

Universality in numerical computation with random data.

Case studies and analytical results

Percy Deift

Courant Institute of Mathematical Sciences

joint with Christian Pfrang, Govind Menon, Sheehan Olver and, mostly,
Tom Trogdon

deift@cims.nyu.edu

Refs: see arXiv 2012–2017

A few years ago, Christian Pfrang, Govind Menon and PD [PDM12], initiated a statistical study of the performance of various standard algorithms to compute the eigenvalues of random real symmetric matrices H . In each case, an initial matrix H_0 is diagonalized either by

a sequence of isospectral iterates H_m

$$H_0 \rightarrow H_1 \rightarrow H_2 \rightarrow \cdots \rightarrow H_m \rightarrow \cdots$$

or by

an isospectral flow

$$t \mapsto H(t) \quad \text{with} \quad H(t=0) = H_0.$$

In the discrete case, as $k \rightarrow \infty$,

H_k converges to a diagonal matrix.

Given $\epsilon > 0$, it follows that for some (first) time k , the off-diagonal entries of H_k are $O(\epsilon)$, and hence the diagonal entries of H_k give the eigenvalues of H_0 to $O(\epsilon)$. The situation is similar for continuous algorithms $t \mapsto H(t)$ as $t \rightarrow \infty$.

Christian W. Pfrang, Percy Deift, and Govind Menon. How long does it take to compute the eigenvalues of a random symmetric matrix? *arXiv Prepr. arXiv1203.4635*, mar 2012

The QR algorithm is a prototypical example of such a discrete algorithm:

1. Write $H_0 = Q_0 R_0$, Q_0 orthogonal, R_0 upper triangular, $(R_0)_{ii} > 0$
2. Set $H_1 = R_0 Q_0 = Q_0^T H_0 Q_0$
3. Write $H_1 = Q_1 R_1$
4. Set $H_2 = R_1 Q_1$
- \vdots

And the Toda algorithm is an example of such a continuous algorithm:

Solve

$$\frac{dH(t)}{dt} = [H(t), B(H(t))] = HB - BH, \quad H(0) = H_0$$

where $B(H) = H_- - H_-^T$.

Both of these are completely integrable Hamiltonian flows, a fact to which we will return later.

The main finding in [PDM12] was that, surprisingly,

the fluctuations in the stopping times were universal (1)

independent of the ensemble considered for the matrices H . More precisely,

- ▶ for $N \times N$ real symmetric matrices H ,
- ▶ chosen from an ensemble \mathcal{E} , and
- ▶ for a given algorithm \mathcal{A} , and
- ▶ a desired accuracy ϵ ,

let

$$T(H) = T_{\epsilon, N, \mathcal{A}, \mathcal{E}}(H) \quad (2)$$

be the stopping time (see later) for the algorithm \mathcal{A} applied to the $N \times N$ matrix H chosen from the ensemble \mathcal{E} , to achieve an accuracy ϵ .

Let $\tilde{T}(H) = \tilde{T}_{\epsilon,N,\mathcal{A},\mathcal{E}}(H)$ be the normalized stopping time

$$\tilde{T}_{\epsilon,N,\mathcal{A},\mathcal{E}}(H) = \frac{T_{\epsilon,N,\mathcal{A},\mathcal{E}}(H) - \langle T_{\epsilon,N,\mathcal{A},\mathcal{E}} \rangle}{\sigma_{\epsilon,N,\mathcal{A},\mathcal{E}}} \quad (3)$$

where $\langle T_{\epsilon,N,\mathcal{A},\mathcal{E}} \rangle$ is the sample average and $\sigma_{\epsilon,N,\mathcal{A},\mathcal{E}}^2 = \langle (T_{\epsilon,N,\mathcal{A},\mathcal{E}} - \langle T_{\epsilon,N,\mathcal{A},\mathcal{E}} \rangle)^2 \rangle$ is the sample variance, taken over a large number (5,000-15,000) of samples of matrices H chosen from \mathcal{E} . Then for a given algorithm \mathcal{A} , and ϵ and N in a suitable scaling region,

the histogram for $\tilde{T}_{\epsilon,N,\mathcal{A},\mathcal{E}}(H)$ is independent of \mathcal{E} . (4)

In general, the histogram will depend on \mathcal{A} , but for a given \mathcal{A} and ϵ and N in the scaling region, the histogram is independent of \mathcal{E} .

— such two-component universality is analogous to the classical central limit theorem for iid $\{X_i\}$,

$$\frac{X_1 + \cdots + X_N - \mu N}{\sigma_N} \xrightarrow{d} \text{standard Gaussian.}$$

Here are two examples, the first is for the QR algorithm and the second is for the Toda algorithm.

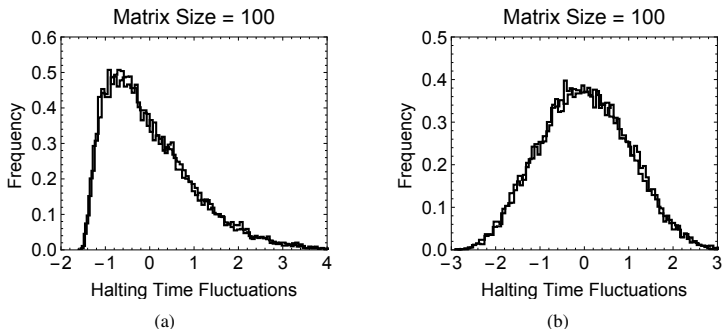


Figure: Universality for $\tilde{T}_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when (a) \mathcal{A} is the QR eigenvalue algorithm and when (b) \mathcal{A} is the Toda algorithm. Panel (a) displays the overlay of two histograms for $\tilde{T}_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ in the case of QR, one for each of the two ensembles $\mathcal{E} = \text{BE}$, consisting of iid mean-zero Bernoulli random variables and $\mathcal{E} = \text{GOE}$, consisting of iid mean-zero normal random variables. Here $\epsilon = 10^{-10}$ and $N = 100$. Panel (b) displays the overlay of two histograms for $\tilde{T}_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ in the case of the Toda algorithm, and again $\mathcal{E} = \text{BE}$ or GOE . And here $\epsilon = 10^{-8}$ and $N = 100$.

The stopping times, or halting times, for a given algorithm can be chosen in various ways, depending on which aspects of the given algorithm one wants to investigate. In the above figures, the stopping times take into account the phenomenon known as deflation, i.e., $T_{\epsilon, N, \mathcal{A}, \mathcal{E}}(H)$ is the first time k (or t in the continuous case) such that H_k (or $H(t)$) has block form

$$H_k = \begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix},$$

with H_{11} $j \times j$, H_{22} $(N - j) \times (N - j)$ such that

$$\|H_{12}\| = \|H_{21}\| \leq \epsilon$$

for some $1 \leq j \leq N - j$. Then the eigenvalues $\{\lambda_j\}$ of H differ from the eigenvalues $\{\hat{\lambda}_j\}$ of the deflated matrix

$$\hat{H} = \begin{pmatrix} H_{11} & 0 \\ 0 & H_{22} \end{pmatrix},$$

by $O(\epsilon)$. The algorithm is then applied to the (smaller) matrices H_{11} and H_{22} , etc.

Subsequent to [PDM12], Govind Menon, Sheehan Olver, Tom Trogdon and PD [DMOT14], raised the question of whether the universality results in the study [PDM12] were limited to eigenvalue algorithms, or whether they were present more generally in numerical computation. And indeed the authors in [DMOT14] found similar universality results for a wide variety of numerical algorithms, including

- (a) more general eigenvalue algorithms such as the Jacobi eigenvalue algorithm, and also algorithms for Hermitian ensembles,
- (b) the conjugate gradient and GMRES algorithms to solve linear $N \times N$ systems $Hx = b$,
- (c) an iterative algorithm to solve the Dirichlet problem $\Delta u = 0$ on a random star-shaped region $\Omega \subset \mathbb{R}^2$ with random boundary data f on $\partial\Omega$,
- (d) a genetic algorithm to compute the equilibrium measure for orthogonal polynomials on the line, and
- (e) decision making algorithms.

Some comments on (a)

(a): In all the calculations in [PDM12], M was real and symmetric with independent entries. Here we considered $N \times N$ Hermitian $M = M^*$ from various unitary invariant ensembles with distributions proportional to

$$e^{-N\text{tr}V(M)} dM$$

where $V(x) : \mathbb{R} \rightarrow \mathbb{R}$ grows sufficiently rapidly. The entries are independent iff V is proportional to x^2 : non-trivial matter to sample ensembles for general V (see Olver, Rao and Trogdon (2015) [ORT15]).

Histograms for the deflation time fluctuations are given below.

Dependent QR Fluctuations




Figure: The observation of two-component universality for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when $\mathcal{A} = \text{QR}$, $\mathcal{E} = \text{QUE, COSH, GUE}$ and $\epsilon = 10^{-10}$. Here we are using deflation time (= halting time), as in [2012]. The left figure displays three histograms, one each for GUE, COSH and QUE, when $N = 70$. The right figure displays the same information for $N = 150$. All histograms are produced with 16,000 samples. Again, we see that two-component universality emerges for N sufficiently large: the histograms follow a universal (independent of \mathcal{E}) law. This is surprising because COSH and QUE have eigenvalue distributions that differ significantly from GUE in that they do not follow the so-called *semi-circle law*.

Some comments on (a)

The Gaussian Unitary Ensemble (GUE) is a complex, unitary invariant ensemble with probability distribution proportional to $e^{-N\text{tr}M^2} dM$.

The Quartic Unitary Ensemble (QUE) is a complex, unitary invariant ensemble with probability distribution proportional to $e^{-N\text{tr}M^4} dM$.

The Cosh Unitary Ensemble (COSH) has its distribution proportional to $e^{-\text{tr} \cosh M} dM$.

Both QUE and COSH do not follow the semi-circle law for the global distribution of the eigenvalues.

Some comments on (b): Conjugate gradient fluctuations

(b): Here the authors started to address the question of whether two-component universality is just a feature of eigenvalue computation, or is present more generally in numerical computation. In particular, the authors considered the solution of the linear system of equations $Wx = b$ where W is a real and positive definite, using the conjugate gradient (CG) method.

The method is iterative and at iteration k of the algorithm an approximate solution x_k of $Wx = b$ is found and the residual $r_k = Wx_k - b$ is computed. For any given $\epsilon > 0$, the method is halted when $\|r_k\|_2 < \epsilon$, and the halting time $k_\epsilon(W, b)$ recorded.

Here the authors considered $N \times N$ matrices A chosen from two different positive definite ensembles \mathcal{E} and vectors $b = (b_j)$ chosen independently with iid entries $\{b_j\}$. Given ϵ (small) and N (large), and $(W, b) \in \mathcal{E}$, the authors record the halting time $k_{\epsilon, N, \mathcal{A}, \mathcal{E}}$, $A = \text{CG}$, and compute the fluctuations $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}(W, b)$. The histograms for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ are given below, and again, two-component universality is evident.

Conjugate Gradient Fluctuations




Figure: The observation of two-component universality for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when $\mathcal{A} = \text{CG}$ and $\mathcal{E} = \text{cLOE}, \text{cPBE}$ with $\epsilon = 10^{-10}$. The left figure displays two histograms, one for cLOE and one for cPBE, when $N = 100$. The right figure displays the same information for $N = 500$. All histograms are produced with 16,000 samples. Again, we see two-component universality emerges for N sufficiently large: the histograms follow a universal (independent of \mathcal{E}) law. With the chosen scaling, we see two-component universality emerge for N sufficiently large: the histograms follow a universal (independent of \mathcal{E}) law.

The critically-scaled Laguerre Orthogonal Ensemble (cLOE) is given by XX^T/m where X is an $N \times m$ matrix with iid Gaussian (mean zero, variance one) entries. The critically-scaled positive definite Bernoulli ensemble (cPBE) is given by XX^T/m where X is an $N \times m$ matrix consisting of iid Bernoulli variables taking the values ± 1 with equal probability.

The critical scaling refers to the choice $m = N + 2\lfloor\sqrt{N}\rfloor$.

Some comments on (b): The GMRES Algorithm

The authors again considered the solution of $Wx = b$ but here W has the form $I + X$ and $X \equiv X_n$ is a random, real non-symmetric matrix and $b = (b_j)$ is independent with uniform iid entries $\{b_j\}$. As $W = I + X$ is (almost surely) no longer positive definite the conjugate gradient algorithm breaks down, and the authors solve $(I + X)x = b$ using the Generalized Minimal Residual (GMRES) algorithm.

Again, the algorithm is iterative and at iteration k of the algorithm an approximate solution x_k of $(I + X)x = b$ is found and the residual $r_k = (I + X)x_k - b$ is computed. As before, for any given $\epsilon > 0$, the method is halted when $\|r_k\|_2 < \epsilon$ and $k_{\epsilon, n, \mathcal{A}, \mathcal{E}}(X, b)$ is recorded. For these computations X is chosen from two distinct ensembles. As in the conjugate gradient problem, the authors compute the histograms for the fluctuations of the halting time $\tau_{\epsilon, n, \mathcal{A}, \mathcal{E}}$ for two ensembles E , where now $A = \text{GMRES}$. The results are given below, where again two component universality is evident.

Some comments on (b)

The critically-scaled shifted Bernoulli Ensemble (cSBE) is given by $I + X/\sqrt{N}$ where X is an $N \times N$ matrix consisting of iid Bernoulli variables taking the values ± 1 with equal probability. The critically-scaled shifted Ginibre Ensemble (cSGE) is given by $I + X/\sqrt{N}$ where X is an $N \times N$ matrix of iid Gaussian variables with mean zero and variance one.

The scaling is chosen so that $\mathbb{P}(\left| \|X/\sqrt{N}\| - 2 \right| > \epsilon)$ tends to zero as $N \rightarrow \infty$.

GMRES Fluctuations

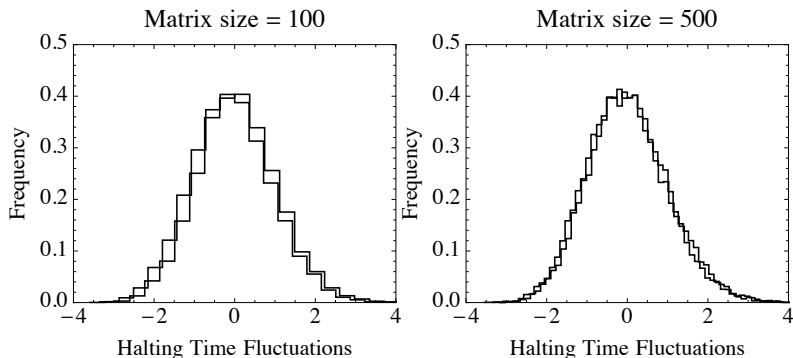


Figure: The observation of two-component universality for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when $\mathcal{A} = \text{GMRES}$, $\mathcal{E} = \text{cSGE}$, cSBE and $\epsilon = 10^{-8}$. The left figure displays two histograms, one for cSGE and one for cSBE, when $N = 100$. The right figure displays the same information for $N = 500$. All histograms are produced with 16,000 samples. We see two-component universality emerge for N sufficiently large: the histograms follow a universal (independent of \mathcal{E}) law.

Some comments on (c): Infinite-dimensional problems

(c): Here the authors raised the issue of whether two-component universality is just a feature of finite-dimensional computation, or is also present in problems which are intrinsically infinite dimensional.

What about PDEs?

In particular, is the universality present in numerical computations for PDEs? As a case study, the authors consider the numerical solution of the Dirichlet problem $\Delta u = 0$ in a star-shaped region $\Omega \subset \mathbb{R}^2$ with $u = f$ on $\partial\Omega$. In this case, the boundary is described by a periodic function of the angle θ , $r = r(\theta)$, and similarly $f = f(\theta)$, $0 \leq \theta \leq 2\pi$.

Two ensembles, BDE and UDE (described below), are derived from a discretization of the problem with specific choices for r , defined by a random Fourier series. The boundary condition f is chosen randomly by letting $\{f(\frac{2\pi j}{N})\}_{j=0}^{N-1}$ be iid uniform on $[-1, 1]$. Histograms for the halting time $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ from these computations are given below and again, two-component universality is evident.

Let Ω be the star-shaped region interior to the curve $(x, y) = (r(\theta) \cos(\theta), r(\theta) \sin(\theta))$ where $r(\theta)$ is given by

$$r(\theta) = 1 + \sum_{j=1}^m (X_j \cos(j\theta) + Y_j \sin(j\theta)), \quad 0 \leq \theta < 2\pi$$

and X_j and Y_j are iid random variables taking values in $[-1/(2m), 1/(2m)]$. Dividing by $2m$ eliminates the possibility that r vanishes. The double-layer potential formulation of the boundary integral equation

$$\pi u(P) - \int_{\partial\Omega} u(Q) \frac{\partial}{\partial n_Q} \log |P - Q| dS_Q = -f(P), \quad P \in \partial\Omega,$$

is solved by discretizing in θ with N points and applying the trapezoidal rule choosing $N = 2m$.

The Bernoulli Dirichlet Ensemble (BDE) is the case where X_m and Y_m are Bernoulli variables taking values $\pm 1/(2m)$ with equal probability. The Uniform Dirichlet Ensemble (UDE) is the case where X_m and Y_m are uniform variables on $[-1/(2m), 1/(2m)]$.

More GMRES Fluctuations




Figure: The observation of two-component universality for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when $\mathcal{A} = \text{GMRES}$, $\mathcal{E} = \text{UDE, BDE}$ and $\epsilon = 10^{-8}$. The left figure displays two histograms, one for UDE and BDE, when $N = 100$. The right figure displays the same information for $N = 500$. All histograms are produced with 16,000 samples. We see two-component universality emerge for N sufficiently large: the histograms follow a universal (independent of \mathcal{E}) law.

The following figure conflates the previous computations from GMRES applied to the shifted ensembles and GMRES applied to the Dirichlet problem given above.

All GMRES Fluctuations


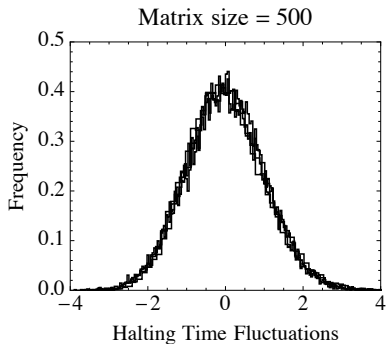


Figure: This figure consists of four histograms, two taken from GMRES applied to the previous shifted ensembles and two taken from GMRES applied to the Dirichlet problem.



What is surprising, and quite remarkable, about these computations is that the histograms in the case of the Dirichlet problem are the **same** as the histograms for the shifted ensembles. In other words, UDE and BDE are structured with random components, whereas cSGE and cSBE have no structure, yet they produce the same statistics.

Some comments on (d): A genetic algorithm

(d): In all the computations discussed so far, the randomness in the computations resides in the initial data. In the next set of computations in [DMOT14], the authors considered an algorithm which is intrinsically stochastic. In particular, they considered a **genetic algorithm**, which they used to compute **Fekete points**. Such points $P^* = (P_1^*, P_2^*, \dots, P_N^*) \in \mathbb{R}^N$ are the global minimizers of the objective function

$$H(P) = \frac{2}{N(N-1)} \sum_{1 \leq i \neq j \leq N} \log |P_i - P_j|^{-1} + \frac{1}{N} \sum_{i=1}^N V(P_i)$$

for real-valued functions $V = V(x)$ which grow sufficiently rapidly as $|x| \rightarrow \infty$. It is well-known that as $N \rightarrow \infty$, the counting measures $\delta_{P^*} = \frac{1}{N} \sum_{i=1}^N \delta_{P_i^*}$ converge to the so-called equilibrium measure μ_V which plays a key role in the asymptotic theory of the orthogonal polynomials generated by the measure $e^{-NV(x)} dx$ on \mathbb{R} . Genetic algorithms are particularly useful for large scale optimization problems, such as those that occur, for example, in the financial industry, and involve two basic components, “mutation” and “crossover”. The authors implemented the genetic algorithm in the following way.

Fix a distribution \mathcal{D} on \mathbb{R} . Draw an initial population $\mathcal{P}_0 = \mathcal{P} = \{P_i\}_{i=1}^n$ consisting of $n = 100$ vectors in \mathbb{R}^N , N large, with elements that are iid uniform on $[-4, 4]$. The random map $F_{\mathcal{D}}(\mathcal{P}) : (\mathbb{R}^N)^n \rightarrow (\mathbb{R}^N)^n$ is defined by one of the following two procedures:

$$\begin{array}{cccc}
 & P_1 & \times & \times & \cdots & \times \\
 & P_2 & \times & \times & \cdots & \times \\
 & & & \vdots & & \\
 & P_{100} & \times & \times & \cdots & \times \\
 & & \underbrace{\hspace{10em}} & & & \\
 & & & N & &
 \end{array}$$

Mutation

Pick one individual $P \in \mathcal{P}$ at random (uniformly). Then pick two integers n_1, n_2 from $\{1, 2, \dots, N\}$ at random (uniformly and independent). Three new individuals are created.

- ▶ \tilde{P}_1 — draw n_1 iid numbers $\{x_1, \dots, x_{n_1}\}$ from \mathfrak{D} and perturb the first n_1 elements : $(\tilde{P}_1)_i = (P)_i + x_i, i = 1, \dots, n_1$, and $(\tilde{P}_1)_i = (P)_i$ for $i > n_1$.
- ▶ \tilde{P}_2 — draw $N - n_2$ iid numbers $\{y_{n_2+1}, \dots, y_N\}$ from \mathfrak{D} and perturb the last $N - n_2$ elements of P : $(\tilde{P}_2)_i = (P)_i + y_i, i = n_2 + 1, \dots, N$, and $(\tilde{P}_2)_i = (P)_i$ for $i \leq n_2$.
- ▶ \tilde{P}_3 — draw $|n_1 - n_2|$ iid numbers $\{z_1, \dots, z_{|n_1 - n_2|}\}$ from \mathfrak{D} and perturb elements $n_1^* = 1 + \min(n_1, n_2)$ through $n_2^* = \max(n_1, n_2)$:
 $(\tilde{P}_3)_i = (P)_i + z_{i-n_1^*+1}, i = n_1^*, \dots, n_2^*$, and $(\tilde{P}_3)_i = (P)_i$ for $i \notin \{n_1^*, \dots, n_2^*\}$.

Pick two individuals P , Q from \mathcal{P} at random (independent and uniformly). Then pick two numbers n_1 , n_2 from $\{1, 2, \dots, N\}$ (independent and uniformly). Two new individuals are created.

- ▶ \tilde{P}_4 — Replace the n_1 th element of P with the n_2 th element of Q and perturb it (additively) with a sample of \mathfrak{D} .
- ▶ \tilde{P}_5 — Replace the n_1 th element of Q with the n_2 th element of T and perturb it (additively) with a sample of \mathfrak{D} .

At each step, the application of either crossover or mutation is chosen with equal probability.

The new individuals are appended to \mathcal{P} (after mutation we have $\tilde{\mathcal{P}} = \mathcal{P} \cup \{\tilde{P}_1, \tilde{P}_2, \tilde{P}_3\}$ and after crossover we have $\tilde{\mathcal{P}} = \mathcal{P} \cup \{\tilde{P}_4, \tilde{P}_5\}$) and $\mathcal{P} \mapsto \mathcal{P}' = F_{\mathcal{D}}(\mathcal{P}) \in (\mathbb{R}^N)^n$ is constructed by choosing the 100 P_i 's in $\tilde{\mathcal{P}}$ which yield the smallest values of $H(P)$. The algorithm produces a sequence of populations $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k, \dots$ in $(\mathbb{R}^N)^n$, $\mathcal{P}_{k+1} = F_{\mathcal{D}}(\mathcal{P}_k)$, $n = 100$, and halts, with halting time recorded, for a given ϵ , when $\min_{P \in \mathcal{P}_k} H(P) - \inf_{P \in \mathbb{R}^N} H(P) < \epsilon$.

The histograms for the fluctuations $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$, with $A = \text{Genetic}$ are given below, for two choices of V , $V(x) = x^2$ and $V(x) = x^4 - 3x^2$, and different choices of $E \simeq \mathcal{D}$. Again, two-component universality is evident.

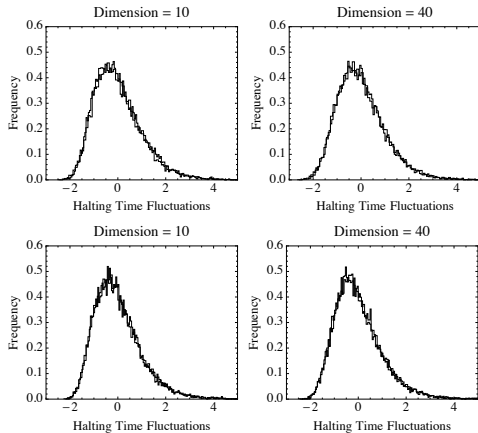


Figure: The observation of two-component universality for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when $\mathcal{A} = \text{Genetic}$, $\epsilon = 10^{-2}$ and $\mathcal{E} \simeq \mathfrak{D}$ where \mathfrak{D} is chosen to be either uniform on $[-1/(10N), 1/(10N)]$ or taking values $\pm 1/(10N)$ with equal probability. The top row is created with the choice $V(x) = x^2$ and the bottom row with $V(x) = x^4 - 3x^2$. Each of the plots in the left column displays two histograms, one for each choice of \mathfrak{D} when $N = 10$. The right column displays the same information for $N = 40$. All histograms are produced with 16,000 samples. The equilibrium measure for $V(x) = x^2$ is supported on one interval whereas the equilibrium measure for $V(x) = x^4 - 3x^2$ is supported on two intervals. It is evident that the histograms collapse onto a universal curve, one for each V .

Some comments on (e): Decision making model

(e): In the final set of computations in [DMOT14], the authors picked up on a common notion in neuroscience that the human brain is a computer with software and hardware. If this is indeed so, then one may speculate that two-component universality should be present certainly in some cognitive actions.

The authors focused on recent work of Bakhtin and Correll [BC12], who have conducted and analyzed the data obtained from experiments with 45 human participants. The participants are shown 200 pairs of images. The images in each pair consist of nine black disks of variable size. The disks in the images within each pair have approximately the same area so that there is no *a priori* bias.

Y Bakhtin and J Correll. A neural computation model for decision-making times. *J. Math. Psychol.*, 56(5):333–340, 2012

Some comments on (e): Decision making model




Fig. 3. An example of a pair of images generated for using in the experiment.

Some comments on (e): Decision making model

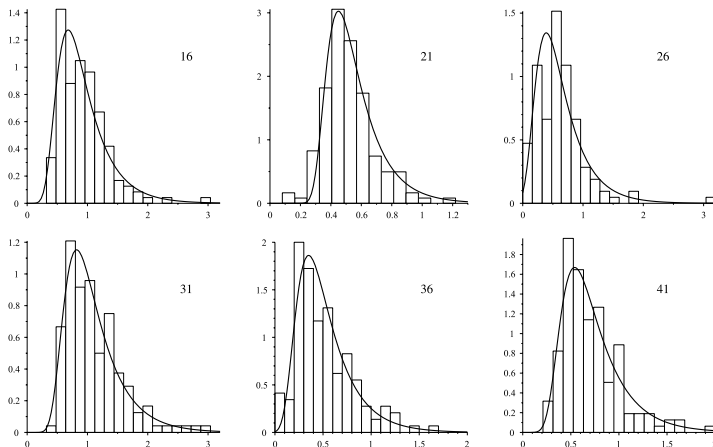
(e): In the final set of computations in [DMOT14], the authors picked up on a common notion in neuroscience that the human brain is a computer with software and hardware. If this is indeed so, then one may speculate that two-component universality should be present certainly in some cognitive actions.

The authors focused on recent work of Bakhtin and Correll [BC12], who have conducted and analyzed the data obtained from experiments with 45 human participants. The participants are shown 200 pairs of images. The images in each pair consist of nine black disks of variable size. The disks in the images within each pair have approximately the same area so that there is no *a priori* bias.

The participants are then asked to decide which of the two images covers the larger (black) area. Bakhtin and Correll then record the time T that it takes for each participant to make a decision. For each participant, the decision times for the 200 pairs are collected and the fluctuation histogram is tabulated. They then compare their experimental results with a dynamical Curie–Weiss model frequently used in describing decision processes, resulting in good agreement.

Y Bakhtin and J Correll. A neural computation model for decision-making times. *J. Math. Psychol.*, 56(5):333–340, 2012

Universality



The solid curve is the (shifted and scaled) Gumbel distribution

$$f_{BC}(x) = \sigma g(\sigma x + \mu), \quad g(z) = \exp(-z - e^{-z})$$

predicted by the Curie–Weiss model of Bakhtin–Correll.

At its essence the model is Glauber dynamics on the hypercube $\{-1, 1\}^N$ with a microscopic approximation of a drift-diffusion process. Consider N variables $\{X_i(t)\}_{i=1}^N$, $X_i(t) \in \{-1, 1\}$. The state of the system at time t is $X(t) = (X_1(t), X_2(t), \dots, X_N(t))$. The transition probabilities are given through the expressions

$$\mathbb{P}(X_i(t + \Delta t) \neq X_i(t) | X(t) = x) = c_i(x)\Delta t + o(\Delta t),$$

where $c_i(x)$ is the spin flip intensity. The observable considered is

$$M(X(t)) = \frac{1}{N} \sum_{i=1}^N X_i(t) \in [-1, 1],$$

and the initial state of the system is chosen so that $M(X(0)) = 0$, a state with no *a priori* bias, as in the case of the experimental setup.

Given $\epsilon \in (0, 1)$, which may not be small, the halting (or decision) time for this model is $k = \inf\{t : |M(X(t))| \geq \epsilon\}$, the time at which the system makes a decision.

Following standard procedures, this model is simulated by first sampling an exponential random variable with mean

$$\left(\sum_i c_i(X(t)) \right)^{-1}$$

to find the time increment Δt at which the system changes state. With probability one, just a single spin flipped.

One determines which spin flips by sampling a random variable Y with distribution

$$\mathbb{P}(Y = i) = \frac{c_i(X(t))}{\sum_i c_i(X(t))}, \quad i = 1, 2, \dots, N,$$

so producing an integer j . Define

$$\begin{aligned} X_i(t + s) &\equiv X_i(t) \text{ if } s \in [0, \Delta t) \text{ for } i = 1, 2, \dots, N, \\ X_i(t + \Delta t) &\equiv X_i(t), \text{ if } i \neq j, \\ X_j(t + \Delta t) &\equiv -X_j(t). \end{aligned}$$

This procedure is repeated with t replaced by $t + \Delta t$ to evolve the system.

Central to the application of the model is the assumption on the statistics of the spin flip intensity $c_i(x)$. The authors in the present paper raised the following question.

If one changes the basic statistics of the c_i 's, will the limiting histograms for the fluctuations of k be affected as N becomes large?

In response to this question the authors considered the following choices for $E \simeq c_i(x)$ ($\beta = 1.3$):

1. $c_i(x) = o_i(x) = e^{-\beta x_i M(x)}$ (the case studied by [BC12]),
2. $c_i(x) = u_i(x) = e^{-\beta x_i (M(x) - M^3(x)/5)}$,
3. $c_i(x) = v_i(x) = e^{-\beta x_i (M(x) + M^8(x))}$.

The histograms for the fluctuations $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ of k are given below for all three choices of c_i . Once again, two-component universality is evident. Thus these computations demonstrate two-component universality for a range of decision process models.




Figure: The observation of two-component universality for $\tau_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ when $\mathcal{A} = \text{Curie-Weiss}$, $\mathcal{E} \simeq o_i$, $u_i, v_i, \epsilon = .5$ and $\beta = 1.3$. The left figure displays three histograms, one for each choice of \mathcal{E} when $N = 50$. The right figure displays the same information for $N = 200$. All histograms are produced with 16,000 samples. The histogram for $\mathcal{E} = o_i$ corresponds to the case studied by [BC12]. It is clear from these computations that the fluctuations collapse on to the universal curve for $\mathcal{E} = o_i$. Thus, reasonable changes in the spin flip intensity do not appear to change the limiting histogram. This indicates why the specific choice made in [BC12] of $\mathcal{E} = o_i$ is perhaps enough to capture the behavior of many individuals.


Two interesting observations

(1) Google searches: When you perform a `Google` search, it actually tells you how long the search took and one can look at the halting time histogram. Trogdon et al considered ensembles of 3000 english words and 3000 Turkish words and they found the following:

L Sagun, T Trogdon, and Y LeCun. Universal halting times in optimization and machine learning. *Q. Appl. Math.*, nov 2017

Two interesting observations

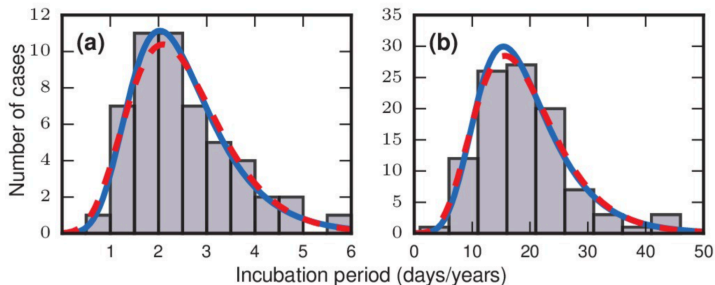
(1) Google searches: When you perform a Google search, it actually tells you how long the search took and one can look at the halting time histogram. Trogdon et al considered ensembles of 3000 english words and 3000 Turkish words and they found the following:



L Sagun, T Trogdon, and Y LeCun. Universal halting times in optimization and machine learning. *Q. Appl. Math.*, nov 2017

Two interesting observations

(2) **Universality in incubation times:** Figure and caption reproduced from [Bah18].



Solid curves are Gumbel densities f_{BC} predicted in [OLSS17] using a graphical model. (a) Data from an outbreak of food-borne streptococcal sore throat, reported by Sartwell (1995) [Sar50], time measured in days. (b) Data from a study of bladder tumors among workers following occupational exposure to a carcinogen in a dye plant, see Goldblatt (1949) [Gol49]. Time measured in years.

Y Bahktin. Universal Statistics of Incubation Periods and Other Detection Times via Diffusion Models. *draft*, 2018

B Ottino-Loffler, J G Scott, and S H Strogatz. Evolutionary dynamics of incubation periods. *Elife*, 6, dec 2017

Bakhtin was able to reproduce the f_{BC} Gumbel distribution using no-called reactive paths for a simple stochastic ODE

$$dx(t) = b(x(t))dt + \sigma(x(t))dW(t)$$

W — Wiener process

So we have an interesting conundrum. What is it exactly that is in common with

- ▶ the neural stochastics of the participants in the Bakhtin–Correll experiment
- ▶ the universality of the Curie–Weiss model
- ▶ the Google searches, and
- ▶ the statistics of incubation times?

All the above results are numerical and experimental. In order to establish universality as a bona fide phenomenon in numerical analysis, and not just an artifact, suggested, however strongly, by certain computations as above, the authors in [DT2016] analyzed a particular algorithm of interest, viz, the Toda algorithm to compute the largest eigenvalue of a random $N \times N$ symmetric matrix. More precisely, they considered the Toda lattice given before

$$\frac{dH}{dt} = [H, B(H)], \quad B(H) = H_- - H_-^T$$

with $H(0) = H$ but now with stopping time $T = T_{\epsilon, N, \mathcal{A}, \mathcal{E}}$ given by

$$E(T) = \sum_{j=2}^N |H_{1j}(T)|^2 = \epsilon^2.$$

By perturbation theory

$$|H_{11}(T) - \lambda_j| \leq \epsilon$$

for some eigenvalue λ_j of $H = H(0)$.

As Toda is an ordering algorithm, $\lambda_j = \lambda_N$, the largest eigenvalue of H , and so with high probability as $N \rightarrow \infty$, T controls the computation of the largest eigenvalue of H . Here $H = H(0)$ is chosen from a very wide variety of invariant ensembles (IEs) and Wigner ensembles (WEs), and it turns out that the analysis of $T = T_{\epsilon, N, \mathcal{A}, \mathcal{E}}(H)$ depends in a crucial way on recent results from random matrix theory (RMT) that are at the forefront of current knowledge.

The gap $\lambda_N - \lambda_{N-1}$ between the two largest eigenvalues of H plays a central role in describing the statistics of T . The following definition quantifies the distribution F_β^{gap} of the inverse of $\lambda_N - \lambda_{N-1}$ on the appropriate scale ($\beta = 1$ real symmetric, $\beta = 2$ complex Hermitian)

$$F_\beta^{\text{gap}}(t) = \lim_{N \rightarrow \infty} \text{Prob} \left(\frac{1}{c_V^{2/3} 2^{-2/3} N^{2/3} (\lambda_N - \lambda_{N-1})} \leq t \right), \quad t > 0, \quad (5)$$

where c_V is an explicit constant which depends on the ensemble. It is a non-trivial result in RMT that the limit (5) exists.

The main result in [DT16] is the following.

Theorem 1

Universality for the Toda algorithm Let $\sigma > 5/3$ be fixed and let (ϵ, N) be in the scaling region

$$\frac{\log \epsilon^{-1}}{\log N} \geq \sigma.$$

Then if H is distributed according to an real ($\beta = 1$) or complex ($\beta = 2$) invariant or Wigner ensemble

$$\lim_{N \rightarrow \infty} \text{Prob} \left(\frac{T}{\tilde{c}_V^{2/3} 2^{-2/3} N^{2/3} (\log \epsilon^{-1} - \frac{2}{3} \log N)} \leq t \right) = F_\beta^{\text{gap}}(t), \quad t \geq 0.$$

Furthermore

$$\epsilon^{-1} |\lambda_N - H_{11}(T)|$$

converges to zero in probability as $N \rightarrow \infty$.

Thus T , suitably scaled, behaves statistically as $N \rightarrow \infty$, like the inverse gap $(\lambda_N - \lambda_{N-1})^{-1}$.

Some comments:

(1): For $\beta = 2$ one can show that for (ϵ, N) in the scaling region

$$\begin{aligned}\mathbb{E}(T) &= \tilde{c}_V^{3/2} 2^{-2/3} N^{2/3} (\log \epsilon^{-1} - \frac{2}{3} \log N) \mathbb{E}(\xi) (1 + o(1)) \\ &= \tilde{c}_V^{3/2} 2^{-2/3} N^{2/3} \log N \left(\frac{\log \epsilon^{-1}}{\log N} - \frac{2}{3} \right) \mathbb{E}(\xi) (1 + o(1))\end{aligned}$$

where ξ is a random variable with distribution $F_{\beta=2}^{\text{gap}}(t)$ and an analogous result for $\text{Var}(T)^{1/2}$.

This kind of result

$$\mathbb{E}(t) \sim N^{2/3} \log N$$

is new: standard results on the statistics of eigenvalue computation give bounds which are typically too large.

Numerical comparisons




Figure: The simulated rescaled histogram for T for both BUE and GUE. Here $\epsilon = 10^{-14}$ and $N = 500$ with 250,000 samples. The solid curve is the rescaled density $f_2^{\text{gap}}(t) = d/dt F_2^{\text{gap}}(t)$.

(2): Note that if $\epsilon = 10^{-16}$ and $N < 10^9$ then

$$\frac{\log \epsilon^{-1}}{\log N} = \frac{16}{9} > \frac{5}{3}$$

so computations that arise in practice typically lie in the scaling region.

(3): Similar theorems have been proved for other algorithms such as QR on ensembles of positive definite matrices [DT17].

(4): The proof of the theorem rests on the fact that the Toda lattice is completely integrable. Using results from Moser from the 1970's one obtains an explicit formula for $E(t)$ in terms of the eigenvalues $\{\lambda_j\}$ and the first components of the associated eigenvectors $\{u_{1j}\}$ of H , $Hu_j = \lambda_j u_j$

$$E(t) = \sum_{j=1}^N (\lambda_j - H_{11}(t))^2 |u_{1j}(t)|^2,$$
$$H_{11}(t) = \sum_{j=1}^N \lambda_j |u_{1j}(t)|^2,$$
$$u_{1j}(t) = u_{1j} e^{\lambda_j t} / \left(\sum_{i=1}^N |u_{1i}(t)|^2 e^{2\lambda_i t} \right)^{1/2}.$$

The analysis of the condition

$$E(T) = \epsilon^2$$

thus reduces to calculus with random variables $\{\lambda_j\}$, $\{u_j\}$, whose precise statistical properties have only been established in the last 3-6 years due to the works of Yau, Erdős, Bourgade and their collaborators.

Most recently, together with Steve Miller, we have been investigating universality properties of cyber algorithms, such as RSA, but that is another story...



Y Bahktin.

Universal Statistics of Incubation Periods and Other Detection Times via Diffusion Models.

draft, 2018.



Y Bakhtin and J Correll.

A neural computation model for decision-making times.

J. Math. Psychol., 56(5):333–340, 2012.



P A Deift, G Menon, S Olver, and T Trogdon.

Universality in numerical computations with random data.

Proc. Natl. Acad. Sci. U. S. A., 111(42):14973–8, oct 2014.



P Deift and T Trogdon.

Universality for Eigenvalue Algorithms on Sample Covariance Matrices.

SIAM J. Numer. Anal., 55(6):2835–2862, jan 2017.



P Deift and T Trogdon.

Universality for the Toda Algorithm to Compute the Largest Eigenvalue of a Random Matrix.

Commun. Pure Appl. Math., 71(3):505–536, mar 2018.



M W Goldblatt.

Vesical tumours induced by chemical compounds.

Br J Indust Med, 6:65–81, 1949.



B Ottino-Loffler, J G Scott, and S H Strogatz.
Evolutionary dynamics of incubation periods.
Elife, 6, dec 2017.



S Olver, N R Rao, and T Trogdon.
Sampling unitary ensembles.
Random Matrices Theory Appl., 4(1):1550002, 2015.



Christian W. Pfrang, Percy Deift, and Govind Menon.
How long does it take to compute the eigenvalues of a random symmetric matrix?
arXiv Prepr. arXiv1203.4635, mar 2012.



P E Sartwell.
The distribution of incubation periods of infectious disease. 1949.
Am. J. Epidemiol., 41(3):310–318, 1950.



L Sagun, T Trogdon, and Y LeCun.
Universal halting times in optimization and machine learning.
Q. Appl. Math., nov 2017.