

Stochastic weight matrix dynamics during learning and Dyson Brownian motion

Gert Aarts ^{1,*}, Biagio Lucini ^{2,†} and Chanju Park ^{1,‡}

¹*Department of Physics, Swansea University, Swansea SA2 8PP, United Kingdom*

²*Department of Mathematics, Swansea University (Bay Campus), Swansea SA1 8EN, United Kingdom*



(Received 14 August 2024; accepted 16 December 2024; published 8 January 2025)

We demonstrate that the update of weight matrices in learning algorithms can be described in the framework of Dyson Brownian motion, thereby inheriting many features of random matrix theory. We relate the level of stochasticity to the ratio of the learning rate and the minibatch size, providing more robust evidence to a previously conjectured scaling relationship. We discuss universal and nonuniversal features in the resulting Coulomb gas distribution and identify the Wigner surmise and Wigner semicircle explicitly in a teacher-student model and in the (near-)solvable case of the Gaussian restricted Boltzmann machine.

DOI: [10.1103/PhysRevE.111.015303](https://doi.org/10.1103/PhysRevE.111.015303)

I. INTRODUCTION

Recent years have seen a dramatic increase of the use of machine learning (ML) in the fundamental sciences, with the adoption and development of many ML applications to improve and speed up scientific analysis [1]. Reversely, there is a growing trend to use the methodology of (theoretical) physics to understand ML algorithms, viewing these as acting on systems with many fluctuating degrees of freedom and hence employing the analogy with statistical physics. In this paper we explore this second direction and argue that learning can be formulated in the framework of Dyson Brownian motion [2], thereby inheriting many features of random matrix theory (RMT) [3–8].

In general terms, ML algorithms aim to minimize some cost function by applying stochastic gradient descent (SGD), or variations thereof, to weight matrices W defined inside the architecture. Stochastic updates of matrices immediately establishes the link with Dyson Brownian motion, which exactly describes those, and yields an equation for the dynamics of the eigenvalues of the symmetric combination $X = W^T W$. This dynamics contains universal aspects, e.g., eigenvalue repulsion due to an induced Coulomb term, as reflected in the Wigner surmise, as well as nonuniversal aspects, related to details of the gradient of the loss function. In many algorithms, stochasticity is inherently present due sampling and minibatch updates. We demonstrate that the distribution of eigenvalues depends on the ratio of the learning rate α and minibatch size $|\mathcal{B}|$, and not on these quantities separately. This linear

scaling rule has been observed previously [9–12] (see also Ref. [13]), but we demonstrate that it is a direct consequence of stochastic matrix dynamics and Dyson Brownian motion. An interesting corollary is that there is no simple limit (e.g., $\alpha \rightarrow 0$, $1/|\mathcal{B}| \rightarrow 0$) in which SGD reduces to a stochastic differential equation (SDE) in continuous time [14,15]. Again, this has been noted before [16], but we emphasize the viewpoint that a tunable ratio $\alpha/|\mathcal{B}|$ in fact prevents an obvious SDE limit (in the limit $\alpha/|\mathcal{B}| \rightarrow 0$, the SGD update reduces to an ordinary differential equation for the eigenvalues). The amount of stochasticity is directly proportional to $\alpha/|\mathcal{B}|$ and therefore determines the strength of the eigenvalue repulsion. It hence sets a fundamental limit on the accuracy of learning, but reversely a larger level of stochasticity leads to better generalization, as it avoids overfitting [9–11]. These concepts find a natural place in the framework of Dyson Brownian motion, as we will demonstrate.

This paper is organized as follows. After setting up the general framework in Sec. II, we demonstrate our arguments explicitly in two cases, namely the Gaussian restricted Boltzmann machine (RBM) and a simple teacher-student model in Sec. III. The conclusions are summarized in Sec. IV, along with directions for future research. Appendix A contains a brief summary of Dyson Brownian motion, while Appendices B, C, and D contain some further comments used in the main text.

With relation to RMT, we note that Refs. [17,18] apply RMT to the distribution of eigenvalues of weight matrices in an empirical manner, but not in the framework of Dyson Brownian motion and the linear scaling rule. In Ref. [19] it was shown that “real-world data” have many features consistent with RMT. Here we apply RMT to weight matrices, not to data.

II. STOCHASTIC WEIGHT MATRIX DYNAMICS

A. Stochastic updates

Let us consider some weight matrix W , which connects layers of nodes in a neural network or a restricted

*Contact author: g.aarts@swansea.ac.uk

†Contact author: b.lucini@swansea.ac.uk

‡Contact author: chanju.b.park@gmail.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

Boltzmann machine. During the learning stage, it is updated using stochastic gradient descent, or variations thereof, by subtracting the change in the loss function, $\mathcal{L}[W]$,

$$W_{ij} \rightarrow W'_{ij} = W_{ij} + \delta W_{ij} = W_{ij} - \alpha \frac{\delta \mathcal{L}}{\delta W_{ij}}, \quad (1)$$

where α is the learning rate. This update is carried out using stochastically chosen minibatches \mathcal{B} of size $|\mathcal{B}|$, with a minibatch average

$$\delta W_{ij}^{\mathcal{B}} = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \delta W_{ij}^b. \quad (2)$$

The entire data set from which the minibatches are drawn is assumed to be sufficiently large to not lead to additional constraints. Since the minibatch size is finite, $\delta W^{\mathcal{B}}$ is a stochastic variable, with the size of the fluctuations set by the central limit theorem. We hence write the update as

$$\delta W_{ij} = \delta W_{ij}^{\mathcal{B}} + \frac{1}{\sqrt{|\mathcal{B}|}} \sqrt{\text{Var}(\delta W_{ij})} \eta_{ij}, \quad (3)$$

where the first term is the deterministic part of the update for a given minibatch and the second part reflects stochastic fluctuations, whose magnitude decreases with increasing minibatch size. The noise $\eta_{ij} \sim \mathcal{N}(0, 1)$ is matrix valued. The variance is defined for each matrix element, i.e.,

$$\text{Var}(\delta W_{ij}) = \langle \delta W_{ij} \delta W_{ij} \rangle - \langle \delta W_{ij} \rangle \langle \delta W_{ij} \rangle. \quad (4)$$

The update (3) is applied to each matrix element and no summation over repeated indices is implied in this section. In terms of the gradient of the loss function, this update reads

$$W'_{ij} = W_{ij} - \alpha \left(\frac{\delta \mathcal{L}}{\delta W_{ij}} \right)_{\mathcal{B}} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{\text{Var} \left(\frac{\delta \mathcal{L}}{\delta W_{ij}} \right)} \eta_{ij}, \quad (5)$$

which makes the learning rate explicit.

Since W is in general a rectangular $M \times N$ matrix, and the connection with random matrix theory is easier for symmetric matrices, we consider the symmetric combination

$$X = W^T W. \quad (6)$$

Without loss of generality we take $N \leq M$; if this is not the case, simply exchange W and W^T . X is a symmetric $N \times N$ matrix with N semipositive real eigenvalues $x_i = \xi_i^2$ ($i = 1, \dots, N$), where ξ_i are the singular values of W , obtained via the singular value decomposition $W = U \Xi V^T$. The focus on the singular values also removes the redundancy of left and right rotations on the weight matrix. The update for X follows from the one for W as

$$X \rightarrow X' = X + \delta W^T W + W^T \delta W \equiv X + \delta X, \quad (7)$$

and for finite batch size, the update is stochastic, with

$$X_{ij} \rightarrow X'_{ij} = X_{ij} + \delta X_{ij}^{\mathcal{B}} + \frac{1}{\sqrt{|\mathcal{B}|}} \sqrt{\text{Var}(\delta X_{ij})} \eta_{ij}, \quad (8)$$

as above, with symmetric noise in this case.

B. Dyson Brownian motion

The framework in which to consider stochastic matrix dynamics for a symmetric matrix X is Dyson Brownian motion

[2], which is summarized in Appendix A. The main feature is that the eigenvalues of X not only evolve stochastically but also repel, due to an induced Coulomb term. The strength of the stochastic term and of the Coulomb term are related to the stochasticity in the original matrix equation for X .

Starting from Eq. (8), the eigenvalues x_i of X evolve according to [see Eq. (A2)]

$$x_i \rightarrow x'_i = x_i + K_i + \sum_{j \neq i} \frac{g_i^2}{x_i - x_j} + \sqrt{2} g_i \eta_i, \quad (9)$$

where K_i and g_i are linked to the deterministic and stochastic terms in Eq. (8), and again $\eta \sim \mathcal{N}(0, 1)$. The term with the summation is the Coulomb term, resulting in eigenvalue repulsion. The learning rate and batch size can be made explicit by writing, cf. Eq. (5),

$$K_i = \alpha \tilde{K}_i, \quad g_i = \frac{\alpha}{\sqrt{|\mathcal{B}|}} \tilde{g}_i, \quad (10)$$

where \tilde{K}_i and \tilde{g}_i are related to the gradient of the loss function and its variance, respectively (quantities with a tilde are independent of the learning rate and batch size at leading order). The eigenvalue update then becomes

$$x_i \rightarrow x'_i = x_i + \alpha \tilde{K}_i + \frac{\alpha^2}{|\mathcal{B}|} \sum_{j \neq i} \frac{\tilde{g}_i^2}{x_i - x_j} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{2} \tilde{g}_i \eta_i. \quad (11)$$

When trying to identify the learning rate or the inverse batch size with a stepsize ϵ , we note here that the drift, including the Coulomb term, and the stochastic term do not scale in the standard (i.e., Itô calculus) manner. Indeed, the difficulties of going from stochastic gradient descent to a stochastic differential equation (SDE) [14,15] are well known [16], with the update reducing to an ordinary differential equation in the limit that, e.g., the learning rate goes to zero. However, as we will demonstrate shortly, the appearance of the learning rate and batch size as in Eq. (11) naturally leads to the linear scaling rule, in which $\alpha/|\mathcal{B}|$ is a parameter whose tunability can be exploited. An obvious SDE limit should therefore in fact not be expected.

We will derive the linear scaling rule [10,11] now, by considering the stationary distribution corresponding to the stochastic process (9). The associated Fokker-Planck equation for the distribution is given in Appendix A. The stationary distribution is known as the Coulomb gas, which reads

$$P_s(\{x_i\}) = \frac{1}{Z} \prod_{i < j} |x_i - x_j| e^{-\sum_i V_i(x_i)/g_i^2}, \quad (12)$$

with

$$Z = \int dx_1 \dots dx_N P_s(\{x_i\}). \quad (13)$$

Here it is assumed that the drift can be derived from a separable potential,

$$K_i(x_i) = -\frac{dV_i(x_i)}{dx_i}. \quad (14)$$

We can again make the learning rate and batch size explicit, by introducing

$$V_i(x_i) = \alpha \tilde{V}_i(x_i). \quad (15)$$

The combination in the exponent then reads

$$\frac{V_i(x_i)}{g_i^2} = \frac{1}{\alpha/|\mathcal{B}|} \frac{\tilde{V}_i(x_i)}{\tilde{g}_i^2}, \quad (16)$$

where the first factor on the right-hand side (RHS) indicates universal scaling with $\alpha/|\mathcal{B}|$ and the second factor depends on the details of the loss function.

If we assume that the potentials V_i have a minimum at $x_i = x_i^s$, such that

$$\tilde{V}_i(x_i) = \tilde{V}_i(x_i^s) + \frac{1}{2} \Omega_i (x_i - x_i^s)^2 + \dots, \quad (17)$$

with Ω_i the curvature around the minimum, the exponentials in Eq. (12) are Gaussians centered at $x_i = x_i^s$ with variance

$$\sigma_i^2 = \frac{\alpha}{|\mathcal{B}|} \frac{\tilde{g}_i^2}{\Omega_i}. \quad (18)$$

The interplay between learning rate, batch size, curvature of the loss function, and the variance of fluctuations is then easy to see. We emphasise that the assumption of a single well-defined minimum is not required in the derivation leading to Eq. (16).

We conclude that the stationary distribution after training depends on the ratio of the learning rate and batch size, and not on these quantities separately. This was observed empirically previously [10,11], but we have demonstrated that it is a direct consequence of stochastic matrix dynamics, when cast in this framework. Reference [13] contains an alternative derivation of the linear scaling rule, using a continuous-time limit of SGD in a weak sense [15], while Ref. [20] contains an application interpreting $\alpha/|\mathcal{B}|$ as an effective temperature.

The Coulomb gas structure of the stationary distribution has further, universal, consequences. Due to the Coulomb term, eigenvalues repel and the spacing between eigenvalues is nonzero. This effect is captured in the Wigner surmise, which is the distribution for the level spacings $S_i = x_{i+1} - x_i$, see Eqs. (A13) and (A15). The spectral density,

$$\rho(x) = \left\langle \frac{1}{N} \sum_{i=1}^N \delta(x - x_i) \right\rangle, \quad (19)$$

takes a universal form as well, called the Wigner semicircle, see Eq. (A17). In the following we explore both universal and nonuniversal features in two examples.

III. APPLICATIONS

In this section we apply the general framework presented above to Gaussian restricted Boltzmann machines (RBMs) and to a teacher-student model. In the former, stochasticity is inherently present, due to sampling and batch updates. In the latter, stochasticity is introduced by hand, but in such a way that it models the stochasticity presented above and observed in the RBMs.

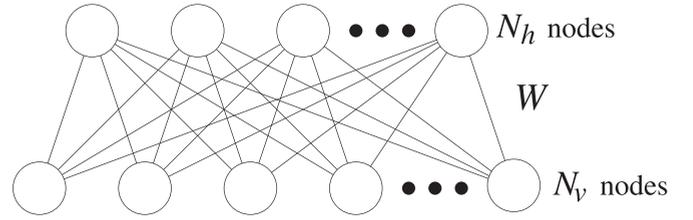


FIG. 1. Sketch of a restricted Boltzmann machine with N_v (N_h) nodes on the visible (hidden) layer, connected by the $N_v \times N_h$ matrix W .

A. Gaussian RBMs

Restricted Boltzmann machines [21,22] consist of one visible layer (with N_v nodes) and one hidden layer (with N_h nodes), see Fig. 1. There are no connections within each layer. The degrees of freedom can be discrete, as in an Ising model, continuous or mixed; Ref. [23] is a useful review. We consider Gaussian RBMs, in which both sets of degrees of freedom are quadratic. The ones on the visible layer are collected in an N_v -dimensional vector ϕ and on the hidden layer in an N_h -dimensional vector h . They are coupled bilocally via the $N_v \times N_h$ weight matrix W . We follow the notation of our previous work [24].

The probability distribution and partition function are given by

$$p(\phi, h) = \frac{1}{Z} e^{-E(\phi, h)}, \quad Z = \int D\phi Dh e^{-E(\phi, h)}, \quad (20)$$

with the energy (or action)

$$E(\phi, h) = \frac{1}{2} \mu^2 \phi^T \phi + \frac{1}{2\sigma_h^2} (h - \eta)^T (h - \eta) - \phi^T W h. \quad (21)$$

The induced distribution on the visible layer is Gaussian as well, and reads

$$p(\phi) = \int Dh p(\phi, h) = \frac{1}{Z} \exp\left(-\frac{1}{2} \phi^T K \phi + J^T \phi\right), \quad (22)$$

with the kernel and source

$$K = \mu^2 \mathbb{1} - \sigma_h^2 W W^T, \quad J = W \eta. \quad (23)$$

Here $1/\mu^2$ is the variance on the visible layer, σ_h^2 is the variance on the hidden layer, and η is a bias, which we will put to zero from now on. This model has been studied in detail from the perspective of lattice field theory (LFT) in Ref. [24], which explains the appearance of μ^2 as a mass parameter for a scalar field. Its generative power is linked explicitly to the number of hidden nodes N_h and the mass parameter μ^2 , which both act as an ultraviolet regulator, using the terminology familiar from LFT. To match to the general notation of the previous section, we identify $N_v = M$ and $N_h = N$ and take $N \leq M$.

In order to have a well-defined model, with a positive kernel $K = \mu^2 \mathbb{1} - \sigma_h^2 W W^T$, the weight matrix W must be initialized in such a way that the eigenvalues of $W W^T$ are bounded between 0 and μ^2/σ_h^2 , for any $0 < N/M \leq 1$. Initialisation is discussed in Appendix B, in terms of the Marchenko-Pastur distribution. Stable initialization is guaranteed provided the matrix elements of W are drawn from a normal distribution with variance $1/M$ and $\mu^2 > 4\sigma_h^2$.

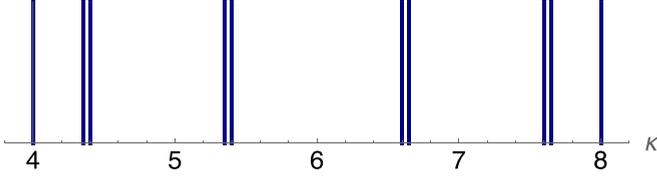


FIG. 2. Target spectrum κ_i ($i = 1, \dots, 10$): each mode, except the lowest and the highest ones, is doubly degenerate.

The model can be trained by maximizing the log-likelihood (or minimizing the Kullback-Leibler divergence), leading to the gradient,

$$\frac{\delta \mathcal{L}}{\delta W_{ia}} = \sigma_h^2 \sum_j (\langle \phi_i \phi_j \rangle_{\text{target}} - \langle \phi_i \phi_j \rangle_{\text{model}}) W_{ja}. \quad (24)$$

Here the first two-point function is evaluated using the target data and the second one is the RBM prediction. The dynamics can be analysed semianalytically [23,24] by performing a singular-value decomposition $W = U \Xi V^T$ in which U and V are orthogonal transformations in M and N dimensions, respectively, and Ξ is a rectangular matrix with N singular values ξ_i ($i = 1, \dots, N$) on the diagonal.

The kernel K on the visible layer is diagonalized as

$$K = \mu^2 \mathbb{1} - \sigma_h^2 W W^T = U D_K U^T, \quad (25)$$

where

$$D_K = \text{diag} \left(\underbrace{\mu^2 - \sigma_h^2 \xi_1^2, \mu^2 - \sigma_h^2 \xi_2^2, \dots, \mu^2 - \sigma_h^2 \xi_N^2}_N, \underbrace{\mu^2, \dots, \mu^2}_{M-N} \right). \quad (26)$$

We are interested in the N eigenvalues $\lambda_i = \mu^2 - \sigma_h^2 \xi_i^2$. It is natural to absorb σ_h^2 in $x_i = \sigma_h^2 \xi_i^2$, such that the eigenvalues read $\lambda_i = \mu^2 - x_i$. As above, we consider the symmetric matrix $X = W^T W$ (times σ_h^2) and focus on the eigenvalues x_i of X , rather than λ_i , since this makes the relation with the Coulomb gas direct. Note that positiveness and stability requires that $0 \leq x_i < \mu^2$.

As the target, we use a simple one-dimensional noninteracting scalar field theory with the spectrum given by the free dispersion relation,

$$\kappa_k = m^2 + p_{\text{lat},k}^2 = m^2 + 2 - 2 \cos \left(\frac{2\pi k}{N_v} \right), \quad (27)$$

with $-N_v/2 < k \leq N_v/2$. This spectrum is shown in Fig. 2 for the case $N_v = 10, m^2 = 4$. Each mode, except the lowest and the highest ones, is doubly degenerate. This is of interest for the Coulomb gas description derived above, since under stochastic matrix dynamics eigenvalues repel and exact degeneracy cannot be reproduced. Given a target eigenvalue κ_i and RBM mass parameter μ^2 , the exact value for x_i is given by $x_i^s = \mu^2 - \kappa_i$.

The numerical training of the RBM is carried out using persistent contrastive divergence (PCD) with minibatches, see Ref. [24] for details. The RBM mass parameter is fixed at $\mu^2 = 9$ and $N_v = N_h = 10$. The learning rate and the minibatch size can be independently varied. For each choice of

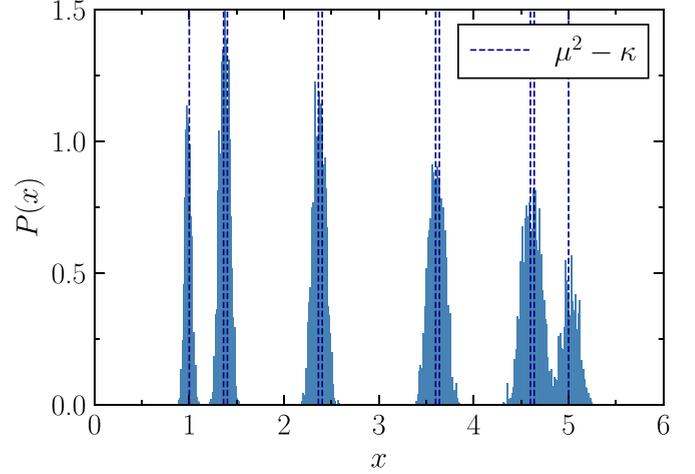


FIG. 3. Learned distributions of eigenvalues $x_i = \mu^2 - \lambda_i$. The target eigenvalues $\mu^2 - \kappa_i$ are shown with dashed vertical lines. All except the lowest and the highest target eigenvalues are doubly degenerate.

learning rate and minibatch size, we have trained 7500 RBMs, to gather statistics. A typical example of the learnt distribution of eigenvalues x_i of X is shown in Fig. 3.

B. Universal predictions

The task is to describe the distributions shown in Fig. 3 using the concepts from random matrix theory presented above. We start with the universal predictions. Due to the eigenvalue repulsion, degenerate modes cannot be learned exactly. This can be analysed using the Wigner surmise: for each of the four doubly degenerate modes we collect data on the level splitting for adjacent eigenvalues, $S_i = x_{i+1} - x_i$. Here we employed the fact that the nondegenerate eigenvalues are well separated. The resulting histograms are shown in Fig. 4 (left). These histograms can be fitted with a single-parameter ansatz [see Eq. (A13)] in terms of the mean level splitting $\langle S_i \rangle = \sqrt{\pi} \sigma_i$. Presenting the same data as a function of $s = S_i / \langle S_i \rangle$ yields the parameter-free Wigner surmise [see Eq. (A15)] leading to a universal description of the level spacing. The linear rise at small s reflects the level repulsion.

Next we turn to the spectral density (A16) for each of the four doubly degenerate modes. The results for two of these are shown in Fig. 5 (note that these are the second-to-lowest and second-to-highest histograms previously shown in Fig. 3). The lines are fits to Eq. (A17) with the position $x_{m,i}$ and width σ_i as free parameters. It is clear that the histograms are not described by simple Gaussians, but are broader and flatter. Non-Gaussianity can be checked by computing the Binder cumulant U_4 of the distribution, which equals (here $\delta x = x - x_m$)

$$U_4 \equiv \frac{\langle \delta x^4 \rangle}{3 \langle \delta x^2 \rangle^2} - 1 = -\frac{4}{27} \approx -0.148 \quad (28)$$

for the Wigner semicircle (with $N = 2$), while it vanishes for a Gaussian distribution. For completeness, we note here that we have also fitted the distributions corresponding to a single, nondegenerate level (the ones farthest to the left and right in

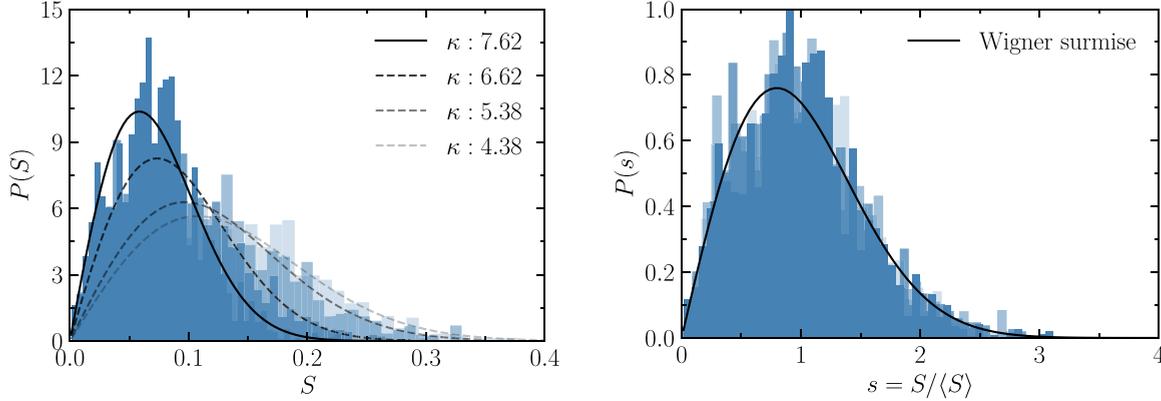


FIG. 4. Wigner surmise for the level spacing S (left) and for the rescaled $s = S/\langle S \rangle$ (right) for the four doubly degenerate modes labeled by κ . The lines on the left are fits with $\langle S \rangle = \sqrt{\pi}\sigma$ as a free parameter. The rescaled histograms on the right collapse to the universal curve, $P(s)$.

Fig. 3) and found these to be described by a Gaussian, as expected.

Importantly, from the general discussion in the previous section it follows that the higher the level of stochasticity (as indicated by g_i^2 and depending on learning rate and batch size), the wider the distributions are expected to be. We have hence repeated the exercise above for a range of learning rates, $0.01 \leq \alpha \leq 0.1$, and minibatch sizes, $4 \leq |\mathcal{B}| \leq 64$. We will demonstrate below that the potentials $V_i(x_i)$ have a global minimum at $x_i^* = \mu^2 - \kappa_i$. The variance in the Coulomb gas σ_i^2 is then proportional to the ratio of the learning rate and batch size, see Eq. (18), and both the mean level spacing $\langle S \rangle$ and the width of the spectral density are expected to scale with $\sqrt{\alpha/|\mathcal{B}|}$ times a model and data dependent function of μ^2 and κ_i .

This is demonstrated in Fig. 6. We show the response of the mean level spacing $\langle S \rangle$ (left) and the width parameter of the spectral density $\sqrt{\pi}\sigma$ (middle) to variation of the learning rate α and the batch size $|\mathcal{B}|$, as follows,

$$\langle S \rangle, \sqrt{\pi}\sigma = a_{\text{fit}} \sqrt{\frac{\alpha}{|\mathcal{B}|}} \kappa_i^2 \Omega_i. \quad (29)$$

A derivation of the nonuniversal factor $\kappa_i^2 \Omega_i = (\mu^2 - \kappa_i)$, using an analysis of the variance of the gradient of the loss

function, can be found in Appendix C. A linear dependence of $\langle S \rangle^2$ and σ^2 on $\alpha/|\mathcal{B}|$ is observed, consistent with the derivation in Sec. II.

Both fits are independent ways to probe the Coulomb gas description of the joint distribution $P(x_i, x_{i+1})$, which depends on one parameter σ_i only. Hence in Fig. 6 (right) we show the expected linear relation between $\langle S \rangle$ and $\sqrt{\pi}\sigma_i$, to provide further support for this description.

The Coulomb term not only breaks the degeneracy of the doubly degenerate modes, but also leads to a repulsion between all eigenvalues. This can be analyzed by determining the mean positions $x_{m,i}$ of the spectral densities as a function of $\alpha/|\mathcal{B}|$. In Fig. 7 we show the ratio of the RBM eigenvalues $\lambda_i = \mu^2 - x_{m,i}$ and the target eigenvalues κ_i as a function of $\alpha/|\mathcal{B}|$, where α and $|\mathcal{B}|$ are independently varied. We conclude that the spectrum of the target theory will only be learnt exactly in the limit that the stochasticity goes to zero, i.e., when $\alpha/|\mathcal{B}| \rightarrow 0$. Of course, as is well known, choosing a small learning rate from the beginning of the training is not recommended, since it may lead to overfitting and limit the scope for generalization.

C. Nonuniversal dynamics

Next we turn to nonuniversal, model-dependent features, which are contained in the drift $K_i(x_i)$ and the potential $V_i(x_i)$.

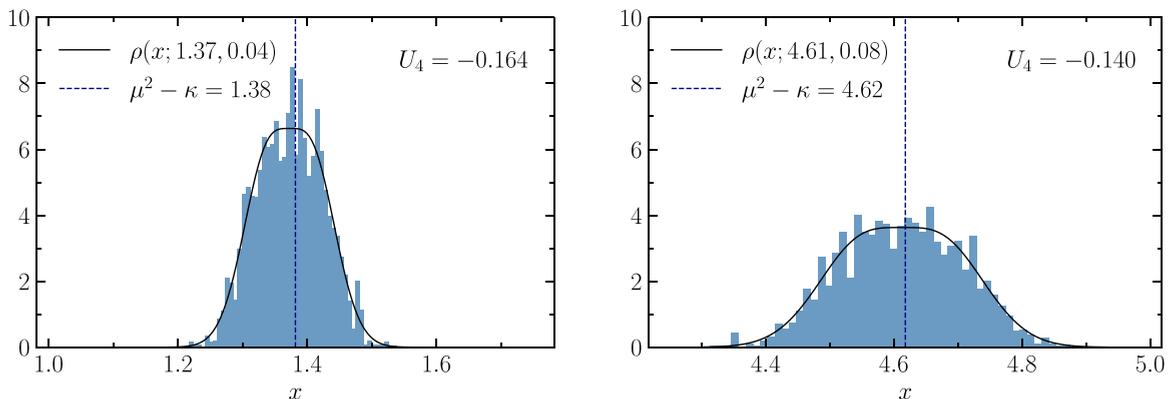


FIG. 5. Spectral density for two values of doubly degenerate target eigenvalue $\kappa = 7.62$ (left) and $\kappa = 4.38$ (right), with fit parameters x_m and σ indicated in the argument of $\rho(x; x_m, \sigma)$. Non-Gaussianity of the distribution is checked by computing the Binder cumulant U_4 .

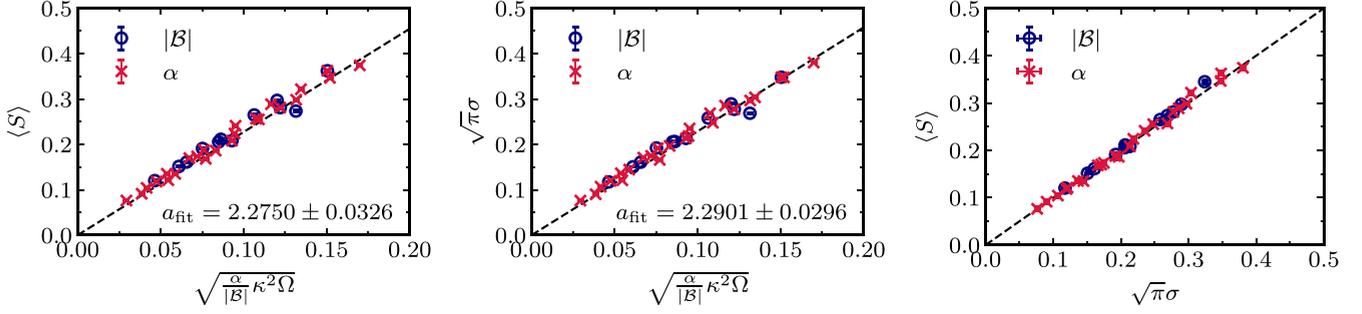


FIG. 6. Response of the mean level spacing $\langle S \rangle$ (left) and the width parameter of the spectral density $\sqrt{\pi}\sigma$ (middle) to variation of the learning rate α and the batch size $|\mathcal{B}|$, presented in the combination $\sqrt{(\alpha/|\mathcal{B}|)\kappa^2\Omega}$, for four doubly degenerate pairs, identified by target eigenvalues κ_i . Expected linear relation between $\langle S \rangle$ and $\sqrt{\pi}\sigma$ upon independent variation of α and $|\mathcal{B}|$ (right).

In Refs. [23,24] it was shown that in the instantaneous eigenbasis the singular values ξ_i of W are subject to the equation

$$\frac{d}{dt}\xi_i = \sigma_h^2 \left(\frac{1}{\kappa_i} - \frac{1}{\mu^2 - \sigma_h^2 \xi_i^2} \right) \xi_i, \quad (30)$$

where it is assumed that a continuous time limit exists. Writing this as an equation for $x_i = \sigma_h^2 \xi_i^2$ (as above) and redefining time as $\tau = 2\sigma_h^2 t$ then yields

$$\frac{d}{d\tau}x_i = \left(\frac{1}{\kappa_i} - \frac{1}{\mu^2 - x_i} \right) x_i. \quad (31)$$

However, it should be clear from the previous discussion that this equation is not relevant for the evolution of the eigenvalues x_i of X in the realistic case, as it lacks both the Coulomb repulsion and the stochasticity. Instead, we should consider [see Eq. (A2)]

$$\frac{d}{d\tau}x_i = K_i(x_i) + \sum_{j \neq i} \frac{g_i^2}{x_i - x_j} + \sqrt{2}g_i\eta_i, \quad (32)$$

with a now determined drift

$$K_i(x_i) = \left(\frac{1}{\kappa_i} - \frac{1}{\mu^2 - x_i} \right) x_i. \quad (33)$$

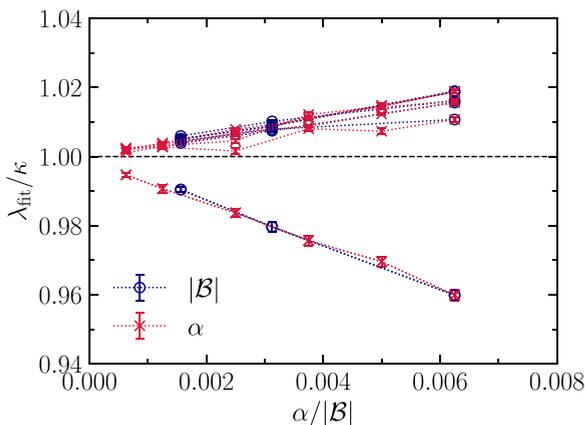


FIG. 7. Ratio of the RBM eigenvalues $\lambda_i = \mu^2 - x_{m,i}$ and the target eigenvalues κ_i as a function of $\alpha/|\mathcal{B}|$, where α and $|\mathcal{B}|$ are independently varied, demonstrating eigenvalue repulsion for nonvanishing stochasticity.

Here we stick to continuous time and it is understood that g_i captures the dependence on α and $|\mathcal{B}|$.

From now on we consider one mode only and drop the index i . This allows us to focus on the properties of the drift $K(x)$, without having to consider the Coulomb term, which will be added again later. The corresponding Fokker-Planck equation (FPE) for one mode reads

$$\partial_\tau P(x, \tau) = \partial_x (g^2 \partial_x - K(x)) P(x, \tau), \quad (34)$$

with the drift

$$K(x) = \left(\frac{1}{\kappa} - \frac{1}{\mu^2 - x} \right) x. \quad (35)$$

A stationary solution exists if the drift can be integrated to yield a potential, using $K(x) = -\partial_x V(x)$, or

$$V(x) = -\int^x dx' K(x') = -\frac{x^2}{2\kappa} - x - \mu^2 \log(\mu^2 - x), \quad (36)$$

where we recall that $0 \leq x < \mu^2$, due to positivity and stability requirements. The stationary distribution for one mode then reads

$$\begin{aligned} P_s(x) &= \frac{1}{Z} e^{-V(x)/g^2} \\ &= \frac{1}{Z} \exp \left[\frac{1}{g^2} \left(\frac{x^2}{2\kappa} + x + \mu^2 \log(\mu^2 - x) \right) \right]. \end{aligned} \quad (37)$$

Z is the normalization factor, such that

$$\int_0^{\mu^2} dx P_s(x) = 1. \quad (38)$$

Note that the distribution is peaked at $x = x_s = \mu^2 - \kappa$, where the drift vanishes, $K(x_s) = 0$.

To analyze the properties of the time-dependent FPE, we cast the dynamics as a quantum-mechanical bound state problem [25]. We factor out the square root of the stationary distribution

$$P(x, \tau) = \sqrt{P_s(x)} \psi(x, \tau), \quad (39)$$

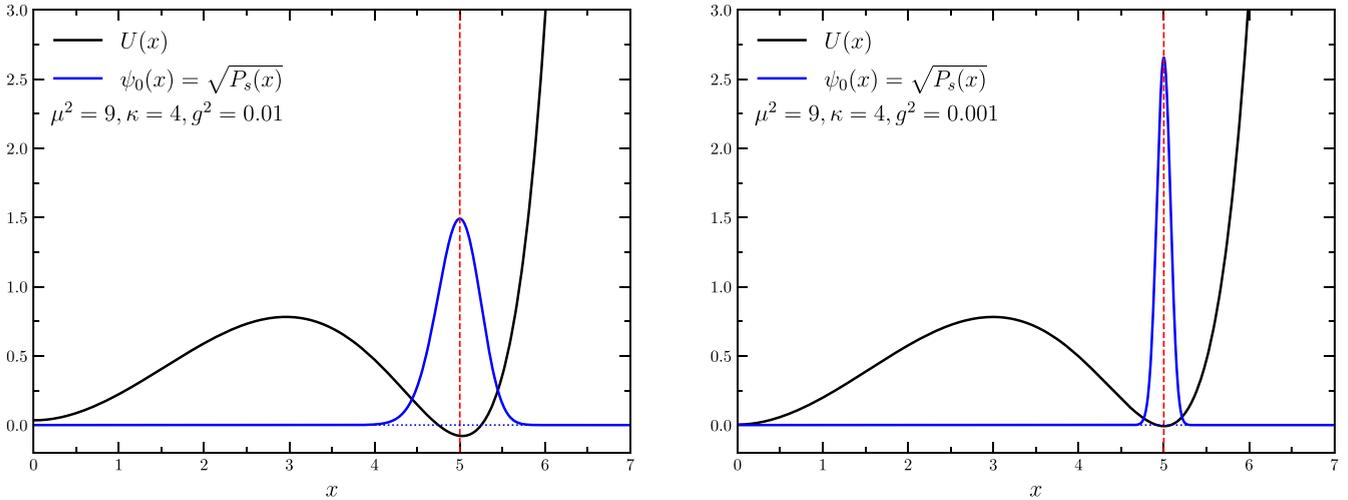


FIG. 8. Analysis of the Fokker-Planck Hamiltonian for a single mode x , with $\mu^2 = 9$, $\kappa = 4$, and $g^2 = 0.01$ (left) and 0.001 (right). Shown are the FP potential $U(x)$ (black line) and the ground-state wave function $\psi_0(x) = \sqrt{P_s(x)}$ (blue line). The vertical dotted line at $x_s = \mu^2 - \kappa = 5$ indicates the expected position of the peak.

and analyze the time dependence of $\psi(x, \tau)$. Taking the time derivative and using Eq. (34) then yields

$$\begin{aligned} \partial_\tau \psi(x, \tau) &= \left(g^2 \partial_x^2 - \frac{1}{4g^2} [\partial_x V(x)]^2 + \frac{1}{2} [\partial_x^2 V(x)] \right) \psi(x, \tau) \\ &\equiv -2H_{\text{FP}} \psi(x, \tau), \end{aligned} \quad (40)$$

where the so-called Fokker-Planck Hamiltonian can be written in a semipositive definite form,

$$H_{\text{FP}} = \frac{1}{2} L^\dagger L, \quad (41)$$

with

$$L^\dagger = -g \partial_x + \frac{1}{2g} \partial_x V(x), \quad L = +g \partial_x + \frac{1}{2g} \partial_x V(x), \quad (42)$$

The eigenvalue problem given by H_{FP} takes the form of a bound state problem, with

$$H_{\text{FP}} \psi_n(x) = E_n \psi_n(x). \quad (43)$$

The ground state, $\psi_0(x)$, with vanishing energy $E_0 = 0$, is determined by $L\psi_0(x) = 0$, which yields $\psi_0(x) = \sqrt{P_s(x)}$, as it should be, see Eqs. (37) and (39).

The solutions with positive energy determine the time-dependent evolution, or learning dynamics, and the time-dependent solution of the FPE can be expressed in terms of

$$\psi(x, \tau) = \psi_0(x) + \sum_{n>0} c_n \psi_n(x) e^{-2E_n \tau}, \quad (44)$$

where the coefficients c_n are determined by the initial distribution $P(x, 0)$, using relation (39).

To analyze the spectrum of H_{FP} in more detail, we write it as follows,

$$H_{\text{FP}} = -\frac{g^2}{2} \partial_x^2 + U(x), \quad U(x) = \frac{1}{g^2} [U_0(x) + g^2 U_1(x)], \quad (45)$$

where

$$U_0(x) = \frac{1}{8} [\partial_x V(x)]^2, \quad U_1(x) = -\frac{1}{4} \partial_x^2 V(x), \quad (46)$$

are both independent of the noise strength g^2 . This way of writing reveals the dual role of $1/g^2$: it plays the role of mass, appearing in the kinetic (∂_x^2) and potential (U_0) terms in a reciprocal manner. Moreover, given that the noise strength is small, g^2 can be treated as a small parameter, such that U_1 can be treated as a perturbative correction. We therefore first consider $U_0(x)$. It has a minimum where the drift vanishes,

$$\partial_x U_0(x)|_{x=x_s} = 0, \quad \partial_x V(x)|_{x=x_s} = -K(x_s) = 0, \quad (47)$$

i.e., at the expected position $x_s = \mu^2 - \kappa$. The curvature at the minimum is given by

$$\partial_x^2 U_0(x)|_{x=x_s} = \frac{\Omega^2}{4}, \quad \Omega \equiv \partial_x^2 V(x)|_{x=x_s} = \frac{\mu^2 - \kappa}{\kappa^2}. \quad (48)$$

Note that $U_0(x_s) = 0$ and $U_1(x_s) = -\Omega/4$. $U_0(x)$ has a local maximum at $x_{\text{max}} = \mu^2 - \sqrt{\kappa \mu^2}$, with $U_0(x_{\text{max}}) > 0$, $U_1(x_{\text{max}}) = 0$. We also note that $U(0) = \Omega \kappa / (4\mu^2) > 0$. The potential $U(x)$ is therefore a nondegenerate double well potential on the line segment $0 \leq x < \mu^2$, with a global minimum at $x = x_s$. Examples are shown in Fig. 8 for two values of g^2 .

Around $x = x_s$, the FP Hamiltonian can hence be written as

$$H_{\text{FP}} = -\frac{g^2}{2} \partial_x^2 + \frac{1}{2g^2} \left(\frac{\Omega}{2} \right)^2 (x - x_s)^2 - \frac{1}{2} \left(\frac{\Omega}{2} \right), \quad (49)$$

making complete the mapping to a harmonic oscillator with mass $1/g^2$, frequency $\Omega/2$, and shifted zero-point energy. Its eigenvalues are therefore

$$E_n = \frac{\Omega}{2} n, \quad n = 0, 1, 2, \dots, \quad (50)$$

and convergence to the stationary limit is determined by the smallest nonzero eigenvalue, $2E_1 = \Omega$, as observed before [24]. Also shown in Fig. 8 is the groundstate wave function, $\psi_0(x) = \sqrt{P_s(x)}$. We have chosen a rather large value of g^2 here, such that it has a substantial width. Indeed, in the harmonic approximation, g^2 does not appear in the spectrum but controls the width of the wave function, and hence the

distribution $P(x, t)$. Less stochasticity results in a narrower distribution centered around $x = x_s$, as one would expect.

The overall dynamics combines the evolution in the potential $V_i(x_i)$ for each eigenvalue with the Coulomb repulsion between the eigenvalues, leading to the spectral densities shown in Figs. 3 and 5.

D. Teacher-student model

Here we implement the findings in a simple teacher-student model, with an emphasis on how the stochasticity can be modelled, when added by hand. This is of interest, since as we have argued above, the noise does not appear in a standard manner (it is multiplied by the learning rate and not its square root, for instance).

We consider a simple teacher-student model, formulated as

$$\mathbf{y}_t = V\mathbf{x}, \quad \mathbf{y}_s = W\mathbf{x}, \quad (51)$$

where V, W are $N \times N$ symmetric matrices and $\mathbf{x} \sim \mathcal{N}(0, 1)$ are N -dimensional Gaussian random vectors. The teacher matrix V is fixed, while the student matrix W is to be learned by minimizing the loss function,

$$\mathcal{L}(W) = \frac{1}{2} \langle \|\mathbf{y}_t - \mathbf{y}_s\|^2 \rangle, \quad (52)$$

where the brackets denote averaging over the input data. In the limit of an infinite amount of input data, $\langle x_i x_j \rangle = \delta_{ij}$, and the dynamics is deterministic. Some stochasticity can be introduced by considering batch updates. The amount of stochasticity is, however, not sufficient for the purpose of this discussion, since it is unavoidably suppressed when $W \rightarrow V$ as the training converges. We therefore introduce additional randomness by hand, in a manner that is consistent with the discussion in Sec. II, see in particular Eq. (5). We consider

$$W_{ij} \rightarrow W'_{ij} = W_{ij} - \alpha \frac{\delta \mathcal{L}}{\delta W_{ij}} + \alpha \sqrt{A_{ij}} \eta_{ij}^{\mathcal{B}}. \quad (53)$$

Here α is the learning rate, which multiplies both the gradient and the stochastic term, A_{ij} indicates the strength of the noise, and $\eta_{ij}^{\mathcal{B}}$ is the stochastic term, averaged over a minibatch of size $|\mathcal{B}|$. The latter is introduced as follows: we consider individual updates with noise $\eta_{ij}^b \sim \mathcal{N}(0, 1)$. A minibatch average is then given by

$$\eta_{ij}^{\mathcal{B}} = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \eta_{ij}^b, \quad (54)$$

such that

$$\langle \eta_{ij}^{\mathcal{B}} \eta_{kl}^{\mathcal{B}} \rangle = \frac{1}{|\mathcal{B}|^2} \sum_{b, b'} \langle \eta_{ij}^b \eta_{kl}^{b'} \rangle = \frac{1}{|\mathcal{B}|} \delta_{ik} \delta_{jl}. \quad (55)$$

We finally write $\eta_{ij}^{\mathcal{B}} = \eta_{ij} / \sqrt{|\mathcal{B}|}$, with $\eta_{ij} \sim \mathcal{N}(0, 1)$, to separate the factor $1/\sqrt{|\mathcal{B}|}$. The noise is explicitly symmetrised, such that throughout the evolution W is symmetric. The final update equation is then

$$W_{ij} \rightarrow W'_{ij} = W_{ij} - \alpha \frac{\delta \mathcal{L}}{\delta W_{ij}} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{A_{ij}} \eta_{ij}, \quad (56)$$

with $\eta_{ij} \sim \mathcal{N}(0, 1)$, which, as stated, is modeled explicitly to be consistent with the arguments given in Sec. II.

We have applied this model to the case of $N = 2$, with the teacher matrix diagonal with eigenvalues $\kappa_{1,2} = \kappa \pm \delta\kappa/2$, and $A_{ij} = A(1 + \delta_{ij})$ (see Appendix D). The eigenvalues of W then satisfy

$$\lambda_1 \rightarrow \lambda'_1 = \lambda_1 - \alpha(\lambda_1 - \kappa_1) + \frac{\alpha^2}{|\mathcal{B}|} \frac{A}{\lambda_1 - \lambda_2} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{2A} \eta_1, \quad (57)$$

$$\lambda_2 \rightarrow \lambda'_2 = \lambda_2 - \alpha(\lambda_2 - \kappa_2) + \frac{\alpha^2}{|\mathcal{B}|} \frac{A}{\lambda_2 - \lambda_1} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{2A} \eta_2. \quad (58)$$

We have solved Eq. (56) numerically and varied the learning rate α , minibatch size $|\mathcal{B}|$ and noise strength A independently. We have constructed the distribution for the level spacing and for the spectral density, and confirmed that in all cases these distributions can be fitted by the Wigner surmise and the Wigner semicircle (for $N = 2$).

We denote the eigenvalue splitting as $S = |\lambda_1 - \lambda_2|$. Taking the difference of the equations above and averaging over the noise, then yields in the stationary limit

$$\langle S \rangle - \delta\kappa - \frac{2\alpha A}{|\mathcal{B}|} \left\langle \frac{1}{S} \right\rangle = 0. \quad (59)$$

Since for the Wigner surmise $\langle 1/S \rangle = \pi/(2\langle S \rangle)$, we finally find that the mean level spacing is given by

$$\langle S \rangle = \frac{1}{2} \left(\delta\kappa + \sqrt{\delta\kappa^2 + \frac{4\pi\alpha A}{|\mathcal{B}|}} \right). \quad (60)$$

In the limit of vanishing stochasticity, this becomes $\langle S \rangle = \delta\kappa$, as it should be (perfect learning), while for the degenerate case this reduces to

$$\langle S \rangle = \sqrt{\frac{\pi\alpha A}{|\mathcal{B}|}}, \quad (61)$$

demonstrating again the linear scaling rule and setting a limit on the precision with which the teacher matrix can be learnt, for finite learning rate, noise strength, and batch size. In Fig. 9 we show the mean level spacing as a function of $\sqrt{\alpha A/|\mathcal{B}|}$, where, as already mentioned, α , $|\mathcal{B}|$, and A are varied independently, and $\delta\kappa = 0$. The fitted slope, $a_{\text{fit}} \approx 1.87$, is close to the analytical prediction $\sqrt{\pi} \approx 1.77$, where the latter is valid when finite learning rate corrections are disregarded (see Appendix C).

IV. SUMMARY AND OUTLOOK

We have considered stochastic weight matrix dynamics in generic learning algorithms and argued that it can be described in the framework of Dyson Brownian motion. This results in universal features dictated by random matrix theory, including eigenvalue repulsion, quantified by the Wigner surmise. The level of stochasticity is determined by the ratio of the learning rate α and the minibatch size $|\mathcal{B}|$ times a nonuniversal function, which encodes details of the loss function and the architecture. The linear scaling rule, i.e., the dependence on

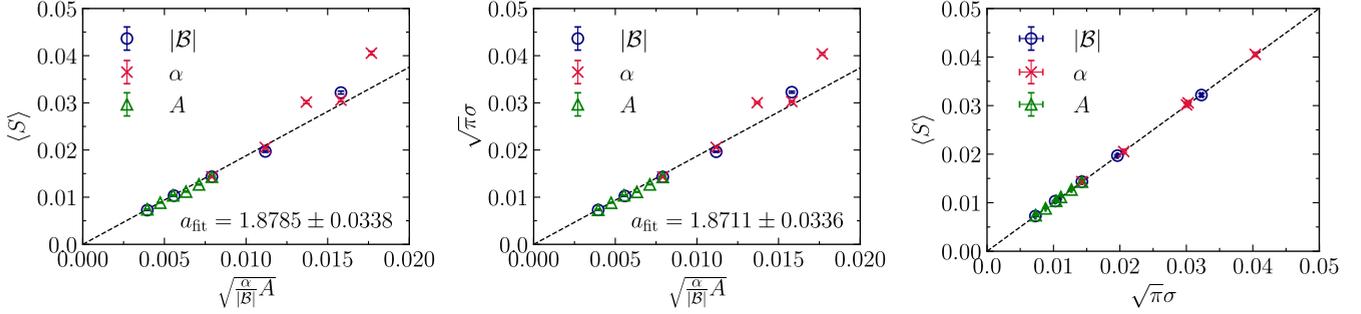


FIG. 9. Response of the eigenvalue splitting $\langle S \rangle = \langle \delta\lambda \rangle$ (left) and the width parameter of the spectral density $\sqrt{\pi}\sigma$ (middle) to independent variations of the learning rate α , the batch size $|\mathcal{B}|$, and the level of stochasticity A , presented in the combination $\sqrt{\alpha A/|\mathcal{B}|}$, in the teacher-student model. Expected linear relation between $\langle S \rangle$ and $\sqrt{\pi}\sigma$ upon independent variation of α , $|\mathcal{B}|$, and A (right).

$\alpha/|\mathcal{B}|$, arises naturally within this framework. One implication is that there is no straightforward limit in which, say, stochastic gradient descent reduces to a stochastic differential equation in continuous time, since the level of stochasticity decreases with decreasing learning rate or increasing mini-batch size. While the inherent stochasticity for finite $\alpha/|\mathcal{B}|$ sets a fundamental limit on the accuracy of learning, it also prevents overfitting and hence helps in decreasing the generalization error.

We have tested our universal and nonuniversal predictions in two simple models: a linear teacher-student model and the Gaussian restricted Boltzmann machine. In the latter the nonuniversal aspects of the dynamics of the eigenvalues can be nicely described as a quantum-mechanical bound state problem in a nondegenerate double well potential, defined on a finite interval. The universal dependence on $\alpha/|\mathcal{B}|$ was confirmed in both models.

Our results are rather general and can be applied to any model in which weight matrices are updated with a stochastic optimization algorithm. The choice of architecture and loss function determine the nonuniversal aspects in the Coulomb potential, as well as the level of stochasticity. In some cases the resulting stochasticity is so small that is hardly noticeable; this was the case in the linear teacher-student model and for that reason we added noise by hand. The Gaussian RBM is stochastic by itself, due to the need to sample, and hence the noise inherent in the model is strong enough to observe the effects without any further addition. Moving forward, it is important to test these ideas in more complicated architectures, including neural networks and transformers. To be able to observe universal features, it may be required to use spectral unfolding, see, e.g., Refs. [19,26,27]. The coupling between different weight matrices across layers as well as the use of adaptive learning rates, such as adam [28], may also influence the manner in which universal features appear. Work along these directions is currently in progress.

ACKNOWLEDGMENTS

The content of this paper was first presented at the ECT* workshop *Machine Learning and the Renormalisation Group* in May 2024. We thank the participants of this meeting, in particular O. Hajizadeh and S. Rezchikov, for discussion, and ECT* for support. We thank S. P. Kumar and N. Talwar for discussion on matrix models. G.A. and B.L. are supported

by STFC Consolidated Grant No. ST/X000648/1. B.L. is further supported by the UKRI EPSRC ExCALIBUR ExaTEPP Project No. EP/X017168/1. C.P. is supported by the UKRI AIMLAC CDT EP/S023992/1.

DATA AVAILABILITY

The code and data used for this paper are available from Ref. [29].

APPENDIX A: DYSON BROWNIAN MOTION AND THE STOCHASTIC COULOMB GAS

In this Appendix we summarize the main concepts of Dyson Brownian motion [2], see in particular the textbook [8] (Ch. 7). For simplicity we use continuous-time notation, but when in doubt, expressions should be understood in the sense of Itô calculus.

Consider a symmetric $N \times N$ matrix X , whose matrix elements are updated according to

$$\frac{dX_{ij}}{dt} = K_{ij}(X) + \sqrt{A_{ij}}\eta_{ij}, \quad (\text{A1})$$

where K_{ij} is the drift term, A_{ij} encodes the strength of the noise, and $\eta_{ij} \sim \mathcal{N}(0, 1)$, all of which are symmetric. Its eigenvalues x_i then evolve according to

$$\begin{aligned} \frac{dx_i}{dt} &= K_i(x_i) + \sum_{j \neq i} \frac{g_i^2}{x_i - x_j} + \sqrt{2}g_i\eta_i \\ &\equiv K_i^{(\text{eff})}(x_i) + \sqrt{2}g_i\eta_i, \end{aligned} \quad (\text{A2})$$

where $K_i(x_i)$ is the drift acting on the eigenvalues, the second term is the induced Coulomb term, and $\eta_i \sim \mathcal{N}(0, 1)$. This result can be derived using second-order perturbation theory with a discrete time step δt , making the standard assumption that the drift and the noise correlator $\langle \eta_{ij}\eta_{kl} \rangle$ scale as δt [8]. The noise strength satisfies $A_{ij} = A(1 + \delta_{ij})$ and we denote the diagonal element as $\sqrt{A_{ii}} = \sqrt{2}g_i$. Strictly speaking the index on g_i should be dropped, but we keep it for future convenience. A derivation in the context of the teacher-student model is given in Appendix D.

The corresponding Fokker-Planck equation (FPE) for the distribution $P(\{x_i\}, t)$ reads

$$\partial_t P(\{x_i\}, t) = \sum_{i=1}^N \partial_{x_i} [(g_i^2 \partial_{x_i} - K_i^{(\text{eff})}(x_i))] P(\{x_i\}, t). \quad (\text{A3})$$

Provided that the drift can be derived from a potential,

$$K_i(x_i) = -\frac{dV_i(x_i)}{dx_i}, \quad (\text{A4})$$

the FPE has the stationary solution

$$P_s(\{x_i\}) = \frac{1}{Z} \prod_{i < j} |x_i - x_j| e^{-\sum_i V_i(x_i)/g_i^2}, \quad (\text{A5})$$

with

$$Z = \int dx_1 \dots dx_N P_s(\{x_i\}). \quad (\text{A6})$$

This distribution is known as the Coulomb gas. The Coulomb repulsion term can also be seen as a Jacobian arising when making the transition from integrating over all matrix elements to just the eigenvalues. In the case that the potentials $V_i(x_i)/g_i^2$ are identical quadratics, the Coulomb gas describes the eigenvalues of the Gaussian orthogonal ensemble.

The presentation above assumes that the potential in the Coulomb gas is separable, see Eqs. (A4) and (A5). If this is not possible, e.g., due to eigenvalues mixing nonlinearly in the drift terms, a more general potential appears, as

$$\sum_i V_i(x_i)/g_i^2 \rightarrow \bar{V}(x_1, x_2, \dots, x_N), \quad (\text{A7})$$

with

$$K_i = -g_i^2 \frac{\partial \bar{V}(x_1, x_2, \dots, x_N)}{\partial x_i}. \quad (\text{A8})$$

While this will complicate the analysis and lead to a richer loss function landscape, the implications are expected to remain.

We specialize to the case with $N = 2$ eigenvalues and assume that the dynamics can be linearized around degenerate minima $x_1^s = x_2^s = \kappa$, such that the drift K_i is proportional to $(\kappa - x_i)$. The partition function is then

$$Z = \frac{1}{N_0} \int dx_1 dx_2 |x_1 - x_2| e^{-V(x_1, x_2)}, \quad (\text{A9})$$

with

$$V(x_1, x_2) = \frac{1}{2\sigma^2} [(x_1 - \kappa)^2 + (x_2 - \kappa)^2]. \quad (\text{A10})$$

The variance σ^2 is chosen to be identical for the two degenerate modes. We assume that the distribution is sufficiently peaked around κ , such that the integral boundaries can be taken as $\pm\infty$. We hence consider

$$Z = \frac{1}{N_0} \int dx_1 dx_2 |x_1 - x_2| e^{-(x_1^2 + x_2^2)/(2\sigma^2)}, \quad (\text{A11})$$

with the normalization constant $N_0 = 4\sqrt{\pi}\sigma^3$.

The Wigner surmise signifies the level spacing $S = x_1 - x_2$. Let us change variables to $x_{1,2} = x \pm S/2$, such that

$$Z = \int_0^\infty dS P(S), \quad (\text{A12})$$

with

$$P(S) = \frac{S}{2\sigma^2} e^{-S^2/(4\sigma^2)}. \quad (\text{A13})$$

The mean level spacing is

$$\langle S \rangle = \int_0^\infty dS S P(S) = \sqrt{\pi}\sigma. \quad (\text{A14})$$

In terms of $s = S/\langle S \rangle$ the surmise is parameter free,

$$P(s) = \frac{\pi}{2} s e^{-\pi s^2/4}. \quad (\text{A15})$$

The spectral density is defined, for arbitrary N , as

$$\rho(x) = \left\langle \frac{1}{N} \sum_{i=1}^N \delta(x - x_i) \right\rangle. \quad (\text{A16})$$

For $N = 2$, it is easily evaluated as

$$\rho(x) = \frac{e^{-x^2/(2\sigma^2)}}{4\sqrt{\pi}\sigma} \left[2e^{-x^2/(2\sigma^2)} + \sqrt{2\pi} \frac{x}{\sigma} \text{Erf} \left(\frac{x}{\sqrt{2}\sigma} \right) \right]. \quad (\text{A17})$$

It is flatter and broader than a simple Gaussian.

APPENDIX B: WEIGHT MATRIX INITIALIZATION

Consider the $M \times N$ random matrix W at initialization, with elements $W \sim \mathcal{N}(0, \sigma_{\text{in}}^2)$. Without loss of generality we take $N \leq M$; if this is not the case, simply exchange W and W^T . The symmetric $N \times N$ matrix $X = W^T W$ has N nonzero eigenvalues $x_i = \xi_i^2$ ($i = 1, \dots, N$), where ξ_i are the singular values of W . The distribution of eigenvalues x of X is given by the Marchenko-Pastur distribution,

$$P_{\text{MP}}(x) = \frac{1}{2\pi\sigma_{\text{in}}^2 M r x} \sqrt{(x_+ - x)(x - x_-)}, \quad (\text{B1})$$

where $r = N/M \leq 1$ and $x_{\pm} = M\sigma_{\text{in}}^2(1 \pm \sqrt{r})^2$, $x_- < x < x_+$. It is desirable for the spectrum of X to only depend on r and not separately on M and N , such that one can take $M, N \rightarrow \infty$ at fixed r . There is some freedom to select σ_{in}^2 , e.g., $\sigma_{\text{in}}^2 = 1/M, 1/N$, or $1/\sqrt{MN}$. By choosing $\sigma_{\text{in}}^2 = 1/M$, the spectrum of X is bounded for all values of $0 < r \leq 1$, which can be advantageous, in particular in the RBM case. The initial distribution then reads

$$P_{\text{MP}}(x) = \frac{1}{2\pi r x} \sqrt{(x_+ - x)(x - x_-)}, \quad (\text{B2})$$

with $x_{\pm} = (1 \pm \sqrt{r})^2$, $0 \leq x_- \leq x \leq x_+ \leq 4$.

APPENDIX C: NONUNIVERSAL PART OF THE LINEAR SCALING RELATION

In Sec. II, we demonstrated that the Coulomb gas potential $V_i(x_i)/g_i^2$ is the product of a universal and a nonuniversal factor, see Eqs. (16) and (18). Here we consider the nonuniversal factor in the case of the restricted Boltzmann machine and sketch a derivation of its parameter dependence.

We are interested in the variance of the gradient of the loss function, $\text{Var}(\delta W)$ or $\text{Var}(\delta X)$. We write the gradient in the

form, see Eq. (24) and putting $\sigma_h^2 = 1$,

$$\frac{\delta \mathcal{L}}{\delta W} W^T = (C^{\text{target}} - C^{\text{RBM}})(W W^T). \quad (\text{C1})$$

The two-point functions on the RHS indicate the target data,

$$C^{\text{target}} = K_{\text{target}}^{-1}, \quad (\text{C2})$$

and the RBM model,

$$C^{\text{RBM}} = \phi^{\text{RBM}} \phi^{\text{RBM}}, \quad (\text{C3})$$

with expectation value

$$\langle C^{\text{RBM}} \rangle = K^{-1}, \quad K = \mu^2 \mathbb{1} - W W^T. \quad (\text{C4})$$

Finally, using the singular-value decomposition $W = U \Xi V^T$,

$$W W^T = U \Xi \Xi^T U^T, \quad (\text{C5})$$

where $\Xi \Xi^T$ contains the squares of the singular values $x_i = \xi_i^2$ on the diagonal. All quantities above are matrix-valued with indices on the visible layer, which we suppress.

To determine the variance, it is important to note that only C^{RBM} is fluctuating during the mini-batch and contrastive divergence updates. The variance is therefore

$$\text{Var} \left(\frac{\delta \mathcal{L}}{\delta W} W^T \right) = (\langle C^{\text{RBM}} C^{\text{RBM}} \rangle - \langle C^{\text{RBM}} \rangle^2) (W W^T) (W W^T). \quad (\text{C6})$$

In the Gaussian RBM only the second moment is nontrivial, see Eq. (C4). Hence

$$\langle C^{\text{RBM}} C^{\text{RBM}} \rangle - \langle C^{\text{RBM}} \rangle^2 = 2K^{-1} K^{-1}. \quad (\text{C7})$$

Going to the eigenbasis for one mode, in which

$$K^{-1} \rightarrow \frac{1}{\mu^2 - x}, \quad W W^T \rightarrow x, \quad (\text{C8})$$

then yields

$$\text{Var} \left(\frac{\delta \mathcal{L}}{\delta W} W^T \right) \sim \frac{2x^2}{(\mu^2 - x)^2}. \quad (\text{C9})$$

After training, $x = \mu^2 - \kappa + \delta x$. At leading order in δx we find therefore

$$\text{Var} \left(\frac{\delta \mathcal{L}}{\delta W} W^T \right) \sim \frac{2(\mu^2 - \kappa)^2}{\kappa^2}. \quad (\text{C10})$$

In the main text we noted that the curvature of the Coulomb gas potential around its minimum is given by

$\Omega_i = (\mu^2 - \kappa_i)/\kappa_i^2$. We therefore write the nonuniversal function representing the variance as

$$\tilde{g}_i^2 \sim \kappa_i^2 \Omega_i^2, \quad (\text{C11})$$

or, see Eq. (18),

$$\sigma_i^2 = \frac{\alpha}{|\mathcal{B}|} \kappa_i^2 \Omega_i. \quad (\text{C12})$$

By analyzing the scaling of the mean level splitting $\langle S \rangle$ with $\alpha/|\mathcal{B}|$ for different κ_i values, we have confirmed this parametric dependence on μ^2 and κ .

As a side note, we also include here the (well-known) statement that stochastic equations suffer from finite-discretization effects, which should be considered when comparing with analytically derived expressions. Consider the simple update for one degree of freedom, which is modeled according to the equations considered in the main paper,

$$x_{n+1} = x_n - \alpha \omega x_n + \sqrt{2\alpha\gamma} \eta_n, \quad (\text{C13})$$

where $\gamma = \alpha A/|\mathcal{B}|$ and $\langle \eta_n \eta_m \rangle = \delta_{nm}$. It is solved by (taking $x_0 = 0$)

$$x_n = \sqrt{2\alpha\gamma} \sum_{i=0}^{n-1} (1 - \alpha\omega)^{n-1-i} \eta_i. \quad (\text{C14})$$

Assuming that $\alpha\omega < 1$ and taking n large, one finds

$$\lim_{n \rightarrow \infty} \langle x_n^2 \rangle = \frac{\alpha A}{|\mathcal{B}|} \frac{1}{\omega} \frac{1}{1 - \alpha\omega/2}. \quad (\text{C15})$$

This illustrates that the variance depends on $\alpha A/|\mathcal{B}|$, as we have emphasised in the main text, but also that there are finite learning rate corrections present, as expected.

APPENDIX D: DERIVATION OF THE EIGENVALUE EQUATION

In this Appendix we add a brief derivation of the eigenvalue equation including the Coulomb term, starting with the matrix update, in the context of the teacher-student model. This follows closely textbook derivations [8], but may be useful to those not familiar with this topic.

We consider a 2×2 symmetric matrix W , updated according to

$$W_{ij} \rightarrow W'_{ij} = W_{ij} - \alpha \frac{\delta \mathcal{L}}{\delta W_{ij}} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{A_{ij}} \eta_{ij}, \quad (\text{D1})$$

where α is the learning rate and $|\mathcal{B}|$ the batch size. Both the noise $\eta_{ij} \sim \mathcal{N}(0, 1)$ and strength A_{ij} are symmetric matrices. The eigenvalues of W are denoted as $\lambda_{1,2}$ and the teacher values are $\kappa_{1,2}$.

We assume that at the current step W is diagonal (or has been diagonalized) and consider one update. The drift is then $-\alpha(\lambda_{1,2} - \kappa_{1,2})$ and the updated matrix reads

$$W' = \begin{pmatrix} \lambda_1 - \alpha(\lambda_1 - \kappa_1) + (\alpha/\sqrt{|\mathcal{B}|})\sqrt{A_{11}}\eta_{11} & (\alpha/\sqrt{|\mathcal{B}|})\sqrt{A_{12}}\eta_{12} \\ (\alpha/\sqrt{|\mathcal{B}|})\sqrt{A_{12}}\eta_{12} & \lambda_2 - \alpha(\lambda_2 - \kappa_2) + (\alpha/\sqrt{|\mathcal{B}|})\sqrt{A_{22}}\eta_{22} \end{pmatrix}. \quad (\text{D2})$$

To compute the eigenvalues of W' , we follow the standard power counting rules and treat the drift as $\mathcal{O}(\epsilon)$ and the noise terms as $\mathcal{O}(\sqrt{\epsilon})$. Expanding to $\mathcal{O}(\epsilon)$ then yields the eigenvalues of W' ,

$$\lambda'_1 = \lambda_1 - \alpha(\lambda_1 - \kappa_1) + \frac{\alpha^2}{|\mathcal{B}|} \frac{A_{12}\eta_{12}^2}{\lambda_1 - \lambda_2} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{A_{11}}\eta_{11}, \quad (\text{D3})$$

$$\lambda'_2 = \lambda_2 - \alpha(\lambda_2 - \kappa_2) + \frac{\alpha^2}{|\mathcal{B}|} \frac{A_{12}\eta_{12}^2}{\lambda_2 - \lambda_1} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{A_{22}}\eta_{22}. \quad (\text{D4})$$

Taking a noise average in the Coulomb term, i.e., replacing η_{12}^2 by 1, denoting $\eta_{11} = \eta_1$, $\eta_{22} = \eta_2$, and using that $A_{ij} = A(1 + \delta_{ij})$ (or $A_{11} = A_{22} = 2A$, $A_{12} = A_{21} = A$) [8], then yields

$$\lambda'_1 = \lambda_1 - \alpha(\lambda_1 - \kappa_1) + \frac{\alpha^2}{|\mathcal{B}|} \frac{A}{\lambda_1 - \lambda_2} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{2A}\eta_1, \quad (\text{D5})$$

$$\lambda'_2 = \lambda_2 - \alpha(\lambda_2 - \kappa_2) + \frac{\alpha^2}{|\mathcal{B}|} \frac{A}{\lambda_2 - \lambda_1} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{2A}\eta_2. \quad (\text{D6})$$

This is the standard equation for the eigenvalues, including the Coulomb and stochastic terms. As mentioned, a detailed discussion of the various steps and an extension to $N \times N$ matrices can be found in the textbook [8].

-
- [1] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborova, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [2] F. J. Dyson, A brownian-motion model for the eigenvalues of a random matrix, *J. Math. Phys.* **3**, 1191 (1962).
- [3] E. P. Wigner, Characteristic vectors of bordered matrices with infinite dimensions, *Ann. Math.* **62**, 548 (1955).
- [4] E. P. Wigner, *Conference on Neutron Physics by Time-of-Flight* (Oak Ridge National Laboratory, Oak Ridge, 1956), p. 67.
- [5] F. J. Dyson, Statistical theory of the energy levels of complex systems. I, *J. Math. Phys.* **3**, 140 (1962).
- [6] F. J. Dyson, Statistical theory of the energy levels of complex systems. II, *J. Math. Phys.* **3**, 157 (1962).
- [7] F. J. Dyson, Statistical theory of the energy levels of complex systems. III, *J. Math. Phys.* **3**, 166 (1962).
- [8] M. L. Mehta, *Random Matrices*, 3rd ed. (Academic Press, New York, 2004).
- [9] P. Goyal, P. Dollár, R. B. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola *et al.*, Accurate, large minibatch SGD: Training ImageNet in 1 hour, [arXiv:1706.02677](https://arxiv.org/abs/1706.02677).
- [10] S. L. Smith and Q. V. Le, A Bayesian perspective on generalization and stochastic gradient descent, *International Conference on Learning Representations* (ICLR, 2018).
- [11] S. L. Smith, P. Kindermans, and Q. V. Le, Don't decay the learning rate, increase the batch size, *International Conference on Learning Representations* (ICLR, 2018).
- [12] S. L. Smith, D. Duckworth, Q. V. Le, and J. Sohl-Dickstein, Stochastic natural gradient descent draws posterior samples in function space, [arXiv:1806.09597](https://arxiv.org/abs/1806.09597).
- [13] P. Chaudhari and S. Soatto, Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks, in *Proceedings of the 2018 Information Theory and Applications Workshop (ITA)* (IEEE, San Diego, CA, 2018), pp. 1–10.
- [14] S. Mandt, M. D. Hoffman, and D. M. Blei, Continuous-time limit of stochastic gradient descent revisited, in *Proceedings of the 8th NIPS Workshop on Optimization for Machine Learning* (NIPS, 2015).
- [15] Q. Li, C. Tai, and W. E, Stochastic modified equations and adaptive stochastic gradient algorithms, in *Proceedings of the 34th International Conference on Machine Learning (ICML, 2017)*, Vol. 70, pp. 2101–2110.
- [16] S. Yaida, Fluctuation-dissipation relations for stochastic gradient descent, in *International Conference on Learning Representations* (ICML, 2019).
- [17] C. H. Martin and M. W. Mahoney, Traditional and heavy-tailed self regularization in neural network models, [arXiv:1901.08276](https://arxiv.org/abs/1901.08276).
- [18] N. P. Baskerville, D. Granzio, and J. P. Keating, Appearance of random Matrix theory in deep learning, *Phys. A: Stat. Mech. Appl.* **590**, 126742 (2022).
- [19] N. Levi and Y. Oz, The underlying scaling laws and universal statistical structure of complex datasets, [arXiv:2306.14975](https://arxiv.org/abs/2306.14975).
- [20] A. Sclocchi and M. Wyart, On the different regimes of stochastic gradient descent, *Proc. Natl. Acad. Sci. USA* **121**, e2316301121 (2024).
- [21] P. Smolensky, Chapter 6: Information processing in dynamical systems: Foundations of harmony theory, in *Parallel Distributed Processing: Volume 1*, edited by D. Rumelhart and J. McClelland (MIT Press, Cambridge, 1986), pp. 194–281.
- [22] G. E. Hinton, Training products of experts by minimizing contrastive divergence, *Neural Comput.* **14**, 1771 (2002).
- [23] A. Decelle and C. Furtlehner, Restricted Boltzmann machine: Recent advances and mean-field theory, *Chin. Phys. B* **30**, 040202 (2021).
- [24] G. Aarts, B. Lucini, and C. Park, Scalar field restricted Boltzmann machine as an ultraviolet regulator, *Phys. Rev. D* **109**, 034521 (2024).
- [25] P. H. Damgaard and H. Hüffel, Stochastic quantization, *Phys. Rep.* **152**, 227 (1987).
- [26] J. J. M. Verbaarschot and T. Wettig, Random matrix theory and chiral symmetry in QCD, *Annu. Rev. Nucl. Part. Sci.* **50**, 343 (2000).
- [27] A. A. Abul-Magd and A. Y. Abul-Magd, Unfolding of the spectrum for chaotic and mixed systems, *Physica A* **396**, 185 (2014).
- [28] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [29] G. Aarts, B. Lucini, and C. Park, `chanjure/Stochastic_weight_matrix_dynamics_during_learning_and_Dyson_Brownian_motion-data_release: v1.0.1` (2024), <https://doi.org/10.5281/zenodo.13294081>.