
The Future is Log-Gaussian: ResNets and Their Infinite-Depth-and-Width Limit at Initialization

Mufan (Bill) Li*
University of Toronto,
Vector Institute

Mihai Nica*
University of Guelph,
Vector Institute

Daniel M. Roy
University of Toronto,
Vector Institute

Abstract

Theoretical results show that neural networks can be approximated by Gaussian processes in the infinite-width limit. However, for fully connected networks, it has been previously shown that for any fixed network width, n , the Gaussian approximation gets worse as the network depth, d , increases. Given that modern networks are deep, this raises the question of how well modern architectures, like ResNets, are captured by the infinite-width limit. To provide a better approximation, we study ReLU ResNets in the infinite-depth-and-width limit, where *both* depth and width tend to infinity as their ratio, d/n , remains constant. In contrast to the Gaussian infinite-width limit, we show theoretically that the network exhibits log-Gaussian behaviour at initialization in the infinite-depth-and-width limit, with parameters depending on the ratio d/n . Using Monte Carlo simulations, we demonstrate that even basic properties of standard ResNet architectures are poorly captured by the Gaussian limit, but remarkably well captured by our log-Gaussian limit. Moreover, our analysis reveals that ReLU ResNets at initialization are hypoactivated: fewer than half of the ReLUs are activated. Additionally, we calculate the interlayer correlations, which have the effect of exponentially increasing the variance of the network output. Based on our analysis, we introduce *Balanced ResNets*, a simple architecture modification, which eliminates hypoactivation and interlayer correlations and is more amenable to theoretical analysis.

1 Introduction

The characterization of infinite-width dynamics of gradient descent (GD) in terms of the so-called Neural Tangent Kernel (NTK) [1–10] represented a major breakthrough in our understanding of deep learning in the large-width regime. Before the identification of infinite-width limits, the theoretical study of deep learning had long been hindered by the apparent analytical intractability of gradient descent and variants acting on the nonconvex objectives used to train neural networks. Despite this progress, evidence suggests that deep neural networks can outperform their infinite-width limits in practice [11], particularly when the depth of the network is large. These observations motivate the study of other approximations that may close the gap.

Several alternative limits have been proposed. Around the time of the discovery of the NTK limit, mean-field limits were also characterized [12–15], and more recently have been linked with the NTK limit [16]. Yang and Hu [17] describe a family of infinite-width limits indexed by the scaling limits of initial weight variance, weight rescaling, and learning rates. This family includes both the NTK and mean field limits. One motivation for studying these alternative limits is that they yield a notion of feature learning, which provably does not occur in the NTK limit [17].

*Equal contribution authors.

Correspondence: mufan.li@mail.utoronto.ca; nicam@uoguelph.ca; daniel.roy@utoronto.ca.

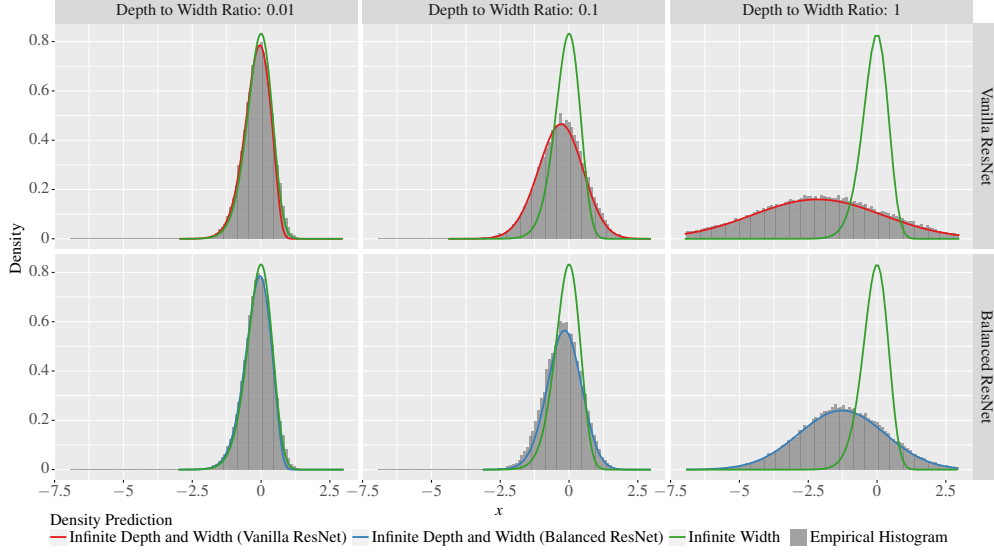


Figure 1: **Probability density function of $\ln \|z^{\text{out}}\|^2$** for six network configurations on initialization. All networks have $n = 100$, $n_{\text{in}} = n_{\text{out}} = 10$, $\alpha = \lambda = 1/\sqrt{2}$. Depth d varies by column. Top row: Vanilla ReLU ResNets. Bottom Row: Balanced ResNets (Section 2.1), which randomize the nonlinearities used at each neuron thereby exponentially reducing the variance. The theoretical curves shown are the result of Theorem 1 and Theorem 4. These predictions converge to the infinite-width prediction when $d/n \rightarrow 0$. The details of the simulations and plot here is in Appendix C.

Despite the variety of these limits, one common feature is that the *depth* of the network (i.e., the number of its layers) is treated as a constant as the width of the network is allowed to grow. Indeed, for fixed width, n , the Gaussian approximations at initialization [18–23] worsen as the depth, d , increases. While real-world networks are fairly wide, their relative depth is not trivial. Hanin and Nica [24] were the first to compute an infinite-depth-and-width limit for fully connected networks. While the training dynamics of this limit are still not completely understood, we now know that, in the infinite-depth-and-width limit, the neural tangent kernel is random, and the derivative of the kernel is nonzero at initialization. This means training does not correspond to that of a linear model, like it does in the NTK limit [25, 26].

In addition to these theoretical corrections to the Gaussian limit, practitioners have begun to notice that even basic properties of standard neural networks do not line up with those predicted by the infinite-width limit. Perhaps the most basic is that the gradients of the network early in training are not Gaussian, but instead are approximately log-Gaussian [27]. In fact, one should see log-Gaussian behaviour agrees with the theoretical predictions of Hanin and Nica [24], although there has yet to be a careful empirical comparison made between the precise predictions coming from infinite-depth-and-width models and real world networks.

In practice, however, fully connected networks are not often used without architectural modifications. In particular, residual connections, ushered into widespread use after the description of the ResNet architecture [28], produced very deep architectures that were practically useful with optimization techniques available at the time. Initialization schemes for ResNets have been studied in the infinite-width limit [19, 29] or with modifications [30–32].

In this work, we consider the infinite-depth-and-width limit of fully connected architectures with residual connections (“Vanilla ResNets”) and standard initializations. Here, the analysis is complicated by the effect of skip connections, which introduce interlayer correlations that have a non-negligible effect in this limit. Surprisingly, we observe a counter-intuitive but fundamental phenomenon, whereby these skip connection cause the network to be **hypoactivated**, meaning that less than half of the neurons are activated on initialization. This fact undermines key assumptions that underlie other infinite-depth-and-width limit studies of architectures without residual connections or with nonstandard modifications, such as post-activation residual connections.

Hypoactivation is a roadblock that all theoretical research into standard ResNet architectures must contend with—it is an unavoidable property of the standard architecture and the root of many technical difficulties. In order to sidestep this roadblock, we introduce a conjecture that bounds the size of hypoactivation and effect on interlayer correlation. The conjecture is inspired by empirical evidence from Monte Carlo simulations, as well as several simplified analyses that ignore certain technical difficulties. The conjecture introduces what we believe to be the minimal assumption necessary to allow rigorous theoretical work to proceed. It also defines an important open problem in the study of limits of residual architectures.

To demonstrate the utility of the conjecture, we show that it leads to precise predictions. In particular, we prove a limit theorem characterizing the exact marginal distribution of the output at initialization in the infinite-depth-and-width limit, up to order $O(dn^{-2})$. Our limit result shows that ResNets have *log-Gaussian behaviour* on initialization, and like fully connected networks [24], the behaviour is determined by the depth-to-width aspect ratio. This corroborates recent empirical observations about deep ResNets [27]. Since real world networks are finite, the question of how well this approximates finite behaviour is of paramount importance. Based on Monte Carlo simulations, we find excellent agreement between our predictions and finite networks (see Figure 1). Moreover, for very deep networks (e.g. $d/n = 1$) the infinite-depth-and-width prediction is extremely different than the infinite-width prediction. More surprisingly, however, is that even at comparatively small depth-to-width ratio (e.g. $d/n = 0.1$) the two limits are already significantly different. Furthermore, we also observe that the effects due to hypoactivation and interlayer correlation are non-negligible; these effects are precisely the difference between Vanilla ResNets and so-called “Balanced Resnets” in Figure 1. Perhaps most importantly, we observe that real network outputs exhibit exponentially larger variance than predicted by infinite-width limits. This type of variance at initialization is known to cause exploding and vanishing gradients and other types of training failures [33, 34]. Our result also implies that the output neurons of the network are not independent as predicted by the Gaussian infinite-width limit. See Figure 3.

In order to maintain the same skip connections from layer to layer, but render the activation patterns completely independent from layer to layer on initialization, we introduce the Balanced ResNet architecture, where the sign of each neuron’s activation function is randomized. We demonstrate that this exponentially decreases the variance of the network output on initialization. Moreover, this independence between neurons also makes the model more amenable to theoretical analysis and opens the door to future understanding of network behaviour. Although it is beyond the scope of this paper, a small preliminary empirical study (Appendix C) suggests that standard training regimes are not negatively affected by replacing standard ResNet architectures with Balanced ones.

We summarize our main contributions as follows:

- We identify and characterize a fundamental property of ResNets that we call **hypoactivation**: less than half of the ReLU neurons are activated. Based on empirical evidence, we formulate a precise minimal conjecture bounding the effect of hypoactivation that permits us to make precise, rigorous estimates for other properties of ResNets.
- We prove a limit theorem which shows that the output of ResNets on initialization exhibits **log-Gaussian** behaviour with parameterization depending on the depth-to-width ratio d/n .
- We provide **empirical evidence** from Monte Carlo simulations showing our theory provides more accurate predictions for simple properties of finite networks compared to the predictions made by the Gaussian infinite-width limit.
- We introduce the **Balanced ResNet architecture**, which corrects the hypoactivation and variance due to layerwise correlation from Vanilla ResNets. We also prove that the output for this architecture is log-Gaussian on initialization with exponentially lower variance. This simple modification can be applied to *any* neural network that uses ReLU activations.

2 Main Results

In terms of the notation in Table 1, a **Vanilla ResNet** with fully connected first/last layers and d hidden layers of width n is defined by

$$z^0 := \frac{1}{\sqrt{n_{\text{in}}}} W^0 x, \quad z^\ell := \alpha z^{\ell-1} + \lambda \sqrt{\frac{2}{n}} W^\ell \varphi_+(z^{\ell-1}) \text{ for } 1 \leq \ell \leq d, \quad z^{\text{out}} := \frac{1}{\sqrt{n}} W^{\text{out}} z^d. \quad (1)$$

Notation	Description	Notation	Description	Table 1: Notation
$n_{\text{in}} \in \mathbb{N}$	Input dimension	$n_{\text{out}} \in \mathbb{N}$	Output dimension	
$n \in \mathbb{N}$	Hidden layer width	$d \in \mathbb{N}$	Number of hidden layers (depth)	
$\varphi_+(\cdot)$	ReLU function $\varphi_+(x) = \max\{x, 0\}$	$\varphi_-(\cdot)$	“Domain Flipped” ReLU $\varphi_-(x) = \max\{-x, 0\}$	
$\alpha \in \mathbb{R}$	Skip connection coefficient	$\lambda \in \mathbb{R}^+$	Feed-forward coefficient	
$x \in \mathbb{R}^{n_{\text{in}}}$	Input	$W^0 \in \mathbb{R}^{n_{\text{in}} \times n}$	Weight matrix at layer 0	
$z^{\text{out}} \in \mathbb{R}^{n_{\text{out}}}$	Network output	$W^{\text{out}} \in \mathbb{R}^{n \times n_{\text{out}}}$	Weight matrix at final layer.	
$z^\ell \in \mathbb{R}^n$	Neurons (pre-activation) for layer $1 \leq \ell \leq d$	$W^\ell \in \mathbb{R}^{n \times n}$	Weight matrix at layer $1 \leq \ell \leq d$ all weights initialized i.i.d. $\sim \mathcal{N}(0, 1)$	

Note that factors of $\sqrt{2n^{-1}}$ in the hidden layer are equivalent to initializing according to the so-called He initialization [35]. Other initializations correspond to changing the coefficient λ . This setup is similar to that of “Stable ResNets” [29], where the infinite-width limit is studied.

In the infinite-depth-and-width limit, the intuition that half of the ReLU units are active (i.e. nonzero) because of symmetry is surprisingly not correct. We find that the following quantities play an important role. We define the **average hypoactivation (of layer ℓ)** and the **total hypoactivation of the network** by

$$h_\ell := \mathbf{E} \left[\|\varphi_+(\hat{z}^\ell)\|^2 \right] - \frac{1}{2}, \quad h_{\text{total}} := \sum_{\ell=1}^d h_\ell, \quad (2)$$

where $\hat{z}^\ell := z^\ell / \|z^\ell\|$ and \mathbf{E} means expectation over the choice of random network weights on initialization. The average hypoactivation is a measure of how many ReLU neurons are activated in layer ℓ ; $h_\ell = 0$ indicates roughly half of the neurons are active. Counter-intuitively, we observe that in a vanilla ResNet, h_ℓ is *negative* and $|h_\ell| = O(1/n)$, indicating slightly less than half the neurons are active.² After compounding over d layers, the total hypoactivation is of order $h_{\text{total}} = O(d/n) = O(1)$ in the infinite-depth-and-width limit. As we will see, this effect has a non-trivial contribution.

At the same time, we also find the covariance between the activations of various layers does not vanish in the infinite-depth-and-width limit. This motivates the definition of the **total interlayer covariance correction**

$$I_{\text{total}} := \sum_{1 \leq \ell \neq \ell' \leq d} \mathbf{Cov} \left(2\|\varphi_+(\hat{z}^\ell)\|^2, 2\|\varphi_+(\hat{z}^{\ell'})\|^2 \right). \quad (3)$$

As with the hypoactivation, skip connections cause this term to be non-trivial. We formulate Conjecture 5, which contains a precise encapsulation of the behaviour of \hat{z}^ℓ we observe.

Conjecture 5 (Informal). In expectation, the layers \hat{z}^ℓ can be approximated by uniform random variables from the sphere up to a relative error of $O(1/n)$.

The conjecture is well supported by Monte-Carlo simulations (see Figure 4). We provide a more detailed discussion and a precise statement of Conjecture 5 in Section 4. Assuming the conjecture holds, we prove a limit theorem about the distribution of z^{out} . Informally, this says that z^{out} is **approximately a log-Gaussian scalar** times an independent Gaussian vector

$$z^{\text{out}} \approx \frac{\|x\|}{\sqrt{n_{\text{in}}}} (\alpha^2 + \lambda^2)^{\frac{d}{2}} \exp \left(\frac{1}{2} \mathcal{N} \left(-\frac{\beta}{2} + 2ch_{\text{total}}, \beta + c^2 I_{\text{total}} \right) \right) \vec{Z}, \quad (4)$$

where \vec{Z} has iid $\mathcal{N}(0, 1)$ entries, and β and c are defined by

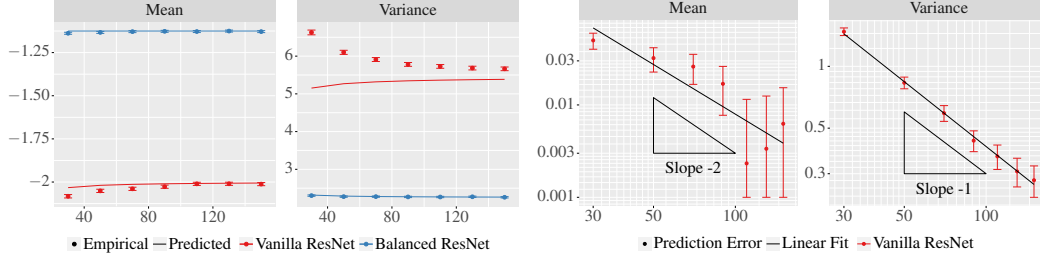
$$\beta := \frac{2}{n} + \frac{d}{n} \cdot \frac{5\lambda^4 + 4\alpha^2\lambda^2}{(\alpha^2 + \lambda^2)^2}, \quad c := \frac{\lambda^2}{\alpha^2 + \lambda^2}. \quad (5)$$

The precise statement, including asymptotic error bounds, is as follows.

Theorem 1. *For any choice of hyperparameters $n_{\text{in}}, n_{\text{out}}, n, d, \alpha, \lambda$, and every input x , the output z^{out} at initialization has a marginal distribution which can be written in the form*

$$z^{\text{out}} \stackrel{d}{=} \frac{\|x\|}{\sqrt{n_{\text{in}}}} (\alpha^2 + \lambda^2)^{\frac{d}{2}} \exp \left(\frac{1}{2} G \right) \vec{Z}, \quad (6)$$

²For a quantity $f = f(n, d)$ whose dependence on width and depth may be implicit, we use the notation $f = O(d^a/n^b)$ to mean that, for all choice of constants $\alpha, \lambda, r_-, r_+ > 0$, there exists a constant $C > 0$ such that $|f(n, d)| \leq Cd^a/n^b$ for all d, n where $r_- < d/n < r_+$. This notation will allow us to state precise limit theorems when $d, n \rightarrow \infty$ with the ratio d/n converging to a constant.



(a) \mathbf{E} and \mathbf{Var} of G from (6) and prediction.

(b) Prediction error for \mathbf{E} and \mathbf{Var} of G from (6).

Figure 2: Empirical mean and variance and infinite-depth-and-width prediction for the random variable $G(n, d, \alpha, \lambda)$ from (6) compared to the results of Theorems 1 & 4. For these simulations, $\alpha = \lambda = 1/\sqrt{2}$, $d/n = 1$ are fixed and width n is varied on the x -axis. The error bars indicate 95% confidence interval (CI) of the Monte-Carlo simulation, with truncation at $1e-3$ for plotting on log-scale. The balanced ResNet predictions fall within the Monte-Carlo CI and so are not plotted on log scale. Since d/n is fixed, Figure 2b indicates the error in the asymptotic prediction of variance (7) is $O(n^{-1}) = O(dn^{-2})$ as claimed. On the other hand, the error in mean (7) is $O(n^{-2}) = O(dn^{-3})$; this is one order smaller than the statement proven in Theorem 1. A possible explanation is that the sub-leading error term in our approximation for G is mean zero.

where $\vec{Z} \in \mathbb{R}^{n_{out}}$ is a Gaussian random vector with iid $\mathcal{N}(0, 1)$ entries, $G = G(n, d, \alpha, \lambda)$ is a random variable which is independent of \vec{Z} and whose distribution does not depend on n_{in} , n_{out} or x .

Consider the limit where both the network depth $d \rightarrow \infty$ and hidden layer width $n \rightarrow \infty$ in such a way that the ratio d/n converges to a non-zero constant. In this limit, assuming that Conjecture 5 holds, then the random variable $G \in \mathbb{R}$ has the following asymptotic behaviour:

$$\mathbf{E}[G] = -\frac{\beta}{2} + 2ch_{total} + O\left(\frac{d}{n^2}\right), \quad \mathbf{Var}[G] = \beta + c^2 I_{total} + O\left(\frac{d}{n^2}\right), \quad (7)$$

where β and c are as in (5), and moreover G converges in distribution to a Gaussian random variable with mean and variance given by (7) in this limit.

The main ideas of the proof of Theorem 1 is given in Section 5 and the detailed proof is given in Appendix B. We also provide more explicit formulas for h_{total} , I_{total} below.

Proposition 2. Assume Conjecture 5 is true. Then in the same infinite-depth-and-width limit as Theorem 1, the total hypoactivation h_{total} and total interlayer covariance I_{total} obey

$$h_{total} = C_{\alpha, \lambda} \frac{d}{n} + O\left(\frac{d}{n^2}\right), \quad I_{total} = \sum_{1 \leq \ell \neq \ell' \leq d} \frac{\bar{J}_2(\theta_{|\ell' - \ell|}) - \bar{J}_2(\pi - \theta_{|\ell' - \ell|})}{n} + O\left(\frac{d}{n^2}\right). \quad (8)$$

Here $C_{\alpha, \lambda}$ is a constant depending on α, λ and

$$\bar{J}_2(\theta) := J_2(\theta) / \pi = 3 \sin(\theta) \cos(\theta) / \pi + (1 - \theta / \pi) (1 + 2 \cos^2 \theta),$$

where $J_2(\theta)$ first appeared in [36], and θ_k is such that $\cos(\theta_k) = \alpha^k / (\alpha^2 + \lambda^2)^{k/2}$.

Remark 3. The Gaussian infinite-width limit predicts that the marginals of z^{out} should have the form of (6) with G being identically zero. As z^{out} depends exponentially on G , the infinite-depth-and-width limit predicts the variance is exponentially larger than the infinite-width limit. See Section 3.1 for a detailed discussion and Figure 3 for verification against finite networks.

Using Monte Carlo simulations, we estimate the constant $C_{\alpha, \lambda}$. The result of Theorem 1 is then compared against finite networks in Figure 2. Full proof can be found Appendix B.

The idea of using a scaling parameter $\alpha = 1/\sqrt{2}$ in the skip connections has been noted in empirical papers [37] and also studied under simplified assumptions on the number of activation in each layer [38]. Our result shows that by choosing $\alpha^2 + \lambda^2 = 1$, the prefactor in (6) does not grow with depth thereby enabling deeper networks to be trained.

In the case that $\alpha = 0$, $\lambda = 1$, the architecture reduces to a fully connected network. In this case, log-Gaussian behaviour of the network was discovered in Hanin and Nica [24]. The fully connected case is simpler because the direction vectors \hat{z}^ℓ are always uniformly distributed on the unit sphere and are independent from layer to layer. This means $h_\ell = 0$ and $\theta_n = 0$ which greatly simplifies the result of Theorem 1.

Furthermore, the proof easily extends to the case where the coefficients $\alpha_\ell, \lambda_\ell$ vary from layer to layer; see Appendix A.1 for the general statement. Hayou et al. [29] and Hanin and Rolnick [33] have studied ResNets where $\alpha_\ell = 1$ in every layer, but λ_ℓ is allowed to vary. The prefactor of our result in this case becomes $\prod_{\ell=1}^d (1^2 + \lambda_\ell^2) \approx \exp\left(\sum_{\ell=1}^d \lambda_\ell^2\right)$, which is akin to the behaviour found in those papers. Additionally, our result precisely quantifies the log-Gaussian behaviour and its dependence on the sequence λ_i through the parameters β and c in the infinite-depth-and-width limit.

2.1 Log-Gaussian Behaviour of Balanced ResNets

Given a collection of *iid uniformly random signs*, $s_i^\ell \in \{+, -\}$, $1 \leq \ell \leq d$, $1 \leq i \leq n$, a **balanced ResNet** is defined much like a Vanilla ResNet, except that a random sign is applied preactivation:

$$z^\ell := \alpha z^{\ell-1} + \lambda \sqrt{\frac{2}{n}} W^\ell \varphi_{s^\ell}(z^{\ell-1}) \text{ for } 1 \leq \ell \leq d,$$

where at each layer the vector function $\varphi_{s^\ell} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ applies either φ_+ or φ_- to the entries according to the random signs s^ℓ . More precisely, the i -th component is

$$\varphi_{s^\ell}(z)_i := \begin{cases} \varphi_+(z_i) = \max(z_i, 0), & \text{if } s_i^\ell = +, \\ \varphi_-(z_i) = \max(-z_i, 0), & \text{if } s_i^\ell = -. \end{cases}$$

An equivalent definition is the entrywise multiplication $\varphi_{s^\ell}(z) = \varphi_+(s^\ell \odot z)$. Note that the random signs s_i^ℓ are *not* trainable parameters; they are frozen on initialization. This same symmetrization was first exploited by Allen-Zhu et al. [39] and Bai and Lee [40] to study a quadratic approximation of the network. We now present a corresponding limiting theorem for Balanced ResNets.

Theorem 4. *For a balanced ResNet, the same result as Theorem 1 given in (6) still holds, but with the mean and variance of G given simply by*

$$\mathbf{E}[G] = -\frac{\beta}{2} + O\left(\frac{d}{n^2}\right), \quad \mathbf{Var}[G] = \beta + O\left(\frac{d}{n^2}\right). \quad (9)$$

Balanced ResNets are constructed so that the activation of each neuron is independent of all others due to the random signs s^ℓ . This eliminates the hypoactivation and variance terms which complicated the analysis of the Vanilla ResNet and necessitated Conjecture 5. Instead, for Balanced ResNets it is straightforward to compute that for any fixed $z, w \in \mathbb{R}^n$ we have

$$\mathbf{E}[\|\varphi_{s^\ell}(z)\|^2] = \frac{\|z\|^2}{2}, \quad \mathbf{Var}[\|\varphi_{s^\ell}(z)\|^2] = \sum_{i=1}^n \frac{z_i^4}{4}, \quad \mathbf{Cov}[\|\varphi_{s^\ell}(z)\|^2, \|\varphi_{s^{\ell'}}(w)\|^2] = 0. \quad (10)$$

Even though the layers $\hat{z}^\ell, \hat{z}^{\ell'}$ are correlated, because the activation functions φ_{s^ℓ} and $\varphi_{s^{\ell'}}$ are set to be independent on initialization, the correlation between layers does not induce a correlation on which neurons are activated from layer to layer. This explains why there is no hypoactivation and interlayer correlation correction in Theorem 4 as there is in Theorem 1.

3 Consequences of Theorems 1 & 4 and Comparison to Infinite-Width Limit

3.1 Vanishing and Exploding Norms

By the basic fact $\mathbf{E}[\exp(\mathcal{N}(\mu, \sigma^2))] = \exp(\mu + \frac{1}{2}\sigma^2)$ it follows from Theorems 1 & 4 that, when the inputs x has $\|x\| = \sqrt{n_{\text{in}}}$, the mean size scale of any neuron z_i^{out} is approximately

$$\mathbf{E}[(z_i^{\text{out}})^2] \approx \begin{cases} (\alpha^2 + \lambda^2)^d \exp(2ch_{\text{total}} + \frac{1}{2}c^2 I_{\text{total}}), & \text{for Vanilla ResNets,} \\ (\alpha^2 + \lambda^2)^d, & \text{for Balanced ResNets.} \end{cases}$$




Figure 3: Behaviour of the mean and variance of the typical output $(z_i^{\text{out}})^2$ and the correlation between $(z_i^{\text{out}})^2, (z_j^{\text{out}})^2$ for two different output neurons. Here $n = 200$, $\alpha = \lambda = 1/\sqrt{2}$ and d/n varies on the x -axis. The infinite width prediction for correlation is zero, but due to plotting on log-scale, we replace zero by 0.01 for display.

(Note that the terms with β cancel out!) When $\alpha^2 + \lambda^2 = 1$, this is constant for Balanced ResNets. In contrast, Vanilla ResNets have a complicated dependence on the network depth d and width n due to the *hypoactivation* and *correlations* terms. This means the behaviour is $\exp(Cd/n)$, which is somewhat surprising. A more serious issue is the variance which Theorems 1 & 4 predict to be

$$\text{Var}[(z_i^{\text{out}})^2] \approx \begin{cases} (\alpha^2 + \lambda^2)^{2d} (3 \exp(\beta + c^2 I_{\text{total}}) - 1) \exp(4ch_{\text{total}} + c^2 I_{\text{total}}), & \text{for Vanilla,} \\ (\alpha^2 + \lambda^2)^{2d} (3 \exp(\beta) - 1), & \text{for Balanced.} \end{cases}$$

Since $\beta \approx Cd/n$, the term $\exp(\beta)$ represents exponentially larger variance for deep nets compared to shallow ones of the same width. In contrast, the variance predicted by the infinite width limit does not grow with depth for this model when $\alpha^2 + \lambda^2 = 1$. This effect means the relative sizes of different network outputs can be widely disparate. Unlike problems with the mean, this issue is harder to resolve. For example, normalization methods that divide all neurons by a constant does nothing to address the large relative disparity between two points. Techniques like batch normalization will be skewed by large outliers. This kind of variance is known to obstruct training [33]. The input-output derivative $\partial_{x_i} z^{\text{out}}$ has the same type of behaviour as z^{out} itself; a simple proof is given in Appendix B. It is expected that the gradient with respect to the weights $\partial_{W_{ij}^\ell} z^{\text{out}}$ will also have the same qualitative behaviour [27] although more investigation is needed to understand this theoretically. Exponentially large variance for gradients is a manifestation of the vanishing-and-exploding gradient problem [34].

Balanced ResNets suffer from this variance problem less because the interlayer correlation term is zero. Since this variance reduction happens at the exponential scale in, the difference can be significant; for networks with $d/n = 1$, the contribution is a factor of $\approx e^{5.5} \approx 250$ for Vanilla ResNets vs. $\approx e^{2.5} \approx 10$ for Balanced ResNets. See Figure 3 for a comparison of these theoretically predicted properties vs experiments with finite networks.

3.2 Correlated Output Neurons

Since the same random variable G multiplies the entire vector z^{out} , the individual neurons in the output layer are *not* independent. For example, Theorem 1 and 4 predict that the squared entries have strictly positive correlation given by $\text{Corr}((z_i^{\text{out}})^2, (z_j^{\text{out}})^2) = (\exp(\sigma^2) - 1)/(3\exp(\sigma^2) - 1)$ for any two neurons $i \neq j$ where $\sigma^2 = \text{Var}(G)$. This tends to $1/3$ as d/n grows. The effect of correlated output neurons persists for Balanced ResNet but is reduced again due to the lower variance. This is very different from the infinite-width limit, which predicts that individual neurons should be independent Gaussians. This prediction of the theorem matches finite networks closely; see Figure 3.

4 Conjecture 5: Hypoactivation and Layerwise Correlations

Vanilla ResNets have a subtle asymmetry in the architecture due to the skip connections. Unlike fully connected networks, the distribution of $\hat{z}^\ell := z^\ell / \|z^\ell\|$ for $\ell \geq 1$ is *not* exactly uniformly distributed

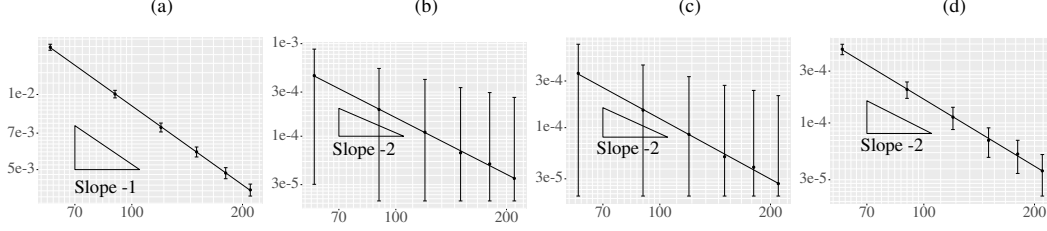


Figure 4: Monte Carlo evidence for Conjecture 5. $d = 200$, $\alpha = \lambda = 1/\sqrt{2}$, and n varies on the x -axis. The plots show the quantities (a) $|\mathbb{E}\|\varphi_+(\hat{z}^\ell)\|^2 - \mathbb{E}\|\varphi_+(u)\|^2|$, (b) $|\mathbf{Var}\|\varphi_+(\hat{z}^\ell)\|^2 - \mathbf{Var}\|\varphi_+(u)\|^2|$, (c) $|\mathbf{Cov}(\|\varphi_+(\hat{z}^\ell)\|^2, \varphi_+(\hat{z}^{\ell-1})\|^2) - C(\theta_1)|$, (d) $|\mathbf{Cov}(\|\varphi_+(\hat{z}^\ell)\|^2, \varphi_+(\hat{z}^{\ell-2})\|^2) - C(\theta_2)|$. $C(\theta_k)$ is the theoretical covariance formula from the term where $\ell' - \ell = k$ in (8). Note also that the absolute error is expected to be $O(n^{-2})$ when the theoretical quantity is $O(n^{-1})$. Figures (a) and (b) verify the conjecture in (11) and (12), and Figures (c) and (d) verify (13) when $k = 1$ and $k = 2$ respectively. For display, we clip the bottom edge of the CI to $2e-5$; otherwise the error bar would go down to $-\infty$ on the log scale.

on the unit sphere. Informally speaking, \hat{z}^ℓ can be thought of as a random walk whose variance at each step is proportional to $\|\varphi_+(\hat{z}^\ell)\|$. More randomness is injected when $\|\varphi_+(\hat{z}^\ell)\|$ is large and less when it is small. The net effect is that the walk moves slower when $\|\varphi_+(\hat{z}^\ell)\|$ is small, thereby spending more time in those locations. Consequently, $\|\varphi_+(\hat{z}^\ell)\|$ is biased toward smaller values.

The size of this effect is limited by entropy; most of the unit sphere \mathbb{S}^{n-1} has $\|\varphi_+(u)\|^2 \approx 1/2$ in the sense that for any $\epsilon > 0$, the measure of the set $\{u : \|\varphi_+(u)\|^2 - \frac{1}{2} > 1/n^{\frac{1}{2}-\epsilon}\}$ vanishes as $n \rightarrow \infty$.

To quantify the effect of the bias, we can think of the evolution of \hat{z}^ℓ , $\ell = 1, 2, \dots$ as a random walk that takes different step sizes at a different points on the sphere. It is reasonable to expect that the behaviour of this processes will be similar to that of a time changed Brownian motion, which is slowed down at the points where \hat{z}^ℓ takes smaller steps. (Proving this comparison precisely is technically difficult since the parameter n simultaneously plays both the role dimension and the step size of the walk.) Based on this heuristic comparison to time changed Brownian motion and on extensive Monte Carlo simulations we conjecture that the expected size of the hypoactivation effect is only $O(1/n)$ in expectation; Conjecture 5 contains a precise statement.

Even if each layer \hat{z}^ℓ is marginally close to the uniform distribution on the unit sphere, the directions \hat{z}^ℓ and $\hat{z}^{\ell+1}$ are not independent because of the skip connections in the network. As above, the exact behaviour is complicated due to fluctuations in the exact number of neurons which are activated in each layer. However, using the idea that $\|\varphi_+(\hat{z}^\ell)\|^2 = \frac{1}{2}(1 + o(1))$, we construct the following approximation. From (16), we have the approximation $\hat{z}^{\ell+1}/\|\hat{z}^\ell\| = \alpha\hat{z}^\ell + \lambda g^{\ell+1}/\sqrt{n}(1 + o(1))$. We observe that the norm of RHS is concentrated around $\sqrt{\alpha^2 + \lambda^2}$ as $n \rightarrow \infty$, so normalizing this to get $\hat{z}^{\ell+1}$ we have

$$\hat{z}^{\ell+1} = \left(\frac{\alpha}{\sqrt{\alpha^2 + \lambda^2}} \hat{z}^\ell + \frac{\lambda}{\sqrt{\alpha^2 + \lambda^2}} \frac{g^{\ell+1}}{\sqrt{n}} \right) (1 + o(1)).$$

Iterating this gives the same relationship for $\hat{z}^{\ell+k}$ where the first coefficient becomes $\alpha^k/\sqrt{\alpha^2 + \lambda^2}^k$. As before, based on Monte Carlo simulations, we conjecture that the size of the error is $O(1/n)$ in expectation. We formalize this as a precise statement in Conjecture 5 below.

Conjecture 5. *The distribution of the unit vector $\hat{z}^\ell = z^\ell / \|z^\ell\|$ is approximately uniformly distributed from the unit sphere $u \in \mathbb{S}^{n-1}$ in the precise sense that the following asymptotics hold*

$$\mathbf{E} [\|\varphi_+(\hat{z}^\ell)\|^2] = \mathbf{E} [\|\varphi_+(u)\|^2] \left(1 + O\left(\frac{1}{n}\right) \right), \quad (11)$$

$$\mathbf{Var} [\|\varphi_+(\hat{z}^\ell)\|^2] = \mathbf{Var} [\|\varphi_+(u)\|^2] \left(1 + O\left(\frac{1}{n}\right) \right), \quad (12)$$

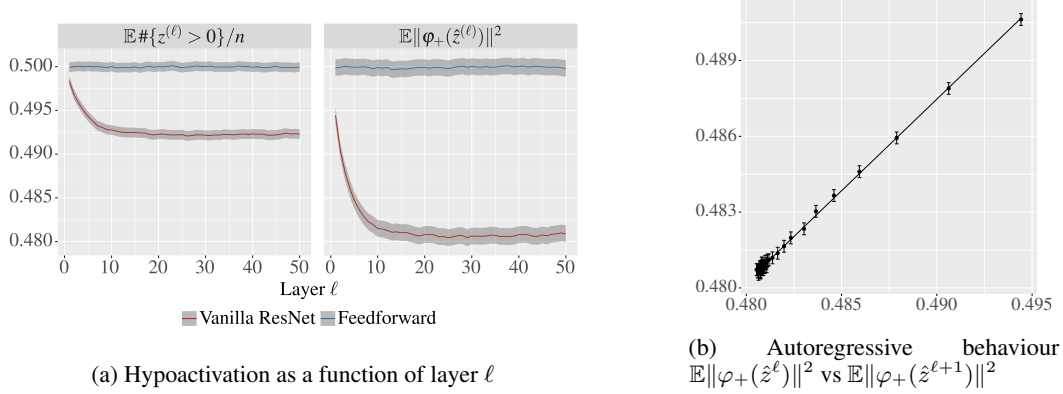


Figure 5: Monte-Carlo simulation for the behaviour of the unit vector \hat{z}^ℓ as a function of layer $0 \leq \ell \leq d$. Here $n = d = 50$, $\alpha = \lambda = \sqrt{2}^{-1}$. Figure 5a shows the mean fraction of neurons which are activated, $\mathbb{E} [\#\{i : z_i^\ell > 0\}] / n$ and the norm of the ReLU $\mathbb{E} [\|\varphi_+(\hat{z}^\ell)\|^2]$. The hypoactivation h_ℓ is how far this is from $\frac{1}{2}$. At layer $\ell = 0$, \hat{z}^ℓ is uniformly distributed from a unit sphere, but approaches a different steady-state as we go deeper into the network. Figure 5b shows evidence that the process $\mathbb{E} [\|\varphi_+(\hat{z}^\ell)\|^2]$ seems to be a linear function of the previous layer $\mathbb{E} [\|\varphi_+(\hat{z}^{\ell-1})\|^2]$. Figure 4 illustrates the dependence as n varies.

where the constants in the big $O(\cdot)$ notation are uniform in ℓ . Moreover, for two layers ℓ, ℓ' , which are $k \geq 1$ layers apart $|\ell' - \ell| = k$, the joint distribution of $\hat{z}^\ell, \hat{z}^{\ell'}$ is approximately equal to the joint distribution of $u, \cos(\theta_k)u + \sin(\theta_k)g/\sqrt{n}$ where g is a Gaussian vector with iid $\mathcal{N}(0, 1)$ entries which is independent of u and θ is such that $\cos(\theta_k) = \alpha^k/(\alpha^2 + \lambda^2)^{k/2}$ in the sense that the following asymptotics hold

$$\mathbf{Cov} \left[\|\varphi_+(\hat{z}^\ell)\|^2, \|\varphi_+(\hat{z}^{\ell'})\|^2 \right] = \mathbf{Cov} \left[\|\varphi_+(u)\|^2, \left\| \varphi_+(\cos(\theta_k)u + \frac{\sin(\theta_k)}{\sqrt{n}}g) \right\|^2 \right] \left(1 + O\left(\frac{1}{n}\right) \right), \quad (13)$$

where the constant in the big $O(\cdot)$ notation is uniform in ℓ, ℓ' .

See Figure 4 for Monte Carlo simulations empirically verifying the conjecture for a fixed depth d , and see Figure 5 for verifying the uniformity in layers ℓ . In particular, we observe that in Figure 5a, we can see the effect of hypoactivation converges rapidly to an equilibrium as the layer ℓ increases. In fact, we can further verify in Figure 5b that hypoactivation appears to be autoregressive, which implies the convergence is exponentially fast. This motivated the uniformity in layers in Conjecture 5.

5 Proof Ideas for Theorems 1 & 4

A key element of the proof is the following property of Gaussian random matrices. If W which has iid $\mathcal{N}(0, 1)$ entries, then for any vector x , we have

$$Wx \stackrel{d}{=} \|x\| g, \quad (14)$$

where g is a vector whose entries are iid $\mathcal{N}(0, 1)$ random variables. Because of the fully connected first and last layer of the network, (14) implies that

$$z^0 \stackrel{d}{=} \frac{\|x\|}{\sqrt{n_{\text{in}}}} g, \quad z^{\text{out}} \stackrel{d}{=} \frac{\|z^d\|}{\sqrt{n}} g'. \quad (15)$$

Hence $G := \ln \left((\|z^d\|^2 / n) \cdot (\|x\|^2 / n_{\text{in}})^{-1} \cdot (\alpha^2 + \lambda^2)^{-d} \right)$ only depends on n, d, α, λ . (Equivalently, G has the distribution of $\ln(\|z^d\|^2 / n \cdot (\alpha^2 + \lambda^2)^{-d})$ when $z^0 = g$.) With this definition, (15) also shows z^{out} is proportional to $\exp(G/2)$, establishing the first part of Theorem 1.

From this construction, the essence of the proof is to understand the distribution of $\|z^d\|$ when $z^0 = g$. To understand $\|z^d\|$, we look at the ratios $\|z^{\ell+1}\|/\|z^\ell\|$ layer by layer. By using the homogeneity property of ReLU $\varphi_+(|c|x) = |c|\varphi_+(x)$, we can divide $z^{\ell+1}$ from (1) by $\|z^\ell\|$ to obtain

$$\frac{z^{\ell+1}}{\|z^\ell\|} = \alpha \hat{z}^\ell + \lambda \sqrt{\frac{2}{n}} W^{\ell+1} \varphi_+(\hat{z}^\ell) \stackrel{d}{=} \alpha \hat{z}^\ell + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(\hat{z}^\ell)\| g^{\ell+1}, \quad (16)$$

where g^ℓ are iid Gaussian vectors with iid $\mathcal{N}(0, 1)$ entries by application of (14). Hence

$$\frac{\|z^{\ell+1}\|}{\|z^\ell\|} \stackrel{d}{=} \left\| \alpha \hat{z}^\ell + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(\hat{z}^\ell)\| g^{\ell+1} \right\| \stackrel{d}{=} \left\| \alpha \vec{e}_1 + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(\hat{z}^\ell)\| g^{\ell+1} \right\|.$$

The last equality follows by applying an orthogonal transformation O such that $O\hat{z}^\ell = \vec{e}_1 = (1, 0, \dots, 0)^T$ inside the norm, and observing that Gaussian random vectors are invariant under orthogonal transformations $Og^\ell \stackrel{d}{=} g^\ell$. Hence we have the telescoping product for $\|z^d\|$:

$$\|z^d\| = \|z^0\| \prod_{\ell=0}^{d-1} \frac{\|z^{\ell+1}\|^2}{\|z^\ell\|^2} \stackrel{d}{=} \|z^0\| \prod_{\ell=0}^{d-1} \left\| \alpha \vec{e}_1 + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(\hat{z}^\ell)\| g^{\ell+1} \right\|. \quad (17)$$

This shows that $\|z^d\|$ is a product of d random variables which are dependent on each other only through the terms $\|\varphi_+(\hat{z}^\ell)\|$. (Note that $\|z^0\|$ is independent of \hat{z}^0 since z^0 is Gaussian.) Since $\|\varphi_+(\hat{z}^\ell)\|^2 \approx 1/2$ with typical fluctuations on the scale $1/\sqrt{n}$, therefore the dependence between terms of (17) is small.

Taking the \ln of (17) exhibits $\ln(\|z^d\|^2/n)$ as a sum of these weakly correlated random variables. Here we note that various tail estimates for the same or related quantities have been developed [3, 41], however these estimates are not precise enough to pinpoint the exact limiting distribution. In contrast, we are able to derive the exact limiting distribution via a Central Limit Theorem (CLT) for weakly correlated sums [42]. The proof of Theorem 1 is completed by computing the mean, variance and covariance of terms using Conjecture 5. For Theorem 4, the final calculation is simplified by (10) which shows the terms are uncorrelated. The detailed proof is given in Appendix B.

Acknowledgement

We would like to thank Blair Bilodeau, Gintare Karolina Dziugaite, Mahdi Haghifam, Yani A. Ioannou, James Lucas, Jeffrey Negrea, Mengye Ren, and Ekansh Sharma for helpful discussions and draft feedback. We would also like to thank the anonymous NeurIPS reviewers for insightful feedback. In particular, one identified numerous relations to existing work, and another helped us identify the uniformity requirement in Conjecture 5. ML is supported by Ontario Graduate Scholarship and the Vector Institute. MN is supported by an NSERC Discovery Grant. DMR is supported in part by an NSERC Discovery Grant, Ontario Early Researcher Award, and a stipend provided by the Charles Simonyi Endowment.

References

- [1] A. Jacot, F. Gabriel, and C. Hongler. “Neural tangent kernel: Convergence and generalization in neural networks.” *Advances in Information Processing Systems (NeurIPS)*. 2018. arXiv: 1806.07572.
- [2] S. Du, J. Lee, H. Li, L. Wang, and X. Zhai. “Gradient descent finds global minima of deep neural networks.” *Int. Conf. Machine Learning (ICML)*. PMLR. 2019, pp. 1675–1685.
- [3] Z. Allen-Zhu, Y. Li, and Z. Song. “A convergence theory for deep learning via over-parameterization.” *Int. Conf. Machine Learning (ICML)*. PMLR. 2019, pp. 242–252.
- [4] D. Zou, Y. Cao, D. Zhou, and Q. Gu. “Gradient descent optimizes over-parameterized deep ReLU networks.” In: *Machine Learning* 109.3 (2020), pp. 467–492.
- [5] L. Chizat, E. Oyallon, and F. Bach. “On Lazy Training in Differentiable Programming.” In: *Advances in Neural Information Processing Systems* 32 (2019), pp. 2937–2947.

- [6] J. Lee, L. Xiao, S. S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington. *Wide neural networks of any depth evolve as linear models under gradient descent*. 2019. arXiv: 1902.06720.
- [7] G. Yang. *Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation*. 2019. arXiv: 1902.04760.
- [8] G. Yang. *Tensor programs ii: Neural tangent kernel for any architecture*. 2020. arXiv: 2006.14548.
- [9] S. Arora, S. S. Du, W. Hu, Z. Li, R. Salakhutdinov, and R. Wang. “On exact computation with an infinitely wide neural net.” *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. 2019, pp. 8141–8150.
- [10] Z. Chen, Y. Cao, D. Zou, and Q. Gu. “How Much Over-parameterization Is Sufficient to Learn Deep Re{LU} Networks?” *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=fgd7we_uZa6.
- [11] S. Arora, S. S. Du, Z. Li, R. Salakhutdinov, R. Wang, and D. Yu. *Harnessing the power of infinitely wide deep nets on small-data tasks*. 2019. arXiv: 1910.01663.
- [12] G. M. Rotskoff and E. Vanden-Eijnden. *Trainability and Accuracy of Neural Networks: An Interacting Particle System Approach*. 2018. arXiv: 1805.00915.
- [13] L. Chizat and F. Bach. *On the Global Convergence of Gradient Descent for Over-parameterized Models using Optimal Transport*. 2018. arXiv: 1805.09545.
- [14] J. Sirignano and K. Spiliopoulos. *Mean Field Analysis of Neural Networks: A Law of Large Numbers*. 2018. arXiv: 1805.01053.
- [15] S. Mei, A. Montanari, and P.-M. Nguyen. “A mean field view of the landscape of two-layer neural networks.” In: *Proceedings of the National Academy of Sciences* 115.33 (2018), E7665–E7671. ISSN: 0027-8424. DOI: 10.1073/pnas.1806579115. eprint: <https://www.pnas.org/content/115/33/E7665.full.pdf>.
- [16] S. Mei, T. Misiakiewicz, and A. Montanari. *Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit*. arXiv: 1902.06015.
- [17] G. Yang and E. J. Hu. “Feature Learning in Infinite-Width Neural Networks.” *Int. Conf. Machine Learning (ICML)*. 2021. arXiv: 2011.14522.
- [18] R. M. Neal. *Bayesian learning for neural networks*. Vol. 118. Springer Science & Business Media, 1995.
- [19] G. Yang and S. S. Schoenholz. “Mean field residual networks: on the edge of chaos.” *Advances in Neural Information Processing Systems*. 2017, pp. 2865–2873.
- [20] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein. “Deep Neural Networks as Gaussian Processes.” *Int. Conf. Learning Representations (ICLR)*. 2018.
- [21] A. De Matthews, J. Hron, M. Rowland, R. Turner, and Z. Ghahramani. “Gaussian process behaviour in wide deep neural networks.” *Int. Conf. Learning Representations (ICLR)*. 2018.
- [22] R. Novak, L. Xiao, Y. Bahri, J. Lee, G. Yang, J. Hron, D. A. Abolafia, J. Pennington, and J. Sohl-dickstein. “Bayesian Deep Convolutional Networks with Many Channels are Gaussian Processes.” *Int. Conf. Learning Representations (ICLR)*. 2018.
- [23] G. Yang. *Tensor programs i: Wide feedforward or recurrent neural networks of any architecture are gaussian processes*. 2019. arXiv: 1910.12478.
- [24] B. Hanin and M. Nica. “Products of many large random matrices and gradients in deep neural networks.” In: *Communications in Mathematical Physics* (2019), pp. 1–36.
- [25] B. Hanin and M. Nica. “Finite Depth and Width Corrections to the Neural Tangent Kernel.” *Int. Conf. Learning Representations (ICLR)*. 2019.
- [26] M. Seleznova and G. Kutyniok. *Analyzing Finite Neural Networks: Can We Trust Neural Tangent Kernel Theory?* 2020. arXiv: 2012.04477.
- [27] B. Chmiel, L. Ben-Uri, M. Shkolnik, E. Hoffer, R. Banner, and D. Soudry. “Neural gradients are near-lognormal: improved quantized and sparse training.” *Int. Conf. Learning Representations (ICLR)*. arXiv: 2006.08173. URL: <https://openreview.net/forum?id=EoFNy62JGd>.
- [28] K. He, X. Zhang, S. Ren, and J. Sun. “Deep residual learning for image recognition.” *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.

- [29] S. Hayou, E. Clerico, B. He, G. Deligiannidis, A. Doucet, and J. Rousseau. “Stable ResNet.” *Int. Conf. Artificial Intelligence and Statistics (AISTATS)*. PMLR. 2021, pp. 1324–1332.
- [30] H. Zhang, Y. N. Dauphin, and T. Ma. “Fixup Initialization: Residual Learning Without Normalization.” *International Conference on Learning Representations*. 2018.
- [31] Y. Blumenfeld, D. Gilboa, and D. Soudry. “Beyond signal propagation: is feature diversity necessary in deep neural network initialization?” *International Conference on Machine Learning*. PMLR. 2020, pp. 960–969.
- [32] D. Arpit, V. Campos Camúñez, and Y. Bengio. “How to initialize your network? Robust initialization for WeightNorm & ResNets.” *Advances in Neural Information Processing Systems 32 (NIPS 2019)*. 2019, pp. 1–10.
- [33] B. Hanin and D. Rolnick. “How to Start Training: The Effect of Initialization and Architecture.” In: *Advances in Neural Information Processing Systems 31* (2018).
- [34] B. Hanin. “Which neural net architectures give rise to exploding and vanishing gradients?” Vol. 31. 2018.
- [35] K. He, X. Zhang, S. Ren, and J. Sun. “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification.” *Proc. IEEE Int. Conf. Computer Vision*. 2015, pp. 1026–1034.
- [36] Y. Cho and L. K. Saul. “Kernel methods for deep learning.” *Advances in Neural Information Processing Systems (NeurIPS)*. 2009, pp. 342–350.
- [37] P. Dhariwal and A. Nichol. *Diffusion Models Beat GANs on Image Synthesis*. 2021. arXiv: 2105.05233.
- [38] D. Balduzzi, M. Frean, L. Leary, J. P. Lewis, K. W.-D. Ma, and B. McWilliams. “The Shattered Gradients Problem: If resnets are the answer, then what is the question?” *Int. Conf. Machine Learning (ICML)*. PMLR. 2017, pp. 342–350.
- [39] Z. Allen-Zhu, Y. Li, and Y. Liang. “Learning and generalization in overparameterized neural networks, going beyond two layers.” In: *Advances in neural information processing systems* (2019).
- [40] Y. Bai and J. Lee. “Beyond Linearization: On Quadratic and Higher-Order Approximation of Wide Neural Networks.” *International Conference on Learning Representations*. 2019.
- [41] S. Buchanan, D. Gilboa, and J. Wright. “Deep Networks and the Multiple Manifold Problem.” *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=0-6Pm_d_Q-.
- [42] M. H. Neumann. “A central limit theorem for triangular arrays of weakly dependent random variables, with applications in statistics.” In: *ESAIM: PS* 17 (2013), pp. 120–134. DOI: 10.1051/ps/2011144.
- [43] A. J. Stam. “Limit Theorems for Uniform Distributions on Spheres in High-Dimensional Euclidean Spaces.” In: *Journal of Applied Probability* 19.1 (1982), pp. 221–228. ISSN: 00219002. URL: <http://www.jstor.org/stable/3213932>.
- [44] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. *JAX: composable transformations of Python+NumPy programs*. Version 0.2.5. 2018. URL: <http://github.com/google/jax>.
- [45] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. “PyTorch: An Imperative Style, High-Performance Deep Learning Library.” In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc., 2019, pp. 8024–8035. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [46] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. “Array programming with NumPy.” In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2. URL: <https://doi.org/10.1038/s41586-020-2649-2>.

- [47] Plotnine Contributors. *has2k1/plotnine: v0.8.0*. Version v0.8.0. Mar. 2021. DOI: 10.5281/zenodo.4636791. URL: <https://doi.org/10.5281/zenodo.4636791>.
- [48] H. Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016. ISBN: 978-3-319-24277-4. URL: <https://ggplot2.tidyverse.org>.
- [49] Pandas Contributors. *pandas-dev/pandas: Pandas 1.0.3*. Version v1.0.3. Mar. 2020. DOI: 10.5281/zenodo.3715232. URL: <https://doi.org/10.5281/zenodo.3715232>.
- [50] *Anaconda Software Distribution*. Version Vers. 2-2.4.0. 2020. URL: <https://docs.anaconda.com/>.
- [51] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing. “Jupyter Notebooks – a publishing format for reproducible computational workflows.” *Positioning and Power in Academic Publishing: Players, Agents and Agendas*. Ed. by F. Loizides and B. Schmidt. IOS Press. 2016, pp. 87–90.
- [52] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization.” *ICLR (Poster)*. 2015. URL: <http://arxiv.org/abs/1412.6980>.
- [53] K. Liu, W. Yang, P. Yang, and F. Ducau. *pytorch-cifar*. <https://github.com/kuangliu/pytorch-cifar>. 2017.
- [54] K. He, X. Zhang, S. Ren, and J. Sun. “Identity mappings in deep residual networks.” *European conference on computer vision*. Springer. 2016, pp. 630–645.
- [55] A. Krizhevsky. “Learning Multiple Layers of Features from Tiny Images.” In: (2009).

A Appendix

A.1 Layer dependent coefficients

We stated our main result with fixed α, λ , but the result easily extends to the case where α, λ vary from layer to layer. This allows comparison between our result and the infinite width limits [29, 33], where they have $\alpha = 1$ and allow λ_i to vary layer by layer. The statement in that setting is modified as follows.

Proposition 6. *Suppose α_i, λ_i are sequences such that λ_i is uniformly bounded away from 0. Define the network by*

$$z^0 = \sqrt{\frac{1}{n_{in}}} W^0 x, \quad z^\ell = \alpha_\ell z^{\ell-1} + \lambda_\ell \sqrt{\frac{2}{n}} W^\ell \varphi_+(z^{\ell-1}) \text{ for } 1 \leq \ell \leq d, \quad z^{out} = \sqrt{\frac{1}{n}} W^{out} z^d \quad (18)$$

Consider the limit where both the network depth $d \rightarrow \infty$ and hidden layer width $n \rightarrow \infty$ in such a way that the ratio $\frac{d}{n}$ converges to a constant. In this limit, the distribution of the network output z^{out} for a given input x is given by

$$z^{out} \stackrel{d}{=} \frac{\|x\|}{\sqrt{n_{in}}} \prod_{i=1}^{\ell} \sqrt{\alpha_i^2 + \lambda_i^2} \exp\left(\frac{1}{2}G\right) \vec{Z} \quad (19)$$

where $\vec{Z} = (Z_1, \dots, Z_{n_{out}}) \in \mathbb{R}^{n_{out}}$ is a Gaussian random vector with iid $\mathcal{N}(0, 1)$ entries and, assuming that Conjecture 5 holds, then the random variable $G \in \mathbb{R}$ converges to a Gaussian random variable in this limit and satisfies

$$\mathbf{E}[G] = -\frac{\beta}{2} + \sum_{\ell=1}^d \frac{2\lambda_\ell^2}{\alpha_\ell^2 + \lambda_\ell^2} h_\ell + O\left(\frac{d}{n^2}\right) \quad (20)$$

$$\mathbf{Var}[G] = \beta + \sum_{\substack{\ell, \ell'=1 \\ \ell \neq \ell'}}^d \frac{\lambda_\ell^4}{(\alpha_\ell^2 + \lambda_\ell^2)^2} \frac{\bar{J}_2(\theta_{\ell, \ell'}) - \bar{J}_2(\pi - \theta_{\ell, \ell'})}{n} + O\left(\frac{d}{n^2}\right), \quad (21)$$

where $\theta_{\ell, \ell'}$ is such that $\cos(\theta_{\ell, \ell'}) = \prod_{i=\ell}^{\ell'-1} \frac{\alpha_i^2}{\alpha_i^2 + \lambda_i^2}$ and $\beta = \frac{2}{n} + \frac{1}{n} \sum_{\ell=1}^d \frac{5\lambda_\ell^4 + 4\alpha_\ell^2 \lambda_\ell^2}{(\alpha_\ell^2 + \lambda_\ell^2)^2}$.

The corresponding limit theorem for balanced ResNets also holds; the hypoactivation and layer-wise covariance term in (20) vanish.

The behaviour is a complicated function of the sequences α_i, λ_i . It would be interesting to use these theoretical results to guide the choice of parameters α_i, λ_i and investigate how this effects training behaviour.

B Proof of main results

To simplify the exposition of the proofs, we will assume without loss of generality that $\alpha^2 + \lambda^2 = 1$. The general case can be reduced to the case $\alpha^2 + \lambda^2 = 1$ by dividing by $\sqrt{\alpha^2 + \lambda^2}$ in each layer and rescaling the parameters α, λ to $\frac{\alpha}{\sqrt{\alpha^2 + \lambda^2}}$ and $\frac{\lambda}{\sqrt{\alpha^2 + \lambda^2}}$.

By the argument of Section 5, the proof reduces to showing that the random variable $G = \ln(\|z^0\|^2/n) + \sum_{\ell=1}^d \ln(X_\ell)$ has the desired asymptotic behaviour where X_ℓ is defined to be

$$X_\ell = \left\| \alpha \vec{e}_1 + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(\hat{z}^\ell)\| g^\ell \right\|^2 = \alpha^2 + \lambda^2 \frac{2}{n} \|\varphi_+(\hat{z}^\ell)\|^2 \|g^\ell\|^2 + 2\alpha\lambda \sqrt{\frac{2}{n}} \|\varphi_+(\hat{z}^\ell)\| g_1^\ell. \quad (22)$$

B.1 Mean Calculation

Lemma 7. $\mathbf{E}[X_\ell] = 1 + 2\lambda^2 h_\ell$ where h_ℓ is the hypoactivation of layer ℓ .

Proof. Taking expectation of both sides of (22), we have

$$\begin{aligned}\mathbf{E}[X_\ell] &= \alpha^2 + \lambda^2 \frac{2}{n} \mathbf{E}[\|\varphi_+(\hat{z}^\ell)\|^2] \mathbf{E}[\|\bar{g}^\ell\|^2] + 2\alpha\lambda \sqrt{\frac{2}{n}} \mathbf{E}[\|\varphi_+(\hat{z}^\ell)\|] \mathbf{E}[\bar{g}_1^\ell] \\ &= \alpha^2 + \lambda^2 \frac{2}{n} \left(\frac{1}{2} + h_\ell\right) n + 0 = 1 + 2\lambda^2 h_\ell,\end{aligned}$$

where we have used $\alpha^2 + \lambda^2 = 1$ in the last line. \square

Corollary 8. Assume Conjecture 5. Then $h_\ell = O(\frac{1}{n})$ and $\mathbf{E}[X_\ell] = 1 + O(\frac{1}{n})$.

Proof. By the conjecture, $h_\ell = \mathbf{E}[\|\varphi_+(\hat{z}^\ell)\|^2] - \frac{1}{2} = \mathbf{E}[\|\varphi_+(\hat{z}^\ell)\|^2] - \mathbf{E}[\|\varphi_+(u)\|^2] = O(\frac{1}{n})$. \square

B.2 Variance Calculation

Lemma 9. Let $z \in \mathbb{R}^n$ be any random vector and let $g \in \mathbb{R}^n$ be an independent Gaussian random vector with iid $\mathcal{N}(0, 1)$ entries. Define

$$X = \left\| \alpha e_1 + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(z)\| g \right\|^2 = \alpha^2 + \lambda^2 2 \|\varphi_+(z)\|^2 \frac{\|g\|^2}{n} + \frac{2\alpha\lambda}{\sqrt{n}} \left\| \sqrt{2} \varphi_+(z) \right\| g_1.$$

Then

$$\mathbf{Var}(X) = \lambda^4 \frac{n+2}{n} \mathbf{Var} \left[2 \|\varphi_+(z)\|^2 \right] + \frac{2\lambda^4}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right]^2 + \frac{4\alpha^2 \lambda^2}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right].$$

Proof. We will use the decomposition $\mathbf{Var}(X) = \mathbf{E}[(X - \alpha^2)^2] + (\mathbf{E}[X] - \alpha^2)^2$ and compute the two terms individually.

In the second term, since $\mathbf{E}[g_1] = 0$ and $\mathbf{E}[\|g\|^2] = n$, we have $\mathbf{E}[X] - \alpha^2 = \lambda^2 \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right]$.

To compute $\mathbf{E}[(X - \alpha^2)^2]$, notice that $X - \alpha^2 = \lambda^2 2 \|\varphi_+(z)\|^2 \cdot \frac{1}{n} \|g\|^2 + 2\alpha\lambda \frac{1}{\sqrt{n}} \left\| \sqrt{2} \varphi_+(z) \right\| g_1$ is a sum of two terms. The two terms are uncorrelated since $\mathbf{Cov}(\|g\|^2, g_1) = 0$. Hence, the expectation of the cross term in $(X - \alpha^2)^2$ is zero and we can compute

$$\begin{aligned}\mathbf{E}[(X - \alpha^2)^2] &= \lambda^4 \mathbf{E} \left[4 \|\varphi_+(z)\|^4 \right] \mathbf{E} \left[\frac{\|g\|^4}{n^2} \right] + \frac{4\alpha^2 \lambda^2}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right] \mathbf{E} [g_1^2] \\ &= \lambda^4 \mathbf{E} \left[4 \|\varphi_+(z)\|^4 \right] \frac{n+2}{n} + \frac{4\alpha^2 \lambda^2}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right].\end{aligned}$$

We have used the fact about χ_n^2 random variables that $\mathbf{E}[\|g\|^4] = n(n+2)$. Finally then:

$$\begin{aligned}\mathbf{Var}(X) &= \mathbf{E} \left[(X - \alpha^2)^2 \right] - \lambda^4 \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right]^2 \\ &= \lambda^4 \mathbf{E} \left[4 \|\varphi_+(z)\|^4 \right] \frac{n+2}{n} - \lambda^4 \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right]^2 + \frac{4\alpha^2 \lambda^2}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right] \\ &= \lambda^4 \frac{n+2}{n} \mathbf{Var} \left[2 \|\varphi_+(z)\|^2 \right] + \frac{2\lambda^4}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right]^2 + \frac{4\alpha^2 \lambda^2}{n} \mathbf{E} \left[2 \|\varphi_+(z)\|^2 \right]\end{aligned}$$

\square

Lemma 10. If $u \in \mathbb{S}^{n-1}$ is a random vector which is distributed uniformly from the unit sphere, then

$$\mathbf{E} \left[2 \|\varphi_+(u)\|^2 \right] = 1, \quad \mathbf{Var} \left[2 \|\varphi_+(u)\|^2 \right] = \frac{3}{n+2}.$$

Proof. This can be calculated directly using properties of the unit sphere, but the proof is complicated by the fact that the entries u_i, u_j are *not* independent. Instead, there is an elementary proof using the following equality in distribution:

$$\varphi_+(u) \|g\| \stackrel{d}{=} \varphi_+(g),$$

where $g \in \mathbb{R}^n$ has iid $\mathcal{N}(0, 1)$ entries and is independent of u . This follows because of the fact that $u \stackrel{d}{=} \frac{g}{\|g\|}$ is independent of $\|g\|$, so $u \|g\| \stackrel{d}{=} g$. By also using the fact that $\varphi_+(\cdot)$ is a positive homogeneous function $\varphi_+(|c|x) = |c| \varphi_+(x)$, we see by applying φ_+ to this equality in distribution that $\varphi_+(u) \|g\| \stackrel{d}{=} \varphi_+(g)$ as desired.

Taking norm and expectation of this equality in distribution gives $\mathbf{E} \left[2 \|\varphi_+(u)\|^2 \|g\|^2 \right] = \mathbf{E} \left[2 \|\varphi_+(g)\|^2 \right]$. Since $\|g\|$ and u are independent, we can factor and rearrange to obtain

$$\mathbf{E} \left[2 \|\varphi_+(u)\|^2 \right] = \frac{\mathbf{E} \left[2 \|\varphi_+(g)\|^2 \right]}{\mathbf{E} \left[\|g\|^2 \right]}.$$

Since the entries of g are independent of each other, it is easier to compute using g and this identity instead of using u . Moreover, because Gaussian distribution are symmetrically distributed, we have that $\{g_1^2, \dots, g_n^2\}$ is independent of $1\{g_1 > 0\}, \dots, 1\{g_n > 0\}$. Hence:

$$\mathbf{E} \left[2 \|\varphi_+(g)\|^2 \right] = \mathbf{E} \left[2 \sum_{i=1}^n g_i^2 1\{g_i > 0\} \right] = 2 \sum_{i=1}^n \mathbf{E} [g_i^2] \mathbf{P} [g_i > 0] = 2n \frac{1}{2} = n$$

and so $\mathbf{E} \left[2 \|\varphi_+(u)\|^2 \right] = \frac{\mathbf{E} [2 \|\varphi_+(g)\|^2]}{\mathbf{E} [\|g\|^2]} = \frac{n}{n} = 1$. Similarly, using $\mathbf{E} \left[4 \|\varphi_+(u)\|^4 \right] = \frac{\mathbf{E} [4 \|\varphi_+(g)\|^4]}{\mathbf{E} [\|g\|^4]}$ we now compute by looking at diagonal and off-diagonal terms as follows

$$\begin{aligned} \mathbf{E} \left[4 \|\varphi_+(g)\|^4 \right] &= 4 \mathbf{E} \left[\sum_{i,j=1}^n g_i^2 g_j^2 1\{g_i > 0\} 1\{g_j > 0\} \right] \\ &= 4 \sum_{i=1}^n \mathbf{E} [g_i^4] \mathbf{P} [g_i > 0] + 4 \sum_{\substack{i,j=1 \\ i \neq j}}^n \mathbf{E} [g_i^2] \mathbf{E} [g_j^2] \mathbf{P} [g_i > 0] \mathbf{P} [g_j > 0] \\ &= 4n \cdot 3 \cdot \frac{1}{2} + 4n(n-1) \frac{1}{2} \cdot \frac{1}{2} = n^2 + 5n. \end{aligned}$$

Using $\mathbf{E} [\|g\|^4] = n(n+2)$, we finally obtain $\mathbf{E} \left[4 \|\varphi_+(u)\|^4 \right] = \frac{n(n+5)}{n(n+2)}$ from which the claimed variance formula follows. \square

Corollary 11. *If u is a uniform from the unit sphere \mathbb{S}^{n-1} , and $X = \left\| \alpha e_1 + \lambda \sqrt{\frac{2}{n}} \|\varphi_+(u)\| g \right\|^2$ as in Lemma 9, then*

$$\mathbf{Var}(X) = \frac{5\lambda^4 + 4\alpha^2\lambda^2}{n}.$$

Proof. Plug in the result of Lemma 10 into Lemma 9. \square

Lemma 12. *Assuming Conjecture 5 is true, and with X_ℓ defined as (22), we have*

$$\mathbf{Var}(X_\ell) = \frac{5\lambda^4 + 4\alpha^2\lambda^2}{n} \left(1 + O\left(\frac{1}{n}\right) \right)$$

Proof. By the conjecture, we have that $\mathbf{E} \left[2 \|\varphi_+(\hat{z}^\ell)\|^2 \right] = \mathbf{E} \left[2 \|\varphi_+(u)\|^2 \right] (1 + O(\frac{1}{n}))$ and $\mathbf{Var} \left[2 \|\varphi_+(\hat{z}^\ell)\|^2 \right] = \mathbf{Var} \left[2 \|\varphi_+(u)\|^2 \right] (1 + O(\frac{1}{n}))$. Hence we can compute $\mathbf{Var}(X_\ell)$ up to a factor of $(1 + O(\frac{1}{n}))$ from Lemma 9 by plugging in the result of Lemma 10. \square

B.3 Uniform distribution on spheres and Gaussian random variables

In this section we develop some approximations which are used in the next section. Let $u \in \mathbb{S}^{n-1} \subset \mathbb{R}^n$ be a uniform random variable from the unit sphere. Let $Z \sim \mathcal{N}(0, 1) \in \mathbb{R}$ be a standard Gaussian. The results in this section concern the error rate in the well known approximation for the marginal distribution of the components u_i , namely $\sqrt{n}u_i \approx Z$.

Lemma 13. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be any bounded function. Then the marginal distribution of any coordinate u_i satisfies*

$$\mathbf{E} [f(\sqrt{n}u_i)] = \mathbf{E} [f(Z)] + O\left(\frac{1}{n}\right)$$

Proof. This is a direct corollary to Theorem 2 of [43], which uses Stirling's formula to show that the total variation distance between the random variables $\sqrt{n}u_i$ and Z is at most $2\left(\sqrt{1 + \frac{3}{n-3}} - 1\right) = O(1/n)$. \square

Lemma 14. *For $p \in \mathbb{N}$, the $2p$ -th moment of the marginal distribution of any coordinate u_i satisfies*

$$\mathbf{E} [(\sqrt{n}u_i)^{2p}] = \mathbf{E} [Z^{2p}] \cdot \left(1 + \frac{2}{n}\right) \cdots \left(1 + \frac{2p-2}{n}\right)^{-1} = \mathbf{E} [Z^{2p}] + O\left(\frac{1}{n}\right)$$

Proof. As in the proof of Lemma 9, we use the equality in distribution $\|g\| u \stackrel{d}{=} g$ where $g \in \mathbb{R}^N$ is a vector whose components are iid $\mathcal{N}(0, 1)$ independent of u . From this it follows that

$$\mathbf{E} [\|g\|^{2p} u_i^{2p}] = \mathbf{E} [g_i^{2p}] = \mathbf{E} [Z^{2p}]$$

The result then follows by using the independence of $\|g\|$ and u , and the formula for the p -th moment of χ_n^2 random variable, namely $\|g\|^{2p} = n(n+2) \cdots (n+2p-2)$. \square

Corollary 15. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be any function that satisfies $|f(x)| \leq Ax^{2p} + B$ for some constants $A, B > 0$ and exponent $p \in \mathbb{N}$. Then*

$$\mathbf{E} [f(\sqrt{n}u_i)] = \mathbf{E} [f(Z)] + O\left(\frac{1}{n}\right)$$

Proof. The proof is immediate writing the difference in expectation as an integral and then comparing $\int f(x) (\rho_{\sqrt{n}u_i}(x) - \rho_Z(x)) dx$ to $\int (Ax^{2p} + B) (\rho_{\sqrt{n}u_i}(x) - \rho_Z(x)) dx$ by the results of the previous two lemmas. \square

B.4 Pairwise covariances

Define the function $\bar{J}_2 : \mathbb{R} \rightarrow \mathbb{R}$ by

$$\bar{J}_2(\theta) := 2\mathbf{E} [\varphi_+^2(Z) \varphi_+^2(\cos(\theta)Z + \sin(\theta)W)], \quad (23)$$

where $Z \in \mathbb{R}, W \in \mathbb{R}$ are iid $\mathcal{N}(0, 1)$ random variables. In Cho and Saul [36] they find an explicit formula for this, namely:

$$\bar{J}_2(\theta) = \frac{J_2(\theta)}{\pi} = \frac{3\sin(\theta)\cos(\theta) + (\pi - \theta)(1 + 2\cos^2\theta)}{\pi}$$

Lemma 16. *Let $u \in \mathbb{S}^{n-1}$ be a uniform random vectors from the unit sphere and let $g \in \mathbb{R}^n$ be a Gaussian vector with iid $\mathcal{N}(0, 1)$ entries which is independent of u . Then*

$$\text{Cov} \left(2\|\varphi_+(u)\|^2, 2\left\|\varphi_+ \left(\cos(\theta)u + \frac{\sin(\theta)}{\sqrt{n}}g \right)\right\|^2 \right) = \frac{\bar{J}_2(\theta) - \bar{J}_2(\pi - \theta)}{n} \left(1 + O\left(\frac{1}{n}\right) \right). \quad (24)$$

Proof. By expanding the norms into sums, $\|x\|^2 = \sum_{i=1}^n x_i^2$, we can compute the covariance by summing over all *pairs* of coordinates $u_i, \cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j$. There are two types of terms to consider. (Note: we use the notation $\varphi_+^2(x) = (\varphi_+(x))^2$.)

Diagonal Terms: $\mathbf{Cov} \left(\varphi_+^2(u_i), \varphi_+^2(\cos(\theta)u_i + \sin(\theta)n^{-\frac{1}{2}}g_i) \right)$ for $1 \leq i \leq n$ We first use the positive homogeneity of φ_+ to extract a factor of \sqrt{n} from both terms

$$\mathbf{Cov} \left(\varphi_+^2(u_i), \varphi_+^2(\cos(\theta)u_i + \frac{\sin(\theta)}{\sqrt{n}}g_i) \right) = \frac{1}{n^2} \mathbf{Cov} \left(\varphi_+^2(\sqrt{n}u_i), \varphi_+^2(\cos(\theta)\sqrt{n}u_i + \sin(\theta)g_i) \right)$$

We now use the approximation of $\sqrt{n}u_i$ by $Z \sim \mathcal{N}(0, 1)$ as in Section B.3 to obtain

$$\begin{aligned} & \frac{1}{n^2} \mathbf{Cov} \left(\varphi_+^2(\sqrt{n}u_i), \varphi_+^2(\cos(\theta)\sqrt{n}u_i + \sin(\theta)g_i) \right) \\ &= \frac{1}{n^2} \mathbf{Cov} \left(\varphi_+^2(Z), \varphi_+^2(\cos(\theta)Z + \sin(\theta)g_i) \right) \left(1 + O\left(\frac{1}{n}\right) \right) \\ &= \frac{\frac{1}{2}\bar{J}_2(\theta) - \frac{1}{4}}{n^2} \left(1 + O\left(\frac{1}{n}\right) \right), \end{aligned}$$

where we have used the definition of $\bar{J}_2(\theta)$ from (23).

Off diagonal terms $\mathbf{Cov} \left(\varphi_+^2(u_i), \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right)$ for $1 \leq i \leq n, 1 \leq j \leq n, i \neq j$

We compute the expectation

$$\mathbf{E} \left[\varphi_+^2(u_i)^2 \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right]$$

by first conditioning on u_j . Conditioned on u_j , the distribution of u_i is $u_i \stackrel{d}{=} \sqrt{1 - u_j^2} \tilde{u}_i$, where \tilde{u} is independent of u and is drawn uniformly from the unit sphere whose dimension is one smaller than that of u , namely $\tilde{u} \in \mathbb{S}^{n-2}$. Since the φ_+ is positive homogeneous, we can factor $\sqrt{1 - u_j^2}$ out to get:

$$\begin{aligned} \mathbf{E} \left[\varphi_+^2(u_i) \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right] &= \mathbf{E} \left[\varphi_+^2(\tilde{u}_i)^2 (1 - u_j^2) \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right] \\ &= \mathbf{E} [\varphi_+^2(\tilde{u}_i)^2] \mathbf{E} \left[(1 - u_j^2) \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right] \\ &= \frac{1}{2(n-1)} \mathbf{E} \left[(1 - u_j^2) \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right] \end{aligned}$$

As in the calculation for the diagonal term, we now again use the approximation that $\sqrt{n}u_j$ is approximately marginally distributed like $Z \sim \mathcal{N}(0, 1)$ and the positive homogeneity of φ_+ to obtain

$$\begin{aligned} & \mathbf{E} \left[\varphi_+^2(u_i) \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right] \\ &= \frac{1}{2(n-1)} \left(\frac{1}{2n} - \mathbf{E} [Z^2 \varphi_+^2(\alpha Z + \lambda g_j)] \right) \left(1 + O\left(\frac{1}{n}\right) \right). \end{aligned}$$

Renaming g_j to W to match the notation of (23), we finally obtain

$$\begin{aligned} & \mathbf{Cov} \left(\varphi_+^2(u_i)^2, \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right) \\ &= \mathbf{E} \left[\varphi_+^2(u_i) \varphi_+^2(\cos(\theta)u_j + \sin(\theta)n^{-\frac{1}{2}}g_j) \right] - \frac{1}{4n^2} \left(1 + O\left(\frac{1}{n}\right) \right) \\ &= \left(\frac{1}{4n^2(n-1)} - \frac{1}{2n^2(n-1)} \mathbf{E} [Z^2 \varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] \right) \left(1 + O\left(\frac{1}{n}\right) \right) \end{aligned}$$

Summing the diagonal and off-diagonal terms, we find the total covariance is:

$$\begin{aligned}
& \mathbf{Cov} \left(2 \|\varphi_+(u)\|^2, 2 \left\| \varphi_+ \left(\cos(\theta)u + \sin(\theta)n^{-\frac{1}{2}}g \right) \right\|^2 \right) \\
&= 4n \times (\text{Diagonal term contribution}) + 4n(n-1) \times (\text{Off Diagonal term contribution}) \\
&= \left(\frac{2\bar{J}_2(\theta) - 1}{n} + \frac{1 - 2\mathbf{E}[Z^2\varphi_+^2(\alpha Z + \lambda W)]}{n} \right) \left(1 + O\left(\frac{1}{n}\right) \right) \\
&= \left(\frac{2\bar{J}_2(\theta) - 2\mathbf{E}[Z^2\varphi_+^2(\alpha Z + \lambda W)]}{n} \right) \left(1 + O\left(\frac{1}{n}\right) \right)
\end{aligned}$$

Finally we notice the identity

$$\begin{aligned}
& 2\bar{J}_2(\theta) - 2\mathbf{E}[Z^2\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] \\
&= 4\mathbf{E}[\varphi_+^2(Z)\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] - 2\mathbf{E}[Z^2\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] \\
&= \mathbf{E}[(2\varphi_+^2(Z) - Z^2)2\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] \\
&= \mathbf{E}[(\varphi_+^2(Z) - \varphi_+^2(-Z))2\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] \\
&= 2\mathbf{E}[\varphi_+^2(Z)\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] - 2\mathbf{E}[\varphi_+^2(-Z)\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] \\
&= 2\mathbf{E}[\varphi_+^2(Z)\varphi_+^2(\cos(\theta)Z + \sin(\theta)W)] - 2\mathbf{E}[\varphi_+^2(Z)\varphi_+^2(-\cos(\theta)Z + \sin(\theta)W)] \\
&= \bar{J}_2(\theta) - \bar{J}_2(\pi - \theta),
\end{aligned}$$

which gives the claimed formula for the covariance. \square

Corollary 17. Assume Conjecture 5 is true. Then

$$\mathbf{Cov} \left(2 \|\varphi_+^2(\hat{z}^\ell)\|^2, 2 \|\varphi_+^2(\hat{z}^{\ell'})\|^2 \right) = \frac{\bar{J}_2(\theta_{|\ell-\ell'|}) - \bar{J}_2(\pi - \theta_{|\ell-\ell'|})}{n} \left(1 + O\left(\frac{1}{n}\right) \right), \quad (25)$$

where θ_k is the angle such that $\theta_k = \cos^{-1}(\alpha^k)$, and the constant in the big $O(\cdot)$ notation is uniform in ℓ, ℓ' .

Proof. This follows immediately from the approximation for the covariance in Conjecture 5 and Lemma 16. \square

Remark 18. By the series expansion for $\arccos(x) = \frac{\pi}{2} - x + O(x^3)$ as $x \rightarrow 0$, it follows that $\bar{J}_2(\theta_k) - \bar{J}_2(\pi - \theta_k) = 8\alpha^k/\pi + O(\alpha^{2k})$ as $k \rightarrow \infty$. This exponential decay in k explains why the total covariance correction I_{total} remains $O(d/n)$ even as $d \rightarrow \infty$.

Lemma 19. Assume Conjecture 5 is true. Then

$$\mathbf{Cov}(X_\ell, X_{\ell'}) = \lambda^4 \frac{\bar{J}_2(\theta_{|\ell-\ell'|}) - \bar{J}_2(\pi - \theta_{|\ell-\ell'|})}{n} \left(1 + O\left(\frac{1}{n}\right) \right), \quad (26)$$

where θ_k is the angle such that $\theta_k = \cos^{-1}(\alpha^k)$, and the constant in the big $O(\cdot)$ notation is uniform in ℓ, ℓ' .

Proof. From the definition of X_ℓ from (22) as a sum of three terms, the covariance can be written as a sum over pairwise combinations of the terms. The terms involving $g_1^\ell, g_1^{\ell'}$ are mean 0 and independent from layer to layer, so these terms has no contribution to the covariance. We remain with

$$\mathbf{Cov}(X_\ell, X_{\ell'}) = \mathbf{Cov} \left(2\lambda^2 \|\varphi_+^2(\hat{z}^\ell)\| \frac{\|g^\ell\|^2}{n}, 2\lambda^2 \|\varphi_+^2(\hat{z}^{\ell'})\| \frac{\|g^{\ell'}\|^2}{n} \right).$$

The result then follows since $\|g^\ell\|^2 = n$, g^ℓ is independent of $g^{\ell'}$, applying the approximation given in Conjecture 5 and the result of Lemma 16. \square

B.5 Central Limit Theorem

Lemma 20. *We have the following asymptotics for the centered moments of $X_\ell - \mathbf{E}[X_\ell]$:*

$$\mathbf{E} \left[(X_\ell - \mathbf{E}[X_\ell])^3 \right] = O \left(\frac{1}{n^2} \right), \mathbf{E} \left[(X_\ell - \mathbf{E}[X_\ell])^{2m} \right] = O \left(\frac{1}{n^m} \right) \text{ for } m \geq 2,$$

where the constant in the big $O(\cdot)$ notation is uniform in ℓ .

Proof. For convenience of notation, let $c = \sqrt{2} \|\varphi_+(\hat{z}^\ell)\|$. Conditionally on the value of the random value c , $X_\ell = \left\| \alpha \vec{e}_1 + c \lambda n^{\frac{1}{2}} g \right\|^2 = \frac{c^2 \lambda^2}{n} \left\| \frac{\alpha \sqrt{n}}{c \lambda} \vec{e}_1 + g \right\|^2$ is a multiple of a non-central $\chi_k^2 \left(\frac{\alpha^2 n}{c^2 \lambda^2} \right)$ random variable. It is a basic fact about non-central chi squared random variables that if $Y \sim \chi_k^2(\mu)$, then

$$\mathbf{E} \left[(Y - \mathbf{E}[Y])^3 \right] = 8(3\mu + k)$$

Hence

$$\mathbf{E} \left[(X_\ell - \mathbf{E}[X_\ell])^3 | c \right] = \left(\frac{c^2 \lambda^2}{n} \right)^3 8 \left(3 \frac{\alpha^2 n}{c^2 \lambda^2} + n \right) = \frac{24c^4 \lambda^4 \alpha^2 + 8c^6 \lambda^6}{n^2}$$

Using the fact that $0 \leq c \leq \sqrt{2}$ almost surely gives $\mathbf{E} \left[(X_\ell - \mathbf{E}[X_\ell])^3 \right] = O \left(\frac{1}{n^2} \right)$ as desired. A similar computation can be carried out to see the bound for the $2m$ -th central moment by using the $2m$ -th moment of a non-central $\chi_n^2(\mu)$ distribution. This can be seen using the formula for the m -th cumulant: if $Y \sim \chi_k^2(\mu)$, then $K_m = 2^{m-1}(m-1)!(m\mu + k)$, from which it follows by the formula to convert from cumulants to central moments. \square

Lemma 21. *Without loss of generality, assume that $\alpha^2 + \lambda^2 = 1$. Let $S = \sum_{\ell=1}^d \ln X_\ell$ then*

$$\mathbf{E}[S] = -\frac{d}{2n} \cdot (5\lambda^4 + 4\alpha^2\lambda^2) + 2\lambda^2 \sum_{\ell=1}^d h_\ell + O \left(\frac{d}{n^2} \right), \quad (27)$$

$$\mathbf{Var}[S] = \frac{d}{n} \cdot (5\lambda^4 + 4\alpha^2\lambda^2) + \lambda^4 \sum_{1 \leq \ell \neq \ell' \leq d} \frac{\bar{J}_2(\theta_{|\ell' - \ell|}) - \bar{J}_2(\pi - \theta_{|\ell' - \ell|})}{n} + O \left(\frac{d}{n^2} \right), \quad (28)$$

where the constant in the big $O(\cdot)$ notation is uniform in d . Moreover, S is asymptotically Gaussian in the infinite width and depth limit.

Proof. By Lemmas 7 and 10, we see that $\mathbf{E}[X_\ell] = 1 + O(\frac{1}{n})$ and $\mathbf{Var}[X_\ell] = O(\frac{1}{n})$. Moreover, conditionally on $\|\varphi_+(\hat{z}^\ell)\|$, X_ℓ has a non-central χ_n^2 distribution. Hence, by Chebyshev's inequality, we know that for any $\epsilon > 0$, $\frac{X_\ell - \mathbf{E}[X_\ell]}{\mathbf{E}[X_\ell]} = O(\frac{1}{n^{\frac{1}{2} - \epsilon}})$ with probability at least $1 - O(n^{-\epsilon})$. On this event, we can hence take the Taylor series expansion of $\ln(x)$ around $x = 1$ to obtain the following

$$\begin{aligned} \ln(X_\ell) &= \ln(\mathbf{E}[X_\ell]) + \ln \left(1 + \left(\frac{X_\ell - \mathbf{E}[X_\ell]}{\mathbf{E}[X_\ell]} \right) \right) \\ &= \ln(\mathbf{E}[X_\ell]) + \left(\frac{X_\ell - \mathbf{E}[X_\ell]}{\mathbf{E}[X_\ell]} \right) - \frac{1}{2} \left(\frac{X_\ell - \mathbf{E}[X_\ell]}{\mathbf{E}[X_\ell]} \right)^2 \\ &\quad + \frac{1}{3} \left(\frac{X_\ell - \mathbf{E}[X_\ell]}{\mathbf{E}[X_\ell]} \right)^3 - \frac{1}{4} \left(\frac{X_\ell - \mathbf{E}[X_\ell]}{\mathbf{E}[X_\ell]} \right)^4 + O \left(\frac{1}{n^{\frac{5}{2} - 5\epsilon}} \right) \end{aligned}$$

Using Lemmas 7, 12, 16, the bounds from Lemma 20, and the fact that $\ln(X_\ell)$ has finite moments, we can take the expectation and variance of this to obtain

$$\begin{aligned} \mathbf{E}[\ln(X_\ell)] &= 1 + 2\lambda^2 h_\ell - \frac{1}{2} \frac{5\lambda^4 + 4\lambda^2\alpha^2}{n} + O \left(\frac{1}{n^2} \right), \\ \mathbf{Var}[\ln(X_\ell)] &= \frac{5\lambda^4 + 4\lambda^2\alpha^2}{n} + O \left(\frac{1}{n^2} \right), \\ \mathbf{Cov}(\ln(X_\ell), \ln(X_{\ell'})) &= \lambda^4 \frac{\bar{J}_2(\theta_{|\ell - \ell'|}) - \bar{J}_2(\pi - \theta_{|\ell - \ell'|})}{n} + O \left(\frac{1}{n^2} \right), \end{aligned}$$

from which the desired mean and variance formula for S follows.

The fact that S is asymptotically Gaussian follows by application of the Central Limit Theorem for weakly dependent triangular arrays [42, Thm. 2.1] using the bounding sequence $\lambda^4(\bar{J}_2(\theta_r) - \bar{J}_2(\pi - \theta_r))$. Note that the sequence $\hat{z}^0, \hat{z}^1, \hat{z}^2, \dots$ is a Markov chain, which simplifies the verification of the covariance condition in this central limit theorem. The fact that the sequence satisfies the Feller condition is also clear because of the independent Gaussian random vectors that appear in the definition of X_ℓ . \square

Proof of Theorem 1. Recall that $G = \ln \|z^0\|^2/n + \sum_{\ell=1}^d X_\ell$. By the telescoping product (17), we have that

$$\ln \left(\frac{\|z^d\|^2}{n} \right) = \ln \left(\frac{\|z^0\|^2}{n} \right) + \sum_{\ell=1}^d \ln X_\ell \quad (29)$$

Also by Section 5, $\ln \frac{\|z^0\|^2}{n} \sim \ln \frac{\|x\|^2}{n_{\text{in}}} + \ln \frac{\chi_n^2}{n}$ can be written in terms of a chi-squared distribution with n degrees of freedom. By standard facts about $\chi^2(n)$ random variables, we have the a central limit theorem as $n \rightarrow \infty$ namely $\ln \frac{\|z^0\|^2}{n} \rightarrow \mathcal{N} \left(\ln \frac{\|x\|^2}{n_{\text{in}}} - \frac{1}{n}, \frac{2}{n} \right)$. We also note that $\|z^0\|$ is independent of X_ℓ for all ℓ . By Lemma 21 the sum $\sum_{\ell=1}^d \ln X_\ell$ converges to an independent Gaussian random variable. The result follows from the fact that a sum of two independent Gaussians is again Gaussian. \square

Proof of Theorem 4. The proof is analogous to the proof of Theorem 1; the calculation for the activation of the layers is simplified by the fact that which neurons activated in each layer are uncorrelated by (10). \square

B.6 Input-Output Gradient for Balanced ResNets

Proposition 22. *For any input x , the output of a Balanced ResNet can be written as*

$$z^{\text{out}} = M(x)x$$

where $M(x)$ is an $n_{\text{in}} \times n_{\text{out}}$ matrix that depends on x and the random network weights. Moreover, the marginal distribution of $M(x)$ is statistically independent of x . Finally, the matrix $M(x)$ has the property that for almost every x , there is an open neighbourhood $\mathcal{A}(x) \subset \mathbb{R}^n$ containing x so that $M(x)$ is constant on $\mathcal{A}(x)$.

Proof. First note that because the Gaussian weights are continuous random variables, the event that any of the entries of z_i^ℓ are exactly 0 is a probability 0 event. Hence, for almost every input x , it makes sense to consider the derivative of the ReLU function φ'_+ evaluated at the neuron values. By writing the action of a ReLU function on a vector $\varphi(z)$ as a matrix multiplication by a diagonal vector of 1's and 0's as $\varphi(z) = \text{diag}(\varphi'(z))z$, the update rule for ResNets can be written as matrix multiplication, namely

$$z^\ell = \alpha z^{\ell-1} + \lambda \sqrt{\frac{2}{n}} W^\ell \varphi_{s^\ell}(z^{\ell-1}) \quad (30)$$

$$= \alpha z^{\ell-1} + \lambda \sqrt{\frac{2}{n}} W^\ell \text{diag}(\varphi'_{s^\ell}(z^{\ell-1})) z^{\ell-1} \quad (31)$$

$$= \left(\alpha I + \lambda \sqrt{\frac{2}{n}} W^\ell \text{diag}(\varphi'_{s^\ell}(z^{\ell-1})) \right) z^{\ell-1} \quad (32)$$

For the Balanced ResNet, the sign s^ℓ used is independent from layer to layer and neuron to neuron. Therefore the derivative, φ'_{s^ℓ} at any input is equally likely to be 0 or 1 independent of everything else (i.e. no matter if the input z_i is positive or negative, $\varphi'_{s^\ell}(z_i) = 1$ exactly half the time, and $\varphi'_{s^\ell}(z_i) = 0$ exactly half the time.) Hence, from (30), we see that the output z^{out} is equal in distribution to

$$z^{\text{out}} \stackrel{d}{=} \sqrt{\frac{1}{n}} W^{\text{out}} \prod_{\ell=1}^d \left(\alpha I + \lambda \sqrt{\frac{2}{n}} W^\ell \text{diag}(B_i^\ell) \right) \sqrt{\frac{1}{n_{\text{in}}}} W^0 x, \quad (33)$$

where B_i^ℓ are fair $\{0, 1\}$ Bernoulli random variables independent of the weights W^ℓ . Equation (33) shows the marginal distribution of $M(x)$ does not depend on x . Finally note that Bernoulli random variables depend only on the sign of the intermediate neurons z^ℓ at the input x and the signs s^ℓ . Therefore, if we find a neighbourhood $\mathcal{A}(x)$ of x such that none of the neurons change sign from positive to negative within the region x , the matrix will remain constant. This is always possible as long as none of z_i^ℓ are exactly 0 since they are continuous functions of x . But $z_i^\ell = 0$ is a probability zero event, so we can find such a neighbourhood $\mathcal{A}(x)$ for almost every x as desired. \square

Corollary 23. *The derivative of z^{out} with respect to any input x_i has the distribution*

$$\frac{\partial}{\partial x_i} z^{out} \stackrel{d}{=} M(x) e_i$$

Proof. This follows immediately since $M(x)$ is constant on the neighbourhood $\mathcal{A}(x)$. \square

Corollary 24. *$\frac{\partial}{\partial x_i} z^{out}$ has the distribution as the output z^{out} at any input x with $\|x\| = 1$.*

Proof. By the previous results, both are equal in distribution to Mu for any unit vector u , where M is the distribution of the random matrix in (33). \square

C Experiments

Throughout the paper, the Monte Carlo simulations were computed on a single NVIDIA Titan-XP GPU. The main tools used in the neural network simulations are the JAX library [44] (Apache 2.0 License) and the PyTorch library [45] (BSD 3-Clause License). Furthermore, the Python libraries numpy [46] (BSD 3-Clause License), plotnine [47] (based on ggplot2 [48], GNU GPLv2 License), and pandas [49] (BSD 3-Clause License) tremendously helpful. We used Python version 3.6.8 from Anaconda 3 [50] (3-clause BSD License) and Jupyter notebook [51] (3-Clause BSD License).

C.1 Vanilla ResNet and Balanced ResNets: MNIST and CIFAR-10 Experiments

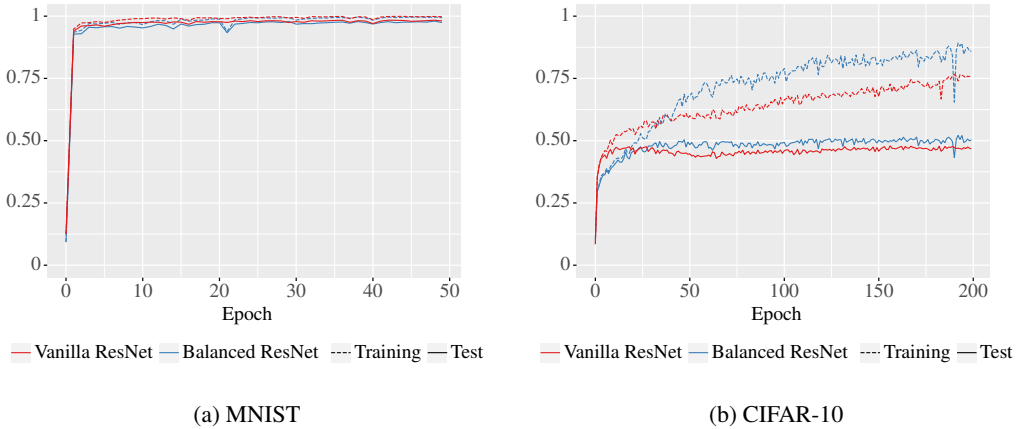


Figure 6: Fully connected Vanilla ResNet and Balanced ResNet accuracies on MNIST & CIFAR-10 using ADAM optimization [52] with learning rate 0.01, $b_1 = 0.9$, $b_2 = 0.999$, $\epsilon = 1e-8$, and batch size 128. Both networks use hidden layer sizes $n = 1000$, $d = 30$ and $\alpha = 1/\sqrt{2}$, $\lambda = 1\sqrt{2}$

These experiments, displayed in Figure 6, investigate the difference between *training performance* of the Vanilla ResNet from (1) and the Balanced ResNet defined in Section 2.1. This is beyond the theory proven in our work which concerns statistical properties of the network on initialization. Both of these architectures are fully-connected networks with skip connections between layers. We observe that both architectures perform similarly in standard training regimes where the networks are much wider than they are deep.

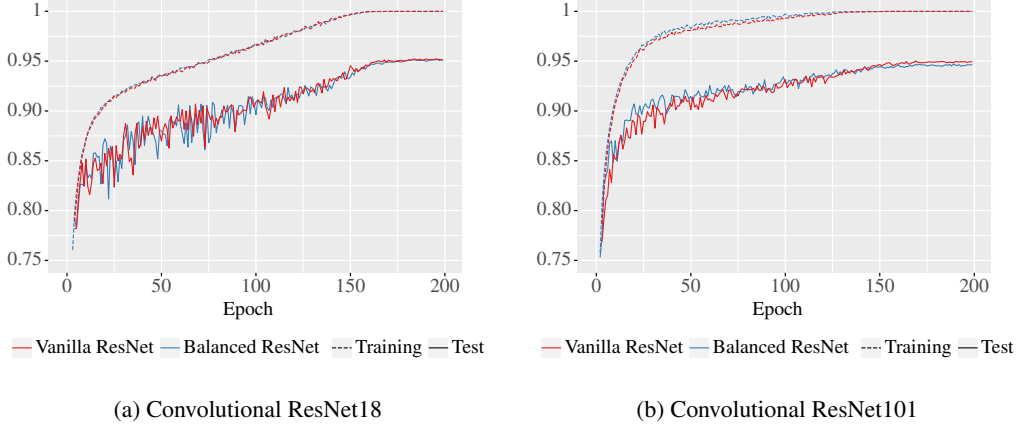


Figure 7: Pre-activation Convolutional ResNet18 and ResNet101 [54] and corresponding balanced versions. Accuracies on CIFAR-10 are reported with learning rate 0.1, momentum 0.9, weight decay $5e-4$, cosine annealing with $T = 200$, and batch size 128.

C.2 Convolutional ResNet CIFAR-10 Experiments

These experiments investigate how the Balanced ResNet architecture modification, namely randomly flipping between φ_+ and φ_- , effects training for deep convolutional ResNets (We call these C-ResNets here to distinguish from the fully connected ones studied in detail in the paper). Although the theory in this paper only are proven for Vanilla ResNets (which are fully connected with skip connections), we expect that the Balanced ResNet tweak does not negatively effect performance in standard training regimes and may allow better initialization for very deep models.

For this experiment, we used the ResNet implementations from the GitHub repository by Liu et al. [53] (MIT License). To create a Balanced Convolutional ResNet (Balanced C-ResNet), we modified the class `PreActBlock` to add flipped ReLUs into the network channel-wise, thinking of channels as the natural generalization of neurons for convolution neural networks. More specifically, we are interested in an input tensor Y of dimension (b, c, h, w) , where b is for batch size, c is for channel, h is for height, and w is for width. Before feeding into a ReLU non-linearity, we will multiply Y by a vector of iid random signs $\{s_j\}_{j \in [c]}$, so that the ReLU output is

$$[\varphi_{s_j}(Y_{i,j,k,l})]_{i,j,k,l} = [\varphi_+(s_j Y_{i,j,k,l})]_{i,j,k,l}. \quad (34)$$

The experiment results are reported in Table 1 and Figure 7. Here, we demonstrate that using the Balanced ResNet architecture does not reduce the performance of existing training regimes. For the deeper ResNet101, the Balanced ResNet seems to perform slightly better early in training but finishes 0.2% worse at the end of training.

Note that even the deeper ResNet101 tested here is still relatively shallow compared to its “width”; most of the layers are either 256 or 1024 channels by 8×8 neurons per channel which represent many more neurons in each hidden layer than the depth of the network, 101. The possible advantage of the Balanced ResNet idea is that it will enable the training architectures which are even *deeper compared to their width*, which are currently not trainable or difficult to train due to initialization issues. A detailed empirical study exploring this idea is needed to study how Balanced ResNets perform beyond the theory proven in this paper.

C.3 Density Plot Calculations

In this section, we describe the calculations required for plotting Figure 1. Firstly, we need to estimate the hypoactivation constant $C_{\alpha,\lambda}$ from Proposition 2 using Monte Carlo simulations. For the choice of $\alpha = \lambda = 1/\sqrt{2}$, we find the constant $C_{\alpha,\lambda} \approx -0.876$, which we use for estimating the mean. See Figure 8 for further simulations with varying α, λ values, and Figure 9 for simulations demonstrating these constants provide accurate prediction for mean and variance.

Architecture	Test Accuracy
Convolutional ResNet18	95.09%
Balanced C-ResNet18	95.08%
Convolutional ResNet101	94.93%
Balanced C-ResNet101	94.70%

Table 1: CIFAR-10 [55] experiment test accuracies with both pre-activation ResNets [54] and the corresponding balanced versions. The ResNet18 architectures used a learning rate 0.05, and the ResNet101 used a learning rate of 0.1. Both architectures used momentum 0.9, weight decay $5e-4$, cosine annealing with $T = 200$, and batch size 128.

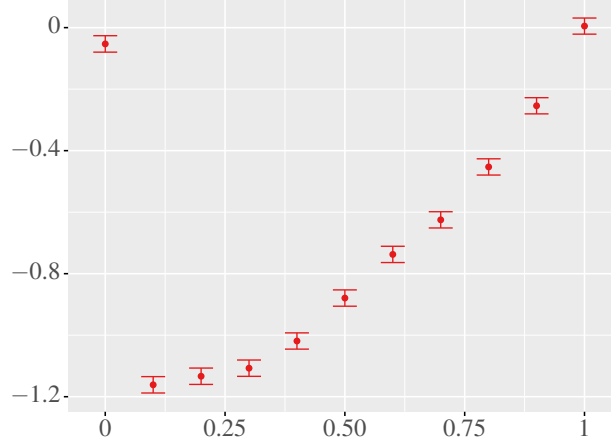


Figure 8: Hypoactivation constant $C_{\alpha, \lambda}$ from Proposition 2 vs. λ^2 . Here we fix $\alpha^2 + \lambda^2 = 1$, and $n = d = 150$.

Next, using the choice of $\|x_{in}\| = 1$ and $\alpha^2 + \lambda^2 = 1$, we can write

$$\ln \|z^{\text{out}}\|^2 \stackrel{d}{=} -\ln n_{in} + G + \ln \|\vec{Z}\|^2.$$

Here we observe that $\|\vec{Z}\|^2 \sim \chi^2(n_{out})$, and therefore we can compute the density of $\ln \|\vec{Z}\|^2$ with a coordinate change

$$\frac{1}{2^{n_{out}/2} \Gamma(n_{out}/2)} \exp\left(\frac{n_{out}x}{2}\right) \exp\left(\frac{-e^x}{2}\right).$$

Finally, since G and $\ln \|\vec{Z}\|^2$ are independent, we can recover the density of $\ln \|z^{\text{out}}\|^2$ via a numerical convolution with a Gaussian density of mean $\mathbb{E}G - \ln n_{in}$ and variance $\text{Var}(G)$.

To compute the density of the infinite width prediction, we first observe that in this limit $z^{\text{out}} \sim \mathcal{N}(0, \sigma^2 I_{n_{out}})$. Therefore it's sufficient to simply compute the variance.

To this goal, we follow the calculations of Hayou et al. [29] with the variance recursion formula (slightly modified to include α)

$$Q_\ell = \alpha^2 Q_{\ell-1} + \lambda^2 \left(\sigma_b^2 + \frac{\sigma_w^2}{2} \left(1 + \frac{f(C_\ell)}{C_\ell} \right) Q_{\ell-1} \right) = Q_{\ell-1},$$

where we plugged in values of $\sigma_b = 0$, $\sigma_w = 2$, and $C_\ell = 1$ and $f(C_\ell) = 0$. To complete the recursion, the initial $Q_0 = \frac{1}{n_{in}} \|x\|^2$. This implies that $z^{\text{out}} \sim \mathcal{N}(0, \frac{1}{n_{in}} I_{n_{out}})$, and therefore

$\|z^{\text{out}}\|^2 \sim \frac{1}{n_{\text{in}}} \chi^2(n_{\text{out}})$, which implies $\ln \|z^{\text{out}}\|^2$ has density

$$\frac{1}{2^{n_{\text{out}}/2} \Gamma(n_{\text{out}}/2)} n_{\text{in}}^{n_{\text{out}}/2} \exp\left(\frac{n_{\text{out}} x}{2}\right) \exp\left(\frac{-n_{\text{in}} e^x}{2}\right).$$

C.4 Additional Monte Carlo Simulations

These additional Monte Carlo simulations provide further comparisons between the infinite depth-and-width limit predictions and finite networks. See Figure 9.




Figure 9: Mean and variance of $G(n, d, \alpha, \lambda)$ from Theorem 1 vs. λ^2 , where $\alpha^2 + \lambda^2 = 1$ and $n = d = 150$. The solid lines indicate our prediction using the infinite-depth-and-width limit. This shows the excess variance appearing in Vanilla vs Balanced ResNets. Also note that fully connected networks with no skip connections, corresponding to $\lambda = 1$, have the highest variance for all balanced ResNets.