
Dyson Brownian motion and random matrix dynamics of weight matrices during learning

Gert Aarts

Department of Physics, Swansea University, Swansea SA2 8PP, United Kingdom
g.aarts@swansea.ac.uk

Ouraman Hajizadeh

Graz, Austria
ou86hajizadeh@gmail.com

Biagio Lucini

Department of Mathematics, Swansea University, Swansea, SA1 8EN, United Kingdom
b.lucini@swansea.ac.uk

Chanju Park

Department of Physics, Swansea University, Swansea SA2 8PP, United Kingdom
chanju.b.park@gmail.com

November 20, 2024

Abstract

During training, weight matrices in machine learning architectures are updated using stochastic gradient descent or variations thereof. In this contribution we employ concepts of random matrix theory to analyse the resulting stochastic matrix dynamics. We first demonstrate that the dynamics can generically be described using Dyson Brownian motion, leading to e.g. eigenvalue repulsion. The level of stochasticity is shown to depend on the ratio of the learning rate and the mini-batch size, explaining the empirically observed linear scaling rule. We verify this linear scaling in the restricted Boltzmann machine. Subsequently we study weight matrix dynamics in transformers (a nano-GPT), following the evolution from a Marchenko-Pastur distribution for eigenvalues at initialisation to a combination with additional structure at the end of learning.

1 Introduction

The dynamics of weight matrices W during training is a topic of considerable interest, as it may shed light on both the capacity and efficiency of machine learning (ML) architectures to learn and generalise. Since most algorithms involve a level of stochasticity, e.g. due to the use of mini-batches and sampling, the appropriate manner to analyse this dynamics is random matrix theory (RMT) [1–6]. In the presence of time evolution, as is the case for learning, one can then employ the framework of Dyson Brownian motion [7] to express the dynamics of $X = W^T W$ in terms of a stochastic evolution equation for the eigenvalues of X [8]. This equation includes an induced Coulomb term, resulting in e.g. eigenvalue repulsion and the Wigner surmise.

Here we address the question how Dyson Brownian motion can be applied to stochastic weight matrix updates, with observable consequences. Reducing the learning dynamics from the space of

matrices to the space of eigenvalues naturally yields the picture in which the eigenvalues evolve from an initial Marchenko-Pastur distribution to one determined by a so-called stochastic Coulomb potential. This Coulomb potential contains both universal and non-universal aspects, which depend on details of the minimisation landscape due to the loss function. Hence a study of the evolution of eigenvalues provides insight into the latter. To demonstrate our analytical findings, we present results in the Gaussian restricted Boltzmann machine (RBM), which can be understood completely, and for transformers (a nano-GPT), where results are currently at an empirical level.

Related work RMT dates back to the analysis of nuclear spectra in the 1960s [1–6], but has been proven to be widely applicable across disciplines. The notion that RMT is useful for the description of weight matrices in ML has been observed previously, see e.g. Refs. [9–12] and the textbook [13]. As far as we know the connection between Dyson Brownian motion and the stochastic Coulomb gas on the one hand and the eigenvalue dynamics of weight matrices on the other hand was first pointed in Ref. [8]. RMT has also been used to study properties of data [14] instead of weight matrices.

2 Dyson Brownian motion and stochastic matrix dynamics

Stochastic gradient descent, or variations thereof, are widely used to minimise loss functions, $\mathcal{L}[W]$, via the update $W \rightarrow W' = W + \delta W = W - \alpha \delta \mathcal{L} / \delta W$, where α is the learning rate. We follow here closely our previous work [8]. Using stochastically chosen mini-batches \mathcal{B} of size $|\mathcal{B}|$ and invoking the central limit theorem, yields the stochastic update for weight matrices W

$$W'_{ij} = W_{ij} - \alpha \left(\frac{\delta \mathcal{L}}{\delta W_{ij}} \right)_{\mathcal{B}} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{\text{Var} \left(\frac{\delta \mathcal{L}}{\delta W_{ij}} \right)} \eta_{ij}, \quad (2.1)$$

where $\eta_{ij} \sim \mathcal{N}(0, 1)$, the subscript \mathcal{B} denotes the true mean of the batch gradient, and there is no summation over repeated indices. To apply RMT, it is beneficial to consider the symmetric combination $X = W^T W$, which has real and semi-positive eigenvalues (the squares of W 's singular values). The update of X can be written as

$$X_{ij} \rightarrow X'_{ij} = X_{ij} + \delta X_{ij}^{\mathcal{B}} + \frac{1}{\sqrt{|\mathcal{B}|}} \sqrt{\text{Var}(\delta X_{ij})} \eta_{ij}, \quad (2.2)$$

where the various terms follow from Eq. (2.1) and the noise η_{ij} is symmetric in this case. Note that the actual update is carried out using Eq. (2.1): the semi-positive and symmetric matrix X is constructed at each iteration by choosing $X = W^T W$. Using now the framework of Dyson Brownian motion [6, 7] yields a stochastic equation for the eigenvalues x_i ($i = 1, \dots, N$) of X , which reads

$$x_i \rightarrow x'_i = x_i + K_i + \sum_{j \neq i} \frac{g_i^2}{x_i - x_j} + \sqrt{2} g_i \eta_i, \quad (2.3)$$

where K_i and g_i are linked to the deterministic and stochastic terms in Eq. (2.2) and again $\eta_i \sim \mathcal{N}(0, 1)$. The Coulomb term results in eigenvalue repulsion. Making the learning rate and batch size explicit by writing, c.f. Eq. (2.1), $K_i = \alpha \tilde{K}_i$ and $g_i = (\alpha / \sqrt{|\mathcal{B}|}) \tilde{g}_i$ (quantities with a tilde are independent of the learning rate and batch size) yields the eigenvalue equation

$$x_i \rightarrow x'_i = x_i + \alpha \tilde{K}_i + \frac{\alpha^2}{|\mathcal{B}|} \sum_{j \neq i} \frac{\tilde{g}_i^2}{x_i - x_j} + \frac{\alpha}{\sqrt{|\mathcal{B}|}} \sqrt{2} \tilde{g}_i \eta_i. \quad (2.4)$$

At first sight, the scaling with α and $|\mathcal{B}|$ looks unwieldy. However, it falls into place when considering the stationary distribution corresponding to the stochastic processes (2.3, 2.4), which is obtained via the associated Fokker-Planck equation. This so-called Coulomb gas distribution reads

$$P_s(\{x_i\}) = \frac{1}{Z} \prod_{i < j} |x_i - x_j| e^{-\sum_i V_i(x_i)/g_i^2}, \quad Z = \int dx_1 \dots dx_N P_s(\{x_i\}). \quad (2.5)$$

Here it is assumed that the drift K_i can be derived from a separable potential $V_i(x_i)$, via $K_i(x_i) = -dV_i(x_i)/dx_i$. By writing $V_i(x_i) = \alpha \tilde{V}_i(x_i)$, the learning rate and batch size can be made explicit again, and the combination in the exponent reads

$$\frac{V_i(x_i)}{g_i^2} = \frac{1}{\alpha/|\mathcal{B}|} \frac{\tilde{V}_i(x_i)}{\tilde{g}_i^2}. \quad (2.6)$$

The first factor on the RHS indicates universal scaling with $\alpha/|\mathcal{B}|$, while the second factor depends on the details of the loss function. The linear scaling rule, in which $\alpha/|\mathcal{B}|$ is a parameter whose tunability can be exploited, was previously observed empirically [15, 16]. Here it is shown that it is a consequence of stochastic matrix dynamics, when cast in the framework of Brownian motion (see Refs. [17, 18] for an alternative derivation). Reducing the level of stochasticity, by reducing $\alpha/|\mathcal{B}|$, leads to a distribution with spectral weight concentrated more closely around the minima of $V_i(x_i)$.

3 Applications

The derivation above is rather general and the actual expressions for the drift K_i , the noise strength g_i and the Coulomb potential V_i may not be readily available. In this section, we apply the description to two systems: the Gaussian RBM, in which the potential $V_i(x_i)$ can be given explicitly [8], and a transformer, in which the analysis is at an empirical level.

3.1 Gaussian restricted Boltzmann machine

The Gaussian RBM, with continuous degrees of freedom ϕ_i ($i = 1, \dots, N_v$) and h_a ($a = 1, \dots, N_h$) on the visible and hidden layer respectively, is determined by the energy (see e.g. the review [19], here we closely follow Refs. [8, 20])

$$E(\phi, h) = \frac{1}{2}\mu^2\phi^T\phi + \frac{1}{2\sigma_h^2}h^T h - \phi^T W h. \quad (3.1)$$

Here μ^2 and σ_h^2 are hyperparameters and we have put a possible bias to zero. The rectangular weight matrix W , with matrix elements W_{ia} , connects the layers. The induced distribution on the visible layer is determined by the matrix $K_{\text{RBM}} = \mu^2\mathbb{1}_{N_v \times N_v} - \sigma_h^2 W W^T$. The target data is summarised in an $N_v \times N_v$ matrix with eigenvalues κ_i . In the stochastic equation for the eigenvalues x_i of $X = \sigma_h^2 W^T W$, the drift is known explicitly [8],

$$\frac{d}{d\tau}x_i = K_i(x_i) + \sum_{j \neq i} \frac{g_i^2}{x_i - x_j} + \sqrt{2}g_i\eta_i, \quad K_i(x_i) = \left(\frac{1}{\kappa_i} - \frac{1}{\mu^2 - x_i} \right) x_i, \quad (3.2)$$

where $\tau = 2\sigma_h^2 t$ (for notational simplicity, we use here continuous time). The parametric dependence of the stochasticity parameter g_i^2 can be determined explicitly as well, with $g_i^2 \sim (\alpha/|\mathcal{B}|) \times \kappa_i^2 \Omega_i$, with $\Omega_i = (\mu^2 - \kappa_i)/\kappa_i^2$, where, as stated above, the first factor is universal and the second factor is model (μ^2) and data (κ_i) dependent [8].

The Coulomb term in Eq. (3.2) leads to eigenvalue repulsion, with a strength proportional to $\alpha/|\mathcal{B}|$. This is demonstrated in Fig. 1 (left). Only in the limit where $\alpha/|\mathcal{B}| \rightarrow 0$ will the target spectrum be




Figure 1: Gaussian RBM: Ratio of the RBM eigenvalues $\lambda_i = \mu^2 - x_i$ 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 non-vanishing stochasticity (left). Response of the mean level spacing $\langle S \rangle$ to variation of α and $|\mathcal{B}|$, presented in the combination $\sqrt{\alpha/|\mathcal{B}|}$ times a non-universal function $\sqrt{\kappa_i^2 \Omega_i} = \sqrt{\mu^2 - \kappa_i}$ (right). Figures from Ref. [8].

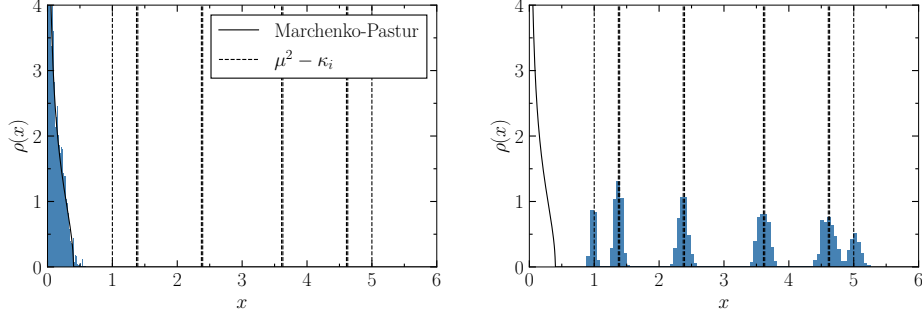


Figure 2: Gaussian RBM: Evolution of eigenvalues of $X = \sigma_h^2 W^T W$ from the Marchenko-Pastur distribution at initialisation (left) to the learned distribution around the target eigenvalues, indicated with the vertical lines, at the end of training (right).

learnt exactly. The distribution of the nonzero level spacing $S_i = x_{i+1} - x_i$ is known as the Wigner surmise. The mean level $\langle S \rangle$ is proportional to $\sqrt{\alpha/|\mathcal{B}|}$, as is confirmed in Fig. 1 (right).

The drift in Eq. (3.2) can be integrated to yield the potential $V_i(x_i) = -\int^{x_i} dx' K_i(x') = -x_i^2/2\kappa_i - x_i - \mu^2 \log(\mu^2 - x_i)$. With the initial matrix elements of W drawn from a normal distribution, the process of learning can then be summarised neatly as the stochastic evolution of the eigenvalues of X from the Marchenko-Pastur distribution initially, as shown in Fig. 2 (left) to the Coulomb gas (2.5) with the potential $V_i(x_i)$ after training, as shown in Fig. 2 (right). See Ref. [8] for more details.

3.2 Transformer

In the Gaussian RBM the drift and potential in the Coulomb gas are known, resulting in full analytic control. In more involved architectures, this is typically not the case. Here we consider a nano-GPT (Generative Pretrained Transformer [21]), with 4 attention blocks each having 4 attention heads. The total number of parameters is around 0.21M. Each attention head contains one key (K), one query (Q) and one value (V) matrix of shape $M \times N = 64 \times 16$, where M is the input feature dimension. The model we use is based on Karpathy’s nano-GPT [22], using the AdamW optimiser [23]. With highly adaptive stepsizes during training, the dependence on learning rate and batch size is more involved. Hence we explore here the evolution of the eigenvalue distribution of $X = W^T W$ through training, where W is one of the attention matrices, at an empirical level.

The initial matrix elements are drawn from a uniform distribution, bounded between $\pm 1/\sqrt{M}$. The eigenvalue distribution of X at initialisation is then given by an Marchenko-Pastur (MP) distribution

$$P_{\text{MP}}(x; \sigma^2, A) = \frac{A}{2\pi\sigma^2 r x} \sqrt{(x_+ - x)(x - x_-)} \theta(x_+ - x)\theta(x - x_-), \quad (3.3)$$

with $x_{\pm} = \sigma^2(1 \pm \sqrt{r})^2$, $r = N/M = 1/4$, area $A = 1$ and $\sigma^2 = 1/3$ for the uniform initialisation.

In Fig. 3 we show the evolution of the eigenvalue distribution of $X = K^T K$, where K is the Key matrix of the transformer’s first layer. To verify universal aspects, we show in the top row the distribution $P(s)$ of the normalised eigenvalue spacing, $s_i = x_{i+1} - x_i$, after spectral unfolding. The histograms are compared to the Wigner surmise, $P(s) = (\pi/2)s \exp(-\pi s^2/4)$, observing good agreement. This indicates that the eigenvalues undergo universal RMT fluctuations.

In the bottom row the evolution of the spectral density $\rho(x)$ is shown. From the initial MP distribution, one observes a shift of the spectrum to larger values, no longer described by an MP distribution (note the different horizontal scales). The appearance of (heavy) tails has been noted also in other architectures [10]. This evolution should be compared to the evolution in the RBM in Fig. 2, with the notable difference that the “exact” spectrum or Coulomb gas potential are not known. To quantify the evolution, we fit the spectral density to the MP distribution with fit parameters σ^2 and area A . The evolution of those is shown in Fig. 4 for the Key matrix of the first head for the four layers. As more spectral weight moves to the tail, the area A of the part of the distribution still described by the

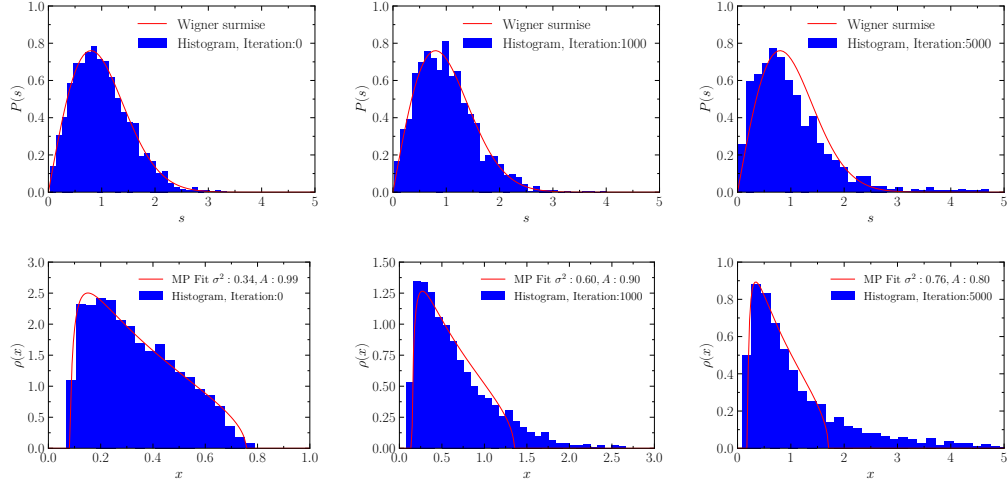


Figure 3: Transformer: Evolution during training of the eigenvalue distribution of $X = K^T K$, where K is the Key matrix of the transformer’s first layer, at initialisation (left), iteration 1000 (middle) and iteration 5000 (right). Above: distribution $P(s)$ of the normalised eigenvalue spacing $s_i = x_{i+1} - x_i$ after spectral unfolding, compared to the Wigner surmise. Below: spectral density $\rho(x)$, compared to fits to the Marchenko-Pastur distribution with fit parameters σ^2 and area A .




Figure 4: Transformer: Evolution of fit parameters area A (left) and σ^2 (right) of the Marchenko-Pastur distribution fit to the spectral density $\rho(x)$ of $X = K^T K$ of the first head for all layers. To determine the statistical uncertainty, training is repeated at least 50 times, using a bootstrap analysis.

MP distribution decreases. This is demonstrated in Fig. 4 (left), from which one deduces that the tail carries 15-25% of the spectral weight. The MP distribution itself evolves, as is shown in Fig. 4 (right), with the σ^2 parameter increasing from $1/3$ to closer to 1, depending on the layer. This may reflect further randomisation of the lower part of the spectrum, with Brownian motion leading to an MP distribution with a larger domain (larger $x_+ \sim \sigma^2$). Work to further understand the spectral form of the tail is in progress, as is a study of the apparent non-monotonic behaviour of the area A .

4 Summary

We have argued that weight matrix dynamics should be viewed in the framework of Dyson Brownian motion. This predicts a universal dependence on the ratio of learning rate over batch size. It elegantly leads to a picture in which the process of learning is described as the stochastic evolution of the eigenvalues of $X = W^T W$ from an initial Marchenko-Pastur distribution to a Coulomb gas distribution with a model-dependent potential. This picture is fully confirmed in the Gaussian restricted Boltzmann machine and supported empirically in a nano-GPT. Further analysis of the latter is in progress.

Acknowledgements – Part of this work was initiated at the ECT* workshop *Machine Learning and the Renormalisation Group* in May 2024. We thank the participants of this meeting for discussion and ECT* for support. GA and BL are supported by STFC Consolidated Grant ST/T000813/1. BL is further supported by the UKRI EPSRC ExCALIBUR ExaTEPP project EP/X017168/1. CP is supported by the UKRI AIMLAC CDT EP/S023992/1.

Research Data and Code Access – The code and data used for the first part of this manuscript are available from Ref. [24].

Open Access Statement – For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

References

- [1] E.P. Wigner, *Characteristic vectors of bordered matrices with infinite dimensions*, *Annals of Mathematics* **62** (1955) 548.
- [2] E.P. Wigner, *Conference on Neutron Physics by Time-of-Flight*, p. 67, 1956.
- [3] F.J. Dyson, *Statistical theory of the energy levels of complex systems. I*, *J. Math. Phys.* **3** (1962) 140.
- [4] F.J. Dyson, *Statistical theory of the energy levels of complex systems. II*, *J. of Math. Phys.* **3** (1962) 157.
- [5] F.J. Dyson, *Statistical Theory of the Energy Levels of Complex Systems. III*, *J. Math. Phys.* **3** (1962) 166.
- [6] M.L. Mehta, *Random Matrices*, Academic Press, New York, 3rd ed. (2004).
- [7] F.J. Dyson, *A Brownian-Motion Model for the Eigenvalues of a Random Matrix*, *J. Math. Phys.* **3** (1962) 1191.
- [8] G. Aarts, B. Lucini and C. Park, *Stochastic weight matrix dynamics during learning and Dyson Brownian motion*, 2407.16427.
- [9] J. Pennington and P. Worah, *Nonlinear random matrix theory for deep learning*, in *Advances in Neural Information Processing Systems*, pp. 2637–2646, 2017, <https://papers.nips.cc/paper/6857-nonlinear-random-matrix-theory-for-deep-learning>.
- [10] C.H. Martin and M.W. Mahoney, *Traditional and Heavy-Tailed Self Regularization in Neural Network Models*, 1901.08276.
- [11] N.P. Baskerville, D. Granzoli and J.P. Keating, *Applicability of Random Matrix Theory in Deep Learning*, 2102.06740.
- [12] N.P. Baskerville, J.P. Keating, F. Mezzadri, J. Najnudel and D. Granzoli, *Universal characteristics of deep neural network loss surfaces from random matrix theory*, *Journal of Physics A: Mathematical and Theoretical* **55** (2022) 494002 [2205.08601].
- [13] R. Couillet and Z. Liao, *Random Matrix Methods for Machine Learning*, Cambridge University Press (2022), 10.1017/9781009128490.
- [14] N. Levi and Y. Oz, *The Universal Statistical Structure and Scaling Laws of Chaos and Turbulence*, 2311.01358.
- [15] S.L. Smith and Q.V. Le, *A Bayesian Perspective on Generalization and Stochastic Gradient Descent*, 1710.06451.
- [16] S.L. Smith, P. Kindermans and Q.V. Le, *Don't Decay the Learning Rate, Increase the Batch Size*, 1711.00489.
- [17] P. Chaudhari and S. Soatto, *Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks*, 1710.11029.
- [18] 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*, vol. 70, pp. 2101–2110, 2017 [1511.06251].
- [19] A. Decelle and C. Furtlehner, *Restricted Boltzmann machine: Recent advances and mean-field theory*, *Chinese Physics B* **30** (2021) 040202 [2011.11307].

- [20] G. Aarts, B. Lucini and C. Park, *Scalar field restricted Boltzmann machine as an ultraviolet regulator*, *Phys. Rev. D* **109** (2024) 034521 [2309.15002].
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez et al., *Attention Is All You Need*, 1706.03762.
- [22] A. Karpathy, *NanoGPT*, <https://github.com/karpathy/nanoGPT.git>, 2023.
- [23] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 1412.6980.
- [24] C. Park, G. Aarts and B. Lucini, *chanjure/Stochastic_weight_matrix_dynamics_during_learning_and_Dyson_Brownian_motion-data_release: v1.0.1*, 2024. 10.5281/zenodo.13310439.