



Baskerville, N. P., Keating, J., Mezzadri, F., Najnudel, J., & Granzio, D. (2022). Universal characteristics of deep neural network loss surfaces from random matrix theory. *Journal of Physics A: Mathematical and Theoretical*, 55(49), Article 494002.
<https://doi.org/10.1088/1751-8121/aca7f5>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1088/1751-8121/aca7f5](https://doi.org/10.1088/1751-8121/aca7f5)

[Link to publication record on the Bristol Research Portal](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via IOP Publishing at <https://doi.org/10.1088/1751-8121/aca7f5>. Please refer to any applicable terms of use of the publisher.

University of Bristol – Bristol Research Portal

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/brp-terms/>

PAPER • OPEN ACCESS

Universal characteristics of deep neural network loss surfaces from random matrix theory





To cite this article: Nicholas P Baskerville *et al* 2022 *J. Phys. A: Math. Theor.* **55** 494002

View the [article online](#) for updates and enhancements.

You may also like

- [Emerging memory technologies for neuromorphic computing](#)
Chul-Heung Kim, Suhwan Lim, Sung Yun Woo et al.
- [Improving the accuracy and robustness of RRAM-based in-memory computing against RRAM hardware noise and adversarial attacks](#)
Sai Kiran Cherupally, Jian Meng, Adnan Siraj Rakin et al.
- [Deep learning in electron microscopy](#)
Jeffrey M Ede

Universal characteristics of deep neural network loss surfaces from random matrix theory

Nicholas P Baskerville^{1,*} , Jonathan P Keating² ,
Francesco Mezzadri¹ , Joseph Najnudel¹ 
and Diego Granzio³ 

¹ School of Mathematics, University of Bristol, Fry Building, Woodland Road, Bristol BS8 1UG, United Kingdom

² Mathematical Institute, University of Oxford, Andrew Wiles Building, Woodstock Road, Oxford OX2 6GG, United Kingdom

³ PureStrength Ltd, London, United Kingdom

E-mail: n.p.baskerville@bristol.ac.uk

Received 12 May 2022; revised 15 November 2022

Accepted for publication 2 December 2022

Published 16 December 2022



CrossMark

Abstract

This paper considers several aspects of random matrix universality in deep neural networks (DNNs). Motivated by recent experimental work, we use universal properties of random matrices related to local statistics to derive practical implications for DNNs based on a realistic model of their Hessians. In particular we derive universal aspects of outliers in the spectra of deep neural networks and demonstrate the important role of random matrix local laws in popular pre-conditioning gradient descent algorithms. We also present insights into DNN loss surfaces from quite general arguments based on tools from statistical physics and random matrix theory.

Keywords: random matrix theory, neural networks, universality, loss surfaces

(Some figures may appear in colour only in the online journal)

* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

The success of large deep neural networks (DNNs) optimised in surprisingly naïve ways, such as stochastic gradient descent, has spawned a considerable amount of interest in characterising their loss surfaces. The *loss* of a neural network is simply a real valued function that measures the network's performance on a task with respect to some reference data set. Neural networks are defined in terms of parameters, or *weights*, of which there are typically a very large number. Fixing the data set, or even the data generating distribution, a neural network's loss can be seen as a surface in high dimensions parametrised by the network's weights.

It was first observed⁴ in a sequence of experimental and theoretical papers [CHM+15, SGAL14] that the loss surfaces of DNNs can be connected with spherical multi-spin glasses from statistical physics [MPV87]. Mathematically, multi-spin glasses are just certain random multivariate polynomials, or equivalently Gaussian processes with specific covariance functions. These works were very influential, providing the first steps towards explaining why DNNs can be trained at all, given the apparent intractability of their optimisation from the viewpoint of classical optimisation theory. The key insight from the spin glass connection was that DNN loss surfaces are indeed very complicated and filled with many local optima, but that, in the high-dimensional limit, the local optima are arranged favourably, so that simple gradient based optimisation methods can be expected to converge to local optima close to the global minimum. This work has been extended in several directions. [BKMN21] extended the scope of the results to more general neural networks, and in so doing uncovered some of the limits of the spin glass model's explanatory power. Departing from the direct connection between glassy systems and neural networks, several authors have used similar high-dimensional random models as toy models for high-dimensional optimisation. Among these works are several that focus on the existence of spurious minima and the problem of recovering signals from high dimensional signal-in-noise models [RABC19, MBAB20, MBC+19]. Similarly, using spin glasses as surrogates for the complex, high-dimensional, random loss surfaces of DNNs, [BKMN22] extended the spin glass analysis to study the non-standard loss surfaces of generative adversarial networks [GPAM+14].

A significant feature common amongst all the work mentioned so far is random matrix theory [Meh04, LXT+18, AGZ10]. When studying the number and configuration of local optima of high dimensional random functions, Kac-Rice formulae are an essential mathematical tool, and random matrices arise as Hessians in these formulae. For spin glasses, this path was first trodden in the theoretical physics literature by Fyodorov [Fyo04, Fyo05] and made rigorous in later work [AAČ13]. Specifically in the case of [CHM+15] and [AAČ13], which lays its mathematical foundation, very detailed calculations can be completed. This is possible because the random matrices that appear belong to the Gaussian Orthogonal Ensemble (GOE), one of the canonical ensembles of random matrix theory. Almost anything one could want to know of the GOE is known and many powerful tools of random matrix theory can be applied to it. In particular, the full joint distribution of its eigenvalues and large deviations results for its eigenvalues and its spectral density, all of which are required by [AAČ13, CHM+15]. Works such as [BKMN21, BKMN22] show how detailed calculations can be completed beyond the standard spin glass case, however these results still depend on important properties of the GOE, as the Hessians in those cases are closely related to the GOE. In a recent work, [GZR20] showed how valuable practical insights about DNN optimisation can be obtained by considering the outliers in the spectrum on the loss surface Hessian. Once again, this work relies on special properties from random matrix theory. Firstly, an expression for the outliers

⁴ Note that early connections between neural networks and spin systems were presented in [Gar88].

follows from a known phase transition result whereby the largest eigenvalue ‘pops out’ of the bulk. This result has been proven only for rotationally invariant matrix ensembles in [BGN11], itself a generalisation of the celebrated BBP phase transition Baik, Ben Arous and Peché [BAP05], though it was conjectured in [BGN11] to be more general (see section 2.3 below). Secondly, the explicit form of a Wigner semi-circle density was used to obtain the concrete outlier expression used in practice.

Challenging the above-mentioned works, an experimental line of work has demonstrated convincingly that special RMT ensembles like the GOE do not appear to be present in DNNs [Pap18, Gra20, BGK22], for example as their Hessians. In addition, there have been challenges in the literature to the practical relevance of spin glass loss surface results for DNNs [BJSJG+19]. In this context, it is natural to question the relevance to DNNs of many of the results discussed above. Indeed, does random matrix theory actually provide insight into real DNNs, or are its powerful tools merely applicable to toy models too divorced from real DNNs to be of any value? Random matrix theory itself may hold the answer to this question, in particular its notion of *universality*. Broadly speaking, universality refers to the phenomenon that certain properties of special random matrix ensembles (such as the GOE) remain true for more general random matrices that share some key feature with the special ensembles. For example, the Wigner semicircle is the limiting spectral density of the Gaussian Wigner ensembles, i.e. matrices with Gaussian entries, independent up to symmetry (symmetric real matrices, Hermitian complex matrices) [Meh04]. The Gaussian case is the simplest to prove, and there are various powerful tools not available in the non-Gaussian case, however the Wigner semicircle has been established as the limiting spectral density for Wigner matrices with quite general distributions on their entries [AGZ10, Tao12]. While surprisingly general in some sense, the Wigner semicircle relies on independence (up to symmetry) of matrix entries, a condition which is not typically satisfied in real systems. The limiting form of the spectral density of a random matrix ensemble is a *macroscopic* property, i.e. the matrix is normalised such that the average distance between adjacent eigenvalues is on the order of $1/\sqrt{N}$, where N is the matrix size. At the opposite end of the scale is the *microscopic*, where the normalisation is such that eigenvalues are spaced on a scale of order 1; at this scale, random matrices display a remarkable universality. For example, any real symmetric matrix has a set of orthonormal eigenvectors and so the set of all real symmetric matrices is closed under conjugation by orthogonal matrices. Wigner conjectured that certain properties of GOE matrices hold for very general random matrices that share the same (orthogonal) symmetry class, namely symmetric random matrices (the same is true of Hermitian random matrices and the unitary symmetry class). The spacings between adjacent eigenvalues should follow a certain explicit distribution, the Wigner surmise, and the eigenvectors should be *delocalised*, i.e. the entries should all be of the same order as the matrix size grows. Both of these properties are true for the GOE and can be proved straightforwardly with quite elementary techniques. Indeed, in the case of 2×2 GOE, it is a standard first exercise in random matrix theory to prove that the eigenvalue spacing distribution is precisely the Wigner surmise (for $N \times N$ GOEs it is only a good approximation and improves as $N \rightarrow \infty$). Microscopic random matrix universality is known to be far more robust than universality on the macroscopic scale. Indeed, such results are well established for invariant ensembles and can be proved using Riemann-Hilbert methods [Dei99]. For more general random matrices, microscopic universality has been proved by quite different methods in a series of works over the last decade or so, of which a good review is [EY17]. Crucial in these results is the notion of a *local law* for random matrices. The technical statement of local laws is given later in the paper, but roughly they assert that the spectrum of a random matrix is, with very high probability, close to the deterministic spectrum defined by its limiting spectral density (e.g. the semicircle law for Wigner matrices). Techniques vary by

ensemble, but generally a local law for a random matrix ensemble provides the control required to demonstrate that certain matrix statistics are essentially invariant under the evolution of the Dyson Brownian motion. In the case of real symmetric matrices, the Dyson Brownian motion converges in finite time to the GOE, hence the statistics preserved under the Dyson Brownian motion must match the GOE. The n -point correlation functions of eigenvalues are one such preserved quantity, from which follows, amongst other properties, that the Wigner surmise is a good approximation to the adjacent spacings distribution.

At the macroscopic scale, there are results relevant to neural networks, for example [PSG18, Pas20] consider random neural networks with Gaussian weights and establish results that are generalised to arbitrary distributions with optimal conditions, so demonstrating universality. On the microscopic scale, [BGK22] provided the first experimental demonstration of the presence of universal local random matrix statistics DNNs, specifically in the Hessians and Gauss-Newton matrices of their loss surfaces. This work has recently been extended to the weight matrices of neural networks [TSR22].

This paper explores the consequences of random matrix universality in DNNs. Our main mathematical result is a significant generalisation of the Hessian spectral outlier result recently presented by [GZR20]. This generalisation removes any need for GOE or Wigner forms of the Hessian and instead leverages much more universal properties of the eigenvectors and eigenvalues of random matrices which we argue are quite likely to hold for real networks. Our results make concrete predictions about the outliers of DNN Hessians which we compare with experiments on several real-world DNNs. These experiments provide indirect evidence of the presence of universal random matrix statistics in the Hessians of large DNNs, which is noteworthy as certainly these DNNs are far too large to permit exact eigendecomposition of their Hessians as in [BGK22]. Along a similar line, we show how local random matrix laws in DNNs can dramatically simplify the dynamics of certain gradient descent optimisation algorithms and may be in part responsible for their success. Finally, we highlight another aspect of random matrix universality relevant to DNN loss surfaces. Recent work [ABM21] has shown that the so-called ‘self averaging’ property of random matrix determinants is very much more universal than previously thought. The self-averaging of random matrix determinants has been used in the spin glass literature both rigorously and non-rigorously (e.g. [Fyo04, Fyo05, AAC13, BKMN21, BKMN22] *inter alia*) and is the key property that produces the exponentially large/small number of local optima repeatedly observed. We argue that insights into the geometry of DNN loss surfaces can be conjectured from quite general assumptions about the Hessian and gradient noise and from the general self-averaging effect of random matrix determinants.

The paper is structured as follows. Section 2 introduces our random matrix theory model for DNN Hessians, derives results for their outliers and compares with experimental results. Section 3 presents the proof of our main result for addition of random matrices, combining quantum unique ergodicity (QUE) and the supersymmetric method. Section 4 proves an extension of the random matrix determinant results of [ABM21] and presents insights into DNN loss surfaces from complexity calculations and random matrix determinants. Section 5 describes the role of random matrix local laws in certain popular DNN optimisation algorithms. Section 6 concludes the paper.

Notation. We adopt the following conventions throughout

- For a probability measure μ , g_μ is its Stieltjes transform.
- $\mu \boxplus \nu$ denotes the free additive convolution between probability measures μ and ν .
- Hats denote empirical quantities unless stated otherwise. For example $\hat{\mu}_X$ is the empirical spectral measure of a matrix X .

- $r(\mu), l(\mu)$ denote the right and left edges of the support of a probability measure μ respectively.
- The bounded-Lipschitz distance between two probability measures is

$$d_{\text{BL}}(\mu, \nu) = \sup \left\{ \left| \int_{\mathbb{R}} f(x)(\mu - \nu)(dx) \right| : \|f\|_{\text{Lip}} + \|f\|_{L^\infty} \leq 1 \right\}, \quad (1.1)$$

where $\|f\|_{\text{Lip}}$ is the Lipschitz constant of a Lipschitz function f .

- The Wasserstein-1 distance between two probability measures is

$$W_1(\mu, \nu) = \sup \left\{ \left| \int_{\mathbb{R}} f(x)(\mu - \nu)(dx) \right| : \|f\|_{\text{Lip}} \leq 1 \right\}. \quad (1.2)$$

2. General random matrix model for loss surface Hessians

2.1. The model

Given a loss function $\mathcal{L} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, a data generating distribution \mathbb{P}_{data} supported on $\mathcal{X} \times \mathcal{Y}$ and a neural network $f_{\mathbf{w}} : \mathcal{X} \rightarrow \mathcal{Y}$ parametrised by $\mathbf{w} \in \mathbb{R}^N$, its batch Hessian is given by

$$H_{\text{batch}} = \frac{1}{b} \sum_{i=1}^b \frac{\partial^2}{\partial \mathbf{w}^2} \mathcal{L}(f_{\mathbf{w}}(\mathbf{x}_i), y_i), \quad (\mathbf{x}_i, y_i) \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_{\text{data}} \quad (2.1)$$

and its true Hessian is given by

$$H_{\text{true}} = \mathbb{E}_{(\mathbf{x}, y) \sim \mathbb{P}_{\text{data}}} \frac{\partial^2}{\partial \mathbf{w}^2} \mathcal{L}(f_{\mathbf{w}}(\mathbf{x}), y). \quad (2.2)$$

Both H_{batch} and H_{true} are $N \times N$ matrix functions of \mathbf{w} ; H_{batch} is random but H_{true} is deterministic. Only in very specific cases and under strong simplifying assumptions can one hope to obtain the distribution of H_{batch} or the value of H_{true} from $\mathcal{L}, \mathbb{P}_{\text{data}}$ and $f_{\mathbf{w}}$. Inspired by the success of many random matrix theory applications, e.g. in Physics, we will instead seek to capture the essential features of DNN Hessians in a sufficiently general random matrix model.

We introduce the following objects:

- A sequence (in N) of random real symmetric $N \times N$ matrices X . X possesses a limiting spectral probability measure μ , i.e. if $\lambda_1, \dots, \lambda_N$ are the eigenvalues of X then

$$\frac{1}{N} \sum_{i=1}^N \delta_{\lambda_i} \rightarrow \mu \quad (2.3)$$

weakly almost surely. We further assume that μ has compact support and admits a smooth density with respect to Lebesgue measure.

- A sequence (in N) of deterministic real symmetric $N \times N$ matrices A with eigenvalues

$$\theta_1, \dots, \theta_p, \xi_1, \dots, \xi_{N-p-q}, \theta'_1, \dots, \theta'_q \quad (2.4)$$

for fixed integers p, q . We assume the existence of limiting measure ν such that, weakly,

$$\frac{1}{N-p-q} \sum_{i=1}^{N-p-q} \delta_{\xi_i} \rightarrow \nu \quad (2.5)$$

where ν is a compactly supported probability measure. The remaining eigenvalues satisfy

$$\theta_1 > \dots > \theta_p > \mathfrak{r}(\nu), \theta'_1 < \dots < \theta'_q < \mathfrak{l}(\nu). \quad (2.6)$$

ν is also assumed to be of the form $\nu = \varepsilon\eta + (1 - \varepsilon)\delta_0$ where η is a compactly supported probability measure which admits a density with respect to Lebesgue measure.

- A decreasing function $\mathfrak{s} : \mathbb{N} \rightarrow (0, 1)$.

With these definitions, we construct the following model for the Hessian:

$$H_{\text{batch}} \equiv H = \mathfrak{s}(b)X + A \quad (2.7)$$

where b is the batch size. We have dropped the subscript on H_{batch} for brevity. Note that H takes the place of the batch Hessian and A taken the place of the true Hessian. $\mathfrak{s}(b)X$ takes the place of the random noise introduced by sampling a finite batch at which to evaluate the Hessian. $\mathfrak{s}(b)$ is an overall scaling induced in X by the batch-wise averaging.

This model is almost completely general. Note that we allow the distribution of X and the value of A to depend on the position in weight space \mathfrak{w} . The only restrictions imposed by the model are

- (a) the existence of ν ;
- (b) the position of θ_i, θ'_j relative to the support of ν ;
- (c) ν may only possess an atom at 0;
- (d) the fixed number of θ_i, θ'_j ;
- (e) the existence of μ ;
- (f) the existence of the scaling $\mathfrak{s}(b)$ in batch size.

All of the above restrictions are discussed later in the section. Finally, we must introduce some properties of the noise model X in order to make any progress. We introduce the assumption that the eigenvectors of X obey QUE [BY17]. The precise meaning of this assumption and a thorough justification and motivation is given later in this section. For now it suffices to say that QUE roughly means that the eigenvectors of X are *delocalised* or that they behave roughly like the rows (or columns) of a uniform random $N \times N$ orthogonal matrix (i.e. a matrix with Haar measure). QUE is known to hold for standard ensembles in random matrix theory, such as quite general Wigner matrices, Wishart matrices, adjacency matrices of certain random graphs etc. Moreover, as discussed further section 2.5 below, QUE can be thought of as a property of quite general random matrix models.

2.2. QUE

It is well known that the eigenvectors of quite general random matrices display a universal property of *delocalisation*, namely

$$|u_k|^2 \sim \frac{1}{N} \quad (2.8)$$

for any component u_k of an eigenvector \mathbf{u} . Universal delocalisation was conjectured by Wigner along with the Wigner surmise for adjacent eigenvalue spacing. Both of these properties, and the more phenomenon of universal correlation functions on the microscopic scale have since been rigorously established for quite a variety of matrix models e.g. [EY17, EY12, EKS19]. [BY17] show that the eigenvectors of generalised Wigner matrices obey QUE, a particular

form of delocalisation, stronger than the above statement. Specifically, they are shown to be approximately Gaussian in the following sense ([BY17] theorem 1.2):

$$\sup_{\substack{\|q\|=1 \\ I \subset [N], \\ |I|=n}} \sup_{|I|=n} \left| \mathbb{E}P \left((N|q^T \mathbf{u}_k|^2)_{k \in I} \right) - \mathbb{E}P \left((|\mathcal{N}_j|^2)_{j=1}^n \right) \right| \leq N^{-\varepsilon}, \quad (2.9)$$

for large enough N , where \mathcal{N}_j are i.i.d. standard normal random variables, $(\mathbf{u}_k)_{k=1}^N$ are the normalised eigenvectors, P is any polynomial in n variables and $\varepsilon > 0$. Note that the set I in this statement is a subset of $[N]$ of fixed size n ; n is not permitted to depend on N .

2.3. Batch Hessian outliers

Let $\{\lambda_i\}$ be the eigenvalues of H . To set the context of our results, let us first simplify and suppose momentarily that $\mathbf{s} = 1$ and, instead of mere QUE, X has eigenvectors distributed with Haar measure, and A is fixed rank, i.e. $\xi_i = 0 \forall i$, then the results of [BGN11] would apply and give

$$\lambda_j \xrightarrow{a.s.} \begin{cases} g_\mu^{-1}(1/\theta_j) & \text{if } \theta_j > 1/g_\mu(\mathbf{r}(\mu)), \\ \mathbf{r}(\mu) & \text{otherwise,} \end{cases} \quad (2.10)$$

for $j = 1, \dots, p$, and

$$\lambda_{N-j+1} \xrightarrow{a.s.} \begin{cases} g_\mu^{-1}(1/\theta'_j) & \text{if } \theta'_j < 1/g_\mu(\mathbf{1}(\mu)), \\ \mathbf{1}(\mu) & \text{otherwise,} \end{cases} \quad (2.11)$$

for $j = 1, \dots, q$.

What follows is our main results for the outliers of H under the general conditions described above.

Theorem 2.1. *Let H be the Hessian matrix model defined in (2.7) and meeting all the conditions in section 2. Then there exist $U_\varepsilon, L_\varepsilon \in \mathbb{R}$ such that, for $j = 1, \dots, p$,*

$$\lambda_j = \begin{cases} \omega^{-1}(\theta_j) & \text{if } \omega^{-1}(\theta_j) > U_\varepsilon, \\ U_\varepsilon & \text{otherwise.} \end{cases} \quad (2.12)$$

and for $j = 1, \dots, q$,

$$\lambda_{N-j+1} = \begin{cases} \omega^{-1}(\theta'_j) & \text{if } \omega^{-1}(\theta'_j) < L_\varepsilon, \\ L_\varepsilon & \text{otherwise,} \end{cases} \quad (2.13)$$

and

$$\omega^{-1}(\theta) = \theta + \mathbf{s}(b)R_\mu(\mathbf{s}(b)\theta^{-1}) + \varepsilon \mathbf{s}(b)^2 d_\eta(\theta) R'_\mu(\mathbf{s}(b)\theta^{-1}) + \mathcal{O}(\varepsilon^2) \quad (2.14)$$

where we define $d_\eta(z) = g_\eta(\theta_j) - \theta_j^{-1}$.

2.3.1. An interlude on prior outlier results. It was conjectured in [BGN11] that (2.10) and (2.11) still hold when X has delocalised eigenvectors in some sense, rather than strictly Haar. Indeed, a careful consideration of the proof in that work does reveal that something weaker than Haar would suffice, for example QUE. See in particular the proof of the critical lemma 9.2 therein which can clearly be repeated using QUE. There is a considerable subtlety here, however, which is revealed best by considering more recent results on deformations of general Wigner matrices. [KY17] shows that very general deterministic deformations of general

Wigner matrices possess an optimal anisotropic local law, i.e. $Y + B$ for Wigner Y and deterministic symmetric B . It is expected therefore that $Y + B$ has delocalised eigenvectors in the bulk. Consider the case where B is diagonal, and say that B has a fixed number of ‘spike’ eigenvalues $\varphi_1 > \dots > \varphi_r$ and remaining eigenvalues $\zeta_1, \dots, \zeta_{N-r}$ where the empirical measure of the ζ_i converges to some measure τ and $\varphi_r > \mathfrak{r}(\tau)$. We can then split $B = B_i + B_o$ where B_i contains only the ζ_j and B_o only the φ_j . The previously mentioned results applies to $Y + B_i$ and then we might expect the generalised result of [BGN11] to apply to give outliers $g_{\mu_{\text{sc}} \boxplus \tau}^{-1}(1/\varphi_i)$ of $Y + B$. This contradicts, however, another result concerning precisely the outliers of such generally deformed Wigner matrices. It was shown in [CDM16] that the outliers of $Y + B$ are $\omega^{-1}(\varphi_j)$ where ω is the subordination function such that $g_{\mu_{\text{sc}} \boxplus \tau}(z) = g_\tau(\omega(z))$. These two expressions coincide when

$$\begin{aligned} \omega^{-1}(z) &= g_{\mu_{\text{sc}} \boxplus \tau}^{-1}(z^{-1}) \\ \iff \omega^{-1}(z) &= \omega^{-1}(g_\tau^{-1}(z^{-1})) \\ \iff g_\tau^{-1}(z^{-1}) &= z \\ \iff g_\tau(z) &= z^{-1} \\ \iff \tau &= \delta_0, \end{aligned} \tag{2.15}$$

i.e. only when B is in fact of negligible rank as $N \rightarrow \infty$. This apparent contradiction is resolved by the observation that the proof in [BGN11] in fact relies implicitly on an *isotropic local law*. Note in particular section 4.1, which translated to our context, would require $\mathbf{v}^T G_{Y+B_i}(z) \mathbf{v} \approx g_{\mu_{\text{sc}} \boxplus \tau}(z)$ with high probability for general unit vectors \mathbf{v} . Such a result holds if and only if $Y + B_i$ obeys an isotropic local law and is violated if its local law is instead anisotropic, as indeed it is, thanks to the deformation.

Proof of theorem 2.1. The conditions on X required to invoke theorem 3.4 from section 3 are satisfied, so we conclude that

$$\hat{g}_H(z) = g_{\mu_b \boxplus \nu}(z) + o(1) = g_\nu(\omega(z)) + o(1) = \hat{g}_A(\omega(z)) + o(1) \tag{2.16}$$

where ω is the subordination function such that $g_{\mu_b \boxplus \nu}(z) = g_\nu(\omega(z))$ and μ_b is the limiting spectral measure of $\mathfrak{s}(b)X$. The reasoning found in [CDM16] then applies regarding the outliers of H . Indeed, suppose that λ is an outlier of H , i.e. λ is an eigenvalue of H contained in $\mathbb{R} \setminus \text{supp}(\mu \boxplus \nu)$. Necessarily \hat{g}_H possesses a singularity at λ , and so \hat{g}_A must have a singularity at $\omega(\lambda)$. For this singularity to persist for all N , $\omega(\lambda)$ must coincide with one of the outliers of A which, unlike the bulk eigenvalues ξ_j , remain fixed for all N . Therefore we have the following expressions for the outliers of H :

$$\{\omega^{-1}(\theta_j) \mid \omega^{-1}(\theta_j) \in \mathbb{R} \setminus \text{supp}(\mu_b \boxplus \nu)\} \cup \{\omega^{-1}(\theta'_j) \mid \omega^{-1}(\theta'_j) \in \mathbb{R} \setminus \text{supp}(\mu_b \boxplus \nu)\}. \tag{2.17}$$

We now consider ε to be small and analyse these outlier locations as a perturbation in ε . Firstly note that

$$g_{\mu_b}(z) = \int \frac{d\mu_b(x)}{z-x} = \int \frac{d\mu(x/\mathfrak{s}(b))}{z-x} = \mathfrak{s}(b) \int \frac{d\mu(x)}{z-\mathfrak{s}(b)x} = g_\mu(z/\mathfrak{s}(b)). \tag{2.18}$$

Also

$$\omega^{-1}(z) = g_{\mu_b \boxplus \nu}^{-1}(g_\nu(z)) \tag{2.19}$$

$$\begin{aligned} &= R_{\mu_b}(g_\nu(z)) + g_\nu^{-1}(g_\nu(z)) \\ &= R_{\mu_b}(g_\nu(z)) + z. \end{aligned} \tag{2.20}$$

We now must take care in computing R_{μ_b} from g_{μ_b} . Recall that the R -transform of a measure is defined as a formal power series [AGZ10]

$$R(z) = \sum_{n=0}^{\infty} k_{n+1} z^n \tag{2.21}$$

where k_n is the n th cumulant of the measure. It is known [AGZ10] that $k_n = C_n$ where the functional inverse of the Stieljtes transform of the measure is given by the formal power series

$$K(z) = \frac{1}{z} + \sum_{n=1} C_n z^{n-1}. \tag{2.22}$$

Now let m_n be the n th moment of μ and similarly let $m_n^{(b)}$ be the n th moment of μ_b , so formally

$$g_{\mu}(z) = \sum_{n \geq 0} m_n z^{-(n+1)}, \quad g_{\mu_b}(z) = \sum_{n \geq 0} m_n^{(b)} z^{-(n+1)}. \tag{2.23}$$

Also let k_n be the n th cumulant of μ and $k_n^{(b)}$ be the n th cumulant of μ_b . Referring to the proof of lemma 5.3.24 in [AGZ10] we find the relations

$$m_n = \sum_{r=1}^n \sum_{\substack{0 \leq i_1, \dots, i_r \leq n-r \\ i_1 + \dots + i_r = n-r}} k_r m_{i_1} \dots m_{i_r}, \tag{2.24}$$

$$m_n^{(b)} = \sum_{r=1}^n \sum_{\substack{0 \leq i_1, \dots, i_r \leq n-r \\ i_1 + \dots + i_r = n-r}} k_r^{(b)} m_{i_1}^{(b)} \dots m_{i_r}^{(b)}. \tag{2.24}$$

Note, in particular, that $m_1 = k_1$. But clearly the moments of μ_b have a simple scaling in $\mathbf{s}(b)$, namely $m_n^{(b)} = \mathbf{s}(b)^n m_n$, hence

$$m_n = \mathbf{s}(b)^{-n} \sum_{r=1}^n \sum_{\substack{0 \leq i_1, \dots, i_r \leq n-r \\ i_1 + \dots + i_r = n-r}} k_r^{(b)} m_{i_1} \dots m_{i_r} \mathbf{s}(b)^{n-r} \tag{2.26}$$

from which we deduce $k_n^{(b)} = \mathbf{s}(b)^n k_n$, which establishes that $R_{\mu_b}(z) = \mathbf{s}(b) R_{\mu}(\mathbf{s}(b)z)$. Recalling (2.20) we find

$$\omega^{-1}(z) = \mathbf{s}(b) R_{\mu}(\mathbf{s}(b)g_{\nu}(z)) + z. \tag{2.27}$$

The form of ν gives

$$g_{\nu}(z) = (1 - \varepsilon) \int \frac{dt}{z-t} \delta_0(t) + \varepsilon \int \frac{d\eta(t)}{t-z} = \frac{1-\varepsilon}{z} + \varepsilon g_{\eta}(z) = \frac{1}{z} + \varepsilon \left(g_{\eta}(z) - \frac{1}{z} \right) \tag{2.28}$$

and so we can expand to give

$$\begin{aligned} \omega^{-1}(\theta_j) &= \theta_j + \mathbf{s}(b) R_{\mu}(\mathbf{s}(b)\theta_j^{-1}) + \varepsilon \mathbf{s}(b)^2 \left(g_{\eta}(\theta_j) - \theta_j^{-1} \right) R'_{\mu}(\mathbf{s}(b)\theta_j^{-1}) + \mathcal{O}(\varepsilon^2) \\ &= \theta_j + \mathbf{s}(b) R_{\mu}(\mathbf{s}(b)\theta_j^{-1}) + \varepsilon \mathbf{s}(b)^2 d_{\eta}(\theta_j) R'_{\mu}(\mathbf{s}(b)\theta_j^{-1}) + \mathcal{O}(\varepsilon^2) \end{aligned} \tag{2.29}$$

where we have defined $d_{\eta}(z) = g_{\eta}(\theta_j) - \theta_j^{-1}$. The argument with the lower outliers $\{\theta'_j\}_{j=1}^q$ is identical.

The problem of determining the support of $\mu_b \boxplus \nu$ is difficult and almost certainly analytically intractable, with [BES20] containing the most advanced results in that direction. However

overall, we have a model for DNN Hessians with a spectrum consisting, with high-probability, of a compactly supported bulk $\mu_b \boxplus \nu$ and a set of outliers given by (2.29) (and similarly for θ'_j) subject to (2.17). The constants $L_\varepsilon, U_\varepsilon$ in the statement (2.12) and (2.13) of the theorem are simply the lower and upper edges of the support of $\text{supp}(\mu_b \boxplus \nu)$. □

Note that (2.29) reduces to outliers of the form $\theta_j + \mathbf{s}(b)^2 R_\mu(\theta_j^{-1})$ if $\varepsilon = 0$ or $d_\eta = 0$, as expected from [BGN11]⁵.

(2.29) is a generalised form of the result used in [GZR20]. We have the power series

$$R_\mu(\mathbf{s}(b)\theta_j^{-1}) = k_1^{(\mu)} + \frac{k_2^{(\mu)} \mathbf{s}(b)}{\theta_j} + \frac{k_3^{(\mu)} \mathbf{s}(b)^2}{\theta_j^2} + \dots, \quad (2.30)$$

$$d_\eta(\theta_j) = \frac{m_1^{(\eta)}}{\theta_j^2} + \frac{m_2^{(\eta)}}{\theta_j^3} + \dots \quad (2.31)$$

where $m_n^{(\eta)}$ are the moments of η and $k_n^{(\mu)}$ are the cumulants of μ . In the case that the spikes θ_j are large enough, we approximate by truncating these power series to give

$$\omega^{-1}(\theta_j) \approx \theta_j + \mathbf{s}(b)m_1^{(\mu)} + \mathbf{s}(b)^2 k_2^{(\mu)} \left(\frac{1}{\theta_j} + \frac{\varepsilon m_1^{(\eta)}}{\theta_j^2} \right) \quad (2.32)$$

where the approximation is more precise for larger b and smaller ε and we have used the fact that the first cumulant of any measure matches the first moment. One could consider for instance a power law for $\mathbf{s}(b)$, i.e.

$$\omega^{-1}(\theta_j) \approx \theta_j + \frac{k_1^{(\mu)}}{b^v} + \frac{k_2^{(\mu)}}{b^{2v}} \left(\frac{1}{\theta_j} + \frac{\varepsilon m_1^{(\eta)}}{\theta_j^2} \right) = \theta_j + \frac{m_1^{(\mu)}}{b^v} + \frac{k_2^{(\mu)}}{b^{2v}} \left(\frac{1}{\theta_j} + \frac{\varepsilon m_1^{(\eta)}}{\theta_j^2} \right) \quad (2.33)$$

for some $v > 0$. In the case that μ is a semicircle, then all cumulants apart from the second vanish, so setting $\varepsilon = 0$ recovers *exactly*

$$\omega^{-1}(\theta_j) = \theta_j + \frac{\sigma^2}{4b^{2v}\theta_j} \quad (2.34)$$

where σ is the radius of the semicircle. To make the link with [GZR20] obvious, we can take $v = 1/2$ and μ to be the semicircle, so giving

$$\omega^{-1}(\theta_j) \approx \theta_j + \frac{\sigma^2}{4b\theta_j} \quad (2.35)$$

where we have truncated $\mathcal{O}(\varepsilon)$ term. We present an argument in favour of the $v = 1/2$ power law below, but we allow for general v when comparing to experimental data.

Remark 2.2. It is quite possible for μ 's density to have a sharp spike at the origin, or even for μ to contain a δ atom at 0, as observed empirically in the spectra of DNN Hessians.

2.4. Experimental results

The random matrix Hessian model introduced above is quite general and abstract. Necessarily the measures μ and η must be allowed to be quite general as it is well established experimentally [Pap18, Gra20, BGK22] that real-world DNN Hessians have spectral bulks

⁵ Note that $d_\eta = 0 \iff \eta = \delta_0$ which is clearly equivalent (in terms of ν) to $\varepsilon = 0$.

that are not familiar as being any standard canonical examples from random matrix theory. That being said, the approximate form in (2.33) gives quite a specific form for the Hessian outliers. In particular, the constants $m_1^{(\mu)}, m_1^{(\eta)}$ and $m_2^{(\mu)}, \varepsilon > 0$ are shared between all outliers at all batch sizes. If the form of the Hessian outliers seen in (2.33) is not observed experimentally, it would suggest at least one of the following does not hold:

- (a) batch sampling induces a simple multiplicative scaling on the Hessian noise (2.7);
- (b) the true Hessian is approximately low-rank (as measured by ε) and has a finite number of outliers;
- (c) the Hessian noise model X has QUE.

In view of this third point, agreement with (2.33) provides an indirect test for the presence of universal random matrix statistics in DNN Hessians.

We can use Lanczos power methods [MS06] to compute good approximations to the top few outliers in the batch Hessian spectra of DNNs [GZR20]. Indeed the so-called Pearlmutter trick [Pea94] enables efficient numerical computation of Hessian-vector products, which is all that one requires for power methods. Over a range of batch sizes, we compute the top five outliers of the batch Hessian for ten different batch seeds. We repeat this procedure at every 25 epochs throughout the training of two standard DNNs for computer vision tasks, VGG16 (named after the Visual Geometry Group at Oxford) and WideResNet28 \times 10, on the CIFAR100 dataset [KH+09] and at every epoch during the training of a simple multi-layer perceptron (MLP) network on the MNIST dataset [LBBH98]. By the end of training each of the models have high test accuracy, specifically the VGG16 architecture which does not use batch normalisation, has a test accuracy of $\approx 75\%$, whereas the WideResNet28 \times 10 has a test accuracy of $\approx 80\%$. The MLP has a test set accuracy of $\approx 95\%$. Full experimental details are given in appendix B.

Remark 2.3. There is a subtlety with regard to obtaining the top outliers using the Lanczos power method. Indeed, since Lanczos provides, in some sense, an approximation to the whole spectrum of a matrix, truncating at m iterations for a $N \times N$ matrix cannot produce good approximations to all of the m top eigenvalues. In reality, experimental results [Pap18, GWG19] show that, for DNNs, and using sufficiently many iterations (m), the top r eigenvalues may be recovered, for $r \ll m$. We display some spectral plots of the full Lanczos results in the appendix C which demonstrate clearly a large number of outliers, and clearly more than 5. As a result, we can have confidence that our numerical procedure is indeed recovering approximations to the top few eigenvalues required for our experiments.

Let $\lambda_b^{(i,j,e)}$ be the top i th empirical outlier (so $i = 1$ is the top outlier) for the j th batch seed and a batch size of b for the model at epoch e . To compare the experimental results to our theoretical model, we propose the following form:

$$\lambda_b^{(i,j,e)} \approx \theta^{(i,e)} + \frac{\alpha^{(e)}}{bv} + \frac{\beta^{(e)}}{b^2v} \left(\frac{1}{\theta^{(i,e)}} + \frac{\gamma^{(e)}}{(\theta^{(i,e)})^2} \right) \quad (2.36)$$

where $\beta^{(e)} > 0$ (as the second cumulant of a any measure of non-negative) and $\theta^{(i,e)} > \theta^{(i+1,e)} > 0$ for all i, e . The parameters $\alpha^{(e)}, \beta^{(e)}, \gamma^{(e)}$ and $\theta^{(i,e)}$ need to be fit to the data, which could be done with standard black-box optimisation to minimise squared error in (2.36), however we propose an alternative approach which reduces the number of free parameters and hence should regularise the optimisation problem. Observe that (2.36) is linear in the parameters $\alpha^{(e)}, \beta^{(e)}, \gamma^{(e)}$ so, neglecting the positivity constraint on $\beta^{(e)}$, we can in fact solve exactly for optimal values. Firstly let us define $\bar{\lambda}_b^{(i,e)}$ to be the empirical mean of $\lambda_b^{(i,j,e)}$ over the batch

seed index j . Each epoch will be treated entirely separately, so let us drop the e superscripts to streamline the notation. We are then seeking to optimise $\alpha, \beta, \gamma, \theta^{(i)}$ to minimise

$$E = \sum_{i,b} \left(\bar{\lambda}_b^{(i)} - \theta^{(i)} - \frac{\alpha}{b^v} - \frac{\beta}{\theta^{(i)} b^{2v}} - \frac{\beta\gamma}{b^{2v} (\theta^{(i)})^2} \right)^2. \quad (2.37)$$

Now make the following definitions

$$y_{ib} = \bar{\lambda}_b^{(i)} - \theta^{(i)}, \quad \mathbf{x}_{ib} = \begin{pmatrix} b^{-v} \\ (\theta^{(i)} b)^{-2v} \\ (b^{2v} (\theta^{(i)})^2)^{-1} \end{pmatrix}, \quad \mathbf{w} = \begin{pmatrix} \alpha \\ \beta \\ \beta\gamma \end{pmatrix}, \quad (2.38)$$

so that

$$E = \sum_{i,b} (y_{ib} - \mathbf{w}^T \mathbf{x}_{ib})^2. \quad (2.39)$$

Finally we can define the n -dimensional vector \mathbf{Y} by flattening the matrix $(y_{ib})_{ib}$, and the $3 \times n$ matrix \mathbf{X} by stacking the vectors \mathbf{x}_{ib} and then flattening of the i, b indices. That done, we have have a standard linear regression problem with design matrix \mathbf{X} and parameters \mathbf{w} . For fixed θ , the global minimum of E is then attained at parameters

$$\mathbf{w}^*(\theta) = (\mathbf{X}\mathbf{X}^T)^{-1} \mathbf{X}\mathbf{Y} \quad (2.40)$$

where the dependence on the parameters θ is through \mathbf{Y} and \mathbf{X} as above. We thus have

$$\alpha = w_1^*, \beta = w_2^*, \gamma = w_3^*/w_2^*$$

and can plug these values back in to (2.37) to obtain an optimisation problem only over the $\theta^{(i)}$. There is no closed form solution for the optimal $\theta^{(i)}$ for this problem, so we fit them using gradient descent. The various settings and hyperparameters of this optimisation were tuned by hand to give convergence and are detailed in appendix B.3. To address the real constraint $\beta > 0$, we add a penalty term to the loss (2.37) which penalises values of $\theta^{(i)}$ leading to negative values of β . The constraint $\theta^{(i)} > \theta^{(i+1)} > 0$ is implemented using a simple differentiable transformation detailed in appendix B.2. Finally, the exponent v is selected by fitting the parameters for each v in $\{-0.1, -0.2, \dots, -0.9\}$ and taking the value with the minimum mean squared error E .

The above process results in 12 fits for VGG and Resnet and ten for MLP (one per epoch). For each of these, we have a theoretical fit for each of the 5 top outliers as a function of batch size which can be compared graphically to the data, resulting in $(2 \times 12 + 10) \times 5 = 170$ plots. Rather than try to display them all, we will select a small subset that illustrates the key features. Figure 1 shows results for the Resnet at epochs 0 (initialisation), 25, 250 and 300 (end of training) and outliers 1, 3 and 5. Between the three models, the Resnet shows consistently the best agreement between the data and the parametric form (2.36). The agreement is excellent at epoch 0 but quickly degrades to that seen in the second row of figure 1, which is representative of the early and middle epochs for the Resnet. Towards the end of training the Resnet returns to good agreement between theory and data, as demonstrated in the third and fourth rows of figure 1 at epochs 250 and 300 respectively.

The VGG16 also has excellent agreement between theory and data at epoch 0, and thereafter is similar to the early epochs of the Resnet, i.e. reasonable, but not excellent, until around epoch 225 where the agreement starts to degrade significantly until the almost complete failure at epoch 300 shown in the first row of figure 2. The MLP has the worst agreement between theory and data, having again excellent agreement at epoch 0, but really quite poor agreement even by epoch 1, as shown in the second row of figure 2.

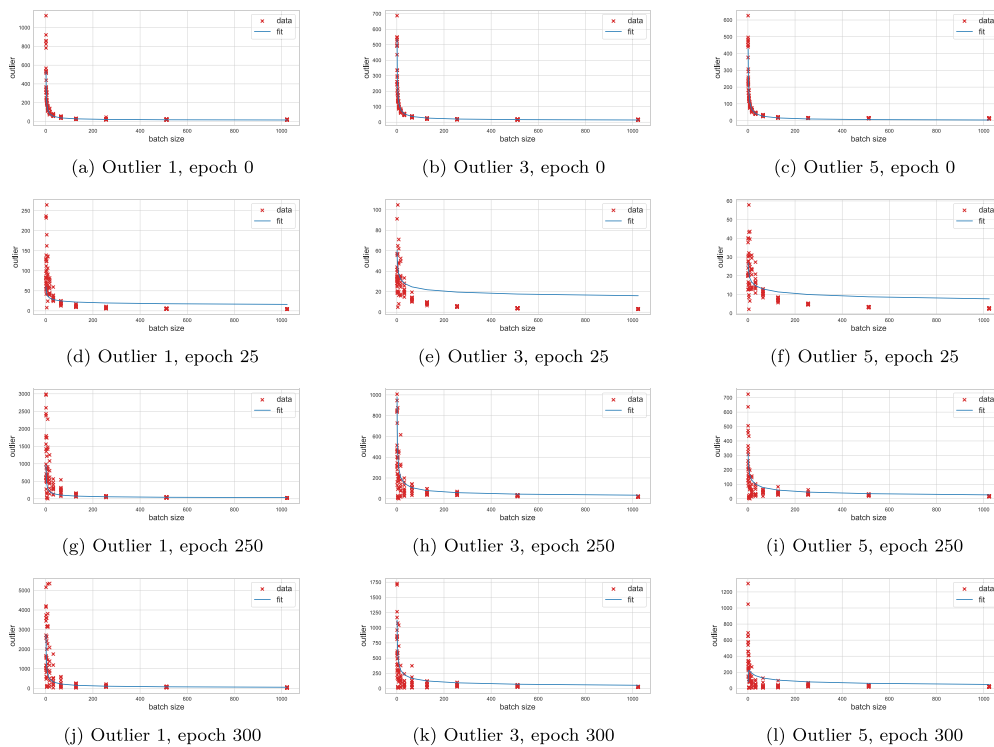


Figure 1. The batch-size scaling of the outliers in the spectra of the Hessians of the Resnet loss on CIFAR100. Training epochs increase top-to-bottom from initialisation to final trained model. Left-to-right the outlier index varies (outlier 1 being the largest). Red cross show results from Lanczos approximations over ten samples (different batches) for each batch size. The blue lines are parametric power law fits of the form (2.36).

The experimental results show an ordering Resnet > VGG > MLP, in terms of how well the random matrix theory loss surface predictions explain the Hessian outliers. We conjecture that this relates to the difficulty of the loss surfaces. Resnets are generally believed to have smoother, simpler loss surfaces [LXT+18] and be easier to train than other architectures, indeed the residual connections were originally introduced for precisely this reason. The VGG is generally more sensitive to training set-up, requiring well-tuned hyperparameters to avoid unstable or unsuccessful training [GB22]. The MLP is perhaps too small to benefit from high-dimensional highly over-parametrised effects.

The parameter values obtained for all models over all epochs are shown in figure 3, with a column for each model. There are several interesting features to draw out of these plots, however note that we cannot meaningfully interpret the parameters for the MLP beyond epoch 0, as the agreement with (2.36) is so poor. Firstly consider the parameter $m_1^{(\mu)}$, which is interpreted as the first moment (i.e. mean) of the spectral density of the noise matrix X . $m_1^{(\mu)} = 0$ is significant, as it is seen in the case of the a symmetric measure μ , such as the Wigner semicircle used by [GZR20]. For the VGG, $m_1^{(\mu)}$ starts close to 0 (figure 3(b)) and generally grows with training epochs (note that the right hand side of this plot is not trustworthy, as we have observed that the agreement with (2.36) does not survive to the end of training). For the Resnet, we see a similar upwards trend (figure 3(a)), with the notable exception that of initialisation




Figure 2. Left-to-right the outlier index varies (outlier 1 being the largest). Red cross show results from Lanczos approximations over ten samples (different batches) for each batch size. The blue lines are parametric power law fits of the form (2.36). This plot show the final epoch (300) for the VGG16 on CIFAR100 and the first epoch for the MLP on MNIST, both being examples of the parametric fit failing to match the data.

(epoch 0). These two observations together, suggest that training encourages a skew in the spectrum of X away from symmetry around 0, however for some structural reason the Resnet is highly skewed at initialisation.

Note that for all models this parameter starts close to 0 and generally grows with training epochs, noting that the right hand side of figure 3(b) at the higher epochs should be ignored owing to the bad fit discussed above.

It is interesting also to observe that $\varepsilon m_1^{(\eta)}$ remains small for all epochs particularly compared to $m_1^{(\mu)}, k_2^{(\mu)}$. This is consistent with the derivation of (2.36), which relies on ε being small, however we emphasise that *this was not imposed as a numerical constraint* but arises naturally from the data. Recall that the magnitude of $\varepsilon m_1^{(\eta)}$ measures the extent of the deviation of A from being exactly low rank, so its small but non-zero values suggest that it is indeed important to allow for the true Hessian to have non-zero rank in the $N \rightarrow \infty$ limit. Finally, we comment that the best exponent is generally not $\nu = 1/2$. Again, the results from the Resnet are the most reliable and they appear to show that the batch scaling, as characterised by ν , is not constant throughout training, particularly comparing epoch 0 and epoch 300, say.

2.5. Justification and motivation of QUE

Work on random matrix universality has shown that a *local law* is the key ingredient in the establishing universal local statistics of eigenvalues [EYY12, EY17] and universal delocalisation of eigenvectors [BY17]. There are several forms of local law, but all provide high probability control on the error between the (random) matrix Green’s function $G(z) = (z - X)^{-1}$ and certain deterministic equivalents. In all cases we use the set

$$\mathcal{S} = \{E + i\eta \in \mathbb{C} \mid |E| \leq \omega^{-1}, N^{-1+\omega} \leq \eta \leq \omega^{-1}\} \tag{2.41}$$

for $\omega \in (0, 1)$ and the local law statements holds for all (large) $D > 0$ and (small) $\xi > 0$ and for all large enough N . The *averaged local law* states:

$$\sup_{z \in \mathcal{S}} \mathbb{P} \left(\left| \frac{1}{N} \text{Tr} G(z) - g_\mu(z) \right| > N^\xi \left(\frac{1}{N\eta} + \sqrt{\frac{\Im g_\mu(z)}{N\eta}} \right) \right) \leq N^{-D}. \tag{2.42}$$



Figure 3. The parameter values produced when fitting experimental neural network Hessian outlier data to (2.36).

The isotropic local law states:

$$\sup_{\|\mathbf{u}\|, \|\mathbf{v}\|=1, z \in \mathcal{S}} \mathbb{P} \left(|\mathbf{u}^T G(z) \mathbf{v} - g_\mu(z)| > N^\xi \left(\frac{1}{N\eta} + \sqrt{\frac{\Im g_\mu(z)}{N\eta}} \right) \right) \leq N^{-D}. \quad (2.43)$$

The *anisotropic local law* states:

$$\sup_{\|\mathbf{u}\|, \|\mathbf{v}\|=1, z \in \mathcal{S}} \mathbb{P} \left(|\mathbf{u}^T G(z) \mathbf{v} - \mathbf{u}^T \Pi(z) \mathbf{v}| > N^\xi \left(\frac{1}{N\eta} + \sqrt{\frac{\Im g_\mu(z)}{N\eta}} \right) \right) \leq N^{-D} \quad (2.44)$$

where $\Pi(\cdot)$ is an $N \times N$ deterministic matrix function on \mathbb{C} . The *entrywise local law* states:

$$\sup_{z \in \mathcal{S}, 1 \leq i, j \leq N} \mathbb{P} \left(|G_{ij}(z) - \Pi_{ij}(z)| > N^\xi \left(\frac{1}{N\eta} + \sqrt{\frac{\Im g_\mu(z)}{N\eta}} \right) \right) \leq N^{-D}. \quad (2.45)$$

The anisotropic local law is a stronger version of the entrywise local law. The anisotropic local law is a more general version of the isotropic local law, which can be recovered in the isotropic case by taking $\Pi = g_\mu I$. The entrywise local law can also be applied in the isotropic case by taking $\Pi = g_\mu I$. The averaged local law is weaker than all of the other laws. General Wigner matrices are known to obey isotropic local semi-circle laws [ES17]. Anisotropic local laws are known for general deformations of Wigner matrices and general covariance matrices [KY17] as well as quite general classes of correlated random matrices [EKS19].

As mentioned above, QUE was proved for general Wigner matrices in [BY17]. It appears that the key ingredient in the proof of QUE (2.9) in [BY17] is the isotropic local semicircle law (2.43) for general Wigner matrices. Indeed, all the intermediate results in Sections 4 of [BY17] take only (2.43) and general facts about the Dyson Brownian Motion eigenvector flow given by

$$d\lambda_k = \frac{dB_{kk}}{\sqrt{N}} + \left(\frac{1}{N} \sum_{\ell \neq k} \frac{1}{\lambda_k - \lambda_\ell} \right) dt, \quad (2.46)$$

$$du_k = \frac{1}{\sqrt{N}} \sum_{\ell \neq k} \frac{dB_{k\ell}}{\lambda_k - \lambda_\ell} u_\ell - \frac{1}{2N} \sum_{\ell \neq k} \frac{dt}{(\lambda_k - \lambda_\ell)^2} u_k. \quad (2.47)$$

This can be generalised to

$$d\lambda_k = \frac{dB_{kk}}{\sqrt{N}} + \left(-V(\lambda_k) + \frac{1}{N} \sum_{\ell \neq k} \frac{1}{\lambda_k - \lambda_\ell} \right) dt, \quad (2.48)$$

$$du_k = \frac{1}{\sqrt{N}} \sum_{\ell \neq k} \frac{dB_{k\ell}}{\lambda_k - \lambda_\ell} u_\ell - \frac{1}{2N} \sum_{\ell \neq k} \frac{dt}{(\lambda_k - \lambda_\ell)^2} u_k. \quad (2.49)$$

where V is a potential function. Note that the eigenvector dynamics are unaffected by the presence of the potential V , so we expect to be able to generalise the proof of [KY17] to any random matrix ensemble with an isotropic local law by defining the potential V so that the invariant ensemble with distribution $Z^{-1} e^{-N\text{Tr}V(X)} dX$ has equilibrium measure μ (Z is a normalisation constant). We show how to construct such a V from μ in appendix A.

The arguments so far suffice to justify a generalisation of the ‘dynamical step’ in the arguments of [BY17], so it remains to consider the ‘comparison step’. The dynamical step establishes QUE for the matrix ensemble with a small Gaussian perturbation, but in the comparison step one must establish that the perturbation can be removed without breaking QUE. To our knowledge no such argument has been articulated beyond generalized Wigner matrices, with the independence of entries and comparable scale of variances being critical to the arguments given by [BY17]. Our guiding intuition is that QUE of the form (2.9) is a general property of random matrices and can reasonably be expected to hold in most, if not all, cases in which there is a local law and universal local eigenvalue statistics are observed. At present, we are

not able to state a precise result establishing QUE in sufficient generality to be relevant for this work, so we shall take it as an assumption.

Assumption 2.4. *Let X be an ensemble of $N \times N$ real symmetric random matrices. Assume that X admits a limiting spectral measure μ with Stieljtes transform m . Suppose that the isotropic local law (2.43) holds for X with μ . Then there is some set $\mathbb{T}_N \subset [N]$ with $|\mathbb{T}_N^c| = o(N)$ such that with $|I| = n$, for any polynomial P in n indeterminates, there exists some $\varepsilon(P) > 0$ such that for large enough N we have*

$$\sup_{\substack{I \subset \mathbb{T}_N, |I|=n, \\ \|\mathbf{q}\|=1}} \left| \mathbb{E} \left(P \left((N(\mathbf{q}^T \mathbf{u}_k)^2)_{k \in I} \right) \right) - \mathbb{E} \left(P \left((|\mathcal{N}_j|^2)_{k \in I} \right) \right) \right| \leq N^{-\varepsilon}. \quad (2.50)$$

Note that the isotropic local law in assumption 2.4 can be obtained from the weaker entrywise law (2.45) as in theorem 2.14 of [BEK+14] provided there exists a $C > 0$ such that $\mathbb{E}|X_{ij}|^2 \leq CN^{-1}$ for all i, j and there exists $C_p > 0$ such that $\mathbb{E}|\sqrt{N}X_{ij}|^p \leq C_p$ for all i, j and integer $p > 0$.

Remark 2.5. In [BY17] the restriction $I \subset \mathbb{T}_N$ is given for the explicit set

$$\mathbb{T}_N = [N] \setminus \{(N^{1/4}, N^{1-\delta}) \cup (N - N^{1-\delta}, N - N^{1/4})\} \quad (2.51)$$

for some $0 < \delta < 1$. In the case of generalised Wigner matrices, this restriction on the indices has since been shown to be unnecessary [BL22, Ben20, BL21]. In our context, we could simply take as an assumption all results holds with $\mathbb{T}_N = [N]$, however our results can in fact be proved using only the above assumption that $|\mathbb{T}_N^c| = o(N)$, so we shall retain this weaker form of the assumptions.

This section is not intended to prove QUE from explicit known properties of DNN Hessians, but rather to provide justification for it as a reasonable modeling assumption in the noise model for Hessians defined in section 2.1. We have shown how QUE can be obtained from an isotropic (or entrywise) local law beyond the Wigner case. It is important to go beyond Wigner or any other standard random matrix ensemble, as we have observed above that the standard macroscopic spectral densities of random matrix theory such as the semicircle law are not observed in practice. That said, we are not aware of any results establishing QUE in the more general case of anisotropic local laws, and this appears to be a very significant technical challenge. We must finally address why a local law assumption, isotropic or otherwise, may be reasonable for the noise matrix X in our Hessian model. Over the last decade or so, universal local statistics of random matrices in the form of k -point correlation functions on the appropriate microscopic scale have been established for a litany of random matrix ensembles. An immediate consequence of such results is that, on the scale of unit mean eigenvalue spacing, Wigner’s surmise holds to a very good approximation, depending only on the symmetry class (orthogonal, unitary or symplectic). Such universality results are rather older for invariant ensembles [Dei99, EY17] and can be established with orthogonal polynomial techniques, however the recent progress focusing on non-invariant ensembles, beginning with Wigner matrices [EYY12] and proceeding to much more general ensembles [EKS19], is built on a very general ‘three step strategy’ (though see [EY12] for connections between universality in invariant and non-invariant ensembles). As with the QUE proof discussed above, the key ingredient in these proofs, as part of the three step strategy [EY17], is establishing a local law. The theoretical picture that has emerged is that, for very general random matrices, when universal local eigenvalue statistics are observed in random matrices, it is due to the mechanism of short time scale relaxation of local statistics under Dyson Brownian Motion made possible by a local law.

It has been observed that universal local eigenvalue statistics do indeed appear to be present in the Hessian of real, albeit quite small, DNNs [BGK22]. Given all of this context, we propose

that a local law assumption of some kind is reasonable for DNN Hessians and not particularly restrictive. As we have shown, if we are willing to make the genuinely restrictive assumption of an isotropic local law for the Hessian noise model, then QUE follows. However an anisotropic local law is arguably more plausible as we expect DNNs Hessians to contain a good deal of dependence between entries, and such correlations are known to generically lead to anisotropic local laws [EKS19].

2.6. Motivation of true Hessian structure

In this section we revisit and motivate the assumptions made about the Hessian in section 2.1. Firstly note that one can always define $A = \mathbb{E}H_{\text{batch}}$ and it is natural then to associate A with the true Hessian H_{true} . In light of (2.1), it is natural to expect some fixed form of the law for $H_{\text{batch}} - A$ for any batch size, but with an overall scaling $\mathfrak{s}(b)$, which must naturally be decreasing in b as experimental results show that the overall spectral width of the batch Hessians of neural networks decreases with increasing batch size. Next we address the assumptions made about the spectrum of A . The first assumption one might think to make is that A has fixed rank relative to N , with spectrum consisting only of the spikes θ_i, θ'_j . Indeed, it has been repeatedly observed, in our own experiments and others [Pap18, GZR20], that neural network Hessians contain a number of spectral outliers separated from the spectral bulk. It is natural to conjecture that such outliers arise from some outliers in an underlying structured deterministic matrix of which the batch Hessian is a noisy version, as in the case of BBP style phase transitions in random matrix theory. The outliers in neural network Hessians have been associated with inter-class separation in the case of classification models [Pap19] and it can be observed that spectra lack (or have smaller and fewer) outliers at the start of training, or if they are intentionally trained to give poor (i.e. random) predictive performance. That being said, in almost any experiment with sensibly trained neural networks, spectral outliers are observed, and over a range of batch sizes (and hence noise levels) suggesting that some of the spike eigenvalues in the true Hessian are above the phase transition threshold.

Behind such an assumption is the intuition that the data distribution does not depend on N and so, in the over-parametrised limit $N \rightarrow \infty$, the overwhelming majority of directions in weight space are unimportant. The form we take for A in the above is a strict generalisation of the fixed rank assumption; A still has a fixed number of spiked directions, but the parameter ε controls the rank of A . Since any experimental investigation is necessarily limited to $N < \infty$, the generalisation to $\varepsilon > 0$ is particularly important. Compact support of the measures μ and η is consistent with experimental observations of DNN Hessian spectra.

2.7. The batch size scaling

Our experimental results considered $\mathfrak{s}(b) = b^{-\nu}$ and $\nu = 1/2$ is the value required to give agreement with [GZR20], a choice which we now justify. From (2.1) we have

$$H_{\text{batch}} = \frac{1}{b} \sum_{i=1}^b (H_{\text{true}} + X^{(i)}) \tag{2.52}$$

where $X^{(i)}$ are i.i.d. samples from the law of X . Suppose that the entries X_{ij} were Gaussian, with $\text{Cov}(X_{ij}, X_{kl}) = \Sigma_{ij,kl}$. Then $Z = X_{ij}^{(p)} + X_{ij}^{(q)}$ has

$$\text{Cov}(Z_{ij}, Z_{kl}) = \mathbb{E}X_{ij}^{(p)} X_{kl}^{(p)} + \mathbb{E}X_{ij}^{(q)} X_{kl}^{(q)} - \mathbb{E}X_{ij}^{(p)} \mathbb{E}X_{kl}^{(p)} - \mathbb{E}X_{ij}^{(q)} \mathbb{E}X_{kl}^{(q)} = 2\Sigma_{ij,kl}. \tag{2.53}$$

In the case of centred X , one then obtains

$$\frac{1}{b} \sum_{i=1}^b X^{(i)} \stackrel{d}{=} b^{-1/2} X. \tag{2.54}$$

Note that this does not quite match the case described in section 2.1, since we do not assume there that $\mathbb{E}X = 0$, however we take this a rough justification for $\mathfrak{s}(b) = b^{-1/2}$ as an ansatz. Moreover, numerical experimentation with $\mathfrak{s}(b) = b^{-\nu}$ for values of $\nu > 0$ shows that $q = 1/2$ gives a reasonable fit to the data (note that the values shown in figures 3(j)–(l) are those producing the best fit, but $\nu = 1/2$ was seen to be not much inferior).

3. Spectral free addition from QUE

3.1. Intermediate results on QUE

This section establishes some intermediate results that follow from assuming QUE for the eigenvectors of a matrix. They will be crucial for our application in the following section.

Lemma 3.1. *Consider a real orthogonal $N \times N$ matrix U with rows $\{\mathbf{u}_i^T\}_{i=1}^N$. Assume that $\{\mathbf{u}_i\}_{i=1}^N$ are the eigenvectors of a real random symmetric matrix with QUE. Let P be a fixed $N \times N$ real orthogonal matrix. Let $V = UP$ and denote the rows of V by $\{\mathbf{v}_i^T\}_{i=1}^N$. Then $\{\mathbf{v}_i\}_{i=1}^N$ also satisfy QUE.*

Proof. Take any unit vector \mathbf{q} , then for any $k = 1, \dots, N$

$$\mathbf{q}^T \mathbf{v}_k = \sum_j q_j V_{kj} = \sum_{j,l} q_j U_{kl} P_{lj} = (P\mathbf{q})^T \mathbf{u}_k.$$

But $\|P\mathbf{q}\|_2 = \|\mathbf{q}\|_2 = 1$ since P is orthogonal, so the statement of QUE for $\{\mathbf{u}_i\}_{i=1}^N$ transfers directly to $\{\mathbf{v}_i\}_{i=1}^N$ thanks to the supremum of all unit \mathbf{q} . \square

Lemma 3.2. *Consider a real orthogonal $N \times N$ matrix U with rows $\{\mathbf{u}_i^T\}_{i=1}^N$. Assume that $\{\mathbf{u}_i\}_{i=1}^N$ are the eigenvectors of a real random symmetric matrix with QUE. Let $\ell_0(\mathbf{q}) = \sum_i \mathbb{1}\{q_i \neq 0\}$ count the non-zero elements of a vector with respect to a fixed orthonormal basis $\{\mathbf{e}_i\}_{i=1}^N$. For any fixed integer $s > 0$, define the set*

$$\mathbb{V}_s = \{\mathbf{q} \in \mathbb{R}^N \mid \|\mathbf{q}\| = 1, \ell_0(\mathbf{q}) = s, q_i = 0 \forall i \in \mathbb{T}_N^c\} \tag{3.1}$$

where, recall the definition

$$\mathbb{T}_N = [N] \setminus \{(N^{1/4}, N^{1-\delta}) \cup (N - N^{1-\delta}, N - N^{1/4})\}.$$

Then the columns $\{\mathbf{u}_i^T\}_{i=1}^N$ of U satisfy a weaker form of QUE (for any fixed $n, s > 0$):

$$\sup_{\substack{\mathbf{q} \in \mathbb{V}_s \\ I \subset \mathbb{T}_N \\ |I|=n}} \sup_{|J|=n} \left| \mathbb{E}P \left((N|\mathbf{q}^T \mathbf{u}_k|^2)_{k \in I} \right) - \mathbb{E}P \left((|\mathcal{N}_j|^2)_{j=1}^m \right) \right| \leq N^{-\varepsilon}. \tag{3.2}$$

We will denote this form of QUE as \widehat{QUE} .

Proof. Take some $\mathbf{q} \in \mathbb{V}_s$. Then there exists some $J \subset \mathbb{T}_N$ with $|J| = s$ and non-zero $\{q_k\}_{k \in J}$ such that

$$\mathbf{q}^T \mathbf{u}'_k = \sum_{j \in J} q_j \mathbf{e}_j^T \mathbf{u}'_k.$$

Take $\{e_i\}_{i=1}^N$ to be a standard basis with $(e_i)_j = \delta_{ij}$, then $e_j^T u'_k = U_{jk} = e_k^T u_j$ so

$$q^T u'_k = \sum_{j \in J} q_j e_k^T u_j$$

but then the coefficients q_j can be absorbed into the definition of the general polynomial in the statement (2.9) of QUE for $\{u_i\}_{i=1}^N$, which completes the proof, noting that the sum only includes indices contained in \mathbb{T}_N owing to the definition of \mathbb{V}_s . \square

Lemma 3.3. Fix some real numbers $\{y_i\}_{i=1}^r$. Fix also a diagonal matrix Λ and an orthonormal set of vectors $\{v_i\}_{i=1}^r$ that satisfies QUE. Then there exists an $\varepsilon > 0$ and $\eta_i \in \mathbb{C}^N$ with

$$\eta_{ij}^2 \in [-1, 1] \quad \forall j \in \mathbb{T}_N, \tag{3.3}$$

$$\eta_{ij}^2 \in [-N^\varepsilon, N^\varepsilon] \quad \forall j \in \mathbb{T}_N^c. \tag{3.4}$$

such that for any integer $l > 0$

$$\mathbb{E} \left(\sum_{i=1}^r y_i v_i^T \Lambda v_i \right)^l - \mathbb{E} \left(\sum_{i=1}^r y_i \frac{1}{N} g_i^T \Lambda g_i \right)^l = N^{-(1+\varepsilon)l} \left(\sum_{i=1}^r y_i \eta_i^T \Lambda \eta_i \right)^l \tag{3.5}$$

where the g_i are i.i.d. Gaussians $N(0, I_N)$.

Proof. Let $\{e_i\}_{i=1}^N$ be the standard orthonormal basis from above. Then

$$\begin{aligned} \mathbb{E} \left(\sum_{i=1}^r y_i v_i^T \Lambda v_i \right)^l &= \mathbb{E} \sum_{i_1, \dots, i_l=1}^r \prod_{k=1}^l y_{i_k} v_{i_k}^T \Lambda v_{i_k} \\ &= \mathbb{E} \sum_{i_1, \dots, i_l=1}^r \sum_{j_1, \dots, j_l=1}^N \prod_{k=1}^l y_{i_k} \lambda_{j_k} (e_{j_k}^T v_{i_k})^2 \end{aligned} \tag{3.6}$$

$$\begin{aligned} \implies & \mathbb{E} \left(\sum_{i=1}^r y_i v_i^T \Lambda v_i \right)^l - \mathbb{E} \left(\sum_{i=1}^r y_i \frac{1}{N} g_i^T \Lambda g_i \right)^l \\ &= N^{-l} \sum_{i_1, \dots, i_l=1}^r \sum_{j_1, \dots, j_l=1}^N \prod_{k=1}^l y_{i_k} \lambda_{j_k} [N \mathbb{E} (e_{j_k}^T v_{i_k})^2 - \mathbb{E} (e_{j_k}^T g_{i_k})^2] \\ &= N^{-l} \sum_{i_1, \dots, i_l=1}^r \sum_{j_1, \dots, j_l \in \mathbb{T}_N} \prod_{k=1}^l y_{i_k} \lambda_{j_k} [N \mathbb{E} (e_{j_k}^T v_{i_k})^2 - \mathbb{E} (e_{j_k}^T g_{i_k})^2] \\ & \quad + N^{-l} \sum_{i_1, \dots, i_l=1}^r \sum_{\substack{j_1 \in \mathbb{T}_N^c, \\ j_2, \dots, j_l \in \mathbb{T}_N}} \prod_{k=1}^l y_{i_k} \lambda_{j_k} [N \mathbb{E} (e_{j_k}^T v_{i_k})^2 - \mathbb{E} (e_{j_k}^T g_{i_k})^2] \\ & \quad + \dots \end{aligned} \tag{3.7}$$

The ellipsis represents the similar terms where further of the j_1, \dots, j_r are in \mathbb{T}_N^c . For $j \in \mathbb{T}_N^c$ the terms

$$[N \mathbb{E} (e_j^T v_{i_k})^2 - \mathbb{E} (e_j^T g_{i_k})^2] \tag{3.8}$$

are excluded from the statement of $\widehat{\text{QUÉ}}$, however we can still bound them crudely. Indeed

$$\sum_{j \in \mathbb{T}_N^c} N(\mathbf{e}_j^T \mathbf{v}_i)^2 = \sum_{j=1}^N N(\mathbf{e}_j^T \mathbf{v}_i)^2 - \sum_{j \in \mathbb{T}_N} N(\mathbf{e}_j^T \mathbf{v}_i)^2 = N - \sum_{j \in \mathbb{T}_N} N(\mathbf{e}_j^T \mathbf{v}_i)^2 \quad (3.9)$$

but since the bound of $\widehat{\text{QUÉ}}$ applies for $j \in \mathbb{T}_N$

$$N\mathbb{E}(\mathbf{e}_j^T \mathbf{v}_i)^2 = \mathbb{E}(\mathbf{e}_j^T \mathbf{g})^2 + o(1) = 1 + o(1) \quad \forall j \in \mathbb{T}_N, \quad (3.10)$$

then

$$\sum_{j \in \mathbb{T}_N^c} N(\mathbf{e}_j^T \mathbf{v}_i)^2 = N - N(1 + o(1)) = o(N) \implies \mathbb{E}(\mathbf{e}_j^T \mathbf{v}_i)^2 = o(1) \quad \forall j \in \mathbb{T}_N^c. \quad (3.11)$$

Note that this error term is surely far from optimal, but is sufficient here. Overall we can now say

$$|\mathbb{E}(\mathbf{e}_j^T \mathbf{v}_i)^2 - \mathbb{E}(\mathbf{e}_j^T \mathbf{g}_i)^2| \leq 1 + o(1) \leq 2 \quad \forall j \in \mathbb{T}_N^c. \quad (3.12)$$

We can apply $\widehat{\text{QUÉ}}$ to the terms in square parentheses to give $\varepsilon_1, \dots, \varepsilon_r > 0$ such that

$$|\mathbb{E}(\mathbf{e}_{j_k}^T \mathbf{v}_{i_k})^2 - \mathbb{E}(\mathbf{e}_{j_k}^T \mathbf{g}_{i_k})^2| \leq N^{-\varepsilon_{i_k}} \quad \forall j_k \in \mathbb{T}_N \quad \forall i_k = 1, \dots, r. \quad (3.13)$$

We can obtain a single error bound by setting $\varepsilon = \min_i \varepsilon_i$, where clearly $\varepsilon > 0$ and then write

$$N\mathbb{E}(\mathbf{e}_{j_k}^T \mathbf{v}_{i_k})^2 - \mathbb{E}(\mathbf{e}_{j_k}^T \mathbf{g}_{i_k})^2 = \eta_{i_k j_k}^2 N^{-\varepsilon} \quad (3.14)$$

where $\eta_{i_k j_k}^2 \in [-1, 1]$. To further include the indices $j \in \mathbb{T}_N^c$, we extended the expression (3.14) to all j_k by saying

$$\eta_{i_k j_k}^2 \in [-1, 1] \quad \forall j_k \in \mathbb{T}_N, \quad (3.15)$$

$$\eta_{i_k j_k}^2 \in [-N^\varepsilon, N^\varepsilon] \quad \forall j_k \in \mathbb{T}_N^c. \quad (3.16)$$

Overall we have

$$\mathbb{E} \left(\sum_{i=1}^r y_i \mathbf{v}_i^T \Lambda \mathbf{v}_i \right)^l - \mathbb{E} \left(\sum_{i=1}^r y_i \frac{1}{N} \mathbf{g}_i^T \Lambda \mathbf{g}_i \right)^l = N^{-l(1+\varepsilon)} \sum_{i_1, \dots, i_l=1}^r \sum_{j_1, \dots, j_l=1}^N \prod_{k=1}^l y_{i_k} \lambda_{j_k} \eta_{i_k j_k}^2 \quad (3.17)$$

but by comparing with (3.6) we can rewrite as

$$\mathbb{E} \left(\sum_{i=1}^r y_i \mathbf{v}_i^T \Lambda \mathbf{v}_i \right)^l - \mathbb{E} \left(\sum_{i=1}^r y_i \frac{1}{N} \mathbf{g}_i^T \Lambda \mathbf{g}_i \right)^l = \left(\sum_{i=1}^r N^{-(1+\varepsilon)} y_i \boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i \right)^l \quad (3.18)$$

where $\boldsymbol{\eta}_i^T = (\eta_{i1}, \dots, \eta_{iN})$.

□

3.2. Main result

Theorem 3.4. *Let X be an $N \times N$ real symmetric random matrix and let D be an $N \times N$ symmetric matrix (deterministic or random). Let $\hat{\mu}_X, \hat{\mu}_D$ be the empirical spectral measures of the sequence of matrices X, D and assume there exist deterministic limit measures μ_X, μ_D . Assume that X has QUE, i.e. assumption 2.4. Assume also the $\hat{\mu}_X$ concentrates in the sense that*

$$\mathbb{P}(W_1(\hat{\mu}_X, \mu_X) > \delta) \lesssim e^{-N^\tau f(\delta)} \tag{3.19}$$

where $\tau > 0$ and f is some positive increasing function. Then $H = X + D$ has a limiting spectral measure and it is given by the free convolution $\mu_X \boxplus \mu_D$.

Remark 3.5. A condition like (3.19) is required so that the Laplace method can be applied to the empirical measure $\hat{\mu}_X$. There are of course other ways to formulate such a condition. Consider for example the conditions used in theorems 2.1 and 4.1 of [ABM21]. There it is assumed the existence of a sequence of deterministic measures $(\mu_N)_{N \geq 1}$ and a constant $\kappa > 0$ such that for large enough N

$$W_1(\mathbb{E}\hat{\mu}_X, \mu_N) \leq N^{-\kappa}, \quad W_1(\mu_N, \mu_X) \leq N^{-\kappa}, \tag{3.20}$$

which is of course just a deterministic version of (3.19). [ABM21] introduce the extra condition around concentration of Lipschitz traces:

$$\begin{aligned} & \mathbb{P}\left(\left|\frac{1}{N}\text{Tr}f(H_N) - \frac{1}{N}\mathbb{E}\text{Tr}f(H_N)\right| > \delta\right) \\ & \leq \exp\left(-\frac{c_\zeta}{N^\zeta} \min\left\{\left(\frac{N\delta}{\|f\|_{\text{Lip}}}\right)^2, \left(\frac{N\delta}{\|f\|_{\text{Lip}}}\right)^{1+\varepsilon_0}\right\}\right), \end{aligned} \tag{3.21}$$

for all $\delta > 0$, Lipschitz f and N large enough, where $\zeta, c_\zeta > 0$ are some constants. As shown in the proof of theorem 2.1, this condition is sufficient to obtain

$$\mathbb{P}\left(\left|\int |\lambda| d\hat{\mu}_X(\lambda) - \int |\lambda| d\mathbb{E}\hat{\mu}_X(\lambda)\right| \leq t\right) \leq \exp\left(-\frac{c_\zeta}{N^\zeta} \min\{(2Nt)^\zeta, (2Nt)^{1+\varepsilon_0}\}\right) \tag{3.22}$$

for any $t > 0$ and for large enough N . Note that [ABM21] prove this instead for integration against a regularised version of $\log|\lambda|$, but the proof relies only the integrand's being Lipschitz, so it goes through just the same here. (3.22) and (3.20) clearly combine to give (3.19). The reader may ignore this remark if they are content to take (3.19) as an assumption. Alternatively, as we have shown, (3.19) can be replaced by (3.20) and (3.21), conditions which have already been used for quite general results in the random matrix theory literature.

Proof. We shall denote use the notation

$$G_H(z) = \frac{1}{N} \text{Tr}(z - H)^{-1}. \tag{3.23}$$

Recall the supersymmetric approach to calculating the expected trace of the resolvent of a random matrix ensemble:

$$\mathbb{E}_H G_H(z) = \frac{1}{N} \frac{\partial}{\partial j} \bigg|_{j=0} \mathbb{E}_H Z_H(j) \tag{3.24}$$

where

$$Z_H(j) = \frac{\det(z + j - H)}{\det(z - H)} = \int d\Psi e^{-i\text{Tr}AH} e^{i\text{Tr}\Psi\Psi^\dagger J}, \tag{3.25}$$

$$A = \phi\phi^\dagger + \chi\chi^\dagger, \tag{3.26}$$

$$J = I_N \otimes \begin{pmatrix} z & 0 \\ 0 & j + z \end{pmatrix}, \tag{3.27}$$

$$d\Psi = \frac{d\phi d\phi^* d\chi d\chi^*}{-(2\pi)^{Ni}}, \tag{3.28}$$

$$\Psi = \begin{pmatrix} \phi \\ \chi \end{pmatrix} \tag{3.29}$$

with $\phi \in \mathbb{C}^N$ and χ, χ^* being N -long vectors of anti-commuting variables. Independence of X and D gives

$$\begin{aligned} \mathbb{E}_H Z_H(j) &= \int d\Psi e^{i\text{Tr}\Psi\Psi^\dagger J} \mathbb{E}_{X,D} e^{-i\text{Tr}A(X+D)} \\ &= \int d\Psi e^{i\text{Tr}\Psi\Psi^\dagger J} \mathbb{E}_D e^{-i\text{Tr}AD} \mathbb{E}_X e^{-i\text{Tr}AX}. \end{aligned} \tag{3.30}$$

\mathbb{E}_D simply means integration against a delta-function density if D is deterministic.

Let us introduce some notation: for $N \times N$ matrices K , $\Phi_X(K) = \mathbb{E}_X e^{-i\text{Tr}XK}$, and similarly Φ_D . We also define a new matrix ensemble $\bar{X} \stackrel{d}{=} O^T \Lambda O$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ are equal in distribution to the eigenvalues of X and O is an entirely independent Haar-distributed orthogonal matrix.

Now

$$\begin{aligned} \mathbb{E}_H Z_H(j) &= \int d\Psi e^{i\text{Tr}\Psi\Psi^\dagger J} \Phi_{\bar{X}}(K) \Phi_D(K) + \int d\Psi e^{i\text{Tr}\Psi\Psi^\dagger J} (\Phi_X(A) - \Phi_{\bar{X}}(A)) \Phi_D(A) \\ \implies \mathbb{E} G_{D+X}(z) &= \mathbb{E} G_{D+\bar{X}}(z) + \frac{1}{N} \frac{\partial}{\partial j} \bigg|_{j=0} \int d\Psi e^{i\text{Tr}\Psi\Psi^\dagger J} (\Phi_X(A) - \Phi_{\bar{X}}(A)) \Phi_D(A) \\ &\equiv \mathbb{E} G_{D+\bar{X}}(z) + E(z) \end{aligned} \tag{3.31}$$

and so we need to analyse the error term $E(z)$.

Now consider $X = U^T \Lambda U$ where the rows of U are the eigenvectors $\{\mathbf{u}_i\}_i$ of X . Say also that $K = Q^T Y Q$ for diagonal $Y = (y_1, \dots, y_r, 0, \dots, 0)$, where we note that K has fixed rank, by construction. Then

$$\text{Tr}XK = Y(UQ^T)^T \Lambda(UQ^T)$$

but lemma 3.1 establishes that the rows of UQ^T obey QUE, since the rows of U do. Further, lemma 3.2 then establishes that the columns of UQ^T obey $\widehat{\text{QUE}}$ as required by lemma 3.3. Let $\{\mathbf{v}_i\}$ be those columns, then we have

$$\text{Tr}XK = \sum_{i=1}^r y_i \mathbf{v}_i^T \Lambda \mathbf{v}_i. \tag{3.32}$$

The expectation over X can be split into eigenvalues and conditional eigenvectors

$$\Phi_X(K) = \mathbb{E}_\Lambda \mathbb{E}_{U|\Lambda} \sum_{l=0}^{\infty} \frac{1}{l!} (-i)^l (\text{Tr}U^T \Lambda U K)^l. \tag{3.33}$$

We can simply bound

$$\left| \sum_{l=0}^n \frac{1}{l!} (-i)^l (\text{Tr}U^T \Lambda U)^l \right| \leq e^{|\text{Tr}U^T \Lambda U K|} \tag{3.34}$$

for any n , but clearly

$$\mathbb{E}_{U|\Lambda} e^{|\text{Tr}U^T \Lambda U K|} < \infty \tag{3.35}$$

since, whatever the distribution of $U | \Lambda$, the integral is over a compact group (the orthogonal group $O(N)$) and the integrand has no singularities. Therefore, by the dominated convergence theorem

$$\Phi_X(K) = \mathbb{E}_\Lambda \sum_{l=0}^\infty \frac{1}{l!} (-i)^l \mathbb{E}_{U|\Lambda} (\text{Tr}U^T \Lambda U K)^l \tag{3.36}$$

and in precisely the same way

$$\Phi_{\bar{X}}(K) = \mathbb{E}_\Lambda \sum_{l=0}^\infty \frac{1}{l!} (-i)^l \mathbb{E}_{O \sim \mu_{\text{Haar}}} (\text{Tr}O^T \Lambda O K)^l. \tag{3.37}$$

Recalling (3.32) we now have

$$\Phi_X(K) = \mathbb{E}_\Lambda \sum_{l=0}^\infty \frac{1}{l!} (-i)^l \mathbb{E}_{U|\Lambda} \left(\sum_{i=1}^r y_i \mathbf{v}_i^T \Lambda \mathbf{v}_i \right)^l \tag{3.38}$$

and similarly

$$\Phi_{\bar{X}}(K) = \mathbb{E}_\Lambda \sum_{l=0}^\infty \frac{1}{l!} (-i)^l \mathbb{E}_{U|\Lambda} \left(\sum_{i=1}^r y_i \bar{\mathbf{v}}_i^T \Lambda \bar{\mathbf{v}}_i \right)^l \tag{3.39}$$

where the $\bar{\mathbf{v}}_i$ are defined in the obvious way from \bar{X} . We would now apply $\widehat{\text{QUE}}$, but to do so we must insist that \mathbb{E}_Λ is taken over the *ordered* eigenvalues of X . Having fixed that convention, lemma 3.3 can be applied to the terms

$$\mathbb{E}_{U|\Lambda} \left(\sum_{i=1}^r y_i \mathbf{v}_i^T \Lambda \mathbf{v}_i \right)^l \tag{3.40}$$

in (3.38). The terms in $\Phi_{\bar{X}}$ can be treated similarly. This results in

$$\begin{aligned} \Phi_X(K) - \Phi_{\bar{X}}(K) = \mathbb{E}_\Lambda \left[\sum_{l=0}^\infty \frac{l!}{l!} \left\{ \mathbb{E}_{\{\mathbf{g}_i\}_{i=1}^r} \left(\sum_{i=1}^r y_i \frac{1}{N} \mathbf{g}_i^T \Lambda \mathbf{g}_i \right)^l + \left(\sum_{i=1}^r N^{-(1+\varepsilon)} y_i \boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i \right)^l \right\} \right. \\ \left. - \sum_{l=0}^\infty \frac{l!}{l!} \left\{ \mathbb{E}_{\{\bar{\mathbf{g}}_i\}_{i=1}^r} \left(\sum_{i=1}^r y_i \frac{1}{N} \bar{\mathbf{g}}_i^T \Lambda \bar{\mathbf{g}}_i \right)^l + \left(\sum_{i=1}^r N^{-(1+\varepsilon)} y_i \bar{\boldsymbol{\eta}}_i^T \Lambda \bar{\boldsymbol{\eta}}_i \right)^l \right\} \right]. \end{aligned} \tag{3.41}$$

The exponential has infinite radius of convergence, so we may re-order the terms in the sums to give cancellation

$$\Phi_X(K) - \Phi_{\bar{X}}(K) = \mathbb{E}_\Lambda \sum_{l=1}^\infty \frac{1}{l!} N^{-(1+\varepsilon)l} (-i)^l \left(\sum_{i=1}^r y_i \boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i \right)^l - \mathbb{E}_\Lambda \sum_{l=1}^\infty \frac{1}{l!} N^{-(1+\varepsilon)l} (-i)^l \left(\sum_{i=1}^r y_i \bar{\boldsymbol{\eta}}_i^T \Lambda \bar{\boldsymbol{\eta}}_i \right)^l.$$

Here $\varepsilon > 0$ and $\boldsymbol{\eta}_i, \bar{\boldsymbol{\eta}}_i \in \mathbb{C}^N$ with

$$-1 \leq [(\boldsymbol{\eta}_i)_j]^2, [(\bar{\boldsymbol{\eta}}_i)_j]^2 \leq 1 \quad \forall i = 1, \dots, r, \quad \forall j \in \mathbb{T}_N, \tag{3.42}$$

$$-N^\varepsilon \leq [(\boldsymbol{\eta}_i)_j]^2, [(\tilde{\boldsymbol{\eta}}_i)_j]^2 \leq N^\varepsilon \quad \forall i = 1, \dots, r, \forall j \in \mathbb{T}_N. \quad (3.43)$$

Simplifying, we obtain

$$\begin{aligned} \Phi_X(K) - \Phi_{\tilde{X}}(K) &= \mathbb{E}_\Lambda \exp \left(-iN^{-(1+\varepsilon)} \sum_{i=1}^r y_i \boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i \right) \\ &\quad - \mathbb{E}_\Lambda \exp \left(-iN^{-(1+\varepsilon)} \sum_{i=1}^r y_i \tilde{\boldsymbol{\eta}}_i^T \Lambda \tilde{\boldsymbol{\eta}}_i \right). \end{aligned} \quad (3.44)$$

Since $|\mathbb{T}_N^c| \leq 2^{\sim} N^{1-\delta}$ we have

$$\sum_{j \in \mathbb{T}_N^c} |\lambda_j| \leq \mathcal{O}(N^{1-d} N^{-1}) \text{Tr}|\Lambda| \quad (3.45)$$

and so

$$|\boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i| \leq \text{Tr}|\Lambda| (1 + \mathcal{O}(N^{\varepsilon-\delta})). \quad (3.46)$$

For any fixed $\delta > 0$, ε can be reduced if necessary so that $\varepsilon < \delta$ and then for sufficiently large N we obtain, say,

$$|\boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i| \leq 2\text{Tr}|\Lambda|. \quad (3.47)$$

Then we can write $\boldsymbol{\eta}_i^T \Lambda \boldsymbol{\eta}_i = \text{Tr}|\Lambda| \xi_i$ for $\xi_i \in [-2, 2]$, and similarly $\tilde{\boldsymbol{\eta}}_i^T \Lambda \tilde{\boldsymbol{\eta}}_i = \text{Tr}|\Lambda| \tilde{\xi}_i$. Now

$$\mathbb{E}_\Lambda \exp \left(-iN^{-(1+\varepsilon)} \sum_{i=1}^r \xi_i y_i \text{Tr}|\Lambda| \right) = \mathbb{E}_\Lambda \exp \left(-iN^{-\varepsilon} \sum_{i=1}^r \xi_i y_i \int d\hat{\mu}_X(\lambda) |\lambda| \right)$$

so we can apply Laplace's method to the empirical spectral measure $\hat{\mu}_X$ to obtain

$$\mathbb{E}_\Lambda \exp \left(-iN^{-(1+\varepsilon)} \sum_{i=1}^r \xi_i y_i \text{Tr}|\Lambda| \right) = \exp \left(-iN^{-\varepsilon} (q + o(1)) \sum_{i=1}^r \xi_i y_i \right) + o(1) \quad (3.48)$$

where the $o(1)$ terms do not depend on the y_i and where we have defined

$$q = \int d\mu_X(\lambda) |\lambda|. \quad (3.49)$$

Further, we can write $\sum_{i=1}^r \xi_i y_i = \zeta \text{Tr}K$, where $\zeta \in [\min_i \{\xi_i\}, \max_i \{\xi_i\}] \subset [-1, 1]$, and similarly $\sum_{i=1}^r \tilde{\xi}_i y_i = \tilde{\zeta} \text{Tr}K$. Then

$$\Phi_X(K) - \Phi_{\tilde{X}}(K) = e^{-iN^{-\varepsilon} \zeta (q + o(1)) \text{Tr}K} - e^{-iN^{-\varepsilon} \tilde{\zeta} (q + o(1)) \text{Tr}K} + o(1) \quad (3.50)$$

but

$$\begin{aligned} \left. \frac{1}{N} \frac{\partial}{\partial j} \right|_{j=0} \int d\Psi e^{i\text{Tr}\Psi \Psi^\dagger J} e^{-iN^{-\varepsilon} \zeta (q + o(1)) \text{Tr}K} \Phi_D(A) &= \mathbb{E}G_{D+N^{-\varepsilon} \zeta (q + o(1)) I}(z) = \mathbb{E}G_D(z + \mathcal{O}(N^{-\varepsilon})) \\ \implies E(z) &= \mathbb{E}G_D(z + \mathcal{O}(N^{-\varepsilon})) + o(1) - \mathbb{E}G_D(z + \mathcal{O}(N^{-\varepsilon})) - o(1) = o(1). \end{aligned} \quad (3.51)$$

We have thus established that

$$\mathbb{E}G_{D+X}(z) = \mathbb{E}G_{D+\tilde{X}}(z) + o(1) \quad (3.52)$$

from which one deduces that $\mu_{D+X} = \mu_{D+\tilde{X}} = \mu_D \boxplus \mu_X = \mu_D \boxplus \mu_X$.

□

Remark 3.6. We have also constructed a non-rigorous argument for theorem 3.4 where the supersymmetric approach is replaced by the replica method. This approach simplifies some of the analysis but at the expense of being not at all rigorous (indeed there are integral expressions in this argument that are manifestly infinite). The supersymmetric methods used here are not fully rigorous (like most of their applications) but we note that recent work is beginning to elevate supersymmetric random matrix calculations to full rigour [SS17, Shc20].

3.3. Experimental validation

Let $U(a, b)$ denote the uniform distribution on the interval (a, b) , and $\Gamma(a)$ the Gamma-distribution with scale parameter a . We consider the following matrix ensembles:

$$\begin{aligned}
 M \sim \text{GOE}^n &: \text{Var}(M_{ij}) = \frac{1 + \delta_{ij}}{2n}, \\
 M \sim \text{UWig}^n &: \sqrt{n}M_{ij} \stackrel{i.i.d.}{\sim} U(0, \sqrt{6}) \text{ up to symmetry,} \\
 M \sim \Gamma\text{Wig}^n &: 2\sqrt{n}M_{ij} \stackrel{i.i.d.}{\sim} \Gamma(2) \text{ up to symmetry,} \\
 M \sim \text{UWish}^n &: M \stackrel{d}{=} \frac{1}{m}XX^T, X_{ij} \stackrel{i.i.d.}{\sim} U(0, \sqrt{12}) \text{ for } X \text{ of size } n \times m, \frac{n}{m} \xrightarrow{n,m \rightarrow \infty} \alpha, \\
 M \sim \text{Wish}^n &: M \stackrel{d}{=} \frac{1}{m}XX^T, X_{ij} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1) \text{ for } X \text{ of size } n \times m, \frac{n}{m} \xrightarrow{n,m \rightarrow \infty} \alpha.
 \end{aligned}$$

All of the $\text{GOE}^n, \text{UWig}^n, \Gamma\text{Wig}^n$ have the same limiting spectral measure, namely μ_{SC} , the semi-circle of radius $\sqrt{2}$. $\text{UWish}^n, \text{Wish}^n$ have a Marcenko–Pastur limiting spectral measure μ_{MP} , and the constant $\sqrt{12}$ is chosen so that the parameters of the MP measure match those of a Gaussian Wishart matrix Wish^n . $\text{GOE}^n, \text{Wish}^n$ are the only ensembles whose eigenvectors are Haar distributed, but all ensembles obey a local law in the sense above. It is known that the sum of GOE^n and any of the other ensembles will have limiting spectral measure given by the free additive convolution of μ_{SC} and the other ensemble’s measure (so either $\mu_{\text{SC}} \boxplus \mu_{\text{MP}}$ or $\mu_{\text{SC}} \boxplus \mu_{\text{SC}}$), indeed this free addition property holds for any invariant ensemble [AGZ10]. Our result implies that the same holds for addition of the non-invariant ensembles. Sampling from the above ensembles is simple, so we can easily generate spectral histograms from multiple independent matrix samples for large n . $\mu_{\text{SC}} \boxplus \mu_{\text{SC}}$ is just another semi-circle measure but with radius 2. $\mu_{\text{SC}} \boxplus \mu_{\text{MP}}$ can be computed in the usual manner with R -transforms and is given by the solution to the polynomial

$$\frac{\alpha}{2}t^3 - \left(\frac{1}{2} + \alpha z\right)t^2 + (z + \alpha - 1)t - 1 = 0.$$

i.e. Say the cubic has roots $\{r_1, r_2 + is_2, r_2 - is_2\}$ for $s_2 \geq 0$, then the density of $\mu_{\text{SC}} \boxplus \mu_{\text{MP}}$ at z is s_2/π . This can all be solved numerically. The resulting plots are in figure 4 and clearly show agreement between the free convolutions and sampled spectral histograms.

We can also test the result in another more complicated case. Consider the case of random d -regular graphs on N vertices. Say $M \sim \text{Reg}^{N,d}$ is the distribution of the adjacency matrix of such random graphs. The limiting spectral density of $M \sim \text{Reg}^{N,d}$ is known in closed form, as is its Stieljtes transform [BHY19] and [BHY19] established a local law of the kind required for our results. Moreover, there are known efficient algorithms for sampling random d -regular graphs [KV03, SW99] along with implementations [HSS08]. Let $\mu_{\text{KM}}^{(d)}$ be the Kesten–McKay law, the limiting spectral measure of d -regular graphs. We could find an explicit degree-6 polynomial for the Stieljtes transform of $\mu_{\text{KM}}^{(d)} \boxplus \mu_{\text{SC}}$ and compare to spectral histograms as above. Alternatively we can investigate agreement with $\mu_{\text{KM}}^{(d)} \boxplus \mu_{\text{SC}}$ indirectly by sampling and




Figure 4. Comparison of theoretical spectral density and empirical from sampled matrices all of size 500×500 . We combine 50 independent matrix samples per plot.




Figure 5. q-q plot comparing the spectrum of samples from $\text{Reg}^{N,d} + UWig^N$ (y-axis) to samples from $\text{Reg}^{N,d} + \text{GOE}^N$ (x-axis).

comparing spectra from say $\text{Reg}^{N,d} + UWig^N$ and also from $\text{Reg}^{N,d} + \text{GOE}^N$. The latter case will certainly yield the distribution $\mu_{\text{KM}}^{(d)} \boxplus \mu_{\text{SC}}$ since the GOE matrices are freely independent from the adjacency matrices. Figure 5 shows a q-q plot⁶ for samples of the spectra from these two matrix distributions and demonstrates near-perfect agreement, thus showing that indeed the spectrum of $\text{Reg}^{N,d} + UWig^N$ is indeed described by $\mu_{\text{KM}}^{(d)} \boxplus \mu_{\text{SC}}$. We reached the same conclusion when repeating the above experiment with $UWish^N + \text{Reg}^{N,d}$ and $Wish^N + \text{Reg}^{N,d}$.

4. Universal complexity of loss surfaces

4.1. Extension of a key result and prevalence of minima

Let’s recall theorem 4.5 from [ABM21]. $H_N(u)$ is our random matrix ensemble with some parametrisation $u \in \mathbb{R}^m$ and its limiting spectral measure is $\mu_\infty(u)$.

Define

$$\mathcal{G}_{-\varepsilon} = \{u \in \mathbb{R}^m \mid \mu_\infty(u)((-\infty, 0)) \leq \varepsilon\}. \tag{4.1}$$

⁶ Recall that a q-q plot shows the quantiles of one distribution on the x axis and another on the y axis. Given two cumulative density functions F_X, F_Y and their percent point functions F_X^{-1}, F_Y^{-1} , the q-q plot is a plot of the parametric curve $(F_X^{-1}(q), F_Y^{-1}(q))$ for $q \in [0, 1]$. Given only finite samples from the random variables X and Y , the empirical percent point functions can be estimated and used in the q-q plot.

So $\mathcal{G}_{-\varepsilon}$ is the event that $\mu_\infty(u)$ is close to being supported only on $(0, \infty)$. Let $l(u), r(u)$ be the left and right edges respectively of the support of $\mu_\infty(u)$.

Theorem 4.1 ([ABM21] theorem 4.5). Fix some $\mathcal{D} \subset \mathbb{R}^m$ and suppose that \mathcal{D} and the matrices $H_N(u)$ satisfy the following.

- For every $R > 0$ and every $\varepsilon > 0$, we have

$$\lim_{N \rightarrow \infty} \frac{1}{N \log N} \log \left[\sup_{u \in B_R} \mathbb{P} \left(d_{\text{BL}}(\hat{\mu}_{H_N(u)}, \mu_\infty(u)) > \varepsilon \right) \right] = -\infty. \tag{4.2}$$

- Several other assumptions detailed in [ABM21].

Then for any $\alpha > 0$ and any fixed $p \in \mathbb{N}$, we have

$$\begin{aligned} & \lim_{N \rightarrow \infty} \frac{1}{N} \log \int_{\mathcal{D}} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) = 0\}] du \\ &= \sup_{u \in \mathcal{D} \cap \mathcal{G}} \left\{ \int_{\mathbb{R}} \log |\lambda| d\mu_\infty(u)(\lambda) - \alpha u^2 \right\}. \end{aligned} \tag{4.3}$$

We claim the following extension

Corollary 4.2. Under the same assumptions as the above theorem and for any integer sequence $k(N) > 0$ such that $k/N \rightarrow 0$ as $N \rightarrow \infty$, we have

$$\begin{aligned} & \lim_{N \rightarrow \infty} \frac{1}{N} \log \int_{\mathcal{D}} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) \leq k\}] du \\ &= \sup_{u \in \mathcal{D} \cap \mathcal{G}} \left\{ \int_{\mathbb{R}} \log |\lambda| d\mu_\infty(u)(\lambda) - \alpha u^2 \right\}. \end{aligned} \tag{4.4}$$

Proof. Firstly note that

$$\begin{aligned} & \frac{1}{N} \log \int_{\mathcal{D}} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) \leq k\}] du \\ & \geq \frac{1}{N} \log \int_{\mathcal{D}} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) = 0\}] du, \end{aligned} \tag{4.5}$$

so it suffices to establish a complementary upper bound. The proof in of theorem 4.5 in [ABM21] establishes an upper bound using

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \int_{(\mathcal{G}_{-\varepsilon})^c} e^{-N\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) = 0\}] du = -\infty \tag{4.6}$$

which holds for all $\varepsilon > 0$. Indeed, $\mathcal{D} = (\mathcal{D} \cap \mathcal{G}_{-\varepsilon}) \cup (\mathcal{D} \cap (\mathcal{G}_{-\varepsilon})^c)$, so

$$\begin{aligned} & \int_{\mathcal{D}} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) \leq k\}] du \\ & \leq \int_{\mathcal{D} \cap \mathcal{G}_{-\varepsilon}} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) \leq k\}] du \\ & \quad + \int_{(\mathcal{G}_{-\varepsilon})^c} e^{-(N+p)\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) \leq k\}] du, \end{aligned} \tag{4.7}$$

so our proof is complete if we can prove the analogous result

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \int_{(\mathcal{G}_{-\varepsilon})^c} e^{-N\alpha u^2} \mathbb{E} [|\det(H_N(u))| \mathbb{1}\{i(H_N(u)) \leq k\}] du = -\infty. \quad (4.8)$$

As in [ABM21], let f_ε be some $\frac{1}{2}$ -Lipschitz function satisfying $\frac{\varepsilon}{2} \mathbb{1}_{x \leq -\varepsilon} \leq f_\varepsilon(x) \leq \frac{\varepsilon}{2} \mathbb{1}_{x \leq 0}$. Suppose $u \in (\mathcal{G}_{-\varepsilon})^c$ and also $i(H_N(u)) \leq k$. Then we have

$$0 \leq \int d\hat{\mu}_{H_N(u)}(x) f_\varepsilon(x) \leq \frac{k\varepsilon}{2N} \quad (4.9)$$

and also

$$\frac{\varepsilon^2}{2} \leq \int d\mu_\infty(u)(x) f_\varepsilon(x) \leq \frac{\varepsilon}{2}. \quad (4.10)$$

We have

$$\begin{aligned} d_{\text{BL}}(\hat{\mu}_{H_N(u)}, \mu_\infty(u)) &\geq \left| \int d\hat{\mu}_{H_N(u)}(x) f_\varepsilon(x) - \int d\mu_\infty(u)(x) f_\varepsilon(x) \right| \\ &\geq \left| \int d\hat{\mu}_{H_N(u)}(x) f_\varepsilon(x) \right| - \left| \int d\mu_\infty(u)(x) f_\varepsilon(x) \right|, \end{aligned} \quad (4.11)$$

so if we can choose

$$\frac{k\varepsilon}{2N} \leq \frac{\varepsilon^2}{2} - \eta \quad (4.12)$$

for some $\eta > 0$, then we obtain $d_{\text{BL}}(\hat{\mu}_{H_N(u)}, \mu_\infty(u)) \geq \eta$. Then applying (4.2) yields the result (4.8). (4.12) can be satisfied if

$$\varepsilon \geq \frac{k}{2N} + \frac{1}{2} \sqrt{\frac{k^2}{N^2} + 8\eta}. \quad (4.13)$$

So, given $\varepsilon > 0$, we can take N large enough such that, say, $\frac{k(N)}{N} < \frac{\varepsilon}{4}$. By taking $\eta < \frac{\varepsilon^2}{128}$ we obtain

$$\frac{k}{2N} + \frac{1}{2} \sqrt{\frac{k^2}{N^2} + 8\eta} < \frac{\varepsilon}{8} + \frac{1}{\sqrt{2}} \max\left(\sqrt{8\eta}, \frac{\varepsilon}{4}\right) < \frac{1 + \sqrt{2}}{8} \varepsilon < \varepsilon \quad (4.14)$$

and so (4.13) is satisfied. Now finally (4.2) can be applied (with η in place of ε) and so we conclude (4.8).

Overall we see that the superexponential BL condition (4.2) is actually strong enough to deal with any $o(N)$ index not just index-0. This matches the GOE (or generally invariant ensemble) case, in which the terms with $\mathbb{1}\{i(H_N(u)) = k\}$ are suppressed compared to the exact minima terms $\mathbb{1}\{i(H_N(u)) = 0\}$. □

Remark 4.3. Note that corollary 4.2 establishes that, on the exponential scale, the number of critical points of any index $k(N) = o(N)$ is no more than the number of exact local minima.

4.2. The dichotomy of rough and smooth regions

Recall the batch loss from section 2.1:

$$\frac{1}{b} \sum_{i=1}^b \mathcal{L}(f_{\mathbf{w}}(\mathbf{x}_i), y_i), \quad (\mathbf{x}_i, y_i) \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_{\text{data}}. \quad (4.15)$$

As with the Hessian in section 2.1, we use the model $L \equiv L_{\text{batch}}(\mathbf{w}) = L_{\text{true}}(\mathbf{w}) + \mathbf{s}(b)V(\mathbf{w})$, where V is a random function $\mathbb{R}^N \rightarrow \mathbb{R}$.

Now let us define the complexity for sets $\mathcal{B} \subset \mathbb{R}^N$

$$C_N(\mathcal{B}) = |\{\mathbf{w} \in \mathcal{B} \mid \nabla L(\mathbf{w}) = 0\}|. \quad (4.16)$$

This is simply the number of stationary points of the training loss in the region \mathcal{B} of weight space. A Kac-Rice formula applied to ∇L gives

$$\mathbb{E}C_N = \int_{\mathcal{B}} d\mathbf{w} \phi_{\mathbf{w}}(-\mathbf{s}(b)^{-1} \nabla L_{\text{true}}) \mathbb{E}|\det(A + \mathbf{s}(b)X)| \quad (4.17)$$

where $\phi_{\mathbf{w}}$ is the density of ∇V at \mathbf{w} . A rigorous justification of this integral formula would, for example, have to satisfy the conditions of the results of [AT+07]. This is likely to be extremely difficult in any generality, though is much simplified in the case of Gaussian V (and X) - see [AT+07] theorem 12.1.1 or [BKMN21] theorem 4.4. Hereafter, we shall take (4.17) as assumed. The next step is to make use of strong self-averaging of the random matrix determinants. Again, we are unable to establish this rigorously at present, but note that this property has been proved in some generality by [ABM21], although we are unable to satisfy all the conditions of those results in any generality here. Self-averaging and using the addition results above gives

$$\frac{1}{N} \log \mathbb{E}|\det(A + \mathbf{s}(b)X)| = \int d(\mu_b \boxplus \nu)(\lambda) \log |\lambda| + o(1)$$

where μ_b, ν depend in principle on \mathbf{w} . We are concerned with $N^{-1} \log \mathbb{E}C_N$, and in particular its sign, which determines the complexity of the loss surface in \mathcal{B} : positive \leftrightarrow exponentially many (in N) critical points, negative \leftrightarrow exponentially few (i.e. none). The natural next step is to apply the Laplace method with large parameter N to determine the leading order term in $\mathbb{E}C_N$, however the integral is clearly not of the right form. Extra assumptions on $\phi_{\mathbf{w}}$ and ∇L_{true} could be introduced, e.g. that they can be expressed as functions of only a finite number of combinations of coordinates of \mathbf{w} .

Suppose that $\phi_{\mathbf{w}}$ has its mode at 0, for any \mathbf{w} , which is arguably a natural property, reflecting in a sense that the gradient noise has no preferred direction in \mathbb{R}^N . The sharp spike at the origin in the spectral density of DNN Hessians suggests that generically

$$\int d(\mu_b \boxplus \nu)(\lambda) \log |\lambda| < 0. \quad (4.18)$$

We claim it is reasonable to expect the gradient (and Hessian) variance to be increasing in $\|\mathbf{w}\|_2$. Indeed, consider the general form of the simplest DNN, a *multi-layer perceptron*:

$$f_{\mathbf{w}}(\mathbf{x}) = \sigma(\mathbf{b}^{(L)} + W^{(L)}\sigma(\mathbf{b}^{(L-1)} + W^{(L-1)}\dots\sigma(\mathbf{b}^{(1)} + W^{(1)}\mathbf{x})\dots)) \quad (4.19)$$

where all of the weight matrices $W^{(l)}$ and bias vectors $\mathbf{b}^{(l)}$ combine to give the weight vector \mathbf{w} . Viewing \mathbf{x} as a random variable, making f a random function of \mathbf{w} , we expect from the above that the variance in $f_{\mathbf{w}}$ is generally increasing in $\|\mathbf{w}\|_2$, and so therefore similarly with L_{batch} .

Overall it follows that $\phi_{\mathbf{w}}(-\mathbf{s}(b)^{-1} \nabla L_{\text{true}})$ is generally decreasing in $\|\nabla L_{\text{true}}\|$, but the maximum value at $\phi_{\mathbf{w}}(0)$ is decreasing in $\|\mathbf{w}\|_2$. The picture is therefore that the loss surface is

simple and without critical points in regions for which ∇L_{true} is far from 0. In neighbourhoods of $\nabla L_{\text{true}} = 0$, the loss surface may become complex, with exponentially many critical points, however if $\|\mathbf{w}\|_2$ is too large then the loss surface may still be without critical points. In addition, the effect of larger batch size (and hence larger $\mathfrak{s}(b)^{-1}$) is to simplify the surface. These considerations indicate that DNN loss surfaces are simplified by over-parametrisation, leading to the spike in the Hessian spectrum and thus (4.18). The simple fact that neural networks' construction leads gradient noise variance to increase with $\|\mathbf{w}\|_2$ has the effect of simplifying the loss landscape far from the origin of weight space, and even precluding the existence of any critical points of the batch loss.

5. Implications for curvature from local laws

Consider a general stochastic gradient update rule with curvature-adjusted preconditioning:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha B_t^{-1} \nabla L(\mathbf{w}_t) \quad (5.1)$$

where recall that $L(\mathbf{w})$ is the batch loss, viewed as a random function on weight space. B_t is some preconditioning matrix which in practice would be chosen to somehow approximate the curvature of L . Such methods are discussed at length in [Mar16] and also describe some of the most successful optimisation algorithms used in practice, such as Adam [KB14]. The most natural choice for B_t is $B_t = \nabla^2 L(\mathbf{w}_t)$, namely the Hessian of the loss surface. In practice, it is standard to include a damping parameter $\delta > 0$ in B_t , avoid divergences when inverting. Moreover, typically B_t will be constructed to be some positive semi-definite approximation to the curvature such as the generalised Gauss Newton matrix [Mar16], or the diagonal gradient variance form used in Adam [KB14]. Let us now suppose that $B_t = B_t(\delta) = \hat{H}_t + \delta$, where \hat{H}_t is some chosen positive semi-definite curvature approximation and $\delta > 0$. We can now identify $B_t(\delta)^{-1}$ as in fact the Green's function of \hat{H}_t , i.e.

$$B_t(\delta)^{-1} = -(-\delta - \hat{H}_t)^{-1} = -G_t(-\delta). \quad (5.2)$$

But G_t is precisely the object used in the statement of a local law on for \hat{H}_t . Note that $\nabla L(\mathbf{w}_t)$ is a random vector and however \hat{H}_t is constructed, it will generally be a random matrix and dependent on $\nabla L(\mathbf{w}_t)$ in some manner that is far too complicated to handle analytically. As we have discussed at length hitherto, we conjecture that a local law is reasonable assumption to make on random matrices arising in DNNs. In particular [BGK22] demonstrated universal local random matrix theory statistics not just for Hessians of deep networks but also for Generalised Gauss-Newton matrices. Our aim here is to demonstrate how a local law on \hat{H}_t dramatically simplifies the statistics of (5.1). Note that some recent work [WHS22] has also made use of random matrix local laws to simplify the calculation of test loss for neural networks.

A local law on \hat{H}_t takes the precise form (for any $\xi, D > 0$

$$\sup_{\|\mathbf{u}\|, \|\mathbf{v}\|=1, z \in \mathcal{S}} \mathbb{P} \left(|\mathbf{u}^T G(z) \mathbf{v} - \mathbf{u}^T \Pi(z) \mathbf{v}| > N^\xi \left(\frac{1}{N\eta} + \sqrt{\frac{\Im g_\mu(z)}{N\eta}} \right) \right) \leq N^{-D} \quad (5.3)$$

where

$$\mathcal{S} = \{E + i\eta \in \mathbb{C} \mid |E| \leq \omega^{-1}, N^{-1+\omega} \leq \eta \leq \omega^{-1}\} \quad (5.4)$$

μ is the limiting spectral measure of \hat{H}_t and, crucially, Π is a *deterministic* matrix. We will use the following standard notation to re-express (5.3)

$$|\mathbf{u}^T G(z) \mathbf{v} - \mathbf{u}^T \Pi(z) \mathbf{v}| \prec \Psi_N(z), \quad \|\mathbf{u}\|, \|\mathbf{v}\| = 1, z \in \mathcal{S}, \quad (5.5)$$

where $\Psi_N(z) = \frac{1}{N\eta} + \sqrt{\frac{\Im g_\mu(z)}{N\eta}}$ and the probabilistic statement, valid for all $\xi, D > 0$ is implicit in the symbol \prec . In fact, we will need the local law outside the spectral support, i.e. at $z = x + i\eta$ where $x \in \mathbb{R} \setminus \text{supp}(\mu)$. In that case $\Psi_N(z)$ is replaced by $\frac{1}{N(\eta + \kappa)}$ where κ is the distance of x from $\text{supp}(\mu)$ on the real axis, i.e.

$$|\mathbf{u}^T G(z) \mathbf{v} - \mathbf{u}^T \Pi(z) \mathbf{v}| \prec \frac{1}{N(\eta + \kappa)}, \quad \|\mathbf{u}\|, \|\mathbf{v}\| = 1, \quad x \in \mathbb{R} \setminus \text{supp}(\mu). \quad (5.6)$$

For $\delta > 0$ this becomes

$$|\mathbf{u}^T G(-\delta) \mathbf{v} - \mathbf{u}^T \Pi(-\delta) \mathbf{v}| \prec \frac{1}{N\delta} \|\mathbf{u}\|_2 \|\mathbf{v}\|_2 \quad (5.7)$$

for $\delta > 0$ and now any \mathbf{u}, \mathbf{v} . Applying this to (5.1) gives

$$|\mathbf{u}^T B_t^{-1} \nabla L(\mathbf{w}_t) - \mathbf{u}^T \Pi_t(-\delta) \nabla L(\mathbf{w}_t)| \prec \frac{1}{N\delta} \|\mathbf{u}\|_2 \|\nabla L(\mathbf{w}_t)\|_2. \quad (5.8)$$

Consider any \mathbf{u} with $\|\mathbf{u}\|_2 = \alpha$, then we obtain

$$|\mathbf{u}^T B_t^{-1} \nabla L(\mathbf{w}_t) - \mathbf{u}^T \Pi_t(-\delta) \nabla L(\mathbf{w}_t)| \prec \frac{\alpha \|\nabla L(\mathbf{w}_t)\|_2}{N\delta}. \quad (5.9)$$

Thus with high probability, for large N , we can replace (5.1) by

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \Pi_t(-\delta) \nabla L(\mathbf{w}_t) \quad (5.10)$$

incurring only a small error, provided that

$$\delta \gg \frac{\|\nabla L(\mathbf{w}_t)\|_2}{N} \alpha. \quad (5.11)$$

Note that the only random variable in (5.10) is $\nabla L(\mathbf{w}_t)$. If we now consider the case $\nabla L(\mathbf{w}_t) = \nabla \bar{L}(\mathbf{w}_t) + \mathbf{g}(\mathbf{w}_t)$ for deterministic \bar{L} , then

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \Pi_t(-\delta) \nabla \bar{L}(\mathbf{w}_t) - \alpha \Pi_t(-\delta) \mathbf{g}(\mathbf{w}_t) \quad (5.12)$$

and so the noise in the parameter update is entirely determined by the gradient noise. Moreover note the *linear* dependence on \mathbf{g} in (5.12). For example, a Gaussian model for \mathbf{g} immediately yields a Gaussian form in (5.12), and e.g. if $\mathbb{E} \mathbf{g} = 0$, then

$$\mathbb{E}(\mathbf{w}_{t+1} - \mathbf{w}_t) = -\alpha \Pi_t(-\delta) \mathbb{E} \nabla L(\mathbf{w}_t). \quad (5.13)$$

A common choice in practice for \hat{H} is a diagonal matrix, e.g. the diagonal positive definite curvature approximation employed by Adam [KB14]. In such cases, \hat{H} is best viewed as an approximation to the eigenvalues of some positive definite curvature approximation. The next result establishes that a local law assumption on a general curvature approximation matrix can be expected to transfer to an analogous result on a diagonal matrix of its eigenvalues.

Proposition 5.1. *Suppose that \hat{H} obeys a local law of the form (5.6). Define the diagonal matrix D such that $D_i \stackrel{d}{=} \lambda_i$ where $\{\lambda_i\}_i$ are the sorted eigenvalues of \hat{H} . Let $G_D(z) = (z - D)^{-1}$ be the resolvent of D . Let $q_j[\mu]$ be the j th quantile of μ , the limiting spectral density of \hat{H} , i.e.*

$$\int_{-\infty}^{q_j[\mu]} d\mu(\lambda) = \frac{j}{N}. \quad (5.14)$$

Then D obeys the local law

$$|(G_D)_{ij} - \delta_{ij}(z - q_j[\mu])^{-1}| \prec \frac{1}{N^{2/3}(\kappa + \eta)^2}, \quad z = x + i\eta, \quad x \in \mathbb{R} \setminus \text{supp}(\mu), \quad (5.15)$$

where κ is the distance of x from $\text{supp}(\mu)$. Naturally, we can redefine $D_i = \lambda_{\sigma i}$ for any permutation $\sigma \in S_N$ and the analogous statement replacing $\mathfrak{q}_j[\mu]$ with $\mathfrak{q}_{\sigma(j)}$ will hold.

Proof. As in [EY17], the local law (5.5), (5.6) is sufficient to obtain rigidity of the eigenvalues in the bulk, i.e. for any $\varepsilon, D > 0$

$$\mathbb{P}\left(\exists j \mid |\lambda_j - \mathfrak{q}_j[\mu]| \geq N^\varepsilon [\min(j, N-j+1)]^{-1/3} N^{-2/3}\right) \leq N^{-D}. \quad (5.16)$$

Then we have

$$\left| \frac{1}{z - \lambda_j} - \frac{1}{z - \mathfrak{q}_j[\mu]} \right| = \left| \frac{\lambda_j - \mathfrak{q}_j[\mu]}{(z - \lambda_j)(z - \mathfrak{q}_j[\mu])} \right|. \quad (5.17)$$

For $z = x + i\eta$ and x at a distance $\kappa > 0$ from $\text{supp}(\mu)$

$$|z - \mathfrak{q}_j[\mu]|^2 \geq \eta^2 + \kappa^2 \geq \frac{1}{2}(\eta + \kappa)^2, \quad (5.18)$$

and the same can be said for $|z - \lambda_j|^2$ with high probability, by applying the rigidity (5.16). A second application of rigidity to $|\lambda_j - \mathfrak{q}_j[\mu]|$ gives

$$\left| \frac{1}{z - \lambda_j} - \frac{1}{z - \mathfrak{q}_j[\mu]} \right| \prec \frac{1}{N^{2/3} \min(j, N-j+1)^{1/3} (\kappa + \eta)^2} \quad (5.19)$$

which yields the result. \square

With this result in hand, we get the generic update rule akin to (5.12), with high probability

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \text{diag} \left(\frac{1}{\pi_j + \delta} \right) \nabla \bar{L}(\mathbf{w}_t) - \alpha \text{diag} \left(\frac{1}{\pi_j + \delta} \right) \mathbf{g}(\mathbf{w}_t) \quad (5.20)$$

where $\{\pi_j\}_{j=1}^N$ are the eigenvalues of $\Pi_t(0)$ and we emphasise again that the π_j are *deterministic*; the only stochastic term is the gradient noise $\mathbf{g}(\mathbf{w}_t)$.

5.1. Implications for preconditioned stochastic gradient descent

The key insight from this section is that generic random matrix theory effects present in preconditioning matrices of large neural networks can be expected to drastically simplify the optimisation dynamics due to high-probability concentration of the pre-conditioning matrices around deterministic equivalents, nullifying the statistical interaction between the pre-conditioning matrices and gradient noise. Moreover, with this interpretation, the damping constant typically added to curvature estimate matrices is more than a simple numerical convenience: it is essential to yield the aforementioned concentration results.

As an example of the kind of analysis that the above makes possible, consider the results of [GBW+21]. The authors consider a Gaussian process model for the noise in the loss surface, resulting in tractable analysis for convergence of stochastic gradient descent in the presence of statistical dependence between gradient noise in different iterations. Such a model implies a specific form of the loss surface Hessian and its statistical dependence on the gradient noise. This situation is a generalisation of the spin glass model exploited in various works [CHM+15, BKMN21, BKMN22], except that in those cases the Hessian can be shown to be independent of the gradients. Absent the very special conditions that lead to independence, one expects the analysis to be intractable, hence why the authors in [GBW+21] restrict to

stochastic gradient descent without preconditioning, or simply assume a high probability concentration on a deterministic equivalent. To make this discussion more concrete, consider a model $L = L_{\text{true}} + V$ where V is a Gaussian process with mean 0 and covariance function

$$K(\mathbf{x}, \mathbf{x}') = k\left(\frac{1}{2}\|\mathbf{x} - \mathbf{x}'\|_2^2\right) q\left(\frac{1}{2}(\|\mathbf{x}\|_2^2 + \|\mathbf{x}'\|_2^2)\right), \tag{5.21}$$

where k is some decreasing function and q some increasing function. The discussion at the end of the previous section suggests that the covariance function for loss noise should not be modelled as stationary, hence the inclusion of the function q in (5.21). For convenience define $\Delta = \frac{1}{2}(\|\mathbf{x} - \mathbf{x}'\|_2^2)$ and $S = \frac{1}{2}(\|\mathbf{x}\|_2^2 + \|\mathbf{x}'\|_2^2)$. Then it is a short exercise in differentiation to obtain

$$\begin{aligned} \text{Cov}(\partial_i V(\mathbf{w}), \partial_j V(\mathbf{w})) &= \text{Cov}(\partial_i V(\mathbf{w}), \partial_j V(\mathbf{w}')) \Big|_{\mathbf{w}=\mathbf{w}'} \\ &= \frac{\partial^2}{\partial w_i \partial w_j'} K(\mathbf{w}, \mathbf{w}') \Big|_{\mathbf{w}=\mathbf{w}'} \\ &= -k'(0)q(\|\mathbf{w}\|_2)\delta_{ij} + k(0)q''(\|\mathbf{w}\|_2^2)w_i w_j'. \end{aligned} \tag{5.22}$$

and moreover

$$\begin{aligned} \text{Cov}(\partial_{ii} V(\mathbf{w}), \partial_j V(\mathbf{w})) &= \text{Cov}(\partial_{ii} V(\mathbf{w}), \partial_j V(\mathbf{w}')) \Big|_{\mathbf{w}=\mathbf{w}'} \\ &= \frac{\partial^3}{\partial w_i \partial w_i \partial w_j'} K(\mathbf{w}, \mathbf{w}') \Big|_{\mathbf{w}=\mathbf{w}'} \\ &= -k'(0)q'(\|\mathbf{w}\|_2^2)w_i \delta_{ij} + q''(\|\mathbf{w}\|_2^2)k(0)w_i w_l w_j' - k'(0)q'(\|\mathbf{x}\|_2)w_i \delta_{jl}. \end{aligned} \tag{5.23}$$

Hence we see that the gradients of L and its Hessian are statistically dependent by virtue of the non-stationary structure of V . Putting aside issues of positive definite pre-conditioning matrices, and taking δ such that $(\nabla^2 L + \delta)^{-1}$ exists (almost surely) for large N , it is clear that the distribution of $(\nabla^2 L + \delta)^{-1} \partial V$ will be complicated and non-Gaussian. This example concretely illustrates our point: even in almost the simplest case, where the gradient noise is Gaussian, the pre-conditioned gradients are generically considerably more complicated and non-Gaussian. Moreover, centred Gaussian noise on gradient is transformed into generically non-centred noise by pre-conditioning. Continuing the differentiation above, it is elementary to obtain the covariance structure of the Hessian $\nabla^2 V$, though the expressions are not instructive. Crucially, however, the Hessian is Gaussian and the covariance of any of its entries is $\mathcal{O}(1)$ (in large N), so the conditions in Example 2.12 of [EKS19] apply to yield an optimal local law on the Hessian, which in turn yields the above high-probability concentration of $(\nabla^2 L + \delta)^{-1}$ provided that δ is large enough.

6. Conclusion

In this paper we have considered several aspects of so-called universal random matrix theory behaviour in DNNs. Motivated by prior experimental results, we have introduced a model for the Hessians of DNNs that is more general than any previously considered and, we argue, actually flexible enough to capture the Hessians observed in real-world DNNs. Our model

is built using random matrix theory assumptions that are more general than those previously considered and may be expected to hold in quite some generality. By proving a new result for the addition of random matrices, using a novel combination of QUE and the supersymmetric method, we have derived expressions for the spectral outliers of our model. Using Lanczos approximation to the outliers of large, practical DNNs, we have compared our expressions for spectral outliers to data and demonstrated strong agreement for some DNNs. As well as corroborating our model, this analysis presents indirect evidence of the presence of universal local random matrix statistics in DNNs, extending earlier experimental results. Our analysis also highlights a possibly interesting distinction between some DNN architectures, as Resnet architectures appear to better agree with our theory than other architectures and Resnets have been previously observed to have better-behaved loss surfaces than many other architectures.

We also presented quite general arguments regarding the number of local optima of DNN loss surfaces and how ‘rough’ or ‘smooth’ such surfaces are. Our arguments build on a rich history of complexity calculations in the statistical physics and mathematics literature but, rather than performing detailed calculations in some specific, highly simplified toy model, we instead present general insights based on minimal assumptions. Finally we highlight an important area where random matrix local laws, an essential aspect of universality, may very directly influence the performance of certain popular optimisation algorithms for DNNs. Indeed, we explain how numerical damping, combined with random matrix local laws, can act to drastically simplify the training dynamics of large DNNs.

Overall it is our hope that this paper demonstrates the relevance of random matrix theory to DNNs beyond highly simplified toy models. Moreover, we have shown how quite general and universal properties of random matrices can be fruitfully employed to derive practical, observable properties of DNN spectra. This work leaves several challenges for future research. All of our work relies on either local laws for e.g. DNN Hessians, or on matrix determinant self-averaging results. Despite the considerable progress towards establishing local laws for random matrices over the last decade or-so, it appears that establishing any such laws for, say, the Hessians of any DNNs is quite out of reach. We expect that the first progress in this direction will come from considering DNNs with random i.i.d. weights and perhaps simple activation functions. Based on the success of recent works on random DNNs [PS20], we conjecture that the Gram matrices of random DNN Jacobians may be the simplest place to establish a local law, adding to the nascent strand of *nonlinear* random matrix theory [PW17, BP19, PS20]. We also believe that there is more to be gained in further studies of forms of random matrix universality in DNNs. For example, our ideas may lead to tractable analysis of popular optimisation algorithms such as Adam [KB14] as the problem is essentially reduced to deriving a local law for the gradient pre-conditioning matrix and dealing with the gradient noise.

Data availability statement

No new data were created or analysed in this study.

Acknowledgement

The authors wish to thank Patrick Lopatto for noticing some technical errors and providing valuable insights into proof strategies and existing results for QUE. They also wish to thank two anonymous reviewers for thorough and helpful reviews of the first draft. For providing generous computing resources, the authors wish to thank Andrew Gittings and the JADE Hartree

facility for generous computing use. JPK is pleased to acknowledge support from European Research Council Advanced Grant 740900 (LogCorRM).

Appendix A. Invariant equivalent ensembles

For an invariant ensemble [Dei99] with potential V we have the following integro-differential equation relating the equilibrium measure μ to the potential V [Unt19]:

$$\frac{\beta}{2} \int \frac{1}{x-y} d\mu(y) = V'(x). \tag{A.1}$$

So in the case of real symmetric matrices we have

$$\frac{1}{2} \bar{g}_\mu(x) = V'(x) \tag{A.2}$$

where g_μ is the Stieljtes transform of μ and the bar over \bar{g}_μ indicates that the principal value has been taken.

Given a sufficiently nice μ (A.1) defines V up-to a constant of integration on $\text{supp}(\mu)$, but V is not determined on $\mathbb{R} \setminus \text{supp}(\mu)$, as is made clear by the following lemma, which we prove for completeness but which has appeared before in various works (e.g. [Dei99]).

Lemma A.1. *For compactly supported probability measure μ on \mathbb{R} and real potential V , define*

$$S_V[\mu](y) = V(y) - \int d\mu(x) \log|y-x|. \tag{A.3}$$

Suppose $S_V[\mu](y) = c$, a constant, for all $y \in \text{supp}(\mu)$ and $S_V[\mu](y) \geq c$ for all $y \in \mathbb{R}$. Then μ is a minimiser amongst all probability measures on \mathbb{R} of the energy

$$\mathcal{E}_V[\mu] = \int d\mu(x)V(x) - \iint_{x<y} d\mu(x)d\mu(y) \log|x-y|. \tag{A.4}$$

Proof. Consider a probability measure that is close to μ in the sense of W_1 distance, say.

For any such measure, one can find an arbitrarily close probability measure μ' of the form

$$\mu' = \mu + \sum_{i=1}^r a_i \mathbb{1}_{[y_i-\delta_i, y_i+\delta_i]} - \sum_{i=1}^s b_i \mathbb{1}_{[z_i-\eta_i, z_i+\eta_i]} \tag{A.5}$$

where all $a_i, b_i > 0$ and $\delta_i, \eta_i, a_i, b_i \leq \varepsilon$ for some small $\varepsilon > 0$. To ensure that μ' is again a probability measure we must impose $\sum_i a_i = \sum_j b_j$. The strategy now is to expand $\mathcal{E}_V[\mu']$ about μ to first order in ε , but first note the symmetrisation

$$\iint_{x<y} d\mu(x)d\mu(y) \log|x-y| = \frac{1}{2} \iint_{x \neq y} d\mu(x)d\mu(y) \log|x-y|. \tag{A.6}$$

Then

$$\begin{aligned} \mathcal{E}_V[\mu'] - \mathcal{E}_V[\mu] &= \sum_{i=1}^r a_i V(y_i) - \sum_{i=1}^s b_i V(z_i) - \sum_{i=1}^r a_i \int d\mu(x) \log|x-y_i| \\ &\quad + \sum_{i=1}^s b_i \int d\mu(x) \log|x-z_i| + \mathcal{O}(\varepsilon^2) \\ &= \sum_{i=1}^r a_i S_V[\mu](y_i) - \sum_{i=1}^s b_i S_V[\mu](z_i) + \mathcal{O}(\varepsilon^2). \end{aligned} \tag{A.7}$$

Observe that if all $y_i, z_i \in \text{supp}(\mu)$ then $S_V[\mu](y_i) = S_V[\mu](z_i) = c$ and so $\mathcal{E}_V[\mu'] = \mathcal{E}_V[\mu]$. Without loss of generality therefore, we take $y_i \notin \text{supp}(\mu)$ and $z_i \in \text{supp}(\mu)$, whence

$$\mathcal{E}_V[\mu'] - \mathcal{E}_V[\mu] \geq c \sum_{i=1}^r a_i - c \sum_{i=1}^s b_i = 0. \quad (\text{A.8})$$

□

The next lemma establishes that, while not unique, a potential V can always be constructed given a measure μ .

Lemma A.2. *Consider a probability measure μ on \mathbb{R} with compact support, absolutely continuous with respect to the Lebesgue measure. Then there exists a potential $V: \mathbb{R} \rightarrow \mathbb{R}$ which yields a well-defined invariant distribution on real symmetric matrices for which the equilibrium measure is μ .*

Proof. (A.2) can be integrated to obtain V and the condition $S_V[\mu] = c$ (a constant) on $\text{supp}(\mu)$ determines V uniquely on $\text{supp}(\mu)$. Next observe that, for $y \in \mathbb{R} \setminus \text{supp}(\mu)$ there exists some constant $R > 0$ such that $|x - y| \leq R + |y|$, since μ is compactly supported, and so $\log|x - y| \leq |y| + R$. Therefore

$$S_V[\mu](y) \geq V(y) - |y| - R. \quad (\text{A.9})$$

V must be chosen on $\mathbb{R} \setminus \text{supp}(\mu)$ to satisfy $S_V[\mu](y) \geq c$, which can be achieved by ensuring

$$V(y) \geq |y| + R + c. \quad (\text{A.10})$$

Additionally, V must be defined for large y such that it defines an legitimate invariant ensemble on symmetric real matrices, i.e. V must decay sufficiently quickly at infinity to give an integrable probability density. Finally, V must be sufficiently smooth, and certainly continuous, so there are boundary conditions at the boundary of $\text{supp}(\mu)$. Suppose $\text{supp}(\mu)$ is composed of K disjoint intervals, then there are $2K$ boundary conditions on V , and the bound (A.10) imposes one further condition. Sufficiently fast decay at infinity can be satisfied by any even degree polynomial V of degree at least 2, therefore a degree $2K + 2$ polynomial can be found with sufficiently fast decay at infinity, satisfying all the boundary conditions and (A.10). □

Appendix B. Experimental details

This section gives full details of the experimental set-up and analysis for the outlier experiments.

B.1. Architectures and training of models

We use the GPU powered Lanczos quadrature algorithm [GPW+18, MS06], with the Pearlmutter trick [Pea94] for Hessian vector products, using the PyTorch [PGC+17] implementation of both Stochastic Lanczos Quadrature and the Pearlmutter. We then train a 16 Layer VGG CNN [SZ14] with $P = 15\,291\,300$ parameters and the 28 Layer Wide Residual Network

[ZK16, HZRS16] architectures on the CIFAR-100 dataset [KH+09] (45 000 training samples and 5000 validation samples) using SGD. We use the following learning rate schedule:

$$\alpha_t = \begin{cases} \alpha_0, & \text{if } \frac{t}{T} \leq 0.5 \\ \alpha_0 \left[1 - \frac{(1-r)(\frac{t}{T}-0.5)}{0.4}\right] & \text{if } 0.5 < \frac{t}{T} \leq 0.9 \\ \alpha_0 r, & \text{otherwise.} \end{cases} \quad (\text{B.1})$$

We use a learning rate ratio $r = 0.01$ and a total number of epochs budgeted $T = 300$. We further use momentum set to $\rho = 0.9$, a weight decay coefficient of 0.0005 and data-augmentation on PyTorch (paszke2017automatic).

B.2. Implementation of constraints

As mentioned in the main text, one of the three weights of the linear model fit in the outlier analysis, β , is constrained to be positive, as it corresponds to a second cumulant, i.e. a variance, of a probability measure. Recall that the linear model's parameters are solved exactly as functions of the unknown $\theta^{(i)}$, and these parameters are in turn optimised using gradient descent. β is unconstrained during the linear solve, but its value is determined by the $\theta^{(i)}$, so to impose the constraint $\beta > 0$ we add to the mean squared error loss the term

$$\beta = 1000 \max(0, -\beta) \quad (\text{B.2})$$

which penalises negative β values and is minimised at any non-negative value. The factor 1000 was roughly tuned by hand to give consistently positive values for β .

There is also the constraint that $\theta^{(i)} > \theta^{(i+1)} > 0$ for all i . This is imposed simply using a re-parametrisation. We introduce unconstrained raw value $t^{(i)}$ taking values in \mathbb{R} and define

$$\theta^{(i)} = \sum_{j=1}^i \log(1 + \exp(t^{(j)})),$$

then the gradient descent optimisation is simply performed over the $t^{(i)}$.

B.3. Fitting of outlier model

We optimise the mean squared error with respect to the raw parameters $t^{(i)}$ using 200 iterations of Adam [KB14] with a learning rate of 0.2. The learning rate was chosen heuristically by increasing in steps until training became unstable. The number of iterations was chosen heuristically as being comfortably sufficient to obtain convergence of Adam. The raw parameters $t^{(i)}$ were initialised by drawing independently from a standard Gaussian. The $t^{(i)}$ were initialised and trained using the above method 20 times and the values with the lowest mean squared error were chosen.

Appendix C. Spectral plots

This appendix contains some plots of the approximate empirical spectral densities of the Hessians of the DNNs considered in our experiments. Figure C.6 shows the spectra for the VGG16 network trained on CIFAR100 at various stages from initialisation to the end of training. Note the clear presence of large outliers present already at epoch 25. These are not intended to be exhaustive and we recommend references such as [Pap18] for detailed discussion of spectral densities like these.




Figure C.6. Approximate empirical spectral densities of the Hessian of the VGG16 network trained on MNIST.

ORCID iDs

Nicholas P Baskerville  <https://orcid.org/0000-0001-5275-3598>
 Jonathan P Keating  <https://orcid.org/0000-0003-0864-038X>
 Francesco Mezzadri  <https://orcid.org/0000-0002-4572-5617>
 Joseph Najnudel  <https://orcid.org/0000-0001-6132-1455>
 Diego Granzio  <https://orcid.org/0000-0003-3169-2081>

References

- [AAĀ13] Auffinger A, Arous G B and Āerný J 2013 Random matrices and complexity of spin glasses *Commun. Pure Appl. Math.* **66** 165–201
- [ABM21] Arous G B, Bourgade P and McKenna B 2021 Exponential growth of random determinants beyond invariance (arXiv:2105.05000)
- [AGZ10] Anderson G W, Guionnet A and Zeitouni O 2010 *An Introduction to Random Matrices (Cambridge Studies in Advanced Mathematics vol 118)* (Cambridge: Cambridge University Press)
- [AT+07] Adler R J and Taylor J E 2007 *Random Fields and Geometry* vol 80 (Berlin: Springer)
- [BAP05] Baik J, Arous G B and Pécché S 2005 Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices *Ann. Probab.* **33** 1643–97
- [BEK+14] Bloemendal A, Erdős L, Knowles A, Yau H-T and Yin J 2014 Isotropic local laws for sample covariance and generalized wigner matrices *Electron. J. Probab.* **19** 1–53
- [Ben20] Benigni L 2020 Eigenvectors distribution and quantum unique ergodicity for deformed wigner matrices *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* **56** 2822–67
- [BES20] Bao Z, Erdős L and Schnelli K 2020 On the support of the free additive convolution *J. Anal. Math.* **142** 323–48
- [BGK22] Baskerville N P, Granzio D and Keating J P 2022 Appearance of random matrix theory in deep learning *Physica A* **590** 126742
- [BGN11] Benaych-Georges F and Nadakuditi R R 2011 The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices *Adv. Math.* **227** 494–521
- [BHY19] Bauerschmidt R, Huang J and Yau H-T 2019 Local kesten–mckay law for random regular graphs *Commun. Math. Phys.* **369** 523–636

- [BJS+19] Baity-Jesi M, Sagun L, Geiger M, Spigler S, Arous G B, Cammarota C, LeCun Y, Wyart M and Biroli G 2019 Comparing dynamics: deep neural networks versus glassy systems *J. Stat. Mech.: Theory Exp.* **2019** 124013
- [BKMN21] Baskerville N P, Keating J P, Mezzadri F and Najnudel J 2021 The loss surfaces of neural networks with general activation functions *J. Stat. Mech.: Theory Exp.* **2021** 064001
- [BKMN22] Baskerville N P, Keating J P, Mezzadri F and Najnudel J 2022 A spin glass model for the loss surfaces of generative adversarial networks *J. Stat. Phys.* **186** 1–45
- [BL21] Benigni L and Lopatto P 2021 Fluctuations in local quantum unique ergodicity for generalized wigner matrices (arXiv:2103.12013)
- [BL22] Benigni L and Lopatto P 2022 Optimal delocalization for generalized wigner matrices *Adv. Math.* **396** 108109
- [BP19] Benigni L and Pécché S 2019 Eigenvalue distribution of nonlinear models of random matrices (arXiv:1904.03090)
- [BY17] Bourgade P and Yau H-T 2017 The eigenvector moment flow and local quantum unique ergodicity *Commun. Math. Phys.* **350** 231–78
- [CDM16] Capitaine M and Donati-Martin C 2016 Spectrum of deformed random matrices and free probability (arXiv:1607.05560)
- [CHM+15] Choromanska A, Henaff M, Mathieu M, Arous G B and LeCun Y 2015 The loss surfaces of multilayer networks *Artificial Intelligence and Statistics (PMLR)* pp 192–204
- [Dei99] Deift P 1999 *Orthogonal Polynomials and Random Matrices: A Riemann-Hilbert Approach* vol 3 (Providence, RI: American Mathematical Society)
- [EKS19] Erdős L, Krüger T and Schröder D 2019 Random matrices with slow correlation decay *Forum of Mathematics, Sigma* vol 7 (Cambridge: Cambridge University Press)
- [ES17] Erdős L and Schnelli K 2017 Universality for random matrix flows with time-dependent density *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* **53** 1606–56
- [EY12] Erdős L and Yau H-T 2012 Universality of local spectral statistics of random matrices *Bull. Am. Math. Soc.* **49** 377–414
- [EY17] Erdős L and Yau H-T 2017 *A Dynamical Approach to Random Matrix Theory (Courant Lecture Notes in Mathematics* vol 28) (Providence, RI: American Mathematical Society)
- [EYY12] Erdős L, Yau H-T and Yin J 2012 Bulk universality for generalized wigner matrices *Probab. Theory Relat. Fields* **154** 341–407
- [Fyo04] Fyodorov Y V 2004 Complexity of random energy landscapes, glass transition and absolute value of the spectral determinant of random matrices *Phys. Rev. Lett.* **92** 240601
- [Fyo05] Fyodorov Y V 2005 Counting stationary points of random landscapes as a random matrix problem *Acta Phys. Pol. B* **36** 2699–707
- [Gar88] Gardner E 1988 The space of interactions in neural network models *J. Phys. A: Math. Gen.* **21** 257,
- [GB22] Granzio D and Baskerville N 2022 A random matrix theory approach to damping in deep learning (arXiv:2011.08181)
- [GBW+21] Granzio D, Baskerville N, Wan X, Albanie S and Roberts S 2021 Iterative averaging in the quest for best test error (arXiv:2003.01247)
- [GPAM+14] Goodfellow I, Pouget-Abadie J, Mirza M, Bing X, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 Generative adversarial nets *Advances in Neural Information Processing Systems* vol 27
- [GPW+18] Gardner J, Pleiss G, Weinberger K Q, Bindel D and Wilson A G 2018 Gpytorch: blackbox matrix-matrix gaussian process inference with GPU acceleration *Advances in Neural Information Processing Systems* pp 7576–86
- [Gra20] Granzio D 2020 Beyond random matrix theory for deep networks (arXiv:2006.07721)
- [GWG19] Granzio D, Wan X and Garipov T 2019 Deep curvature suite (arXiv:1912.09656)
- [GZR20] Granzio D, Zohren S and Roberts S 2020 Learning rates as a function of batch size: a random matrix theory approach to neural network training (arXiv:2006.09092)
- [HSS08] Hagberg A A, Schult D A and Swart P J 2008 Exploring network structure, dynamics and function using networkx *Proc. 7th Python in Science Conf. (Pasadena, CA. USA)* ed G Varoquaux, T Vaught and J Millman pp 11–15
- [HZRS16] Kaiming H, Zhang X, Ren S and Sun J 2016 Deep residual learning for image recognition *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* pp 770–8
- [KB14] Kingma D P and Jimmy B 2014 Adam: a method for stochastic optimization (arXiv:1412.6980)

- [KH+09] Krizhevsky A *et al* 2009 Learning multiple layers of features from tiny images
- [KV03] Kim J H and Vu V H 2003 Generating random regular graphs *Proc. 35th Annual ACM Symp. on Theory of Computing* pp 213–22
- [KY17] Knowles A and Yin J 2017 Anisotropic local laws for random matrices *Probab. Theory Relat. Fields* **169** 257–352
- [LBBH98] LeCun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition *Proc. IEEE* **86** 2278–324
- [LXT+18] Hao Li, Zheng X, Taylor G, Studer C and Goldstein T 2018 Visualizing the loss landscape of neural nets *Advances in Neural Information Processing Systems* vol 31
- [Mar16] Martens J 2016 *Second-Order Optimization for Neural Networks* (Canada: University of Toronto)
- [MBAB20] Maillard A, Arous G B and Biroli G 2020 Landscape complexity for the empirical risk of generalized linear models *Proc. First Mathematical and Scientific Machine Learning Conf. (Proc. Machine Learning Research* vol 107) (July 2020) ed L Jianfeng and R Ward (Princeton, NJ: Princeton University) pp 287–327
- [MBC+19] Mannelli S S, Biroli G, Cammarota C, Krzakala F and Zdeborová L 2019 Who is afraid of big bad minima? Analysis of gradient-flow in spiked matrix-tensor models *Advances in Neural Information Processing Systems* pp 8676–86
- [Meh04] Mehta M L 2004 *Random Matrices* (Amsterdam: Elsevier)
- [MPV87] Mézard M, Parisi G and Virasoro M 1987 *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* vol 9 (Singapore: World Scientific Publishing Company)
- [MS06] Meurant G and Strakoš Z 2006 The Lanczos and conjugate gradient algorithms in finite precision arithmetic *Acta Numer.* **15** 471–542
- [Pap18] Pappan V 2018 The full spectrum of deepnet Hessians at scale: dynamics with sgd training and sample size (arXiv:1811.07062)
- [Pap19] Pappan V 2019 Measurements of three-level hierarchical structure in the outliers in the spectrum of deepnet Hessians (arXiv:1901.08244)
- [Pas20] Pastur L 2020 On random matrices arising in deep neural networks. Gaussian case (arXiv:2001.06188)
- [Pea94] Pearlmutter B A 1994 Fast exact multiplication by the Hessian *Neural Comput.* **6** 147–60
- [PGC+17] Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L and Lerer A 2017 Automatic differentiation in pytorch
- [PS20] Pastur L and Slavin V 2020 On random matrices arising in deep neural networks: general I.I.D case (arXiv:2011.11439)
- [PSG18] Pennington J, Schoenholz S and Ganguli S 2018 The emergence of spectral universality in deep networks *Int. Conf. on Artificial Intelligence and Statistics (PMLR)* pp 1924–32
- [PW17] Pennington J and Worah P 2017 Nonlinear random matrix theory for deep learning *Advances in Neural Information Processing Systems* vol 30
- [RABC19] Ros V, Arous G B, Biroli G and Cammarota C 2019 Complex energy landscapes in spiked-tensor and simple glassy models: ruggedness, arrangements of local minima and phase transitions *Phys. Rev. X* **9** 011003
- [SGAL14] Sagun L, Ugur Guney V, Arous G B and Yann L 2014 Explorations on high dimensional landscapes (arXiv:1412.6615)
- [Shc20] Shcherbina T 2020 Characteristic polynomials for random band matrices near the threshold *J. Stat. Phys.* **179** 920–44
- [SS17] Shcherbina M and Shcherbina T 2017 Characteristic polynomials for 1d random band matrices from the localization side *Commun. Math. Phys.* **351** 1009–44
- [SW99] Steger A and Wormald N C 1999 Generating random regular graphs quickly *Comb. Probab. Comput.* **8** 377–96
- [SZ14] Simonyan K and Zisserman A 2014 Very deep convolutional networks for large-scale image recognition (arXiv:1409.1556)
- [Tao12] Tao T 2012 *Topics in Random Matrix Theory* vol 132 (Providence, RI: American Mathematical Society)
- [TSR22] Thamm M, Staats M and Rosenow B 2022 Random matrix analysis of deep neural network weight matrices (arXiv:2203.14661)

- [Unt19] Unterberger J 2019 Global fluctuations for 1d log-gas dynamics. Covariance kernel and support *Electron. J. Probab.* **24** 1–28
- [WHS22] Wei A, Wei H and Steinhardt J 2022 More than a toy: random matrix models predict how real-world neural representations generalize (arXiv:[2203.06176](https://arxiv.org/abs/2203.06176))
- [ZK16] Zagoruyko S and Komodakis N 2016 Wide residual networks (arXiv:[1605.07146](https://arxiv.org/abs/1605.07146))