB to continuous functions and operators. *SIAM J. Sci. Comput.*, 25:1743–1770, 2004.
- [4] J.-P. Berrut and L. N. Trefethen. Barycentric Lagrange interpolation. *SIAM Review*, 46(3):501–517, 2004.
- [5] T. Claeys, I. Krasovsky, and A. Its. Higher-order analogues of the Tracy–Widom distribution and the Painlevé II hierarchy. *Comm. Pure Appl. Math.*, 63:362–412, 2010.
- [6] T. Claeys and S. Olver. Numerical study of higher order analogues of the Tracy–Widom distribution. *Cont. Maths*, 578:83–99, 2011.
- [7] T. Claeys and M. Vanlessen. Universality of a double scaling limit near singular edge points in random matrix models. *Comm. Math. Phys.*, 273(2):499–532, 2007.
- [8] C. Clenshaw. A note on the summation of chebyshev series. *Mathematics of Computation*, 9(51):118–120, 1955.
- [9] P. Deift. *Orthogonal Polynomials and Random Matrices: a Riemann–Hilbert Approach*. AMS, 2000.
- [10] P. Deift and D. Gioev. *Random Matrix Theory: Invariant Ensembles and Universality*, volume 18. American Mathematical Society, 2009.
- [11] P. Deift, T. Kriecherbauer, K. T.-R. McLaughlin, S. Venakides, and X. Zhou. Strong asymptotics of orthogonal polynomials with respect to exponential weights. *Comm. Pure Appl. Math.*, 52(12):1491–1552, 1999.
- [12] I. Dumitriu and A. Edelman. Matrix models for beta ensembles. *Journal of Mathematical Physics*, 43(11):5830–5847, 2002.
- [13] R. Durrett. *Probability: theory and examples*. Cambridge Univ. Press, 2010.
- [14] A. Edelman and B. D. Sutton. The beta-Jacobi matrix model, the CS decomposition, and generalized singular value problems. *Found. Comput. Maths*, 8(2):259–285, 2008.
- [15] L. Erdős, S. Péché, J. A. Ramírez, B. Schlein, and H.-T. Yau. Bulk universality for Wigner matrices. *Comm. Pure Appl. Maths*, 63(7):895–925, 2010.
- [16] W. Gautschi. *Orthogonal Polynomials: Applications and Computation*. Oxford University Press, 2004.
- [17] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág. Determinantal processes and independence. *Probability Surveys*, 3:206–229, 2006.
- [18] X. H. Li and G. Menon. Numerical solution of Dyson Brownian motion and a sampling scheme for invariant matrix ensembles. *J. Stat. Phys.*, 153, 2013.
- [19] D. S. Lubinsky. Bulk universality holds in measure for compactly supported measures. *Journal d’Analyse Mathématique*, 116(1):219–253, 2012.
- [20] V. A. Marčenko and L. A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Sbornik: Mathematics*, 1(4):457–483, 1967.
- [21] J. C. Mason and D. C. Handscomb. *Chebyshev Polynomials*. CRC Press, 2002.
- [22] S. Olver. Computation of equilibrium measures. *J. Approx. Theory*, 163:1185–1207, 2011.
- [23] S. Olver and R. R. Nadakuditi. Numerical computation of convolutions in free probability theory. *arXiv:1203.1958*, 2012.
- [24] S. Olver and A. Townsend. Fast inverse transform sampling in one and two dimensions. *arXiv:1307.1223*, 2013.
- [25] S. Olver and T. Trogdon. Numerical solution of Riemann–Hilbert problems: random matrix theory and orthogonal polynomials. *Const. Approx.*, 39:101–149, 2013.
- [26] L. A. Pastur. On the spectrum of random matrices. *Theoretical and Mathematical Physics*, 10(1):67–74, 1972.

- [27] N. R. Rao and A. Edelman. The polynomial method for random matrices. *Foundations of Computational Mathematics*, 8(6):649–702, 2008.
- [28] A. Scardicchio, C. E. Zachary, and S. Torquato. Statistical properties of determinantal point processes in high-dimensional euclidean spaces. *Physical Review E*, 79(4):041108, 2009.
- [29] G. W. Stewart. The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM Journal on Numerical Analysis*, 17(3):pp. 403–409, 1980.
- [30] T. Tao and V. Vu. Random matrices: Universality of local eigenvalue statistics up to the edge. *Communications in Mathematical Physics*, 298(2):549–572, 2010.
- [31] T. Tao and V. Vu. Random matrices: universality of local eigenvalue statistics. *Acta Mathematica*, 206(1):127–204, 2011.
- [32] T. Trogdon and S. Olver. A Riemann–Hilbert approach to Jacobi operators and Gaussian quadrature. *arXiv:1311.5838*, 2013.
- [33] E. P. Wigner. On the distribution of the roots of certain symmetric matrices. *Annals Maths*, 67(2):325–327, 1958.
- [34] N. Witte, P. Forrester, and C. M. Cosgrove. Gap probabilities for edge intervals in finite gaussian and jacobi unitary matrix ensembles. *Nonlinearity*, 13(5):1439, 2000.