
Dynamical Isometry is Achieved in Residual Networks in a Universal Way for any Activation Function

Wojciech Tarnowski

Piotr Warchoł

Stanisław Jastrzębski

Jacek Tabor

Maciej A. Nowak

Jagiellonian University in Kraków

Jagiellonian University in Kraków

Abstract

We demonstrate that in residual neural networks (ResNets) dynamical isometry is achievable irrespective of the activation function used. We do that by deriving, with the help of Free Probability and Random Matrix Theories, a universal formula for the spectral density of the input-output Jacobian at initialization, in the large network width and depth limit. The resulting singular value spectrum depends on a single parameter, which we calculate for a variety of popular activation functions, by analyzing the signal propagation in the artificial neural network. We corroborate our results with numerical simulations of both random matrices and ResNets applied to the CIFAR-10 classification problem. Moreover, we study consequences of this universal behavior for the initial and late phases of the learning processes. We conclude by drawing attention to the simple fact, that initialization acts as a confounding factor between the choice of activation function and the rate of learning. We propose that in ResNets this can be resolved based on our results by ensuring the same level of dynamical isometry at initialization.

1 Introduction

Deep Learning has achieved unparalleled success in fields such as object detection and recognition, language translation, and speech recognition (LeCun et al., 2015). At the same time, models achieving these state-of-the-art results are increasingly deep and complex (Canziani et al., 2016), which often leads to optimization challenges such as vanishing gradients. Many solutions to this problem have been proposed. In particular, Residual Neural Networks

remedy this to some extent (He et al., 2016; Veit et al., 2016) by using skip connections in the network architecture, which improve gradient flow. As a result, Residual Neural Networks outmatched other competing models in the 2015 ILSVRC and COCO competitions. Yet another approach towards solving this problem is to tailor fit the networks weight initialization to facilitate training, for example by ensuring dynamical isometry (Pennington et al., 2017). In this latter case, the insights are based on an analysis of the statistical properties of information propagation in the network and a study of the full singular spectrum of a particular matrix, namely the input-output Jacobian, via the techniques of Free Probability and Random Matrix Theories (FPT & RMT). This perspective has recently led to successfully training a 10000 layer vanilla convolutional neural network (Xiao et al., 2018).

RMT is a versatile tool that, since its inception, saw a substantial share of applications, from the earliest in nuclear physics (Wigner, 1993) to the latest in game theory (Carmona et al., 2018) (see (Akemann et al., 2011) for some of the use cases discovered in the mean time). It is thus not surprising that it found its way to be used to understand artificial neural networks. In particular, to study their loss surface (Choromanska et al., 2015; Pennington and Bahri, 2017), the associated Gram matrix (Louart et al., 2018; Pennington and Worah, 2017) and in the case of single layer networks, their dynamics (Liao and Couillet, 2018). Our main contribution is extending the theoretical analysis of (Pennington et al., 2017; Schoenholz et al., 2016; Pennington et al., 2018) to residual networks. In particular, we find that residual networks can achieve dynamical isometry for many different activation functions provided that the variance of weight initialization scale is inversely proportional to the number of skip-connections. This is in contrast to feedforward networks, where orthogonal weights and antisymmetric sigmoidal activation functions (like tanh) are required. These theoretical results are supported by an empirical investigation on the popular CIFAR-10 benchmark.

1.1 Related work

The framework of dynamical mean field theory, we will apply to study signal propagation in neural networks, was first used in this context in (Poole et al., 2016). There, the authors showed the existence of an order-to-chaos expressivity transition for deep feedforward neural networks with random initial weights, on the plane spanned by the variances of the network weights and biases. This in turn led to the insight of (Schoenholz et al., 2016), that arbitrary deep networks can be trained as long as they are close to the criticality associated with that transition. The techniques developed in these works, together with methods of FPT and RMT, allowed, for the first time, to analytically compute the singular value distribution of the input-output Jacobian of a deep feedforward network with nonlinear activation function and at criticality (Pennington et al., 2017). Finally, (Pennington et al., 2018) showed, that for feedforward neural networks, in their large depth limit and at a special point of the above mentioned critical line, the singular spectrum of the Jacobian is given by a universal distribution depending on the form of the activation function used. In particular they distinguish the *Bernoulli* and the *smooth* universality classes corresponding to piecewise linear and some nonlinear activation functions. In fact, in this paper, we take the approach of that last work and apply it to fully connected residual neural networks. We find a single universality class for this architecture.

Let us also mention some recent, important developments in the area of residual neural network initialization. One of the earlier developments, is the introduction of layer-sequential unit-variance (LSUV) initialization (Mishkin and Matas, 2015). The two step process involved normalizing the outputs of the neurons on the first forward run and showed promising results. In another, very relevant paper (Taki, 2017), analyzing the signal propagation in a similar manner to that mentioned in the paragraph above, shows for ResNets, with piecewise linear, symmetric as well as ReLU activation functions, that the proper variance for network weight initialization is of order $\frac{1}{NL}$, where L is the number of layers and N the number of neurons in each layer. A similar conclusion is reached by (Balduzzi et al., 2017). We corroborate this result with our analysis. Another contribution shows that adding skip connections to the network, eliminates the critical behavior described above (Yang and Schoenholz, 2017). Finally, the importance of initialization in ResNets is shown by (Zhang et al., 2019), where it is demonstrated that initializing to a zero function enables training state of the art residual networks without the use of batch normalization (Ioffe and Szegedy, 2015). Note that ResNet with this initialization achieves in fact an ideal isometry.

When finishing this manuscript, we have learned of a recent paper tackling the same problem of ResNets initialization by studying the singular spectrum properties of the Jaco-

bian with the tools of Free Probability. While the analysis of (Ling and Qiu, 2018) is related, it is crucial to note that the authors do not observe the universal character of the singular spectrum - the main result of our paper, and treat only piecewise linear activation functions. It is also worth mentioning that, similar to us, they rediscover the importance of $\frac{1}{LN}$ scaling of (Balduzzi et al., 2017) and (Taki, 2017).

1.2 Our results

Our contributions are the following. We show that the singular spectrum of the input-output Jacobian, in the the networks large width and depth limit, is given by a universal formula - with the dependence on the type of activation function encapsulated in a single parameter. Furthermore, we calculate the layer dependent statistical properties of the pre-activations for a variety of activation functions. All together, this gives the associated singular spectra of the Jacobian, which we compare with random matrix and artificial neural network numerical simulations corroborating our theoretical results. The singular values of the input-output Jacobian concentrate around 1 for a wide range of parameters, which shows that fine-tuning the initialization is not required for achieving dynamical isometry in ResNets. Interestingly, the universal formula for the singular spectrum of the Jacobian is valid also in the case when batch normalization is used. Even though the final results of the theoretical calculations are derived in the limit $L, N \rightarrow \infty$, the numerical experiments match them already for $L = 10$ (with $N = 500$). As a practical application of our work and the universality property it uncovers, we propose a framework for setting up weight initialization in experiments with residual neural networks.

1.3 Structure of the paper

We follow this introductory section by defining the model of ResNets we will work with and with a short note on the relevance of the input-output Jacobian. Then, in subsection 3.1, we derive the equation governing the Green's function and hence the spectrum of the Jacobian, which depends on a single parameter, which we denote by c . Proceeding is the analysis of the propagation of the information in the network via an analysis of the probability density function describing the pre-activations across the layers at network initialization. This allows us to calculate c for many different activation functions in Appendix D. We close the second section of the paper by revealing the random matrix experiments confirming our results. Sec. 4 is devoted to the outcome of associated residual neural network numerical calculations. There, we showcase the resulting, experimental, universal spectrum of the Jacobian and the outcomes of the learning processes. We close the paper with a discussion section. In Appendix A we give a brief comment on the influence of batch normalization on the presented

setup. In the rest of the Appendices, we show the results of numerical experiments validating the signal propagation recurrence relations and some baseline (based on using the same weight matrix variances, irrespective of the choice of activation functions) simulations of the learning process.

2 The model

In this paper, we consider a deep, residual network of L layers of a constant width of N neurons. We follow the typical nomenclature of the literature and therefore the real-valued, synaptic matrix for the l -th layer is denoted by \mathbf{W}^l , whereas the real-valued bias vectors are \mathbf{b}^l . The information propagates in this network according to:

$$\mathbf{x}^l = \phi(\mathbf{h}^l) + a\mathbf{x}^{l-1}, \quad \mathbf{h}^l = \mathbf{W}^l\mathbf{x}^{l-1} + \mathbf{b}^l, \quad (1)$$

where \mathbf{h}^l and \mathbf{x}^l are pre- and post-activations respectively and ϕ is the activation function itself, acting entry-wise on the vector of pre-activations. We have introduced the parameter a to track the influence of skip connections in the calculations, however we do not study its influence on the Jacobian's spectrum or learning in general. By \mathbf{x}^0 we denote the input of the network and by \mathbf{x}^L its output. Our primary interest will lay in exploring the singular value spectral properties of the input-output Jacobian:

$$J_{ik} = \frac{\partial x_i^L}{\partial x_k^0}, \quad (2)$$

known to be useful in studying initialization schemes of neural networks at least since the work of (Glorot and Bengio, 2010). It in particular holds the information on the severity of the exploding gradients problem.

2.1 Relevance of the input-output Jacobian

To understand why we are interested in the Jacobian, consider the neural network adjusting its weights during the learning process. In a simplified, example by example scenario, this happens according to

$$\Delta W_{ij}^l = -\eta \frac{\partial E(\mathbf{x}^L, \mathbf{y})}{\partial W_{ij}^l}, \quad (3)$$

where $E(\mathbf{x}^L, \mathbf{y})$ is the error function depending on \mathbf{x}^L - the output of the network, \mathbf{y} - the correct output value associated with that example and, implicitly through \mathbf{x}^L , on the parameters of the model, namely the weights and biases. Here, for simplicity, we consider only the adjustments of the weights - an analogous reasoning applies to the biases. η is the learning rate. By use of the chain rule we can rewrite this as:

$$\Delta W_{ij}^l = -\eta \sum_{k,t} \frac{\partial x_t^L}{\partial W_{ij}^l} \frac{\partial x_k^L}{\partial x_t^L} \frac{\partial E(\mathbf{x}^L, \mathbf{y})}{\partial x_k^L}, \quad (4)$$

For the learning process to be stable, all three terms need to be bounded. Out of those, the middle one can become problematic if a poor choice of the initialization scheme is made. We can rewrite it as:

$$\frac{\partial x_k^L}{\partial x_t^l} = \left[\prod_{i=l+1}^L (\mathbf{D}^i \mathbf{W}^i + \mathbf{1}a) \right]_{kt} \quad (5)$$

and see the larger the difference between L and l , the more terms we have in the product, and (in general) the less control there is over its behavior. Here $\mathbf{1}$ is an identity matrix and, \mathbf{D}^l is a diagonal matrix such that $D_{ij}^l = \phi'(h_i^l)\delta_{ij}$. Indeed, it was proposed by (Glorot and Bengio, 2010), that learning in deep feed-forward neural networks can be improved by keeping the mean singular value of the Jacobian associated with layer i (in our setup $\mathbf{J}^i = \mathbf{D}^i \mathbf{W}^i + \mathbf{1}a$), close to 1 for all i 's. It is also important for the dynamics of learning to be driven by data, not by the random initialization of the network. The latter may take place if the Jacobian to the l -th layer possesses very large singular values which dominate the learning or very small singular values suppressing it. In the optimal case all singular values should be concentrated around 1 regardless of how deep is the considered layer. One therefore examines the case of $l = 0$, namely $\partial x_k^L / \partial x_t^0$ - the input-output Jacobian, as the most extreme object of (5). The feature that in the limit of large depth all singular values of \mathbf{J} concentrate around 1, irrespective of the depth of the network, was coined as dynamical isometry (Saxe et al., 2013).

Note, that the spectral problem for the full Jacobian

$$\mathbf{J} = \prod_{l=1}^L (\mathbf{D}^l \mathbf{W}^l + \mathbf{1}a) \quad (6)$$

belongs to the class of matrix-valued diffusion processes (Gudowska-Nowak et al., 2003; Janik and Wieczorek, 2004), leading to a complex *eigenvalue* spectrum. We note that the large N limit, spectral properties of (6) with $\mathbf{D} = \mathbf{1}$ (deep linear networks), and different symmetry classes of \mathbf{W} , was derived already by (Gudowska-Nowak et al., 2003). Due to non-normality of the Jacobian, singular values cannot be easily related to eigenvalues. Therefore we follow (Pennington et al., 2017, 2018) and tackle the full *singular spectrum* of the Jacobian (or equivalently the eigenvalue spectrum of $\mathbf{J}\mathbf{J}^T$), extending these works to the case of the Residual Neural Network model.

3 Spectral properties of the Jacobian

3.1 Spectral analysis

Free Probability Theory, or Free Random Variable (FRV) Theory (Voiculescu et al., 1992), is a powerful tool for the spectral analysis of random matrices in the limit of their large size. It is a counterpart of the classical Probability

Theory for the case of non-commuting observables. For a pedagogical introduction to the subject, see (Mingo and Speicher, 2017) - here we start by laying out the basics useful in the derivations of this subsection. The fundamental objects of the theory are the Green's functions (a.k.a. Stieltjes transforms in mathematical literature):

$$G_H(z) = \left\langle \frac{1}{N} \text{Tr}(z\mathbf{1} - \mathbf{H})^{-1} \right\rangle = \int_{-\infty}^{\infty} \frac{\rho_H(\lambda) d\lambda}{z - \lambda}, \quad (7)$$

which generate spectral moments and where the subscript H indicates, that his formulation is proper for self-adjoint matrices. The eigenvalue density can be recovered via the Sochocki-Plemelj formula

$$\rho_H(x) = -\frac{1}{\pi} \lim_{\epsilon \rightarrow 0} G_H(x + i\epsilon). \quad (8)$$

The associated free cumulants are generated by the so-called R -transform, which plays the role of the logarithm of the characteristic function in the classical probability. By this correspondence, the R -transform is additive under addition, i.e. $R_{X+Y}(z) = R_X(z) + R_Y(z)$ for mutually free, but non-commuting random ensembles X and Y . Moreover, it is related to G via the functional equations

$$G\left(R(z) + \frac{1}{z}\right) = z, \quad R(G(z)) + \frac{1}{G(z)} = z. \quad (9)$$

On the other hand, the so-called S -transform facilitates calculations of the spectra of products of random matrices, as it satisfies $S_{AB}(z) = S_A(z)S_B(z)$, provided A and B are mutually free and at least one is positive definite. If additionally, the ensemble has a finite mean, the S -transform can be easily obtained from the R -transform, and vice versa, through a pair of the following, mutually inverse maps $z = yS(y)$ and $y = zR(z)$. Explicitly:

$$S(zR(z)) = \frac{1}{R(z)}, \quad R(zS(z)) = \frac{1}{S(z)}. \quad (10)$$

Denoting now \mathbf{J}_L the Jacobian across L layers and $\mathbf{Y}_l = (a\mathbf{1} + \mathbf{D}^l \mathbf{W}^l)$, one can write the recursion relation $\mathbf{J}_L \mathbf{J}_L^T = \mathbf{Y}_L \mathbf{J}_{L-1} \mathbf{J}_{L-1}^T \mathbf{Y}_L$. The latter matrix is isospectral to $\mathbf{Y}_L^T \mathbf{Y}_L \mathbf{J}_{L-1} \mathbf{J}_{L-1}^T$, which leads to the equation for the S -transform $S_{\mathbf{J}_L \mathbf{J}_L^T}(z) = S_{\mathbf{Y}_L^T \mathbf{Y}_L}(z) S_{\mathbf{J}_{L-1} \mathbf{J}_{L-1}^T}(z)$. Proceeding inductively, we arrive at

$$S_{\mathbf{J} \mathbf{J}^T}(z) = \prod_{l=1}^L S_{\mathbf{Y}_l \mathbf{Y}_l^T}(z). \quad (11)$$

To find the S -transform of the single layer Jacobian, we will first consider its Green's function

$$G(z) = \left\langle \frac{1}{N} \text{Tr}(z\mathbf{1} - \mathbf{Y}_l \mathbf{Y}_l^T)^{-1} \right\rangle, \quad (12)$$

with the averaging over the ensemble of weight matrices \mathbf{W}^l . To facilitate the study of G , in particular to cope with

$\mathbf{Y} \mathbf{Y}^T$, one linearizes the problem by introducing matrices of size $2N \times 2N$

$$\mathcal{Z} := \begin{pmatrix} -a & 1 \\ z & -a \end{pmatrix}, \quad \mathcal{X} := \begin{pmatrix} \mathbf{X} & 0 \\ 0 & \mathbf{X}^T \end{pmatrix}, \quad (13)$$

with $\mathbf{X} = \mathbf{D}^l \mathbf{W}^l$. Another crucial ingredient is the block trace operation (bTr), which is the trace applied to each $N \times N$ block. The generalized Green's function is defined as a block trace of the generalized resolvent $(\mathcal{Z} \otimes \mathbf{1} - \mathcal{X})^{-1}$

$$\mathcal{G} := \begin{pmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{pmatrix} = \left\langle \frac{1}{N} \text{bTr} \begin{pmatrix} -a - \mathbf{X} & 1 \\ z & -a - \mathbf{X}^T \end{pmatrix}^{-1} \right\rangle. \quad (14)$$

Remarkably, the Green's function of $\mathbf{Y} \mathbf{Y}^T$ is the G_{12} entry of the generalized Green's function. This construction is a slight modification of the quaternionization approach to large non-Hermitian matrices developed by (Janik et al., 1997), therefore we adapt these concepts here for calculations in the large width limit of the network. Furthermore, the generalized Green's function (14) is given implicitly by the solution of the Schwinger-Dyson equation

$$\mathcal{G}(\mathcal{Z}) = (\mathcal{Z} - \mathcal{R}(\mathcal{G}(\mathcal{Z})))^{-1}. \quad (15)$$

Here \mathcal{R} is the generalized R -transform of FRV theory. This construction is a generalization of standard FRV tools to the matrix-valued functions. In particular, (15) is such a generalization of (9) to 2×2 matrices.

To study two common weight initializations, Gaussian and scaled orthogonal, on the same footing, we assume that \mathbf{W} belongs to the class of biunitarily invariant random matrices, i.e. its pdf is invariant under multiplication by two orthogonal matrices, $P(\mathbf{U} \mathbf{W} \mathbf{V}^T) = P(\mathbf{W})$ for $\mathbf{U}, \mathbf{V} \in O(N)$. In the large N limit these matrices are known in free probability as R -diagonal operators (Nica and Speicher, 1996). A product of R -diagonal operator with an arbitrary operator remains R -diagonal (Nica and Speicher, 2006), thus the matrix \mathbf{X} is R -diagonal too.

The generalized \mathcal{R} -transform of R -diagonal operators takes a remarkably simple form (Nowak and Tarnowski, 2017)

$$\mathcal{R}(\mathcal{G}) = A(G_{12} G_{21}) \begin{pmatrix} 0 & G_{12} \\ G_{21} & 0 \end{pmatrix}. \quad (16)$$

Here, $A(x) = \sum_{k=1}^{\infty} c_{2k} x^{k-1}$ is the determining sequence, which generates cumulants c_{2k} , it is a Taylor expansion of $A(x)$ at 0. For the later use we mention a simple relation for the determining sequence of a scaled matrix $A_{aX}(z) = a^2 A_X(a^2 z)$, which generalizes the Hermitian case $G_{aH}(z) = \frac{1}{a} G_H(\frac{z}{a})$ or, equivalently, $R_{aH}(z) = a R_H(az)$.

We derive the equation for the Green's function (with $G(z) = G_{12}$) by substituting the \mathcal{R} -transform (16) into (15) and eliminating irrelevant variables. It thus reads:

$$G = \frac{GA(zG^2) - 1}{a^2 - z(1 - GA(zG^2))^2}, \quad (17)$$

where for clarity we omitted the argument of the Green's function. In the next step we substitute $z \rightarrow R(z) + \frac{1}{z}$ and use (9) to obtain

$$z = \frac{zA(z^2R + z) - 1}{a^2 - (R + \frac{1}{z})(1 - zA(z^2R + z))}. \quad (18)$$

Then, we substitute $z \rightarrow zS(z)$ and use (10), which leads us to

$$1 = \frac{zSA(z(z+1)S) - 1}{a^2zS - (z+1)(1 - zSA(z(z+1)S))}. \quad (19)$$

This equation is exact. To incorporate the additional scaling of weights variances by $1/L$ in our considerations, as proposed by (Taki, 2017) and (Balduzzi et al., 2017), we rescale $\mathbf{X} \rightarrow \mathbf{X}/\sqrt{L}$ and since we are interested in deep networks, we keep only the leading term in $1/L$ (see also (Janik and Wieczorek, 2004)). This leads to $A(z(z+1)S) \rightarrow \frac{1}{L}A(\frac{1}{L}z(z+1)S) = \frac{c_2}{L} + \mathcal{O}(\frac{1}{L^2})$, which simplifies (19) to a quadratic equation for S . Choosing the appropriate branch of the solution, we see that

$$S_{Y_i Y_i^T}(z) = \frac{1}{a^2} \left(1 - \frac{c_2^l}{a^2 L} (1 + 2z) + \mathcal{O}\left(\frac{1}{L^2}\right) \right). \quad (20)$$

Here

$$c_2^l = \left\langle \frac{1}{N} \text{Tr} \mathbf{W}^l \mathbf{D}^l \mathbf{D}^l (\mathbf{W}^l)^T \right\rangle = \frac{\sigma_w^2}{N} \sum_i^N (\phi'(h_i^l))^2 \quad (21)$$

is the squared spectral radius of the matrix $\mathbf{D}^l \mathbf{W}^l$. In general, c_2^l can vary across the depth of the network due to non-constant variance of preactivations. Assuming that this variability is bounded, we can consider the logarithm of (11) and write:

$$\ln S_{JJ^T}(z) = -2L \ln a - \frac{(1+2z)}{a^2} c, \quad (22)$$

where we defined the effective cumulant $c = \frac{1}{L} \sum_{l=1}^L c_2^l$ and used $\ln(1+x) \approx x$. This allows us to deduce the form of the S -transform, assuming that a does not scale with L

$$S_{JJ^T}(z) = \frac{1}{a^2 L} e^{-\frac{c}{a^2}(1+2z)}. \quad (23)$$

Substituting $z \rightarrow zR(z)$ and using (10), we obtain

$$a^{2L} = R(z) \exp \left[-\frac{c}{a^2} (1 + 2zR(z)) \right]. \quad (24)$$

Then, we substitute $z \rightarrow G(z)$ and use (9) to finally get

$$a^{2L} G(z) = (zG(z) - 1) e^{\frac{c}{a^2}(1-2zG(z))}, \quad (25)$$

an equation for the Green's function characterizing the square singular values of the Jacobian, which can be solved numerically. We do that for a range of different activation functions and present the results with numerical simulations to corroborate them in Fig. 1.

We close this section with a remark that the above analysis is not restricted only to the model (1), but analogous reasoning can be performed for networks in which skip connections bypass more than one fully connected block. The qualitative results remain unaltered provided that L is replaced by the number of skip connections.

3.2 Signal propagation

The formulas we have derived until now were given in terms of a single parameter c , which is the squared derivative of the activation function averaged within each layer and across the depth of the network. Thus, we now need to address the behavior of preactivations. In the proceeding paragraph, we closely follow a similar derivation done in (Schoenholz et al., 2016), for fully connected feed forward networks.

For the simplicity of our arguments, we consider here W_{ij}^l and b_i^l as independent identically distributed (iid) Gaussian random variables with 0 mean and variances $\frac{(\sigma_w)^2}{LN}$ and $(\sigma_b)^2$, respectively. Here, $(\sigma_w)^2$ is of order one, and the additional scaling is meant to reflect those introduced in the previous paragraphs. At the end of this section we provide an argument that the same results hold for scaled orthogonal matrices.

In this subsection, we will denote the averaging over variables W_{ij}^l and b_i^l , at a given layer l , by $\langle \cdot \rangle_{wbl}$. By $\langle u \rangle_l$ we denote the sample average, of some variable u in the l -th layer: $\langle u \rangle_l \equiv \frac{1}{N} \sum_{i=1}^N u_i^l$. Note that the width (N) is independent of the layer number, however the derivation can be easily generalized to the opposite case, when the architecture is more complicated. Unless stated otherwise explicitly, all integrals are calculated over the real line.

We are interested in the distribution of h_i^l in our model, depending on the input vectors and the probability distributions of W_{ij}^l and b_i^l . If we assume they are normal (as can be argued using the Central Limit Theorem), we just need the first two moments. It is clear that $\langle h_i^l \rangle_{wbl} = 0$. Furthermore, we assume ergodicity, i.e. that averaging some quantity over a layer of neurons is equivalent to averaging this quantity for one neuron over an ensemble of neural networks with random initializations. We assume this is true for h_i , x_i , W_{ij} and b_i . Thus, we can say that $\langle h \rangle_l = 0$ and moreover, as we work in the limit of wide networks, $\langle f(h) \rangle_l$ (where f is some function of h_l) can be replaced with an averaging over a normal distribution of variance $q^l \equiv \langle (h^l)^2 \rangle_{wbl}$. This is the crux of the dynamical mean field theory approach (Poole et al., 2016) for feed-forward neural networks. We have in particular:

$$c_2^l = \sigma_w^2 \langle (\phi'(h))^2 \rangle_l = \sigma_w^2 \int \mathcal{D}z \phi'^2 \left(\sqrt{q^l z} \right), \quad (26)$$

where $\mathcal{D}z = \exp(-z^2/2)dz/\sqrt{2\pi}$. To calculate the effective cumulant, we need to know how the variance of the distribution of the preactivations changes as the input information propagates across consecutive layers of the network. It is shown in Appendix C, that q^l satisfy the recurrence equation

$$q^{l+1} = a^2 q^l + (1 - a^2) \sigma_b^2 + \frac{(\sigma_w)^2}{L} \int \mathcal{D}z \phi^2 \left(\sqrt{q^l} z \right) + 2 \frac{(\sigma_w)^2}{L} \left[\sum_{k=1}^{l-1} a^k \int \mathcal{D}z \phi \left(\sqrt{q^{l-k}} z \right) \right] \int \mathcal{D}z \phi \left(\sqrt{q^l} z \right), \quad (27)$$

with the initial condition $q^1 = \sigma_b^2 + \frac{\sigma_w^2}{L}$.

We remark here that the above reasoning concerning signal propagation holds also when the weights are scaled orthogonal matrices, i.e. $\mathbf{W}\mathbf{W}^T = \frac{\sigma_w^2}{L}\mathbf{1}$. In such a case $\langle W_{ij} \rangle = 0$ and $\langle W_{ij} W_{kl} \rangle = \frac{\sigma_w^2}{NL} \delta_{ik} \delta_{jl}$ (Collins and Śniady, 2006) and the entries of \mathbf{W} can be approximated as independent Gaussians (Chatterjee and Meckes, 2007).

3.3 Random matrix simulations

To thoroughly test the theoretical predictions of Section 3, we run numerical simulations using *Mathematica*. The initial condition, input vector \mathbf{x}^0 of length $N = 500$, filled with iid Gaussian random variables of zero mean and unit variance, is propagated according to the recurrence (1), for various activation functions. The network weights and biases are generated from normal distribution of zero mean and σ_w^2/NL and σ_b^2 variances, respectively, with $N = 500$. The propagation of variance of preactivations, post-activations as well as the calculation of the second cumulants c^l for the studied activation functions, across the network, is presented in Appendix E. All numerical simulations corroborate our theoretical results. Here, for clarity and as a generic example, in Fig. 1 (upper), we show the distribution of singular values of the input-output Jacobian (defined in (6)) for the tanh nonlinearity for various network depths. In this example the Jacobian is not independent of the signal propagation, contrary to the case of piecewise linear activation functions. Similarly, in the lower panel of Fig. 1, we showcase the outcome of numerical experiments and the associated, matching, theoretical results for the most popular ReLU activation function, for various initializations resulting in different values of the effective cumulant c .

4 Experiments on image classification

The goal of this section is to test our theoretical predictions on real data via the popular CIFAR-10 benchmark (Krizhevsky, 2009). To this end we will use a single representation fully connected residual network, see Fig. 2.

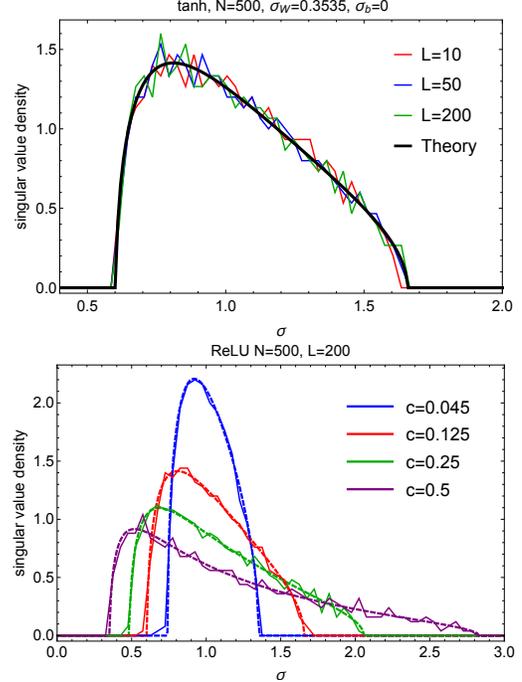


Figure 1: (Top) density of singular values of the input-output Jacobian for the residual network with tanh nonlinearity. Note that the asymptotic theoretical result describes remarkably well not very deep ($L = 10$) networks. (Bottom) asymptotic distribution of singular values for various values of parameter c (dashed) juxtaposed with the numerical simulations for ReLU nonlinearity (solid). Note that histograms were calculated from a single random initialization. The smaller c , the narrower the spectrum and the closer to the ideal isometry.

This simplified version of the model of (He et al., 2016) does not use (i) multiple stages with different dimension of hidden representation, and (ii) two layers within residual block. We leave study of a more general version of ResNets for future work.

4.1 Achieving dynamical isometry for any activation function

Perhaps the most interesting prediction of our theory is that ResNets, in contrast to fully connected networks, can achieve dynamical isometry for many different activation functions. We will study this empirically by looking at \mathbf{J} , at initialization, for different activation functions and number of residual blocks. Please note that by \mathbf{J} we refer to Jacobian of the output of the last residual block with respect to the input of the first one, see also Fig. 2.

We consider the following popular activation functions: ReLU (Nair and Hinton, 2010), Tanh, Hard Tanh, Sigmoid, SeLU (Klambauer et al., 2017) and Leaky ReLU (Maas

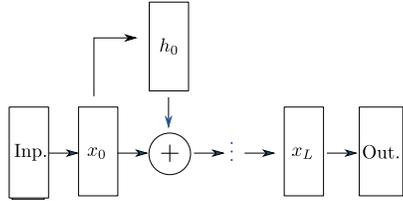


Figure 2: Residual network architecture used in the paper.

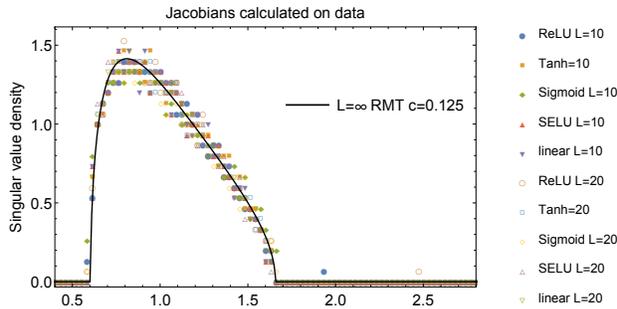


Figure 3: Singular spectra obtained for various activation functions and depth $L = 10, 20$. The network was fed with examples from CIFAR10 dataset.

et al., 2013) with the leaking constant 0.05 and 0.25. For each activation function we consider the number of blocks L to be 10 and 20. All weights of the network are initialized from a zero-centered normal distribution whereas biases are initialized to zero. The weights of the residual blocks are initialized using standard deviation σ_W / \sqrt{NL} , other weights are initialized as by (Taki, 2017). For the given activation function and the number of blocks L , we set σ_W in such a way that the effective cumulant $c = 0.125$, which ensures the concentration of eigenvalues of the Jacobian around one, and hence dynamical isometry (see Appendix D for more details and Fig. 1 for the shape of the singular value densities).

For each pair of activation function and number of blocks we compute the empirical spectrum of \mathbf{J} at initialization, the results are reported in Fig. 3. Indeed, we observe that upon scaling the initializations standard deviation, in such a way that c is kept constant, the empirical spectrum of \mathbf{J} is independent of the number of residual blocks or the choice of activation functions.

4.2 Learning dynamics are more similar at universality under dynamical isometry

Our next objective is to investigate whether networks achieving dynamical isometry share similar learning dynamics. While this is outside of the scope of our theoretical investigation, it is inspired by studies such as (Pennington et al., 2017), which demonstrate the importance of dynamical isometry.

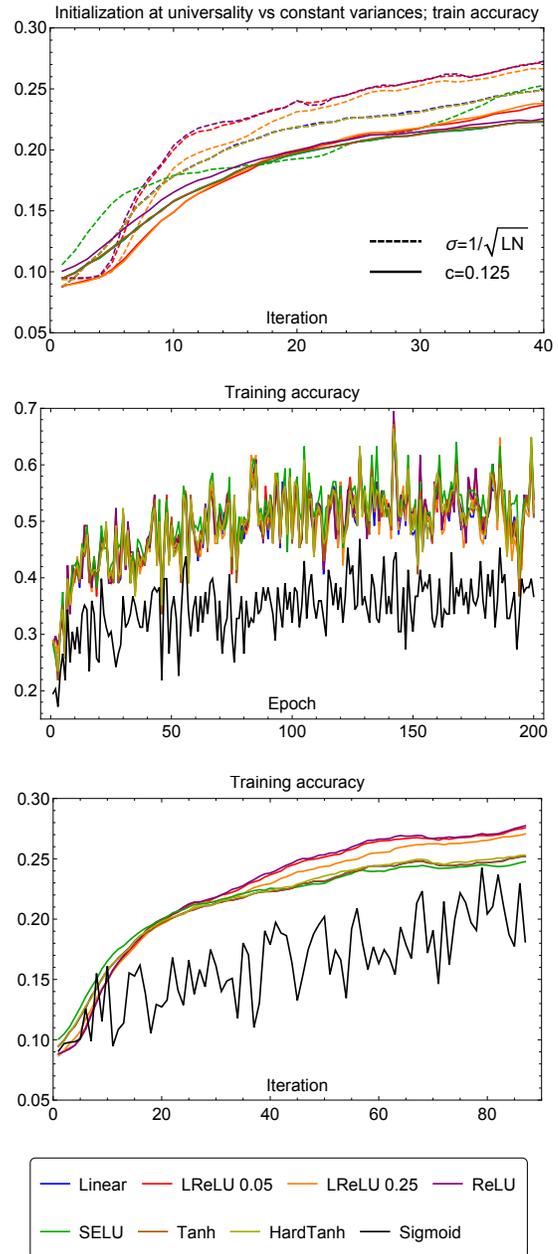


Figure 4: Training accuracy during first 200 epochs (middle) and first 100 iterations (bottom) of residual networks with various activation functions. The weight initialization was chosen for each activation function in such a way that the effective cumulant is $c = 0.125$. In the top panel, the dynamics with this initialization was juxtaposed with analogous training of networks in which the variance of weights was chosen to be $\frac{1}{LN}$ for all activation functions. We used leaky ReLU with $\alpha = 0.05, 0.25$.

ical isometry at the initialization for the subsequent optimization.

We consider the same set of experiments as in the previous section, and follow a similar training protocol to (He et al., 2016). We train for 200 epochs and drop the learning rate by a factor of 10 after epochs 80, 120, and 160. We use batch-size 128 and a starting learning rate of either 10^{-3} or 10^{-4} ¹.

First, we look at the learning dynamics on the training set. We can observe that most of the activation functions exhibit similar training accuracy evolution, see Fig. 4 (middle). Using the sigmoid activation, led however to significantly slower optimization. This is due to a faster growth of the variances of post- and pre- activations (which can be observed in Fig. 5), which exacerbates the neuron saturation problem.

Overall our results suggest that the singular spectrum of \mathbf{J} at initialization does not fully determine generalization and training performance. Nonetheless, setting the same effective cumulant for the experiments with different activation functions results in a markedly coinciding behavior of neural networks using activation functions of similar characteristics. This is in contrast to a setup in which the variances of the weight matrix entries are set to be equal. To demonstrate this we run another set of training experiments, this time with all standard deviations $\sigma = 1/\sqrt{LN}$. The plots depicting the full results are relegated to Fig. 7 in Appendix F. Here, in Fig. 4 (top) we showcase the training accuracy during the first 40 iterations for these two setups (excluding, for clarity, the networks with the sigmoid activation function). With different effective cumulants, the network learning dynamics, differs among experiments with different activation functions, especially at the beginning of learning.

This suggests that the spectrum of the input-output Jacobian at initialization can be treated as a confounding variable in such experiments. Ensuring that the level of dynamic isometry, and hence the value of the effective cumulant is kept the same, provides the possibility of a more meaningful comparison of the effect of activation functions on learning dynamics.

5 Synopsis and discussion

The main focus of this paper was the singular spectrum of the input-output Jacobian of a simple model of residual neural networks. We have shown that in the large network depth limit, it is described by a single, universal equation. This holds irrespective of the activation function used, for biunitarily invariant weight initializations matrices, a set covering Gaussian and scaled orthogonal initialization

¹We use relatively low learning rates, largely because we omit batch normalization layers in the architecture.

schemes. The singular value density depends on a single parameter called the effective cumulant, which can be calculated by considering the propagation of information in the network, via a dynamical mean field theory approach. This parameter depends on the activation function used, variance of biases and the entries of the weight matrices, and, for some activation functions, also on the depth of the network. We demonstrated the validity of our theoretical results in numerical experiments, both by generating random matrices adhering to the assumptions of the model and by evaluating the Jacobians of residual networks (at initialization) on the CIFAR10 dataset.

For a given activation function and/or network depth, it is always possible to set the weight matrix entries variances in such a way, that the resulting singular spectra of the Jacobians not only fulfill the conditions for dynamical isometry, but also are exactly the same, irrespective of the activation function used. This observation allows us to eliminate the singular spectrum of the Jacobian treated as a confounding factor in experiments with the learning process of simple residual neural networks for different activation functions. As an example of how this approach can be applied, we examined how accuracies of simple residual neural networks, employing a variety of activation functions, change during the learning process. When using the same variances of weight matrices entries, the learning curves of similar activation functions differed between each other more than when the networks were initialized with the same input-output Jacobian spectra. This allows, in our opinion, for a more meaningful comparison between the effects of choosing the activation function. We hope this observation will help with the research of deep neural networks.

Acknowledgements

PW would like to acknowledge the funding of the Polish National Science Centre, through the project SONATA number 2016/21/D/ST2/01142. WT appreciates the financial support from the Polish Ministry of Science and Higher Education through "Diamond Grant" 0225/DIA/2015/44. SJ was supported by the Polish National Science Centre through ETIUDA stipend No. 2017/24/T/ST6/00487.

References

- Akemann, G., Baik, J., and Di Francesco, P. (2011). *The Oxford handbook of random matrix theory*. Oxford University Press.
- Balduzzi, D., Frean, M., Leary, L., Lewis, J., Ma, K. W.-D., and McWilliams, B. (2017). The shattered gradients problem: If resnets are the answer, then what is the question? *arXiv preprint arXiv:1702.08591*.
- Canziani, A., Paszke, A., and Culurciello, E. (2016). An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*.

- Carmona, R., Cerenzia, M., and Palmer, A. Z. (2018). The dyson game. *arXiv preprint arXiv:1808.02464*.
- Chatterjee, S. and Meckes, E. (2007). Multivariate normal approximation using exchangeable pairs. *arXiv preprint math/0701464*.
- Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. (2015). The loss surfaces of multilayer networks. In *Artificial Intelligence and Statistics*, pages 192–204.
- Collins, B. and Śniady, P. (2006). Integration with respect to the haar measure on unitary, orthogonal and symplectic group. *Communications in Mathematical Physics*, 264(3):773–795.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256.
- Gudowska-Nowak, E., Janik, R. A., Jurkiewicz, J., and Nowak, M. A. (2003). Infinite products of large random matrices and matrix-valued diffusion. *Nuclear Physics B*, 670(3):479–507.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Bach, F. R. and Blei, D. M., editors, *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 448–456. JMLR.org.
- Janik, R. A., Nowak, M. A., Papp, G., Wambach, J., and Zahed, I. (1997). Non-hermitian random matrix models: Free random variable approach. *Physical Review E*, 55(4):4100.
- Janik, R. A. and Wieczorek, W. (2004). Multiplying unitary random matrices—universality and spectral properties. *Journal of Physics A: Mathematical and General*, 37(25):6521.
- Klambauer, G., Unterthiner, T., Mayr, A., and Hochreiter, S. (2017). Self-normalizing neural networks. In *Advances in Neural Information Processing Systems*, pages 971–980.
- Krizhevsky, A. (2009). Learning Multiple Layers of Features from Tiny Images. Master’s thesis.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436.
- Liao, Z. and Couillet, R. (2018). The dynamics of learning: A random matrix approach. *arXiv preprint arXiv:1805.11917*.
- Ling, Z. and Qiu, R. C. (2018). Spectrum concentration in deep residual learning: a free probability approach. *arXiv preprint arXiv:1807.11694*.
- Louart, C., Liao, Z., Couillet, R., et al. (2018). A random matrix approach to neural networks. *The Annals of Applied Probability*, 28(2):1190–1248.
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. (2013). Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3.
- Mingo, J. A. and Speicher, R. (2017). *Free probability and random matrices*, volume 4. Springer.
- Mishkin, D. and Matas, J. (2015). All you need is a good init. *arXiv preprint arXiv:1511.06422*.
- Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.
- Nica, A. and Speicher, R. (1996). *r*-diagonal pairs—a common approach to haar unitaries and circular elements. *arXiv preprint funct-an/9604012*.
- Nica, A. and Speicher, R. (2006). *Lectures on the combinatorics of free probability*, volume 13. Cambridge University Press.
- Nowak, M. A. and Tarnowski, W. (2017). Complete diagrammatics of the single-ring theorem. *Physical Review E*, 96(4):042149.
- Pennington, J. and Bahri, Y. (2017). Geometry of neural network loss surfaces via random matrix theory. In *International Conference on Machine Learning*, pages 2798–2806.
- Pennington, J., Schoenholz, S., and Ganguli, S. (2017). Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. In *Advances in Neural Information Processing Systems*, pages 4785–4795.
- Pennington, J., Schoenholz, S. S., and Ganguli, S. (2018). The emergence of spectral universality in deep networks. *arXiv preprint arXiv:1802.09979*.
- Pennington, J. and Worah, P. (2017). Nonlinear random matrix theory for deep learning. In *Advances in Neural Information Processing Systems*, pages 2637–2646.
- Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., and Ganguli, S. (2016). Exponential expressivity in deep neural networks through transient chaos. In *Advances in neural information processing systems*, pages 3360–3368.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. (2013). Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. (2016). Deep information propagation. *arXiv preprint arXiv:1611.01232*.

- Taki, M. (2017). Deep residual networks and weight initialization. *arXiv preprint arXiv:1709.02956*.
- Veit, A., Wilber, M. J., and Belongie, S. (2016). Residual networks behave like ensembles of relatively shallow networks. In *Advances in Neural Information Processing Systems*, pages 550–558.
- Voiculescu, D. V., Dykema, K. J., and Nica, A. (1992). *Free random variables*. Number 1. American Mathematical Soc.
- Wigner, E. P. (1993). Characteristic vectors of bordered matrices with infinite dimensions i. In *The Collected Works of Eugene Paul Wigner*, pages 524–540. Springer.
- Xiao, L., Bahri, Y., Sohl-Dickstein, J., Schoenholz, S. S., and Pennington, J. (2018). Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. *arXiv preprint arXiv:1806.05393*.
- Yang, G. and Schoenholz, S. (2017). Mean field residual networks: On the edge of chaos. In *Advances in Neural Information Processing Systems*, pages 7103–7114.
- Zhang, H., Dauphin, Y. N., and Ma, T. (2019). Residual learning without normalization via better initialization. In *International Conference on Learning Representations*.