

Random Fully Connected Neural Networks as Perturbatively Solvable Hierarchies

Boris Hanin*

Department of Operations Research
and Financial Engineering
Princeton University
bhanin@princeton.edu

January 18, 2023

Abstract

This article considers fully connected neural networks with Gaussian random weights and biases as well as L hidden layers, each of width proportional to a large parameter n . For polynomially bounded non-linearities we give sharp estimates in powers of $1/n$ for the joint cumulants of the network output and its derivatives. Moreover, we show that network cumulants form a perturbatively solvable hierarchy in powers of $1/n$ in that k -th order cumulants in one layer have recursions that depend to leading order in $1/n$ only on j -th order cumulants at the previous layer with $j \leq k$. By solving a variety of such recursions, however, we find that the depth-to-width ratio L/n plays the role of an effective network depth, controlling both the scale of fluctuations at individual neurons and the size of inter-neuron correlations. Thus, while the cumulant recursions we derive form a hierarchy in powers of $1/n$, contributions of order $1/n^k$ often grow like L^k and are hence non-negligible at positive L/n . We use this to study a somewhat simplified version of the exploding and vanishing gradient problem, proving that this particular variant occurs if and only if L/n is large. Several key ideas in this article were first developed at a physics level of rigor in a recent monograph of Daniel A. Roberts, Sho Yaida, and the author. This article not only makes these ideas mathematically precise but also significantly extends them, opening the way to obtaining corrections to all orders in $1/n$.

1 Introduction

We live in an era of big data and cheap computation. This has led to remarkable progress in domains ranging from self-driving cars [35] to automatic drug discovery [33] and machine translation [8]. Underpinning many of these exciting practical developments is a class of computational models called neural networks. While they were originally developed in the 1940's and 1950's [26, 54], the complexity of state-of-the-art neural nets is unprecedented. And yet, despite their empirical utility, a theoretical understanding of how they work and how to make them better is nascent. In fact, it is sometimes said that neural networks are simply too complicated to allow for a rigorous understanding of their key features.

*BH gratefully acknowledges support by the NSF through DMS-2143754, DMS-1855684, and DMS-2133806 as well support from an ONR MURI on Foundations of Deep Learning.

This article adds to the growing body of literature to the contrary. Namely, in the simplest important setting of fully connected networks, we develop a flexible set of probabilistic tools for studying correlation functions of *random fully connected neural networks* at finite width.

A random fully connected neural network, defined precisely in §2.1, is a random field whose law is determined by a fixed and typically non-linear function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ as well as two structural parameters: a depth $L \in \mathbb{N}$ and a width $n \in \mathbb{N}$. For a fixed depth L , in the limit of infinite width $n \rightarrow \infty$ random neural networks converge in distribution to Gaussian processes (see Theorem 2.2). To study finite width effects we therefore consider the higher cumulants for the distribution of the outputs. Our approach is recursive in the network depth L and perturbatively in the reciprocal $1/n$ of the network width. Our main results are:

- We give in Theorem 3.1 sharp estimates in powers of $1/n$ for the size of all joint cumulants of networks outputs and their derivatives. These can be viewed as quantitative results refinements of the central limit theorem at fixed L and large n (cf Theorems 2.2 and Theorem 3.2).
- We derive exact recursions with respect to depth that describe cumulants at one layer in terms of cumulants at the previous layer (see e.g. Corollary 3.4 and Proposition 11.2). These recursions take the form of perturbatively solvable hierarchies in the sense that the recursion for the k -th cumulant at layer $\ell + 1$ involves, at leading order in $1/n$ only the j -th cumulants at layer ℓ with $j \leq k$ (see Corollaries 3.3 and 3.4). This is distinct from but similar in spirit to the remarkable article [29], which derived a dynamical perturbative hierarchy for the so-called neural tangent kernel. As in our work, the perturbative parameter is $1/n$. However, the emphasis in [29] is primarily on training dynamics while ours is on understanding the effect of depth.
- Solving some special cases of the cumulant recursions to leader order in $1/n$ reveals that it is the depth-to-width ratio L/n , which we term the *effective network depth*, rather than the apparent depth L , that is a more informative measure of neural network depth and complexity. Mathematically, this suggests a non-trivial double scaling limit for random neural networks in which

$$n, L \rightarrow \infty \quad \text{and} \quad L/n \rightarrow \xi \in [0, \infty).$$

For non-linear networks this scaling limit has only started to be considered [17, 21, 53, 40]. Even in the very special case of product of L iid random $n \times n$ matrices (sometimes called deep linear networks) the simultaneous large n, L regime has revealed a range of interesting and not fully understood properties (see e.g. references in [1, 2] as well as [16, 21, 22, 42]). In contrast to the $\xi = 0$ regime typically considered in previous work on neural networks (c.f. e.g. [12, 13, 32, 41]), we show that when $\xi > 0$ our double scaling limit is capable of exhibiting non-Gaussian and non-linear effects. We find the following effects to leading order in $1/n$:

- We prove in Corollary 4.1 that for any non-linearity σ from the $K_* = 0$ universality class (defined in §4.2.2) and for $k = 2, 3, 4$, the $2k^{\text{th}}$ cumulants of the output of a random neural network with non-linearity σ grow like $(L/n)^{k/2-1}$. This implies, for instance, that both the correlations between neurons and the fluctuations of a single neuron grow like the effective depth L/n at large L, n (see Remark 4.2

and just after Theorem 4.4). Since the components of the output of a depth 0 neural network are simply iid Gaussians, we see that the effective depth L/n can be thought of as a measure of how close a random fully connected network is to a Gaussian process. Related questions were considered in [4, 15].

- We show in Theorems 4.4 and 4.5 that, for random neural networks initialized as in practice (see §4.1), the variance of the gradient of the network output with respect to either its input or a trainable parameter in its first layer grows like L/n at large L . As explained in §4.4, this gives the first mathematical characterization for fairly general fully connected networks of the so-called exploding and vanishing gradient problem (EVGP). Previously, this problem was solved rigorously in the important special case of random ReLU networks [17, 21] and solved at a physics level of rigor by a somewhat different but related set of ideas in the monograph [53] of Roberts, Yaida, and the author.

The remainder of the introduction is structured as follows. We begin by giving in §2.1 the precise definition of random fully connected neural networks. We then formulate and motivate the main question taken up in this article in §2.2.

2 Background and Motivation

2.1 What is a (Random) Fully Connected Neural Network?

Neural networks are parameterized families of functions. The simplest kind of networks are called fully connected. Each such network is specified by its *architecture*, which consists of an input dimension $n_0 \in \mathbb{Z}_+$, a network depth $L \in \mathbb{Z}_+$, hidden layer widths $n_1, \dots, n_L \in \mathbb{Z}_+$, an output dimension $n_{L+1} \in \mathbb{Z}_+$, and a non-linearity $\sigma : \mathbb{R} \rightarrow \mathbb{R}$. The functions computed by a fully connected network with a given architecture are all maps from \mathbb{R}^{n_0} to $\mathbb{R}^{n_{L+1}}$ that associate to each network input $x_\alpha \in \mathbb{R}^{n_0}$ an output $z_\alpha^{(L+1)} \in \mathbb{R}^{n_{L+1}}$ through a sequence of intermediate representations $z_\alpha^{(\ell)} \in \mathbb{R}^{n_\ell}$ as follows:

$$z_\alpha^{(\ell+1)} := \begin{cases} b^{(\ell+1)} + W^{(\ell+1)}\sigma(z_\alpha^{(\ell)}), & \ell \geq 1 \\ b^{(1)} + W^{(1)}x_\alpha, & \ell = 0 \end{cases}, \quad W^{(\ell+1)} \in \mathbb{R}^{n_{\ell+1} \times n_\ell}, b^{(\ell+1)} \in \mathbb{R}^{n_{\ell+1}}. \quad (2.1)$$

In this recursion, the univariate function σ applied to a vector $z_\alpha^{(\ell)} \in \mathbb{R}^{n_\ell}$ is short-hand for applying it separately to each component. The entries of the matrices $W^{(\ell)}$ and the components of the vectors $b^{(\ell)}$ are called the *weights and biases* in layer ℓ , respectively. One typically refers to

$$z_\alpha^{(\ell)} = \left(z_{1;\alpha}^{(\ell)}, \dots, z_{n_\ell;\alpha}^{(\ell)} \right) \in \mathbb{R}^{n_\ell}$$

as the *vector of pre-activations at layer ℓ* corresponding to the input x_α . The most popular choices of σ in practice include $\text{ReLU}(t) := \max\{0, t\}$ as well the hyperbolic tangent and their variations. We will analyze these cases in detail later (see §4.2), but for our general results make only the following mild assumption

Assumption 2.1. *There exists $r \geq 1$ so that the r -th derivative of σ exists almost everywhere and grows at most polynomially:*

$$\exists k \geq 1 \text{ s.t. } \sup_{x \in \mathbb{R}} \left| (1 + |x|)^{-k} \frac{d^r}{dx^r} \sigma(x) \right| < \infty.$$

The primary objects of study in this article are *random fully connected neural networks*, obtained by choosing network weights and biases to be independent centered Gaussians:

$$W_{ij}^{(\ell)} \sim \mathcal{N}(0, C_W/n_{\ell-1}), \quad b_i^{(\ell)} \sim \mathcal{N}(0, C_b) \quad \text{independent.} \quad (2.2)$$

Here $C_b \geq 0, C_W > 0$ are fixed constants. The $1/n_{\ell-1}$ scaling in the weight variance ensures that the moments of the outputs $z_\alpha^{(L+1)}$ remain uniformly bounded as $n_1, \dots, n_L \rightarrow \infty$ (see e.g. Theorem 2.2 and (3.6)). While we carry out a substantial portion of our analysis for any choice of C_b, C_W , as we will see in §4.1, there often exist distinguished σ -dependent settings of C_b, C_W for which random neural networks at infinite width $n_1, \dots, n_L \rightarrow \infty$ are well-behaved at large depth L .

2.2 Statement and Motivation for Questions Addressed in this Article

The main problem we take up in the present article is to characterize the finite dimensional distributions of a random neural network $x_\alpha \mapsto z_\alpha^{(L+1)}$ (and its derivatives with respect to x_α) in the regime where the input dimension n_0 is arbitrary with the inputs x_α satisfying

$$\frac{1}{n_0} \|x_\alpha\|_2^2 < \infty,$$

the output dimension n_{L+1} is fixed, and the hidden layer widths n_ℓ are large but finite:

$$\exists c, C > 0 \text{ s.t.} \quad cn \leq n_1, \dots, n_L \leq Cn, \quad n \gg 1. \quad (2.3)$$

Our approach will be to describe the random field $x_\alpha \mapsto z_\alpha^{(L+1)}$ perturbatively in $1/n$ and recursively in L . Before proceeding to the technical statements of our results, we pause to address below the following motivational questions:

- Why study random neural networks?
- Why treat $1/n$ as a perturbative parameter?
- Why study networks at finite width?
- What role does σ play?

The primary use of a neural network with a fixed architecture is to find a setting of its parameters $W^{(\ell)}, b^{(\ell)}$ giving rise to a mapping $x_\alpha \mapsto z_\alpha^{(L+1)}$ as in (2.1) that matches a dataset $\{(x_i, f(x_i))\}$ of values for some otherwise unknown function $f : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_{L+1}}$. The optimization procedure for finding these weights and biases is almost always some variant of gradient descent starting with $W^{(\ell)}, b^{(\ell)}$ drawn at random from the distribution (2.2). Thus, studying the properties of random neural networks gives direct insights into the starting conditions for neural network optimization. For instance, understanding the behavior of neural networks at the start of training gives a principled way to set optimization hyperparameters (e.g. the variances C_b, C_W and the step size used for gradient descent). We refer the interested reader our discussion of criticality in §4 as well as to §3.2 in [53] the articles [25, 43, 52, 67] for more on this point.

Next, to address why $1/n$ can reasonably be treated perturbatively, we recall that networks used in practice often have a very large number of parameters. This is reflected in the fact that both the layer widths n_ℓ and the network depth L are often big. Thus, it is

sensible to first understand various limits of neural networks in which the number of parameters tends to infinity. The most well-studied regime of this type (though not the only option cf eg [45, 55, 57, 58, 64]) is the infinite width limit, also known as the *NTK regime*. By definition, this regime is accessed by fixing the depth L , the input and output dimensions n_0, n_{L+1} , and the non-linearity σ , the initialization scheme (2.2) and considering the limit when $n_1, \dots, n_L \rightarrow \infty$. From the random matrix theory point of view this is the free probability regime. In view of the relation (2.3) the NTK regime is obtained by taking $n \rightarrow \infty$ at fixed L and has two salient features:

- At the start of training, neural networks converge to a Gaussian process (see Theorem 2.2 below) [36, 44, 48, 49, 62, 63, 65].
- For the purposes of optimization of a squared loss, the network can be replaced by *its linearization at the start of training* (see [9, 12, 32, 37, 41]).

Taken together, these two points show that at least at infinite width and finite depth, it is the structure of the network at initialization that determines not only the start of training but really the entire training trajectory. However, the infinite width limit is too rigid to capture the ability of real work networks to learn data-dependent features (see e.g. [20, 29, 53, 64]). Only finite width networks can capture these effects! Since the starting point for our analysis of random neural networks at finite width is their infinite width behavior, we take this opportunity to record the following result about random neural networks in the NTK regime.

Theorem 2.2 (Random Networks at fixed L and Infinite Width are Gaussian Processes). *Fix a non-negative integer $r \geq 0$, and suppose $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is r -times differentiable and that its r -th derivative is polynomially bounded:*

$$\exists k \geq 1 \text{ s.t. } \sup_{x \in \mathbb{R}} \left| (1 + |x|)^{-k} \frac{d^r}{dx^r} \sigma(x) \right| < \infty.$$

Then the finite-dimensional distributions of the stochastic process $x_\alpha \mapsto z_\alpha^{(L+1)}$ and its derivatives of order up to r converge to those of a centered Gaussian process with n_{L+1} iid components. The limiting variance of each component

$$K_{\alpha\beta}^{(L+1)} := \lim_{n_1, \dots, n_L \rightarrow \infty} \text{Cov} \left(z_{i;\alpha}^{(L+1)}, z_{i;\beta}^{(L+1)} \right), \quad x_\alpha, x_\beta \in \mathbb{R}^{n_0},$$

satisfies the recursion

$$K_{\alpha\beta}^{(\ell+1)} = \begin{cases} C_b + C_W \langle \sigma(z_\alpha) \sigma(z_\beta) \rangle_{K^{(\ell)}}, & \ell \geq 1 \\ C_b + \frac{C_W}{n_0} \sum_{j=1}^{n_0} x_{j;\alpha} x_{j;\beta}, & \ell = 0 \end{cases}. \quad (2.4)$$

In the statement of Theorem 2.2 and henceforth we reserve the symbol $\langle f(z_\alpha, z_\beta) \rangle_\kappa$ to denote the expectation of $f(z_\alpha, z_\beta)$ with respect to the Gaussian distribution

$$(z_\alpha, z_\beta) \sim \mathcal{N} \left(0, \begin{pmatrix} \kappa_{\alpha\alpha} & \kappa_{\alpha\beta} \\ \kappa_{\alpha\beta} & \kappa_{\beta\beta} \end{pmatrix} \right),$$

where $\kappa_{\alpha\beta} = \kappa(x_\alpha, x_\beta)$ is a given covariance function. The conclusion in Theorem 2.2 is not new, having been obtained many times and under a variety of different assumptions (including for more general architectures) [19, 36, 44, 51, 63]. We refer the interested reader to [19] for a discussion of prior work and note only that convergence of the derivatives of the field $z_\alpha^{(L+1)}$ to its Gaussian limit does not seem to have been previously considered. We give a short proof that includes convergence of derivatives along the lines of the arguments in [19, 36] in Appendix §A.

Let us remark that the relation between $K^{(\ell+1)}$ to $K^{(\ell)}$ supplied by the recursion (2.4) is in general non-linear due to the presence of σ and depends crucially on the weight and bias variances C_W, C_b . Given σ , finding a choice of C_b, C_W so that this recursion is well-behaved (e.g. not exponentially growing or decaying) at large ℓ is an important practical matter [30, 36, 50]. We explain this in §4.1, where we also point out that the different possible large ℓ limits suggest the existence of universality classes of random neural networks.

Finally, as we alluded to above, the tremendous simplification of random neural networks in the NTK regime comes at a steep cost in terms of the descriptive power of the resulting model in the sense that such models cannot learn data-dependent features. We again refer to reader to articles such as [10, 45, 55, 57, 58, 64] in which a different initialization scheme leads to infinite width limits capable of feature learning. Since we do not treat feature learning in this article, we will not elaborate further on this point, referring the interested reader to [9, 20, 29, 53, 60, 64] instead.

Finally, before formulating the new results in this article, we point the interested to a final tranche of physics-style references such as [3, 11, 14, 34, 38, 46, 47, 56], which analyze a variety of questions related to either very wide or infinitely wide networks.

3 Results

In this section we give precise formulations of our results on random neural networks. We start in §3.1 by stating a structural result, Theorem 3.1, on the size of the joint cumulants of a random neural networks and its derivatives in powers of $1/n$. We then extend this order of magnitude estimate to obtain in Theorem 3.2 a general prescription for obtaining series expansions in powers of $1/n$ for expectation values of functions of a random depth $\ell + 1$ neural network in terms of those of a random depth ℓ network. An application of Theorem 3.2 yields explicit recursions for the 4th, 6th, 8th joint cumulants for components $z_{i;\alpha}^{(\ell)}$ of the vector of pre-activations at layer ℓ corresponding to a fixed network input x_α . These recursions are recorded in Corollary 3.4. See §5 for an overview of the proof of Theorems 3.1 and 3.2 as well as Corollary 3.4.

After stating these results, we describe in §4 the important notion of criticality in random neural networks, which corresponds to choosing values of the weight and bias variance parameters C_b, C_W in (2.2) depending on σ in such a way that the infinite width covariance $K_{\alpha\beta}^{(\ell)}$ (and in fact the higher cumulants at finite width as well) is well-behaved at large ℓ . The variety of resulting large ℓ behaviors determines universality classes of random neural networks, as we explain in §4.2.

We then turn in §4.2.1 and §4.2.2 to solving to the $2k$ -th cumulant recursions for $k = 1, 2, 3, 4$ from Corollary 3.4 in random networks tuned to criticality with non-linearities either from the universality class of ReLU or of tanh. We will see that, at large network depth L and leading order in $1/n$, these cumulants depend only on the effective network depth L/n . We

formalize our belief that this is a universal phenomenon in Conjecture 4.3. We then solve in §4.4 a simplified version of the so-called exploding and vanishing gradient problem (EVGP), which concerns the empirical variance of parameter gradients in a random neural network. More precisely, for non-linearities in the $K_* = 0$ universality class we prove (see Theorem 4.5) that to leading order in $1/n$ the empirical variance over weights in the first layer of network gradients grows linearly in the effective network depth L/n . These results depend on Theorem 4.4, which gives for non-linearities in the $K_* = 0$ universality class asymptotics at large ℓ for joint fourth cumulants between the values of partial derivatives of the output of a random neural network.

3.1 Cumulants of Random Neural Networks at Finite Width

Since in the infinite n limit, the field $z_\alpha^{(L+1)}$ is Gaussian (see Theorem 2.2), it is natural to measure perturbations around this regime by considering the behavior of the cumulants of $z_\alpha^{(L+1)}$ and its derivatives. Let us therefore agree that, given random variables X_1, \dots, X_k with finite moments defined on the same probability space, we will denote their mixed cumulant by

$$\kappa(X_1, \dots, X_k) := i^k \frac{\partial^k}{\partial t_1 \cdots \partial t_k} \Big|_{t=0} \log \mathbb{E} [\exp [-i(t_1 X_1 + \cdots + t_k X_k)]] . \quad (3.1)$$

Thus, for example, $\kappa(X_1) = \mathbb{E}[X_1]$ and $\kappa(X_1, X_2) = \text{Cov}(X_1, X_2)$. We refer the reader to §6.1 for background on cumulants. Our first result, Theorem 3.1, gives estimates on the order in $1/n$ of the cumulants of $z_\alpha^{(L+1)}$ and its derivatives. To state it, we fix a finite collection

$$\{x_\alpha, \alpha \in \mathcal{A}\} \subseteq \mathbb{R}^{n_0},$$

of $|\mathcal{A}|$ distinct network inputs. Moreover, we fix a collection of p directional derivatives:

$$D = (d_1, \dots, d_p), \quad d_j := \nabla_{v_j} = \sum_{i=1}^{n_0} v_{ij} \partial_{x_i}. \quad (3.2)$$

and for any multi-index $J = (j_1, \dots, j_p) \in \mathbb{N}^p$ set

$$D_\alpha^J := d_1^{j_1} \cdots d_p^{j_p} \Big|_{x=x_\alpha}$$

for the corresponding differential operator of order $|J| := j_1 + \cdots + j_p$.

Theorem 3.1 (order of magnitude for cumulants of random neural networks). *Fix $r, L \geq 1$ and suppose that $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ satisfies Assumption (2.1) with this value of r . Suppose further that one of the following two conditions holds:*

- σ is smooth
- the limiting covariance matrix

$$\left(\lim_{n \rightarrow \infty} \text{Cov} \left(D_{\alpha_1}^{J_1} z_{1;\alpha_1}^{(\ell)}, D_{\alpha_2}^{J_2} z_{1;\alpha_2}^{(\ell)} \right) \right)_{\substack{|J_1|, |J_2| \leq r \\ \alpha_1, \alpha_2 \in \mathcal{A}}} \quad (3.3)$$

of derivatives of order at most r in the directional derivatives d_1, \dots, d_p of the scalar field $z_{1;\alpha}^{(\ell)}$ is strictly positive definite in the infinite width limit for all $\ell \leq L$.

Then, for each $k, \ell \geq 1$, as $n \rightarrow \infty$

$$\kappa \left(D_{\alpha_1}^{J_1} z_{i_1; \alpha_1}^{(\ell)}, \dots, D_{\alpha_k}^{J_k} z_{i_k; \alpha_k}^{(\ell)} \right) = \begin{cases} 0, & k \text{ odd} \\ O(n^{-\frac{k}{2}+1}), & k \text{ even} \end{cases}, \quad (3.4)$$

where the implicit constant in the error term depends on k , the inputs $x_{\alpha_1}, \dots, x_{\alpha_k}$, the multi-indices J_1, \dots, J_k , the weight and bias variances C_b, C_W , the non-linearity σ , and the layer index ℓ .

We prove Theorem 3.1 in §7. At a physics level of rigor, Theorem 3.1 with $k = 4$ and no derivatives was already derived in the breakthrough work of Yaida [61]. In fact, Yaida's original article went much further: it obtained a recursive formula with respect to ℓ for the fourth cumulant $\kappa(z_{i_1; \alpha_1}^{(\ell)}, \dots, z_{i_4; \alpha_4}^{(\ell)})$ at layer ℓ in terms of the second and fourth cumulants at layer $\ell - 1$. This is analogous to the recursion (2.4) for the infinite width covariance $K_{\alpha_1 \alpha_2}^{(\ell)}$. This theme was then picked up and significantly expanded upon in the physics monograph [53], which computes, among other things, at order $1/n$ the leading corrections to the field $z_{\alpha}^{(\ell)}$ and its derivatives with respect to both x_{α} and model parameters. We will reproduce some of these recursions and obtain new ones of a similar flavor below.

Compared to this prior work the main novelty of Theorem 3.1 is two-fold. First, it gives sharp estimates in powers of $1/n$ for cumulants of all orders (the sharpness can already be seen when $\ell = 2$). Second, it treats in a uniform way the cumulants for not only the values but also all derivatives of $z_{\alpha}^{(\ell)}$.

In order to put Theorem 3.1 into context, we take this opportunity to make an important remark. Namely, it is no accident that the order of magnitude $O(n^{-k+1})$ for the $2k$ -th cumulant in Theorem 3.1 is the same as that of the k -th cumulant of an average of n iid random variables. To see why, let us denote by $\mathcal{F}^{(\ell)}$ the sigma algebra generated by the weights and biases in layers up to and including ℓ . Since we've assumed weights and biases to be Gaussian and independent for different neurons, we find that conditional on $\mathcal{F}^{(\ell)}$ the vectors

$$z_{i; \mathcal{A}}^{(\ell+1)} := \left(z_{i; \alpha}^{(\ell+1)}, \alpha \in \mathcal{A} \right)$$

are independent centered Gaussians with the covariance

$$\Sigma_{\alpha \alpha'}^{(\ell)} := \text{Cov} \left(z_{i; \alpha}^{(\ell+1)}, z_{i; \alpha'}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) = C_b + \frac{C_W}{n_{\ell}} \sum_{j=1}^{n_{\ell}} \sigma \left(z_{j; \alpha}^{(\ell)} \right) \sigma \left(z_{j; \alpha'}^{(\ell)} \right). \quad (3.5)$$

Put another way, we have the following equality in distribution:

$$\left(z_{i; \alpha}^{(\ell+1)} \right)_{i=1}^{n_{\ell+1}} \stackrel{d}{=} \left(\left(\Sigma^{(\ell)} \right)^{1/2} Z_i \right)_{i=1}^{n_{\ell+1}}, \quad (3.6)$$

where Z_i are i.i.d. standard Gaussians and are independent of the random covariance matrix

$$\Sigma^{(\ell)} := \left(\Sigma_{\alpha \alpha'}^{(\ell)} \right)_{\alpha, \alpha' \in \mathcal{A}}.$$

Inspecting (3.5) shows this conditional covariance has the structure of an average of n_{ℓ} identically distributed random matrices, each depending only on the pre-activations $z_{i; \mathcal{A}}^{(\ell)}$ of a single neuron at layer ℓ . We will refer to such random variables as *collective observables*. The key

point is that while the $z_{i;\mathcal{A}}^{(\ell)}$ are not independent for different i at finite width, we will show that they are sufficiently weakly correlated that the cumulants of $\Sigma_{\alpha\alpha'}^{(\ell)}$ have the same order in n as if they were exactly independent (see Theorem 7.3 and Lemma 7.5). This reasoning applies equally well to derivatives of $z_{i;\alpha}^{(\ell+1)}$ and explains the order of magnitude of the estimates in Theorem 3.1.

Collective observables such as $\Sigma_{\alpha\alpha'}^{(\ell)}$ are a convenient book-keeping device for studying the fully distribution of

$$D^{\leq r} z_{i;\mathcal{A}}^{(\ell+1)} := \left(D_{\alpha}^J z_{i;\alpha}^{(\ell+1)}, \quad \alpha \in \mathcal{A}, J \in \mathbb{N}^p, |J| \leq r \right).$$

Indeed, the conditional Gaussian structure (3.6) means the cumulants of $D^{\leq r} z_{i;\mathcal{A}}^{(\ell+1)}$ are easily expressible in terms of those of $D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha\alpha'}^{(\ell)}$. For example, using Wick's Theorem and the multi-linearity of cumulants reveals that the fourth cumulant

$$\kappa \left(D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}^{(\ell+1)}, D_{\alpha_2}^{J_2} z_{i_2;\alpha_2}^{(\ell+1)}, D_{\alpha_3}^{J_3} z_{i_3;\alpha_3}^{(\ell+1)}, D_{\alpha_4}^{J_4} z_{i_4;\alpha_4}^{(\ell+1)} \right)$$

equals

$$\delta_{i_1 i_2} \delta_{i_3 i_4} \kappa_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(J_1 J_2)(J_3 J_4),(\ell)} + \delta_{i_1 i_3} \delta_{i_2 i_4} \kappa_{(\alpha_1 \alpha_3)(\alpha_2 \alpha_4)}^{(J_1 J_3)(J_2 J_4),(\ell)} + \delta_{i_1 i_4} \delta_{i_2 i_3} \kappa_{(\alpha_1 \alpha_4)(\alpha_2 \alpha_3)}^{(J_1 J_4)(J_2 J_3),(\ell)},$$

where we've abbreviated

$$\kappa_{(\alpha_1 \alpha'_1), \dots, (\alpha_k \alpha'_k)}^{(J_1 J'_1), \dots, (J_k J'_k), (\ell)} := \kappa \left(D_{\alpha_1}^{J_1} D_{\alpha'_1}^{J'_1} \Sigma_{\alpha_1 \alpha'_1}^{(\ell)}, \dots, D_{\alpha_k}^{J_k} D_{\alpha'_k}^{J'_k} \Sigma_{\alpha_k \alpha'_k}^{(\ell)} \right). \quad (3.7)$$

The general pattern is that odd mixed cumulants in the $z_{\alpha}^{(\ell)}$ and its derivatives vanish by symmetry and $2k$ -th cumulants are finite sums of cumulants $\kappa_{(\alpha_1 \alpha'_1), \dots, (\alpha_k \alpha'_k)}^{(J_1 J'_1), \dots, (J_k J'_k), (\ell)}$. In Corollary 3.4 and Proposition 10.2 below we will obtain a range of recursions for these cumulants and thereby implicitly for the cumulants of the original field $z_{\alpha}^{(\ell)}$ as well.

To obtain such recursions note that, at first glance, the estimate (3.4) only gives the order of magnitude in powers of $1/n$ of the cumulants of $\kappa \left(D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}^{(\ell)}, \dots, D_{\alpha_k}^{J_k} z_{i_k;\alpha_k}^{(\ell)} \right)$ but does not seem to provide information about their structural dependence on the remaining model parameters C_b, C_W, σ, ℓ . However, such information can be obtained by combining (3.4) with an additional argument. To state the result, let us agree to write

$$\kappa_{\alpha_1 \alpha_2}^{(\ell)} := \text{Cov} \left(z_{i;\alpha_1}^{(\ell)}, z_{i;\alpha_2}^{(\ell)} \right), \quad \kappa_{\alpha_1 \alpha_2}^{J_1 J_2, (\ell)} := D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \kappa_{\alpha_1 \alpha_2}^{(\ell)}. \quad (3.8)$$

Physicists might refer to $\kappa_{\alpha_1 \alpha_2}^{(\ell)}$ as a dressed two point function. Moreover, let us denote by $\langle \cdot \rangle_{\kappa^{(\ell)}}$ the expectation with respect to a collection of centered jointly Gaussian random vectors

$$D_{\mathcal{A}}^{\leq r} z_i = \left(D_{\alpha}^J z_{i;\alpha}, \quad \alpha \in \mathcal{A}, |J| \leq r \right)$$

with the same covariance

$$\text{Cov} \left(D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}, D_{\alpha_2}^{J_2} z_{i_2;\alpha_2} \right) := \text{Cov} \left(D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}^{(\ell)}, D_{\alpha_2}^{J_2} z_{i_2;\alpha_2}^{(\ell)} \right) = \delta_{i_1 i_2} \kappa_{\alpha_1 \alpha_2}^{J_1 J_2, (\ell)}$$

as the true vectors of derivatives $D_{\mathcal{A}}^{\leq r} z_{i;\mathcal{A}}^{(\ell)}$ in each component separately but zero covariance for different i .

Theorem 3.2 (perturbative expansions of expectations of observables at finite width). *Fix $r \geq 0$ and suppose that f is both a continuous function and a tempered distribution. Then for any $q_* \geq 0$ we have*

$$\begin{aligned} & \mathbb{E} \left[f \left(D^{\leq r} z_{1,\mathcal{A}}^{(\ell+1)}, \dots, D^{\leq r} z_{m,\mathcal{A}}^{(\ell+1)} \right) \right] \\ &= \sum_{q=0}^{2q_*} \frac{(-1)^q}{2^q q!} \left\langle \mathbb{E} \left[\left(\sum_{\substack{|J|,|J'| \leq r \\ \alpha, \alpha' \in \mathcal{A}}} \Delta_{\alpha\alpha'}^{JJ',(\ell)} \sum_{j=1}^m \partial_{D_{\alpha}^{J} z_{j;\alpha}} \partial_{D_{\alpha'}^{J'} z_{j;\alpha'}} \right)^q \right] f \left(D_{\mathcal{A}}^{\leq r} z_1, \dots, D_{\mathcal{A}}^{\leq r} z_m \right) \right\rangle_{\kappa^{(\ell+1)}} \\ &+ O(n^{-q_*-1}), \end{aligned} \quad (3.9)$$

where the sum is over multi-indices $J, J' \in \mathbb{N}^p$ of order at most r , we've set

$$\Delta_{\alpha\alpha'}^{JJ',(\ell)} := D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha\alpha'}^{(\ell)} - \mathbb{E} \left[D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha\alpha'}^{(\ell)} \right], \quad (3.10)$$

and the derivatives $\partial_{D_{\alpha}^J z_{j;\alpha}}$ are interpreted in the weak sense if f is not differentiable.

We prove Theorem 3.2 in §9 and give the main idea of the proof in §5. By substituting various polynomials in $z_{\alpha}^{(\ell+1)}$ and its derivatives for f into the perturbative expansion (3.9), it is now straightforward in principle to deduce recursions for the cumulants $\kappa_{(\alpha_1 \alpha'_1), \dots, (\alpha_k \alpha'_k)}^{(J_1 J'_1), \dots, (J_k J'_k), (\ell+1)}$ at layer $\ell + 1$ in terms of objects of the same type at layer ℓ . In particular, we have the following

Corollary 3.3 (cumulants in random neural networks form a hierarchy to leading order in $1/n$). *With the assumptions of Theorem 3.1, the mixed cumulant*

$$\kappa \left(D_{\alpha_1}^{J_1} z_{i_1; \alpha_1}^{(\ell+1)}, \dots, D_{\alpha_{2k}}^{J_{2k}} z_{i_{2k}; \alpha_{2k}}^{(\ell+1)} \right)$$

equals

$$\sum_{\substack{j \leq k \\ J'_i, i=1, \dots, 2j \\ |J'_1| + \dots + |J'_{2j}| \\ \leq |J_1| + \dots + |J_{2k}|}} C(J'_i, K_{\alpha_i \alpha'_i}^{(\ell)}, i, i' = 1, \dots, 2j) \kappa \left(D_{\alpha_1}^{J'_1} z_{1; \alpha_1}^{(\ell)}, \dots, D_{\alpha_{2j}}^{J'_{2j}} z_{2j; \alpha_{2j}}^{(\ell)} \right) + O(n^{-k}), \quad (3.11)$$

where the sum is over multi-indices J'_i , the constants $C(J'_i, K_{\alpha_i \alpha'_i}^{(\ell)}, i, j = 1, \dots, 2k)$ depend only on the multi-indices J'_i and the infinite width covariance $K^{(\ell)}$, while the implicit constant in the error term depends on k , the inputs $x_{\alpha_1}, \dots, x_{\alpha_k}$, the multi-indices J_1, \dots, J_k , the weight and bias variances C_b, C_W , the non-linearity σ , and the layer index ℓ .

We do not know how to efficiently compute the coefficients C in the recursion (3.11) for arbitrary k . Instead, we will compute them by hand for small values of k in the two settings:

- We fix a single input $x_{\alpha} \in \mathbb{R}^{n_0}$ and obtain a recursion for the cumulants:

$$\kappa_{2k; \alpha}^{(\ell)} := \kappa_{(\alpha\alpha) \dots (\alpha\alpha)}^{(00) \dots (00), (\ell)} = \frac{1}{(2k-1)!!} \kappa \left(\underbrace{z_{i; \alpha}^{(\ell+1)}, \dots, z_{i; \alpha}^{(\ell+1)}}_{2k \text{ times}} \right) \quad (3.12)$$

when $k = 2, 3, 4$ (the pairs $(\alpha\alpha)$ and (00) both appear k times). See Corollary 3.4.

- We again fix a single input $x_\alpha \in \mathbb{R}^{n_0}$ and consider any two partial derivatives

$$d_j := \sum_{i=1}^{n_0} c_{i,j} \partial_{x_j} \Big|_{x=x_\alpha}, \quad j = 1, 2.$$

We obtain recursions for

$$\kappa_{(i_1 i'_1)(i_2 i'_2); \alpha}^{(\ell)} := \kappa \left(d_{i_1} d_{i'_1} \Delta_{\alpha\alpha}^{(\ell)}, d_{i_2} d_{i'_2} \Delta_{\alpha\alpha}^{(\ell)} \right). \quad (3.13)$$

These cumulants are the building blocks for understanding the fourth mixed cumulants of $z_{i;\alpha}^{(\ell+1)}, d_1 z_{i;\alpha}^{(\ell+1)}, d_2 z_{i;\alpha}^{(\ell+1)}$. For example,

$$\kappa \left(d_1 z_{i;\alpha}^{(\ell+1)}, d_1 z_{i;\alpha}^{(\ell+1)}, d_2 z_{i;\alpha}^{(\ell+1)}, d_2 z_{i;\alpha}^{(\ell+1)} \right) = \kappa_{(11)(22); \alpha}^{(\ell)} + 2\kappa_{(12)(12); \alpha}^{(\ell)}.$$

This will be done in the course of proving Theorem 4.5. We refer the interested reader to Proposition 10.2 for the precise recursions and to Proposition 11.2 their solutions for non-linearities from the $K_* = 0$ universality class (including tanh).

In order to facilitate a compact form for the recursions described in first bullet point, let us write

$$T_{i,j;\alpha}^{(\ell)} := C_W^j \left\langle \partial_z^i \left\{ \left(\sigma^2(z) - \langle \sigma^2(z) \rangle_{K_{\alpha\alpha}^{(\ell)}} \right)^j \right\} \right\rangle_{K_{\alpha\alpha}^{(\ell)}}, \quad (3.14)$$

with the derivatives interpreted in the weak sense if σ is not sufficiently smooth and where we remind the reader of our standing notation

$$\langle f(z) \rangle_K = \int_{-\infty}^{\infty} f(zK^{1/2}) e^{-\frac{z^2}{2}} \frac{dz}{\sqrt{2\pi}}, \quad K \geq 0.$$

Corollary 3.4. *Fix $r \geq 1$ and suppose that $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ satisfies Assumption (2.1) with this value of r . Consider a depth L random neural network with input dimension n_0 , hidden layer widths n_1, \dots, n_L , output dimension n_{L+1} and non-linearity σ . Fix $x_\alpha \in \mathbb{R}^{n_0}$ and define*

$$\chi_{||;\alpha}^{(\ell)} := \frac{1}{2} T_{2,1;\alpha}^{(\ell)} = \frac{C_W}{2} \langle \partial_z^2 \sigma(z)^2 \rangle_{K_{\alpha\alpha}^{(\ell)}},$$

where the second derivative is interpreted in the weak sense if σ is not twice differentiable. For each $\ell = 1, \dots, L$, in the notation of (3.12), the fourth cumulant satisfies

$$\kappa_{4;\alpha}^{(\ell+1)} = \frac{T_{0,2;\alpha}^{(\ell)}}{n_\ell} + \left(\chi_{||;\alpha}^{(\ell)} \right)^2 \kappa_{4;\alpha}^{(\ell)} + O(n^{-2}). \quad (3.15)$$

Further, the 6-th cumulant satisfies

$$\kappa_{6;\alpha}^{(\ell+1)} = \frac{T_{0,3;\alpha}^{(\ell)}}{n_\ell^2} + \frac{3T_{2,2;\alpha}^{(\ell)}}{2n_\ell} \chi_{||;\alpha}^{(\ell)} \kappa_{4;\alpha}^{(\ell)} - \frac{3T_{4,1;\alpha}^{(\ell)}}{8} \left(\chi_{||;\alpha}^{(\ell)} \kappa_{4;\alpha}^{(\ell)} \right)^2 + \left(\chi_{||;\alpha}^{(\ell)} \right)^3 \kappa_{6;\alpha}^{(\ell)} + O(n^{-3}). \quad (3.16)$$

Finally, the 8-th cumulant satisfies:

$$\begin{aligned}
\kappa_{8;\alpha}^{(\ell+1)} &= \frac{1}{n_\ell^3} \left(T_{0,4;\alpha}^{(\ell)} - 3 \left(T_{0,2;\alpha}^{(\ell)} \right)^2 \right) \\
&+ \frac{1}{n_\ell^2} \left[2T_{2,3;\alpha}^{(\ell)} \chi_{||;\alpha}^{(\ell)} - 12T_{0,2;\alpha}^{(\ell)} \left(\chi_{||;\alpha}^{(\ell)} \right)^2 + \frac{3}{2} \left(T_{2,2;\alpha}^{(\ell)} \right)^2 - \frac{3}{2} T_{4,1;\alpha}^{(\ell)} T_{0,2;\alpha}^{(\ell)} \right] \kappa_{4;\alpha}^{(\ell)} \\
&- \frac{1}{n_\ell} \left[2T_{2,2;\alpha}^{(\ell)} T_{4,1;\alpha}^{(\ell)} \chi_{||;\alpha}^{(\ell)} - \frac{1}{2} T_{4,2;\alpha}^{(\ell)} \left(\chi_{||;\alpha}^{(\ell)} \right)^2 + \left(\chi_{||;\alpha}^{(\ell)} \right)^4 \right] \left(\kappa_{4;\alpha}^{(\ell)} \right)^2 \\
&+ \frac{1}{n_\ell} \left[5T_{0,2;\alpha}^{(\ell)} T_{4,1;\alpha}^{(\ell)} \chi_{||;\alpha}^{(\ell)} + 12T_{2,2;\alpha}^{(\ell)} \left(\chi_{||;\alpha}^{(\ell)} \right)^2 \right] \kappa_{6;\alpha}^{(\ell)} \\
&+ \frac{3}{32} \left(T_{4,1;\alpha}^{(\ell)} \right)^2 \left(\chi_{||;\alpha}^{(\ell)} \right)^2 \left(\kappa_{4;\alpha}^{(\ell)} \right)^3 - \frac{1}{2} \left(\chi_{||;\alpha}^{(\ell)} \right)^3 T_{4,1;\alpha}^{(\ell)} \kappa_{4;\alpha}^{(\ell)} \kappa_{6;\alpha}^{(\ell)} \\
&+ \left(\chi_{||;\alpha}^{(\ell)} \right)^4 \kappa_{8;\alpha}^{(\ell)} + O(n^{-4}). \tag{3.17}
\end{aligned}$$

The initial condition for the recursions (3.15)-(3.17) is that $\kappa_{2k;\alpha}^{(1)} = 0$ for all $k \geq 2$.

Remark 3.5. Note that for $k = 2, 3, 4$, we therefore see that to leading order in $1/n$ the recursions for $\kappa_{2k;\alpha}^{(\ell+1)}$ depends only on $\kappa_{2j;\alpha}^{(\ell)}$ for $j \leq k$. This allows us to interpret (3.15) - (3.17) as a forming the start of hierarchy in powers of $1/n$ describing the cumulants of the output of a random neural network.

Let us briefly compare Corollary 3.4 to results in prior work:

- In the special case when σ is 1-homogeneous (i.e. is linear, ReLU, leaky ReLU, etc, see (4.5)), the full distribution of a neuron pre-activation $z_{i;\alpha}^{(\ell)}$ can be worked out in closed form. Namely, as we explain in §4.2.1 and Appendix C, it has the same distribution as a Gaussian with an independent random variance given by a product of independent weighted chi-squared random variables. This was first pointed out in [17, 21] and described in the language of special functions (namely Meijer G functions) in [66]. For such non-linearities obtaining the recursions (3.15)-(3.17) is not new.
- The breakthrough work of Yaida [61] was the first to obtain, at a physics level of rigor, the recursion (3.15) and probe its solutions at large ℓ .
- The ideas of Yaida in [61] then seeded the development in the monograph [53] of Roberts, Yaida, and the author a much richer analysis, producing at a physics level of rigor many recursions similar in flavor to (3.15)-(3.17) that describe the the behavior of objects such as network derivatives at the start of training, the NTK at the start of training, and even the change in the NTK and the resulting output of a *trained* network. Many of these results go far beyond what we are currently capable of doing mathematically.
- The analysis in [53] never required studying cumulants $\kappa_{2k;\alpha}^{(\ell)}$ for $k \geq 3$, and while the techniques developed there can certainly be used to obtain the recursions (3.16) and (3.17) we take a rather different approach in this article that produces such recursions more directly.

The functional $\chi_{||;\alpha}^{(\ell)}$ plays a fundamental role in the recursive description of random neural networks supplied by Corollary 3.4, whose proof is in §9. In the following section we explain

a principled procedure, called *tuning to criticality*, that reveals its origin (as well as that of a similar object we denote $\chi_{\perp;\alpha}^{(\ell)}$) and explains how to choose C_b, C_W so that these functionals are approximately equal to 1 at large ℓ . As we will see, such a choice will ensure that the recursions in Corollary 3.4 and their infinite width counterpart (2.4) have well-behaved solutions at large ℓ . We will then return in §4.2.1 and §4.2.2 to solving the recursions from Corollary 3.4 in random networks tuned to criticality (see (4.7) and Corollary 4.1).

4 Criticality and Universality in Random Neural Networks at Large Depth

In the previous section we presented two kinds of results about the structure of random neural networks at large but finite width. The first, Theorem 3.1, concerned the order of magnitude for cumulants of the output of such a random network and its derivatives. The second, Theorem 3.2 and Corollary 3.4, spelled out recursions with respect to the layer index ℓ that describe, to leading order in $1/n$, network cumulants at layer $\ell + 1$ in terms of those at layer ℓ . Our purpose in forthcoming sections is to analyze these recursions at large ℓ and to apply this analysis to obtain results about the structure of gradients in deep fully connected networks. Before doing this, we must take a step back and ask: for which σ, C_b, C_W are the recursions (2.4) describing the infinite width covariance $K^{(\ell)}$ well-behaved at large ℓ ?

In section §4.1, we recall a more or less canonical answer to this question whose roots are in the early articles [51, 52] and that was recently spelled out in the generality presented here in [53]. This procedure, called *tuning to criticality*, prescribes combinations of σ, C_b, C_W for which $K^{(\ell)}$ is indeed well-behaved at large ℓ . As we shall see below, the term *criticality* is meant to be evocative of its use in the analysis of 2d systems in statistical mechanics in that *tuning to criticality* consists of choosing C_b, C_W so that the infinite width covariance function $K^{(\ell)}$ is as close to constant as a function ℓ as possible.

At a high level, there are two reasons to ask that $K^{(\ell)}$ be slowly varying as a function of ℓ . First, it arguably only makes sense to study perturbative corrections in $1/n$ recursively in ℓ if the limiting $n \rightarrow \infty$ covariance structure does not change too rapidly between consecutive layers. Second, and perhaps more importantly, as explained and thoroughly validated in [50, 52], deep fully connected networks (without residual connections [24], batch normalization [31], etc) are numerically stable enough for gradient-based optimization to succeed only if they are tuned to criticality.

We discuss in §4.2 how considerations underlying criticality naturally give rise to a notion of universality classes for random neural networks. Even the correct definition of universality is still not fully understood. Unlike in random matrix theory, universality for random neural networks depends not on the statistics of the individual weights and biases (though this is also an interesting direction to consider e.g. [19]) but rather on the effect of the non-linearity σ on the behavior of the infinite width covariance $K^{(\ell)}$ at large values of the depth ℓ .

Before giving the details, we take this opportunity to emphasize, as we have elsewhere, that the definitions of criticality and universality, the approach to solving the recursions for $\kappa_{2k;\alpha}^{(\ell)}$ from Corollary 3.4, and the resulting lessons learned about the role of the effective network depth L/n closely follow the ideas developed in the monograph [53]. Though we pursue them in a somewhat different way, the author would nonetheless like to acknowledge that his co-authors Dan Roberts and Sho Yaida in the book deserve significant credit.

4.1 Tuning to Criticality

As originally explained in [51, 52] and recently spelled out in a definitive way in [53], tuning a neural network to criticality means seeking choices of (C_b, C_W) that lead to critical fixed points of the form (K_*, K_*, K_*) for the recursion (2.4), viewed as a dynamical system describing $(K_{\alpha\alpha}^{(\ell)}, K_{\beta\beta}^{(\ell)}, K_{\alpha\beta}^{(\ell)})$ with time parameter ℓ . Specifically, criticality requires

$$\exists K_* \geq 0 \quad \text{s.t.} \quad K_* = C_b + C_W \langle \sigma^2(z) \rangle_{K_*} \quad (*)$$

$$\forall \ell \geq 1 \quad \left. \frac{\partial K_{\alpha\alpha}^{(\ell)}}{\partial K_{\alpha\alpha}^{(1)}} \right|_{K_{\alpha\alpha}^{(1)}=K_*} = 1 \quad (**)$$

$$\forall \ell \geq 1 \quad \left. \frac{\partial K_{\alpha\beta}^{(\ell)}}{\partial K_{\alpha\beta}^{(1)}} \right|_{K_{\alpha\alpha}^{(1)}=K_{\alpha\alpha}^{(1)}=K_{\alpha\beta}^{(1)}=K_*} = 1, \quad (**)$$

where

$$K_{\alpha\alpha}^{(1)} = C_b + C_W K_{\alpha\beta}^{(0)}, \quad K_{\alpha\beta}^{(0)} := \frac{1}{n_0} \sum_{j=1}^{n_0} x_{j;\alpha} x_{j;\beta}, \quad x_\alpha, x_\beta \in \mathbb{R}^{n_0}.$$

The intuitive meaning of these conditions is as follows. Due to Theorem 2.2, the first guarantees the existence of a fixed point K_* for of the recursion

$$K_{\alpha\alpha}^{(\ell+1)} = C_b + C_W \langle \sigma^2(z) \rangle_{K_{\alpha\alpha}^{(\ell)}} \quad (4.1)$$

of the infinite width variance. In particular, (*) implies

$$K_{\alpha\alpha}^{(1)} = C_b + \frac{C_W}{n_0} \|x_\alpha\|^2 = K_* \quad \implies \quad K_{\alpha\alpha}^{(\ell)} = \lim_{n \rightarrow \infty} \text{Var} [z_{i;\alpha}^{(\ell)}] = K_* \quad \forall \ell \geq 1.$$

Thus, if a network is tuned to criticality, there is a critical radius

$$K_{\text{crit}}^2 := n_0 C_W^{-1} (K_* - C_b)$$

such that for inputs x_α on the sphere of radius K_{crit} the variance of $z_{i;\alpha}^{(\ell)}$ is independent of ℓ in the infinite width limit. In non-critical networks, we expect this variance to either grow or decay exponentially in ℓ , leading to numerical instabilities. The second condition (**) considers the infinite width limit of the variance of $z_{i;\alpha}^{(\ell)}$ for an input x_α for which $K_{\alpha\alpha}^{(1)}$ is close to K_* . Specifically, condition (**) requires for all $\ell \geq 1$ that

$$K_{\alpha\alpha}^{(1)} = \text{Var}[z_{i;\alpha}^{(1)}] = K_* + \delta K \quad \implies \quad K_{\alpha\alpha}^{(\ell)} = \lim_{n \rightarrow \infty} \text{Var} [z_{i;\alpha}^{(\ell)}] = K_* + \delta K + o(\delta K).$$

This guarantees that the fixed point K_* of the recursion (4.1) is critical and hence that for inputs near the sphere of radius K_{crit} the variance of the resulting pre-activations $z_{i;\alpha}^{(\ell)}$ is approximately constant in ℓ at large n . The final condition (\perp) considers instead the covariance between two inputs on the sphere of radius K_{crit} . Namely, given two nearby network inputs $x_\alpha, x_\beta \in \mathbb{R}^{n_0}$ with

$$K_{\alpha\alpha}^{(1)} = K_{\beta\beta}^{(1)} = K_*, \quad K_{\alpha\beta}^{(1)} = C_b + \frac{C_W}{n_0} \sum_{j=1}^{n_0} x_{j;\alpha} x_{j;\beta} = K_* - \delta K,$$

the third condition asks that

$$K_{\alpha\beta}^{(\ell)} = \lim_{n \rightarrow \infty} \text{Cov} \left(z_{i;\alpha}^{(\ell)}, z_{i;\beta}^{(\ell)} \right) = K_* - \delta K + o(\delta K), \quad \forall \ell.$$

This ensures that the covariance between pre-activations $z_{i;\alpha}^{(\ell)}$ and $z_{i;\beta}^{(\ell)}$ corresponding to two nearby inputs on the K_{crit} -sphere are approximately independent of ℓ at large n . A simple computation directly from the recursion (2.4) shows that

$$\chi_{\parallel}(K) := \left. \frac{\partial K_{\alpha\alpha}^{(\ell+1)}}{\partial K_{\alpha\alpha}^{(\ell)}} \right|_{K_{\alpha\alpha}^{(\ell)}=K} = \frac{C_W}{2} \langle \partial_z^2(\sigma^2(z)) \rangle_K \quad (4.2)$$

$$\chi_{\perp}(K) := \left. \frac{\partial K_{\alpha\beta}^{(\ell+1)}}{\partial K_{\alpha\beta}^{(\ell)}} \right|_{K_{\alpha\alpha}^{(\ell)}=K_{\beta\beta}^{(\ell)}=K_{\alpha\beta}^{(\ell)}=K} = C_W \langle (\partial_z \sigma(z))^2 \rangle_K. \quad (4.3)$$

Hence, all together, tuning to criticality requires

$$\boxed{K_* \geq 0 \text{ s.t. } K_* = C_b + C_W \langle \sigma^2(z) \rangle_{K_*} \quad \text{and} \quad \chi_{\parallel}(K_*) = \chi_{\perp}(K_*) = 1.} \quad (4.4)$$

4.2 Universality Classes of Random Neural Networks: Two Examples In Search of a General Definition

We now turn to discussing the notion of universality classes for random neural networks. To start, recall from Theorem 2.2 that the behavior at large depth ℓ of random fully connected neural networks at infinite width is completely specified by the asymptotics of the limiting covariance function $K^{(\ell)}$. Observe, moreover, that the coefficients in the recursions for $k = 2, 3, 4$ of the cumulants $\kappa_{2k;\alpha}^{(\ell+1)}$ from Corollary 3.4, which by Theorem 3.1 determine the behavior of random neural networks at finite width to the first four orders in $1/n$, are completely determined by σ , the infinite width covariance $K^{(\ell)}$, and cumulants $\kappa_{2j;\alpha}^{(\ell)}$, $j \leq k$. It is therefore in terms of the large ℓ behavior of $K^{(\ell)}$ that we should hope to define universality classes of random neural networks at large depth. At present it is not clear what the correct general definition of such a universality class should be. We content ourselves instead with studying two important classes of examples.

4.2.1 The Universality Class of ReLU

The most popular non-linearities used in practice are positively homogeneous of degree 1, i.e. have the form

$$\sigma(t) = (a_- \mathbf{1}_{\{t < 0\}} + a_+ \mathbf{1}_{\{t > 0\}})t, \quad a_-, a_+ \in \mathbb{R}, \quad a_- \neq a_+, \quad a_-^2 + a_+^2 \neq 0. \quad (4.5)$$

Such non-linearities include the ReLU ($a_- = 0, a_+ = 1$) and the leaky ReLU ($a_- \in (0, 1), a_+ = 1$). A direct computation, left to the reader, shows that criticality is achieved if and only if

$$K_* \geq 0 \text{ is arbitrary} \quad \text{and} \quad C_b = 0, \quad C_W = \frac{2}{a_+^2 + a_-^2}.$$

Thus, the first property of the ReLU universality class is that setting $(C_b, C_W) = (0, 2/(a_+^2 + a_-^2))$ allows all non-negative K_* to satisfy (*). In fact, at criticality, a simple symmetrization

argument shows that the variance of neuron pre-activations is preserved exactly *even at finite width*

$$\text{Var} [z_{i;\alpha}^{(\ell)}] = \text{Var} [z_{i;\alpha}^{(1)}] = \frac{C_W}{n_0} \|x_\alpha\|^2 \quad \forall \ell, n_0, \dots, n_\ell \geq 1, x_\alpha \in \mathbb{R}^{n_0} \quad (4.6)$$

and, relatedly, that we have

$$\chi_{\parallel;\alpha}^{(\ell)} := \chi_{\parallel}(K_{\alpha\alpha}^{(\ell)}) = 1 = \chi_{\perp;\alpha}^{(\ell)} := \chi_{\perp}(K_{\alpha\alpha}^{(\ell)}) =: \chi_{\perp;\alpha}^{(\ell)}, \quad \forall \ell \geq 1, x_\alpha \in \mathbb{R}^{n_0}.$$

The remarkable property (4.6) is much stronger than the criticality condition (*), which requires only that this condition holds for *some* value of $n_0^{-1} \|x_\alpha\|^2$ and only in the limit when $n \rightarrow \infty$. It implies that the cumulant recursions from Corollary 3.4 for 1-homogeneous non-linearities have constant coefficients and are therefore particularly simple to solve. For instance, we find at leading order in $1/n$

$$\begin{aligned} \kappa_{4;\alpha}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle \sigma(z)^4 \rangle_{K_{\alpha\alpha}^{(\ell)}} - \langle \sigma(z)^2 \rangle_{K_{\alpha\alpha}^{(\ell)}}^2 \right] + \left(\chi_{\parallel;\alpha}^{(\ell)} \right)^2 \kappa_{4;\alpha}^{(\ell)} \\ &= \left(\frac{2}{(a_+^2 + a_-^2)n_0} \|x_\alpha\|^2 \right)^2 \left(6 \frac{a_+^4 + a_-^4}{(a_+^2 + a_-^2)^2} - 1 \right) \sum_{\ell'=1}^{\ell} \frac{1}{n_{\ell'}}, \end{aligned}$$

which shows that while $\kappa_{4;\alpha}^{(\ell)}$ is suppressed by one power of $1/n$ relative to the infinite width variance $K_{\alpha\alpha}^{(\ell)}$, it also grows one order faster in ℓ . This illustrates an important and general theme: depth amplifies finite width effects. It is the effective depth ℓ/n of neurons at layer ℓ that measures the distance to the infinite width Gaussian regime.

Moreover, in the special setting of 1-homogeneous non-linearities there is a simple method for obtaining the full distribution of the pre-activation vector $z_\alpha^{(\ell)}$ at a single input at any finite values of n_0, \dots, n_ℓ . This was first pointed out in [17, 21, 66] and is briefly reviewed in Appendix C. A key takeaway is that if we take the hidden layer widths $n_1 = \dots = n_L = n$, then we have following convergence in distribution to product of independent normal and log-normal random variables:

$$\lim_{\substack{n, L \rightarrow \infty \\ L/n \rightarrow \xi \in [0, \infty)}} z_{i;\alpha}^{(L)} \stackrel{d}{=} \left(\frac{2 \|x_\alpha\|^2}{(a_+^2 + a_-^2)n_0} \right)^{1/2} Z_1 \exp[-\mu(\xi, a_+, a_-) + \sigma(\xi, a_+, a_-)Z_2], \quad (4.7)$$

where

$$\mu(\xi, a_+, a_-) = \sigma^2(\xi, a_+, a_-) := \frac{\xi}{4} \left(6 \frac{a_+^4 + a_-^4}{(a_+^2 + a_-^2)^2} - 1 \right), \quad Z_1, Z_2 \sim \mathcal{N}(0, 1) \text{ iid.}$$

The convergence (4.7) reveals that for a fixed input the distribution of the output of a random with 1-homogeneous non-linearities at large depth and width depends in a simple way on the limiting effective network depth ξ . This bolsters the claim that they are all part of the same universality class. It also means that increasing the network depth L drives it away from the infinite width Gaussian behavior observed at $\xi = 0$ and that the outputs of *deep and wide* networks are not well-approximated by a Gaussian at all, unless ξ is infinitesimal, in which case the log-normal term $\exp[-\mu(\xi, a_+, a_-) + \sigma(\xi, a_+, a_-)Z_2]$ is negligible.

Prior work [20, 21, 23] of the author shows that when $\sigma = \text{ReLU}$ (or any other 1-homogeneous non-linearity), the distribution at large n, L of not only the network output

$z_{i;\alpha}^{(L+1)}$ but also is derivatives with respect to inputs x_α and model parameters (e.g. weights and biases) depends only on the effective depth L/n . We will return to this theme in our discussion of exploding and vanishing gradients (Theorem 4.5). Finally, we note that it has also been observed that log-normal random variables describe the structure of gradients in residual networks, even during/after training [39].

To complete our discussion of the ReLU universality class, we make two final remarks. First, a direct computation (reviewed briefly in Proposition C.1 of Appendix C) shows that at criticality for any non-zero inputs $x_{\alpha_1}, x_{\alpha_2} \in \mathbb{R}^{n_0}$ with the same norm we have

$$\lim_{n \rightarrow \infty} \text{Corr} \left(z_{i;\alpha_1}^{(\ell)}, z_{i;\alpha_2}^{(\ell)} \right) = 1 - \frac{2(a_+ - a_-)^2}{3\pi(a_+^2 + a_-^2)} \ell^{-2} (1 + o(1)). \quad (4.8)$$

The power law exponent 2 that appears in this estimate is common to all 1-homogeneous non-linearities and is another reason to believe they fall into the same universality class. In contrast, this exponent equals one for non-linearities in the $K_* = 0$ universality class presented below. The estimate (4.8) suggests that in order to define a double scaling limit $n, L \rightarrow \infty$ and $L/n \rightarrow \xi$ in which the entire field $x_\alpha \mapsto z_\alpha^{(L+1)}$ is non-degenerate (rather than just its value at a single input) one must rescale distances in the input space to prevent the collapse of correlations otherwise guaranteed in (4.8). We leave this as an interesting direction for future work.

4.2.2 The Universality Class of Hyperbolic Tangent

The second class of non-linearities we study is what [53] termed the $K_* = 0$ universality class, which we take to mean non-linearities σ such that

- σ is a smooth, odd function satisfying Assumption 2.1.
- σ satisfies

$$\sigma_1 \sigma_3 < 0, \quad \sigma_j := \left. \frac{1}{j!} \frac{d^j}{dt^j} \right|_{t=0} \sigma(t). \quad (4.9)$$

- $K_* = 0$ is the unique fixed point of equation (*).
- At criticality, for every non-zero network input $x_\alpha \in \mathbb{R}^{n_0}$ and each $\delta \in (0, 1)$ we have as $L \rightarrow \infty$ that

$$K_{\alpha\alpha}^{(L)} = \frac{1}{aL} \left(1 + O(L^{-1+\delta}) \right), \quad (4.10)$$

where the implicit constant depends on δ and x_α and we've set

$$a := -6 \frac{\sigma_3}{\sigma_1}.$$

This specific value of a , which is positive by (4.9), is the only possible candidate for decay of the form (4.10) that is consistent with the recursion (2.4).

Some remarks are in order. First, if $K_* = 0$ is the unique fixed point for (*), then a simple computation shows that criticality is achieved if and only if

$$K_* = 0, \quad C_b = 0, \quad C_W = \sigma_1^{-2}. \quad (4.11)$$

Next, our definition of the $K_* = 0$ universality class does not make apparent whether it is empty. As we will see in Proposition B.1, however, the $K_* = 0$ universality class is in fact quite large and contains for example the hyperbolic tangent and more generally any non-linearity that is tanh-like in the sense that is smooth with $\sigma_1 \neq 0$, has the opposite sign as its second derivative

$$\text{for almost every } z, \text{sgn}(\sigma(z)\sigma''(z)) = -1,$$

is sub-linear

$$\exists C > 0 \text{ s.t. } \forall z \in \mathbb{R} \quad |\sigma(z)| \leq |\sigma_1 z|,$$

and is controlled by its first few non-zero Taylor series coefficients at 0:

$$\exists C \geq 0 \text{ s.t. } \forall z \geq 0, \quad \sigma_1 z + \sigma_3 z^3 \leq \sigma(z) \leq \sigma_1 z + \sigma_3 z^3 + C z^4.$$

Further, by definition, for the $K_* = 0$ universality class, the infinite width variance $K_{\alpha\alpha}^{(\ell)}$ of neuron pre-activations $z_{i;\alpha}^{(\ell)}$ is qualitatively different from that of 1-homogeneous non-linearities. Indeed, $K_{\alpha\alpha}^{(L)}$ depends on L , decaying polynomially to 0. Moreover, at large L , the value of $K_{\alpha\alpha}^{(L)}$ is independent of the initial condition $K_{\alpha\alpha}^{(0)}$ to leading order in L . As a final remark let us point out that searching for non-linearities σ so that $K_* = 0$ at criticality is quite natural. Indeed, for any σ that is twice differentiable, we have

$$\chi_{\parallel}(K) = \chi_{\perp}(K) + C_W \langle \sigma(z)\sigma''(z) \rangle_K$$

Hence, if $K > 0$, then

$$\chi_{\parallel}(K) = 1, \chi_{\perp}(K) = 1 \quad \implies \quad \langle \sigma(z)\sigma''(z) \rangle_K = 0.$$

But if σ is a sigmoidal function such as tanh, then $\sigma(z)\sigma''(z) < 0$ for all $z \neq 0$. Hence, $\langle \sigma(z)\sigma''(z) \rangle_K = 0$ can only occur when $K = 0$.

As in the monograph [53], let us now probe the role of network depth by studying the large L behavior of the cumulants $\kappa_{2k;\alpha}^{(L)}$, $k = 2, 3, 4$, in networks with non-linearities from the $K_* = 0$ universality class tuned to criticality. Note that in (4.10) the limiting behavior of the variance $K_{\alpha\alpha}^{(L)}$ depends (mildly) on the non-linearity σ in terms of its first few Taylor coefficients at 0. As we are about to see, however, the behavior of the higher cumulants $\kappa_{2k;\alpha}^{(L)}$, $k = 2, 3, 4$, when normalized by the appropriate power of $K_{\alpha\alpha}^{(L)}$, is independent of σ at leading order in n and L and depends only on universal constants and the effective network depth L/n .

Corollary 4.1. *Suppose σ is a non-linearity from the $K_* = 0$ universality class, and consider a depth L random neural network with input dimension n_0 , output dimension n_{L+1} , hidden layer widths satisfying*

$$n_1, \dots, n_L = n \gg 1,$$

and non-linearity σ that has been tuned to criticality as in (4.11). Write $\xi = L/n$ and define the normalized cumulants

$$\widehat{\kappa}_{2k;\alpha}^{(L)} := \frac{\kappa_{2k;\alpha}^{(\ell)}}{(K_{2;\alpha}^{(\ell)})^k}.$$

We have for $k = 2, 3, 4$ that

$$\widehat{\kappa}_{2k;\alpha}^{(L)} = C_{2k} \xi^{k-1} (1 + O(L^{-1})) + O(n^{-k}), \quad (4.12)$$

where

$$C_4 = \frac{2}{3}, \quad C_6 = \frac{28}{15}, \quad C_8 = \frac{8756}{315}$$

are some positive universal constants. The implicit constants in error terms $O(L^{-1})$ depend on σ, C_b, C_W and constants in $O(n^{-j})$, $j = 2, 3, 4$ may depend in addition on L .

Remark 4.2. Combining estimate (4.12) for $k = 2$ with the definition (4.10) of the $K_* = 0$ universality class and (3.12) yields that in the setting of Corollary 3.4 we have to first order in $1/n$ and to leading order in $1/L$ that

$$\widehat{\text{Var}} \left[\left(z_{1;\alpha}^{(\ell+1)} \right)^2 \right] = 2(1 + \xi), \quad \text{Corr} \left(\left(z_{1;\alpha}^{(\ell+1)} \right)^2, \left(z_{2;\alpha}^{(\ell+1)} \right)^2 \right) = \frac{2}{3}\xi, \quad \xi := \frac{L}{n},$$

where $\widehat{\text{Var}}[X] = \text{Var}[X]/\mathbb{E}[X^2]$. Thus, both the fluctuations of a single neuron pre-activation and the correlation between different neurons is controlled to first order in $1/n, 1/L$ by the effective network depth ξ .

We prove Corollary 3.4 in §11.2. The formula (4.12) derived in [61] for $k = 2$ at a physics level of rigor. Moreover, Corollary 3.4 represents a partial analog of the convergence result (4.7), which lends credence to the following

Conjecture 4.3 (Existence of Double Scaling Limit for Random Neural Networks). *Consider a random depth L neural network with input dimension n_0 , hidden layer widths*

$$n_1, \dots, n_L = n \gg 1,$$

output dimension n_{L+1} and non-linearity σ . Suppose further that this network is tuned to criticality in the sense that (4.4) is satisfied. Fix a non-zero network input $x_\alpha \in \mathbb{R}^{n_0}$ and write $\xi = L/n$. For each $k \geq 1$ there exists $C_{2k} > 0$ depending on the universality class of σ so that

$$\frac{\kappa_{2k;\alpha}^{(L)}}{\left(K_{2;\alpha}^{(L)} \right)^k} = C_{2k} \xi^{k-1} + O\left(\xi^k\right).$$

Moreover, for each $\xi \in [0, \infty)$ there exists a probability distribution $\mathbb{P}_{\xi, \sigma}$ on \mathbb{R} , depending only on ξ and σ , such that in the double scaling limit

$$n, L \rightarrow \infty, \quad \frac{L}{n} \rightarrow \xi,$$

the random variable $z_{i;\alpha}^{(\ell)}$ converges in distribution to a random variable with law $\mathbb{P}_{\xi, \sigma}$.

4.3 Gradients in Random Neural Networks

We presented in §3 and §4 a range of results about random fully connected neural networks. The high-level takeaway was that such networks are succinctly described when n, L are large by the effective network depth $\xi = L/n$. We saw that when $\xi = 0$ the network outputs are a Gaussian process (Theorem 2.2). When $\xi > 0$, however, a different picture emerges. Namely, the higher order cumulants of the network output grow like powers of ξ (see Corollary 3.4 and

Conjecture 4.3). In this section we continue to consider a random neural network $x_\alpha \mapsto z_\alpha^{(L+1)}$ and study for any ℓ the partial derivatives

$$\partial_{x_{j;\alpha}} z_{i;\alpha}^{(\ell)}, \quad j = 0, \dots, n_0,$$

where by definition when $j = 0$ we set

$$\partial_{x_{0;\alpha}} z_{i;\alpha}^{(\ell)} := z_{i;\alpha}^{(\ell)}.$$

In §10.2 - §10.4 we obtain recursions with respect to ℓ for both the infinite width covariances

$$K_{(ij)}^{(\ell)} := \lim_{n \rightarrow \infty} \text{Cov} \left(\frac{\partial z_{i;\alpha}^{(\ell+1)}}{\partial x_{i;\alpha}}, \frac{\partial z_{j;\alpha}^{(\ell+1)}}{\partial x_{j;\alpha}} \right), \quad i, j = 1, 2,$$

and the fourth cumulants

$$\begin{aligned} \kappa_{(ii)(jj)}^{(\ell)} &:= \left(1 - \frac{2}{3} \delta_{ij}\right) \kappa \left(\frac{\partial z_{i;\alpha}^{(\ell+1)}}{\partial x_{i;\alpha}}, \frac{\partial z_{i;\alpha}^{(\ell+1)}}{\partial x_{i;\alpha}}, \frac{\partial z_{j;\alpha}^{(\ell+1)}}{\partial x_{j;\alpha}}, \frac{\partial z_{j;\alpha}^{(\ell+1)}}{\partial x_{j;\alpha}} \right), \quad i, j = 0, 1, 2. \\ &= \kappa \left(\frac{\partial z_{1;\alpha}^{(\ell+1)}}{\partial x_{i;\alpha}}, \frac{\partial z_{1;\alpha}^{(\ell+1)}}{\partial x_{i;\alpha}}, \frac{\partial z_{2;\alpha}^{(\ell+1)}}{\partial x_{j;\alpha}}, \frac{\partial z_{2;\alpha}^{(\ell+1)}}{\partial x_{j;\alpha}} \right). \end{aligned}$$

These recursions are valid for any non-linearity and any choice of C_b, C_W . Then, in §11.2 we solve these recursions in the context of a random neural network with a non-linearity from the $K_* = 0$ universality class tuned to criticality. We record several of these solutions in the following result.

Theorem 4.4 (Second and Fourth Joint Cumulants for Values and Derivatives of Random Neural Networks). *Fix $r \geq 1$ and suppose that $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ satisfies Assumption 2.1 with this value of r . Fix also an $x_\alpha \neq 0 \in \mathbb{R}^{n_0}$. Consider a random neural network $z_\alpha^{(L+1)}$ with input dimension n_0 , output dimension n_{L+1} , hidden layer widths $n_1, \dots, n_L = n$, and non-linearity σ belonging to the $K_* = 0$ universality class that has been tuned to criticality as in (4.11). We have for any $\delta \in (0, 1)$*

$$K_{(00)}^{(\ell+1)} = \frac{1}{a\ell} + O_\delta(\ell^{-2+\delta}) \tag{4.13}$$

$$K_{(10)}^{(\ell+1)} = \frac{C_W e^{-2\gamma} x_{1;\alpha}}{n_0 \ell^2} (1 + O(\ell^{-1})) \tag{4.14}$$

$$K_{(11)}^{(\ell+1)} = \frac{C_W e^{-\gamma}}{n_0 \ell} (1 + O(\ell^{-1})) + O\left(\frac{(x_{1;\alpha})^2}{n_0^2 \ell^3}\right) \tag{4.15}$$

$$K_{(12)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha} x_{2;\alpha}}{n_0^2 \ell^3}\right), \tag{4.16}$$

where γ is the Euler-Mascheroni constant and the implicit constant in the $O_\delta(\ell^{-2+\delta})$ error

depends on δ . Moreover, denoting by $\xi = L/n$ the effective network depth, we have

$$\frac{\kappa_{(11)(00)}^{(\ell+1)}}{K_{(11)}^{(\ell)} K_{(00)}^{(\ell)}} = -\frac{1}{3}\xi(1 + O(\ell^{-1})) + O(n^{-2}) \quad (4.17)$$

$$\frac{\kappa_{(11)(11)}^{(\ell+1)}}{(K_{(11)}^{(\ell)})^2} = \frac{8}{3}\xi(1 + O(\ell^{-1})) + O(n^{-2}) \quad (4.18)$$

$$\frac{\kappa_{(11)(22)}^{(\ell+1)}}{K_{(11)}^{(\ell)} K_{(22)}^{(\ell)}} = \frac{2}{3}\xi(1 + O(\ell^{-1})) + O(n^{-2}), \quad (4.19)$$

where the implicit constants in the $O(\ell^{-1})$ error terms depend only on x_α, σ and the implicit constants in the $O(n^{-2})$ error terms may depend in addition on ℓ .

Theorem 4.4 reveals several interesting phenomena. First, we find from (4.13) and (4.14) that at infinite width values and derivatives become uncorrelated at large L :

$$\lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} \text{Corr} \left(z_{1;\alpha}^{(L+1)}, \partial_{x_{1;\alpha}} z_{1;\alpha}^{(L+1)} \right) = 0.$$

Moreover, at finite width, the effective network depth ξ controls the distribution of gradients both in terms of their fluctuations at a single neuron and their correlations between neurons. For instance, to leading order in n, L :

$$\frac{\text{Var} \left[\left(\partial_{x_{1;\alpha}} z_{1;\alpha}^{(L+1)} \right)^2 \right]}{\mathbb{E} \left[\left(\partial_{x_{1;\alpha}} z_{1;\alpha}^{(L+1)} \right)^2 \right]} = \frac{8}{3}\xi, \quad \text{Corr} \left(\left(\frac{\partial z_{1;\alpha}^{(L+1)}}{\partial x_{1;\alpha}} \right)^2, \left(\frac{\partial z_{1;\alpha}^{(L+1)}}{\partial x_{2;\alpha}} \right)^2 \right) = \frac{2}{3}\xi.$$

In the next section, we will apply such estimates to understand a simple version of the so-called exploding and vanishing gradient problem.

4.4 Exploding and Vanishing Gradient Problem for First Layer Weights

On its face, Theorem 4.4 concerns only derivatives of $z_{i;\alpha}^{(L+1)}$ with respect to the network input x_α . However, it is the derivatives of $z_{i;\alpha}^{(L+1)}$ with respect to the model parameters (weights and biases) that are arguably more important. To explain this more precisely, consider a fully connected neural network $x_\alpha \mapsto z_\alpha^{(L+1)}$ with input dimension n_0 , output dimension n_{L+1} , and hidden layer widths n_1, \dots, n_L . For any $m \in \mathbb{N}$ let us set

$$[m] := \{1, \dots, m\}$$

and denote by

$$\theta = (\theta_\mu, \mu \in [\#\text{params}]) = \left(W_{ij}^{(\ell)}, b_i^{(\ell)}, \ell \in [L+1], i \in [n_\ell], j \in [n_{\ell-1}] \right)$$

the vector of its trainable parameters, where we've set

$$\#\text{params} := \#\text{weights and biases} = \sum_{\ell=1}^{L+1} n_\ell (1 + n_{\ell-1}).$$

As we briefly discussed in §2.2 the typical use case for a network is to optimize its parameters θ by some variant of gradient descent

$$\theta_\mu(t+1) = \theta_\mu(t) - \eta \partial_\mu \mathcal{L}(\theta(t)), \quad \mu \in [\#\text{params}] \quad (4.20)$$

starting from a random initialization $\theta(0)$ drawn from the distribution (2.2). The goal of this procedure is to minimize a loss

$$\mathcal{L}(\theta) = \mathcal{L}(z_{\mathcal{A}_{\text{train}}}^{(L+1)}(\theta)), \quad z_{\mathcal{A}_{\text{train}}}^{(L+1)} := \left(z_\alpha^{(L+1)}, \alpha \in \mathcal{A}_{\text{train}} \right),$$

which often depends only on the network output and measures for each θ how well the resulting neural net function $x_\alpha \mapsto z_\alpha^{(L+1)}(\theta)$ matches given training data $\{(x_\alpha, f(x_\alpha)), \alpha \in \mathcal{A}_{\text{train}}\}$ for a function $f: \mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_{L+1}}$. For example, we might have

$$\mathcal{L}(\theta) = \frac{1}{|\mathcal{A}_{\text{train}}|} \sum_{\alpha \in \mathcal{A}_{\text{train}}} \left\| f(x_\alpha) - z_\alpha^{(L+1)} \right\|_2^2,$$

though many different losses are used in practice. The parameter η is called the learning rate, or step size. An important and well-known [5, 17, 27] numerical stability issue that comes up in the course of using a first order method such as (4.20) to optimize the parameters of a deep neural network is called the exploding and vanishing gradient problem (EVGP). Informally, the EVGP occurs when the update (4.20) is numerically unstable:

$$\text{EVGP} \quad \longleftrightarrow \quad \text{for many parameters } \theta_\mu \quad \frac{|\eta \partial_\mu \mathcal{L}(\theta)|}{|\theta_\mu|} \approx 0 \text{ or } \infty.$$

The presence of the EVGP causes optimization to break down since if $|\partial_\mu \mathcal{L}(\theta)| / |\theta_\mu| \approx 0$ or ∞ , then the relative change in the parameter θ_μ is either too small to be useful or so large that it amounts to a large random step in the space of parameters and therefore is unlikely to decrease the loss.

It has long been known that sufficiently deep fully connected neural networks are prone to suffer from the EVGP [5, 27, 28]. Indeed, many important practical innovations such as residual connections [24] were invented in part to address such numerical issues, though by now they are seen as useful for many other reasons. To see why deep enough neural nets tend to have numerically unstable gradients note that since the loss only depends on the network outputs on a finite training set $z_{\mathcal{A}_{\text{train}}}^{(L+1)}$, we may use the chain rule to write

$$\partial_\mu \mathcal{L}(\theta) = \sum_{\alpha \in \mathcal{A}_{\text{train}}} J_{\theta_\mu} z_\alpha^{(L+1)} J_{z_\alpha^{(L+1)}} \mathcal{L}(z_{\mathcal{A}_{\text{train}}}^{(L+1)}),$$

where for any function f we denote by $J_x f(x)$ its Jacobian with respect to x . Moreover, suppose that θ_μ is a weight or bias in layer ℓ_0 . Then, again by the chain rule, the Jacobian of the network output with respect to μ equals

$$J_{\theta_\mu} z_\alpha^{(L+1)} = J_{\theta_\mu} z_\alpha^{(\ell_0)} J_{z_\alpha^{(\ell_0)}} z_\alpha^{(L+1)}.$$

Both of these terms may have large fluctuations. Perhaps most importantly, the second term can be rewritten as a product

$$J_{z_\alpha^{(\ell_0)}} z_\alpha^{(L+1)} = \prod_{\ell=\ell_0}^L J_{z_\alpha^{(\ell)}} z_\alpha^{(\ell+1)},$$

of $L - \ell_0 + 1$ layer-to-layer Jacobian matrices (which are random at the start of training), with each term having the following explicit representation:

$$J_{z_\alpha^{(\ell)}} z_\alpha^{(\ell+1)} = D^{(\ell+1)} W^{(\ell+1)}, \quad D^{(\ell+1)} := \text{Diag} \left(\sigma' \left(z_{i;\alpha}^{(\ell+1)} \right), i = 1, \dots, n_{\ell+1} \right).$$

Except in the simple case when σ is the identity and we require $W^{(\ell+1)}$ to be orthogonal, the top singular value of each layer-to-layer Jacobian will tend to deviate from 1, causing the norms of their products to either grow or decay exponentially with L . This poor conditioning for large matrix products is a key culprit behind the EVGP.

Often, though not always, the gradients $\partial_\mu \mathcal{L}$ have the largest fluctuations at the start of training. Thus, the EVGP is primarily a property of random neural networks, and the techniques in this article can be used to shed light on when it occurs. At the start of training the entries of the Jacobian $J_{z_\alpha^{(\ell_0)}} z_\alpha^{(L+1)}$ are precisely the derivatives of a random neural network with depth $L - \ell_0 + 1$ with respect to its input. One can therefore hope to use Theorems 3.1 and 3.2 to address the following more precise formulation of the EVGP:

EVGP \longleftrightarrow for a typical random draw of weights and biases the size of the relative Jacobians of the network output

$$|J_{\theta_\mu} z_{j;\alpha}^{(L+1)}| / |\theta_\mu|$$

has a large variance over μ (i.e. over network parameters).

This characterization of the EVGP says in essence that if we can choose only a single learning rate η for a group of otherwise indistinguishable parameters, such as the weights in a single layer, then it make sense to set

$$\eta = \left(\text{average relative gradient } |J_{\theta_\mu} z_{j;\alpha}^{(L+1)}| / |\theta_\mu| \right)^{-1}.$$

The EVGP will then occur when the size of relative gradients have a large variance across parameters in a typical random draw of weights and biases, so that for some parameters our setting of η is too small whereas for others it is too large.

We will leave the study of the EVGP in this general form as pertaining to future work and consider here a special case. Specifically, we compute in Theorem 4.5 below the average of the empirical variance of the squared gradients $|J_{\theta_\mu} z_{j;\alpha}^{(L+1)}|^2$ when $\theta_\mu = W_{ij}^{(1)}$ varies over all weights in the first layer. By construction, the typical size of $|\theta_\mu|$ is the same for all such μ . The EVGP will thus occur for weights in the first layer if and only if the empirical fluctuations over μ of the raw Jacobians $|J_{W_{ij}^{(1)}} z_{j;\alpha}^{(L+1)}|$ is large. To state the precise result we fix $x_\alpha \neq 0 \in \mathbb{R}^{n_0}$ and $q \in [n_{L+1}]$. We then write

$$\text{Grad Mean}^{(1)} := \frac{1}{n_0 n_1} \sum_{\substack{1 \leq i \leq n_1 \\ 1 \leq j \leq n_0}} \left(\partial_{W_{ij}^{(1)}} z_{q;\alpha}^{(L+1)} \right)^2$$

for the empirical mean of the squared gradients of $z_{q;\alpha}^{(L+1)}$ with respect to the weights $W_{ij}^{(1)}$ in the first layer and analogously

$$\text{Grad Var}^{(1)} := \frac{1}{n_0 n_1} \sum_{\substack{1 \leq i \leq n_1 \\ 1 \leq j \leq n_0}} \left(\partial_{W_{ij}^{(1)}} z_{q;\alpha}^{(L+1)} \right)^4 - \left(\text{Grad Mean}^{(1)} \right)^2$$

for the empirical variance of the squared gradients.

Theorem 4.5 (EVGP for First Layer Weights in Fully Connected Networks). *Let $x_\alpha \mapsto z_\alpha^{(L+1)}$ be a random neural network with input dimension n_0 , output dimension n_{L+1} , hidden layer widths*

$$n_1, \dots, n_L = n \gg 1,$$

and non-linearity σ . Suppose also that

- σ satisfies Assumption. 2.1
- σ belongs to the $K_* = 0$ universality class in the sense described in §4.1 as well as §4.2.2.
- C_b, C_W are tuned to criticality in the sense described as in (4.11).

Then, for any non-zero $x_\alpha \in \mathbb{R}^{n_0}$ we have

$$\frac{\mathbb{E} \left[\text{Grad Var}^{(1)} \right]}{\mathbb{E} \left[\text{Grad Mean}^{(1)} \right]^2} = C_{\sigma, \alpha} \left(1 + \frac{8}{3} \xi + O(L^{-1}) + O(n^{-1}) + O_L(n^{-2}) \right),$$

where $\xi = \frac{L}{n}$ is the effective network depth and

$$C_{\sigma, \alpha} := 3 \frac{\langle (\sigma')^4 \rangle_{K_{\alpha\alpha}^{(1)}}}{\langle (\sigma')^2 \rangle_{K_{\alpha\alpha}^{(1)}}^2} \frac{n_0^{-1} \|x_\alpha\|_4^4}{(n_0^{-1} \|x_\alpha\|_2^2)^2} - 1$$

is a strictly positive constant depending only on x_α, σ , the implicit constants in the error terms $O(n^{-1}), O(L^{-1})$ depend only on x_α, σ whereas the implicit constant in the error term $O_L(n^{-2})$ depends on x_α, σ, L .

Remark 4.6. *Theorem 4.5 shows that, at least to leading order in $1/n$, the exploding and vanishing gradient problem occurs for first layer weights in a critically tuned random fully connected with non-linearities from the $K_* = 0$ universality class if and only if the effective network depth L/n is large.*

We prove Theorem 4.5 in §11. To put this result into the context, let us make several remarks. First, the analog of Theorem 4.5 for 1-homogeneous non-linearities was derived in prior work [17, 21]. As we've alluded to before, the behavior of a random fully connected neural network with such non-linearities can be studied by more specialized methods that reveal not only the leading order dependence in $1/n$ but actually the full dependence on L/n plus errors of size L/n^2 (see Appendix C). The conclusion, just as for non-linearities from the $K_* = 0$ universality class, is that the EVGP occurs if and only if L/n is large.

Second, Theorem 4.5 is the first time the EVGP has been characterized mathematically in random fully connected networks for a broad class of non-linearities beyond those that are 1-homogeneous. However, as throughout, the author would like to acknowledge the prior work [53] of Roberts, Yaida, and the author which also studies the EVGP and shows, at a physics level of rigor, that the variance *over random initialization* of a single squared gradient $(\partial_\mu z_{q;\alpha}^{(L+1)})^2$ also scales like L/n at large L . This is conceptually different from showing that in a typical initialization the empirical variance of $(\partial_\mu z_{q;\alpha}^{(L+1)})^2$ over μ in the first layer is large, which is what Theorem 4.5 establishes. Nonetheless, the conclusions are the essentially the same and the basic techniques used to derive the results are similar.

5 Overview of Proofs

In this section, we present the essential idea for how we analyze a random fully connected neural network $x_\alpha \mapsto z_\alpha^{(L+1)}$ at finite width. Our approach is based on the following structural properties:

- The sequence of fields $z_\alpha^{(\ell)}$ is a Markov Chain with respect to ℓ .
- Conditional on the sigma algebra $\mathcal{F}^{(\ell)}$ defined by $z_\alpha^{(\ell)}$ is a Gaussian field with independent components $z_{i;\alpha}^{(\ell+1)}$. See Lemma 7.1.
- The variance of each component $z_{i;\alpha}^{(\ell+1)}$ depends on $z_\alpha^{(\ell)}$ only through random variables of the form

$$\mathcal{O}_f^{(\ell)} := \frac{1}{n_\ell} \sum_{j=1}^{n_\ell} f(z_{j;\alpha}^{(\ell)}),$$

which we refer to as collective observables. See (3.6).

- Collective observables are self-averaging to a similar extent as if the random variables $f(z_{i;\alpha}^{(\ell)})$ were independent in the sense that for any $q \geq 0$, we have

$$\mathbb{E} \left[\left(\mathcal{O}_f^{(\ell)} - \mathbb{E} \left[\mathcal{O}_f^{(\ell)} \right] \right)^q \right] = O_q \left(n^{-\lceil \frac{q}{2} \rceil} \right). \quad (5.1)$$

See Theorem 7.3 and Lemma 7.5.

Let us briefly explain, mostly dispensing with rigor, how these four ideas come together to obtain a recursive description of the distribution of the field $z_\alpha^{(\ell+1)}$ in terms of that of $z_\alpha^{(\ell)}$. To keep the notation to a minimum, we fix a network input x_α and focus on describing the joint distribution of $z_{i;\alpha}^{(\ell+1)}$, $i = 1, \dots, m$. Extensions to multiple inputs and derivatives proceed along very similar lines. Denoting by $\xi = (\xi_1, \dots, \xi_m)$ dual variables, consider the characteristic function

$$p^{(\ell+1)}(\xi) := \mathbb{E} \left[\exp \left[-i \sum_{i=1}^m \xi_i z_{i;\alpha}^{(\ell+1)} \right] \right].$$

Conditioning on $z_\alpha^{(\ell)}$ and using (3.6) allows us to write

$$p^{(\ell+1)}(\xi) := \mathbb{E} \left[\exp \left[-\frac{1}{2} \|\xi\|^2 \Sigma_{\alpha\alpha}^{(\ell)} \right] \right],$$

where we remind the reader that

$$\Sigma_{\alpha\alpha}^{(\ell)} = \text{Var} \left[z_{i;\alpha}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right] = C_b + \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \sigma(z_{j;\alpha}^{(\ell)})^2$$

is a collective observable *at the previous layer*. Writing

$$\kappa_{\alpha\alpha}^{(\ell)} := \mathbb{E} \left[\Sigma_{\alpha\alpha}^{(\ell)} \right], \quad \Delta_{\alpha\alpha}^{(\ell)} := \Sigma_{\alpha\alpha}^{(\ell)} - \mathbb{E} \left[\Sigma_{\alpha\alpha}^{(\ell)} \right],$$

we find

$$p^{(\ell+1)}(\xi) := \mathbb{E} \left[\exp \left[-\frac{1}{2} \|\xi\|^2 \Delta_{\alpha\alpha}^{(\ell)} \right] \right] \exp \left[-\frac{1}{2} \|\xi\|^2 \kappa_{\alpha\alpha}^{(\ell)} \right].$$

The second term is precisely the characteristic function of a centered m -dimensional Gaussian with iid components of variance $\kappa_{\alpha\alpha}^{(\ell)}$. Moreover, at least heuristically, the first term can be written

$$\mathbb{E} \left[\exp \left[-\frac{1}{2} \|\xi\|^2 \Delta_{\alpha\alpha}^{(\ell)} \right] \right] = \sum_{q \geq 0} \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^q \frac{(-1)^q}{2^q q!} \|\xi\|^{2q} \right].$$

The concentration estimates (5.1) ensure that this series converges. Moreover, since the Fourier transform turns polynomials into derivatives we have

$$-\|\xi\|^2 = \text{Laplacian in the variables } z_{i;\alpha}^{(\ell+1)}.$$

Hence, we obtain for any reasonable test function f that

$$\mathbb{E} \left[f(z_{i;\alpha}^{(\ell+1)}, i = 1, \dots, m) \right] = \sum_{q=0}^{\infty} \frac{1}{2^q q!} \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^q \left\langle \left(\sum_{i=1}^m \partial_{z_{i;\alpha}}^2 \right)^q f(z_{i;\alpha}, i = 1, \dots, m) \right\rangle_{\kappa_{\alpha\alpha}^{(\ell)}} \right],$$

where $(z_{i;\alpha}, i = 1, \dots, m)$ is a vector of iid centered Gaussians with variance $\kappa_{\alpha\alpha}^{(\ell)}$. The concentration estimates (5.1) ensure that this expression is a power series in $1/n$. In particular,

$$\begin{aligned} \mathbb{E} \left[f(z_{i;\alpha}^{(\ell+1)}, i = 1, \dots, m) \right] &= \langle f(z_{i;\alpha}, i = 1, \dots, m) \rangle_{\kappa_{\alpha\alpha}^{(\ell)}} \\ &+ \frac{\mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right]}{8} \left\langle \left(\sum_{i=1}^m \partial_{z_{i;\alpha}}^2 \right)^2 f(z_{i;\alpha}, i = 1, \dots, m) \right\rangle_{\kappa_{\alpha\alpha}^{(\ell)}} + O(n^{-2}). \end{aligned} \quad (5.2)$$

This is the essence of Theorem 3.2. To derive usable recursions for cumulants of $z_{i;\alpha}^{(\ell+1)}$, note for instance that, in the notation of Corollary 3.4,

$$\kappa_{4;\alpha}^{(\ell)} := \frac{1}{3} \kappa \left(z_{i;\alpha}^{(\ell+1)}, z_{i;\alpha}^{(\ell+1)}, z_{i;\alpha}^{(\ell+1)}, z_{i;\alpha}^{(\ell+1)} \right) = \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right].$$

Writing

$$X_j := \sigma(z_{j;\alpha}^{(\ell+1)})^2 - \mathbb{E} \left[\sigma(z_{j;\alpha}^{(\ell+1)})^2 \right]$$

we thus have

$$\kappa_{\alpha\alpha}^{(\ell+1)} = \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right] = \frac{C_W^2}{n_\ell} \mathbb{E} [X_1^2] + C_W^2 (1 - n_\ell^{-1}) \mathbb{E} [X_1 X_2].$$

Applying the expansion (5.2) to both these terms and a bit of algebra already yields

$$\begin{aligned} \kappa_{\alpha\alpha}^{(\ell+1)} &= \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right] \\ &= \frac{C_W^2}{n_\ell} \left(\langle \sigma^4 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}} - \langle \sigma^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}}^2 \right) \\ &+ C_W^2 (1 - n_\ell^{-1}) \left(\left(\langle \sigma^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}} - \mathbb{E} \left[\sigma(z_{1;\alpha}^{(\ell)})^2 \right] \right)^2 + \frac{1}{4} \langle \partial^2 \sigma^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}}^2 \kappa_{4;\alpha}^{(\ell)} \right) + O(n^{-2}). \end{aligned}$$

A short argument supplied in §9 shows that $\langle \sigma^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}} = \mathbb{E} \left[\sigma(z_{1;\alpha}^{(\ell)})^2 \right] + O(n^{-1})$ and that we may replace $\kappa_{\alpha\alpha}^{(\ell)}$ by its infinite width limit $K_{\alpha\alpha}^{(\ell)}$ in all remaining expectations at the cost of an $O(n^{-1})$ error. This already yields the recursion (3.15) of Corollary 3.4.

6 Background

6.1 Properties of Cumulants

Recall that, given random variables X_1, \dots, X_k on the same probability space, we denote their mixed cumulant by

$$\kappa(X_1, \dots, X_k) := i^k \frac{\partial^k}{\partial t_1 \cdots \partial t_k} \Big|_{t=0} \log \mathbb{E} [\exp [-i(t_1 X_1 + \cdots + t_k X_k)]] .$$

In the following result, we recall the key properties of these mixed cumulants that we will need.

Proposition 6.1 (See Theorem 2.3.1 in [7]). *Mixed cumulants satisfy the following properties.*

1. *Suppose $X = (X_1, \dots, X_k)$ is a random vector with finite moments of all orders. Then*

$$\kappa(X_1, \dots, X_k) = \sum_{\pi=(\pi_1, \dots, \pi_b)} \kappa(\kappa(X_{\pi_1} | \mathcal{F}), \dots, \kappa(X_{\pi_b} | \mathcal{F})), \quad (6.1)$$

where the sum is over all partitions π of $[k]$ and for each $a = 1, \dots, b$

$$X_{\pi_a} := (X_i, i \in \pi_a).$$

This is known as the law of total cumulance. See [6].

2. *Suppose $X = (X_1, \dots, X_k)$ is a random vector with finite moments of all orders. When X can be partitioned into two independent subsets, the mixed cumulant $\kappa(X)$ vanishes. More precisely, suppose $I \subseteq [k]$ and that $I, I^c \neq \emptyset$. Then*

$$X_I := (X_i, i \in I) \perp X_{I^c} = (X_i, i \notin I) \implies \kappa(X_1, \dots, X_k) = 0. \quad (6.2)$$

3. *Mixed cumulants are multi-linear. More precisely if*

$$\{X_{i,j}, 1 \leq j \leq k, i \leq T_j\}$$

are random variables with finite moments defined on the same probability space, then

$$\kappa \left(\sum_{i_1=1}^{T_1} a_{i_1,1} X_{i_1,1}, \dots, \sum_{i_k=1}^{T_k} a_{i_k,k} X_{i_k,k} \right) = \sum_{i_1=1}^{T_1} \cdots \sum_{i_k=1}^{T_k} a_{i_1,1} \cdots a_{i_k,k} \kappa(X_{i_1,1}, \dots, X_{i_k,k}) \quad (6.3)$$

for any $a_{i,j} \in \mathbb{R}$.

4. *Suppose $X = (X_1, \dots, X_j) \sim \mathcal{N}(0, \Sigma)$ is a centered Gaussian with covariance Σ . Then for any $i_1, \dots, i_k \in [j]$*

$$\kappa(X_{i_1}, \dots, X_{i_k}) = \begin{cases} \Sigma_{i_1 i_2}, & k = 2 \\ 0, & \text{otherwise} \end{cases}. \quad (6.4)$$

5. *Moments are polynomials in cumulants. Specifically, suppose $X = (X_1, \dots, X_k)$ is a random vector with finite moments of all orders. Then,*

$$\mathbb{E}[X_1 \cdots X_k] = \sum_{\pi=(\pi_1, \dots, \pi_b)} \prod_{a=1}^b \kappa(X_{\pi_a}), \quad (6.5)$$

where the sum is over all partitions of $[k]$ and for each $a \in [b]$ we've written

$$X_{\pi_a} := (X_i, i \in \pi_a).$$

6. *Cumulants are polynomials in moments. Specifically,*

$$\kappa(X_1, \dots, X_k) = \sum_{\pi=(\pi_1, \dots, \pi_b)} (-1)^{b-1} (b-1)! \prod_{a=1}^b \mathbb{E} \left[\prod_{i \in \pi_a} X_i \right], \quad (6.6)$$

where the sum is over all partitions of $[k]$ and for each $a \in [b]$ we've written

$$X_{\pi_a} := (X_i, i \in \pi_a).$$

6.2 Lemmata

In this section, we collect two simple auxiliary results that we will need. The first is a Lemma for Solving Certain Recursions.

Lemma 6.2. *Fix $C_1, C_2, \psi > 0$ satisfying*

$$C_2 \geq 1, \quad \psi \neq C_2 + 1$$

as well as $*$ $\in \{\leq, \geq\}$. Suppose also that for each $\ell \geq 0$ we have

$$a_{\ell+1} * \xi_{\ell} + (1 - \zeta_{\ell})a_{\ell}, \quad \zeta_{\ell} \in [0, 1] \quad (6.7)$$

with $a_0 \in \mathbb{R}$ given and that there exist $C'_1, C'_2 > 0$ so that

$$\left| \xi_{\ell} - C_1 \ell^{-\psi} \right| \leq C'_1 \ell^{-1-\psi}, \quad \left| \zeta_{\ell} - C_2 \ell^{-1} \right| \leq C'_2 \ell^{-2}.$$

Then

$$a_{\ell+1} * \frac{\ell^{1-\psi}}{1 - \psi + C_2} (1 + O(\ell^{-1})) + e^{-C_2 \gamma} \ell^{-C_2} a_0 (1 + O(\ell^{-1})) \quad (6.8)$$

where γ is the Euler-Mascheroni constant and the implied constants depend only C_1, C_2, C'_1, C'_2 .

Proof. By unfolding the recursion (6.7) we find

$$a_{\ell+1} * \sum_{\ell'=1}^{\ell} \xi_{\ell'} \prod_{\ell''=\ell'+1}^{\ell} (1 - \zeta_{\ell''}) + a_0 \prod_{\ell''=0}^{\ell} (1 - \zeta_{\ell''}).$$

We have

$$\begin{aligned}
\prod_{\ell''=1}^{\ell} (1 - \zeta_{\ell''}) &= \exp \left[\sum_{\ell''=1}^{\ell} \log (1 - C_2(\ell'')^{-1} + O(\ell^{-2})) \right] \\
&= \exp \left[O(\ell^{-1}) + \sum_{\ell''=1}^{\ell} -C_2(\ell'')^{-1} \right] \\
&= \exp [O(\ell^{-1}) - C_2 \log(\ell) - C_2 \gamma] \\
&= e^{-C_2 \gamma} \ell^{-C_2} (1 + O(\ell^{-1})).
\end{aligned}$$

This gives the second term in (6.8). For the first term, we write

$$\begin{aligned}
\sum_{\ell'=1}^{\ell} \xi_{\ell'} \prod_{\ell''=\ell'+1}^{\ell} (1 - \zeta_{\ell}) &= \sum_{\ell'=1}^{\ell} \xi_{\ell'} \exp \left[\sum_{\ell''=\ell'+1}^{\ell} \log (1 - \zeta_{\ell}) \right] \\
&= \sum_{\ell'=1}^{\ell} \xi_{\ell'} \exp \left[\sum_{\ell''=\ell'+1}^{\ell} -C_2(\ell'')^{-1} + O((\ell'')^{-2}) \right] \\
&= \sum_{\ell'=1}^{\ell} C_1(\ell')^{-\psi} (1 + O(\ell')^{-1}) \exp \left[-C_2 \log \left(\frac{\ell}{\ell'} \right) + O((\ell')^{-1}) \right] \\
&= \ell^{-C_2} \sum_{\ell'=1}^{\ell} C_1(\ell')^{-\psi+C_2} (1 + O(\ell')^{-1}) \\
&= \frac{C_1}{1 + C_2 - \psi} \ell^{1-\psi} (1 + O(\ell^{-1})).
\end{aligned}$$

This completes the proof of (6.8). \square

The second result we will need is the following simple Lemma about Gaussian Integrals.

Lemma 6.3. *Fix $r \geq 1$, a $r \times r$ matrix Σ and measurable function $g : \mathbb{R}^r \rightarrow \mathbb{R}$ that is polynomially bounded:*

$$\exists k \geq 1 \text{ s.t. } \sup_{x \in \mathbb{R}^k} \left| (1 + \|x\|)^{-k} g(x) \right| < \infty.$$

If X is a standard Gaussian random vector in \mathbb{R}^r , then the function

$$\Sigma \mapsto \mathbb{E} \left[g \left(\Sigma^{1/2} X \right) \right] \tag{6.9}$$

is smooth on the open set of strictly positive definite $k \times k$ matrices. Further, if g is a smooth function and each of its derivatives is polynomially bounded, then the map (6.9) is extends to a smooth function on the closed set of positive semi-definite matrices and, in particular,

$$\frac{\partial}{\partial \Sigma_{ij}} \mathbb{E} \left[g \left(\Sigma^{1/2} X \right) \right] = \mathbb{E} \left[(\partial_i \partial_j g)(\Sigma^{1/2} X) \right]. \tag{6.10}$$

Proof. On the open set of strictly positive definite matrices, the Gaussian density

$$\Sigma \mapsto \exp \left[-\frac{1}{2} x^T \Sigma^{-1} x - \frac{1}{2} \log \det(2\pi \Sigma) \right]$$

is a smooth function of Σ with derivatives that are polynomials in x and the entries of Σ, Σ^{-1} . The assumption that f is polynomially bounded shows that we may differentiate under the integral sign and see that that

$$\mathbb{E} \left[g(\Sigma^{1/2} X) \right] = \int_{\mathbb{R}^r} g(x) \exp \left[-\frac{1}{2} x^T \Sigma^{-1} x - \frac{1}{2} \log \det(2\pi \Sigma) \right] dx$$

is indeed a smooth function of Σ . Suppose instead that g is a smooth function and that its derivatives are all polynomially bounded. Suppose first that g is in fact a Schwartz function. Then, writing \hat{g} for its Fourier transform we have

$$\mathbb{E} \left[g(\Sigma^{1/2} X) \right] = \int_{\mathbb{R}^r} \hat{g}(\xi) \exp \left[-\frac{1}{2} \xi^T \Sigma \xi \right] d\xi.$$

Since \hat{g} is also Schwartz, we may differentiate under the integral sign to obtain

$$\frac{\partial}{\partial \Sigma_{ij}} \mathbb{E} \left[g(\Sigma^{1/2} X) \right] = - \int_{\mathbb{R}^r} \xi_i \xi_j \hat{g}(\xi) \exp \left[-\frac{1}{2} \xi^T \Sigma \xi \right] d\xi = \mathbb{E} \left[\partial_{x_i} \partial_{x_j} \Big|_{x=\Sigma^{1/2} X} g(x) \right]. \quad (6.11)$$

Finally, if g is not Schwartz but is smooth with all derivatives being polynomially bounded, we consider the convolution

$$g_\epsilon(x) := (g * \psi_\epsilon)(x), \quad \psi_\epsilon(y) = \exp \left[-\frac{\|y\|^2}{2\epsilon} - \frac{1}{2} \log(2\pi\epsilon) \right].$$

Then, g_ϵ is Schwartz for all $\epsilon > 0$. Moreover, note that $g_\epsilon(\Sigma^{1/2} x)$ is also polynomially bounded for any PSD matrix Σ . Specifically, for any fixed PSD matrix Σ_0 we have for any $k \geq 1$

$$\begin{aligned} & \sup_{\epsilon \in [0,1]} \sup_{\substack{\|\Sigma - \Sigma_0\| \leq 1 \\ \Sigma \text{ PSD}}} \sup_{x \in \mathbb{R}^r} \left| (1 + \|x\|)^{-k} g_\epsilon(\Sigma^{1/2} x) \right| \\ &= \sup_{\epsilon \in [0,1]} \sup_{\substack{\|\Sigma - \Sigma_0\| \leq 1 \\ \Sigma \text{ PSD}}} \sup_{x \in \mathbb{R}^r} \left| (1 + \|x\|)^{-k} \int_{\mathbb{R}^r} g(\Sigma^{1/2}(x-y)) \psi_\epsilon(y) dy \right| \\ &\leq \sup_{\epsilon \in [0,1]} \sup_{\substack{\|\Sigma - \Sigma_0\| \leq 1 \\ \Sigma \text{ PSD}}} \sup_{x \in \mathbb{R}^r} \left\{ (1 + \|x\|)^{-k} \int_{\mathbb{R}^r} \left(1 + \left\| \Sigma^{1/2}(x-y) \right\|^k \right) \psi_\epsilon(y) dy \right\} \\ &< \infty, \end{aligned} \quad (6.12)$$

Note that there exists $K > 0$ depending only k, r, Σ_0 so that

$$\sup_{\|\Sigma - \Sigma_0\| \leq 1} \left\| \Sigma^{1/2}(x-y) \right\|^k \leq K \left(1 + \left\| \Sigma_0^{1/2} \right\|^k \right) (\|x\|^k + \|y\|^k).$$

Hence, since

$$\sup_{\epsilon \in [0,1]} \int_{\mathbb{R}^r} \|y\|^k \psi_\epsilon(y) dy < \infty$$

we find that

$$\sup_{\epsilon \in [0,1]} \sup_{\substack{\|\Sigma - \Sigma_0\| \leq 1 \\ \Sigma \text{ PSD}}} \sup_{x \in \mathbb{R}^r} \left| (1 + \|x\|)^{-k} g_\epsilon(\Sigma^{1/2} x) \right| < \infty. \quad (6.13)$$

The estimate above allows us to use dominate convergence to see that for any PSD Σ

$$\mathbb{E} \left[g(\Sigma^{1/2} X) \right] = \lim_{\epsilon \rightarrow 0} \mathbb{E} \left[g_\epsilon(\Sigma^{1/2} X) \right]. \quad (6.14)$$

To complete the proof we note that g_ϵ and $\partial_i \partial_j g_\epsilon$ are both Schwartz for any positive ϵ . Moreover, $\partial_i \partial_j \partial_k \partial_m g_\epsilon$ satisfies (6.13). Hence, we conclude by applying (6.11) that for any PSD matrix Σ_0 there exists $C > 0$ so that

$$\begin{aligned} & \sup_{\substack{\|\Sigma - \Sigma_0\| \leq 1 \\ \Sigma \text{ PSD}}} \sup_{\epsilon \in [0,1]} \frac{\left| \mathbb{E} \left[g_\epsilon(\Sigma^{1/2} X) \right] - \mathbb{E} \left[g_\epsilon(\Sigma_0^{1/2} X) \right] - \sum_{i,j=1}^r \mathbb{E} \left[(\partial_i \partial_j g_\epsilon)(\Sigma_0^{1/2} X) \right] (\Sigma - \Sigma_0)_{ij} \right|}{\|\Sigma - \Sigma_0\|^2} \\ & \leq \sup_{\epsilon \in [0,1]} \sup_{\|\Sigma - \Sigma_0\| \leq 1} \sum_{i,j,k,m=1,\dots,r} \left| \mathbb{E} \left[(\partial_i \partial_j \partial_k \partial_m) g_\epsilon(\Sigma^{1/2} X) \right] \right| \\ & \leq C. \end{aligned}$$

Thus, if $\Sigma - \Sigma_0 / \|\Sigma - \Sigma_0\| \rightarrow \Sigma_1$, we find by applying (6.14) to $\partial_i \partial_j g$ that

$$\begin{aligned} \lim_{\Sigma \rightarrow \Sigma_0} \frac{\mathbb{E} \left[g(\Sigma^{1/2} X) \right] - \mathbb{E} \left[g(\Sigma_0^{1/2} X) \right]}{\|\Sigma - \Sigma_0\|} &= \lim_{\Sigma \rightarrow \Sigma_0} \lim_{\epsilon \rightarrow 0} \frac{\mathbb{E} \left[g_\epsilon(\Sigma^{1/2} X) \right] - \mathbb{E} \left[g_\epsilon(\Sigma_0^{1/2} X) \right]}{\|\Sigma - \Sigma_0\|} \\ &= \lim_{\Sigma \rightarrow \Sigma_0} \lim_{\epsilon \rightarrow 0} \left\{ \sum_{i,j=1}^r \mathbb{E} \left[(\partial_i \partial_j g_\epsilon)(\Sigma_0^{1/2} X) \right] \frac{(\Sigma - \Sigma_0)_{ij}}{\|\Sigma - \Sigma_0\|} \right\} \\ &= \sum_{i,j=1}^r \mathbb{E} \left[\partial_i \partial_j g(\Sigma_0^{1/2} X) \right] (\Sigma_1)_{ij}. \end{aligned}$$

This shows that (6.10) holds for any g that is smooth, completing the proof of Lemma 6.3. \square

7 Proof of Theorem 3.1

Let us first recall the notation. We fix $r \geq 1$ and assume that $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ satisfies assumption 2.1 with this value of r . We also fix a finite collection $x_{\mathcal{A}} = \{x_\alpha, \alpha \in \mathcal{A}\} \subseteq \mathbb{R}^{n_0}$ of distinct network inputs and p directional derivatives d_1, \dots, d_p as in (3.2). We denote by

$$N(p, r) = \# \{J = (j_1, \dots, j_p) \in \mathbb{N}^p \mid j_1 + \dots + j_p \leq r\},$$

which computes the number of possible derivatives of order at most r in the p directional derivatives d_j . We also denote by $\mathcal{F}^{(\ell)}$ the sigma algebra generated by the weights and biases in layers up to and including ℓ . The starting point for proving Theorem 3.1 is the following simple but fundamental observation.

Lemma 7.1. *For each $\ell \geq 0$, conditional on $\mathcal{F}^{(\ell)}$,*

$$\left\{ \left(D_\alpha^J z_{i;\alpha}^{(\ell+1)}, \alpha \in \mathcal{A}, |J| \leq r \right) \right\}_{i=1}^{n_{\ell+1}}$$

is a collection of $n_{\ell+1}$ iid centered Gaussians of dimension $N(p, r)$.

Proof. The defining recursion (2.1) of a fully connected network yields for each α, J

$$D_\alpha^J z_{i;\alpha}^{(\ell+1)} = D_\alpha^J \left\{ b_i^{(\ell+1)} + \sum_{j=1}^{n_\ell} W_{ij}^{(\ell+1)} \sigma \left(z_{j;\alpha}^{(\ell)} \right) \right\} = \delta_{|J|=0} b_i^{(\ell+1)} + \sum_{j=1}^{n_\ell} W_{ij}^{(\ell+1)} D_\alpha^J \sigma \left(z_{j;\alpha}^{(\ell)} \right). \quad (7.1)$$

Note that $D_\alpha^J \sigma(z_{j;\alpha}^{(\ell)})$ are measurable with respect to $\mathcal{F}^{(\ell)}$. The conclusion now follows since the weights $W_{ij}^{(\ell+1)}$, $j = 1 \dots, n_\ell$ and bias $b_i^{(\ell+1)}$ are centered Gaussians and are independent for different i . \square

Thus, the structure of $z_\alpha^{(\ell+1)}$ and its derivatives is always that of a Gaussian field after conditioning on $\mathcal{F}^{(\ell)}$. To ease the notation in what comes given $f : \mathbb{R}^{|\mathcal{A}| \times N(n_0, r)} \rightarrow \mathbb{R}$, let us abbreviate

$$f \left(z_{j;\mathcal{A}}^{(\ell)} \right) := f \left(D_\alpha^J z_{j;\alpha}^{(\ell)}, \alpha \in \mathcal{A}, |J| \leq r \right), \quad j \in [n_\ell].$$

Next, we remind the reader that given $f : \mathbb{R}^{|\mathcal{A}| \times N(n_0, r)} \rightarrow \mathbb{R}$, which is measurable and polynomially bounded, the corresponding collective observable $\mathcal{O}_f^{(\ell)}$ at layer ℓ is

$$\mathcal{O}_f^{(\ell)} = \frac{1}{n_\ell} \sum_{j=1}^{n_\ell} f \left(z_{j;\mathcal{A}}^{(\ell)} \right)$$

and that the statement (A.3) in Proposition A.1 ensures

$$\sup_{n \geq 1} \mathbb{E} \left[\left| \mathcal{O}_f^{(\ell)} \right| \right] < \infty. \quad (7.2)$$

Recall also our notation for the conditional covariance

$$\Sigma_{\alpha_1 \alpha_2}^{(\ell)} := \text{Cov} \left(z_{i;\alpha_1}^{(\ell+1)}, z_{i;\alpha_2}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) = C_b + \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \sigma \left(z_{j;\alpha_1}^{(\ell)} \right) \sigma \left(z_{j;\alpha_2}^{(\ell)} \right)$$

and note that both it and its derivatives

$$\begin{aligned} D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} &= \text{Cov} \left(D_{\alpha_1}^{J_1} z_{i;\alpha_1}^{(\ell+1)}, D_{\alpha_2}^{J_2} z_{i;\alpha_2}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) \\ &= D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \left(C_b + \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \sigma \left(z_{j;\alpha_1}^{(\ell)} \right) \sigma \left(z_{j;\alpha_2}^{(\ell)} \right) \right) \end{aligned}$$

are collective observables at layer ℓ . Our first application of Lemma 7.1 is the following reduction of the study of cumulants of $D_\alpha^J z_{i;\alpha}^{(\ell+1)}$ to the cumulants of certain collective observables.

Proposition 7.2. *Fix $k, \ell \geq 1$ and p -dimensional multi-indices J_1, \dots, J_k with $|J_i| \leq r$. If k is odd, then*

$$\kappa \left(D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}^{(\ell+1)}, \dots, D_{\alpha_k}^{J_k} z_{i_k;\alpha_k}^{(\ell+1)} \right) = 0$$

In contrast, if k is even

$$\kappa \left(D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}^{(\ell+1)}, \dots, D_{\alpha_k}^{J_k} z_{i_k;\alpha_k}^{(\ell+1)} \right) = \text{finite sums of } \kappa \left(\mathcal{O}_{f_1}^{(\ell)}, \dots, \mathcal{O}_{f_{k/2}}^{(\ell)} \right),$$

where $\mathcal{O}_{f_j}^{(\ell)}$ are collective observables of the form

$$D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)}, \quad |J_1|, |J_2| \leq r. \quad (7.3)$$

Proof. Using (6.1) and recalling that $\mathcal{F}^{(\ell)}$ is the sigma algebra generated by weights and biases in layers up to and including ℓ , we have that $\kappa\left(D_{\alpha_1}^{J_1} z_{i_1; \alpha_1}^{(\ell+1)}, \dots, D_{\alpha_k}^{J_k} z_{i_k; \alpha_k}^{(\ell+1)}\right)$ equals

$$\sum_{\pi=(\pi_1, \dots, \pi_B)} \kappa\left(\kappa\left(\left(D^J z^{(\ell+1)}\right)_{\pi_1} \mid \mathcal{F}^{(\ell)}\right), \dots, \kappa\left(\left(D^J z^{(\ell+1)}\right)_{\pi_B} \mid \mathcal{F}^{(\ell)}\right)\right), \quad (7.4)$$

where the sum is over partitions π of $[k]$ and for $b = 1, \dots, B$ we've abbreviated

$$\left(D^J z^{(\ell+1)}\right)_{\pi_b} := \left(D_{\alpha_t}^{J_t} z_{i_t; \alpha_t}^{(\ell+1)}, \quad t \in \pi_b\right).$$

By Lemma 7.1, $\{(D_{\alpha}^J z_{i; \alpha}^{(\ell+1)}, \alpha \in \mathcal{A}, |J| \leq d), i = 1, \dots, n_{\ell+1}\}$ are iid centered Gaussians conditional on $\mathcal{F}^{(\ell)}$. Hence, by the properties (6.2) and (6.3) and (6.4) from Proposition 6.1, in the sum over partitions above, a term is non-zero only if

$$\forall b \in [B], \quad |\pi_b| = 2 \quad \text{and} \quad i_{\pi_b(1)} = i_{\pi_b(2)}$$

This proves that $\kappa\left(D_{\alpha_1}^{J_1} z_{i_1; \alpha_1}^{(\ell+1)}, \dots, D_{\alpha_k}^{J_k} z_{i_k; \alpha_k}^{(\ell+1)}\right)$ vanishes if k is odd. To treat the case when k is even observe that by (7.1)

$$\kappa\left(D_{\alpha_1}^{J_1} z_{i_1; \alpha_1}^{(\ell+1)}, D_{\alpha_2}^{J_2} z_{i_2; \alpha_2}^{(\ell+1)} \mid \mathcal{F}^{(\ell)}\right) = \delta_{i_1 i_2} D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)}.$$

Substituting this into (7.4) completes the proof. \square

When $k = 2$, Proposition 7.2 and our assumption (2.1) shows that for each $\ell \geq 1$, any $i_1, i_2 \in [n_{\ell+1}]$, $\alpha \in \mathcal{A}$, and multi-indices J_1, J_2 of order at most d , there exists a polynomially bounded function $f : \mathbb{R}^{|\mathcal{A}| \times N(n_0, d)} \rightarrow \mathbb{R}$ for which

$$\kappa\left(D_{\alpha_1}^{J_1} z_{i_1; \alpha_1}^{(\ell+1)}, D_{\alpha_2}^{J_2} z_{i_2; \alpha_2}^{(\ell+1)}\right) = \mathbb{E}\left[\mathcal{O}_f^{(\ell)}\right]$$

In light of (7.2) this proves Theorem 3.1 when $k = 2$. Further, since the cumulant of 2 or more random variables is shift-invariant, we may assume for $k \geq 3$ that the collective observables $D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)}$ in Proposition 7.2 are replaced by their zero mean versions:

$$\Delta_{\alpha_1 \alpha_2}^{J_1, J_2, (\ell)} := D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} - \mathbb{E}\left[D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)}\right]. \quad (7.5)$$

Hence, Theorem 3.1 is a special case of the following result.

Theorem 7.3. *Fix $k, m \geq 1$. Consider any m -tuple $F = (f_1, \dots, f_m)$ consisting of measurable, functions*

$$f_i : \mathbb{R}^{|\mathcal{A}| \times N(n_0, r)} \rightarrow \mathbb{R}, \quad i = 1, \dots, m$$

that are polynomially bounded and satisfy

$$\mathbb{E}\left[\mathcal{O}_{f_i}^{(\ell)}\right] = \mathbb{E}\left[f_i\left(z_{1; \mathcal{A}}^{(\ell)}\right)\right] = 0, \quad i = 1, \dots, m.$$

Define the m -tuple of collective observables

$$\vec{\mathcal{O}}_F^{(\ell)} := \left(\mathcal{O}_{f_i}^{(\ell)}, i = 1, \dots, m\right).$$

Consider further any measurable polynomially bounded functions

$$g_j : \mathbb{R}^m \rightarrow \mathbb{R}, \quad j = 1, \dots, k.$$

which are smooth in a neighborhood of 0. If f_i and σ are in fact smooth, then, for every $\ell \geq 1$

$$\sup_{n \geq 1} \left| n^{k-1} \kappa \left(g_1 \left(\vec{\mathcal{O}}_F^{(\ell)} \right), \dots, g_k \left(\vec{\mathcal{O}}_F^{(\ell)} \right) \right) \right| < \infty \quad (7.6)$$

Moreover, (7.6) holds without the assumption that f_i, σ are smooth provided that for each ℓ the vector of iterated directional derivatives $(D_{\alpha}^J z_{i;\alpha}^{(\ell)}, |J| \leq r, \alpha \in \mathcal{A})$ of order at most r is non-degenerate in the sense of (3.3).

Proof. Our starting point is a reduction of Theorem 7.3 to the case when g_j are polynomials. This is related to a technique called the delta method in some parts of the mathematical statistics literature [59].

Proposition 7.4 (Polynomials are Enough for Theorem 7.3). *Fix $m \geq 1$ and suppose that for each $n \geq 1$ we have an m -tuple $X_n = (X_{n,1}, \dots, X_{n,m})$ of mean 0 random variables that possess bounded moments of all orders:*

$$\sup_{n \geq 1} \left| \mathbb{E} \left[X_{n,1}^{q_1} \cdots X_{n,m}^{q_m} \right] \right| < \infty, \quad \forall q_1, \dots, q_m \geq 0. \quad (7.7)$$

Suppose for any given polynomials p_1, \dots, p_k in m variables we have

$$\sup_{n \geq 1} \left| n^{k-1} \kappa (p_1(X_n), \dots, p_k(X_n)) \right| < \infty. \quad (7.8)$$

Then, for any measurable, polynomially bounded functions $g_j : \mathbb{R}^m \rightarrow \mathbb{R}$, $j = 1, \dots, k$, which are smooth in some fixed neighborhood of 0

$$\sup_{n \geq 1} \left| n^{k-1} \kappa (g_1(X_n), \dots, g_k(X_n)) \right| < \infty. \quad (7.9)$$

Proof. We begin with the following simple Lemma, which allows us to translate between the cumulants bounds (7.8) and high probability bounds.

Lemma 7.5. *For any $q \geq 1$*

$$\sup_{n \geq 1} \sup_{1 \leq i \leq m} \left| n^{\lceil \frac{q}{2} \rceil} \mathbb{E} \left[X_{i;n}^q \right] \right| < \infty.$$

Proof. We have by property (6.5) from Proposition 6.1 that

$$\mathbb{E} \left[X_{i;n}^q \right] = \sum_{\substack{\pi = (\pi_1, \dots, \pi_B) \\ \pi \in P(m)}} \prod_{b=1}^B \kappa \left(\underbrace{X_{i;n}, \dots, X_{i;n}}_{|\pi_b| \text{ times}} \right).$$

Since by assumption $X_{i;n}$ has mean 0, we have

$$\kappa(X_{i;n}) = \mathbb{E} [X_{i;n}] = 0.$$

Thus, the only partitions $\pi = (\pi_1, \dots, \pi_B) \in S(m)$ that give rise to non-zero terms in the expression above must have $B \leq \lfloor \frac{q}{2} \rfloor$. Moreover, for any such partition, we have

$$\left\lfloor \frac{q}{2} \right\rfloor = q - \left\lfloor \frac{q}{2} \right\rfloor = - \left\lfloor \frac{q}{2} \right\rfloor + \sum_{b=1}^B |\pi_b| \leq \sum_{b=1}^B (|\pi_b| - 1).$$

Hence, we find

$$\sup_{n \geq 1} \left| n^{\lfloor \frac{q}{2} \rfloor} \mathbb{E} \left[X_{i;n}^q \right] \right| \leq \sum_{\substack{\pi = (\pi_1, \dots, \pi_B) \\ \pi \in P(m), |\pi_b| \geq 2}} \prod_{b=1}^B \sup_{n \geq 1} \left| n^{|\pi_b| - 1} \kappa \left(\underbrace{X_{i;n}, \dots, X_{i;n}}_{|\pi_b| \text{ times}} \right) \right| < \infty,$$

where the final inequality follows from the assumption (7.8). \square

Applying Markov's inequality and Lemma 7.5 shows that for any $q \geq 1$ we have

$$\sup_{n \geq 1} n^q \mathbb{P}(S_n^c) < \infty, \quad S_n := \left\{ |X_{i;n}| \leq n^{-1/4}, \quad i = 1, \dots, m \right\}. \quad (7.10)$$

This localization estimate allows us to replace each g_i by its Taylor expansion around 0. Indeed, note that

$$\kappa(g_1(X_n), \dots, g_k(X_n)) = P(\mathbb{E}[g_1(X_n)^{q_1} \dots g_k(X_n)^{q_k}], \quad q_1 + \dots + q_k \leq k)$$

for some universal polynomial P evaluated at the mixed moments of X_n (the formula for this polynomial is given in (6.6) but is not important). Moreover, using the growth assumption (7.7) on X and the fact that g_i are polynomially bounded we find that

$$\sup_{n \geq 1} \mathbb{E}[g(X_n)^{q_1} \dots g_k(X_n)^{q_k}] < \infty, \quad \forall q_1, \dots, q_k \geq 1. \quad (7.11)$$

This, in combination with the localization estimate (7.10) applied with $q = k - 1$ yields

$$\kappa(g_1(X_n), \dots, g_k(X_n)) = P(\mathbb{E}[\mathbf{1}_{S_n} g_1(X_n)^{q_1} \dots g_k(X_n)^{q_m}], \quad q_1 + \dots + q_m \leq k) + O(n^{-k+1}).$$

Note that for n sufficiently large, on the event S_n , the argument X_n is any fixed neighborhood of 0. Hence, we may write

$$g_j(X_n) = p_j(X_n) + O(n^{-k+1}),$$

where p_j represents the q -th order Taylor expansion of g_j around 0 with q sufficiently large and the constant in the error term is uniformly bounded. This yields

$$\kappa(g_1(X_n), \dots, g_k(X_n)) = P(\mathbb{E}[\mathbf{1}_{S_n} p_1(X_n)^{q_1} \dots p_k(X_n)^{q_m}], \quad q_1 + \dots + q_m \leq k) + O(n^{-k+1}).$$

Finally, using the mixed moment estimates (7.7) and the localization estimate (7.10), we conclude

$$\begin{aligned} \kappa(g_1(X_n), \dots, g_k(X_n)) &= P(\mathbb{E}[p_1(X_n)^{q_1} \dots p_k(X_n)^{q_m}], \quad q_1 + \dots + q_m \leq k) + O(n^{-k+1}) \\ &= \kappa(p_1(X_n), \dots, p_k(X_n)) + O(n^{-k+1}). \end{aligned}$$

Recalling (7.8) completes the proof. \square

Proposition 7.4 shows that, in establishing the conclusion (7.6) of Theorem 7.3, it is sufficient to assume that g_j are polynomials. The remainder of the proof of Theorem 7.3 is by induction on ℓ , starting with $\ell = 1$. In view of Proposition 7.4, the following result establishes the base case.

Proposition 7.6 (Base Case: Theorem 7.3 holds for polynomials at $\ell = 1$). *Fix $k, m \geq 1$ and suppose $f_i, i = 1, \dots, m$ are as in the statement of Theorem 7.3. Then, if p_1, \dots, p_k are any polynomials in m variables, we have*

$$\sup_{n \geq 1} \left| n^{k-1} \kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(1)} \right) \right) \right| < \infty.$$

Proof. Since cumulants are multi-linear, we may and shall assume that p_a are monomials:

$$p_a(x) = x^{Q^{(a)}} := x_1^{q_1^{(a)}} \cdots x_m^{q_m^{(a)}}, \quad x = (x_1, \dots, x_m), \quad Q^{(a)} = \left(q_1^{(a)}, \dots, q_m^{(a)} \right). \quad (7.12)$$

Recall that

$$\mathcal{O}_{f_i}^{(1)} := n_1^{-1} \sum_{j=1}^{n_1} f_i \left(z_{j;\mathcal{A}}^{(1)} \right).$$

Therefore, writing $q^{(a)} := q_1^{(a)} + \dots + q_m^{(a)}$ we find

$$p_a \left(\vec{\mathcal{O}}_F^{(1)} \right) = n_1^{-q^{(a)}} \sum_{J^{(a)}} f_{J^{(a)}}, \quad f_{J^{(a)}} := \prod_{i=1}^m \prod_{q=1}^{q_i^{(a)}} f_i \left(z_{j_{q,i};\mathcal{A}}^{(1)} \right),$$

where the sum is over tuples of multi-indices

$$J^{(a)} = \left(J_1^{(a)}, \dots, J_m^{(a)} \right), \quad J_i^{(a)} = \left(j_{q,i}^{(a)} \in [n_1], i \in [m], q \in [q_i^{(a)}] \right). \quad (7.13)$$

Hence, using that cumulants are multi-linear (see (6.3)), we obtain

$$\kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(1)} \right) \right) = n_1^{-(q^{(1)} + \dots + q^{(k)})} \sum_{J^{(1)}, \dots, J^{(k)}} \kappa \left(f_{J^{(1)}}, \dots, f_{J^{(k)}} \right),$$

where the sum extends over ordered collections $(J^{(a)}, 1 \leq a \leq k)$ of multi-indices as in (7.13). The expression on the right hand side can be interpreted as an average. Namely, we can think of the indices $j_{q,i}^{(a)} \in [n_1]$ are chosen uniformly from $[n_1]$ and independently for all i, q, a . Writing \mathcal{E} for the average with respect to this distribution, we obtain

$$\kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(1)} \right) \right) = \mathcal{E} \left[\kappa \left(f_{J^{(1)}}, \dots, f_{J^{(k)}} \right) \right].$$

Our goal is to show that this average is small. To quantify this, let us associate to each collection $(J^{(a)}, a \in [k])$ a graph

$$\mathcal{G} \left(J^{(a)}, a \in [k] \right) = \left([k], E \left(J^{(a)}, a \in [k] \right) \right), \quad (7.14)$$

with vertex set $[k]$ and edge set defined by

$$(a, a') \in \mathcal{E} \left(J^{(a)}, a \in [k] \right) \iff \exists i, i' \in [m], q \in [q_i^{(a)}], q' \in [q_{i'}^{(a')}] \text{ s.t. } j_{q,i}^{(a)} = j_{q',i'}^{(a')}.$$

The key point is that in light of the vanishing property (6.2) of cumulants and the fact that neurons at layer 1 are independent

$$\mathcal{G}\left(J^{(a)}, a \in [k]\right) \text{ disconnected} \implies \kappa\left(f_{J^{(1)}}^{(1)}, \dots, f_{J^{(k)}}^{(1)}\right) = 0.$$

Hence,

$$\kappa\left(p_1\left(\vec{\mathcal{O}}_F^{(1)}\right), \dots, p_k\left(\vec{\mathcal{O}}_F^{(1)}\right)\right) = \mathcal{E}\left[\mathbf{1}_{\{\mathcal{G}(J^{(a)}, a \in [k]) \text{ connected}\}} \kappa\left(f_{J^{(1)}}^{(1)}, \dots, f_{J^{(k)}}^{(1)}\right)\right].$$

Since f_i are assumed to be polynomially bounded and the distribution of the neuron pre-activations $z_{i;\alpha}^{(1)}$ is that of centered Gaussians with mean 0 and covariance

$$\text{Cov}\left(z_{i_1;\alpha_1}^{(1)}, z_{i_2;\alpha_2}^{(1)}\right) = \delta_{i_1 i_2} \left(C_b + \frac{C_W}{n_0} \sum_{j=1}^{n_0} x_{j;\alpha_1} x_{j;\alpha_2} \right),$$

we have for any fixed k that

$$\sup_{n \geq 1} \sup_{J^{(1)}, \dots, J^{(k)}} \left| \kappa\left(f_{J^{(1)}}^{(1)}, \dots, f_{J^{(k)}}^{(1)}\right) \right| < \infty.$$

Hence,

$$\kappa\left(p_1\left(\vec{\mathcal{O}}_F^{(1)}\right), \dots, p_k\left(\vec{\mathcal{O}}_F^{(1)}\right)\right) = O\left(\mathcal{P}\left(\mathcal{G}\left(J^{(a)}, a \in [k]\right) \text{ connected}\right)\right),$$

where \mathcal{P} is the probability measure associated to our random draw of $J^{(1)}, \dots, J^{(k)}$. To complete the proof, note that since $m, q_i^{(a)}$ are fixed, by a simple union bound, we obtain

$$\mathcal{P}\left(\mathcal{G}\left(J^{(a)}, a \in [k']\right) \text{ connected} \mid \mathcal{G}\left(J^{(a)}, a \in [k' - 1]\right) \text{ connected}\right) = O(n^{-1}).$$

Hence,

$$\begin{aligned} & \mathcal{P}\left(\mathcal{G}\left(J^{(a)}, a \in [k]\right) \text{ connected}\right) \\ &= \prod_{k'=2}^k \mathcal{P}\left(\mathcal{G}\left(J^{(a)}, a \in [k']\right) \text{ connected} \mid \mathcal{G}\left(J^{(a)}, a \in [k' - 1]\right) \text{ connected}\right) \\ &= O(n^{-k+1}). \end{aligned} \tag{7.15}$$

Thus,

$$\kappa\left(p_1\left(\vec{\mathcal{O}}_F^{(1)}\right), \dots, p_k\left(\vec{\mathcal{O}}_F^{(1)}\right)\right) = O(n^{-k+1}),$$

as desired. \square

Propositions 7.4 and 7.6 together show that the conclusion (7.6) of Theorem 7.3 holds at layer 1. In conjunction with Proposition 7.4, the following result establishes that if the conclusion (7.6) of Theorem 7.3 holds at some layer $\ell \geq 1$ then it also holds at layer $\ell + 1$. This will complete the proof by inductive of Theorem 7.3.

Proposition 7.7 (Inductive Step: Reducing polynomial cumulants in layer $\ell + 1$ to smooth cumulants in layer ℓ). *Fix $\ell \geq 1$.*

Case 1: *Suppose that σ is smooth. Assume that for any collection*

$$F' = \left(f'_i : \mathbb{R}^{|\mathcal{A}| \times N(n_0, r)} \rightarrow \mathbb{R}, i = 1, \dots, m \right)$$

of smooth and polynomially bounded functions and any g_j as in the statement of Theorem 7.3 the conclusion (7.6) of Theorem 7.3 holds at layer ℓ :

$$\sup_{n \geq 1} \left| n^{k-1} \kappa \left(g_1 \left(\vec{\mathcal{O}}_{F'}^{(\ell)} \right), \dots, g_k \left(\vec{\mathcal{O}}_{F'}^{(\ell)} \right) \right) \right| < \infty.$$

Then, if p_1, \dots, p_k are any polynomials in m variables, and $F = (f_i, i = 1, \dots, m)$ is an arbitrary collection of smooth and polynomially bounded functions $f_i : \mathbb{R}^{|\mathcal{A}| \times N(n_0, r)} \rightarrow \mathbb{R}$, then

$$\sup_{n \geq 1} \left| n^{k-1} \kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right) \right) \right| < \infty.$$

Case 2: *Suppose σ is not smooth but satisfies Assumption 2.1 and that $(D_\alpha^J z_{i;\alpha}^{(\ell)}, \alpha \in \mathcal{A}, |J| \leq r)$ is non-degenerate in the infinite width in the sense of (3.3). Assume that for any collection*

$$F' = \left(f'_i : \mathbb{R}^{|\mathcal{A}| \times N(n_0, r)} \rightarrow \mathbb{R}, i = 1, \dots, m \right)$$

of measurable and polynomially bounded functions and any g_j as in the statement of Theorem 7.3 the conclusion (7.6) of Theorem 7.3 holds at layer ℓ :

$$\sup_{n \geq 1} \left| n^{k-1} \kappa \left(g_1 \left(\vec{\mathcal{O}}_{F'}^{(\ell)} \right), \dots, g_k \left(\vec{\mathcal{O}}_{F'}^{(\ell)} \right) \right) \right| < \infty.$$

Then, if p_1, \dots, p_k are any polynomials in m variables, and $F = (f_i, i = 1, \dots, m)$ is an arbitrary collection of measurable and polynomially bounded functions $f_i : \mathbb{R}^{|\mathcal{A}| \times N(n_0, d)} \rightarrow \mathbb{R}$, then

$$\sup_{n \geq 1} \left| n^{k-1} \kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right) \right) \right| < \infty.$$

Proof. The proof of Proposition 7.7 is similar but somewhat more involved than that of Proposition 7.6. Moreover, the two cases are proved in essentially the same way, except that we will employ the different cases in Lemma 6.3. We give the details in the case when σ is smooth and indicate where the proof is modified slightly to handle the non-smooth case.

To start, as in the proof of Proposition 7.6, note that since cumulants are multi-linear (see (6.3)), it is enough to assume that p_j are monomials. Thus, borrowing the notation from the proof of Proposition 7.6 (see starting (7.12)), we find

$$\kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right) \right) = n_{\ell+1}^{-(q^{(1)} + \dots + q^{(a)})} \sum_{J^{(1)}, \dots, J^{(k)}} \kappa \left(f_{J^{(1)}}^{(\ell+1)}, \dots, f_{J^{(k)}}^{(\ell+1)} \right),$$

where

$$f_{J^{(a)}}^{(\ell+1)} := \prod_{i=1}^m \prod_{q=1}^{q_i^{(a)}} f_{j_\alpha^{(a)}}^{(\ell+1)}, \quad f_j^{(\ell+1)} := f \left(z_{j; \mathcal{A}}^{(\ell+1)} \right).$$

Note that, as in Proposition A.1, the polynomially bounded assumption on f_j and the non-linearity σ together with the Gaussianity of weights and biases show that

$$\sup_{n \geq 1} \left| \kappa \left(f_{J^{(1)}}^{(\ell+1)}, \dots, f_{J^{(k)}}^{(\ell+1)} \right) \right| < \infty. \quad (7.16)$$

Using the law of total cumulance (6.1), we find that $\kappa \left(p_1(\vec{\mathcal{O}}_F^{(\ell+1)}), \dots, p_k(\vec{\mathcal{O}}_F^{(\ell+1)}) \right)$ equals

$$\sum_{\pi=(\pi_1, \dots, \pi_B)} n_{\ell+1}^{-(q^{(1)} + \dots + q^{(a)})} \sum_{J^{(1)}, \dots, J^{(k)}} \kappa \left(\kappa \left(f_{J^{(\pi_1)}}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right), \dots, \kappa \left(f_{J^{(\pi_B)}}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) \right),$$

where π is any partition of $[k]$ and

$$f_{J^{(\pi_b)}} := (f_{J^{(a)}}), \quad a \in \pi_b.$$

Just as in the proof of Proposition 7.6, we may interpret the sum over $J^{(1)}, \dots, J^{(k)}$ as an average over the distribution in which $j_{q;i}^{(a)}$ are drawn iid uniformly on $[n_{\ell+1}]$. Writing \mathcal{E} for averages with respect to this distribution yields

$$\kappa \left(p_1(\vec{\mathcal{O}}_F^{(\ell+1)}), \dots, p_k(\vec{\mathcal{O}}_F^{(\ell+1)}) \right) = \sum_{\pi=(\pi_1, \dots, \pi_b)} \mathcal{E} \left[\kappa \left(\kappa \left(f_{J^{(\pi_1)}}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right), \dots, \kappa \left(f_{J^{(\pi_b)}}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) \right) \right]$$

As in (7.14), we may associate to each collection $J^{(\pi_t)}$ the graph $\mathcal{G}(J^{(\pi_t)})$. Recall that by Lemma 7.1, the neurons pre-activations $D_\alpha^J z_{i;\alpha}^{(\ell+1)}$ in layer $\ell+1$ are independent for different i conditional on $\mathcal{F}^{(\ell)}$. Hence, in view of the vanishing property (6.2) of cumulants, we obtain that $\kappa \left(p_1(\vec{\mathcal{O}}_F^{(\ell+1)}), \dots, p_k(\vec{\mathcal{O}}_F^{(\ell+1)}) \right)$ equals

$$\sum_{\pi=(\pi_1, \dots, \pi_B)} \mathcal{E} \left[\mathbf{1}_{\{\mathcal{G}(J^{(\pi_b)}) \text{ connected } \forall b \in [B]\}} \kappa \left(\kappa \left(f_{J^{(\pi_1)}}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right), \dots, \kappa \left(f_{J^{(\pi_B)}}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) \right) \right]$$

Since we've assumed that f_i are smooth and polynomially bounded, Lemma 6.3 shows that for each $b \in [B] = \{1, \dots, B\}$ the conditional cumulant $\kappa \left(f_{J^{(\pi_b)}}^{(\ell+1)} \mid z^{(\ell)} \right)$ is a smooth function of the centered entries $D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Delta_{\alpha_1 \alpha_2}^{(\ell)}$ of the conditional covariance of $\left(D_{\alpha}^J z_{i;\mathcal{A}}^{(\ell+1)}, \alpha \in \mathcal{A}, |J| \leq r \right)$ given $\mathcal{F}^{(\ell)}$. Thus, since these entries are collective observables at layer ℓ we may apply the inductive hypothesis of Case 1 to find that

$$\kappa \left(p_1(\vec{\mathcal{O}}_F^{(\ell+1)}), \dots, p_k(\vec{\mathcal{O}}_F^{(\ell+1)}) \right) = \sum_{\pi=(\pi_1, \dots, \pi_B)} \mathcal{P} \left[\mathcal{G}(J^{(\pi_b)}) \text{ connected } \forall b \in [B] \right] O(n^{-B+1}).$$

Combining this with the estimate (7.15) shows

$$\kappa \left(p_1(\vec{\mathcal{O}}_F^{(\ell+1)}), \dots, p_k(\vec{\mathcal{O}}_F^{(\ell+1)}) \right) = \sum_{\pi=(\pi_1, \dots, \pi_B)} O(n^{-B+1}) \prod_{b=1}^B O(n^{-|\pi_b|+1}) = O(n^{-k+1}),$$

as desired. The proof in Case 2 is almost identical. The only difference is that, we must introduce the event

$$S_n = \left\{ \left| \Delta_{\alpha_1 \alpha_2}^{J_1 J_2, (\ell)} \right| < n^{-1/4} \right\}. \quad (7.17)$$

Precisely as in the proof of Lemma 7.5 we find that

$$\mathbb{P}(S_n^c) = O(n^{-\infty}).$$

Hence,

$$\kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right) \right) = \kappa \left(p_1 \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right), \dots, p_k \left(\vec{\mathcal{O}}_F^{(\ell+1)} \right) \mid S_n \right) + O(n^{-\infty}),$$

where we've implicitly used (7.16). Moreover, since in Case 2 we assume that the vector $\left(D_\alpha^J z_{i;\alpha}^{(\ell+1)}, |J| \leq r, \alpha \in \mathcal{A} \right)$ is non-degenerate in the infinite width limit in the sense of (3.3), we see that for n sufficiently large the covariance of $\left(D_\alpha^J z_{i;\alpha}^{(\ell+1)}, |J| \leq r, \alpha \in \mathcal{A} \right)$ given $\mathcal{F}^{(\ell)}$, which is the matrix with entries

$$\mathbb{E} \left[D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} \right], \quad \alpha_1, \alpha_2 \in \mathcal{A}, |J_1|, |J_2| \leq r$$

is also non-degenerate. On the event S_n , the conditional covariance of $\left(D_\alpha^J z_{i;\alpha}^{(\ell+1)}, |J| \leq r, \alpha \in \mathcal{A} \right)$ given $\mathcal{F}^{(\ell)}$, which is a matrix with entries $D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)}$, is also non-degenerate for all n sufficiently large. Hence, we again conclude by Lemma 6.3 that conditional on S_n (which is measurable with respect to $\mathcal{F}^{(\ell)}$) for each $b \in [B]$ the conditional cumulant $\kappa \left(f_{J(\pi_b)}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right)$ is a smooth function of $D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)}$, which are collective observables at layer ℓ . The remainder of the proof now proceeds in the same way as for Case 1. \square

\square

8 Proof of Theorem 3.2

Let us recall the notation. We consider a random depth L neural network with layer widths n_0, \dots, n_{L+1} with

$$\exists c, C > 0 \text{ s.t.} \quad cn \leq n_1, \dots, n_L \leq Cn,$$

and a non-linearity σ that satisfies (2.1) for some $r \geq 1$. We also fix $p \geq 1$ directional derivatives d_1, \dots, d_p as in (3.2) and the corresponding vectors of iterated directional derivatives

$$D^{\leq r} z_{i,\mathcal{A}}^{(\ell+1)} := \left(D_\alpha^J z_{i;\alpha}^{(\ell+1)}, \alpha \in \mathcal{A}, J = (j_1, \dots, j_p) \in \mathbb{N}^p, |J| \leq r \right).$$

Theorem 3.2 concerns, for each fixed $m, \ell \geq 1$, the expectation of a function f of the form

$$f \left(D^{\leq r} z_{1,\mathcal{A}}^{(\ell+1)}, \dots, D^{\leq r} z_{m,\mathcal{A}}^{(\ell+1)} \right),$$

which depends on all directional derivatives in d_i of order at most r in any m neuron pre-activations at layer $\ell + 1$. We seek to show that if f is both continuous and a tempered distribution, then for all $q_* \geq 1$ we have

$$\begin{aligned} \mathbb{E} \left[f \left(D^{\leq r} z_{1,\mathcal{A}}^{(\ell+1)}, \dots, D^{\leq r} z_{m,\mathcal{A}}^{(\ell+1)} \right) \right] &= O(n^{-q_*-1}) + \\ &+ \sum_{q=0}^{2q_*} \frac{(-1)^q}{2^q q!} \mathbb{E} \left[\left\langle \left(\sum_{\substack{|J|, |J'| \leq r \\ \alpha, \alpha' \in \mathcal{A}}} \Delta_{\alpha \alpha'}^{JJ', (\ell)} \sum_{j=1}^m \partial_{D_\alpha^J z_{j;\alpha}} \partial_{D_{\alpha'}^{J'} z_{j;\alpha'}} \right)^q f \left(D_{\mathcal{A}}^{\leq r} z_1, \dots, D_{\mathcal{A}}^{\leq r} z_m \right) \right\rangle_{\kappa^{(\ell)}} \right]. \end{aligned} \quad (8.1)$$

We remind the reader the notation in this formula. First, we continue to denote by $\langle \cdot \rangle_{\kappa^{(\ell)}}$ the expectation with respect to a collection of centered jointly Gaussian random vectors

$$D_{\mathcal{A}}^{\leq r} z_i = (D_{\alpha}^J z_{i;\alpha}, \alpha \in \mathcal{A}, |J| \leq r)$$

with the same covariance

$$\text{Cov} (D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}, D_{\alpha_2}^{J_2} z_{i_2;\alpha_2}) = \text{Cov} (D_{\alpha_1}^{J_1} z_{i_1;\alpha_1}^{(\ell)}, D_{\alpha_2}^{J_2} z_{i_2;\alpha_2}^{(\ell)}) = \delta_{i_1 i_2} \kappa_{\alpha_1 \alpha_2}^{J_1 J_2, (\ell)}$$

as the true vectors of derivatives $D_{\mathcal{A}}^{\leq r} z_{i;\mathcal{A}}^{(\ell)}$ in each component separately but zero covariance for different i . Second,

$$\kappa^{(\ell)} = \mathbb{E} \left[\Sigma^{\leq r, (\ell)} \right], \quad \Sigma^{\leq r, (\ell)} = \left(D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha \alpha'}^{(\ell)} \right)_{\substack{|J|, |J'| \leq r \\ \alpha, \alpha' \in \mathcal{A}}}, \quad (8.2)$$

is an average of the conditional covariances

$$D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha \alpha'}^{(\ell)} := \text{Cov} \left(D_{\alpha}^J z_{1;\alpha}^{(\ell+1)}, D_{\alpha'}^{J'} z_{1;\alpha'}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) = D_{\alpha}^J D_{\alpha'}^{J'} \left\{ C_b + \frac{C_W}{n_{\ell}} \sum_{j=1}^{n_{\ell}} \sigma(z_{j;\alpha}^{(\ell)}) \sigma(z_{j;\alpha'}^{(\ell)}) \right\}.$$

Finally, $\Delta_{\alpha \alpha'}^{JJ', (\ell)}$ measures the corresponding fluctuations:

$$\Delta_{\alpha \alpha'}^{JJ', (\ell)} := D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha \alpha'}^{(\ell)} - \mathbb{E} \left[D_{\alpha}^J D_{\alpha'}^{J'} \Sigma_{\alpha \alpha'}^{(\ell)} \right],$$

and we collect $\Delta_{\alpha \alpha'}^{JJ', (\ell)}$ into a matrix as follows:

$$\Delta^{\leq r, (\ell)} := \left(\Delta_{\alpha \alpha'}^{JJ', (\ell)} \right)_{\substack{|J|, |J'| \leq r \\ \alpha, \alpha' \in \mathcal{A}}}.$$

Our first step is to note that since the weights and biases in layer $\ell + 1$ are Gaussian, independent of one another, and independent of the sigma algebra $\mathcal{F}^{(\ell)}$ generated by all prior weights and biases, we may write

$$\mathbb{E} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)}, \dots, D^{\leq r} z_{m;\mathcal{A}}^{(\ell+1)} \right) \right] = \mathbb{E} \left[f \left(\left(\Sigma^{\leq r, (\ell)} \right)^{1/2} Z_1, \dots, \left(\Sigma^{\leq r, (\ell)} \right)^{1/2} Z_m \right) \right],$$

where Z_1, \dots, Z_m are standard Gaussians which are independent of one another and of $\Sigma^{\leq r, (\ell)}$. Moreover, because $\Sigma^{\leq r, (\ell)}$ is PSD the relation (8.2) ensures

$$\ker(\kappa^{(\ell)}) \subseteq \ker(\Sigma^{\leq r, (\ell)}) \text{ a.s.}$$

By decomposing

$$Z_i = Z_{i;\parallel} + Z_{i;\perp}, \quad Z_{i;\parallel} \in \ker(\kappa^{(\ell)}), \quad Z_{i;\perp} \in \ker(\kappa^{(\ell)})^{\perp}$$

and writing $\Sigma_{\perp}^{\leq r, (\ell)}$ for the compression of $\Sigma^{\leq r, (\ell)}$ onto $\ker(\kappa^{(\ell)})^{\perp}$ we obtain by a slight abuse of notation that

$$\mathbb{E} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)}, \dots, D^{\leq r} z_{m;\mathcal{A}}^{(\ell+1)} \right) \right] = \mathbb{E} \left[f \left(\left(\Sigma_{\perp}^{\leq r, (\ell)} \right)^{1/2} Z_{1,\perp}, \dots, \left(\Sigma_{\perp}^{\leq r, (\ell)} \right)^{1/2} Z_{m,\perp} \right) \right].$$

The key point is now that $Z_{i,\perp}$ are standard Gaussian vectors supported on a subspace on which $\kappa^{(\ell)}$ is strictly positive definite and that $\Sigma_{\perp}^{\leq r,(\ell)}$ maps this subspace into itself. Consider the event

$$S_n = \left\{ \left| \Delta_{\alpha\alpha'}^{JJ',(\ell)} \right| < n^{-1/4}, \alpha, \alpha' \in \mathcal{A}, |J|, |J'| \leq r \right\} = \left\{ \left\| \kappa^{(\ell)} - \Sigma^{\leq r,(\ell)} \right\|_{\infty} < n^{-1/4} \right\}.$$

Note that, by applying Theorem 7.3 and arguing exactly as in Lemma 7.5, we find that since $\Delta_{\alpha\alpha'}^{JJ',(\ell)}$ are centered collective observables,

$$\mathbb{P}(S_n^c) = O(n^{-\infty}).$$

Since f is a tempered distribution and a continuous function, its expectation against any Gaussian is finite and we therefore have

$$\mathbb{E} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)}, \dots, D^{\leq r} z_{m;\mathcal{A}}^{(\ell+1)} \right) \right] = \mathbb{E} \left[\mathbf{1}_{S_n} f \left(\left(\Sigma_{\perp}^{\leq r,(\ell)} \right)^{1/2} Z_{1,\perp}, \dots, \left(\Sigma_{\perp}^{\leq r,(\ell)} \right)^{1/2} Z_{m,\perp} \right) \right],$$

plus an error of size $O(n^{-\infty})$. Let us denote by $\widehat{f}(\xi_1, \dots, \xi_m)$ the Fourier transform of f and abbreviate

$$\xi = (\xi_1, \dots, \xi_m), \quad \|\xi\|^2 := \sum_{i=1}^m \|\xi_i\|^2, \quad d\xi := d\xi_1 \cdots d\xi_m$$

For a $C > 0$ that we will choose later let us write

$$\begin{aligned} & \mathbb{E} \left[\mathbf{1}_{S_n} f \left(\left(\Sigma_{\perp}^{\leq r,(\ell)} \right)^{1/2} Z_{1,\perp}, \dots, \left(\Sigma_{\perp}^{\leq r,(\ell)} \right)^{1/2} Z_{m,\perp} \right) \right] \\ &= \int \widehat{f}(\xi) \mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Sigma_{\perp}^{\leq r,(\ell)} \xi_i \right] \right] d\xi \\ &= \int_{\|\xi\|^2 > C \log(n)} \widehat{f}(\xi) \mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Sigma_{\perp}^{\leq r,(\ell)} \xi_i \right] \right] d\xi \\ &+ \int_{\|\xi\|^2 \leq C \log(n)} \widehat{f}(\xi) \mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Sigma_{\perp}^{\leq r,(\ell)} \xi_i \right] \right] d\xi \\ &=: I_C + II_C. \end{aligned}$$

Let us now check that

$$\forall q \geq 1 \exists C = C(q) \text{ s.t. } I_C = O(n^{-q}). \quad (8.3)$$

By the fundamental structure theorem of tempered distributions (see e.g. [?]), there exist bounded continuous function $u_{I,J}$ and an integer $o(f)$, called the order of f , such that

$$\widehat{f}(\xi) = \sum_{\substack{I,J \\ |I,J| \leq o(f)}} \xi^I D^J u_{I,J}(\xi), \quad (8.4)$$

where where the derivatives D^J with respect to ξ_1, \dots, ξ_m are defined in the weak sense and ξ raised to a multi-index I denotes the corresponding monomial. Thus, we may use (8.4) to

write

$$\begin{aligned} |I_C| &= \left| \sum_{\substack{I,J \\ |I|,|J| \leq o(f)}} \int_{\|\xi\|^2 > C \log(n)} u_{I,J}(\xi) D^J \left(\xi^I \mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Sigma_{\perp}^{\leq r,(\ell)} \xi_i \right] \right] \right) d\xi \right| \\ &\leq \sum_{\substack{I,J \\ |I|,|J| \leq o(f)}} \|u_{I,J}\|_{\infty} \int_{\|\xi\|^2 > C \log(n)} \mathbb{E} \left[\mathbf{1}_{S_n} |p_{o(f)}(\xi)| \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Sigma_{\perp}^{\leq r,(\ell)} \xi_i \right] \right] d\xi, \end{aligned}$$

where $p_{o(f)}$ is some polynomial of degree at most $2o(f)$ in the variables ξ_1, \dots, ξ_m in which the coefficients are themselves polynomials the entries of $\Sigma_{\perp}^{\leq r,(\ell)}$. On the event S_n , entries of $\Sigma_{\perp}^{\leq r,(\ell)}$ are uniformly bounded in n since by Theorem 7.3 the entries of $\kappa^{(\ell)}$ are uniformly bounded in n and the event S_n guarantees that the difference $\kappa^{(\ell)} - \Sigma^{\leq r,(\ell)}$ is small for all large n . In particular, for some $T > 0$ we may write

$$\mathbf{1}_{S_n} |p_{o(f)}(\xi_1, \dots, \xi_m)| \leq \mathbf{1}_{S_n} T \left(1 + \|\xi\|^2\right)^{o(f)}.$$

Note moreover that for all n sufficiently large, on the event S_n , we have that for some $\lambda_0 > 0$ and any $\xi \in \ker(\kappa^{(\ell)})^{\perp}$ that

$$\frac{1}{2} \xi^T \Sigma_{\perp}^{\leq r,(\ell)} \xi \geq \lambda_0 \|\xi\|^2.$$

Hence, passing to polar coordinates, we find that

$$I_C \leq T \sum_{\substack{I,J \\ |I|,|J| \leq o(f)}} \|u_{I,J}\|_{\infty} \int_{r^2 > C \log(n)} (1+r^2)^{o(f)+mN(r,p)|\mathcal{A}|-1} e^{-\lambda_0 r^2} dr,$$

where we recall that $N(r,p)$ is the number of derivatives of order at most r in the p vector fields d_1, \dots, d_p . Thus, we conclude that that for any $q \geq 1$ there indeed exists $C = C(q), C' = C'(q)$ such that

$$I_C \leq C' n^{-q},$$

confirming (8.3). We therefore define $C := C(q_* + 1)$ and rewrite II_C as follows:

$$\begin{aligned} II_C &= \int_{\|\xi\|^2 \leq C \log(n)} \widehat{f}(\xi) \mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Sigma_{\perp}^{\leq r,(\ell)} \xi_i \right] \right] d\xi \\ &= \int_{\|\xi\|^2 \leq C \log(n)} \widehat{f}(\xi) \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \kappa^{(\ell)} \xi_i \right] \mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Delta^{\leq r,(\ell)} \xi_i \right] \right] d\xi. \end{aligned}$$

Note that on the event S_n there exists $T > 0$ so that

$$\sup_{\|\xi\|^2 \leq C \log(n)} \sum_{i=1}^m \xi_i^T \Delta_{\mathcal{A}}^{\leq r,(\ell)} \xi_i \leq CTm |\mathcal{A}|^2 \frac{\log(n)}{n^{1/4}}.$$

Hence, we may choose $Q^* = Q^*(q^*, C, |\mathcal{A}|) \geq 1$ so that

$$\mathbb{E} \left[\mathbf{1}_{S_n} \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \Delta_{\mathcal{A}}^{\leq r,(\ell)} \xi_i \right] \right] = \mathbb{E} \left[\mathbf{1}_{S_n} \sum_{q=0}^{Q^*} \frac{(-1)^q}{2^q q!} \left(\sum_{i=1}^m \xi_i^T \Delta_{\mathcal{A}}^{\leq r,(\ell)} \xi_i \right)^q \right] + O(n^{-q^*-1}). \quad (8.5)$$

We thus conclude that II_C equals

$$\sum_{q=0}^{Q^*} \frac{(-1)^q}{2^q q!} \int_{\|\xi\|^2 \leq C \log(n)} \mathbb{E} \left[\left(\sum_{i=1}^m \xi_i^T \Delta_{\mathcal{A}}^{\leq r, (\ell)} \xi_i \right)^q \right] \widehat{f}(\xi) \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \kappa^{(\ell)} \xi_i \right] d\xi$$

plus an error of size $O(n^{-q^*-1})$. Note also that by applying Lemma 7.5 we have

$$\mathbb{E} \left[\mathbf{1}_{S_n} \left(\sum_{i=1}^m \xi_i^T \Delta_{\mathcal{A}}^{(\ell)} \xi_i \right)^q \right] = O \left(\|\xi\|^{2q} n^{-\lceil \frac{q}{2} \rceil} \right).$$

Hence, since \widehat{f} is a tempered distribution and $\kappa^{(\ell)}$ is strictly positive definite on $\ker(\kappa^{(\ell)})^\perp$, the terms corresponding to $2q^* + 1 \leq q \leq Q^*$ in (8.5) are of size $O(n^{-q^*-1})$. Moreover, applying the same reasoning as we used to bound I_C , by incurring another error of order $O(n^{-q^*-1})$ we may drop the restriction in II_C that $\|\xi\|^2 \leq C \sqrt{\log(n)}$. All together, II_C therefore equals

$$\sum_{q=0}^{2q^*} \frac{(-1)^q}{2^q q!} \int \mathbb{E} \left[\left(\sum_{i=1}^m \xi_i^T \Delta_{\mathcal{A}}^{\leq r, (\ell)} \xi_i \right)^q \right] \widehat{f}(\xi) \exp \left[-\frac{1}{2} \sum_{i=1}^m \xi_i^T \kappa^{(\ell)} \xi_i \right] d\xi.$$

plus an error of size $O(n^{-q^*-1})$. Using that multiplication by components of ξ_i acting on the Fourier transform corresponds to differentiation of with respect to the variables $\{D_\alpha^J z_{i;\alpha}, \alpha \in \mathcal{A}, |J| \leq r\}$, yields the desired expression (8.1) and completes the proof of Theorem 3.2. \square

9 Proof of Corollary 3.4

The goal of this section is to derive recursions for

$$\kappa_{2k;\alpha}^{(\ell+1)} = \kappa_k \left(\underbrace{\Delta_{\alpha\alpha}^{(\ell)}, \dots, \Delta_{\alpha\alpha}^{(\ell)}}_{k \text{ times}} \right),$$

where we defined $\Delta_{\alpha\alpha}^{(\ell)}$ in (3.10). Let us write

$$X_j := \sigma \left(z_{j;\alpha}^{(\ell)} \right)^2 - \mathbb{E} \left[\sigma \left(z_{j;\alpha}^{(\ell)} \right)^2 \right]$$

so that

$$\Delta_{\alpha\alpha}^{(\ell)} = \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} X_j.$$

By symmetry, we then have

$$\kappa_{4;\alpha}^{(\ell+1)} = \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right] = \frac{C_W^2}{n_\ell} \mathbb{E} [X_1^2] + C_W^2 (1 - n_\ell^{-1}) \mathbb{E} [X_1 X_2] \quad (9.1)$$

$$\begin{aligned} \kappa_{6;\alpha}^{(\ell+1)} &= \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^3 \right] \\ &= \frac{C_W^3}{n_\ell^2} \mathbb{E} [X_1^3] + \frac{3C_W^3}{n_\ell} \left(1 - \frac{1}{n_\ell} \right) \mathbb{E} [X_1^2 X_2] + C_W^3 \left(1 - \frac{1}{n_\ell} \right) \left(1 - \frac{2}{n_\ell} \right) \mathbb{E} [X_1 X_2 X_3] \end{aligned} \quad (9.2)$$

$$\begin{aligned} \kappa_{8;\alpha}^{(\ell+1)} &= \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^4 \right] - 3\mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right]^2 = \frac{C_W^4}{n_\ell^3} \left(\mathbb{E} [X_1^4] - 3\mathbb{E} [X_1^2]^2 \right) \\ &\quad + \frac{C_W^4}{n_\ell^2} \left(1 - \frac{1}{n_\ell} \right) \left[\binom{4}{2} \left\{ \mathbb{E} [X_1^2 X_2]^2 - \mathbb{E} [X_1^2] \mathbb{E} [X_2^2] - 2\mathbb{E} [X_1 X_2]^2 \right\} \right. \\ &\quad \quad \left. + \binom{4}{1} \left\{ \mathbb{E} [X_1^3 X_2] - \mathbb{E} [X_1^3] \mathbb{E} [X_2] - 2\mathbb{E} [X_1^2] \mathbb{E} [X_1 X_2] \right\} \right] \\ &\quad + \frac{C_W^4}{n_\ell} \left(1 - \frac{1}{n_\ell} \right) \left(1 - \frac{2}{n_\ell} \right) \binom{4}{2} \left\{ \mathbb{E} [X_1^2 X_2 X_3] - \mathbb{E} [X_1^2] \mathbb{E} [X_1 X_2] - 2\mathbb{E} [X_1 X_2]^2 \right\} \\ &\quad + C_W^4 \left(1 - \frac{1}{n_\ell} \right) \left(1 - \frac{2}{n_\ell} \right) \left(1 - \frac{3}{n_\ell} \right) \mathbb{E} [X_1 X_2 X_3 X_4]. \end{aligned} \quad (9.3)$$

To evaluate the mixed moments of X_i that appear in (9.1)-(9.3), we use Theorem 3.2 in the case $g \equiv 1, r = 0, q_* = 1$. In this setting, if f is a continuous function and a tempered distribution, we find

$$\begin{aligned} \mathbb{E} \left[f \left(z_{1;\alpha}^{(\ell)}, \dots, z_{m;\alpha}^{(\ell)} \right) \right] &= \langle f(z_{1;\alpha}, \dots, z_{m;\alpha}) \rangle_{\kappa^{(\ell)}} \quad (9.4) \\ &\quad + \frac{\kappa_{4;\alpha}^{(\ell)}}{2^2 \cdot 2!} \left\langle \left(\sum_{j=1}^m \partial_{z_{j;\alpha}}^2 \right)^2 f(z_{1;\alpha}, \dots, z_{m;\alpha}) \right\rangle_{\kappa^{(\ell)}} \\ &\quad + \frac{\kappa_{6;\alpha}^{(\ell)}}{2^3 \cdot 3!} \left\langle \left(\sum_{j=1}^m \partial_{z_{j;\alpha}}^2 \right)^3 f(z_{1;\alpha}, \dots, z_{m;\alpha}) \right\rangle_{\kappa^{(\ell)}} \\ &\quad + \frac{\kappa_{8;\alpha}^{(\ell)} + 3 \left(\kappa_{4;\alpha}^{(\ell)} \right)^2}{2^4 \cdot 4!} \left\langle \left(\sum_{j=1}^m \partial_{z_{j;\alpha}}^2 \right)^4 f(z_{1;\alpha}, \dots, z_{m;\alpha}) \right\rangle_{\kappa^{(\ell)}} + O(n^{-4}). \end{aligned}$$

We remind the reader that, by definition, $z_{1;\alpha}, \dots, z_{m;\alpha}$ are iid centered Gaussians with variance $\kappa_{\alpha\alpha}^{(\ell)}$. Since the derivations of (3.15)-(3.17) are very similar, let us give the details for only cases of $\kappa_{4;\alpha}^{(\ell)}$ and $\kappa_{6;\alpha}^{(\ell)}$. We have, using (9.4), that

$$\begin{aligned} \kappa_{\alpha\alpha}^{(\ell+1)} &= \mathbb{E} \left[\left(\Delta_{\alpha\alpha}^{(\ell)} \right)^2 \right] = \frac{C_W^2}{n_\ell} \left(\langle \sigma^4 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}} - \langle \sigma^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}}^2 \right) \\ &\quad + C_W^2 (1 - n_\ell^{-1}) \left(\langle X_1 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}}^2 + \frac{1}{4} \langle \partial^2 \sigma^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}}^2 \kappa_{4;\alpha}^{(\ell)} \right) + O(n^{-2}). \end{aligned} \quad (9.5)$$

Next, we will show in §10.4 below that

$$\kappa_{\alpha\alpha}^{(\ell)} = K_{\alpha\alpha}^{(\ell)} + O(n^{-1}).$$

Moreover, note that since $x_\alpha \neq 0$, we have $K_{\alpha\alpha}^{(\ell)}$ is non-zero. Hence, Gaussian integration by parts yields that for any measurable polynomially bounded f we have

$$\langle f \rangle_{\kappa_{\alpha\alpha}^{(\ell)}} = \langle f \rangle_{K_{\alpha\alpha}^{(\ell)}} + O(n^{-1}). \quad (9.6)$$

Note also that for all i by applying (9.4) we have

$$\langle X_i \rangle_{\kappa^{(\ell)}} = -\frac{1}{8} \kappa_{4;\alpha}^{(\ell)} \langle \partial^4 X_1 \rangle_{K^{(\ell)}} + O(n^{-2}). \quad (9.7)$$

Thus, recalling the definition (4.3) of $\chi_{||;\alpha}^{(\ell)}$ the estimate (9.5) immediately yields (3.15). Next, using (9.4) as well as (9.6) and (9.7) that

$$C_W^3 \mathbb{E} [X_1^3] = C_W^3 \langle X_1^3 \rangle_{K_{\alpha\alpha}^{(\ell)}} + O(n^{-1}) = T_{0,3;\alpha}^{(\ell)} + O(n^{-1}).$$

Further, we seek to evaluate $\mathbb{E} [X_1^2 X_2]$ up to errors of size $O(n^{-2})$. We apply (9.4) as well as (9.6) and (9.7) to obtain

$$\begin{aligned} C_W^3 \mathbb{E} [X_1^2 X_2] &= C_W^3 \langle X_1^2 \rangle_{\kappa^{(\ell)}} \langle X_2 \rangle_{\kappa^{(\ell)}} + O(n^{-2}) \\ &+ \frac{C_W^3}{8} \kappa_{4;\alpha}^{(\ell)} [\langle X_1^2 \rangle_{K^{(\ell)}} \langle \partial^4 X_2 \rangle_{K^{(\ell)}} + 2 \langle \partial^2 X_1^2 \rangle_{K^{(\ell)}} \langle \partial^2 X_2 \rangle_{K^{(\ell)}}] \\ &+ \frac{1}{2} T_{2,2;\alpha}^{(\ell)} \chi_{||;\alpha}^{(\ell)} \kappa_{4;\alpha}^{(\ell)} + O(n^{-2}). \end{aligned}$$

Finally, we must evaluate $\mathbb{E} [X_1 X_2 X_3]$ up to errors of size $O(n^{-3})$. Again using (9.7), we find

$$\begin{aligned} C_W^3 \mathbb{E} [X_1 X_2 X_3] &= C_W^3 \langle X_1 \rangle_{\kappa^{(\ell)}} \langle X_2 \rangle_{\kappa^{(\ell)}} \langle X_3 \rangle_{\kappa^{(\ell)}} \\ &+ \frac{C_W^3}{8} \kappa_{4;\alpha}^{(\ell)} \left[6 \langle X_1 \rangle_{\kappa^{(\ell)}} \langle \partial^2 X_2 \rangle_{\kappa^{(\ell)}}^2 + O(n^{-2}) \right] \\ &+ \frac{C_W^3}{48} \kappa_{6;\alpha}^{(\ell)} \left[6 \langle \partial^2 X_1 \rangle_{\kappa^{(\ell)}}^3 + O(n^{-1}) \right] + O(n^{-3}) \\ &= -\frac{3}{8} T_{4,1;\alpha}^{(\ell)} \left(\chi_{||;\alpha}^{(\ell)} \kappa_{4;\alpha}^{(\ell)} \right)^2 \langle \partial^4 X_1 \rangle_{K^{(\ell)}} + \left(\chi_{||;\alpha}^{(\ell)} \right)^3 \kappa_{6;\alpha}^{(\ell)} + O(n^{-3}). \end{aligned}$$

This completes the derivation of the recursion for $\kappa_{6;\alpha}^{(\ell)}$. \square

10 Recursions for 2nd, 4th Cumulants of $z_{i;\alpha}^{(\ell)}$ and its Derivatives

In this section, we consider a random depth L fully connected network with input dimension n_0 , hidden layer widths, n_1, \dots, n_L , and non-linearity σ satisfying Assumption (2.1). We also fix a network input $x_\alpha \in \mathbb{R}^{n_0}$. Our goal in this section is to give a recursive description of the second and fourth cumulants of $z_{i;\alpha}^{(\ell+1)}$ and its derivatives in terms of those of $z_{i;\alpha}^{(\ell)}$. The starting point is to establish recursions for

$$K_{(ij)}^{(\ell)} := \partial_{x_{i;\alpha_1}} \partial_{x_{j;\alpha_2}} K_{\alpha_1 \alpha_2}^{(\ell)} \Big|_{x_{\alpha_1} = x_{\alpha_2} = x_\alpha}$$

where $i, j = 1, \dots, n_0$ (note that $K_{(ij)}^{(\ell)} = K_{\alpha\alpha}^{(\ell)}$). This is done in Lemma 11.1 in §10.2. We then proceed in §10.3 to detail a number of recursions for the fourth cumulants of $z_{i;\alpha}^{(\ell)}$ and its first derivatives. Finally, §10.4 provides recursions to first order in $1/n$ for the finite width correction of the variance of the covariance of $z_{i;\alpha}^{(\ell)}$ and its first derivative. All these recursions hold for any σ satisfying (2.1). In §11.2 we will solve these recursions for σ belonging to the $K_* = 0$ universality class defined in §4 and will then use these solutions to prove Theorem 4.5.

10.1 Notation

We will abbreviate

$$d_i z_{i;\alpha}^{(\ell)} := \partial_{x_{i;\alpha}} z_{i;\alpha}^{(\ell)}, \quad i = 1, \dots, n_0.$$

Moreover, to keep the notation as uniform as possible, we will also write

$$d_0 := \text{identity} \tag{10.1}$$

so that $d_0 z_{i;\alpha}^{(\ell)} = z_{i;\alpha}^{(\ell)}$. With this notation, let us agree to denote by

$$\kappa_{(j_1 j_2)}^{(\ell)} := \text{Cov} \left(d_{j_1} z_{i;\alpha}^{(\ell)}, d_{j_2} z_{i;\alpha}^{(\ell)} \right) \tag{10.2}$$

the finite width covariance of $z_{i;\alpha}^{(\ell)}$ and its derivatives and by

$$K_{(j_1 j_2)}^{(\ell)} := \lim_{n \rightarrow \infty} \kappa_{(j_1 j_2)}^{(\ell)}$$

the corresponding infinite width covariance. Note that with this notation we have $K_{(00)}^{(\ell)} = K_{\alpha\alpha}^{(\ell)}$. We also write

$$S_{(ij)}^{(\ell)} := K_{(ij)}^{(\ell)} - \kappa_{(ij)}^{(\ell)}. \tag{10.3}$$

Further, for a function f we will write

$$\langle f \rangle_{(\ell)} := \langle f(d_j z_i), \quad i = 1, \dots, m, j = 0, 1, \dots, n_0 \rangle_{K^{(\ell)}}, \tag{10.4}$$

for the expectation of a function f with respect to the centered Gaussian satisfying

$$\text{Cov}(d_{j_1} z_{i_1}, d_{j_2} z_{i_2}) = \delta_{i_1 i_2} K_{(j_1 j_2)}^{(\ell)}, \quad j_1, j_2 = 0, 1, 2.$$

For the sake of recording more manageable formulas, we will often drop the argument z in expressions such as $\sigma(z)$ so that

$$\langle \sigma \sigma' d_1 z (d_2 z)^2 \rangle_{(\ell)} = \langle \sigma(z) \sigma'(z) d_1 z (d_2 z)^2 \rangle_{(\ell)}$$

as so on. Note that in the expectation (10.4) variables such as z_{i_1}, z_{i_2} with different neural indices are independent. Thus, for example

$$\langle \sigma \sigma' \rangle_{(\ell)} = \langle \sigma(z) \sigma'(z) \rangle_{(\ell)} = \int_{\mathbb{R}} \sigma \left(z (K_{\alpha\alpha}^{(\ell)})^{1/2} \right) \sigma' \left(z (K_{\alpha\alpha}^{(\ell)})^{1/2} \right) e^{-z^2/2} \frac{dz}{\sqrt{2\pi}}.$$

In analogy with (10.4), we will denote by $\langle f(d_j z_i, i = 1, \dots, m, j = 0, 1, \dots, n_0) \rangle_{\kappa^{(\ell)}}$ the expectation of f with respect to a centered Gaussian satisfying

$$\text{Cov}(d_{j_1} z_{i_1}, d_{j_2} z_{i_2}) = \delta_{i_1 i_2} \kappa_{(j_1 j_2)}^{(\ell)}, \quad j_1, j_2 = 0, 1, 2.$$

in which all neurons are again independent centered Gaussians but this time with the finite width covariance. Further, when f depends only a single variable $z = d_0 z_i$, we will denote by ∂ derivatives with respect to z , so that for any function $f(z)$ we have

$$\partial^i f = \partial^i f(z) = \frac{\partial^i}{\partial z^i} f(z).$$

Here and elsewhere all derivatives are in the weak sense if f is not sufficiently smooth. Finally, as in Corollary 3.4 and §4.1, we also set

$$\begin{aligned}\chi_{\parallel}^{(\ell)} &:= \chi_{\parallel;\alpha}^{(\ell)} = \chi_{\parallel}^{(\ell)}(K_{\alpha\alpha}^{(\ell)}) = \frac{C_W}{2} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \\ \chi_{\perp}^{(\ell)} &:= \chi_{\perp;\alpha}^{(\ell)} = \chi_{\perp}^{(\ell)}(K_{\alpha\alpha}^{(\ell)}) = C_W \langle (\sigma')^2 \rangle_{(\ell)}.\end{aligned}$$

10.2 Recursions at Infinite Width

The following Lemma gives recursions with respect to ℓ for the infinite width covariance $K_{(ij)}^{(\ell)}$.

Lemma 10.1. *We have*

$$K_{(00)}^{(\ell+1)} = C_b + C_W \langle \sigma^2 \rangle_{(\ell)} \quad (10.5)$$

$$K_{(10)}^{(\ell+1)} = \chi_{\parallel}^{(\ell)} K_{(10)}^{(\ell)} \quad (10.6)$$

$$K_{(11)}^{(\ell+1)} = C_W \langle \partial^2 (\sigma')^2 \rangle_{(\ell)} \left(K_{(10)}^{(\ell)} \right)^2 + \chi_{\perp}^{(\ell)} K_{(11)}^{(\ell)} \quad (10.7)$$

$$K_{(12)}^{(\ell+1)} = C_W \langle \partial^2 (\sigma')^2 \rangle_{(\ell)} K_{(10)}^{(\ell)} K_{(20)}^{(\ell)} + \chi_{\perp}^{(\ell)} K_{(12)}^{(\ell)}. \quad (10.8)$$

Proof. The proof of these recursions consists of integrating out the weights and biases in layer ℓ and sometimes using Gaussian integration by parts. For instance,

$$K_{(00)}^{(\ell+1)} = \lim_{n \rightarrow \infty} \mathbb{E} \left[\left(z_{i;\alpha}^{(\ell+1)} \right)^2 \right] = \lim_{n \rightarrow \infty} \left(C_b + C_W \mathbb{E} \left[\sigma \left(z_{i;\alpha}^{(\ell)} \right)^2 \right] \right) = C_b + C_W \langle \sigma^2 \rangle_{(\ell)},$$

gives (10.5), which is a restatement of the recursion (2.4) we already say in Theorem 2.2. Further, using the definition of $K_{(10)}^{(\ell+1)}$, Theorem 2.2, and integrating by parts we have

$$\begin{aligned}K_{(10)}^{(\ell+1)} &= \lim_{n \rightarrow \infty} \mathbb{E} \left[z_{i;\alpha}^{(\ell+1)} d_1 z_{i;\alpha}^{(\ell+1)} \right] \\ &= \lim_{n \rightarrow \infty} C_W \mathbb{E} \left[\sigma \left(z_{i;\alpha}^{(\ell)} \right) \sigma' \left(z_{i;\alpha}^{(\ell)} \right) d_1 z_{i;\alpha}^{(\ell)} \right] \\ &= C_W \langle \sigma(z) \sigma'(z) d_1 z \rangle_{(\ell)} \\ &= C_W K_{(10)}^{(\ell)} \langle \sigma(z) \sigma'(z) \rangle_{(\ell)} \\ &= \chi_{\parallel}^{(\ell)} K_{(10)}^{(\ell)}.\end{aligned}$$

This proves (10.6). Similar reasoning gives

$$\begin{aligned}K_{(11)}^{(\ell+1)} &= C_W \langle (\sigma'(z) d_1 z)^2 \rangle_{(\ell)} \\ &= C_W \left[K_{(11)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} + \left(K_{(10)}^{(\ell)} \right)^2 \langle \partial^2 (\sigma')^2 \rangle_{(\ell)} \right] \\ &= C_W \langle \partial^2 (\sigma')^2 \rangle_{(\ell)} \left(K_{(10)}^{(\ell)} \right)^2 + \chi_{\perp}^{(\ell)} K_{(11)}^{(\ell)},\end{aligned}$$

proving (10.7). The derivation of (10.8) is analogous. \square

10.3 4th Cumulant Recursions at Finite Width

In this section we obtain a recursive description for the fourth cumulants of $z_{i;\alpha}^{(\ell)}$ and its first derivatives. In keep with our convention (10.1), we will denote by $\partial_{x_{0;\alpha}}$ the identity operator. By explicitly integrating out the weights and biases in layer $\ell + 1$ we find

$$\begin{aligned} & \kappa \left(\partial_{x_{j_1;\alpha}} z_{i_1;\alpha}^{(\ell+1)}, \partial_{x_{j_2;\alpha}} z_{i_2;\alpha}^{(\ell+1)}, \partial_{x_{j_3;\alpha}} z_{i_3;\alpha}^{(\ell+1)}, \partial_{x_{j_4;\alpha}} z_{i_4;\alpha}^{(\ell+1)} \right) \\ &= \delta_{i_1 i_2} \delta_{i_3 i_4} \kappa_{(j_1 j_2)(j_3 j_4)}^{(\ell)} + \delta_{i_1 i_3} \delta_{i_2 i_4} \kappa_{(j_1 j_3)(j_2 j_4)}^{(\ell)} + \delta_{i_1 i_4} \delta_{i_2 i_3} \kappa_{(j_1 j_4)(j_2 j_3)}^{(\ell)}, \end{aligned} \quad (10.9)$$

where

$$\begin{aligned} \kappa_{(j_1 j_2)(j'_1 j'_2)}^{(\ell)} &= \kappa \left(\partial_{x_{j_1;\alpha}} z_{i_1;\alpha}^{(\ell+1)}, \partial_{x_{j_2;\alpha}} z_{i_1;\alpha}^{(\ell+1)}, \partial_{x_{j'_1;\alpha}} z_{i_2;\alpha}^{(\ell+1)}, \partial_{x_{j'_2;\alpha}} z_{i_2;\alpha}^{(\ell+1)} \right) \\ &= \text{Cov} \left(\partial_{x_{j_1;\alpha_1}} \partial_{x_{j_2;\alpha_2}} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} \Big|_{\alpha_1 = \alpha_2 = \alpha}, \partial_{x_{j'_1;\alpha_1}} \partial_{x_{j'_2;\alpha_2}} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} \Big|_{\alpha_1 = \alpha_2 = \alpha} \right). \end{aligned} \quad (10.10)$$

We remind the reader that the quantity $\Sigma_{\alpha_1 \alpha_2}^{(\ell)}$ appearing in (10.10) is the conditional covariance

$$\text{Cov} \left(z_{i_1;\alpha_1}^{(\ell+1)}, z_{i_2;\alpha_2}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) = \delta_{i_1 i_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} = \delta_{i_1 i_2} \left(C_b + \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \sigma \left(z_{j;\alpha_1}^{(\ell)} \right) \sigma \left(z_{j;\alpha_2}^{(\ell)} \right) \right)$$

that we introduced in (3.5). Our next result gives recursions with respect to ℓ directly for the covariances $\kappa_{(j_1 j_2)(j'_1 j'_2)}^{(\ell)}$ and hence via (10.9) for the fourth cumulants of $z_{i;\alpha}^{(\ell)}$ and its first derivatives as well.

Proposition 10.2. *We have*

$$\kappa_{(00)(00)}^{(\ell+1)} = \frac{C_W^2}{n_\ell} \left(\langle \sigma^4 \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)}^2 \right) + \left(\chi_{||}^{(\ell)} \right)^2 \kappa_{(00)(00)}^{(\ell)} + O(n^{-2}).$$

Further,

$$\begin{aligned} \kappa_{(10)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left(\langle \sigma^3 \sigma' d_1 z \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)} \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \right) \\ &+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \right] + O(n^{-2}). \end{aligned} \quad (10.11)$$

Moreover,

$$\begin{aligned} \kappa_{(10)(10)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma \sigma' d_1 z)^2 \rangle_{(\ell)} - \langle \sigma \sigma' d_1 z \rangle_{(\ell)}^2 \right] \\ &+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)}^2 + 8\kappa_{(10)(00)}^{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} \right. \\ &\quad \left. + 8\kappa_{(10)(10)}^{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)}^2 \right] + O(n^{-2}). \end{aligned}$$

Similarly,

$$\begin{aligned}\kappa_{(10)(20)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} - \langle \sigma\sigma' d_1 z \rangle_{(\ell)} \langle \sigma\sigma' d_2 z \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2(\sigma\sigma') d_1 z \rangle_{(\ell)} \langle \partial^2(\sigma\sigma') d_2 z \rangle_{(\ell)} \right. \\ &\quad + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial(\sigma\sigma') \rangle_{(\ell)} \langle \partial^2(\sigma\sigma') d_2 z \rangle_{(\ell)} + 4\kappa_{(20)(00)}^{(\ell)} \langle \partial(\sigma\sigma') \rangle_{(\ell)} \langle \partial^2(\sigma\sigma') d_1 z \rangle_{(\ell)} \\ &\quad \left. + 8\kappa_{(10)(20)}^{(\ell)} \langle \partial(\sigma\sigma') \rangle_{(\ell)}^2 \right] + O(n^{-2}).\end{aligned}$$

In addition,

$$\begin{aligned}\kappa_{(11)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma\sigma' d_1 z)^2 \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)} \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + 8\kappa_{(10)(00)}^{(\ell)} \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right. \\ &\quad \left. + 4\kappa_{(11)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right] + O(n^{-2}).\end{aligned}$$

Also,

$$\begin{aligned}\kappa_{(12)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^2 \sigma^2 d_1 z d_2 z \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \sigma^2 \rangle_{(\ell)} \right] \\ &= \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right. \\ &\quad + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial(\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} + 4\kappa_{(20)(00)}^{(\ell)} \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \\ &\quad \left. + 4\kappa_{(12)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right] + O(n^{-2}).\end{aligned}$$

Additionally,

$$\begin{aligned}\kappa_{(11)(10)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle \sigma(\sigma' d_1 z)^3 \rangle_{(\ell)} - \langle \sigma\sigma' d_1 z \rangle_{(\ell)} \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2(\sigma\sigma') d_1 z \rangle_{(\ell)} \langle \partial^2(\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \right. \\ &\quad + 4\kappa_{(10)(00)}^{(\ell)} \left(\langle \partial(\sigma\sigma') \rangle_{(\ell)} \langle \partial^2(\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + 2 \langle \partial^2(\sigma\sigma') d_1 z \rangle_{(\ell)} \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} \right) \\ &\quad + 16\kappa_{(10)(10)}^{(\ell)} \langle \partial(\sigma\sigma') \rangle_{(\ell)} \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} + 4\kappa_{(11)(00)}^{(\ell)} \langle \partial^2(\sigma\sigma') d_1 z \rangle_{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \\ &\quad \left. + 8\kappa_{(11)(10)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial(\sigma\sigma') \rangle_{(\ell)} \right] + O(n^{-2}).\end{aligned}$$

Also,

$$\begin{aligned}
\kappa_{(12)(10)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^3 \sigma (d_1 z)^2 d_2 z \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \right] \\
&= \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} \right. \\
&\quad + 4\kappa_{(10)(00)}^{(\ell)} \left\{ \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} + \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \right\} \\
&\quad + 4\kappa_{(20)(00)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} + 4\kappa_{(12)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} \\
&\quad + 8\kappa_{(10)(10)}^{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} + 8\kappa_{(10)(20)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \\
&\quad \left. + 8\kappa_{(12)(10)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \right] + O(n^{-2}).
\end{aligned}$$

Further,

$$\begin{aligned}
\kappa_{(11)(20)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^2 \sigma (d_1 z)^2 d_2 z \rangle_{(\ell)} - \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \langle \sigma \sigma' d_2 z \rangle_{(\ell)} \right] \\
&\quad + \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} \right. \\
&\quad + 8\kappa_{(10)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} + 4\kappa_{(20)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \\
&\quad + 4\kappa_{(11)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} + 16\kappa_{(10)(20)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \\
&\quad \left. + 8\kappa_{(11)(20)}^{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \right] + O(n^{-2}).
\end{aligned}$$

Additionally,

$$\begin{aligned}
\kappa_{(11)(11)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma' dz)^4 \rangle_{(\ell)} - \langle (\sigma' dz)^2 \rangle_{(\ell)}^2 \right] \\
&\quad + \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)}^2 + 16\kappa_{(10)(00)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \right. \\
&\quad + 8\kappa_{(11)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + 32\kappa_{(10)(10)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)}^2 \\
&\quad \left. + 32\kappa_{(11)(10)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} + 8\kappa_{(11)(11)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)}^2 \right].
\end{aligned}$$

Also,

$$\begin{aligned}
\kappa_{(11)(22)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^4 (d_1 z d_2 z)^2 \rangle_{(\ell)} - \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \langle (\sigma' d_2 z)^2 \rangle_{(\ell)} \right] \\
&\quad + \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_2 z)^2 \rangle_{(\ell)} \right. \\
&\quad + 8\kappa_{(10)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_2 z)^2 \rangle_{(\ell)} + 8\kappa_{(20)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \\
&\quad + 4\kappa_{(11)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_2 z)^2 \rangle_{(\ell)} + 4\kappa_{(22)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma') (d_1 z)^2 \rangle_{(\ell)} \\
&\quad + 32\kappa_{(10)(20)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \\
&\quad + 16\kappa_{(11)(20)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} + 16\kappa_{(22)(10)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \\
&\quad \left. + 8\kappa_{(11)(22)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)}^2 \right] + O(n^{-2}).
\end{aligned}$$

Finally,

$$\begin{aligned}
\kappa_{(12)(12)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^4 (d_1 z d_2 z)^2 \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)}^2 \right] \\
&+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)}^2 \right. \\
&\quad + 8\kappa_{(10)(00)}^{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} + 8\kappa_{(20)(00)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \\
&\quad + 8\kappa_{(12)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} + 26\kappa_{(10)(20)}^{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \\
&\quad + 8\kappa_{(10)(10)}^{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)}^2 + 8\kappa_{(20)(20)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)}^2 \\
&\quad + 16\kappa_{(12)(10)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} + 16\kappa_{(12)(20)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \\
&\quad \left. + 8\kappa_{(12)(12)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)}^2 \right] + O(n^{-2}).
\end{aligned}$$

Proof. The derivations of the recursions in Proposition 10.2 are all similar. The idea is to apply Theorem 3.2 in the special case of $q_* = 1$, $r = 1$, a single network input x_α , and directional derivatives

$$d_j = \partial_{x_{j;\alpha}}, \quad j = 1, 2.$$

To keep the notation as uniform as possible, we will continue to write

$$d_0 = \text{identity}$$

for the identity operator. We will have occasion to apply the result to compute expressions of the form

$$\mathbb{E} \left[f \left(d_j z_{i;\alpha}^{(\ell)}, \quad i = 1, \dots, m, \quad j = 0, 1, 2 \right) \right], \quad (10.12)$$

where the function f depends on the pre-activations $z_{i;\alpha}^{(\ell)}$, $i = 1, \dots, m$ at finitely many neurons and is a polynomial of total degree between 0 and 4 in the derivatives $d_j z_{i;\alpha}^{(\ell)}$. Theorem 3.2 allows us to expand (10.12) into a perturbative series in powers of $1/n$ with coefficients that are Gaussian integrals of the form

$$\langle \mathcal{P} f(d_j z_i, i = 1, \dots, m, j = 1, 2) \rangle_{\kappa^{(\ell)}},$$

where \mathcal{P} is a differential operator consisting of finite sums of partial derivatives with respect to the variables $d_j z_i$, and $\langle \cdot \rangle_{\kappa^{(\ell)}}$ denotes an expectation in which $d_j z_{i;\alpha}$ are centered Gaussian with covariance

$$\delta_{i_1 i_2} \kappa_{(j_1 j_2)}^{(\ell)} := \text{Cov} \left(d_{j_1} z_{i_1}, d_{j_2} z_{i_2} \right) := \delta_{i_1 i_2} \text{Cov} \left(d_{j_1} z_{i_1;\alpha}^{(\ell)}, d_{j_2} z_{i_2;\alpha}^{(\ell)} \right), \quad j_1, j_2 = 0, 1, 2.$$

Specifically, taking $g \equiv 1$ in Theorem 3.2 we find that (10.12) equals

$$\langle f \rangle_{\kappa^{(\ell)}} + \frac{1}{8} \sum_{i,i'=1}^m \sum_{j_1, j_2, j'_1, j'_2=0}^2 \kappa_{(j_1 j_2)(j'_1 j'_2)}^{(\ell)} \left\langle \frac{\partial^2}{\partial d_{j_1} z_i \partial d_{j_2} z_i} \frac{\partial^2}{\partial d_{j'_1} z_{i'} \partial d_{j'_2} z_{i'}} f \right\rangle_{K^{(\ell)}} + O(n^{-2})$$

Note that in the expectations inside the sum the integrals are with respect to a centered Gaussian with the infinite width covariance $K^{(\ell)}$ defined in §10.1 instead of the finite width covariance $\kappa^{(\ell)}$. The reason is that, as we'll show in §10.4, we have $\kappa^{(\ell)} = K^{(\ell)} + O(n^{-1})$ and

since $\kappa_{(00)(00)}^{(\ell)} = O(n^{-1})$ by Theorem 3.1, the replacement of $\kappa^{(\ell)}$ by $K^{(\ell)}$ incurs an error of size $O(n^{-2})$, which we are neglecting anyway.

With that said, note that if f is a polynomial of degree 0 in the derivatives $d_j z_{i;\alpha}^{(\ell)}$ for $j = 1, 2$ (i.e. doesn't depends on them), then we find that (10.12) equals

$$\langle f \rangle_{\kappa^{(\ell)}} + \frac{1}{8} \sum_{i,i'=1}^m \kappa_{(00)(00)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial z_i^2} \frac{\partial^2}{\partial z_{i'}^2} f \right\} \right\rangle_{K^{(\ell)}} + O(n^{-2}). \quad (10.13)$$

If instead f is a polynomial of degree 1 in the derivatives $d_j z_{i;\alpha}^{(\ell)}$ then (10.12) equals (10.13) plus the expression

$$\frac{1}{8} \sum_{i,i'=1}^m 2\kappa_{(10)(00)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial z_i \partial d_1 z_i} \frac{\partial^2}{\partial z_{i'}^2} + \frac{\partial^2}{\partial z_{i'} \partial d_1 z_{i'}} \frac{\partial^2}{\partial z_i^2} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2), \quad (10.14)$$

where $(1 \leftrightarrow 2)$ means the same expression with d_1 exchanged with d_2 , $\kappa_{(10)(00)}^{(\ell)}$ replaced by $\kappa_{(20)(00)}^{(\ell)}$, and so on. If f is a polynomial of degree 2 in the derivatives $d_j z_{i;\alpha}^{(\ell)}$ then we must add to (10.13) and (10.14) in addition

$$\begin{aligned} & \frac{1}{8} \sum_{i,i'=1}^m \left[\kappa_{(11)(00)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i^2} \frac{\partial^2}{\partial z_{i'}^2} + \frac{\partial^2}{\partial d_1 z_{i'}^2} \frac{\partial^2}{\partial z_i^2} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \right. \\ & + \\ & + 4\kappa_{(10)(10)}^{(\ell)} \left\langle \frac{\partial^2}{\partial z_i \partial d_1 z_i} \frac{\partial^2}{\partial z_{i'} \partial d_1 z_{i'}} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \\ & \left. + 4\kappa_{(10)(20)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial z_i \partial d_1 z_i} \frac{\partial^2}{\partial z_{i'} \partial d_2 z_{i'}} + \frac{\partial^2}{\partial z_i \partial d_2 z_i} \frac{\partial^2}{\partial z_{i'} \partial d_1 z_{i'}} \right\} f \right\rangle_{K^{(\ell)}} \right] \end{aligned} \quad (10.15)$$

Similarly, if f is a polynomial of degree 3 in the derivatives $d_j z_{i;\alpha}^{(\ell)}$ then we must add to (10.13)-(10.15) in also

$$\begin{aligned} & \frac{1}{8} \sum_{i,i'=1}^m \left[2\kappa_{(11)(10)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i^2} \frac{\partial^2}{\partial z_{i'} \partial d_1 z_{i'}} + \frac{\partial^2}{\partial d_1 z_{i'}^2} \frac{\partial^2}{\partial z_i \partial d_1 z_i} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \right. \\ & + 2\kappa_{(11)(20)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i^2} \frac{\partial^2}{\partial z_{i'} \partial d_2 z_{i'}} + \frac{\partial^2}{\partial d_1 z_{i'}^2} \frac{\partial^2}{\partial z_i \partial d_2 z_i} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \\ & \left. + 4\kappa_{(12)(10)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i \partial d_2 z_i} \frac{\partial^2}{\partial z_{i'} \partial d_1 z_{i'}} + \frac{\partial^2}{\partial d_1 z_{i'} \partial d_2 z_{i'}} \frac{\partial^2}{\partial z_i \partial d_1 z_i} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \right]. \end{aligned} \quad (10.16)$$

Finally, if f is a polynomial of degree 4 in the derivatives $d_j z_{i;\alpha}^{(\ell)}$ then we must add to (10.13)-

(10.16) also

$$\begin{aligned}
& \frac{1}{8} \sum_{i,i'=1}^m \left[\kappa_{(11)(11)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i^2} \frac{\partial^2}{\partial d_1 z_{i'}^2} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \right. \\
& + 2\kappa_{(11)(12)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i^2} \frac{\partial^2}{\partial d_1 z_{i'} \partial d_2 z_{i'}} + \frac{\partial^2}{\partial d_1 z_{i'}^2} \frac{\partial^2}{\partial d_1 z_i \partial d_2 z_i} \right\} f \right\rangle_{K^{(\ell)}} + (1 \leftrightarrow 2) \\
& + \kappa_{(11)(22)}^{(\ell)} \left\langle \left\{ \frac{\partial^2}{\partial d_1 z_i^2} \frac{\partial^2}{\partial d_2 z_{i'}^2} + \frac{\partial^2}{\partial d_1 z_{i'}^2} \frac{\partial^2}{\partial d_2 z_i^2} \right\} f \right\rangle_{K^{(\ell)}} \\
& \left. + 4\kappa_{(12)(12)}^{(\ell)} \left\langle \frac{\partial^2}{\partial d_1 z_i \partial d_2 z_i} \frac{\partial^2}{\partial d_1 z_{i'} \partial d_2 z_{i'}} f \right\rangle_{K^{(\ell)}} \right]. \tag{10.17}
\end{aligned}$$

With these formulas in hand, the derivations of the recursions in Proposition 10.2 are straight forward and we provide the details for the case of $\kappa_{(10)(00)}^{(\ell+1)}$, which already illustrates all the main ideas of the derivation. We have

$$\kappa_{(10)(00)}^{(\ell+1)} = \mathbb{E} \left[\left(\Sigma_{\alpha\alpha}^{(\ell)} - \mathbb{E} \left[\Sigma_{\alpha\alpha}^{(\ell)} \right] \right) \cdot \left(\partial_{x_{1;\alpha_1}} \Sigma_{\alpha_1\alpha_2}^{(\ell)} \Big|_{\alpha_1=\alpha_2=\alpha} - \mathbb{E} \left[\partial_{x_{1;\alpha_1}} \Sigma_{\alpha_1\alpha_2}^{(\ell)} \right] \Big|_{\alpha_1=\alpha_2=\alpha} \right) \right].$$

Recall that

$$\Sigma_{\alpha_1\alpha_2}^{(\ell)} = \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \sigma \left(z_{j;\alpha_1}^{(\ell)} \right) \sigma \left(z_{j;\alpha_2}^{(\ell)} \right)$$

Hence,

$$\Sigma_{\alpha\alpha}^{(\ell)} - \mathbb{E} \left[\Sigma_{\alpha\alpha}^{(\ell)} \right] = \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} X_{j;\alpha},$$

where

$$X_{j;\alpha} := \sigma \left(z_{j;\alpha}^{(\ell)} \right)^2 - \mathbb{E} \left[\sigma \left(z_{j;\alpha}^{(\ell)} \right)^2 \right].$$

Similarly

$$\partial_{x_{\alpha_1}} \Sigma_{\alpha_1\alpha_2}^{(\ell)} \Big|_{\alpha_1=\alpha_2=\alpha} - \mathbb{E} \left[\partial_{x_{\alpha_1}} \Sigma_{\alpha_1\alpha_2}^{(\ell)} \right] \Big|_{\alpha_1=\alpha_2=\alpha} = \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} Y_{j;\alpha},$$

where

$$Y_{j;\alpha} = \sigma \left(z_{j;\alpha}^{(\ell)} \right) \sigma' \left(z_{j;\alpha}^{(\ell)} \right) \partial_{x_{1;\alpha}} z_{j;\alpha}^{(\ell)} - \mathbb{E} \left[\sigma \left(z_{j;\alpha}^{(\ell)} \right) \sigma' \left(z_{j;\alpha}^{(\ell)} \right) \partial_{x_{1;\alpha}} z_{j;\alpha}^{(\ell)} \right].$$

By symmetry, we obtain

$$\kappa_{(10)(00)}^{(\ell+1)} = \frac{C_W^2}{n_\ell} \mathbb{E} [X_{1;\alpha} Y_{1;\alpha}] + C_W^2 \left(1 - \frac{1}{n_\ell} \right) \mathbb{E} [X_{1;\alpha} Y_{2;\alpha}].$$

We have by Theorem 3.2 that

$$\frac{C_W^2}{n_\ell} \mathbb{E} [X_{1;\alpha} Y_{1;\alpha}] = \frac{C_W^2}{n_\ell} \left[\langle \sigma^3 \sigma' d_1 z \rangle_{K^{(\ell)}} - \langle \sigma^2 \rangle_{K^{(\ell)}} \langle \sigma \sigma' d_1 z \rangle_{K^{(\ell)}} \right] + O(n^{-2}),$$

giving the first term in (10.11). Further, note that $X_{1;\alpha} Y_{2;\alpha}$ is a polynomial of total degree one in the derivative $d_1 z_{i;\alpha}^{(\ell)}$. Hence, applying (10.13) and (10.14), we find that $C_W^2 \mathbb{E} [X_{1;\alpha} Y_{2;\alpha}]$ equals

$$C_W^2 \langle X_{1;\alpha} \rangle_{\kappa^{(\ell)}} \langle Y_{2;\alpha\alpha} \rangle_{\kappa^{(\ell)}}$$

plus

$$\frac{C_W^2}{4} \left[2\kappa_{(00)(00)}^{(\ell)} \left\langle \frac{\partial^2}{\partial z_1^2} X_{1;\alpha} \frac{\partial^2}{\partial z_2^2} Y_{2;\alpha} \right\rangle_{K^{(\ell)}} + 2\kappa_{(10)(00)}^{(\ell)} \left\langle \frac{\partial^2}{\partial z_1^2} X_{1;\alpha} \frac{\partial^2}{\partial z_2 \partial d_1 z_2} Y_{2;\alpha} \right\rangle_{K^{(\ell)}} \right]$$

plus an error of size $O(n^{-2})$. Note that by Theorem 3.2 we have

$$\langle X_{1;\alpha} \rangle_{\kappa^{(\ell)}}, \mathbb{E}[Y_{1;\alpha}] = O(n^{-1}).$$

Hence,

$$C_W^2 \langle X_{1;\alpha} \rangle_{\kappa^{(\ell)}} \langle Y_{2;\alpha} \rangle_{\kappa^{(\ell)}} = O(n^{-2}).$$

Further,

$$\left\langle \frac{\partial^2}{\partial z_1^2} X_{1;\alpha} \frac{\partial^2}{\partial z_2^2} Y_{2;\alpha} \right\rangle_{K^{(\ell)}} = \langle \partial^2(\sigma^2) \rangle_{K^{(\ell)}} \langle \partial^2(\sigma\sigma') d_1 z \rangle_{K^{(\ell)}}$$

and

$$\left\langle \frac{\partial^2}{\partial z_1^2} X_{1;\alpha} \frac{\partial^2}{\partial z_2 \partial d_1 z_2} Y_{2;\alpha} \right\rangle_{K^{(\ell)}} = \langle \partial^2 \sigma^2 \rangle_{K^{(\ell)}} \langle \partial(\sigma\sigma') \rangle_{K^{(\ell)}}.$$

Thus, all together, we find

$$\begin{aligned} \kappa_{(10)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle \sigma^3 \sigma' d_1 z \rangle_{K^{(\ell)}} - \langle \sigma^2 \rangle_{K^{(\ell)}} \langle \sigma \sigma' d_1 z \rangle_{K^{(\ell)}} \right] \\ &+ \frac{C_W^2}{4} \kappa_{(00)(00)}^{(\ell)} \langle \partial^2(\sigma^2) \rangle_{K^{(\ell)}} \langle \partial^2(\sigma\sigma') d_1 z \rangle_{K^{(\ell)}} \\ &+ \frac{C_W^2}{2} \kappa_{(10)(00)}^{(\ell)} \langle \partial^2(\sigma^2) \rangle_{K^{(\ell)}} \langle \partial(\sigma\sigma') \rangle_{K^{(\ell)}} \\ &+ O(n^{-2}), \end{aligned}$$

completing our derivation of (10.11). □

10.4 Recursions for Finite Width Corrections to 2nd Cumulant

We remind the reader of our standing notation. Namely, we fix a network input $x_\alpha \neq 0 \in \mathbb{R}^{n_0}$ and define

$$\kappa_{(ij)}^{(\ell)} := \text{Cov} \left(\partial_{x_{i;\alpha}} z_{i;\alpha}^{(\ell)}, \partial_{x_{j;\alpha}} z_{j;\alpha}^{(\ell)} \right) = \partial_{x_{i;\alpha_1}} \partial_{x_{j;\alpha_2}} \kappa_{\alpha_1 \alpha_2}^{(\ell)} \Big|_{x_{\alpha_1} = x_{\alpha_2} = x_\alpha}, \quad i, j = 1, \dots, n_0.$$

We also recall that

$$K_{(ij)}^{(\ell)} := \lim_{n \rightarrow \infty} \kappa_{(ij)}^{(\ell)}, \quad S_{(ij)}^{(\ell)} := \kappa_{(ij)}^{(\ell)} - K_{(ij)}^{(\ell)}.$$

The purpose of this section is to derive recursions for $S_{(ij)}^{(\ell)}$ to leading order in $1/n$. We have

Proposition 10.3. *We have*

$$S_{(00)}^{(\ell+1)} = \frac{C_W}{8} \kappa_{4;\alpha}^{(\ell)} \langle \partial^4 \sigma^2 \rangle_{K_{(00)}^{(\ell)}} + \chi_{\parallel;\alpha}^{(\ell)} S_{(00)}^{(\ell)} + O(n^{-2}). \quad (10.18)$$

Further,

$$\begin{aligned} S_{(10)}^{(\ell+1)} &= \chi_{\parallel}^{(\ell)} S_{(10)}^{(\ell)} + K_{(10)}^{(\ell)} \frac{C_W}{2} S_{(00)}^{(\ell)} \langle \partial^3(\sigma\sigma') \rangle_{K^{(\ell)}} \\ &\quad + \frac{C_W}{8} \left[\kappa_{(00)(00)}^{(\ell)} \langle \partial^4(\sigma\sigma') d_1 z \rangle_{\kappa^{(\ell)}} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^3(\sigma\sigma') \rangle_{\kappa^{(\ell)}} \right] + O(n^{-2}). \end{aligned} \quad (10.19)$$

Moreover,

$$\begin{aligned} S_{(11)}^{(\ell+1)} &= \chi_{\perp}^{(\ell)} S_{(11)}^{(\ell)} + \frac{C_W}{2} K_{(11)}^{(\ell)} S_{(00)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} + O(n^{-2}) \\ &\quad + \frac{C_W}{2} \left(K_{(10)}^{(\ell)} \right)^2 S_{(00)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + 2C_W K_{(10)}^{(\ell)} S_{(10)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \\ &\quad + \frac{C_W}{8} \left\{ \kappa_{(00)(00)}^{(\ell)} \left(K_{(11)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + \left(K_{(10)}^{(\ell)} \right)^2 \langle \partial^6(\sigma')^2 \rangle_{K^{(\ell)}} \right) \right. \\ &\quad \left. + 8\kappa_{(10)(00)}^{(\ell)} K_{(10)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + 4 \left(2\kappa_{(10)(10)}^{(\ell)} + \kappa_{(11)(00)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right\}. \end{aligned} \quad (10.20)$$

Finally,

$$\begin{aligned} S_{(12)}^{(\ell+1)} &= \chi_{\perp}^{(\ell)} S_{(12)}^{(\ell)} + \frac{C_W}{2} S_{(00)}^{(\ell)} \left(K_{(12)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{(\ell)} + K_{(10)}^{(\ell)} K_{(20)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{(\ell)} \right) \\ &\quad + C_W \left(K_{(10)}^{(\ell)} S_{(20)}^{(\ell)} + K_{(20)}^{(\ell)} S_{(10)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{(\ell)} + O(n^{-2}) \\ &\quad + \frac{1}{8} \left\{ \kappa_{(00)(00)}^{(\ell)} \left(K_{(12)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{(\ell)} + K_{(10)}^{(\ell)} K_{(20)}^{(\ell)} \langle \partial^6(\sigma')^2 \rangle_{(\ell)} \right) \right. \\ &\quad + 4 \left(\kappa_{(10)(00)}^{(\ell)} K_{(20)}^{(\ell)} + \kappa_{(20)(00)}^{(\ell)} K_{(10)}^{(\ell)} \right) \langle \partial^4(\sigma')^2 \rangle_{(\ell)} \\ &\quad \left. + 4 \left(\kappa_{(12)(00)}^{(\ell)} + 2\kappa_{(10)(20)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right\}. \end{aligned} \quad (10.21)$$

Proof. The derivation of these recursion is very similar, so we provide the details for the most involved case of $S_{(12)}^{(\ell+1)}$. By definition, we have

$$\kappa_{(12)}^{(\ell+1)} = \mathbb{E} \left[d_1 z_{i;\alpha}^{(\ell+1)} d_2 z_{i;\alpha}^{(\ell+1)} \right] = C_W \mathbb{E} \left[\sigma'(z_{i;\alpha}^{(\ell)})^2 d_1 z_{i;\alpha}^{(\ell)} d_2 z_{i;\alpha}^{(\ell)} \right].$$

Applying Theorem 3.2 (or more precisely equations (10.13) -(10.15)) we therefore find

$$\begin{aligned} \kappa_{(12)}^{(\ell+1)} &= C_W \langle (\sigma')^2 d_1 z d_2 z \rangle_{\kappa^{(\ell)}} + O(n^{-2}) \\ &\quad + \frac{C_W}{8} \left\{ \kappa_{(00)(00)}^{(\ell)} \langle \partial^4(\sigma')^2 d_1 z d_2 z \rangle_{K^{(\ell)}} \right. \\ &\quad + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^3(\sigma')^2 d_2 z \rangle_{K^{(\ell)}} + 4\kappa_{(20)(00)}^{(\ell)} \langle \partial^3(\sigma')^2 d_1 z \rangle_{K^{(\ell)}} \\ &\quad \left. + 4 \left(\kappa_{(12)(00)}^{(\ell)} + 2\kappa_{(10)(20)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right\}. \end{aligned} \quad (10.22)$$

Next, integrating by parts yields

$$\begin{aligned} \langle (\sigma')^2 d_1 z d_2 z \rangle_{\kappa^{(\ell)}} &= \kappa_{(12)}^{(\ell)} \langle (\sigma')^2 \rangle_{\kappa^{(\ell)}} + \kappa_{(10)}^{(\ell)} \kappa_{(20)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{\kappa^{(\ell)}} \\ &= \kappa_{(12)}^{(\ell)} \left(\langle (\sigma')^2 \rangle_{K^{(\ell)}} + \frac{1}{2} S_{(00)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right) \\ &\quad + \left(K_{(10)}^{(\ell)} + S_{(10)}^{(\ell)} \right) \left(K_{(20)}^{(\ell)} + S_{(20)}^{(\ell)} \right) \left(\langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} + \frac{1}{2} S_{(00)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right). \end{aligned}$$

Similarly, for $j = 1, 2$ we find

$$\begin{aligned} \langle \partial^4(\sigma')^2 d_1 z d_2 z \rangle_{K^{(\ell)}} &= K_{(12)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + K_{(10)}^{(\ell)} K_{(20)}^{(\ell)} \langle \partial^6(\sigma')^2 \rangle_{K^{(\ell)}} + O(n^{-1}) \\ \langle \partial^3(\sigma')^2 d_j z \rangle_{K^{(\ell)}} &= K_{(j0)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}}. \end{aligned}$$

Substituting this into (10.22), comparing with the recursion (10.8), and recalling that $\kappa_{(j_1 j_2)(j'_1 j'_2)}^{(\ell)} = O(n^{-1})$ yields the recursion (10.21). \square

11 Proof of Theorem 4.5

In this section we prove Theorem 4.5. The key technical step will be to solve the recursions for all possible second and fourth cumulants of the value and partial derivatives of some component of the output of a random fully connected network with respect to two different components of the network input that were obtained in §10. Before diving into this part of the proof in §11.2, we give the proof of Theorem 4.5 modulo solving the recursions in §10.

11.1 Reduction of Theorem 4.5 to Cumulant Recursions from §10

We first recall the setup of Theorem 4.5. Namely, we fix a non-linearity σ satisfying (2.1) that belongs to the $K_* = 0$ universality class, as defined in §4.2.2, and consider a depth L random fully connected network with input dimension n_0 , hidden layer widths $n_1 = \dots = n_L = n$, output dimension n_{L+1} and non-linearity σ . We assume this network is tuned to criticality (see §4.1). We then fix a network input $x_\alpha \neq 0$ and seek to compute the averages

$$\mathbb{E} \left[\text{Grad Mean}^{(1)} \right] \quad \text{and} \quad \mathbb{E} \left[\text{Grad Var}^{(1)} \right]$$

of the empirical mean

$$\text{Grad Mean}^{(1)} = \frac{1}{n_0 n_1} \sum_{j=1}^{n_0} \sum_{i=1}^{n_1} \left(\frac{\partial z_{q;\alpha}^{(L+1)}}{\partial W_{ij}^{(1)}} \right)^2 \quad (11.1)$$

and the empirical variance

$$\text{Grad Var}^{(1)} = \frac{1}{n_0 n_1} \sum_{j=1}^{n_0} \sum_{i=1}^{n_1} \left(\frac{\partial z_{q;\alpha}^{(L+1)}}{\partial W_{ij}^{(1)}} \right)^4 - \left(\text{Grad Mean}^{(1)} \right)^2 \quad (11.2)$$

of the squared gradients $(\partial z_{q;\alpha}^{(L+1)} / \partial W_{ij}^{(1)})^2$ over all weights in layer 1. To evaluate these quantities to leading order in n, ℓ , note that by symmetry the mean of the expressions $\text{Grad Mean}^{(1)}$ and $\text{Grad Var}^{(1)}$ (11.1) and (11.2) are independent of the choice of q . Next, note that by the chain rule we have

$$\frac{\partial z_{q;\alpha}^{(L+1)}}{\partial W_{ij}^{(1)}} = x_{j;\alpha} \sigma'(z_{i;\alpha}^{(1)}) \frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{i;\alpha}^{(1)})}.$$

Hence, by symmetry,

$$\mathbb{E} \left[\text{Grad Mean}^{(1)} \right] = \frac{1}{n_0} \|x_\alpha\|^2 \mathbb{E} \left[\left(\sigma' \left(z_{1;\alpha}^{(1)} \right) \frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{i;\alpha}^{(1)})} \right)^2 \right].$$

Recalling that $\mathcal{F}^{(\ell)}$ is the sigma algebra generated by all weights and biases in layers up to and including ℓ , we may thus write

$$\mathbb{E} \left[\text{Grad Mean}^{(1)} \right] = \frac{1}{n_0} \|x_\alpha\|^2 \mathbb{E} \left[\left(\sigma' \left(z_{1;\alpha}^{(1)} \right) \right)^2 \mathbb{E} \left[\left(\frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{i;\alpha}^{(1)})} \right)^2 \mid \mathcal{F}^{(1)} \right] \right]. \quad (11.3)$$

Observe that the conditional expectation in the previous line is precisely the square of the partial derivative of one component of the output in a random depth $L - 1$ network tuned to criticality with respect to one component of the input $\sigma_\alpha^{(\ell)}$. In the notation (10.2) above, we therefore have

$$\mathbb{E} \left[\left(\frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{1;\alpha}^{(1)})} \right)^2 \mid \mathcal{F}^{(\ell)} \right] = \kappa_{(11)}^{(L-1)}.$$

We will see by combining Lemma 11.1 and Proposition 11.3 below that

$$\kappa_{(11)}^{(L-1)} = \frac{Cw e^{-\gamma}}{n_1 L} (1 + O(L^{-1})) + O(n^{-2}),$$

where the implicit constant in $O(L^{-1})$ is independent of L, n . Since at leading order in n, L this expression is a (non-random) constant, we write by a slight abuse of notation

$$\begin{aligned} \mathbb{E} \left[\text{Grad Mean}^{(1)} \right] &= \frac{1}{n_0} \|x_\alpha\|^2 \mathbb{E} \left[\left(\sigma' \left(z_{1;\alpha}^{(\ell)} \right) \right)^2 \right] \kappa_{(11)}^{(L-1)} \\ &= \frac{1}{n_0} \|x_\alpha\|^2 \langle (\sigma')^2 \rangle_{K^{(1)}} \kappa_{(11)}^{(L-1)} \end{aligned} \quad (11.4)$$

where in the last step we use that pre-activations in layer 1 are iid Gaussian even at finite width. We now similarly simplify the mean of the empirical gradient variance. Namely, we first note that by symmetry we have the decomposition

$$\mathbb{E} \left[\text{Grad Var}^{(1)} \right] = I - II$$

where

$$\begin{aligned} I &:= \left(\frac{1}{n_0} \|x_\alpha\|_4^4 - \frac{1}{n_1} \left(\frac{1}{n_0} \|x_\alpha\|^2 \right)^2 \right) \mathbb{E} \left[\left(\sigma' \left(z_{1;\alpha}^{(1)} \right) \frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{1;\alpha}^{(1)})} \right)^4 \right] \\ II &:= \left(\frac{1}{n_0} \|x_\alpha\|^2 \right)^2 \left(1 - \frac{1}{n_1} \right) \mathbb{E} \left[\left(\sigma' \left(z_{1;\alpha}^{(1)} \right) \sigma' \left(z_{2;\alpha}^{(1)} \right) \left(\frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{1;\alpha}^{(1)})} \frac{\partial z_{q;\alpha}^{(L+1)}}{\partial \sigma(z_{2;\alpha}^{(1)})} \right)^2 \right)^2 \right]. \end{aligned}$$

By conditioning on $\mathcal{F}^{(1)}$ as before (and again using that at leading order in n, L the cumulants $\kappa_{(ij)(i'j')}^{(\ell)}$ are independent of the initial condition at large ℓ) we have

$$I := 3 \left(\kappa_{(11)(11)}^{(L-1)} + \left(\kappa_{(11)}^{(L-1)} \right)^2 \right) \langle (\sigma')^4 \rangle_{K^{(\ell)}} \left(\frac{1}{n_0} \|x_\alpha\|_4^4 - \frac{1}{n_1} \left(\frac{1}{n_0} \|x_\alpha\|^2 \right)^2 \right) \quad (11.5)$$

$$\begin{aligned} II &:= \left(2\kappa_{(12)(12)}^{(L-1)} + 2 \left(\kappa_{(12)}^{(L-1)} \right)^2 + \kappa_{(11)(22)}^{(L-1)} + \kappa_{(11)}^{(L-1)} \kappa_{(22)}^{(L-1)} \right) \\ &\quad \times \langle (\sigma')^2 \rangle_{K^{(1)}}^2 \left(\frac{1}{n_0} \|x_\alpha\|^2 \right)^2 \left(1 - \frac{1}{n_1} \right). \end{aligned} \quad (11.6)$$

Combining the expressions (11.4), (11.5), and (11.6), we find that

$$\begin{aligned} \frac{\mathbb{E} \left[\text{Grad Var}^{(1)} \right]}{\mathbb{E} \left[\text{Grad Mean}^{(1)} \right]^2} &= 3 \frac{\langle (\sigma')^4 \rangle_{K^{(1)}}}{\langle (\sigma')^2 \rangle_{K^{(1)}}^2} \left(\frac{n_0^{-1} \|x_\alpha\|_4^4}{(n_0^{-1} \|x_\alpha\|_2^2)^2} - \frac{1}{n_1} \right) \left(\widehat{\kappa}_{(11)(11)}^{(L-1)} + 1 \right) \\ &\quad - \left(1 - \frac{1}{n_1} \right) \left(2\widehat{\kappa}_{(12)(12)}^{(L-1)} + 2 \left(\widehat{\kappa}_{(12)}^{(L-1)} \right)^2 + \widehat{\kappa}_{(11)(22)}^{(L-1)} + 1 \right), \end{aligned}$$

plus errors of size $O(n^{-2})$, where we've abbreviated for $i, j \in \{1, 2\}$

$$\widehat{\kappa}_{(ii)(jj)}^{(\ell)} := \frac{\kappa_{(ii)(jj)}^{(\ell)}}{\kappa_{(ii)}^{(\ell)} \kappa_{(jj)}^{(\ell)}}, \quad \widehat{\kappa}_{(12)(12)}^{(\ell)} := \frac{\kappa_{(12)(12)}^{(\ell)}}{\kappa_{(11)}^{(\ell)} \kappa_{(22)}^{(\ell)}}, \quad \widehat{\kappa}_{(12)}^{(\ell)} := \frac{\kappa_{(12)}^{(\ell)}}{\left(\kappa_{(11)}^{(\ell)} \kappa_{(22)}^{(\ell)} \right)^{1/2}}.$$

Appealing to Propositions 11.2 and 11.3 we conclude

$$\begin{aligned} \widehat{\kappa}_{(11)(11)}^{(L-1)} + 1 &= 1 + \frac{8L}{3n} + O(n^{-1}) + O(n^{-2}) \\ 2\widehat{\kappa}_{(12)(12)}^{(L-1)} + 2 \left(\widehat{\kappa}_{(12)}^{(L-1)} \right)^2 + \widehat{\kappa}_{(11)(22)}^{(L-1)} + 1 &= 1 + \frac{8L}{3n} + O(n^{-1}) + O(n^{-2}), \end{aligned}$$

where the implicit constants in $O(n^{-1})$ depend only on σ whereas the constant in $O(n^{-2})$ depends on both σ and L . Combining these expressions with the following estimates completes the derivation of Theorem 4.5, modulo solving the recursions for the 2nd cumulants $\kappa_{(ij)}^{(\ell)}$ and the fourth cumulants $\kappa_{(i_1 j_1)(i_2 j_2)}^{(\ell)}$, which we supply in the next section. \square

11.2 Completion of Proof of Theorem 4.4: 2nd, 4th Cumulants at Large Depth for the $K_* = 0$ Universality Class

In this section, we complete the proof of Theorem 4.5 by solving the 2nd and 4th cumulant recursion obtained in §10 when σ belongs to the $K_* = 0$ universality class. We start in §11.3 by solving the infinite width variance recursions for $K_{(ij)}^{(\ell)}$. We then use this for solving to leading order in $1/n$ the fourth cumulant recursions for $\kappa_{(ij)(i'j')}^{(\ell)}$ in §11.4. Finally, we combine in §11.5 the results of the preceding two sections to solve the recursions for $S_{(ij)}^{(\ell)}$ again at leading order in n, ℓ .

11.3 Two Point Recursions

We remind the reader of our standing notation. Namely, we let σ be a non-linearity from the $K_* = 0$ universality class (see §4.2.2). We also fix a network input $x_\alpha \neq 0 \in \mathbb{R}^{n_0}$ and define

$$K_{(ij)}^{(\ell)} := \lim_{n \rightarrow \infty} \text{Cov} \left(\partial_{x_{i;\alpha}} z_{1;\alpha}^{(\ell)}, \partial_{x_{j;\alpha}} z_{1;\alpha}^{(\ell)} \right) = \partial_{x_{i;\alpha_1}} \partial_{x_{j;\alpha_2}} K_{\alpha_1 \alpha_2}^{(\ell)} \Big|_{x_{\alpha_1} = x_{\alpha_2} = x_\alpha}, \quad i, j = 1, \dots, n_0.$$

The recursive description for $K_{(ij)}^{(\ell)}$ was obtained in Lemma 10.1. These recursions held for any non-linearity and any values of C_b, C_W . The following Lemma records the large ℓ behavior of $K_{(ij)}^{(\ell)}$ when σ belongs to the $K_* = 0$ universality class and the network is tuned to criticality.

Lemma 11.1. *We have*

$$K_{(00)}^{(\ell+1)} = \frac{1}{a\ell} + O_\delta(\ell^{-2+\delta}) \quad (11.7)$$

$$K_{(10)}^{(\ell+1)} = \frac{C_W e^{-2\gamma} x_{1;\alpha}}{n_0 \ell^2} (1 + O(\ell^{-1})) \quad (11.8)$$

$$K_{(11)}^{(\ell+1)} = \frac{C_W e^{-\gamma}}{n_0 \ell} (1 + O(\ell^{-1})) + O\left(\frac{(\sigma_1^{(0)})^2}{n_0^2 \ell^3}\right) \quad (11.9)$$

$$K_{(12)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha} x_{2;\alpha}}{n_0^2 \ell^3}\right) \quad (11.10)$$

plus errors of size $O(n^{-1})$.

Proof. From Lemma 10.1 we have, up to errors of size $O(n^{-1})$, that

$$K_{(00)}^{(\ell+1)} = C_W \langle \sigma^2 \rangle_{(\ell)} \quad (11.11)$$

$$K_{(10)}^{(\ell+1)} = \chi_{\parallel}^{(\ell)} K_{(10)}^{(\ell)} \quad (11.12)$$

$$K_{(11)}^{(\ell+1)} = C_W \langle \partial^2(\sigma')^2 \rangle_{(\ell)} \left(K_{(10)}^{(\ell)} \right)^2 + \chi_{\perp}^{(\ell)} K_{(11)}^{(\ell)} \quad (11.13)$$

$$K_{(12)}^{(\ell+1)} = C_W \langle \partial^2(\sigma')^2 \rangle_{(\ell)} K_{(10)}^{(\ell)} K_{(20)}^{(\ell)} + \chi_{\perp}^{(\ell)} K_{(12)}^{(\ell)}. \quad (11.14)$$

We already proved (11.7) in §B. Next, (11.12) yields

$$K_{(10)}^{(\ell+1)} = K_{(10)}^{(1)} \prod_{\ell'=1}^{\ell} \chi_{\parallel}^{(\ell')}.$$

By direct computation we have

$$K_{(10)}^{(1)} = \lim_{n \rightarrow \infty} \mathbb{E} \left[z_{i;\alpha}^{(1)} d_1 z_{i;\alpha}^{(1)} \right] = \frac{C_W}{n_0} x_{1;\alpha}.$$

Moreover,

$$\begin{aligned} \prod_{\ell'=1}^{\ell} \chi_{\parallel}^{(\ell')} &= \prod_{\ell'=1}^{\ell} \left(1 - \frac{2}{\ell'} + O((\ell')^{-2}) \right) \\ &\quad \exp \left[\sum_{\ell'=1}^{\ell} \log \left(1 - \frac{2}{\ell'} + O((\ell')^{-2}) \right) \right] \\ &= \exp \left[O(\ell^{-1}) + -2 \sum_{\ell'=1}^{\ell} \frac{1}{\ell'} \right] \\ &= e^{-2\gamma} \ell^{-2} (1 + O(\ell^{-1})), \end{aligned}$$

where γ is the Euler-Mascheroni constant. This yields (11.8). Further,

$$C_W \langle \partial^2(\sigma')^2 \rangle_{(\ell)} \left(K_{(10)}^{(\ell)} \right)^2 = O\left(\frac{(\sigma_1^{(0)})^2}{n_0^2 \ell^4}\right).$$

Hence, using Lemma 6.2, we find to leading order in ℓ, n that

$$K_{(11)}^{(\ell+1)} = K_{11}^{(1)} \prod_{\ell'=1}^{\ell} \chi_{\perp}^{(\ell')} + O\left(\frac{(\sigma_1^{(0)})^2}{n_0^2 \ell^3}\right) = K_{11}^{(1)} \prod_{\ell'=1}^{\ell} (1 - \ell^{-1} + O(\ell^{-2})) + O\left(\frac{(\sigma_1^{(0)})^2}{n_0^2 \ell^3}\right)$$

Noting that

$$K_{11}^{(1)} = \frac{C_W}{n_0}$$

shows

$$K_{(11)}^{(\ell+1)} = K_{11}^{(1)} \prod_{\ell'=1}^{\ell} \chi_{\perp}^{(\ell')} + O\left(\frac{(\sigma_1^{(0)})^2}{n_0^2 \ell^3}\right) = \frac{C_W e^{-\gamma}}{n_0 \ell} (1 + O(\ell^{-1})) + O\left(\frac{(\sigma_1^{(0)})^2}{n_0^2 \ell^3}\right),$$

completing the derivation of (11.9). The derivation of (11.30) is similar except that $K_{12}^{(1)} = 0$. \square

11.4 Four Point Recursions

We remind the reader of our standing notation. Namely, we let σ be a non-linearity from the $K_* = 0$ universality class (see §4.2.2). We also fix a network input $x_{\alpha} \neq 0 \in \mathbb{R}^{n_0}$ and recall that

$$\begin{aligned} \kappa_{(ij)(i'j')}^{(\ell)} &:= \text{Cov}\left(\partial_{x_{i;\alpha}} \partial_{x_{j;\alpha}} \Delta_{\alpha\alpha}^{(\ell)}, \partial_{x_{i';\alpha}} \partial_{x_{j';\alpha}} \Delta_{\alpha\alpha}^{(\ell)}\right) \\ &= \kappa\left(\partial_{x_{i;\alpha}} z_{1;\alpha}^{(\ell+1)}, \partial_{x_{j;\alpha}} z_{1;\alpha}^{(\ell+1)}, \partial_{x_{i';\alpha}} z_{2;\alpha}^{(\ell+1)}, \partial_{x_{j';\alpha}} z_{2;\alpha}^{(\ell+1)}\right), \end{aligned}$$

where $\Delta_{\alpha\alpha}^{(\ell)}$ is defined in (3.10).

The recursive description for $\kappa_{(ij)(i'j')}^{(\ell)}$ was obtained in Proposition 10.2. These recursions held for any non-linearity and any values of C_b, C_W . The following Lemma records the large ℓ behavior of $\kappa_{(ij)(i'j')}^{(\ell)}$ when σ belongs to the $K_* = 0$ universality class and the network is tuned to criticality.

Proposition 11.2. *We have*

$$\kappa_{(00)(00)}^{(\ell)} = \frac{2}{3n\ell a^2}(1 + O(\ell^{-1})) = \frac{2\ell}{3n} \left(K_{(00)}^{(\ell)} \right)^2 (1 + O(\ell^{-1})) \quad (11.15)$$

$$\kappa_{(10)(00)}^{(\ell+1)} = \frac{C_W e^{-2\gamma} x_{1;\alpha}}{3ann_0\ell^2} (1 + O(\ell^{-1})) = \frac{\ell}{3n} K_{(10)}^{(\ell)} K_{(00)}^{(\ell)} (1 + O(\ell^{-1})) \quad (11.16)$$

$$\kappa_{(10)(10)}^{(\ell+1)} = \frac{C_W e^{-\gamma}}{3ann_0\ell} (1 + O(\ell^{-1})) = \frac{\ell}{3n} K_{(00)}^{(\ell)} K_{(11)}^{(\ell)} (1 + O(\ell^{-1})) \quad (11.17)$$

$$\kappa_{(10)(20)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha} x_{2;\alpha}}{nn_0^2\ell^3}\right) \quad (11.18)$$

$$\kappa_{(11)(00)}^{(\ell+1)} = -\frac{C_W e^{-\gamma}}{3ann_0\ell} (1 + O(\ell^{-1})) = -\frac{\ell}{3n} K_{(11)}^{(\ell)} K_{(00)}^{(\ell)} (1 + O(\ell^{-1})) \quad (11.19)$$

$$\kappa_{(12)(00)}^{(\ell+1)} = O\left(\frac{1}{nn_0^2\ell^3}\right) \quad (11.20)$$

$$\kappa_{(11)(10)}^{(\ell+1)} = O\left(\frac{1}{nn_0^2\ell^2}\right) \quad (11.21)$$

$$\kappa_{(12)(10)}^{(\ell+1)} = O\left(\frac{x_{2;\alpha}}{nn_0^2\ell}\right) \quad (11.22)$$

$$\kappa_{(11)(20)}^{(\ell+1)} = O\left(\frac{x_{2;\alpha}}{nn_0^2\ell}\right) \quad (11.23)$$

$$\kappa_{(11)(11)}^{(\ell+1)} = \frac{8C_W^2 e^{-2\gamma}}{3nn_0^2\ell} (1 + O(\ell^{-1})) = \frac{8\ell}{3n} \left(K_{(11)}^{(\ell)} \right)^2 (1 + O(\ell^{-1})) \quad (11.24)$$

$$\kappa_{(11)(22)}^{(\ell+1)} = \frac{2C_W^2 e^{-2\gamma}}{3nn_0^2\ell} (1 + O(\ell^{-1})) = \frac{2\ell}{3n} K_{(11)}^{(\ell)} K_{(22)}^{(\ell)} (1 + O(\ell^{-1})) \quad (11.25)$$

$$\kappa_{(12)(12)}^{(\ell+1)} = \frac{C_W^2 e^{-2\gamma}}{nn_0^2\ell} (1 + O(\ell^{-1})) = \frac{\ell}{n} K_{(11)}^{(\ell)} K_{(22)}^{(\ell)} (1 + O(\ell^{-1})) \quad (11.26)$$

plus errors of size $O(n^{-2})$. Here, γ denotes the Euler-Mascheroni constant.

Proof. This proof is a long but straightforward calculation using Theorem 3.2, Lemma 6.2, and the estimates from Lemma 11.1. We will spell this out for the first few recursions and then simply write the recursion and its solution. For example, Theorem 3.2 shows

$$\kappa_{(00)(00)}^{(\ell+1)} = \frac{C_W^2}{n_\ell} \left(\langle \sigma^4 \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)}^2 \right) + \left(\chi_{||}^{(\ell)} \right)^2 \kappa_{(00)(00)}^{(\ell)}$$

plus errors of size $O(n^{-2})$. A direction computation then yields

$$\frac{C_W^2}{n_\ell} \left(\langle \sigma^4 \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)}^2 \right) = \frac{2}{n_\ell} \left(\kappa_{(00)}^{(\ell)} \right)^2 (1 + O(\ell^{-1})).$$

Hence, setting $n_\ell = n$ we may apply Lemma 6.2 to obtain

$$\kappa_{(00)(00)}^{(\ell)} = \frac{2}{3n\ell a^2} (1 + O(\ell^{-1})).$$

Similarly, we have

$$\begin{aligned}\kappa_{(10)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left(\langle \sigma^3 \sigma' d_1 z \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)} \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \right) \\ &\quad + \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \right]\end{aligned}$$

plus errors of size $O(n^{-2})$. A direct computation shows that

$$\frac{C_W^2}{n_\ell} \left(\langle \sigma^3 \sigma' d_1 z \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)} \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \right) = \frac{2\kappa_{00}^{(\ell)} \kappa_{10}^{(\ell)}}{n_\ell} (1 + O(\ell^{-1}))$$

Hence, setting $n_\ell = n$ and

$$\kappa_{(10)(00)}^{(\ell+1)} = -\frac{2C_W x_{1;\alpha} e^{-2\gamma}}{3ann_0 \ell^3} (1 + O(\ell^{-1})) + \left(1 - \frac{4}{\ell} + O(\ell^{-2})\right) \kappa_{(10)(00)}^{(\ell)}.$$

Thus, Lemma 6.2 yields

$$\kappa_{(10)(00)}^{(\ell+1)} = -\frac{e^{-2\gamma} (x_{1;\alpha})^2}{3ann_0 \ell^2} = O\left(\frac{x_{1;\alpha}}{nn_0 \ell^2}\right).$$

Further, applying Theorem 3.2 gives

$$\begin{aligned}\kappa_{(10)(10)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma \sigma' d_1 z)^2 \rangle_{(\ell)} - \langle \sigma \sigma' d_1 z \rangle_{(\ell)}^2 \right] \\ &\quad + \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)}^2 + \kappa_{(10)(00)}^{(\ell)} 8 \langle \partial (\sigma \sigma') \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} \right. \\ &\quad \left. + 8\kappa_{(10)(10)}^{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)}^2 \right].\end{aligned}$$

A direct computation shows

$$\frac{C_W^2}{n_\ell} \left[\langle (\sigma \sigma' d_1 z)^2 \rangle_{(\ell)} - \langle \sigma \sigma' d_1 z \rangle_{(\ell)}^2 \right] = \frac{\kappa_{11}^{(\ell)} \kappa_{00}^{(\ell)}}{n} (1 + O(\ell^{-1})).$$

Setting $n_\ell = n$ and using the formulas above shows

$$\kappa_{(10)(10)}^{(\ell+1)} = \frac{\kappa_{11}^{(\ell)} \kappa_{00}^{(\ell)}}{n} (1 + O(\ell^{-1})) + \left(1 - \frac{4}{\ell}\right) \kappa_{(10)(10)}^{(\ell)}$$

and hence by Lemma 6.2 we see that

$$\kappa_{(10)(10)}^{(\ell+1)} = \frac{C_W e^{-\gamma}}{3nn_0 a \ell} (1 + O(\ell^{-1})).$$

Next, we obtain

$$\begin{aligned}\kappa_{(10)(20)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma \sigma')^2 d_1 z d_2 z \rangle_{(\ell)} - \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \langle \sigma \sigma' d_2 z \rangle_{(\ell)} \right] \\ &\quad + \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} \right. \\ &\quad \left. + \kappa_{(10)(00)}^{(\ell)} 4 \langle \partial (\sigma \sigma') \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} + \kappa_{(20)(00)}^{(\ell)} 4 \langle \partial (\sigma \sigma') \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_1 z \rangle_{(\ell)} \right. \\ &\quad \left. + \kappa_{(10)(20)}^{(\ell)} 8 \langle \partial (\sigma \sigma') \rangle_{(\ell)}^2 \right].\end{aligned}$$

Keeping the leading order in ℓ we therefore have

$$\kappa_{(10)(20)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha}x_{2;\alpha}}{nn_0^2\ell^4}\right) + \left(1 - \frac{4}{\ell} + O(\ell^{-2})\right) \kappa_{(10)(20)}^{(\ell)}.$$

Thus,

$$\kappa_{(10)(20)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha}x_{2;\alpha}}{nn_0^2\ell^3}\right)$$

Next,

$$\begin{aligned} \kappa_{(11)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma\sigma' d_1 z)^2 \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)} \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[2\kappa_{(00)(00)}^{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right. \\ &\quad \left. \kappa_{(11)(00)}^{(\ell)} 4 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right]. \end{aligned}$$

We have

$$\frac{C_W^2}{n_\ell} \left[\langle (\sigma\sigma' d_1 z)^2 \rangle_{(\ell)} - \langle \sigma^2 \rangle_{(\ell)} \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \right] = O\left(\frac{1}{nn_0\ell^4}\right).$$

Keeping the leading order in ℓ gives

$$\kappa_{(11)(00)}^{(\ell+1)} = \kappa_{11}^{(\ell)} \kappa_{(00)(00)}^{(\ell)} \frac{C_W}{2} \langle \partial^2 (\sigma')^2 \rangle_{(\ell)} (1 + O(\ell^{-1})) + \chi_\perp^{(\ell)} \chi_\parallel^{(\ell)} \kappa_{(11)(00)}^{(\ell)}$$

The solution is therefore

$$\kappa_{(11)(00)}^{(\ell+1)} = -\frac{2C_W e^{-\gamma}}{3ann_0\ell} (1 + O(\ell^{-1})).$$

Next,

$$\begin{aligned} \kappa_{(12)(00)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^2 \sigma^2 d_1 z d_2 z \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \sigma^2 \rangle_{(\ell)} \right] \\ &= \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right. \\ &\quad + \kappa_{(10)(00)}^{(\ell)} 4 \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} + \kappa_{(20)(00)}^{(\ell)} 4 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \\ &\quad \left. + \kappa_{(12)(00)}^{(\ell)} 4 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 \sigma^2 \rangle_{(\ell)} \right]. \end{aligned}$$

A direct computation shows that

$$\frac{C_W^2}{n_\ell} \left[\langle (\sigma')^2 \sigma^2 d_1 z d_2 z \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \sigma^2 \rangle_{(\ell)} \right] = -a\kappa_{11}^{(\ell)} \left(\kappa_{00}^{(\ell)} \right)^2 + O\left(\frac{1}{n_0^2 n \ell^4}\right).$$

Thus, to leading order in ℓ we have

$$\kappa_{(12)(00)}^{(\ell+1)} = -a\kappa_{11}^{(\ell)} \left(\kappa_{00}^{(\ell)} \right)^2 (1 + O(\ell^{-1})) + \chi_\parallel^{(\ell)} \chi_\perp^{(\ell)} \kappa_{(12)(00)}^{(\ell)}$$

and we conclude that

$$\kappa_{(12)(00)}^{(\ell+1)} = O\left(\frac{1}{nn_0\ell^2}\right).$$

Next,

$$\begin{aligned} \kappa_{(11)(10)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle \sigma(\sigma' d_1 z)^3 \rangle_{(\ell)} - \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} \langle \partial^2(\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \right. \\ &\quad + \kappa_{(10)(00)}^{(\ell)} 4 \left(\langle \partial(\sigma \sigma') \rangle_{(\ell)} \langle \partial^2(\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + 2 \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} \right) \\ &\quad + \kappa_{(10)(10)}^{(\ell)} 16 \langle \partial(\sigma \sigma') \rangle_{(\ell)} \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} + \kappa_{(11)(00)}^{(\ell)} 2 \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \\ &\quad \left. + \kappa_{(11)(10)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial(\sigma \sigma') \rangle_{(\ell)} \right]. \end{aligned}$$

A direct computation shows that to leading order in ℓ

$$\frac{C_W^2}{n_\ell} \left[\langle \sigma(\sigma' d_1 z)^3 \rangle_{(\ell)} - \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \right] = \frac{\kappa_{10}^{(\ell)} \kappa_{11}^{(\ell)}}{n_\ell} = O\left(\frac{1}{n_0^2 n \ell^3}\right).$$

We since all other terms are at this order in ℓ or lower we obtain

$$\kappa_{(11)(10)}^{(\ell+1)} = O\left(\frac{1}{nn_0^2\ell^2}\right).$$

Next,

$$\begin{aligned} \kappa_{(12)(10)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^3 \sigma (d_1 z)^2 d_2 z \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \right] \\ &= \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2(\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} \right. \\ &\quad + \kappa_{(10)(00)}^{(\ell)} 4 \left\{ \langle \partial(\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} + \langle \partial^2(\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \partial(\sigma \sigma') \rangle_{(\ell)} \right\} \\ &\quad + \kappa_{(20)(00)}^{(\ell)} 4 \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} + \kappa_{(12)(00)}^{(\ell)} 4 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2(\sigma \sigma') d_1 z \rangle_{(\ell)} \\ &\quad + \kappa_{(10)(10)}^{(\ell)} 8 \langle \partial(\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial(\sigma \sigma') \rangle_{(\ell)} + \kappa_{(10)(20)}^{(\ell)} 8 \langle \partial(\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial(\sigma \sigma') \rangle_{(\ell)} \\ &\quad \left. + \kappa_{(12)(10)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial(\sigma \sigma') \rangle_{(\ell)} \right]. \end{aligned}$$

A direct computation shows that to leading order in ℓ we have

$$\frac{C_W^2}{n_\ell} \left[\langle (\sigma')^3 \sigma (d_1 z)^2 d_2 z \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \langle \sigma \sigma' d_1 z \rangle_{(\ell)} \right] = O\left(\frac{x_{2;\alpha}}{nn_0\ell^3}\right)$$

and

$$\kappa_{(12)(10)}^{(\ell+1)} = O\left(\frac{x_{2;\alpha}}{nn_0^2\ell^3}\right) + \left(1 - \frac{2}{\ell} + O(\ell^{-2})\right) \kappa_{(12)(10)}^{(\ell)}.$$

Hence,

$$\kappa_{(12)(10)}^{(\ell+1)} = O\left(\frac{x_{2;\alpha}}{nn_0^2\ell^2}\right).$$

Next,

$$\begin{aligned} \kappa_{(11)(20)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^2 \sigma (d_1 z)^2 d_2 z \rangle_{(\ell)} - \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \langle \sigma \sigma' d_2 z \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} \right. \\ &\quad + \kappa_{(10)(00)}^{(\ell)} 8 \langle \partial^2 (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} + \kappa_{(20)(00)}^{(\ell)} 4 \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \\ &\quad + \kappa_{(11)(00)}^{(\ell)} 4 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma \sigma') d_2 z \rangle_{(\ell)} + \kappa_{(10)(20)}^{(\ell)} 16 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial (\sigma \sigma') \rangle_{(\ell)} \\ &\quad \left. + \kappa_{(11)(20)}^{(\ell)} 8 \langle \partial (\sigma \sigma') \rangle_{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \right]. \end{aligned}$$

We therefore have

$$\kappa_{(11)(20)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha} x_{2;\alpha}}{nn_0^2\ell^3}\right) + \left(1 - \frac{4}{\ell} + O(\ell^{-2})\right) \kappa_{(11)(20)}^{(\ell)}.$$

Hence,

$$\kappa_{(11)(20)}^{(\ell+1)} = O\left(\frac{x_{1;\alpha} x_{2;\alpha}}{nn_0^2\ell^2}\right).$$

Next,

$$\begin{aligned} \kappa_{(11)(11)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma' dz)^4 \rangle_{(\ell)} - \langle (\sigma' dz)^2 \rangle_{(\ell)}^2 \right] \\ &+ \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)}^2 + \kappa_{(10)(00)}^{(\ell)} 16 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \right. \\ &\quad + \kappa_{(11)(00)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + \kappa_{(10)(10)}^{(\ell)} 32 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)}^2 \\ &\quad \left. + \kappa_{(11)(10)}^{(\ell)} 32 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} + \kappa_{(11)(11)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)}^2 \right] \\ &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma' dz)^4 \rangle_{(\ell)} - \langle (\sigma' dz)^2 \rangle_{(\ell)}^2 \right] \\ &+ C_W^2 \kappa_{(11)(00)}^{(\ell)} \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} + \kappa_{(11)(11)}^{(\ell)} \left(\chi_{\perp}^{(\ell)} \right)^2. \end{aligned}$$

A direction computation shows that to leading order in ℓ

$$\frac{C_W^2}{n_\ell} \left[\langle (\sigma' dz)^4 \rangle_{(\ell)} - \langle (\sigma' dz)^2 \rangle_{(\ell)}^2 \right] = \frac{2 \left(\kappa_{11}^{(\ell)} \right)^2}{n_\ell}.$$

Hence, to leading order in ℓ we simply have

$$\kappa_{(11)(11)}^{(\ell+1)} = \left(\frac{2}{n_\ell} \left(\kappa_{11}^{(\ell)} \right)^2 + (-2a) \kappa_{11}^{(\ell)} \kappa_{(11)(00)}^{(\ell)} \right) (1 + O(\ell^{-1})) + \left(1 - \frac{2}{\ell} \right) \kappa_{(11)(11)}^{(\ell)}.$$

We thus find

$$\kappa_{(11)(11)}^{(\ell+1)} = \frac{7C_W^2 e^{-2\gamma}}{3nn_0^2\ell} (1 + O(\ell^{-1})).$$

Next,

$$\begin{aligned} \kappa_{(11)(22)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^4 (d_1 z d_2 z)^2 \rangle_{(\ell)} - \langle (\sigma' d_1 z)^2 \rangle_{(\ell)} \langle (\sigma' d_2 z)^2 \rangle_{(\ell)} \right] \\ &+ \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial^2 (\sigma')^2 (d_1 z)^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_2 z)^2 \rangle_{(\ell)} \right. \\ &\quad + \left(\kappa_{(10)(00)}^{(\ell)} + \kappa_{(20)(00)}^{(\ell)} \right) 8 \langle \partial^2 (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_2 z \rangle_{(\ell)} \\ &\quad + \kappa_{(11)(00)}^{(\ell)} 4 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 (d_2 z)^2 \rangle_{(\ell)} + \kappa_{(22)(00)}^{(\ell)} 4 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma') (d_1 z)^2 \rangle_{(\ell)} \\ &\quad + \kappa_{(10)(20)}^{(\ell)} 32 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \\ &\quad + \kappa_{(11)(20)}^{(\ell)} 16 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} + \kappa_{(22)(10)}^{(\ell)} 16 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \\ &\quad \left. + \kappa_{(11)(22)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)}^2 \right]. \end{aligned}$$

Substituting the solutions above we find

$$\kappa_{(11)(22)}^{(\ell+1)} = \frac{4C_W^2 e^{-2\gamma}}{3nn_0^2\ell} (1 + O(\ell^{-1})).$$

Finally,

$$\begin{aligned} \kappa_{(12)(12)}^{(\ell+1)} &= \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^4 (d_1 z d_2 z)^2 \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)}^2 \right] \\ &+ \frac{C_W^2}{8} \left[\kappa_{(00)(00)}^{(\ell)} 2 \langle \partial (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)}^2 \right. \\ &\quad + \kappa_{(10)(00)}^{(\ell)} 8 \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} + \kappa_{(20)(00)}^{(\ell)} 8 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} \\ &\quad + \kappa_{(12)(00)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial^2 (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)} + \kappa_{(10)(20)}^{(\ell)} 16 \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \\ &\quad + \kappa_{(10)(10)}^{(\ell)} 8 \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)}^2 + \kappa_{(20)(20)}^{(\ell)} 8 \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)}^2 \\ &\quad + \kappa_{(12)(10)}^{(\ell)} 16 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_2 z \rangle_{(\ell)} + \kappa_{(12)(20)}^{(\ell)} 16 \langle (\sigma')^2 \rangle_{(\ell)} \langle \partial (\sigma')^2 d_1 z \rangle_{(\ell)} \\ &\quad \left. + \kappa_{(12)(12)}^{(\ell)} 8 \langle (\sigma')^2 \rangle_{(\ell)}^2 \right]. \end{aligned}$$

To leading order in ℓ we have

$$\kappa_{(12)(12)}^{(\ell+1)} = \frac{C_W^2}{n_\ell} \left[\langle (\sigma')^4 (d_1 z d_2 z)^2 \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)}^2 \right] (1 + O(\ell^{-1})) + \left(\chi_\perp^{(\ell)} \right)^2 \kappa_{(12)(12)}^{(\ell)}.$$

A direct computation shows

$$\frac{C_W^2}{n_\ell} \left[\langle (\sigma')^4 (d_1 z d_2 z)^2 \rangle_{(\ell)} - \langle (\sigma')^2 d_1 z d_2 z \rangle_{(\ell)}^2 \right] = \frac{\kappa_{11}^{(\ell)} \kappa_{22}^{(\ell)}}{n} (1 + O(\ell^{-1})).$$

Hence,

$$\kappa_{(12)(12)}^{(\ell+1)} = \frac{C_W^2 e^{-2\gamma}}{nn_0^2\ell} (1 + O(\ell^{-1})).$$

□

11.5 Finite Width Corrections to 2nd Cumulant

We remind the reader of our standing notation. Namely, we let σ be an non-linearity from the $K_* = 0$ universality class (see §4.2.2). We also fix a network input $x_\alpha \neq 0 \in \mathbb{R}^{n_0}$ and recall that

$$\kappa_{(ij)}^{(\ell)} := \text{Cov} \left(\partial_{x_{i;\alpha}} z_{i;\alpha}^{(\ell)}, \partial_{x_{j;\alpha}} z_{j;\alpha}^{(\ell)} \right) = \partial_{x_{i;\alpha_1}} \partial_{x_{j;\alpha_2}} \kappa_{\alpha_1 \alpha_2}^{(\ell)} \Big|_{x_{\alpha_1} = x_{\alpha_2} = x_\alpha}, \quad i, j = 1, \dots, n_0.$$

The recursive description for $\kappa_{(ij)}^{(\ell)}$ and its correction

$$S_{(ij)}^{(\ell)} := \kappa_{(ij)}^{(\ell)} - K_{(ij)}^{(\ell)}$$

relative to the infinite width limit was obtained in Proposition 10.3. These recursions held for any non-linearity and any values of C_b, C_W . The following Lemma records the large ℓ behavior of $\kappa_{(ij)}^{(\ell)}$ when σ belongs to the $K_* = 0$ universality class and the network is tuned to criticality.

Proposition 11.3. *Fix $\delta \in (0, 1)$. We have*

$$\kappa_{(00)}^{(\ell)} = \left[1 - \frac{\ell}{3n} + O(n^{-1}) + O_\delta(\ell^{-1+\delta}) + O(n^{-2}) \right] K_{(00)}^{(\ell)} \quad (11.27)$$

$$\kappa_{(10)}^{(\ell)} = \left[1 + O(n^{-1}) + O(\ell^{-1}) + O(n^{-2}) \right] K_{(10)}^{(\ell)} \quad (11.28)$$

$$\kappa_{(11)}^{(\ell)} = \left[1 + \frac{\ell}{3n} + O(n^{-1}) + O(\ell^{-1}) + O\left(\frac{(\sigma_1^{(0)})^2}{n_0 \ell^2}\right) + O(n^{-2}) \right] K_{(11)}^{(\ell)} \quad (11.29)$$

$$\kappa_{(12)}^{(\ell)} = O\left(\left[1 + \frac{\ell}{n} + O(\ell^{-1}) + O(n^{-2})\right] K_{(12)}^{(\ell)}\right), \quad (11.30)$$

where the implicit constants in the error terms $O(n^{-1}), O(\ell^{-1}), O((\sigma_1^{(0)})^2/n_0 \ell^2), O(\ell/n)$ depend only on σ , the implicit constant in $O_\delta(\ell^{-1+\delta})$ depends on δ and σ , and the implicit constants in $O(n^{-2})$ depend both on σ and on ℓ .

Remark 11.4. *The results of Proposition 11.3 make clear that the error terms $O(n^{-2})$ must depend on ℓ . For instance, note that $\kappa_{(00)}^{(\ell)}$ must be non-negative and hence the $O(n^{-2})$ term is negligible only if $n \gg \ell$.*

Proof. Recall that in Lemma 11.1 we already the large ℓ behavior of

$$K_{(ij)}^{(\ell)} := \lim_{n \rightarrow \infty} \kappa_{(ij)}^{(\ell)}, \quad i, j \in \{0, 1, 2\}.$$

Define

$$S_{(ij)}^{(\ell)} := \kappa_{(ij)}^{(\ell)} - K_{(ij)}^{(\ell)}, \quad i, j \in \{0, 1, 2\}.$$

By direct computation and Theorem 3.2 we obtain

$$\kappa_{(00)}^{(\ell+1)} = C_W \mathbb{E} \left[\sigma(z_{i;\alpha}^{(\ell)})^2 \right] = C_W \left\{ \langle \sigma^2 \rangle_{\kappa_{(00)}^{(\ell)}} + \frac{1}{8} \kappa_{4;\alpha}^{(\ell)} \langle \partial^4 \sigma^2 \rangle_{\kappa_{(00)}^{(\ell)}} + O(n^{-2}) \right\}.$$

Comparing this with the recursion (10.5) and using that by Theorem 3.1 we know $\kappa_{4;\alpha}^{(\ell)} = O(n^{-1})$ allows us to conclude

$$S_{(00)}^{(\ell)} = O(n^{-1}), \quad (11.31)$$

where at the moment the implicit constant in the error term depends in an unknown way on ℓ . However, substituting this estimate back into the recursion for $\kappa_{(00)}^{(\ell+1)}$ and integrating by parts yields

$$\begin{aligned} S_{(00)}^{(\ell+1)} &= C_W \left\{ \langle \sigma^2 \rangle_{\kappa_{(00)}^{(\ell)}} - \langle \sigma^2 \rangle_{K_{(00)}^{(\ell)}} \right\} + \frac{C_W}{8} \kappa_{4;\alpha}^{(\ell)} \langle \partial^4 \sigma^2 \rangle_{K_{(00)}^{(\ell)}} + O(n^{-2}) \\ &= \frac{C_W}{8} \kappa_{4;\alpha}^{(\ell)} \langle \partial^4 \sigma^2 \rangle_{K_{(00)}^{(\ell)}} + \chi_{\parallel;\alpha}^{(\ell)} S_{(00)}^{(\ell)} + O(n^{-2}) \end{aligned}$$

Using (11.7) we obtain

$$S_{(00)}^{(\ell+1)} = -\frac{2}{3na\ell} (1 + O(\ell^{-1})) + \left(1 - \frac{2}{\ell}\right) S_{(00)}^{(\ell)} + O(n^{-2}).$$

Noting that $S_{(00)}^{(1)} = 0$ and applying Lemma 6.2, we conclude

$$S_{(00)}^{(\ell+1)} = -\frac{1}{3an} (1 + O(\ell^{-1})) + O(n^{-2}),$$

where as usual the implicit constant in the $O(\ell^{-1})$ error is independent of n, ℓ but the constant in $O(n^{-2})$ may depend on ℓ . Hence, by (11.7) we find

$$S_{(00)}^{(\ell+1)} = \left(-\frac{\ell}{3n} (1 + O(\ell^{-1})) + O(n^{-2}) \right) K_{(00)}^{(\ell)}.$$

Combining this with the definition (10.3) of $S_{(00)}^{(\ell)}$ yields (11.27). To obtain (11.28), we proceed similarly. Namely, Theorem 3.2, the estimate $S_{(00)}^{(\ell)} = O(n^{-1})$, and integration by parts yield

$$\begin{aligned} \kappa_{(10)}^{(\ell+1)} &= C_W \mathbb{E} \left[\sigma(z_{i;\alpha}^{(\ell)}) \sigma'(z_{i;\alpha}^{(\ell)}) d_1 z_{i;\alpha}^{(\ell)} \right] \\ &= C_W \left\{ \langle \sigma \sigma' d_1 z \rangle_{\kappa^{(\ell)}} + \frac{1}{8} \left[\kappa_{(00)(00)}^{(\ell)} \langle \partial^4 (\sigma \sigma') d_1 z \rangle_{\kappa^{(\ell)}} + 2\kappa_{(10)(00)}^{(\ell)} \langle \partial^3 (\sigma \sigma') \rangle_{\kappa^{(\ell)}} \right] \right\} \\ &= \kappa_{(10)}^{(\ell)} \left[\chi_{\parallel}^{(\ell)} + \frac{C_W}{2} S_{(00)}^{(\ell)} \langle \partial^3 (\sigma \sigma') \rangle_{K^{(\ell)}} \right] + O(n^{-2}) \\ &\quad + \frac{C_W}{8} \left[\kappa_{(00)(00)}^{(\ell)} \langle \partial^4 (\sigma \sigma') d_1 z \rangle_{\kappa^{(\ell)}} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^3 (\sigma \sigma') \rangle_{\kappa^{(\ell)}} \right]. \end{aligned}$$

Hence,

$$\begin{aligned} S_{(10)}^{(\ell+1)} &= S_{(10)}^{(\ell)} \left[\chi_{\parallel}^{(\ell)} + \frac{C_W}{2} S_{(00)}^{(\ell)} \langle \partial^3 (\sigma \sigma') \rangle_{K^{(\ell)}} \right] + K_{(10)}^{(\ell)} \frac{C_W}{2} S_{(00)}^{(\ell)} \langle \partial^3 (\sigma \sigma') \rangle_{K^{(\ell)}} \\ &\quad + \frac{C_W}{8} \left[\kappa_{(00)(00)}^{(\ell)} \langle \partial^4 (\sigma \sigma') d_1 z \rangle_{\kappa^{(\ell)}} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^3 (\sigma \sigma') \rangle_{\kappa^{(\ell)}} \right] + O(n^{-2}). \end{aligned}$$

This recursion shows that $S_{(10)}^{(\ell)} = O(n^{-1})$ and hence we obtain

$$\begin{aligned} S_{(10)}^{(\ell+1)} &= \chi_{\parallel}^{(\ell)} S_{(10)}^{(\ell)} + K_{(10)}^{(\ell)} \frac{C_W}{2} S_{(00)}^{(\ell)} \langle \partial^3(\sigma\sigma') \rangle_{K^{(\ell)}} \\ &\quad + \frac{C_W}{8} \left[\kappa_{(00)(00)}^{(\ell)} \langle \partial^4(\sigma\sigma') d_1 z \rangle_{\kappa^{(\ell)}} + 4\kappa_{(10)(00)}^{(\ell)} \langle \partial^3(\sigma\sigma') \rangle_{\kappa^{(\ell)}} \right] + O(n^{-2}). \end{aligned}$$

The asymptotics from the formula above for $S_{(00)}^{(\ell)}$ and Proposition 11.2 we see that

$$\kappa_{(10)(00)}^{(\ell)} = -K_{(10)}^{(\ell)} \frac{C_W}{2} S_{(00)}^{(\ell)} + O\left(\frac{1}{n\ell^3}\right).$$

Hence, we find all together that

$$S_{(10)}^{(\ell+1)} = \left(1 - \frac{2}{\ell} + O(\ell^{-2})\right) S_{(10)}^{(\ell)} + O\left(\frac{x_{1;\alpha}}{nn_0\ell^3}\right) + O(n^{-2}) = O\left(\frac{x_{1;\alpha}}{nn_0\ell^2}\right) + O(n^{-2}).$$

Combining this with (11.8) confirms (11.28). We now proceed to deriving (11.29). By using Theorem 3.2 and integrating by parts

$$\begin{aligned} \kappa_{(11)}^{(\ell+1)} &= \mathbb{E} \left[\left(d_1 z_{i;\alpha}^{(\ell+1)} \right)^2 \right] \\ &= \mathbb{E} \left[\left(\sigma'(z_{i;\alpha}^{(\ell)}) d_1 z_{i;\alpha}^{(\ell)} \right)^2 \right] \\ &= \kappa_{(11)}^{(\ell)} \left[\chi_{\perp}^{(\ell)} + \frac{C_W}{2} S_{(00)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right] + C_W \left(\kappa_{(10)}^{(\ell)} \right)^2 \left[\langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} + \frac{1}{2} S_{(00)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} \right] \\ &\quad + \frac{C_W}{8} \left\{ \kappa_{(00)(00)}^{(\ell)} \left[K_{(11)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + \left(K_{(10)}^{(\ell)} \right)^2 \langle \partial^6(\sigma')^2 \rangle_{K^{(\ell)}} \right] \right. \\ &\quad \left. + 4\kappa_{(10)(00)}^{(\ell)} K_{(10)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + \left(2\kappa_{(10)(10)}^{(\ell)} + 4\kappa_{(11)(00)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right\} + O(n^{-2}). \end{aligned}$$

Comparing this with the recursion (10.7) for $K_{(11)}^{(\ell)}$ shows that $S_{(11)}^{(\ell)} = O(n^{-1})$ and yields the recursion

$$\begin{aligned} S_{(11)}^{(\ell+1)} &= \chi_{\perp}^{(\ell)} S_{(11)}^{(\ell)} + \frac{C_W}{2} K_{(11)}^{(\ell)} S_{(00)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} + O(n^{-2}) \\ &\quad + C_W \left(K_{(10)}^{(\ell)} \right)^2 \left[\langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} + \frac{1}{2} S_{(00)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} \right] + 2C_W K_{(10)}^{(\ell)} S_{(10)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \\ &\quad + \frac{C_W}{8} \left\{ \kappa_{(00)(00)}^{(\ell)} \left[K_{(11)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + \left(K_{(10)}^{(\ell)} \right)^2 \langle \partial^6(\sigma')^2 \rangle_{K^{(\ell)}} \right] \right. \\ &\quad \left. + 4\kappa_{(10)(00)}^{(\ell)} K_{(10)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + \left(2\kappa_{(10)(10)}^{(\ell)} + 4\kappa_{(11)(00)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right\}. \end{aligned}$$

Recalling the asymptotics from Lemma 11.1 and Proposition 11.2 we find to leading order in ℓ and n that

$$S_{(11)}^{(\ell+1)} = \chi_{\perp}^{(\ell)} S_{(11)}^{(\ell)} + \frac{C_W}{2} K_{(11)}^{(\ell)} S_{(00)}^{(\ell)} \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}},$$

and hence by (11.9) we obtain

$$S_{(11)}^{(\ell)} = \frac{C_W e^{-\gamma}}{3nn_0} (1 + O(\ell^{-1})) + O(n^{-2}) = \left[\frac{\ell}{3n} + O(n^{-1}) + O(n^{-2}) \right] K_{(11)}^{(\ell)}.$$

Combining this with (11.9) completes our derivation of (11.29). A very similar computation reveals

$$\begin{aligned} S_{(12)}^{(\ell+1)} &= \chi_{\perp}^{(\ell)} S_{(12)}^{(\ell)} + C_W \kappa_{(10)}^{(\ell)} \kappa_{(20)}^{(\ell)} \left[\langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} + \frac{1}{2} S_{(00)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} \right] \\ &+ \frac{C_W}{8} \left\{ \kappa_{(00)(00)}^{(\ell)} \left[K_{(12)}^{(\ell)} \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} + K_{(10)}^{(\ell)} K_{(20)}^{(\ell)} \langle \partial^6(\sigma')^2 \rangle_{K^{(\ell)}} \right] \right. \\ &\quad + 2 \left(\kappa_{(10)(00)}^{(\ell)} K_{(20)}^{(\ell)} + \kappa_{(20)(00)}^{(\ell)} K_{(10)}^{(\ell)} \right) \langle \partial^4(\sigma')^2 \rangle_{K^{(\ell)}} \\ &\quad \left. + 2 \left(2\kappa_{(12)(00)}^{(\ell)} + \kappa_{(10)(20)}^{(\ell)} \right) \langle \partial^2(\sigma')^2 \rangle_{K^{(\ell)}} \right\} + O(n^{-2}). \end{aligned}$$

Using the asymptotics from Proposition 11.2 and Lemma 6.2 we obtain

$$S_{(12)}^{(\ell)} = O\left(\frac{x_{1;\alpha} x_{2;\alpha}}{nn_0^2 \ell}\right) + O(n^{-2}).$$

Combined with (11.10) this proves (11.30). \square

12 Statements and Declarations

As required by Springer, we provide here statements on funding and competing interests. This research was partially funded by the National Science Foundation through Grants DMS-2143754, DMS-1855684, and DMS-2133806 as well support from an ONR MURI on Foundations of Deep Learning. Finally, the author declares that he has no conflicts of interest.

References

- [1] Gernot Akemann and Zdzislaw Burda. Universal microscopic correlation functions for products of independent ginibre matrices. *Journal of Physics A: Mathematical and Theoretical*, 45(46):465201, 2012.
- [2] Gernot Akemann, Zdzislaw Burda, and Mario Kieburg. Universality of local spectral statistics of products of random matrices. *Physical Review E*, 102(5):052134, 2020.
- [3] Anders Andreassen and Ethan Dyer. Asymptotics of wide convolutional neural networks. *arXiv preprint arXiv:2008.08675*, 2020.
- [4] Andrea Basteri and Dario Trevisan. Quantitative gaussian approximation of randomly initialized deep neural networks. *arXiv preprint arXiv:2203.07379*, 2022.
- [5] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
- [6] David R Brillinger. The calculation of cumulants via conditioning. *Annals of the Institute of Statistical Mathematics*, 21(1):215–218, 1969.
- [7] David R Brillinger. *Time series: data analysis and theory*. SIAM, 2001.

-
- [8] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
 - [9] Lenaic Chizat and Francis Bach. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*, 2018.
 - [10] Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pages 3036–3046, 2018.
 - [11] Omry Cohen, Or Malka, and Zohar Ringel. Learning curves for overparametrized deep neural networks: A field theory perspective. *Physical Review Research*, 3(2):023034, 2021.
 - [12] Simon S Du, Jason D Lee, Yuandong Tian, Barnabas Poczos, and Aarti Singh. Gradient descent learns one-hidden-layer cnn: Don’t be afraid of spurious local minima. *ICML*, 2018.
 - [13] Simon S. Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019.
 - [14] Ethan Dyer and Guy Gur-Ari. Asymptotics of wide networks from feynman diagrams. In *2019 ICML Workshop on Physics for Deep Learning*.
 - [15] Ronen Eldan, Dan Mikulincer, and Tselil Schramm. Non-asymptotic approximations of neural networks by gaussian processes. In *Conference on Learning Theory*, pages 1754–1775. PMLR, 2021.
 - [16] Vadim Gorin and Yi Sun. Gaussian fluctuations for products of random matrices. *arXiv preprint arXiv:1812.06532*, 2018.
 - [17] Boris Hanin. Which neural net architectures give rise to exploding and vanishing gradients? In *Advances in Neural Information Processing Systems*, 2018.
 - [18] Boris Hanin. Random neural networks in the infinite width limit as gaussian processes. *arXiv preprint arXiv:2107.01562*, 2021.
 - [19] Boris Hanin, Ryan Jeong, and David Rolnick. Deep relu networks preserve expected length. *ICLR*, 2022.
 - [20] Boris Hanin and Mihai Nica. Finite depth and width corrections to the neural tangent kernel. *ICLR 2020 and arXiv:1909.05989*, 2019.
 - [21] Boris Hanin and Mihai Nica. Products of many large random matrices and gradients in deep neural networks. *Communications in Mathematical Physics*, 376(1):287–322, 2020.
 - [22] Boris Hanin and Grigoris Paouris. Non-asymptotic results for singular values of gaussian matrix products. *Geometric and Functional Analysis*, 31(2):268–324, 2021.

- [23] Boris Hanin and David Rolnick. How to start training: The effect of initialization and architecture. In *Advances in Neural Information Processing Systems*, pages 571–581, 2018.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [25] Tianyu He, Darshil Doshi, and Andrey Gromov. Understanding layer normalization in deep neural networks. *Bulletin of the American Physical Society*, 2022.
- [26] Donald O. Hebb. *The organization of behavior: A neuropsychological theory*. Wiley, 1949.
- [27] Sepp Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116, 1998.
- [28] Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, Jürgen Schmidhuber, et al. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- [29] Jiaoyang Huang and Horng-Tzer Yau. Dynamics of deep neural networks and neural tangent hierarchy. In *International Conference on Machine Learning*, pages 4542–4551. PMLR, 2020.
- [30] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- [31] Sergey Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. In *Advances in neural information processing systems*, pages 1945–1953, 2017.
- [32] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- [33] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [34] Maciej Koch-Janusz and Zohar Ringel. Mutual information, neural networks and the renormalization group. *Nature Physics*, 14(6):578–582, 2018.
- [35] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [36] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. *ICML 2018 and arXiv:1711.00165*, 2018.

-
- [37] Jaehoon Lee, Lechao Xiao, Samuel S Schoenholz, Yasaman Bahri, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *arXiv preprint arXiv:1902.06720*, 2019.
- [38] Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari. The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218*, 2020.
- [39] Mufan Li, Mihai Nica, and Dan Roy. The future is log-gaussian: Resnets and their infinite-depth-and-width limit at initialization. *Advances in Neural Information Processing Systems*, 34, 2021.
- [40] Mufan Bill Li, Mihai Nica, and Daniel M Roy. The neural covariance sde: Shaped infinite depth-and-width networks at initialization. *arXiv preprint arXiv:2206.02768*, 2022.
- [41] Chaoyue Liu, Libin Zhu, and Mikhail Belkin. Toward a theory of optimization for over-parameterized systems of non-linear equations: the lessons of deep learning. *arXiv preprint arXiv:2003.00307*, 2020.
- [42] Dang-Zheng Liu, Dong Wang, and Yanhui Wang. Lyapunov exponent, universality and phase transition for products of random matrices. *arXiv preprint arXiv:1810.00433*, 2018.
- [43] James Martens, Andy Ballard, Guillaume Desjardins, Grzegorz Swirszcz, Valentin Dalibard, Jascha Sohl-Dickstein, and Samuel S Schoenholz. Rapid training of deep neural networks without skip connections or normalization layers using deep kernel shaping. *arXiv preprint arXiv:2110.01765*, 2021.
- [44] Alexander G de G Matthews, Mark Rowland, Jiri Hron, Richard E Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*, 2018.
- [45] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- [46] Gadi Naveh, Oded Ben David, Haim Sompolinsky, and Zohar Ringel. Predicting the outputs of finite deep neural networks trained with noisy gradients. *Physical Review E*, 104(6):064301, 2021.
- [47] Gadi Naveh and Zohar Ringel. A self consistent theory of gaussian processes captures feature learning effects in finite cnns. *Advances in Neural Information Processing Systems*, 34, 2021.
- [48] Radford M Neal. Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pages 29–53. Springer, 1996.
- [49] Roman Novak, Lechao Xiao, Jaehoon Lee, Yasaman Bahri, Greg Yang, Jiri Hron, Daniel A Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian deep convolutional networks with many channels are gaussian processes. *arXiv preprint arXiv:1810.05148*, 2018.

-
- [50] Daniel S Park, Jascha Sohl-Dickstein, Quoc V Le, and Samuel L Smith. The effect of network width on stochastic gradient descent and generalization: an empirical study. *arXiv preprint arXiv:1905.03776*, 2019.
- [51] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *Advances in neural information processing systems*, pages 3360–3368, 2016.
- [52] Maithra Raghu, Ben Poole, Jon M. Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, pages 2847–2854, 2017.
- [53] Daniel A Roberts, Sho Yaida, and Boris Hanin. *The Principles of Deep Learning Theory: An Effective Theory Approach to Understanding Neural Networks*. Cambridge University Press, 2022.
- [54] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [55] Grant M Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *arXiv preprint arXiv:1805.00915*, 2018.
- [56] Inbar Seroussi and Zohar Ringel. Separation of scales and a thermodynamic description of feature learning in some cnns. *arXiv preprint arXiv:2112.15383*, 2021.
- [57] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020.
- [58] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of deep neural networks. *Mathematics of Operations Research*, 2021.
- [59] Jay M Ver Hoef. Who invented the delta method? *The American Statistician*, 66(2):124–127, 2012.
- [60] Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in over-parametrized models. *arXiv preprint arXiv:2002.09277*, 2020.
- [61] Sho Yaida. Non-gaussian processes and neural networks at finite widths. *MSML*, 2020.
- [62] Greg Yang. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv preprint arXiv:1902.04760*, 2019.
- [63] Greg Yang. Tensor programs i: Wide feedforward or recurrent neural networks of any architecture are gaussian processes. *arXiv preprint arXiv:1910.12478*, 2019.
- [64] Greg Yang and Edward J Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In *International Conference on Machine Learning*, pages 11727–11737. PMLR, 2021.

- [65] Greg Yang and Sam S Schoenholz. Deep mean field theory: Layerwise variance and width variation as methods to control gradient explosion. 2018.
- [66] Jacob Zavatore-Veth and Cengiz Pehlevan. Exact marginal prior distributions of finite bayesian neural networks. *Advances in Neural Information Processing Systems*, 34, 2021.
- [67] Guodong Zhang, Aleksandar Botev, and James Martens. Deep learning without shortcuts: Shaping the kernel with tailored rectifiers. *arXiv preprint arXiv:2203.08120*, 2022.

A Proof of Theorem 2.2

Our proof of Theorem 2.2 closely follows the proof of Theorem 1.2 in [18]. Let us first recall the notation and assumptions. We fix $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ such that

- There exists $r \geq 1$ so that the r -th derivative of σ belongs to L^∞ .
- There exist $c, c' > 0$ so that

$$\left\| e^{-cx^2-c'} \frac{d^r}{dx^r} \sigma(x) \right\|_{L^\infty} < \infty.$$

We take $C_b \geq$ and $C_W > 0$ and consider a random depth L neural network with input dimension n_0 , output dimension n_{L+1} , hidden layer widths satisfying

$$\exists c, C > 0 \text{ s.t. } cn \leq n_1, \dots, n_L \leq Cn, \quad n \gg 1,$$

non-linearity σ and random weights and biases as in (2.2). We also fix a finite collection

$$x_{\mathcal{A}} := \{x_\alpha, \quad \alpha \in \mathcal{A}\}$$

of distinct network inputs as well as an integer m and study for each ℓ the random vectors

$$D^{\leq r} z_{\mathcal{A}}^{(\ell)} := (D^{\leq r} z_{i;\mathcal{A}}, \quad i = 1, \dots, m),$$

where

$$D^{\leq r} z_{i;\mathcal{A}} := \left(D_\alpha^J z_{i;\alpha}^{(\ell)}, \quad \alpha \in \mathcal{A}, i = 1, \dots, m, |J| \leq r \right)$$

are the derivatives of $z_{i;\mathcal{A}}^{(\ell)}$ of order at most r . Our goal is to show that, as $n \rightarrow \infty$, the joint distribution of the random vectors $D^{\leq r} z_{i;\mathcal{A}}^{(\ell)}$ converges to that of centered jointly Gaussian vectors that are independent for different i and satisfy

$$\lim_{n \rightarrow \infty} \text{Cov} \left(D_{\alpha_1}^{J_1} z_{i;\alpha_1}^{(\ell)}, D_{\alpha_2}^{J_2} z_{i;\alpha_2}^{(\ell)} \right) = D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} K_{\alpha_1 \alpha_2}^{(\ell)},$$

where

$$K_{\alpha_1 \alpha_2}^{(\ell+1)} = C_b + C_W \langle \sigma(z_\alpha) \sigma(z_\beta) \rangle_{K^{(\ell)}}, \quad K_{\alpha_1 \alpha_2}^{(1)} = C_b + C_W \sum_{j=1}^{n_0} x_{j;\alpha_1} x_{j;\alpha_2}$$

is the infinite width covariance from Theorem 2.2. To prove this, let us denote by $\mathcal{F}^{(\ell)}$ the sigma algebra generated by the weights and biases in layer up to and including ℓ . Observe that, conditional on $\mathcal{F}^{(\ell)}$, we have that $D^{\leq r} z_{i;\mathcal{A}}^{(\ell)}$ are already independent for different i and

that, since the weights and biases are Gaussian, each is a centered Gaussian with conditional covariance

$$\text{Cov} \left(D_{\alpha_1}^{J_1} z_{i;\alpha_1}^{(\ell+1)}, D_{\alpha_2}^{J_2} z_{i;\alpha_2}^{(\ell+1)} \mid \mathcal{F}^{(\ell)} \right) = D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)},$$

where

$$\Sigma_{\alpha_1 \alpha_2}^{(\ell)} = C_b + \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \sigma \left(z_{j;\alpha_1}^{(\ell)} \right) \sigma \left(z_{j;\alpha_2}^{(\ell)} \right).$$

Thus, by the continuous mapping theorem, it suffices to show that for any multi-indices J_1, J_2 with $|J_i| \leq r$ and any $\alpha_1, \alpha_2 \in \mathcal{A}$ we have

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} \right] \quad \text{exists and is finite} \quad (\text{A.1})$$

and

$$\lim_{n \rightarrow \infty} \text{Var} \left[D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} \right] = 0. \quad (\text{A.2})$$

We establish (A.1) and (A.2) by induction on ℓ the following more general statement.

Proposition A.1. *Denote by $N(n_0, r)$ the number of derivatives of order at most r in n_0 variables. Consider any measurable function $f : \mathbb{R}^{N(n_0, r) \times |\mathcal{A}|} \rightarrow \mathbb{R}$ that is polynomially bounded, and define*

$$\mathcal{O}_f^{(\ell)} := \frac{1}{n_\ell} \sum_{j=1}^{n_\ell} f \left(D^{\leq r} z_{j;\mathcal{A}} \right).$$

Then,

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\mathcal{O}_f^{(\ell)} \right] \quad \text{exists and is finite} \quad (\text{A.3})$$

and

$$\lim_{n \rightarrow \infty} \text{Var} \left[\mathcal{O}_f^{(\ell)} \right] = 0. \quad (\text{A.4})$$

Proof. We proceed by induction, starting with $\ell = 1$. Since weights and biases are Gaussian, the vectors $D^{\leq r} z_{i;\mathcal{A}}^{(1)}$ are independent for all i and jointly Gaussian. The polynomial growth assumption on f show the moments of $f(x)$ are finite if x is Gaussian. This allows us to apply the SLLN to conclude both (A.3) and (A.4).

Let us now assume we have proved (A.1) and (A.2) for layers $1, \dots, \ell$. We start by fixing any polynomially bounded f and establishing (A.1) at layer $\ell + 1$. We have

$$\mathbb{E} \left[\mathcal{O}_f^{(\ell+1)} \right] = \mathbb{E} \left[f \left(D^{\leq r} z_{1;\mathcal{A}} \right) \right].$$

As above, conditionl on $\mathcal{F}^{(\ell)}$, we have the following equality in distribution:

$$D^{\leq r} z_{1;\mathcal{A}} \stackrel{d}{=} \left(\Sigma^{\leq r, (\ell)} \right)^{1/2} Z, \quad Z \sim \mathcal{N} \left(0, I_{N(n_0, r) \times |\mathcal{A}|} \right) \quad (\text{A.5})$$

where Z is independent of the conditional covariance matrix

$$\Sigma^{\leq r, (\ell)} = \left(D_{\alpha_1}^{J_1} D_{\alpha_2}^{J_2} \Sigma_{\alpha_1 \alpha_2}^{(\ell)} \right)_{\alpha_1, \alpha_2 \in \mathcal{A}}.$$

The key observation is that each entry of $\Sigma^{\leq r,(\ell)}$ is of the form $\mathcal{O}_f^{(\ell)}$ for polynomially bounded f . Hence, we may apply the inductive hypothesis to conclude that there exists a matrix $\bar{\Sigma}^{\leq r,(\ell)}$ such that the following convergence in distribution holds

$$\Sigma^{\leq r,(\ell)} \xrightarrow{d} \bar{\Sigma}^{\leq r,(\ell)} \quad \text{as } n \rightarrow \infty.$$

The polynomial growth assumption on f together with the Skorohod representation theorem and dominated convergence show that

$$\lim_{n \rightarrow \infty} \mathbb{E} [f(D^{\leq r} z_{1;\mathcal{A}})] = \mathbb{E} \left[f \left(\left(\bar{\Sigma}^{\leq r,(\ell)} \right)^{1/2} Z \right) \right] =: \bar{\mathcal{O}}_f^{(\ell+1)} \quad \text{exists and is finite.}$$

This proves (A.3). To show (A.4), we proceed similarly. Namely, we have

$$\text{Var} \left[\mathcal{O}_f^{(\ell+1)} \right] = \frac{1}{n_{\ell+1}} \text{Var} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right) \right] + \left(1 - \frac{1}{n_{\ell+1}} \right) \text{Cov} \left(f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right), f \left(D^{\leq r} z_{2;\mathcal{A}}^{(\ell+1)} \right) \right).$$

Note that

$$\text{Var} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right) \right] \leq \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^{n_\ell} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right) \right]^2 \right].$$

Hence, since f^2 is also polynomially bounded we have already shown that (A.3) holds at layer $\ell + 1$, we see that

$$\text{Var} \left[\mathcal{O}_f^{(\ell+1)} \right] = \text{Cov} \left(f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right), f \left(D^{\leq r} z_{2;\mathcal{A}}^{(\ell+1)} \right) \right) + O(n^{-1}).$$

Next, using that conditional on $\mathcal{F}^{(\ell)}$ the vectors $D^{\leq r} z_{i;\mathcal{A}}^{(\ell+1)}$ are independent for different i we conclude from the law of total covariance that

$$\text{Cov} \left(f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right), f \left(D^{\leq r} z_{2;\mathcal{A}}^{(\ell+1)} \right) \right) \leq \text{Var} \left[\mathbb{E} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right) \mid \mathcal{F}^{(\ell)} \right] \right].$$

Combining the equality in distribution (A.5) with the polynomial growth condition on f and the dominated convergence theorem we find

$$\lim_{n \rightarrow \infty} \text{Var} \left[\mathbb{E} \left[f \left(D^{\leq r} z_{1;\mathcal{A}}^{(\ell+1)} \right) \mid \mathcal{F}^{(\ell)} \right] \right] = \text{Var} \left[\mathbb{E} \left[f \left(\left(\bar{\Sigma}^{\leq r,(\ell)} \right)^{1/2} Z \right) \right] \right] = 0.$$

This completes the proof that (A.4) holds at infinite width, establishing Proposition A.1. \square

B Infinite Width Analysis of Tanh-like Non-linearities

The purpose of this section is to derive some basic properties of the infinite width variance recursion

$$\kappa_{\alpha\alpha}^{(\ell+1)} = C_b + C_W \langle \sigma(z)^2 \rangle_{\kappa_{\alpha\alpha}^{(\ell)}}. \quad (\text{B.1})$$

We abbreviate

$$\sigma_j := \frac{1}{j!} \frac{d^j}{dx^j} \Big|_{x=0} \sigma(x)$$

and consider here the case when σ that is a tanh-like non-linearity in the sense that σ satisfies:

- σ is smooth at 0 with $\sigma_1 \neq 0$
- σ has the opposite sign as its second derivative

$$\text{for almost every } z, \operatorname{sgn}(\sigma(z)\sigma''(z)) = -1. \quad (\text{B.2})$$

Note that this forces $\sigma_2 = 0$ and

$$a := -\frac{6\sigma_3}{\sigma_1} > 0.$$

- σ is sub-linear:

$$\exists C > 0 \text{ s.t. } \forall z \in \mathbb{R} \quad |\sigma(z)| \leq |\sigma_1 z|, \quad (\text{B.3})$$

- σ is controlled by its first few non-zero Taylor series coefficients at 0:

$$\exists C \geq 0 \text{ s.t. } \forall z \geq 0, \quad \sigma_1 z + \sigma_3 z^3 \leq \sigma(z) \leq \sigma_1 z + \sigma_3 z^3 + C z^4 \quad (\text{B.4})$$

We will be interested in understanding the recursion (B.1) at criticality in the sense defined in §4.1. Specifically, we remind the reader that this means we choose C_b, C_W so that

$$\begin{aligned} \exists K_* \geq 0 \quad \text{s.t.} \quad & K_* = C_b + C_W \langle \sigma^2(z) \rangle_{K_*} \\ & \chi_{\parallel}(K_*) = \frac{C_W}{2} \langle \partial^2(\sigma(z)^2) \rangle_{K_*} = 1 \\ & \chi_{\perp}(K_*) = C_W \langle (\sigma'(z))^2 \rangle_{K_*} = 1. \end{aligned}$$

Before stating our main result (Proposition B.1), let us explain intuitively what we expect. First of all, as we shall see in Proposition B.1, tanh-like non-linearities requires $K_* = 0$ for criticality. Second, by Taylor expanding the recursion (2.4) around small values of $K_{\alpha\alpha}^{(\ell)}$ we find

$$K_{\alpha\alpha}^{(\ell+1)} = K_{\alpha\alpha}^{(\ell)} - a \left(K_{\alpha\alpha}^{(\ell)} \right)^2 + O \left(\left(K_{\alpha\alpha}^{(\ell)} \right)^3 \right).$$

This is well-approximated by the ODE

$$\frac{d}{dt} K(t) = -a K(t)^2,$$

whose solution is

$$K(t) = \left(at + \frac{1}{K(0)} \right)^{-1}.$$

This form for the solution has two important properties that we will check in Proposition B.1 hold for the actual solution $K_{\alpha\alpha}^{(\ell)}$ to the discrete difference equation (2.4):

- At large t , $K(t)$ tends to zero like $1/at$ plus an error of size roughly $O(t^{-2})$.
- The leading order behavior of $K(t)$ at large t is independent of the initial condition.

Proposition B.1. *If σ is a tanh-like non-linearity in the sense defined above then criticality is achieved for σ only with*

$$K_* = 0, \quad C_b = 0, \quad \text{and} \quad C_W = \sigma_1^{-2}. \quad (\text{B.5})$$

Moreover, for every $\delta \in (0, 1)$ we have

$$K_{\alpha\alpha}^{(0)} > 0 \quad \Rightarrow \quad \sup_{\ell \geq 1} \ell^{2-\delta} \left| K_{\alpha\alpha}^{(\ell)} - \frac{1}{a\ell} \right| < \infty. \quad (\text{B.6})$$

Proof. Note that for any $K \geq 0$ we have

$$\chi_{\parallel}(K) = \chi_{\perp}(K) + C_W \langle \sigma(z)\sigma''(z) \rangle_K. \quad (\text{B.7})$$

Hence, at criticality, we must have

$$\langle \sigma(z)\sigma''(z) \rangle_{K_*} = 0.$$

But due to assumption (B.2) we have

$$K > 0 \quad \Rightarrow \quad \langle \sigma(z)\sigma''(z) \rangle_K < 0.$$

Thus, we indeed find that we must have $K_* = 0$ at criticality. Hence, in light of (B.7) criticality is equivalent to the system of equations

$$K_* = 0 = C_b + C_W \sigma(0)^2, \quad \chi_{\parallel}(0) = \chi_{\perp}(0) = C_W \langle (\sigma'(z))^2 \rangle_0 = C_W \sigma_1^2 = 1.$$

This system has a unique solution:

$$C_b = 0, \quad C_W = \sigma_1^{-2},$$

completing the proof of the criticality conditions (B.5). Let us now establish (B.6). First note that at criticality the sub-linearity condition (B.3) guarantees that for all $\delta > 0$ there exists $c_{\delta} \in (0, 1)$ such that

$$K > \delta \quad \Rightarrow \quad C_W \langle \sigma(z)^2 \rangle_K < (1 - c_{\delta}) \langle z^2 \rangle_K = (1 - c_{\delta})K.$$

Hence, for all $K, \delta > 0$ there exists $\ell_0 \geq 1$ such that

$$K_{\alpha\alpha}^{(0)} \leq K \quad \Rightarrow \quad K_{\alpha\alpha}^{(\ell)} \leq \delta \quad \forall \ell \geq \ell_0. \quad (\text{B.8})$$

In particular, $K_{\alpha\alpha}^{(\ell)}$ is monotonically decreasing and converges to $K_* = 0$ as ℓ grows. Let us now define for each $\ell \geq 1$

$$K_{\alpha\alpha}^{(\ell)} =: \frac{1}{a\ell} + \epsilon^{(\ell)}, \quad a := -6 \frac{\sigma_3}{\sigma_1} > 0,$$

where a is positive due to (B.2). Note that since $K_{\alpha\alpha}^{(\ell)}$ tends to zero with ℓ , so does $\epsilon^{(\ell)}$. Let us agree that for any $t \in \mathbb{R}$ the symbol t^+ (resp. t^-) means that for ℓ sufficiently large we may make the constant t^+ (resp. t^-) arbitrary close to t from above (resp. below). In order to prove (B.6), we start with the following elementary estimate.

Lemma B.2. *For all $\ell \geq 1$, we have*

$$\epsilon^{(\ell+1)} \geq -\frac{1}{a\ell^2(\ell+1)} + \epsilon^{(\ell)} \left(1 - \frac{2}{\ell} - a\epsilon^{(\ell)}\right). \quad (\text{B.9})$$

Further, there exists a constant $C > 0$ depending only on σ with the following property. For all $K > 0$ there exists a constant $\ell_0 \geq 1$ so that if $K_{\alpha\alpha}^{(0)} \leq K$, then for all $\delta \in (0, 1)$ we have

$$\epsilon^{(\ell+1)} \leq \frac{C}{\ell^3} + \epsilon^{(\ell)} \left(1 - \frac{2-\delta}{\ell}\right), \quad \forall \ell \geq \ell_\delta := \max \left\{ \frac{C}{\delta}, \frac{2C}{a}, \ell_0 \right\}. \quad (\text{B.10})$$

Proof. Plugging in the estimates (B.4) into the recursion (B.1) yields for some $C > 0$ depending only on σ

$$\epsilon^{(\ell+1)} \leq \frac{C}{\ell^3} + \epsilon^{(\ell)} \left[1 - \frac{2}{\ell} + \frac{C}{\ell^2}\right] + \left(\epsilon^{(\ell)}\right)^2 \left(-a + \frac{C}{\ell}\right) + C(\epsilon^{(\ell)})^3$$

Note that for all $\ell \geq 2C/a$ we have $-a + C/\ell \leq 0$. Hence, for all $\ell \geq \max \{2C/a, C/\delta\}$ we have

$$\epsilon^{(\ell+1)} \leq \frac{C}{\ell^3} + \epsilon^{(\ell)} \left[1 - \frac{2-\delta}{\ell}\right] + C(\epsilon^{(\ell)})^3.$$

Moreover, if $\epsilon^{(\ell)} \leq 0$, then $(\epsilon^{(\ell)})^3 \leq 0$. If on the other hand $\epsilon^{(\ell)} \geq 0$, then from (B.8) we find that there for all $K > 0$ there exists ℓ_0 so that $(\epsilon^{(\ell)})^3 \leq a(\epsilon^{(\ell)})^2/4$ for all $\ell \geq \ell_0$. Hence, in all cases, for each $\delta \in (0, 1)$ if $\ell \geq \max \{2C/a, C/\delta, \ell_0\}$, we find

$$\epsilon^{(\ell+1)} \leq \frac{C}{\ell^3} + \epsilon^{(\ell)} \left[1 - \frac{2-\delta}{\ell}\right],$$

as claimed. The lower bound follows from a similar but simpler computation. \square

Fix $\delta \in (0, 1)$. The relation (B.10), together with Lemma 6.2, show that for all $K > 0$ there exists some $C' > 0$ depending on δ, σ, K such that if $K_{\alpha\alpha}^{(0)} \leq K$ then

$$\epsilon^{(\ell+1)} \leq \sum_{\ell'=\ell_\delta}^{\ell} \frac{C}{\ell'^3} \prod_{\ell''=\ell'+1}^{\ell} \left(1 - \frac{2-\delta}{\ell''}\right) \leq C' \left[\frac{1}{\ell^2} + \epsilon^{(\ell_\delta)} \frac{1}{\ell^{2-\delta}}\right].$$

This shows that

$$\forall \delta \in (0, 1) \exists \ell_\delta \geq 1 \text{ s.t. } \epsilon^{(\ell)} \leq \frac{1}{\ell^{2-\delta}} \quad \forall \ell \geq \ell_\delta. \quad (\text{B.11})$$

To conclude (B.6) it therefore remains to deduce that

$$\forall K_1, K_2 > 0, \delta \in (0, 1) \exists \ell_\delta \geq 1 \text{ s.t. } K_1 < K_{\alpha\alpha}^{(0)} < K_2 \implies \epsilon^{(\ell)} \geq -\frac{1}{\ell^{2-\delta}} \quad \forall \ell \geq \ell_\delta. \quad (\text{B.12})$$

To aid with this, we will need the following

Lemma B.3. *For any $\delta \in (0, 1)$ there exists $\ell_\delta \geq 1$ with the property that if $\ell \geq \ell_\delta$ then*

$$\epsilon^{(\ell)} \geq -\ell^{-2+\delta} \implies \epsilon^{(\ell+1)} \geq -(\ell+1)^{-2+\delta}.$$

Proof. Suppose $\epsilon^{(\ell)} \geq -\ell^{-2+\delta}$. The lower bound in (B.9) yields for some $C, C' > 0$

$$\begin{aligned} \epsilon^{(\ell+1)} + (\ell+1)^{-2+\delta} &\geq (\ell+1)^{-2+\delta} \left[1 - (1+\ell^{-1})^{2-\delta} \right] - 2\ell^{-3+\delta} - C\ell^{-4+2\delta} \\ &\geq \delta\ell^{-3+\delta} - C'(\ell^{-3} + \ell^{-4+2\delta}), \end{aligned}$$

which is non-negative for all ℓ sufficiently large. \square

We are now in a position to establish (B.12). In light of the previous Lemma we need only consider the case when

$$\forall \delta \in (0, 1) \exists \ell_\delta \geq 1 \quad \text{s.t.} \quad \epsilon^{(\ell_\delta)} < -\ell_\delta^{2-\delta}.$$

Note that in light of the upper bound (B.11) we find that for all $\delta \in (0, 1)$ there exists $\ell_\delta \geq 1$ and $C_\delta > 0$ so that for all $\ell \geq \ell_\delta$ we have

$$K_{\alpha\alpha}^{(\ell+1)} \geq K_{\alpha\alpha}^{(\ell)} \left(1 - aK_{\alpha\alpha}^{(\ell)} \right) \geq K_{\alpha\alpha}^{(\ell)} \left(1 - a \left(-\frac{1}{a\ell} + C_\delta \ell^{-2+\delta} \right) \right) = K_{\alpha\alpha}^{(\ell)} \left(1 - \frac{1}{\ell} - aC_\delta \ell^{-2+\delta} \right).$$

Hence, assuming $K_2 \geq K_{\alpha\alpha}^{(0)} \geq K_1 > 0$, we may iterate this inequality to find that there exists $c > 0$ depending on K_1, K_2 and $\ell_0 \geq 1$ so that

$$K_{\alpha\alpha}^{(\ell)} \geq \frac{c}{a\ell} \quad \forall \ell \geq \ell_0.$$

Hence, since $\epsilon^{(\ell)} < 0$ for all $\ell \geq \ell_\delta$ we find for all $\ell \geq \max\{\ell_0, \ell_\delta\}$ that

$$-a(\epsilon^{(\ell)})^2 \geq \epsilon^{(\ell)} \frac{1-c}{\ell}$$

Substituting this into the lower bound (B.9), we find that for all $\ell \geq \max\{\ell_0, \ell_\delta\}$

$$\epsilon^{(\ell+1)} \geq -\frac{C'}{\ell^3} + \epsilon^{(\ell)} \left(1 - \frac{1+c}{\ell} \right).$$

Since $\epsilon^{(\ell_\delta)} < 0$, we see by applying Lemma 6.2 that there exists $C > 0$ so that for all $\ell \geq \max\{\ell_0, \ell_\delta\}$

$$\epsilon^{(\ell+1)} \geq -\frac{C}{\ell^{1+c}}.$$

But now we can bootstrap this estimate. Indeed, for any $\delta \in (0, 1)$ we substitute this into the lower bound (B.9) to find that for all ℓ sufficiently large

$$\epsilon^{(\ell+1)} \geq -\frac{C'}{\ell^3} + \epsilon^{(\ell)} \left(1 - \frac{2-\delta}{\ell} \right).$$

Again applying Lemma 6.2 yields that for all ℓ sufficiently large

$$\epsilon^{(\ell+1)} \geq -\frac{C}{\ell^{2-\delta}}.$$

This completes the proof. \square

C Exact Solutions for 1-homogeneous activations

In this appendix, we collect several known computations related to the distribution of neuron activations in random fully connected networks with 1-homogeneous activations. Specifically, we fix a one homogeneous non-linearity

$$\sigma(t) = (a_+ \mathbf{1}_{t>0} + a_- \mathbf{1}_{t<0})t$$

and consider a random fully connected neural network with input dimension n_0 , output dimension n_{L+1} , hidden layer widths n_1, \dots, n_ℓ , and non-linearity σ that is tuned to criticality in the sense that

$$C_b = 0, \quad C_W = \frac{2}{a_+^2 + a_-^2}.$$

Our first task is to derive in §C.1 a known exact formula for the infinite width covariance $K_{\alpha\beta}^{(\ell+1)}$ as a function of $K_{\alpha\alpha}^{(\ell)}, K_{\alpha\beta}^{(\ell)}, K_{\beta\beta}^{(\ell)}$. Then, in Section §C.2, we sketch a derivation of the limiting distribution (4.7) of a neuron pre-activation in the double scaling limit $n, L \rightarrow \infty, L/n \rightarrow \gamma$.

C.1 Covariance Propagation in Random Fully Connected 1-homogeneous Networks

In this section, we consider two network inputs x_α, x_β of the same norm:

$$K_{\alpha\alpha}^{(0)} = \frac{1}{n_0} \|x_\alpha\|^2 = K = \frac{1}{n_0} \|x_\beta\|^2 = K_{\beta\beta}^{(0)}, \quad K > 0. \quad (\text{C.1})$$

Let us define

$$\epsilon_{\alpha\beta}^{(\ell)} := \frac{1 - \text{Corr}_{\alpha\beta}^{(\ell)}}{2}, \quad \text{Corr}_{\alpha\beta}^{(\ell)} := \frac{K_{\alpha\beta}^{(\ell)}}{\left(K_{\alpha\alpha}^{(\ell)} K_{\beta\beta}^{(\ell)}\right)^{1/2}}$$

Our goal is to derive the following explicit recursion for $\epsilon_{\alpha\beta}^{(\ell+1)}$ in terms of $\epsilon_{\alpha\beta}^{(\ell)}$. This derivation follows the approach in §5.5 [53]. To the author's knowledge, the following formula (or really something equivalent) was first derived in [?].

Proposition C.1 (Correlation propagation for 1-homogeneous activation functions). *At criticality, we have the following exact formula:*

$$1 - 2\epsilon_{\alpha\beta}^{(\ell+1)} = \frac{2C_W(a_+ - a_-)^2}{\pi} \left[\frac{1}{2} \sqrt{\epsilon_{\alpha\beta}^{(\ell)}(1 - \epsilon_{\alpha\beta}^{(\ell)})} + \left(\frac{1}{2} - \epsilon_{\alpha\beta}^{(\ell)} \right) \cos^{-1} \left(\sqrt{\epsilon_{\alpha\beta}^{(\ell)}} \right) \right] + C_W a_+ a_- (1 - 2\epsilon_{\alpha\beta}^{(\ell)}) \quad (\text{C.2})$$

In particular, taking $\epsilon_{\alpha\beta}^{(\ell)}$ small we find

$$\epsilon_{\alpha\beta}^{(\ell+1)} = \epsilon_{\alpha\beta}^{(\ell)} - \frac{4}{3\pi} \left(\epsilon_{\alpha\beta}^{(\ell)} \right)^{3/2} + O \left(\left(\epsilon_{\alpha\beta}^{(\ell)} \right)^{5/2} \right).$$

Hence, as $\ell \rightarrow \infty$,

$$\epsilon_{\alpha\beta}^{(\ell)} = \frac{2}{3\pi} \ell^{-2} (1 + o(1)).$$

Proof. We have from Theorem 2.2 that

$$K_{\alpha\beta}^{(\ell+1)} = C_b + C_W \langle \sigma(z_\alpha) \sigma(z_\beta) \rangle_{K^{(\ell)}}, \quad (\text{C.3})$$

where we recall that the brackets above mean the average with respect to the Gaussian distribution

$$\begin{pmatrix} z_\alpha \\ z_\beta \end{pmatrix} \sim \mathcal{N} \left(0, \begin{pmatrix} K_{\alpha\alpha}^{(\ell)} & K_{\alpha\beta}^{(\ell)} \\ K_{\alpha\beta}^{(\ell)} & K_{\beta\beta}^{(\ell)} \end{pmatrix} \right).$$

Since we are at criticality, we have

$$C_b = 0, \quad C_W = \frac{2}{a_+^2 + a_-^2}$$

and that moreover

$$K_{\alpha\alpha}^{(\ell)} = K_{\beta\beta}^{(\ell)} = K,$$

where K is the constant from (C.1). Our first step is to change from the Gaussian variables z_α, z_β to the new Gaussian variables

$$\xi = \frac{z_\alpha + z_\beta}{2\sqrt{K}}, \quad \eta = \frac{z_\alpha - z_\beta}{2\sqrt{K}}.$$

We have

$$z_\alpha = \sqrt{K}(\xi + \eta), \quad z_\beta = \sqrt{K}(\xi - \eta).$$

Moreover, writing

$$\epsilon := \epsilon_{\alpha\beta}^{(\ell)} = \frac{1}{2} \left(1 - \frac{K_{\alpha\beta}^{(\ell)}}{(K_{\alpha\alpha}^{(\ell)} K_{\beta\beta}^{(\ell)})^{1/2}} \right)$$

we find

$$\text{Var}[\xi] = 1 - \epsilon, \quad \text{Var}[\eta] = \epsilon, \quad \text{Cov}[\xi, \eta] = 0.$$

Hence, the right hand side of the recursion (C.3) reads

$$C_W K \int_{\mathbb{R}} \int_{\mathbb{R}} \sigma \left((1 - \epsilon)^{1/2} \xi + \epsilon^{1/2} \eta \right) \sigma \left((1 - \epsilon)^{1/2} \xi - \epsilon^{1/2} \eta \right) \exp \left[-\frac{1}{2} (\xi^2 + \eta^2) \right] \frac{d\xi d\eta}{2\pi}.$$

Using the definition of σ yields

$$\begin{aligned} & \sigma \left((1 - \epsilon)^{1/2} \xi + \epsilon^{1/2} \eta \right) \sigma \left((1 - \epsilon)^{1/2} \xi - \epsilon^{1/2} \eta \right) \\ &= (a_+ \mathbf{1}_{\xi+\eta>0} + a_- \mathbf{1}_{\xi+\eta<0}) (a_+ \mathbf{1}_{\xi-\eta>0} + a_- \mathbf{1}_{\xi-\eta<0}) ((1 - \epsilon)\xi^2 - \epsilon\eta^2). \end{aligned}$$

Changing variables $(\xi, \eta) \rightarrow (-\xi, -\eta)$ inside the integral and averaging yields

$$\begin{aligned} K_{\alpha\beta}^{(\ell+1)} &= C_W K a_+ a_- (1 - 2\epsilon) \\ &+ \frac{C_W K (a_+ - a_-)^2}{2} \int_{\mathbb{R}^2} \mathbf{1}_{(1-\epsilon)\xi^2 - \epsilon\eta^2 > 0} ((1 - \epsilon)\xi^2 - \epsilon\eta^2) \exp \left[-\frac{1}{2} (\xi^2 + \eta^2) \right] \frac{d\xi d\eta}{2\pi}. \end{aligned}$$

Passing to polar coordinates and explicitly computing the resulting integral is now straightforward and completes the derivation of (C.2). \square

C.2 Full Distribution of Neuron Pre-activations at a Single Input and the Derivation of (4.7)

Our purpose in this section is to briefly recall an exact formula for the full distribution of a neuron pre-activation $z_{i;\alpha}^{(L+1)}$. For this, note that since $x_\alpha \mapsto z_\alpha^{(L+1)}$ is piecewise linear and the event that the Jacobian $J_{x_\alpha} z_\alpha^{(L+1)}$ is not well-defined at x_α has probability zero, we may write

$$z_\alpha^{(L+1)} = J_{x_\alpha} z_\alpha^{(L+1)} x_\alpha.$$

Next,

$$J_{x_\alpha} z_\alpha^{(L+1)} = W^{(L+1)} D^{(L)} W^{(L)} \dots D^{(1)} W^{(1)}, \quad (\text{C.4})$$

where $W^{(\ell)}$ are simply the weight matrices and

$$D^{(\ell)} := \text{Diag} \left(\sigma'(z_{i;\alpha}^{(\ell)}), i = 1, \dots, n_\ell \right).$$

Arguing exactly as in Proposition 2 of [21], we have the following equality in distribution:

$$D^{(L)} W^{(L)} \dots D^{(1)} W^{(1)} \stackrel{d}{=} A \widehat{D}^{(L)} W^{(L)} \dots \widehat{D}^{(1)} W^{(1)},$$

where A is a diagonal matrix with iid ± 1 entries on the diagonal that is independent of $W^{(\ell)}$ and the diagonal matrices

$$\widehat{D}^{(\ell)} = \text{Diag} \left(\underbrace{a_+ \xi_i^{(\ell)} + a_- (1 - \xi_i^{(\ell)})}_{=: d_i^{(\ell)}}, i = 1, \dots, n_\ell \right), \quad \xi_i^{(\ell)} \sim \text{Bernoulli}(1/2) \text{ iid}.$$

Combining this with (C.4) and recalling that the entries of $W^{(L+1)}$ are iid centered Gaussians with variance C_W/n_L yields

$$z_{i;\alpha}^{(L+1)} \stackrel{d}{=} Z_1 \cdot \left(\frac{C_W}{n_L} \right)^{1/2} \left\| \widehat{D}^{(L)} W^{(L)} \dots \widehat{D}^{(1)} W^{(1)} x_\alpha \right\|,$$

where $Z_1 \sim \mathcal{N}(0, 1)$ is independent of $\widehat{D}^{(\ell)}, W^{(\ell)}, i = 1, \dots, L$. Further, due to the right orthogonal invariance of the Gaussian matrices $W^{(\ell)}$ and the normalization that the variance of the entries of $W^{(\ell)}$ is $C_W/n_{\ell-1}$, we have that

$$\begin{aligned} & \log \left[\left(\frac{C_W}{n_L} \right)^{1/2} \left\| \widehat{D}^{(L)} W^{(L)} \dots \widehat{D}^{(1)} W^{(1)} x_\alpha \right\| \right] \\ & \stackrel{d}{=} \frac{1}{2} \log \left[\frac{C_W}{n_0} \|x_\alpha\|^2 \right] + \sum_{\ell=1}^L \frac{1}{2} \log \left[\frac{C_W}{n_L} \left\| \widehat{D}^{(\ell)} \widehat{W}^{(\ell)} u^{(\ell)} \right\|^2 \right] \end{aligned}$$

where $u^{(\ell)} \in \mathbb{R}^{n_{\ell-1}}$ is collection of deterministic unit vectors and $\widehat{W}^{(\ell)}$ are independent random matrices with iid standard Gaussian entries. The summands on the previous line are independent and are each distributed like the logarithm of a randomly weighted χ^2 random viable:

$$\frac{C_W}{n_\ell} \left\| \widehat{D}^{(\ell)} W^{(\ell)} u^{(\ell)} \right\|^2 \stackrel{d}{=} \frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \left(d_j^{(\ell)} \right)^2 \left(Z_j^{(\ell)} \right)^2,$$

where $Z_i^{(\ell)} \sim \mathcal{N}(0, 1)$ are iid and independent of $d_i^{(\ell)}$. Putting this all together, we find that

$$z_{i;\alpha}^{(L+1)} \stackrel{d}{=} \left(\frac{C_W}{n_0} \|x_\alpha\|^2 \right)^{1/2} \cdot Z_1 \cdot \prod_{\ell=1}^L \left(\frac{C_W}{n_\ell} \right)^{1/2} \left\| \widehat{D}^{(\ell)} W^{(\ell)} u^{(\ell)} \right\|$$

is a product of $L + 1$ independent random variables. Moreover, a direct computation shows that

$$\begin{aligned} \mathbb{E} \left[\log \left[\frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \left(d_i^{(\ell)} \right)^2 \left(Z_i^{(\ell)} \right)^2 \right] \right] &= -\frac{1}{2} \text{Var} \left[\frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \left(d_i^{(\ell)} \right)^2 \left(Z_i^{(\ell)} \right)^2 \right] + O(n_\ell^{-2}) \\ &= -\frac{1}{2n_\ell} \left(6 \frac{a_+^4 + a_-^4}{(a_+^2 + a_-^2)^2} - 1 \right) + O(n_\ell^{-2}) \end{aligned}$$

and also that

$$\begin{aligned} \text{Var} \left[\log \left[\frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \left(d_i^{(\ell)} \right)^2 \left(Z_i^{(\ell)} \right)^2 \right] \right] &= \text{Var} \left[\frac{C_W}{n_\ell} \sum_{j=1}^{n_\ell} \left(d_i^{(\ell)} \right)^2 \left(Z_i^{(\ell)} \right)^2 \right] + O(n_\ell^{-2}) \\ &= \frac{1}{n_\ell} \left(6 \frac{a_+^4 + a_-^4}{(a_+^2 + a_-^2)^2} - 1 \right) + O(n_\ell^{-2}). \end{aligned}$$

Combining the preceding two estimates, taking $n, L \rightarrow \infty$ with $L/n \rightarrow \gamma$ and applying the CLT yields

$$\lim_{\substack{n, L \rightarrow \infty \\ L/n \rightarrow \gamma \in [0, \infty)}} z_{i;\alpha}^{(L)} \xrightarrow{d} \left(\frac{C_W}{n_0} \|x_\alpha\|^2 \right)^{1/2} Z_1 \exp \left[-\mu(\gamma, a_+, a_-) + \sigma(\gamma, a_+, a_-) Z_2 \right],$$

where

$$\mu(\gamma, a_+, a_-) = \sigma^2(\gamma, a_+, a_-) := \frac{\gamma}{4} \left(6 \frac{a_+^4 + a_-^4}{(a_+^2 + a_-^2)^2} - 1 \right), \quad Z_1, Z_2 \sim \mathcal{N}(0, 1) \text{ iid.}$$

This is precisely the statement of (4.7).