



PAPER

OPEN ACCESS

Feature learning and generalization in deep networks with orthogonal weights

RECEIVED
25 March 2025REVISED
11 June 2025ACCEPTED FOR PUBLICATION
21 July 2025PUBLISHED
7 August 2025

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.

Hannah Day^{1,5} , Yonatan Kahn^{1,2,5,*} and Daniel A Roberts^{3,4} ¹ Department of Physics, University of Illinois Urbana-Champaign, Urbana, IL 61801, United States of America² Department of Physics, University of Toronto; and Vector Institute, Toronto, Ontario, Canada³ Center for Theoretical Physics and Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America⁴ Sequoia Capital, Menlo Park, CA 94025, United States of America⁵ Equal contribution.

* Author to whom any correspondence should be addressed.

E-mail: yf.kahn@utoronto.ca, hjday2@illinois.edu and drob@mit.edu**Keywords:** neural tangent kernel, feature learning, orthogonal weights**Abstract**

Fully-connected deep neural networks with weights initialized from independent Gaussian distributions can be tuned to criticality, which prevents the exponential growth or decay of signals propagating through the network. However, such networks still exhibit fluctuations that grow linearly with the depth of the network, which may impair the training of networks with width comparable to depth. We show analytically that rectangular networks with tanh activations and weights initialized from the ensemble of orthogonal matrices have corresponding preactivation fluctuations which are independent of depth, to leading order in inverse width. Moreover, we demonstrate numerically that, at initialization, all correlators involving the neural tangent kernel (NTK) and its descendants at leading order in inverse width—which govern the evolution of observables during training—saturate at a depth of ~ 20 , rather than growing without bound as in the case of Gaussian initializations. We speculate that this structure preserves finite-width feature learning while reducing overall noise, thus improving both generalization and training speed in deep networks with depth comparable to width. We provide some experimental justification by relating empirical measurements of the NTK to the superior performance of deep non-linear orthogonal networks trained under full-batch gradient descent on the MNIST and CIFAR-10 classification tasks.

1. Introduction

The initialization distributions for neural network parameters are key hyperparameters that can affect training and performance. A large body of work has focused on how *critical initialization hyperparameters* can mitigate the exploding and vanishing gradient problem by preventing the exponential growth or decay of preactivations. Poole *et al* [1] first noted the exponential growth with depth, [2] introduced the notion of criticality, [3–5] studied finite-width effects in correlation functions, [6, 7] related the notions of criticality to training dynamics, [8] performed empirical studies of different initialization schemes, [9] proposed a critical initialization which allowed feature learning in the infinite-width limit, and [10] identified a one-parameter family of initializations with various perturbative and non-perturbative limits. Modern developments have given rise to optimizers such as Adam [11] and AdamW [12] that are relatively insensitive to initializations, as well as techniques like LayerNorm [13] that explicitly enforce order-one values for propagating signals, guaranteeing critical initializations when implemented with residual connections [14] (see also [15]).

The minimal model of a wide-but-finite multi-layer perceptron (MLP) trained with full-batch gradient descent (GD) provides a theoretically-tractable playground to study the effect of initializations on the statistics and dynamics of ensembles of networks. In such a setup, studied in detail in [7] for the case of weights and biases drawn i.i.d. from zero-mean Gaussian distributions, MLPs of depth L and width n tuned

to criticality still exhibit fluctuations which scale as L/n , characterized by the connected 4-point correlator of preactivations $V \sim \mathbb{E}[z^4]_{\text{conn.}}$, where the expectation is taken over an ensemble of networks. If $L/n \sim 1$, these fluctuations grow to order-one sizes, potentially inhibiting the efficacy of training. On the other hand, finite-width effects are crucial for feature learning, as encapsulated by the evolution of the neural tangent kernel [16], (NTK) \hat{H} , which governs GD dynamics [17, 18]. Correlators of the NTK fluctuations and derivatives of the NTK also scale like L/n . The competition between the ‘good’ fluctuations which allow feature learning and the ‘bad’ fluctuations which simply inject noise leads [7] to speculate that there exists an optimal aspect ratio $L/n \ll 1$ for Gaussian initializations, which can be calculated using techniques from information theory.

In this paper, we provide both analytical and numerical evidence that initializing the weights of an MLP from a distribution of orthogonal matrices yields drastically different behavior than Gaussian initializations for the same network architectures. In particular, orthogonal initializations yield correlators which saturate at moderate depth, rather than continuing to scale linearly with depth. Our study builds on previous work in [19] for linear networks and [20, 21] for nonlinear networks, which demonstrated the depth independence of the singular value variance of the input–output Jacobian for orthogonal networks at criticality, but only when paired with linear or tanh activation functions. Subsequent work confirmed that orthogonal initializations can also improve training in deep convolutional neural networks (CNNs) [22], have provably superior training dynamics to Gaussian initializations in deep linear MLPs [23], and can achieve improved training in the $L/n \sim 1$ regime [24]. Our work goes beyond these studies by showing the depth independence of the normalized single-input 4-point correlator, which measures a different facet of network fluctuations than the input–output Jacobian, while working entirely in the finite-width regime to leading order in $1/n$. Furthermore, we study for the first time the correlators of the NTK and its descendants, which govern feature learning at finite width, and empirically demonstrate that these (suitably normalized) correlators become independent of depth at $L \sim 20$. This provides a plausible theoretical link between the previously-observed superior generalization of deep orthogonal networks, and the dynamics of the NTK under GD at finite width, confirmed by our numerical experiments which show that the GD training dynamics become independent of depth for $L \sim 20$.

This paper is organized as follows. In section 2, we briefly motivate the relevance of studying MLPs—despite the success of more elaborate architectures—both as a unit of modern architectures such as transformers, and as stand-alone neural networks in situations where speed or interpretability is paramount. In section 3, we set out our notation and review the relevant parts of the formalism of [7]. In section 4, we warm up by computing the exact single-input preactivation distribution for orthogonal MLPs with linear activation functions, finding the satisfying result that these networks preserve the norm of layer outputs for all depths. In section 5, we compute the dimensionless single-input connected 4-point correlator at leading order in $1/n$ and show that it grows as L/n for ReLU activation functions, but is depth-independent for activation functions in the $K^* = 0$ universality class of [7]. In section 6, we empirically measure the full set of single-input NTK correlators at $\mathcal{O}(1/n)$ at initialization in an ensemble of MLPs and compare their behavior for Gaussian and orthogonal weights. In section 7, we perform numerical experiments by training MLPs of various L and n on the MNIST and CIFAR-10 classification problems using full-batch GD with weight and bias learning rates scaled according to the prescriptions of [7]. Consistent with previous studies, we find that orthogonal initializations perform essentially identically to Gaussian initializations in the $L/n \ll 1$ regime, but have superior performance in the $L/n \sim 1$ regime, where we also observe suppressed NTK fluctuations during training for orthogonal networks. We conclude in section 8 with an outlook and plans for future work. Further measurements of multi-input correlators are provided in appendix 8, and variations on our training experiments comprise appendix B.

2. Motivation for studying MLP training dynamics

We focus on MLPs in this work as a tractable model of a deep learning architecture, amenable to analysis with techniques borrowed from high-energy physics [7] or statistical physics [25]. While the MLP is one of the oldest examples of a neural network, the deep learning revolution began in earnest with the development of architectures such as the CNN [26] for image recognition, followed a few years later by the transformer [27] for natural language processing. Both CNNs and transformers incorporate inductive biases about their training data through architectural choices such as weight-sharing, self-attention, and masking, but the same principles of criticality at initialization are relevant for training both CNNs [22] and transformers [28]. Furthermore, MLPs are a crucial building block of transformers, and thus a better understanding of the dynamics of MLP training can complement studies which focus on the dynamics and interpretability of the attention layers (see e.g. [29]). Even with the development of more advanced architectures, MLPs are still a demonstrated tool for discovery in physics [30, 31]. As such, we expect that the

approach we use to studying MLPs, and in particular the effect of orthogonal initializations on training dynamics, can be relevant to both MLP-only architectures and MLPs as components of larger networks.

3. Notation and preliminaries

Throughout this paper, we will use the notation of [7], which we briefly review here. The MLP architecture is defined by a forward iteration equation

$$z_i^{(1)} = b_i^{(1)} + \sum_{j=1}^{n_0} W_{ij}^{(1)} x_j, \quad z_i^{(\ell+1)} = b_i^{(\ell+1)} + \sum_{j=1}^{n_\ell} W_{ij}^{(\ell+1)} \sigma(z_j^{(\ell)}), \quad (3.1)$$

where $z_j^{(\ell)}$ is the preactivation at layer ℓ , $W_{ij}^{(\ell+1)}$ and $b_j^{(\ell+1)}$ are the weights and biases at layer $\ell + 1$ with $1 \leq \ell \leq L - 1$, n_ℓ is the width of layer ℓ , and σ is the element-wise activation function. The input layer ($\ell = 0$) will be denoted $\mathbf{x} \equiv \mathbf{z}^{(0)}$; all boldface vectors at layer ℓ have n_ℓ components, and $\|\mathbf{z}\| \equiv \sqrt{\mathbf{z} \cdot \mathbf{z}}$. Neural indices at each layer are denoted with Latin subscripts running from 1 to n_ℓ , and we abbreviate $\sigma(z_j^{(\ell)}) \equiv \sigma_j^{(\ell)}$ and similarly with derivatives of activation functions, $\sigma'(z_j^{(\ell)}) \equiv \sigma_j'^{(\ell)}$. Greek indices will refer to samples from the dataset $\mathcal{D} \equiv \{\mathbf{x}_\alpha\}$, with $z_{i;\alpha} \equiv z_i(\mathbf{x}_\alpha)$ the preactivations corresponding to input \mathbf{x}_α .

When we consider statistics of networks at initialization, we will envision \mathcal{D} as fixed, such that $\mathbb{E}[\dots]$ denotes an expectation over the only source of stochasticity in the model, namely the initialization distributions. All probability distributions are taken to be conditional on the data, which may be a single input or the entire set \mathcal{D} . The weights and biases in each layer will be taken to have zero mean, with variances defined by

$$\mathbb{E} \left[W_{i_1 j_1}^{(\ell)} W_{i_2 j_2}^{(\ell)} \right] = \delta_{i_1 i_2} \delta_{j_1 j_2} \frac{C_W}{n_{\ell-1}}, \quad (3.2)$$

$$\mathbb{E} \left[b_{i_1}^{(\ell)} b_{i_2}^{(\ell)} \right] = \delta_{i_1 i_2} C_b, \quad (3.3)$$

where for simplicity, we assume C_W and C_b are constant across layers. In particular, we will use equation (3.2) to define scaled $n \times n$ orthogonal matrices O satisfying $O^T O = C_W \mathbb{I}_{n \times n}$. In order to have square weight matrices to more easily compare orthogonal and Gaussian weights, we will take all hidden layers to have width n , and abbreviate $n_\ell \equiv n$ for $\ell = 1, 2, \dots, L - 1$. Consequently, we will drop the layer label on W and b when computing expectations since their distributions are identical at each layer.

If the weight distribution were Gaussian, then equation (3.2) would fully specify the initialization hyperparameters because all higher-point moments of weights in the same layer can be expressed in terms of C_W by Wick's theorem. In particular, the fourth moment would take the form

$$\mathbb{E} \left[W_{i_1 j_1} W_{i_2 j_2} W_{i_3 j_3} W_{i_4 j_4} \right] = \frac{C_W^2}{n^2} \left[(12)(34)_i (12)(34)_j + (13)(24)_i (13)(24)_j + (14)(23)_i (14)(23)_j \right] \quad (3.4)$$

where we are using an abbreviated notation $(ab)(cd)_i \equiv \delta_{i_a i_b} \delta_{i_c i_d}$. However, for weights drawn uniformly from the Haar measure on scaled orthogonal matrices with $n > 2$, the fourth moment is instead

$$\mathbb{E} \left[W_{i_1 j_1} W_{i_2 j_2} W_{i_3 j_3} W_{i_4 j_4} \right] = C_W^2 \begin{pmatrix} (12)(34)_i \\ (13)(24)_i \\ (14)(23)_i \end{pmatrix}^T \begin{pmatrix} \mathcal{W}[1, 1] & \mathcal{W}[2] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[1, 1] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[2] & \mathcal{W}[1, 1] \end{pmatrix} \begin{pmatrix} (12)(34)_j \\ (13)(24)_j \\ (14)(23)_j \end{pmatrix}, \quad (3.5)$$

where

$$\mathcal{W}[1, 1] = \frac{n+1}{(n-1)n(n+2)}, \quad \mathcal{W}[2] = \frac{-1}{(n-1)n(n+2)} \quad (3.6)$$

are combinatorial factors computed from the Weingarten functions \mathcal{W} , first introduced in [32] and further studied in [33, 34] (see [35, 36] for reviews). Note that for large n , $\mathcal{W}[1, 1] \sim 1/n^2$ and $\mathcal{W}[2] \sim 1/n^3$, so that at large width, orthogonal and Gaussian weights with the same variance have identical fourth moments up to $1/n$ corrections⁶.

⁶ Throughout this paper, we will avoid discussing the infinite-width limit whenever possible because of a rather subtle order-of-limits issue. Orthogonal matrices at all finite n retain higher-point correlations between entries, but in the strict $n \rightarrow \infty$ limit, the matrix entries become independent Gaussian-distributed variables and the matrix is no longer orthogonal. In other words, the orthogonality constraint involves summing n terms which becomes ill-defined as $n \rightarrow \infty$ unless $1/n$ effects are included. To avoid this complication, we will keep n finite in all calculations.

The kernel $K^{(\ell)}$ and the 4-point vertex $V^{(\ell)}$ are defined as 2- and 4-point connected correlators of preactivations at layer ℓ to leading order in $1/n$ across an ensemble of MLPs⁷,

$$\mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell)} z_{i_2; \alpha_2}^{(\ell)} \right] = \delta_{i_1 i_2} K_{\alpha_1 \alpha_2}^{(\ell)} + \mathcal{O}(1/n), \quad (3.7)$$

$$\mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell)} z_{i_2; \alpha_2}^{(\ell)} z_{i_3; \alpha_3}^{(\ell)} z_{i_4; \alpha_4}^{(\ell)} \right]_{\text{conn.}} = \frac{1}{n} \left[(12)(34)_i V_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell)} + (2 \text{ sim.}) \right] + \mathcal{O}(1/n^2). \quad (3.8)$$

In equation (3.7), the $\mathcal{O}(1/n)$ refers to the fact that K is the leading part of the 2-point correlator; we will refer to the full 2-point correlator to all orders in $1/n$ as the metric, discussed further in section 5.4. The neural index tensor structure of equation (3.7) follows from equation (3.2) because it involves an expectation over two weights, which is identical for Gaussian and orthogonal initializations. Similarly, the tensor structure of equation (3.8) follows from equation (3.4) for Gaussian weights, where the two similar terms in equation (3.8) have the neural indices and sample indices matched up as in the displayed term, and we will show in section 4 below that it also holds for orthogonal weights. The structure of the forward equation (3.1) allows a type of dimensional analysis (known in the physics literature as ‘power counting’) where we treat the preactivations z as dimensional variables. Thus, when computing the 4-point vertex as a function of depth, we will use the normalized combination

$$\tilde{V}^{(\ell)} \equiv \frac{V^{(\ell)}}{(K^{(\ell)})^2}, \quad (3.9)$$

which is dimensionless because it has 4 powers of z in the numerator and 4 in the denominator⁸.

Finally, because the weight distribution $p(W)$ is still nearly-Gaussian up to $1/n$ corrections, we can follow the prescription of [7] to integrate out the weights and biases and derive marginal distributions $p(\mathbf{z}^{(\ell+1)} | \mathbf{z}^{(\ell)})$ which are also nearly-Gaussian. In doing so we will encounter Gaussian expectations with covariance matrix g , which we denote with angle brackets,

$$\langle F(z_{\alpha_1}, \dots, z_{\alpha_m}) \rangle_g \equiv \int \left[\frac{\prod_{\alpha \in \mathcal{D}} dz_{\alpha}}{\sqrt{\det(2\pi g)}} \right] \exp \left(-\frac{1}{2} \sum_{\beta_1, \beta_2 \in \mathcal{D}} (g^{-1})^{\beta_1 \beta_2} z_{\beta_1} z_{\beta_2} \right) F(z_{\alpha_1}, \dots, z_{\alpha_m}). \quad (3.10)$$

4. Exact preactivation distribution for linear orthogonal networks

Here we derive the fact that the probability distribution for preactivations of a linear orthogonal network of width n , with zero biases ($b_i^{(\ell)} = 0$) and a single input, is a uniform distribution on the sphere S^{n-1} with radius equal to the norm of the input. Though this result is rather obvious from the definition of an orthogonal matrix, the calculational techniques serve as a useful warm-up for the nonlinear networks considered in section 5, where a full closed-form expression is not available. We begin by explicitly calculating the 2- and 4-point correlators, then generalize to m -point correlators in order to construct the characteristic function of the probability distribution, and finally Fourier-transform the characteristic function to obtain the probability distribution itself. We will find it instructive to compare this derivation to the results of [39] for the identical architecture with Gaussian initializations.

For orthogonal weights, the 2-point correlator of preactivations of the first hidden layer, with a single input \mathbf{x} taken to have the same width n as the hidden layers, is given by

$$\mathbb{E} \left[z_{i_1}^{(1)} z_{i_2}^{(1)} \right] = \sum_{j_1, j_2} \mathbb{E} [W_{i_1 j_1} W_{i_2 j_2}] x_{j_1} x_{j_2} = \delta_{i_1 i_2} \frac{C_W}{n} \|\mathbf{x}\|^2, \quad (4.1)$$

where as mentioned we have set all biases to zero, and the second equality makes use of equation (3.2). Noting the index structure, we make the ansatz

$$\mathbb{E} \left[z_{i_1}^{(\ell)} z_{i_2}^{(\ell)} \right] = \delta_{i_1 i_2} K^{(\ell)}, \quad (4.2)$$

⁷ The nomenclature ‘vertex’ for V refers to a useful diagrammatic notation for calculations, see [37, 38] for examples.

⁸ We have deliberately kept equation (3.9) schematic by leaving off the sample indices, because in general, for $\alpha_1 \neq \alpha_2 \neq \alpha_3 \neq \alpha_4$, dimensional analysis provides no unique prescription to choose the indices in K . We will show in appendix A.1 that for the purposes of determining scaling with width and depth, these choices do not matter.

so that 2-point correlator for an arbitrary layer is

$$\mathbb{E} \left[z_{i_1}^{(\ell+1)} z_{i_2}^{(\ell+1)} \right] = \sum_{j_1, j_2} \mathbb{E} [W_{i_1 j_1} W_{i_2 j_2}] \mathbb{E} \left[z_{j_1}^{(\ell)} z_{j_2}^{(\ell)} \right] = \delta_{i_1 i_2} \frac{C_W}{n} \sum_j K^{(\ell)} = C_W \delta_{i_1 i_2} K^{(\ell)}. \quad (4.3)$$

Thus we have the kernel recursion

$$K^{(\ell+1)} = C_W K^{(\ell)} = (C_W)^{\ell+1} K^{(0)} \quad (4.4)$$

with initial condition

$$K^{(0)} \equiv \frac{1}{n} \|\mathbf{x}\|^2, \quad (4.5)$$

which is identical to the recursion for Gaussian-initialized linear networks [7]. In particular, choosing $C_W = 1$ ensures criticality where $K^{(\ell)} = K^{(0)}$.

The 4-point correlator of preactivations of the first hidden layer is given by

$$\begin{aligned} \mathbb{E} \left[z_{i_1}^{(1)} z_{i_2}^{(1)} z_{i_3}^{(1)} z_{i_4}^{(1)} \right] &= \sum_{j_1, j_2, j_3, j_4} \mathbb{E} [W_{i_1 j_1} W_{i_2 j_2} W_{i_3 j_3} W_{i_4 j_4}] x_{j_1} x_{j_2} x_{j_3} x_{j_4} \\ &= C_W^2 \sum_{j_1, j_2, j_3, j_4} x_{j_1} x_{j_2} x_{j_3} x_{j_4} \\ &\quad \times \begin{pmatrix} (12) (34)_i \\ (13) (24)_i \\ (14) (23)_i \end{pmatrix}^T \begin{pmatrix} \mathcal{W}[1, 1] & \mathcal{W}[2] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[1, 1] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[2] & \mathcal{W}[1, 1] \end{pmatrix} \begin{pmatrix} (12) (34)_j \\ (13) (24)_j \\ (14) (23)_j \end{pmatrix} \\ &= C_W^2 ((12) (34)_i + (13) (24)_i + (14) (23)_i) (\mathcal{W}[1, 1] + 2\mathcal{W}[2]) \left(\|\mathbf{x}\|^2 \right)^2 \\ &= \frac{C_W^2}{n^2 + 2n} ((12) (34)_i + (13) (24)_i + (14) (23)_i) \left(\|\mathbf{x}\|^2 \right)^2, \end{aligned} \quad (4.6)$$

where the second equality makes use of equation (3.5). Again noting the index structure, we make the ansatz

$$\mathbb{E} \left[z_{i_1}^{(\ell)} z_{i_2}^{(\ell)} z_{i_3}^{(\ell)} z_{i_4}^{(\ell)} \right] = ((12) (34)_i + (13) (24)_i + (14) (23)_i) G_4^{(\ell)} \quad (4.7)$$

so that the 4-point correlator of an arbitrary layer is

$$\begin{aligned} \mathbb{E} \left[z_{i_1}^{(\ell+1)} z_{i_2}^{(\ell+1)} z_{i_3}^{(\ell+1)} z_{i_4}^{(\ell+1)} \right] &= \sum_{j_1, j_2, j_3, j_4} \mathbb{E} [W_{i_1 j_1} W_{i_2 j_2} W_{i_3 j_3} W_{i_4 j_4}] \mathbb{E} \left[z_{j_1}^{(\ell)} z_{j_2}^{(\ell)} z_{j_3}^{(\ell)} z_{j_4}^{(\ell)} \right] \\ &= C_W^2 \sum_{j_1, j_2, j_3, j_4} \left((12) (34)_j + (13) (24)_j + (14) (23)_j \right) G_4^{(\ell)} \\ &\quad \times \begin{pmatrix} (12) (34)_i \\ (13) (24)_i \\ (14) (23)_i \end{pmatrix}^T \begin{pmatrix} \mathcal{W}[1, 1] & \mathcal{W}[2] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[1, 1] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[2] & \mathcal{W}[1, 1] \end{pmatrix} \begin{pmatrix} (12) (34)_j \\ (13) (24)_j \\ (14) (23)_j \end{pmatrix} \\ &= C_W^2 ((12) (34)_i + (13) (24)_i + (14) (23)_i) (\mathcal{W}[1, 1] + 2\mathcal{W}[2]) (n^2 + 2n) G_4^{(\ell)} \\ &= C_W^2 ((12) (34)_i + (13) (24)_i + (14) (23)_i) G_4^{(\ell)}. \end{aligned} \quad (4.8)$$

Note that the $1/(n^2 + 2n)$ factor from the sum of rows of the Weingarten matrix exactly cancels with the $n^2 + 2n$ from summing over the Kronecker deltas from the Wick contraction pairings, whereas the corresponding result for a Gaussian initialization would instead have a remaining factor of $(n + 2)/n$. Thus, for orthogonal initializations, the G_4 recursion for an orthogonal distribution is independent of width:

$$G_4^{(\ell+1)} = C_W^2 G_4^{(\ell)} = (C_W^2)^{\ell+1} G_4^{(0)}, \quad (4.9)$$

with initial condition

$$G_4^{(0)} \equiv \frac{1}{n(n+2)} \left(\|\mathbf{x}\|^2 \right)^2 = \frac{n}{n+2} \left(K^{(0)} \right)^2. \quad (4.10)$$

Furthermore, the 4-point vertex (i.e. the connected 4-point correlator) now takes on the simple form

$$V^{(\ell)} \equiv G_4^{(\ell)} - \left(K^{(\ell)}\right)^2 = -\frac{2}{n+2} \left(K^{(\ell)}\right)^2 \quad (\text{orthogonal}). \tag{4.11}$$

Compared with the analogous result for Gaussian initializations [7],

$$V^{(\ell)} = \frac{2(\ell-1)}{n} \left(K^{(\ell)}\right)^2 + \mathcal{O}(1/n^2) \quad (\text{Gaussian}), \tag{4.12}$$

the 4-point vertex for orthogonal initializations is independent of depth if K is tuned to a fixed point, $K^{(\ell)} = K^*$, while the same quantity grows linearly with depth for Gaussian initializations.

To generalize the above derivation to higher-point correlators, we use the fact that any 2 m -point expectation of orthogonal matrix elements can be written as an analogous Weingarten matrix [35], whose rows or columns each sum to $[n(n+2)(n+4)\cdots(n+2m-2)]^{-1}$. The inverse of this factor is proportional to the combinatorial factor which sums Wick contractions of m indices with values from 1 to n , defined as

$$c_{2m}(n) = \frac{n(n+2)(n+4)\cdots(n+2m-2)}{n^m} \tag{4.13}$$

in [7]. Thus, these two factors will always cancel when evaluating 2 m -point expectations of preactivations after the first layer. This somewhat magical cancellation occurs because the Weingarten matrix is, by definition, the pseudo-inverse of the Gram matrix of the vector space spanned by all $(2m-1)!!$ pair partitions of 2 m indices [35].

Now, defining the 2 m -point correlator $G_{2m}^{(\ell)}$ through the ansatz

$$\mathbb{E} \left[z_{i_1}^{(\ell)} z_{i_2}^{(\ell)} \dots z_{i_{2m}}^{(\ell)} \right] = \left(\sum_{\text{all pairings}} \delta_{i_{k_1} i_{k_2}} \delta_{i_{k_3} i_{k_4}} \dots \delta_{i_{k_{2m-1}} i_{k_{2m}}} \right) G_{2m}^{(\ell)}, \tag{4.14}$$

such that $G_2^{(\ell)} \equiv K^{(\ell)}$ (and $G_{2m+1}^{(\ell)} = 0$), we can compute the 2 m -point expectation for all m :

$$\begin{aligned} \mathbb{E} \left[z_{i_1}^{(\ell+1)} z_{i_2}^{(\ell+1)} \dots z_{i_{2m}}^{(\ell+1)} \right] &= \sum_{j_1, j_2, \dots, j_{2m}} \mathbb{E} \left[W_{i_1 j_1} W_{i_2 j_2} \dots W_{i_{2m} j_{2m}} \right] \mathbb{E} \left[z_{j_1}^{(\ell)} z_{j_2}^{(\ell)} \dots z_{j_{2m}}^{(\ell)} \right] \\ &= C_W^m \left(\sum_{\text{all pairings}} \delta_{i_{k_1} i_{k_2}} \delta_{i_{k_3} i_{k_4}} \dots \delta_{i_{k_{m-1}} i_{k_m}} \right) G_{2m}^{(\ell)}. \end{aligned} \tag{4.15}$$

Thus we have the recursion

$$G_{2m}^{(\ell+1)} = C_W^m G_{2m}^{(\ell)} = C_W^{m(\ell+1)} G_{2m}^{(0)} \quad (\text{orthogonal}), \tag{4.16}$$

where the Weingarten combinatorial factor appears in the initial condition:

$$G_{2m}^{(0)} = \frac{n^m}{n(n+2)(n+4)\cdots(n+2m-2)} \left(K^{(0)}\right)^m. \tag{4.17}$$

Critical initialization with $C_W = 1$ ensures that this recursion is trivial, $G_{2m}^{(\ell)} = G_{2m}^{(0)}$, while other choices of C_W lead to exponential growth or decay. The analogous recursion for Gaussian initializations, from Ch. 3 of [7], involves a product of $c_{2m}(n)$ and thus is considerably more complicated.

Having computed all of the moments of $z^{(\ell)}$, the characteristic function $\varphi_{z^{(\ell)}}(u)$ may now be constructed:

$$\varphi_{z^{(\ell)}}(\mathbf{u}) \equiv \mathbb{E} \left[e^{i\mathbf{u} \cdot \mathbf{z}^{(\ell)}} \right] = 1 - \frac{1}{2!} \sum_{i_1, i_2} u_{i_1} u_{i_2} \mathbb{E} \left[z_{i_1}^{(\ell)} z_{i_2}^{(\ell)} \right] + \frac{1}{4!} \sum_{i_1, i_2, i_3, i_4} u_{i_1} u_{i_2} u_{i_3} u_{i_4} \mathbb{E} \left[z_{i_1}^{(\ell)} z_{i_2}^{(\ell)} z_{i_3}^{(\ell)} z_{i_4}^{(\ell)} \right] - \dots \tag{4.18}$$

All terms contain only even powers of z since the odd moments vanish, and each expectation yields $(2m-1)!!$ Wick pairings and a Weingarten combinatorial factor from the initial condition of the G_{2m} recursion. Rearranging some factors, we obtain the power series representation

$$\varphi_{z^{(\ell)}}(u) = \sum_{m=0}^{\infty} \frac{1}{(n/2)(n/2+1)\cdots(n/2+(m-1))} \frac{(-nC_W^\ell K^{(0)} u^2/4)^m}{m!}, \tag{4.19}$$

where $u = \|\mathbf{u}\|$. Note that the characteristic function is purely radial, depending only on u , which makes sense because φ should be invariant under orthogonal transformations of \mathbf{z} . Furthermore, the power series takes the form of a confluent hypergeometric limit function,

$$\varphi_{\mathbf{z}^{(\ell)}}(u) = {}_0F_1\left(\frac{n}{2}; -\frac{1}{4}u^2 n C_W^\ell K^{(0)}\right), \tag{4.20}$$

which can also be expressed as a Bessel function,

$$\varphi_{\mathbf{z}^{(\ell)}}(u) = \left(\frac{u}{2}\sqrt{n C_W^\ell K^{(0)}}\right)^{-\frac{n-2}{2}} \Gamma(n/2) J_{\frac{n-2}{2}}\left(u\sqrt{n C_W^\ell K^{(0)}}\right). \tag{4.21}$$

This Bessel form is useful for computing the integral in the next paragraph.

Proceeding on, we take the Fourier transform to recover the probability distribution for $\mathbf{z}^{(\ell)}$. Since the characteristic function is radial, we may perform the $n - 1$ angular integrals in an identical fashion to [39] using Hankel transform relations. Letting $z^{(\ell)} = \|\mathbf{z}^{(\ell)}\|$, we have

$$\begin{aligned} p_n(\mathbf{z}^{(\ell)}|\mathbf{x}) &= (2\pi)^{-n/2} \left(z^{(\ell)}\right)^{\frac{2-n}{2}} \int_0^\infty du u u^{\frac{n-2}{2}} J_{\frac{n-2}{2}}(z^{(\ell)}u) \varphi_{\mathbf{z}^{(\ell)}}(u) \\ &= \pi^{-\frac{n}{2}} \left(C_W^\ell K^{(0)}\right)^{-\frac{n-2}{4}} n^{-\frac{n+2}{4}} \Gamma\left(1 + \frac{n}{2}\right) \left(z^{(\ell)}\right)^{-\frac{n}{2}} \delta\left(z^{(\ell)} - \sqrt{n C_W^\ell K^{(0)}}\right) \text{ (orthogonal),} \end{aligned} \tag{4.22}$$

where the delta function arises from the orthogonality relation for Bessel functions of like order. Thus, the exact distribution of preactivations at layer ℓ is simply a uniform distribution on the sphere S^{n-1} of radius $\sqrt{n C_W^\ell K^{(0)}}$, since all the factors apart from the delta function rearrange to cancel the volume factor $d\Omega_{n-1}$. If we further choose critical initialization, $C_W = 1$, then $p_n(\mathbf{z}^{(\ell)}|\mathbf{x})$ becomes independent of ℓ and preserves the norm of the input data, since $\sqrt{n K^{(0)}} = \|\mathbf{x}\|$. Successive layers of the linear network simply apply random rotations and/or reflections to the preactivations without changing their magnitude. Of course, this is exactly what one should expect from multiplying an input vector \mathbf{x} by successive random Haar-distributed orthogonal matrices, which is what this linear network does.

It is interesting to compare the simple interpretation of equation (4.22) with the analogous result for a rectangular network with Gaussian initializations from [39], translated into our notation:

$$p_n(\mathbf{z}^{(\ell)}|\mathbf{x}) = \frac{1}{(\Gamma(n/2))^{\ell-1} (2^\ell \pi n C_W^\ell K^{(0)})^{n/2}} G_{0,\ell}^{\ell,0} \left(\frac{(z^{(\ell)})^2}{2^\ell n C_W^\ell K^{(0)}} \middle| \begin{matrix} - \\ 0, 0, \dots, 0 \end{matrix} \right) \text{ (Gaussian)} \tag{4.23}$$

where $G_{0,\ell}^{\ell,0}$ is a Meijer G-function. In this case, even at criticality with $C_W = 1$, the dependence on depth is explicit in the order of the G-function, and expanding perturbatively in $1/n$ yields non-Gaussianities that grow with ℓ , as derived in Ch. 3 of [7]:

$$G_{2m}^{(\ell)} = (c_{2m}(n))^{\ell-1} \left(K^{(0)}\right)^m \text{ (Gaussian),} \tag{4.24}$$

which should be contrasted with the depth-independent recursion for orthogonal initializations, equation (4.16).

5. Critical initialization and fluctuations for nonlinear orthogonal networks

While deep linear networks provide a nice warmup, they are necessarily limited in that they are only able to compute linear functions of their input. To study more generally useful networks, we need to consider networks with nonlinear activation functions. In this section, we follow the logic of [7] and use the recursion relation for the kernel to determine the orthogonal matrix scaling factors C_W (and bias variance C_b) that ensure critical behavior for orthogonal initializations with various activations. We then show that the normalized single-input 4-point vertex \tilde{V} is depth-independent for linear and tanh activations, but scales linearly with depth for ReLU activations, consistent with the results of [20] for the singular value variance of the input-output Jacobian. We further show that at leading order in $1/n$, finite-width corrections to the kernel with tanh activations do not require further tunings of C_W to maintain criticality, in contrast to the situation for Gaussian initializations.

5.1. Kernel recursion and critical initialization hyperparameters

To derive a recursion relation for the kernel, we compute the 2-point expectation

$$\mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell+1)} z_{i_2; \alpha_2}^{(\ell+1)} \right] = \mathbb{E} \left[\left(b_{i_1}^{(\ell+1)} + \sum_{j_1} W_{i_1 j_1}^{(\ell+1)} \sigma_{j_1; \alpha_1}^{(\ell)} \right) \left(b_{i_2}^{(\ell+1)} + \sum_{j_2} W_{i_2 j_2}^{(\ell+1)} \sigma_{j_2; \alpha_2}^{(\ell)} \right) \right] \quad (5.1)$$

$$= C_b \delta_{i_1 i_2} + \frac{C_W}{n} \delta_{i_1 i_2} \sum_j \mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_2}^{(\ell)} \right], \quad (5.2)$$

where we have used equations (3.2) and (3.3) and the statistical independence of the $(\ell + 1)$ th layer weights from the ℓ th layer activations. Now, we recall that the non-Gaussianity of the orthogonal weight distribution is suppressed by $1/n$, as detailed in section 3. Thus, the preactivation distribution $p_n(\mathbf{z}^{(\ell)} | \mathcal{D})$, which is used to compute the remaining expectation in equation (5.2), only receives $1/n$ corrections compared to initializations with independent Gaussian weights. By the results of [7], $p_n(\mathbf{z}^{(\ell)} | \mathcal{D})$ with Gaussian weights is also Gaussian up to $1/n$ corrections, so to leading order in $1/n$ we can replace the expectation by a Gaussian expectation with respect to the kernel $K_{\alpha_1 \alpha_2}^{(\ell)}$ in layer ℓ . Likewise, evaluating the left-hand side to leading order in $1/n$ using equation (3.7), we obtain the recursion

$$K_{\alpha_1 \alpha_2}^{(\ell+1)} = C_b + C_W \langle \sigma_{\alpha_1} \sigma_{\alpha_2} \rangle_{K^{(\ell)}}, \quad (5.3)$$

where the brackets mean a Gaussian expectation with respect to $K^{(\ell)}$ as defined in equation (3.10). Equation (5.3) is identical to the kernel recursion for Gaussian weights, which means the criticality conditions are also identical. We could have anticipated this result, since the kernel recursion is a statement about the infinite-width behavior, and the suppression of non-Gaussianities by $1/n$ in the orthogonal distribution means that orthogonal and Gaussian initializations are the same in the limit $n \rightarrow \infty$.

As derived in Ch. 5 of [7], the single-input susceptibilities $\chi_{\parallel}(K)$ and $\chi_{\perp}(K)$ control the behavior of the kernel for two nearby inputs. They are defined as

$$\chi_{\parallel}(K) \equiv \frac{C_W}{2K^2} \langle \sigma(z) \sigma(z) (z^2 - K) \rangle_K, \quad (5.4)$$

$$\chi_{\perp}(K) \equiv C_W \langle \sigma'(z) \sigma'(z) \rangle_K. \quad (5.5)$$

The criticality conditions are $\chi_{\parallel}(K) = \chi_{\perp}(K) = 1$, which imposes that the kernel neither grows nor decays exponentially, and thus nearby inputs stay close as they propagate through the network⁹. The value of K which solves these conditions is the fixed point K^* , which implicitly satisfies

$$K^* = C_b + C_W \langle \sigma(z) \sigma(z) \rangle_{K^*} \quad (5.6)$$

for any single input vector \mathbf{x} .

The value of K^* is activation-function dependent and, except for some edge cases, can be divided into two universality classes where the behavior of correlators is the same for all activation functions in the universality class up to order-one factors. The first class contains scale-invariant (piecewise-linear) activation functions such as linear and ReLU, for which

$$(C_b, C_W)^{\text{critical}} = \left(0, \frac{1}{A} \right); \quad K^* = \frac{1}{n_0 A} \|\mathbf{x}\|^2 \quad (\text{scale-invariant}), \quad (5.7)$$

with $A = 1$ for linear and $A = 1/2$ for ReLU. The second class contains smooth nonlinear activation functions with $\sigma(0) = 0$, where

$$(C_b, C_W)^{\text{critical}} = \left(0, \frac{1}{(\sigma'(0))^2} \right); \quad K^* = 0 \quad (5.8)$$

and the kernel decays to zero like $1/\ell$. In particular, $C_W = 1$ is critical for tanh activation functions.

⁹ Banta et al [38] showed that $\chi_{\parallel} = \chi_{\perp} = 1$ is actually a sufficient condition for preventing exponential growth/decay of all higher-point correlators for multiple inputs, in MLP architectures with Gaussian weights.

5.2. 4-point vertex recursion

We now follow the same logic to compute the recursion relation for the 4-point vertex V . Since both universality classes considered above have the biases initialized to zero ($C_b = 0$) at criticality, we will drop the biases for simplicity in the calculations that follow. Consider the expectation of the fourth moment of preactivations:

$$\begin{aligned} \mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell+1)} z_{i_2; \alpha_2}^{(\ell+1)} z_{i_3; \alpha_3}^{(\ell+1)} z_{i_4; \alpha_4}^{(\ell+1)} \right] &= \sum_{j_1, j_2, j_3, j_4} \mathbb{E} [W_{i_1 j_1} W_{i_2 j_2} W_{i_3 j_3} W_{i_4 j_4}] \mathbb{E} \left[\sigma_{j_1; \alpha_1}^{(\ell)} \sigma_{j_2; \alpha_2}^{(\ell)} \sigma_{j_3; \alpha_3}^{(\ell)} \sigma_{j_4; \alpha_4}^{(\ell)} \right] \\ &= C_W^2 \begin{pmatrix} (12)(34)_i \\ (13)(24)_i \\ (14)(23)_i \end{pmatrix}^T \sum_{j_1, j_2, j_3, j_4} \begin{pmatrix} \mathcal{W}[1, 1] & \mathcal{W}[2] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[1, 1] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[2] & \mathcal{W}[1, 1] \end{pmatrix} \begin{pmatrix} (12)(34)_j \\ (13)(24)_j \\ (14)(23)_j \end{pmatrix} \mathbb{E} \left[\sigma_{j_1; \alpha_1}^{(\ell)} \sigma_{j_2; \alpha_2}^{(\ell)} \sigma_{j_3; \alpha_3}^{(\ell)} \sigma_{j_4; \alpha_4}^{(\ell)} \right] \\ &= C_W^2 \begin{pmatrix} (12)(34)_i \\ (13)(24)_i \\ (14)(23)_i \end{pmatrix}^T \begin{pmatrix} \mathcal{W}[1, 1] & \mathcal{W}[2] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[1, 1] & \mathcal{W}[2] \\ \mathcal{W}[2] & \mathcal{W}[2] & \mathcal{W}[1, 1] \end{pmatrix} \begin{pmatrix} \sum_{j, k} \mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_2}^{(\ell)} \sigma_{k; \alpha_3}^{(\ell)} \sigma_{k; \alpha_4}^{(\ell)} \right] \\ \sum_{j, k} \mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_3}^{(\ell)} \sigma_{k; \alpha_2}^{(\ell)} \sigma_{k; \alpha_4}^{(\ell)} \right] \\ \sum_{j, k} \mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_4}^{(\ell)} \sigma_{k; \alpha_2}^{(\ell)} \sigma_{k; \alpha_3}^{(\ell)} \right] \end{pmatrix}, \end{aligned} \tag{5.9}$$

where in the second line we used equation (3.5) for the expectation of four orthogonal weights. Let us now expand this expectation in powers of $1/n$, beginning with an expansion of the Weingarten factors in equation (3.6):

$$\mathcal{W}[2] = \frac{-1}{(n-1)n(n+2)} = -\frac{1}{n^3} + \mathcal{O}(1/n^4); \tag{5.10}$$

$$\mathcal{W}[1, 1] = \frac{n+1}{(n-1)n(n+2)} = \frac{1}{n^2} + \frac{2}{n^4} + \mathcal{O}(1/n^5). \tag{5.11}$$

Since $\mathcal{W}[1, 1] \sim 1/n^2$ but $\mathcal{W}[2] \sim 1/n^3$, the leading terms are on the diagonal, in which case the Weingarten matrix is proportional to the identity matrix and we have

$$\mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell+1)} z_{i_2; \alpha_2}^{(\ell+1)} z_{i_3; \alpha_3}^{(\ell+1)} z_{i_4; \alpha_4}^{(\ell+1)} \right] = \frac{C_W^2}{n^2} \left((12)(34)_i \sum_{j, k} \mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_2}^{(\ell)} \sigma_{k; \alpha_3}^{(\ell)} \sigma_{k; \alpha_4}^{(\ell)} \right] + (2 \text{ sim.}) + \mathcal{O}(1/n) \right) \tag{5.12}$$

where the two remaining terms have the i indices paired with the corresponding α samples indices. At this stage, the remainder of the calculation of the expectation (and the subtraction of the various 2-point correlators to obtain the connected 4-point correlator) is identical to [7] for Gaussian weights.

The sole effect of orthogonal weights, to leading order, is to add an additional $1/n$ term to the connected correlator from $\mathcal{W}[2]$, as the subleading part of $\mathcal{W}[1, 1]$ is higher-order. Since the tensor structure of the correlator is already identical to the 4-point correlator as defined in equation (3.8), we can specialize to the coefficient of $(12)(34)_i$, which is $\frac{1}{n} V_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell+1)}$. Letting ΔV refer to the additional terms in V specific to orthogonal initializations, we have

$$\frac{1}{n} \Delta V_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell+1)} \equiv -\frac{C_W^2}{n^3} \sum_{j, k} \left(\mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_3}^{(\ell)} \sigma_{k; \alpha_2}^{(\ell)} \sigma_{k; \alpha_4}^{(\ell)} \right] + \mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_4}^{(\ell)} \sigma_{k; \alpha_2}^{(\ell)} \sigma_{k; \alpha_3}^{(\ell)} \right] \right). \tag{5.13}$$

To leading order, ΔV is $\mathcal{O}(1)$, so we must take the off-diagonal terms in the sum with $j \neq k$, of which there are $n^2 - n \sim n^2$ terms. To this order we may also replace the expectations with Gaussian expectations [7]:

$$\mathbb{E} \left[\sigma_{j; \alpha_1}^{(\ell)} \sigma_{j; \alpha_3}^{(\ell)} \sigma_{k; \alpha_2}^{(\ell)} \sigma_{k; \alpha_4}^{(\ell)} \right] = \langle \sigma_{\alpha_1} \sigma_{\alpha_3} \rangle_{K^{(\ell)}} \langle \sigma_{\alpha_2} \sigma_{\alpha_4} \rangle_{K^{(\ell)}} + \mathcal{O}(1/n), \tag{5.14}$$

using the bracket notation of equation (3.10) for Gaussian expectations. Each term in the sum is then identical, so evaluating the sums just gives a factor of n^2 , which combines with the $1/n^3$ prefactor in equation (5.13) to give

$$\Delta V_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell+1)} = -C_W^2 \left[\langle \sigma_{\alpha_1} \sigma_{\alpha_3} \rangle_{K^{(\ell)}} \langle \sigma_{\alpha_2} \sigma_{\alpha_4} \rangle_{K^{(\ell)}} + \langle \sigma_{\alpha_1} \sigma_{\alpha_4} \rangle_{K^{(\ell)}} \langle \sigma_{\alpha_2} \sigma_{\alpha_3} \rangle_{K^{(\ell)}} \right] + \mathcal{O}(1/n). \tag{5.15}$$

Combining this contribution from orthogonal weights with the remaining parts of the 4-point vertex arising from the non-Gaussianity of the preactivation distribution [7] gives the full recursion to leading order in $1/n$:

$$\begin{aligned}
V_{(\alpha_1\alpha_2)(\alpha_3\alpha_4)}^{(\ell+1)} &= C_W^2 [\langle 1234 \rangle_{K^{(\ell)}} - \langle 12 \rangle_{K^{(\ell)}} \langle 34 \rangle_{K^{(\ell)}} - \langle 13 \rangle_{K^{(\ell)}} \langle 24 \rangle_{K^{(\ell)}} - \langle 14 \rangle_{K^{(\ell)}} \langle 23 \rangle_{K^{(\ell)}}] \\
&+ \frac{C_W^2}{4} \sum_{\beta_1, \dots, \beta_4 \in \mathcal{D}} V_{(\beta_1\beta_2)(\beta_3\beta_4)}^{(\ell)} \left\langle \sigma_{\alpha_1} \sigma_{\alpha_2} \left(z_{\beta_1} z_{\beta_2} - K_{\beta_1\beta_2}^{(\ell)} \right) \right\rangle_{K^{(\ell)}} \\
&\times \left\langle \sigma_{\alpha_3} \sigma_{\alpha_4} \left(z_{\beta_3} z_{\beta_4} - K_{\beta_3\beta_4}^{(\ell)} \right) \right\rangle_{K^{(\ell)}}
\end{aligned} \tag{5.16}$$

where $V_{(\beta_1\beta_2)(\beta_3\beta_4)}^{(\ell)} \equiv \sum_{\gamma_1, \dots, \gamma_4 \in \mathcal{D}} K_{(\beta_1\beta_2)}^{\gamma_1\beta_1} K_{(\beta_2\beta_3)}^{\gamma_2\beta_2} K_{(\beta_3\beta_4)}^{\gamma_3\beta_3} K_{(\beta_4\beta_1)}^{\gamma_4\beta_4} V_{(\gamma_1\gamma_2)(\gamma_3\gamma_4)}^{(\ell)}$ and we have used the abbreviation $\langle ij \dots k \rangle_{K^{(\ell)}} \equiv \langle \sigma_{\alpha_i} \sigma_{\alpha_j} \dots \sigma_{\alpha_k} \rangle_{K^{(\ell)}}$, with all activations and preactivations on the RHS at layer ℓ . Note that the index structure of the inhomogeneous term in equation (5.16) is exactly that of a Wick contraction, such that for linear activations, the inhomogeneous term vanishes identically as $\langle 1234 \rangle = \langle 12 \rangle \langle 34 \rangle + \langle 13 \rangle \langle 24 \rangle + \langle 14 \rangle \langle 23 \rangle$. This is a special case of the depth-independence of the full preactivation distribution derived in section 4.

5.3. Depth dependence of V for various activations

While equation (5.16) is somewhat opaque for four general inputs, it simplifies greatly in various limits. First, as noted above, for linear activations $\sigma(z) = z$ the term in square brackets has precisely the index structure of a *connected* 4-point correlator of preactivations $z^{(\ell)}$, which vanishes in a Gaussian expectation. Therefore, for linear activations, the recursion for V is homogeneous, $V^{(\ell+1)} \propto C_W^2 V^{(\ell)}$. At criticality with $C_W = 1$, V is thus independent of ℓ and equal to its value at the first layer,

$$\begin{aligned}
V_{(\alpha_1\alpha_2)(\alpha_3\alpha_4)}^{(\ell)} &= V_{(\alpha_1\alpha_2)(\alpha_3\alpha_4)}^{(1)} = -\frac{1}{n^2} [(\mathbf{x}_{\alpha_1} \cdot \mathbf{x}_{\alpha_3})(\mathbf{x}_{\alpha_2} \cdot \mathbf{x}_{\alpha_4}) + (\mathbf{x}_{\alpha_1} \cdot \mathbf{x}_{\alpha_4})(\mathbf{x}_{\alpha_2} \cdot \mathbf{x}_{\alpha_3})] \\
&= -[K_{\alpha_1\alpha_3}^* K_{\alpha_2\alpha_4}^* + K_{\alpha_1\alpha_4}^* K_{\alpha_2\alpha_3}^*] \quad (\text{linear}).
\end{aligned} \tag{5.17}$$

This is simply the leading-order non-Gaussianity from the exact preactivation distribution derived in section 4; since the exact distribution is depth-independent, its connected fourth moment is necessarily depth-independent too. In fact, it is possible to make $V^{(\ell)}$ vanish identically for a finite-width linear network by drawing the first-layer weights i.i.d. from a *Gaussian* distribution with $C_W = 1$, which we will refer to as ‘mixed’ initializations. In that case, the initial condition for the V recursion is $V^{(1)} = 0$ [7], and depth-independence from subsequent orthogonal weight layers ensures $V^{(\ell)} = 0$.

To see the general behavior of the V recursion for nonlinear activation functions, it suffices to consider a single input $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = \alpha$, where we abbreviate $V_{(\alpha\alpha)(\alpha\alpha)} \equiv V$ and $K_{\alpha\alpha} \equiv K$. In that case, evaluating the remaining expectations as in [7], the single-input recursion is

$$V^{(\ell+1)} = C_W^2 [\langle \sigma^4 \rangle_{K^{(\ell)}} - 3\langle \sigma^2 \rangle_{K^{(\ell)}}^2] + \chi_{\parallel}^2 (K^{(\ell)}) V^{(\ell)}, \tag{5.18}$$

which has an additional factor of 3 in the second term in brackets compared to the case for Gaussian weights.

Consistent with previous studies [20, 21], nothing special happens for orthogonal initializations paired with ReLU activations. In that case, at the fixed point $K = K^*$, the Gaussian expectations are $\langle \sigma^2 \rangle = K^*/2$ and $\langle \sigma^4 \rangle = 3(K^*)^2/2$. With the critical values $C_W = 2$ and $\chi_{\parallel} = 1$, the recursion for the normalized 4-point vertex (see equation (3.9)) has the solution

$$\tilde{V}^{(\ell)} = 3\ell - 5 \quad (\text{ReLU}), \tag{5.19}$$

which grows linearly with depth. The slope is 3 compared to 5 for Gaussian weights, which makes the fluctuations overall a bit smaller but does not qualitatively change any behavior during training, as we show in appendix B.

In stark contrast, the orthogonal initializations exhibit similar behavior to linear networks when paired with nonlinear activation functions in the $K^* = 0$ universality class, such as tanh. Because the fixed point of the kernel is at zero, asymptotically all of the preactivations will be small (but not exponentially small, due to the critical tuning), and we may linearize the preactivations around the fixed point to compute the expectations. Using the Taylor expansion $\tanh(z) = z - \frac{z^3}{3} + \mathcal{O}(z^5)$, and working to $\mathcal{O}(K^3)$, we have

$$\langle \sigma^4 \rangle_K \approx \left\langle \left(z - \frac{z^3}{3} \right)^4 \right\rangle_K = 3K^2 - 20K^3 + \mathcal{O}(K^4), \tag{5.20}$$

$$\langle \sigma^2 \rangle_K^2 \approx \left\langle \left(z - \frac{z^3}{3} \right)^2 \right\rangle_K^2 = K^2 - 4K^3 + \mathcal{O}(K^4). \tag{5.21}$$

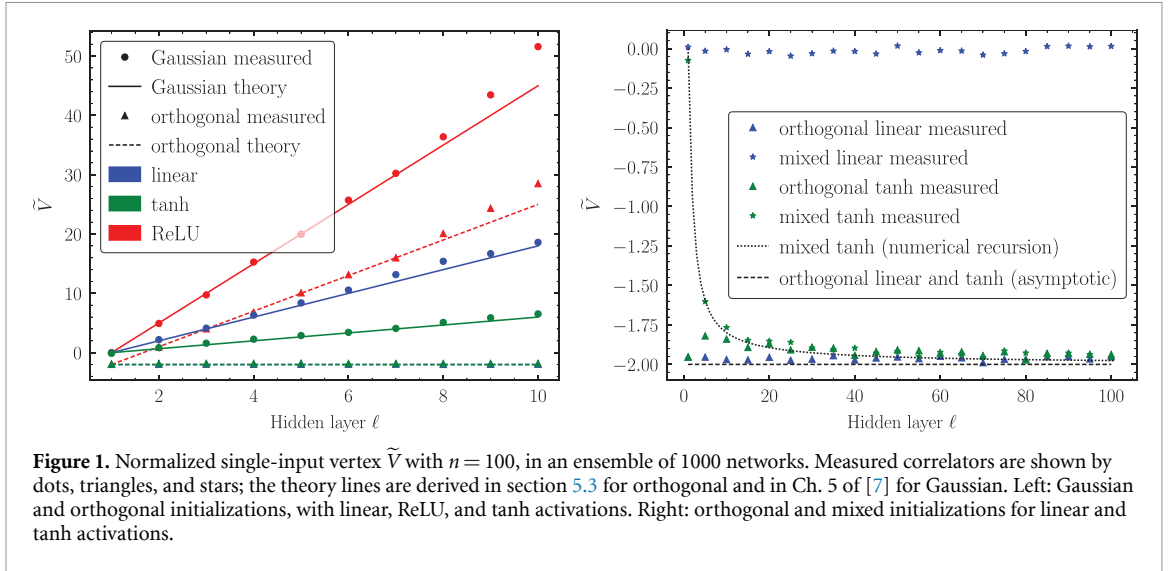


Figure 1. Normalized single-input vertex \tilde{V} with $n = 100$, in an ensemble of 1000 networks. Measured correlators are shown by dots, triangles, and stars; the theory lines are derived in section 5.3 for orthogonal and in Ch. 5 of [7] for Gaussian. Left: Gaussian and orthogonal initializations, with linear, ReLU, and tanh activations. Right: orthogonal and mixed initializations for linear and tanh activations.

Similarly, at criticality with $C_W = 1$, linearizing the susceptibility yields [7]

$$\chi_{\parallel}(K) = 1 - 4K + 17K^2 + \mathcal{O}(K^3). \quad (5.22)$$

Linearizing the kernel around the fixed point, $K^{(\ell)} = K^* + \Delta K^{(\ell)}$ with $K^* = 0$, the V recursion becomes

$$V^{(\ell+1)} = V^{(\ell)} \left[1 - 8\Delta K^{(\ell)} + 50(\Delta K^{(\ell)})^2 \right] - 8(\Delta K^{(\ell)})^3 + \mathcal{O}(\Delta K^{(\ell)})^4. \quad (5.23)$$

Since the kernel recursion is the same for Gaussian and orthogonal initializations, we may use the asymptotic solution for the kernel $\Delta K^{(\ell)} = 1/(2\ell) + \mathcal{O}(1/\ell^2)$ as in Ch. 5.3.3. of [7]. With some benefit of hindsight, we take the ansatz $V^{(\ell)} = c/\ell^2$. The left-hand side of the recursion (5.23) is

$$\frac{c}{(\ell+1)^2} = \frac{c}{\ell^2} - \frac{2c}{\ell^3} + \mathcal{O}(1/\ell^4),$$

while the right-hand side of equation (5.23) reduces to

$$\frac{c}{\ell^2} - \frac{4c+1}{\ell^3} + \mathcal{O}(1/\ell^4).$$

Equating these to leading order in $1/\ell$ fixes $c = -1/2$, which yields a normalized correlator

$$\tilde{V}^{(\ell)} \equiv \frac{V^{(\ell)}}{(\Delta K^{(\ell)})^2} = -2 \quad (\text{tanh}). \quad (5.24)$$

Remarkably, $\tilde{V}^{(\ell)}$ for tanh is depth-independent to leading order in ℓ/n , and takes the same value as for a linear activation function, as can be seen from the single-input limit of equation (5.17).

We have verified equations (5.17), (5.19) and (5.24) with empirical measurements over an ensemble of MLPs. We initialized an ensemble of 1000 rectangular networks of width $n = 100$ and hidden depth $L - 1 = 10$ with linear, ReLU, and tanh activations, and measured the single-input V for both Gaussian and orthogonal initializations. The input was taken to be a random vector of length n with entries drawn uniformly from $[0, 1]$, and we used the same input for each network configuration. The results are shown in figure 1 (left) along with the theoretical results for Gaussian initializations from [7]. We find excellent agreement between measurement and theory for both initializations at small depth, with deviations at large depth arising from next-to-leading-order ℓ^2/n^2 corrections. These corrections appear to be largest for ReLU activations, while tanh activations provide a good fit even out to $\ell = 10$.

Figure 1 (right) extends the architecture to hidden depth $L - 1 = 100$ and verifies that with mixed initializations (Gaussian in the first layer, orthogonal in subsequent layers) and linear activations, V can be made to vanish identically (up to order $1/n$ corrections) for all depths, as anticipated below equation (5.17). Intriguingly, we see that with tanh activations, the mixed initializations do result in a vanishing $V^{(1)}$ (as for the linear network), but in deeper layers \tilde{V} asymptotes to -2 . This validates our linearized analysis since it

verifies that $\tilde{V} = -2$ is indeed a fixed point of the V recursion, independent of initial conditions. We further confirm this by numerically solving the single-input recursion (5.18), which shows a good match to the measured values. Even for purely orthogonal initializations, V has a slow decay to the fixed point at -2 because of higher-order $1/\ell^2$ corrections to the asymptotic analysis.

5.4. next-to-leading-order (NLO) metric

To complete our analytic derivations, we also compute the NLO metric $G^{\{1\}}$, namely the $1/n$ piece of equation (3.7):

$$\mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell)} z_{i_2; \alpha_2}^{(\ell)} \right] = \delta_{i_1 i_2} \left[K_{\alpha_1 \alpha_2}^{(\ell)} + \frac{1}{n} G^{\{1\}(\ell)} + \dots \right]. \quad (5.25)$$

Because this correlator only involves a 2-point expectation of preactivations, the recursion will be identical to that derived in [7] for Gaussian weights, with all non-Gaussianities encoded in V . The recursion for the single-input NLO metric and tanh activations, expanded in powers of $1/\ell$, is (see Ch. 5.4.2 of [7])

$$G^{\{1\}(\ell+1)} = G^{\{1\}(\ell)} \left[1 - 4\Delta K^{(\ell)} + \dots \right] + V^{(\ell)} [-2 + \dots]. \quad (5.26)$$

Similar to the analysis of V , we make a power law ansatz for the depth dependence of $G^{\{1\}(\ell)}$. For the case of Gaussian weights, $G^{\{1\}(\ell)}$ turns out to be constant in ℓ to leading order, such that $G^{\{1\}(\ell)}/K^{(\ell)} \sim \ell/n$. However, for orthogonal weights with $V^{(\ell)} = -1/(2\ell^2)$, matching powers of $1/\ell$ yields $G^{\{1\}(\ell)} = 1/\ell$ to leading order, so that $G^{\{1\}(\ell)}$ is a true perturbative correction to the metric that does *not* grow with depth:

$$K^{(\ell)} + \frac{1}{n} G^{\{1\}(\ell)} = \frac{1}{2\ell} \left(1 + \frac{2}{n} + \dots \right). \quad (5.27)$$

Equivalently, we can consider the ‘renormalized’ perspective of [7], where NLO metric corrections imply finite-width corrections to the initialization hyperparameters:

$$G^{\{1\}(\ell+1)} = c_b^{\{1\}(\ell)} + c_W^{\{1\}(\ell)} \left[K^{(\ell)} - 2 \left(K^{(\ell)} \right)^2 + \dots \right] + G^{\{1\}(\ell)} \left[1 - 4K^{(\ell)} + \dots \right] - 2V^{(\ell)} + \dots \quad (5.28)$$

where $C_b^{(\ell)} = c_b^{\{0\}(\ell)} + \frac{1}{n} c_b^{\{1\}(\ell)} + \dots$ and $C_W^{(\ell)} = c_W^{\{0\}(\ell)} + \frac{1}{n} c_W^{\{1\}(\ell)} + \dots$. In the case of Gaussian weights, maintaining depth-independence of the two-point function at finite width requires fine-tuning the weight initialization to $\mathcal{O}(1/n)$ at all layers, $c_W^{\{1\}(\ell)} = \frac{2}{3}$. However, for orthogonal weights, after inserting the series expansions into equation (5.28), matching powers of ℓ to $\mathcal{O}(1/\ell^2)$ requires that $c_b^{\{1\}(\ell)} = c_W^{\{1\}(\ell)} = 0$. In other words, the infinite-width hyperparameters $C_W = 1, C_b = 0$ still maintain criticality at leading order in $1/n$, independent of the depth of the network. This shows that the two perspectives are the same for orthogonal weights, where depth-independence of the initialization hyperparameters is equivalent to the depth-independence of the NLO metric correction from equation (5.27).

5.5. Relation to previous work

The result of equation (5.24) offers an alternative perspective on the meaning of the ‘linear regime’ of critically-initialized networks. Early work in [19] showed that deep orthogonal networks have depth-independent training times. Pennington and Schoenholz [20] studied the singular value spectrum of the input–output Jacobian, and found that by tuning to criticality, as $K^* \rightarrow 0$ the fourth moment of the singular value spectrum becomes depth-independent and the maximum singular value approaches its value for a linear orthogonal network. The interpretation was that since the preactivations were so small, most of the neurons in the network were operating as if the activation function were linear; this is precisely the sense in which equation (5.24) was derived by an asymptotic analysis which linearized the activation function about zero. Our results support this interpretation because the single-input \tilde{V} takes the same value for critical tanh networks as for critical linear networks. Indeed, as noted in [7], the variance of the input–output Jacobian is controlled by a 2-input component of the full multi-input $V_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell)}$. In appendix A.1, we show measurements of $\tilde{V}_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell)}$ for various choices of sample indices, all of which exhibit the same depth-independence. These results strongly suggest that the *entire* normalized 4-input tensor $\tilde{V}_{(\alpha_1 \alpha_2)(\alpha_3 \alpha_4)}^{(\ell)}$ is depth-independent. We plan to study this further analytically in future work.

Finally, we note that $V < 0$ is somewhat unusual in light of the discussion of optimal depth-to-width ratios from tripartite mutual information considerations in appendix A of [7]. In Gaussian networks, $V > 0$ at all layers, which corresponds to a redundant storing of mutual information among neurons in the final

layer. This leads to a competition with the growth of mutual information with depth at large width, resulting in a theoretically-optimal aspect ratio. By contrast, with $V < 0$ in orthogonal tanh networks, mutual information is stored synergistically, and maximum-entropy arguments do not lead to an optimal aspect ratio. As we will see in the measurements and numerical experiments which follow, this suggests that the aspect ratio L/n does not determine the behavior of orthogonal networks, but rather that the behavior depends only weakly on L and is largely independent of n . We leave to future work the theoretical implications of this result for feature learning.

6. Measurement of NTK statistics

While the results of section 5 provide some evidence that orthogonal weights reduce network fluctuations, they only concern the statistics of preactivations at initialization and cannot yet be directly linked to training and/or generalization. To understand the behavior of the model during training, we will study the complete set of leading-order higher-point NTK correlators which are only non-vanishing at finite width and, according to the arguments of [7], govern the evolution of the NTK during training and thus the phenomenon of feature learning at finite width. Section 6.1 reviews the relevant results from [7] in order to keep this paper self-contained, while section 6.2 reports the results of our measurements of the relevant correlators.

6.1. NTK and learning rate definitions

The NTK at layer ℓ is defined (in the notation of [7]) as

$$\widehat{H}_{i_1 i_2; \alpha_1 \alpha_2}^{(\ell)} \equiv \sum_{\mu, \nu} \lambda_{\mu\nu} \frac{dz_{i_1; \alpha_1}^{(\ell)}}{d\theta_\mu} \frac{dz_{i_2; \alpha_2}^{(\ell)}}{d\theta_\nu}, \quad (6.1)$$

where μ indexes the complete set of model parameters θ_μ (i.e. weights and biases) and $\lambda_{\mu\nu}$ is the learning rate tensor. The hat on \widehat{H} emphasizes that it is a stochastic variable at finite width. Under full batch GD, the parameters are updated according to

$$\theta_\mu \leftarrow \theta_\mu - \eta \sum_\nu \lambda_{\mu\nu} \frac{d\mathcal{L}_A}{d\theta_\nu}, \quad (6.2)$$

where η is the global learning rate, and \mathcal{L}_A is the loss function evaluated on the training set $\mathcal{A} \subset \mathcal{D}$. We will adopt the prescription of [7] and use the learning rate tensor to scale the weight and bias learning rates differently:

$$\lambda_{b_{i_1}^{(\ell)} b_{i_2}^{(\ell)}} = \delta_{i_1 i_2} \lambda_b^{(\ell)}, \quad \lambda_{W_{i_1 j_1}^{(\ell)} W_{i_2 j_2}^{(\ell)}} = \delta_{i_1 i_2} \delta_{j_1 j_2} \frac{\lambda_W^{(\ell)}}{n}, \quad (6.3)$$

where we have also allowed for a layer dependence to the learning rates. With the learning rate tensor (6.3), the NTK obeys a forward iteration equation

$$\begin{aligned} \widehat{H}_{i_1 i_2; \alpha_1 \alpha_2}^{(\ell+1)} &= \delta_{i_1 i_2} \left[\lambda_b^{(\ell+1)} + \lambda_W^{(\ell+1)} \left(\frac{1}{n} \sum_j \sigma_{j_1; \alpha_1}^{(\ell)} \sigma_{j_2; \alpha_2}^{(\ell)} \right) \right] \\ &+ \sum_{j_1, j_2} W_{i_1 j_1}^{(\ell+1)} W_{i_2 j_2}^{(\ell+1)} \sigma_{j_1; \alpha_1}^{(\ell)} \sigma_{j_2; \alpha_2}^{(\ell)} \widehat{H}_{j_1 j_2; \alpha_1 \alpha_2}^{(\ell)}, \end{aligned} \quad (6.4)$$

which can be derived by substituting the definition of \widehat{H} into the MLP forward equation (3.1) and evaluating the derivatives using the chain rule.

As is well-known in the literature [16–18], at infinite width the NTK is frozen during training and becomes deterministically equal to its mean,

$$\mathbb{E} \left[\widehat{H}_{i_1 i_2; \alpha_1 \alpha_2}^{(\ell)} \right] \equiv \delta_{i_1 i_2} \Theta_{\alpha_1 \alpha_2}^{(\ell)} + \mathcal{O}(1/n), \quad (6.5)$$

which depends on the initialization hyperparameters and learning rate tensor. As with the kernel, the frozen NTK takes the same value with Gaussian and orthogonal weights [24]. To see the distinction between the two initializations, we will be interested in the fluctuations

$$\widehat{\Delta H}_{i_1 i_2; \alpha_1 \alpha_2}^{(\ell)} \equiv \widehat{H}_{i_1 i_2; \alpha_1 \alpha_2}^{(\ell)} - \delta_{i_1 i_2} \Theta_{\alpha_1 \alpha_2}^{(\ell)}, \quad (6.6)$$

which appear in two correlators to leading order in $1/n^{10}$:

$$\mathbb{E} \left[\widehat{\Delta H}_{i_1 i_2; \alpha_1 \alpha_2}^{(\ell)} \widehat{\Delta H}_{i_3 i_4; \alpha_3 \alpha_4}^{(\ell)} \right] \equiv \frac{1}{n} \left[(12) (34)_i A_{\alpha_1 \alpha_2 \alpha_3 \alpha_4}^{(\ell)} + (13) (24)_i B_{\alpha_1 \alpha_3 \alpha_2 \alpha_4}^{(\ell)} + (14) (23)_i B_{\alpha_1 \alpha_4 \alpha_2 \alpha_3}^{(\ell)} \right], \quad (6.7)$$

$$\mathbb{E} \left[z_{i_1; \alpha_1}^{(\ell)} z_{i_2; \alpha_2}^{(\ell)} \widehat{\Delta H}_{i_3 i_4; \alpha_3 \alpha_4}^{(\ell)} \right] \equiv \frac{1}{n} \left[(12) (34)_i D_{\alpha_1 \alpha_2 \alpha_3 \alpha_4}^{(\ell)} + (13) (24)_i F_{\alpha_1 \alpha_3 \alpha_2 \alpha_4}^{(\ell)} + (14) (23)_i F_{\alpha_1 \alpha_4 \alpha_2 \alpha_3}^{(\ell)} \right], \quad (6.8)$$

with $(ab)(cd)_i \equiv \delta_{ia} \delta_{ib} \delta_{ic} \delta_{id}$ as in section 3. Recursion relations for A , B , D , and F are derived in [7] for Gaussian weights; we leave a parallel derivation for orthogonal weights to future work, and focus on empirical measurements of these correlators as we describe below. We may apply similar dimensional analysis reasoning as in section 3 to define \widehat{H} as a second dimensionful quantity, and consider the dimensionless normalized correlators

$$\widetilde{A}^{(\ell)} \equiv \frac{A^{(\ell)}}{(\Theta^{(\ell)})^2}, \quad \widetilde{B}^{(\ell)} \equiv \frac{B^{(\ell)}}{(\Theta^{(\ell)})^2}; \quad (6.9)$$

$$\widetilde{D}^{(\ell)} \equiv \frac{D^{(\ell)}}{K^{(\ell)} \Theta^{(\ell)}}, \quad \widetilde{F}^{(\ell)} \equiv \frac{F^{(\ell)}}{K^{(\ell)} \Theta^{(\ell)}}, \quad (6.10)$$

which all have equal powers of z and/or \widehat{H} in the numerator and denominator¹¹.

The training dynamics of MLP architectures are governed by additional correlators at leading order in $1/n$ that involve the descendants of the NTK obtained by expanding a GD update to higher order in the learning rate η . The dNTK¹²

$$\widehat{dH}_{i_0 i_1 i_2; \alpha_0 \alpha_1 \alpha_2}^{(\ell)} = \sum_{\ell_1, \ell_2=1}^{\ell} \sum_{\mu_1, \nu_1, \mu_2, \nu_2} \lambda_{\mu_1 \nu_1}^{(\ell_1)} \lambda_{\mu_2 \nu_2}^{(\ell_2)} \frac{d^2 z_{i_0; \alpha_0}^{(\ell)}}{d\theta_{\mu_1}^{(\ell_1)} d\theta_{\mu_2}^{(\ell_2)}} \frac{dz_{i_1; \alpha_1}^{(\ell)}}{d\theta_{\nu_1}^{(\ell_1)}} \frac{dz_{i_2; \alpha_2}^{(\ell)}}{d\theta_{\nu_2}^{(\ell_2)}} \quad (6.11)$$

acts as a metafeature, updating the NTK at order η and the preactivations at order η^2 . Likewise, there are two ddNTK meta-metafeatures, $\widehat{dd}_I H$ and $\widehat{dd}_{II} H$ —each carrying four sample indices and involving six derivatives of the preactivations¹³—which update the dNTK at order η , the NTK at order η^2 , and the preactivations at order η^3 . The hierarchy terminates there, such that the ddNTKs are frozen at initialization and do not update to leading order in $1/n$. The perturbative expansion in the learning rate breaks down for activation functions with discontinuous derivatives such as ReLU, so we will focus on tanh activations in what follows, which are (fortunately) the activations which give the most interesting phenomenology for orthogonal initializations.

For our purposes, we are interested in the correlators at leading order in $1/n$,

$$\mathbb{E} \left[\widehat{dH}_{i_0 i_1 i_2; \alpha_0 \alpha_1 \alpha_2}^{(\ell)} z_{i_3; \delta_3}^{(\ell)} \right] \equiv \frac{1}{n} \left[(03) (12)_i P_{\alpha_0 \alpha_1 \alpha_2 \alpha_3}^{(\ell)} + (01) (23)_i Q_{\alpha_0 \alpha_1 \alpha_2 \alpha_3}^{(\ell)} + (02) (13)_i Q_{\alpha_0 \alpha_2 \alpha_1 \alpha_3}^{(\ell)} \right], \quad (6.12)$$

$$\mathbb{E} \left[\widehat{dd}_I H_{i_1 i_2 i_3 i_4; \alpha_1 \alpha_2 \alpha_3 \alpha_4}^{(\ell)} \right] \equiv \frac{1}{n} \left[(12) (34)_i R_{\alpha_1 \alpha_2 \alpha_3 \alpha_4}^{(\ell)} + (13) (24)_i R_{\alpha_1 \alpha_3 \alpha_2 \alpha_4}^{(\ell)} + (14) (23)_i R_{\alpha_1 \alpha_4 \alpha_2 \alpha_3}^{(\ell)} \right], \quad (6.13)$$

$$\mathbb{E} \left[\widehat{dd}_{II} H_{i_1 i_2 i_3 i_4; \alpha_1 \alpha_2 \alpha_3 \alpha_4}^{(\ell)} \right] \equiv \frac{1}{n} \left[(12) (34)_i S_{\alpha_1 \alpha_2 \alpha_3 \alpha_4}^{(\ell)} + (13) (24)_i T_{\alpha_1 \alpha_3 \alpha_2 \alpha_4}^{(\ell)} + (14) (23)_i U_{\alpha_1 \alpha_4 \alpha_2 \alpha_3}^{(\ell)} \right]. \quad (6.14)$$

Notice that the index labeling for \widehat{dH} is slightly different because the first index is distinguished, such that the tensor decomposition does not exactly follow that of a Wick contraction. Equations (6.7)–(6.8) and (6.12)–(6.14) (along with equation (3.8)) comprise the full set of MLP correlators which appear at $\mathcal{O}(1/n)$, all of which carry four neural indices and four sample indices. The dimensionless normalized correlators are

$$\widetilde{P}^{(\ell)} \equiv \frac{P^{(\ell)}}{(\Theta^{(\ell)})^2}, \quad \widetilde{Q}^{(\ell)} \equiv \frac{Q^{(\ell)}}{(\Theta^{(\ell)})^2}; \quad (6.15)$$

¹⁰ While the tensor decomposition of these correlators was derived in [7] for the case of Gaussian weights, we have verified numerically that the same structure holds for orthogonal weights, analogous to the case of V in section 5.

¹¹ Note that to leading order in $1/n$, these normalized correlators take the same value whether K or G is used in the denominator, and likewise Θ or \widehat{H} .

¹² Since the dNTK mixes learning rates at different layers, and we are allowing the weight and bias learning rates to change with layer, we have written $\lambda_{\mu_1 \nu_1}^{(\ell_1)}$ to emphasize that the parameter indices μ_1 and ν_1 belong to layer ℓ_1 .

¹³ Schematically, $\widehat{dd}_I H \sim \lambda^3 \frac{d^3 z}{d\theta^3} \left(\frac{dz}{d\theta} \right)^3$ and $\widehat{dd}_{II} H \sim \lambda^3 \left(\frac{d^2 z}{d\theta^2} \right)^2 \left(\frac{dz}{d\theta} \right)^2$; the full expressions may be found in [7].

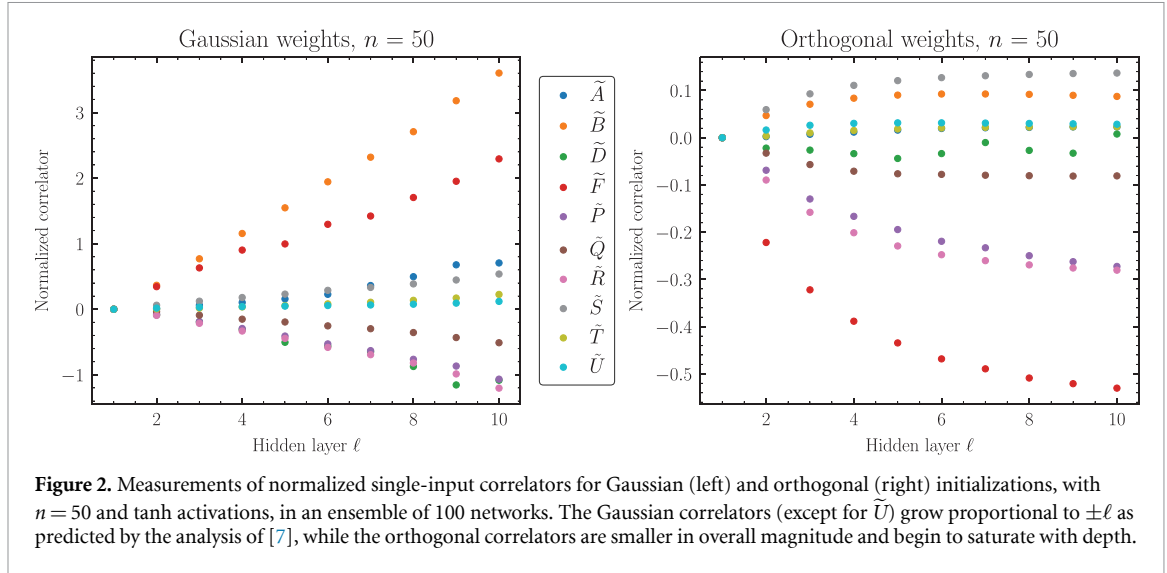


Figure 2. Measurements of normalized single-input correlators for Gaussian (left) and orthogonal (right) initializations, with $n = 50$ and tanh activations, in an ensemble of 100 networks. The Gaussian correlators (except for \tilde{U}) grow proportional to $\pm\ell$ as predicted by the analysis of [7], while the orthogonal correlators are smaller in overall magnitude and begin to saturate with depth.

$$\tilde{R}^{(\ell)} \equiv \frac{R^{(\ell)}K^{(\ell)}}{(\Theta^{(\ell)})^3}, \quad \tilde{S}^{(\ell)} \equiv \frac{S^{(\ell)}K^{(\ell)}}{(\Theta^{(\ell)})^3}, \quad \tilde{T}^{(\ell)} \equiv \frac{T^{(\ell)}K^{(\ell)}}{(\Theta^{(\ell)})^3}, \quad \tilde{U}^{(\ell)} \equiv \frac{U^{(\ell)}K^{(\ell)}}{(\Theta^{(\ell)})^3}. \quad (6.16)$$

For Gaussian weights, all dimensionless correlators grow linearly with ℓ for tanh activations, except for \tilde{U} which decays as $1/\ell$.

6.2. Measurements of NTK, dNTK, and ddNTK statistics at initialization

Analogous to our measurements of V in section 5, we initialized an ensemble of 100 rectangular networks of width $n = 50$ and hidden layers $L - 1 = 10$ with tanh activations, and measured all single-input NTK correlators from section 6.1 for both Gaussian and orthogonal initializations. The input was taken to be a random vector of length n with entries drawn uniformly from $[0, 1]$. The weight and bias learning rates which enter the definitions of \widehat{H} , \widehat{dH} , and $\widehat{dd}_{I,II}H$ were taken to be

$$\lambda_b^{(\ell)} = \frac{1}{\ell}, \quad \lambda_W^{(\ell)} = 1, \quad (6.17)$$

where the scaling of the bias learning rate with depth follows the ‘learning rate equivalence principle’ prescription of [7] to ensure roughly equal contributions to the NTK at each layer for tanh activations¹⁴. The results for the dimensionless correlators are shown in figure 2, using the same random input for both sets of initializations¹⁵. It is clear that for the Gaussian initialization (left), all correlators except \tilde{U} are growing approximately linearly with depth, while for orthogonal initialization (right), the growth is much slower and appears to saturate at finite depth.

Due to memory limitations, it is difficult to measure the dNTK and ddNTK correlators at very large width and depth. However, to investigate the behavior of the orthogonal correlators further, we reduced the width to 20 and increased the number of hidden layers to $L - 1 = 30$. In this regime, with $L/n \gtrsim 1$, the effective theory of the Gaussian-initialized network begins to break down and we expect to see the largest deviations between Gaussian and orthogonal initializations. The results are shown in figure 3. For Gaussian initializations, the correlators become noisy and begin to grow exponentially for $\ell/n \gtrsim 1$ (including \tilde{U}), consistent with the effective theory arguments of [7]. However, for orthogonal weights, the correlators saturate at a depth $\ell \sim 20$, and take approximately the same asymptotic values as they do for $n = 50$ in figure 2. Intriguingly, the A and D correlators, whose recursions for Gaussian initializations involve the 4-point vertex V [7], seem to be consistently smaller than B and F at all depths for orthogonal initializations, which is not the case for Gaussian initializations. While the recursions may well change for orthogonal weights (a derivation which we leave to future work), this is at least suggestive of the fact that the dominant

¹⁴ We demonstrate in appendix B that these prescriptions, designed to ensure that each layer yields an order-one contribution to the NTK, give essentially identical generalization at small L/n compared to ordinary (non-tensorial) GD, but seem to lead to worse generalization for deep orthogonal networks. We speculate on the reasons for this behavior in appendix B.

¹⁵ The code to generate these plots may be found at <https://github.com/ykahn/OrthogonalInits>. Note that the RAM requirements for measuring ddH in particular are quite substantial due to the proliferation of indices; for the chosen architecture hyperparameters and an ensemble of size 100, about 7 GB of RAM are required.

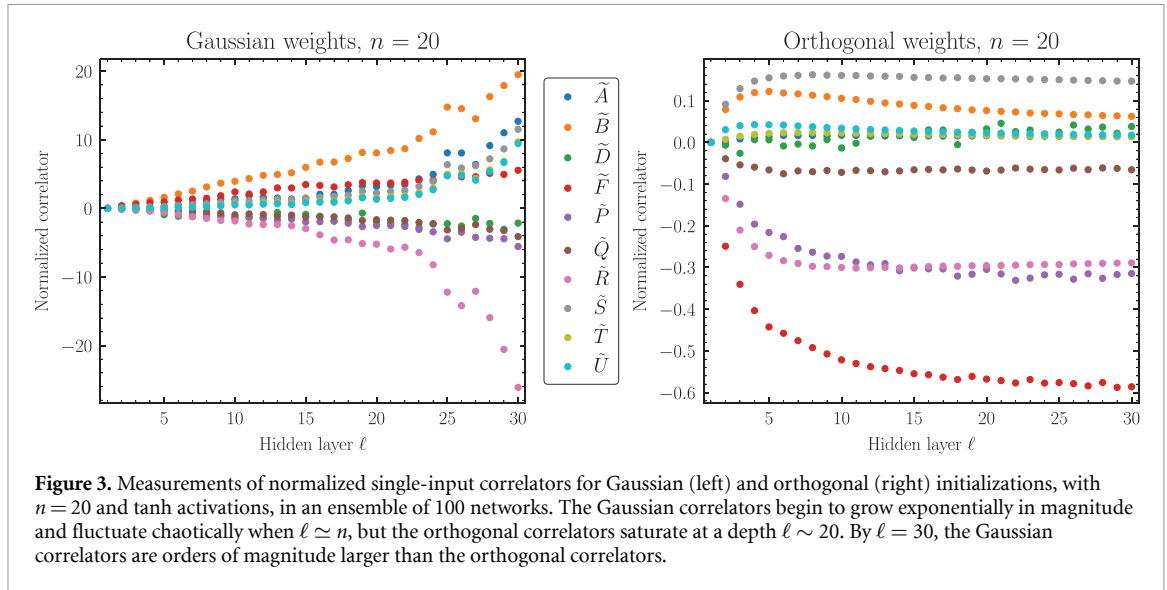


Figure 3. Measurements of normalized single-input correlators for Gaussian (left) and orthogonal (right) initializations, with $n = 20$ and tanh activations, in an ensemble of 100 networks. The Gaussian correlators begin to grow exponentially in magnitude and fluctuate chaotically when $\ell \simeq n$, but the orthogonal correlators saturate at a depth $\ell \sim 20$. By $\ell = 30$, the Gaussian correlators are orders of magnitude larger than the orthogonal correlators.

correlators that drive feature learning (i.e. the evolution of the NTK and the dNTK) are *not* the ones associated with noisy ensemble fluctuations.

All of the above observations suggest that, unlike MLPs with Gaussian weights, orthogonal networks do *not* have an L/n cutoff for the perturbative behavior of correlators at initializations, but rather are only weakly sensitive to width and are roughly depth-independent for $L \gtrsim 20$ regardless of width. Empirically, the behavior of the normalized orthogonal correlators seems characteristic of a logarithmic *decay* $a_0 + a_1 \log \ell / \ell$ (to be contrasted with the linear *growth* of the Gaussian correlators), where a full calculation of the orthogonal initialization recursions would permit us to predict the constants a_0 and a_1 as [7] does for Gaussian initializations.

Informed by the measurements of the correlators at initialization, we conjecture that the smaller values of the ‘noise’ correlators A and D , whose recursions involve the 4-point vertex V , will lead to faster training for orthogonal initializations compared to Gaussian initializations for networks of the same depth and width. Furthermore, we can predict that the performance of an orthogonal network at fixed width should be roughly independent of L for $L \gtrsim 20$, and generalization at the end of training should be superior to Gaussian networks in the regime $L/n \sim 1$, where Gaussian networks suffer from large fluctuations. We now proceed to test these conjectures experimentally.

7. Training and generalization experiments

The results of section 6 are suggestive but not persuasive, because (among other reasons) they only refer to the ensemble statistics at initialization. In this section, using the toy examples of the MNIST and CIFAR-10 classification problems, we will provide more evidence for the superior training behavior of orthogonal initializations conjectured above by comparing the final test loss on trained networks from both initializations. These experiments are very similar to those performed in [20] and [24], but we will make some slightly different choices in order to stay as close as possible to the theoretical framework of this paper.

7.1. Experimental setup

The setup of our experiments was as follows. We initialized MLP networks of hidden layer width $n = 30$ or $n = 100$, tanh activations on each hidden layer, and the number of hidden layers ranging from $L - 1 = 1$ to $L - 1 = 50$. The initial biases were set to zero and the initial weights were sampled from either Gaussian or orthogonal distributions with critical initialization $C_W = 1$. Note that an orthogonal distribution is not well-defined for the first layer because its weight matrix is not a square matrix (the input dimension is 784 for MNIST and 3072 for CIFAR-10), so instead, the first layer weights are sampled from the QR decomposition of a rectangular matrix with random Gaussian entries, as implemented by the default Keras [40] orthogonal

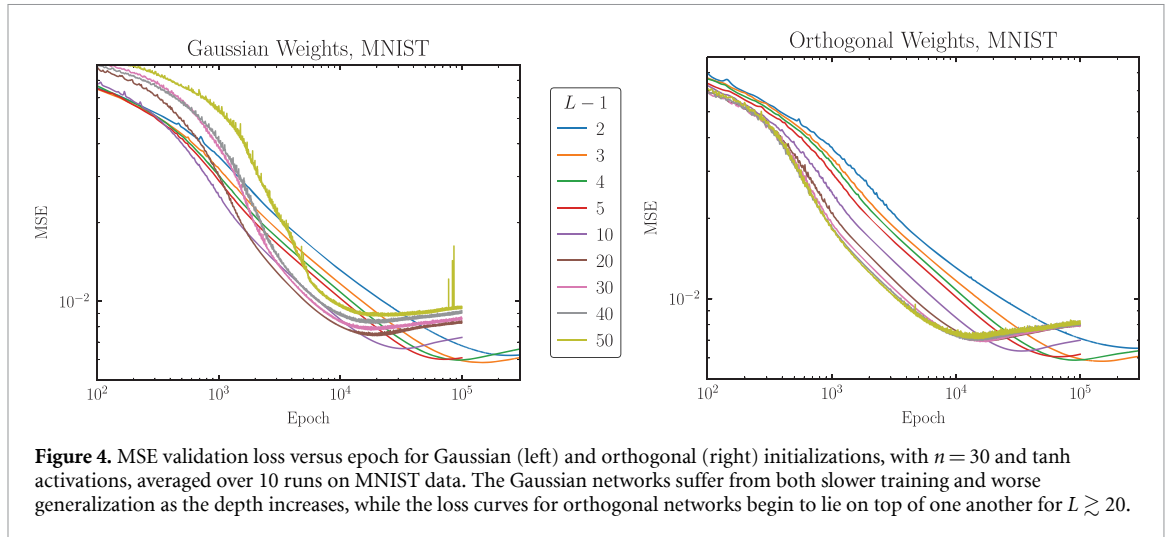


Figure 4. MSE validation loss versus epoch for Gaussian (left) and orthogonal (right) initializations, with $n = 30$ and tanh activations, averaged over 10 runs on MNIST data. The Gaussian networks suffer from both slower training and worse generalization as the depth increases, while the loss curves for orthogonal networks begin to lie on top of one another for $L \gtrsim 20$.

initializer¹⁶. When the weight matrix is a square matrix, this sampling algorithm results in a Haar-distributed orthogonal matrix as desired [41]. However, as discussed in section 5.3 for the case of the 4-point vertex V , and in appendix A.2 for the NTK correlators A , B , D , and F , this choice simply affects the initial conditions of the various correlator recursions and not the large-depth asymptotics. Consequently, we expect the analysis of the previous section (in particular the depth independence for $L \gtrsim 20$) to apply. For the MNIST dataset, where most pixels are zero in the raw dataset, we performed a standard rescaling so that the entire dataset has zero mean and unit variance. The loss function was taken to be mean squared error (MSE) loss,

$L = \|\mathbf{z}_\alpha^{(L)} - \mathbf{y}_\alpha\|^2$ where \mathbf{y}_α is the 10-dimensional one-hot label encoding the class of the training point \mathbf{x}_α ¹⁷.

For MNIST, we trained all networks with full-batch tensorial GD (6.2) with learning rate tensor given by equation (6.3), weight and bias learning rate rescalings given in equation (6.17), and a global learning rate of $\eta = 10^{18}$. We impose implicit regularization via early stopping, using a validation set of size 10^4 , and take our final trained network to be the one corresponding to the best validation loss. We verified that for architectures of hidden depth greater than or equal to 5, 10^5 epochs is sufficient to identify the onset of overfitting, with more training required to identify overfitting for very shallow networks: for hidden depths of 1, 2, 3, and 4, we extended training to 10^6 , 5×10^5 , 3×10^5 , and 3×10^5 epochs, respectively. We did not attempt to optimize the global learning rate for the best overall performance, but rather chose η large enough for training to not exceed 10^5 epochs for deep networks, but small enough to prevent divergences in the training loss from large step sizes. We have checked that varying η gives qualitatively similar results for all experiments in this section. Finally, to compare the performance of trained networks across different architectures, we averaged the test loss over 10 instantiations of each network at the epoch determined by early stopping for each instantiation. For CIFAR-10, we used a validation set of size 7×10^3 partitioned from the training set. We used the same learning rate $\eta = 10$, but found that overfitting began later, so we trained for 3×10^5 epochs for hidden depths greater than or equal to 5 and 6×10^6 epochs for smaller depths. Due to the slower training times for CIFAR-10 and the compute limitations for this study, we only used 5 CIFAR-10 training runs for each architecture.

7.2. Results for training dynamics

In figures 4 and 5 we show the validation loss as a function of training epoch for MNIST (averaged over 10 runs) and CIFAR-10 (averaged over 5 runs), respectively. For orthogonal initializations, the loss curves for MNIST (and to a lesser extent for CIFAR) appear to converge for depths of 20 or greater, as predicted based on the behavior of the NTK correlators as a function of depth in section 6.2. By contrast, for Gaussian initializations with both datasets, deeper networks appear to train more slowly and achieve a worse

¹⁶ With an eye toward practical implications, we note that the computational complexity of QR decomposition is $\mathcal{O}(n^3)$, while filling a matrix with independent Gaussian-distributed entries is $\mathcal{O}(n^2)$, so there is a mild compute penalty for orthogonal initializations. Furthermore, as implemented in `Keras`, it is not clear if orthogonal initializations can be parallelized across GPUs for extremely large networks, though in this limit we expect Gaussian and orthogonal networks to perform very similarly.

¹⁷ The use of MSE loss is somewhat unconventional in the classification context, but corresponds directly to the analysis of [7] and has been shown to have essentially equivalent or even superior performance to cross-entropy loss across a wide variety of tasks [42]; see appendix B for results using cross-entropy loss.

¹⁸ With the conventions of equation (6.3), this is not a parametrically large learning rate, since the weight learning rate is effectively $10/n$ which is less than 1 for both widths we consider.

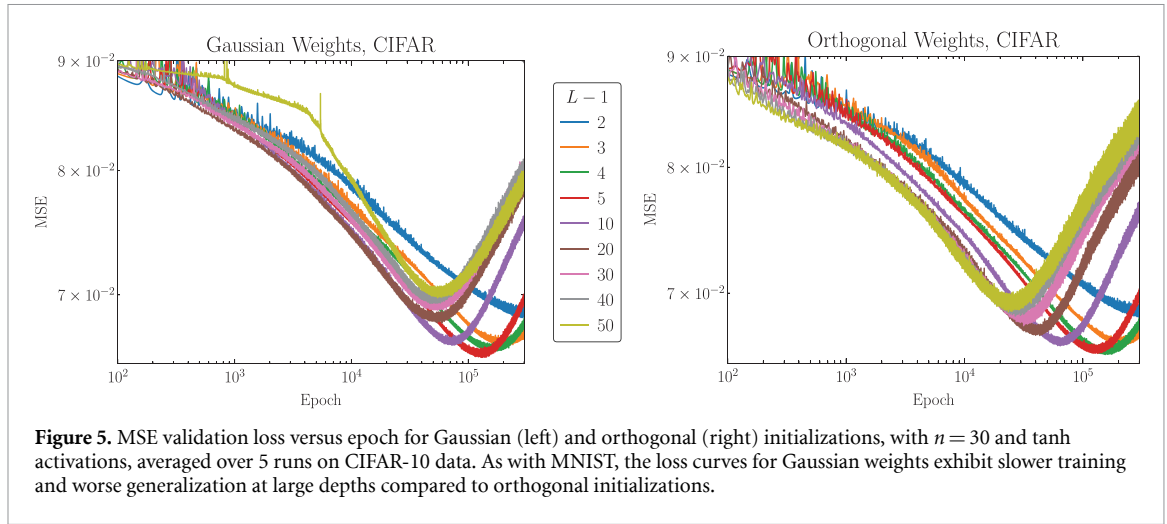


Figure 5. MSE validation loss versus epoch for Gaussian (left) and orthogonal (right) initializations, with $n = 30$ and tanh activations, averaged over 5 runs on CIFAR-10 data. As with MNIST, the loss curves for Gaussian weights exhibit slower training and worse generalization at large depths compared to orthogonal initializations.

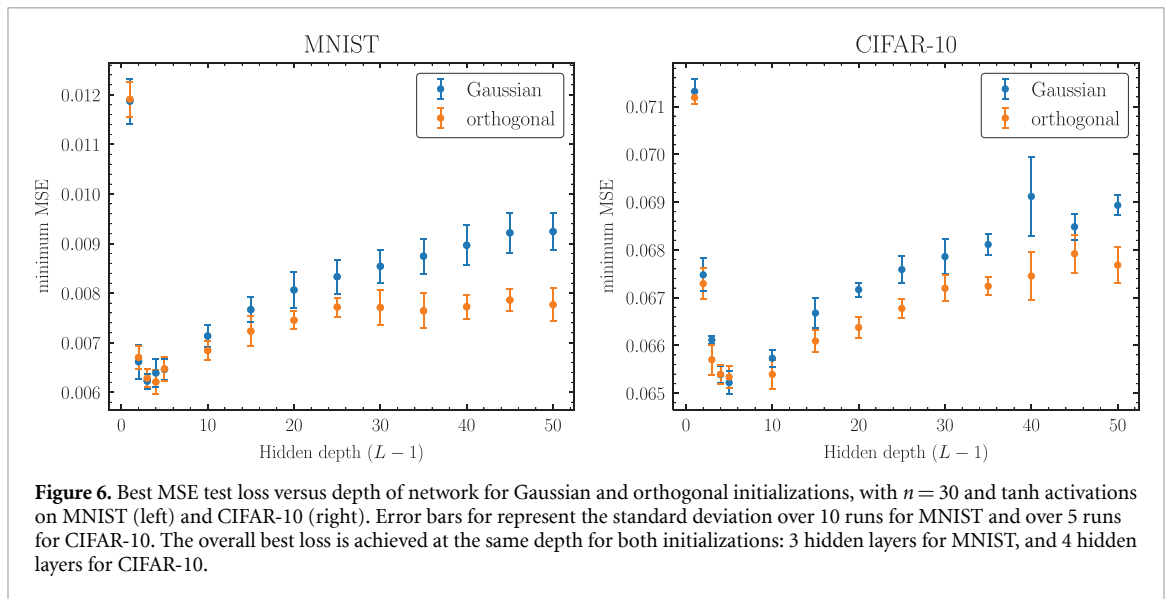


Figure 6. Best MSE test loss versus depth of network for Gaussian and orthogonal initializations, with $n = 30$ and tanh activations on MNIST (left) and CIFAR-10 (right). Error bars for represent the standard deviation over 10 runs for MNIST and over 5 runs for CIFAR-10. The overall best loss is achieved at the same depth for both initializations: 3 hidden layers for MNIST, and 4 hidden layers for CIFAR-10.

validation loss. Note that the non-monotonic oscillations in the loss, which are especially apparent in the CIFAR data, are characteristic of the ‘edge of stability’ GD dynamics first identified in [43], and are a generic feature of full-batch GD with fixed learning rate. In general, for the same network architecture, orthogonal networks appear to train faster than Gaussian networks, as measured both by the steepness of the loss curve early in training and the epoch at which the minimum of the validation loss curve is reached. In appendix B we verify that the superior training performance of orthogonal tanh networks holds even as the optimizer and loss function are varied, but that there only seems to be a minor improvement in performance with ReLU activations, as predicted by our analysis in section 5.3.

In figure 6, we compare the averaged best test losses for Gaussian and orthogonal initializations for MNIST (left) and CIFAR-10 (right) for a network width $n = 30$, with error bars indicating the standard deviation over our runs. In figure 7, we compare the averaged best test losses for $n = 30$ and $n = 100$ for both initializations. Consistent with previous studies [20, 24], orthogonal initialization results in a smaller test loss (to within the quoted error bars), and hence better generalization, for *all* architecture hyperparameters. Away from the infinite-width limit (when $L \gtrsim 20$), the value of the best loss for MNIST is roughly independent of depth for the orthogonal initialization, despite the increasing number of model parameters, but continues to degrade with depth for Gaussian initializations. This behavior is also seen to some extent in CIFAR-10 (figure 6, right): the loss appears to flatten with depth around $L \sim 30$, though the trend at large depth is more difficult to see with the larger variance compared to MNIST. These results further validate our predictions in section 6.2, as the plateau in orthogonal best-loss values for MNIST begins at approximately the same depth

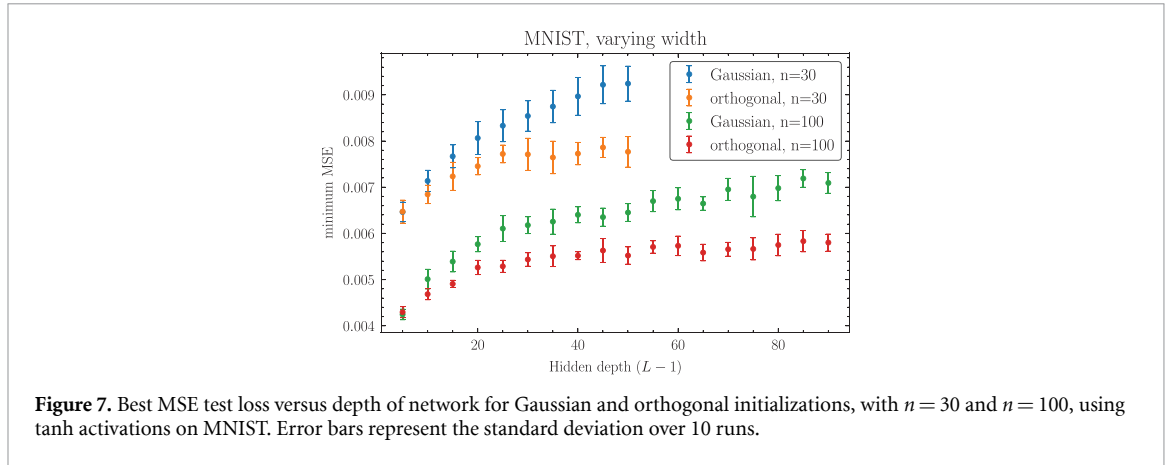


Figure 7. Best MSE test loss versus depth of network for Gaussian and orthogonal initializations, with $n = 30$ and $n = 100$, using tanh activations on MNIST. Error bars represent the standard deviation over 10 runs.

as when the NTK correlators become depth-independent. From figure 7, we also see that the plateau occurs at the same depth regardless of network width. This is consistent with our observation in section 6.2 that the magnitude of the orthogonal correlators is roughly independent of width.

For $L \lesssim 20$, the results enter the regime of validity of the infinite-width approximation, so the differences between orthogonal and Gaussian initializations become increasingly small. For $n = 30$, as we approach very small values of L , we find the best hidden depth for the MNIST classification task is $L - 1 = 3$ for both initializations; for CIFAR-10 the best hidden depth is $L - 1 = 4$ for both initializations. The information-theoretic arguments in [7] predict an optimal depth of 3.6 for width-30 tanh MLPs with an output layer of width $n_L = 10$, which is consistent with our results across both classification tasks and initializations. This is remarkable given the opposite sign of V for orthogonal initializations, as discussed in section 5.5. However, for $n = 100$, the predicted optimal depth is 12, while it is clear from figure 7 that the true optimum occurs for shallower networks. While we have demonstrated that it is possible to successfully train both orthogonal networks for $L \gtrsim 20$ (analogous to the results of [22]), since the best architecture for both widths is firmly in the $L/n \ll 1$ regime where both initializations give similar results, it is unclear if there is a practical benefit to using these very deep networks or orthogonal initializations. That said, in future work we intend to study whether the smaller correlator fluctuations in orthogonal networks lead to smaller variance in the trained network predictions, which could in fact provide a practical benefit for these initializations even in the $L/n \ll 1$ regime.

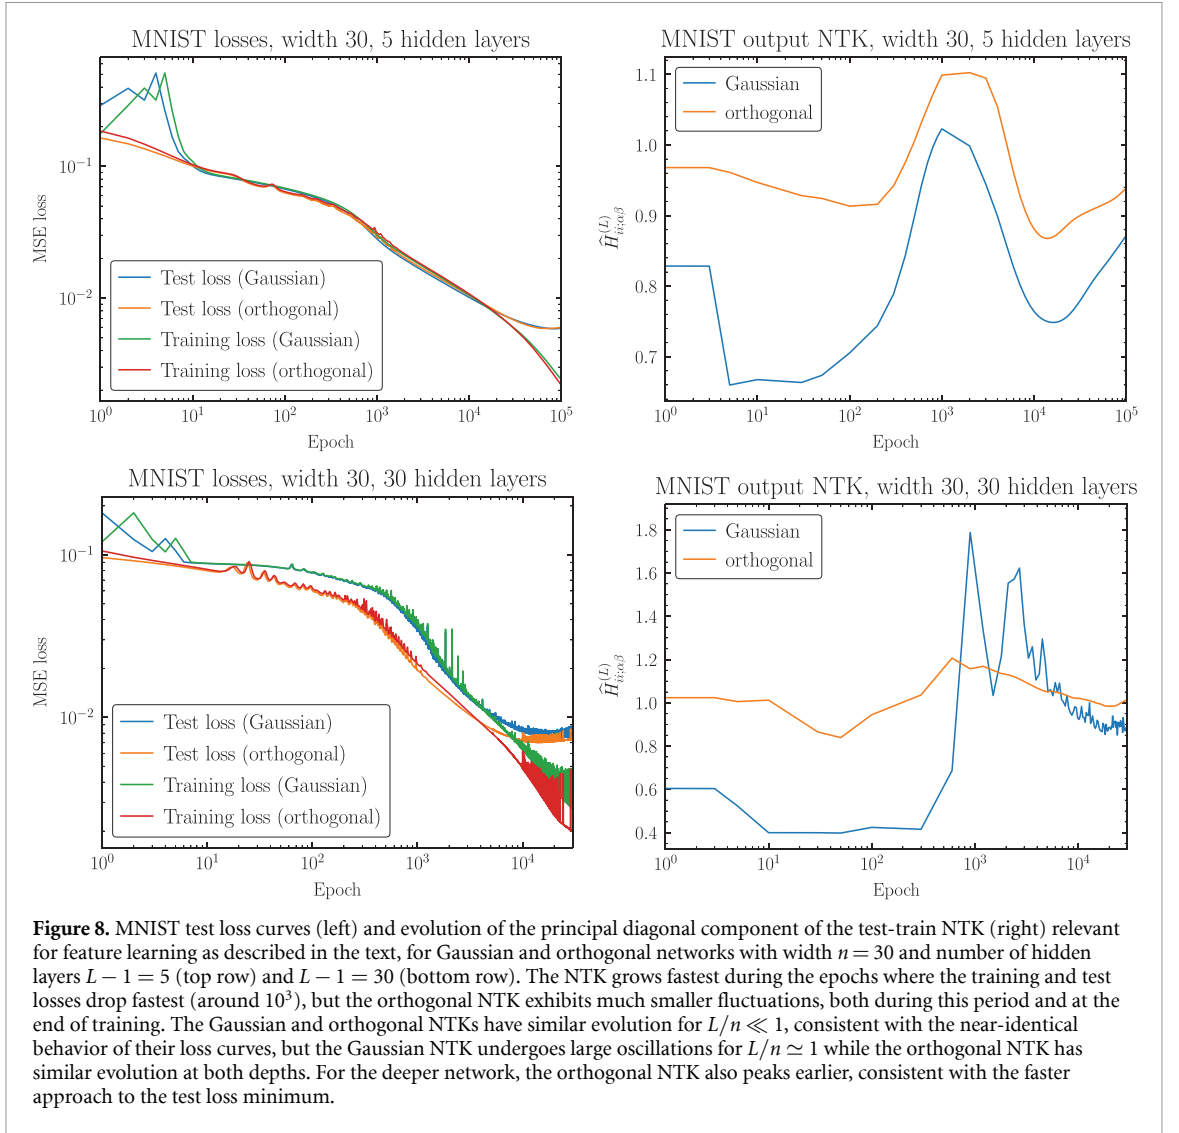
7.3. Empirical NTK dynamics, generalization, and feature learning

Using the same network architectures and training algorithms described above, we also measured components of the NTK relevant for generalization and feature learning over the course of training. To illustrate the typical behavior, we trained single instantiations of width-30 networks with hidden depths $L - 1 = 5$ and $L - 1 = 30$, for both choices of initialization. First, we measured the component of the output NTK $\widehat{H}_{ii;\alpha\beta}^{(L)}$ which is diagonal in neural indices but off-diagonal in sample indices, by taking one sample index α to be a random member of class ‘6’ from the MNIST test set, and averaging the second sample index β over all members of class ‘6’ in the training set¹⁹. This observable is related to generalization because it updates how training examples in a given class affect the prediction of a test example in the same class. In figure 8, we plot the evolution of the principal diagonal component in neural indices, namely $\widehat{H}_{ii;\alpha\beta}^{(L)}$ with $i = 7$, since this is the output node which labels class ‘6’ (in the standard MNIST convention with $i = 1$ labeling class ‘0’) and thus should be most strongly correlated with the chosen data points. To elucidate the relationship between the NTK evolution and feature learning, we also plot the test and training loss curves for the same architectures²⁰.

Several features are immediately apparent in figure 8. First, for shallow networks with $L/n \ll 1$ (top row), the loss curves for Gaussian and orthogonal initializations are essentially identical, as are the shapes of the

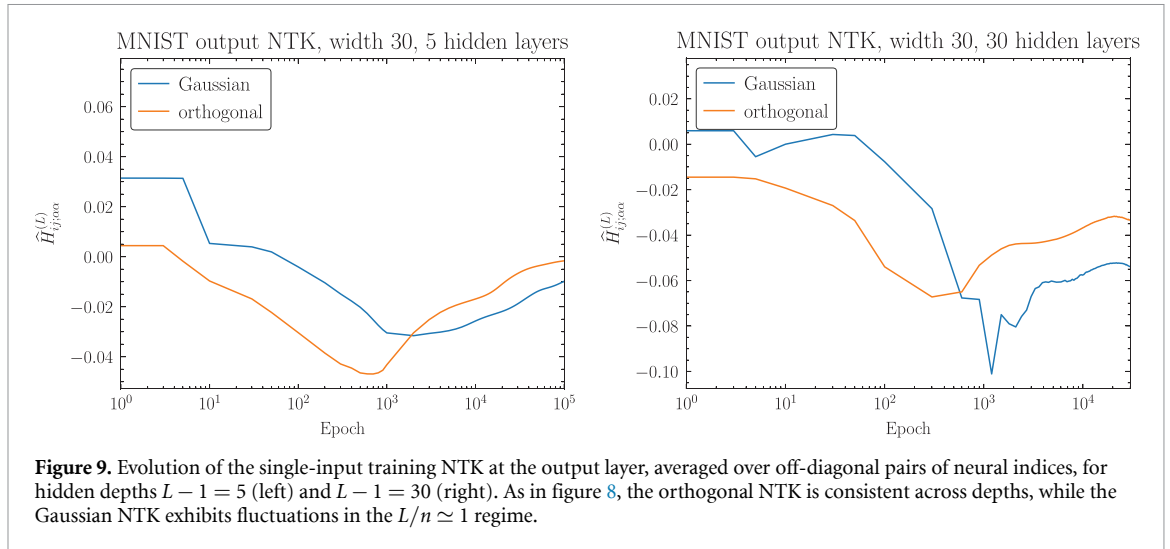
¹⁹ There is nothing special about the number 6, but at least from the perspective of topology, the handwritten representation of the digit 6 has more structure than the more obvious choices of 0 or 1.

²⁰ In this experiment, in order to directly compare with the analogous plots in section 7.2 and because we are not invoking early stopping, we use the validation set of size 10^4 as our test set.



NTK curves. The NTK does evolve over training at the $\approx 10\%$ level, consistent with the finite-width expectation, but the magnitude of the variation is larger for Gaussian than for orthogonal, due to the L/n growth of the dNTK correlators which drive the NTK evolution. The NTK peaks roughly around the epoch where the loss curve is steepest, and peaks slightly earlier for this instantiation of the Gaussian network than for the orthogonal network; the relative speed of training is also visible in the loss curves. This similarity between the two initializations stands in sharp contrast to the case of $L/n \approx 1$ (bottom row). There, the orthogonal network trains much faster (reflected in an NTK which peaks earlier) and exhibits markedly smaller variation over training than the Gaussian NTK, which has large oscillations around the epoch of fastest learning, followed by continued small oscillations as overfitting sets in, which are completely absent in the orthogonal case. Furthermore, the overall magnitude of the orthogonal NTK in the deep network is within $\sim 10\%$ of the analogous values for the shallow network, while the Gaussian NTK varies by factors of ~ 2 . All of these observations support our conjecture that large NTK fluctuations inhibit both training and generalization, leading to superior performances for orthogonal networks compared to Gaussian networks with the same architecture.

Finally, to illustrate feature learning in the same models, we measure the output NTK which is diagonal in sample indices but off-diagonal in neural indices, $\hat{H}_{ij;\alpha\alpha}^{(L)}$. Specifically, we take the sample index to be a random member of class ‘6’ from the training set, and average over all off-diagonal pairs of neural indices. The results are shown in figure 9. The overall size of the NTK is smaller than the diagonal component in figure 8, because the infinite-width NTK $\Theta_{\alpha\alpha}$ is purely diagonal in neural indices, so all deviations from zero are finite-width effects indicative of feature learning. Analogous to figure 8, the Gaussian and orthogonal NTKs behave similarly for $L/n \ll 1$. For $L/n \approx 1$, the Gaussian NTK exhibits fluctuations around the epoch the loss curve is steepest, while the orthogonal NTK is very similar to its behavior for small depths.



8. Conclusion and outlook

In this paper, we used the formalism of neural network correlation functions to help establish a firmer link between previous studies on the advantages of orthogonal initializations, and the important practical phenomena of feature learning and generalization. We identified the empirical depth-independence of the normalized NTK, dNTK, and ddNTK correlators—and the corresponding suppression of L/n fluctuations—as possible explanations for the faster training and superior generalization of orthogonal networks with tanh activations, and corroborated those explanations with measurements of the NTK during training. In these examples, the advantages of orthogonal initializations are small (improvements in test loss at the percent level, and improvements in training speed by an order-one factor in epochs), but very consistent across a wide variety of measurements. To our knowledge, there has been relatively little follow-up work in this direction since [24] (though see [44] for a recent application to sparse networks), perhaps because the theoretical connection to feature learning away from the infinite-width limit was lacking. We hope this work rekindles some interest in the benefits of alternative initialization schemes to improve neural network performance.

There are numerous aspects of our analysis which we intend to pursue in future work. Most importantly, we plan to derive the full expressions for the NTK and dNTK recursions, analogous to the Gaussian versions in [7], in order to demonstrate the depth independence of the correlators analytically. Such a derivation might also be able to shed light on the origin of the depth saturation scale $L \sim 20$: for instance, the components of the architecture or data which determines this number. Our numerical results seem to suggest that orthogonal networks have no L/n cutoff; while we worked to leading order in $1/n$ in this paper, it may be possible to prove or at least motivate depth-independence to all orders in $1/n$, as we calculated for deep linear networks in section 4. Furthermore, since the same NTK correlators we studied control the variance of the trained network predictions, we intend to study whether orthogonal networks exhibit reduced variance compared to Gaussian networks, which could be of use in applications in physics where uncertainty quantification is paramount [45]. It would also be interesting to see if a theoretically-optimal global learning rate for MLPs can be derived analytically from studying the NTK correlators, and perhaps relate this to the $\eta = \mathcal{O}(1/L)$ prescription and the training time $\tau = \mathcal{O}(\sqrt{L})$ under (non-tensorial) SGD from [20]. Finally, we intend to generalize our work to architectures and optimization algorithms of more practical use in modern applications. In particular, in light of the results of [14, 15], studying MLPs with LayerNorm and residual connections would be the next step in understanding the effect of orthogonal initializations on transformer architectures.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: [10.5281/zenodo.16688148](https://doi.org/10.5281/zenodo.16688148). Additional data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

We thank Yasaman Bahri, Nathaniel Craig, Jeffrey Pennington, Sho Yaida, and Kevin Zhang for stimulating conversations. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award Number DE-SC0023704. This work utilizes the computing resources of the HAL cluster [46] supported by the National Science Foundation's Major Research Instrumentation program, Grant #1725729, as well as the University of Illinois Urbana-Champaign. YK and DR thank the Aspen Center for Physics, which is supported by National Science Foundation Grant PHY-2210452, for hospitality during the completion of this work. DR acknowledges support from the National Science Foundation under Cooperative Agreement PHY-2019786 (the NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>) and appreciates the sanction of Sequoia Capital.

Appendix A. Measurements of multi-input correlators

Since the behavior of practical neural networks during training relies on correlations between multiple inputs from the training set, we measured the multi-input 4-point vertex V and the multi-input NTK correlators A , B , D , and F on random input data to verify that they exhibit the same behavior with depth as the single-input correlators studied in the main text. In particular, these NTK correlators are leading order in the global learning rate η and govern the leading finite-width corrections to the mean trained network predictions [7].

A.1. Measurements of V

We initialized an ensemble of 1000 rectangular networks of width $n = 100$ and depth $L = 10$ with linear, ReLU, and tanh activations, and measured V for both Gaussian and orthogonal initializations, as well as a mixed initialization with a Gaussian first layer and all other layers orthogonal. The input was taken to be a random vector of length n with entries drawn uniformly from $[0, 1]$, and we used the same input for each network configuration. We studied correlations between one, two, three, and four inputs; since V accepts four inputs as arguments, for the case of two distinct inputs, we considered both the case of three identical arguments and one different (3–1) and the case of two pairs of identical arguments (2–2). For two distinct arguments with sample indices α and β we measured

$$\tilde{V}_{(\alpha\alpha)(\beta\beta)} \equiv \frac{V_{(\alpha\alpha)(\beta\beta)}}{\frac{1}{2} (K_{\alpha\alpha}K_{\beta\beta} + (K_{\alpha\beta})^2)}, \quad \tilde{V}_{(\alpha\alpha)(\alpha\beta)} \equiv \frac{V_{(\alpha\alpha)(\alpha\beta)}}{K_{\alpha\alpha}K_{\alpha\beta}}. \quad (\text{A.1})$$

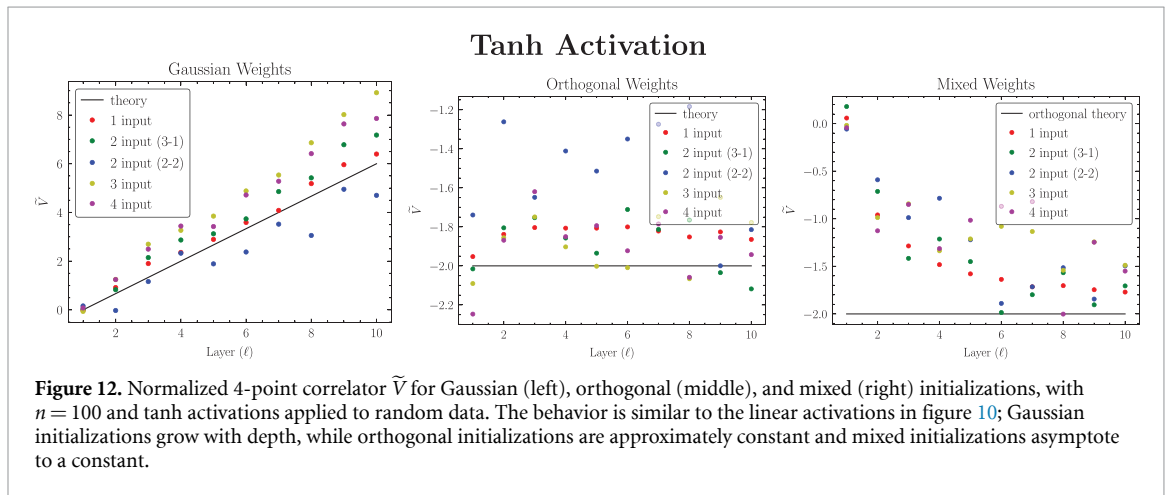
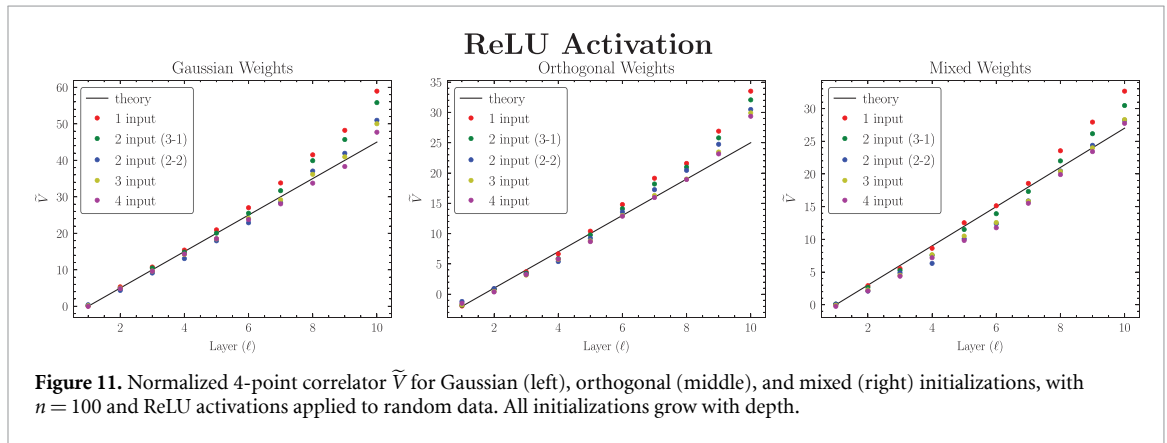
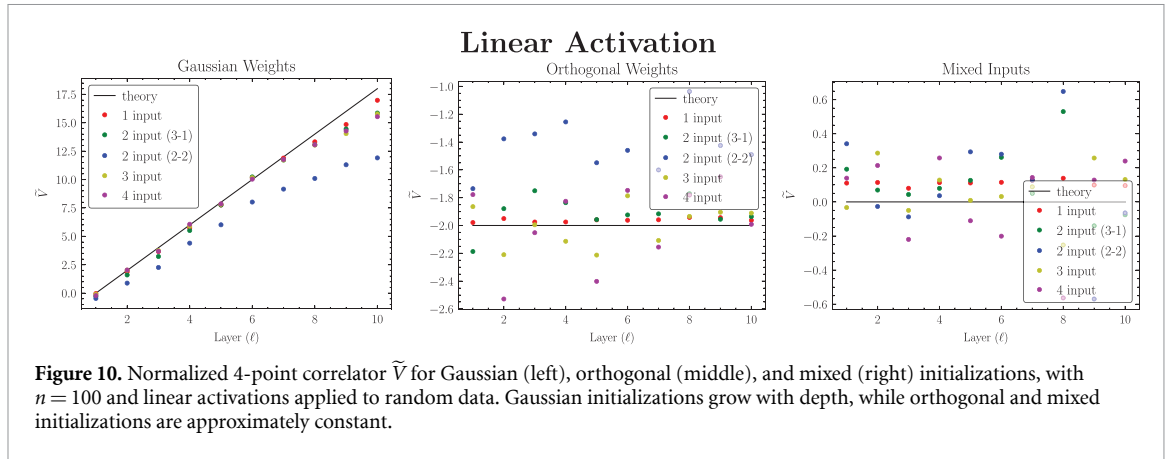
For three distinct arguments α, β, γ , we measured

$$\tilde{V}_{(\alpha\beta)(\gamma\alpha)} \equiv \frac{V_{(\alpha\beta)(\gamma\alpha)}}{\frac{1}{2} (K_{\alpha\beta}K_{\alpha\gamma} + K_{\alpha\alpha}K_{\beta\gamma})}, \quad (\text{A.2})$$

and for four distinct arguments $\alpha, \beta, \gamma, \delta$, we measured

$$\tilde{V}_{(\alpha\beta)(\gamma\delta)} \equiv \frac{V_{(\alpha\beta)(\gamma\delta)}}{\frac{1}{3} (K_{\alpha\beta}K_{\gamma\delta} + K_{\alpha\gamma}K_{\beta\delta} + K_{\alpha\delta}K_{\beta\gamma})}. \quad (\text{A.3})$$

The results are shown in figures 10–12, where for reference, we also show the theoretical values for the single-input correlators in solid black. Consistent with the single-input behavior, the multi-input correlators grow approximately linearly with depth for Gaussian initializations with all activations. The orthogonal correlators also mirror the single-input behavior: they grow linearly for ReLU activations, are approximately constant for tanh and linear activations, and asymptote to a constant for mixed initializations. The larger scatter for the multi-input orthogonal correlators is likely due to our choice of normalization; the kernel $K_{\alpha\beta}$ for $\alpha \neq \beta$ is typically quite small since random high-dimensional vectors are approximately orthogonal, and a more detailed analysis using the decomposition of the multi-input correlator into irreducible representations of its symmetries would likely lead to more stable behavior.



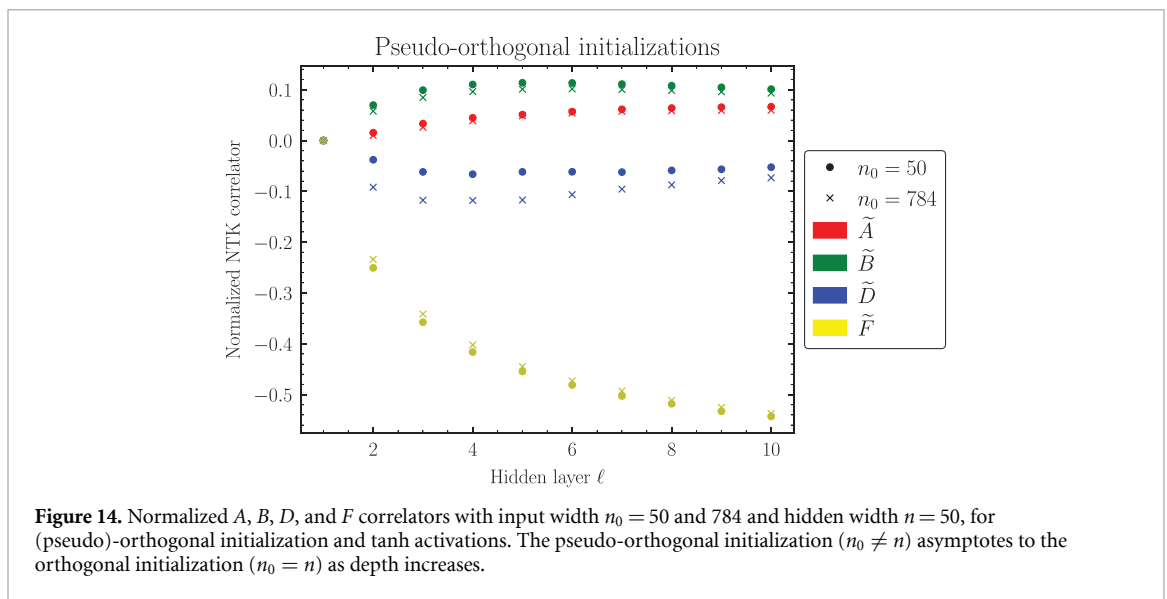
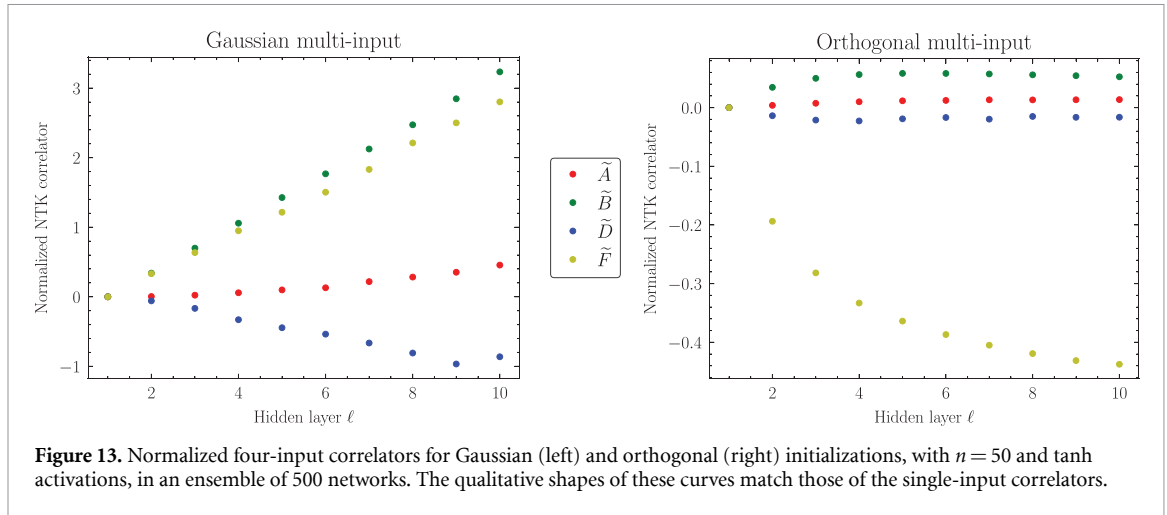
A.2. Measurements of NTK correlators

We initialized an ensemble of 500 rectangular networks of width $n = 50$ and depth $L = 10$ with tanh activations, and measured the normalized correlators of equations (6.7) and (6.8) with four distinct random inputs $\alpha_1, \alpha_2, \alpha_3, \alpha_4$, with entries uniform on $[0, 1]$. As with the 4-input V , we normalized the correlator by Θ^2 or ΘK as appropriate with sample index structure analogous to equation (A.3):

$$\left(\tilde{A}, \tilde{B}\right)_{\alpha\beta\gamma\delta} \equiv \frac{(A, B)_{\alpha\beta\gamma\delta}}{\frac{1}{3} (\Theta_{\alpha\beta}\Theta_{\gamma\delta} + \Theta_{\alpha\gamma}\Theta_{\beta\delta} + \Theta_{\alpha\delta}\Theta_{\beta\gamma})} \tag{A.4}$$

and

$$\left(\tilde{D}, \tilde{F}\right)_{\alpha\beta\gamma\delta} \equiv \frac{(D, F)_{\alpha\beta\gamma\delta}}{\frac{1}{6} (\Theta_{\alpha\beta}K_{\gamma\delta} + \Theta_{\alpha\gamma}K_{\beta\delta} + \Theta_{\alpha\delta}K_{\beta\gamma} + \Theta_{\gamma\delta}K_{\alpha\beta} + \Theta_{\beta\delta}K_{\alpha\gamma} + \Theta_{\beta\gamma}K_{\alpha\delta})}. \tag{A.5}$$



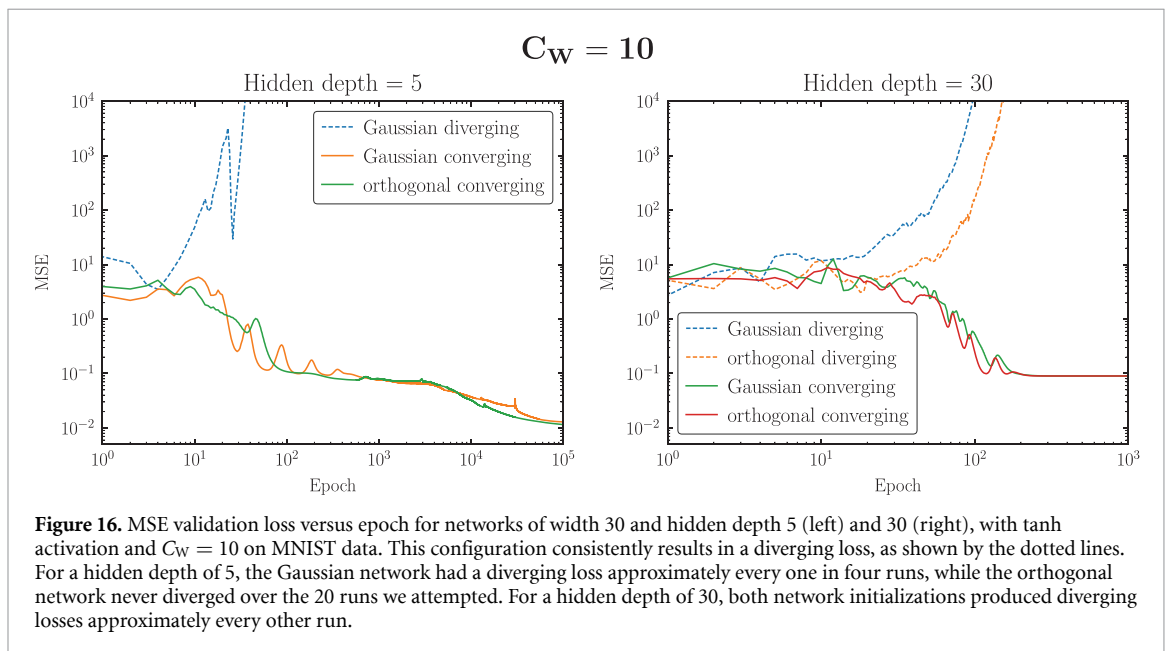
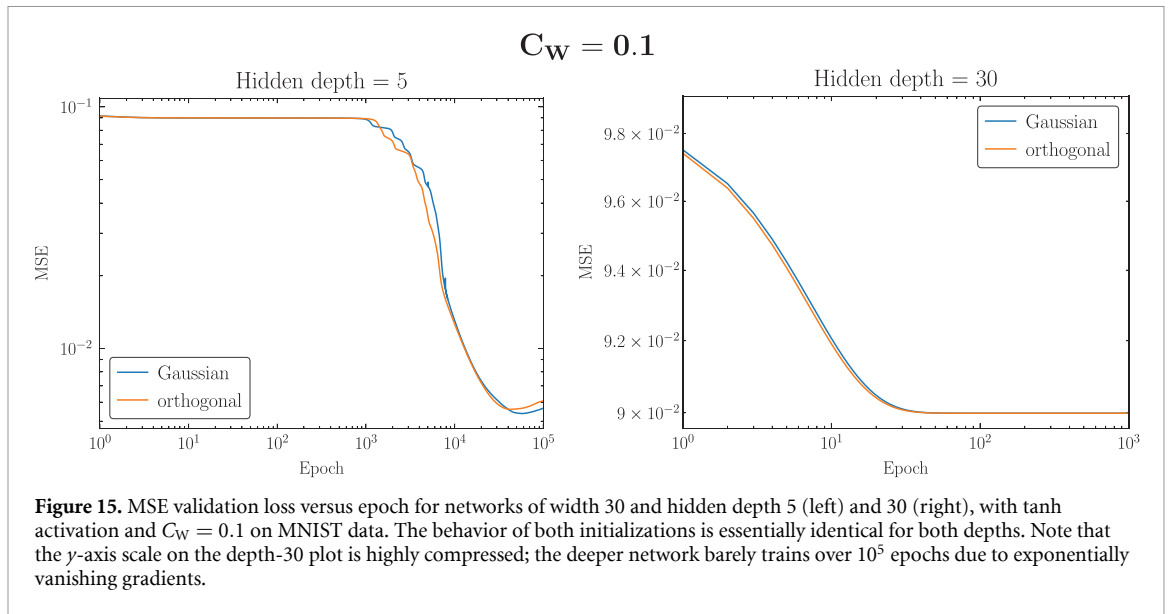
The results are shown in figure 13 for Gaussian and orthogonal initializations. As was the case with V , we find that these multi-input NTK correlators at initialization have same qualitative behavior as the single-input correlators, with Gaussian initializations leading to linearly-growing normalized correlators, but orthogonal initializations saturating to a finite value. We intend to explore the behavior of these multi-input correlators through analytical calculations in future work.

Finally, we verified that the asymptotic behavior of the NTK correlators is independent of the particular choice of mixed initialization in the first layer. Specifically, the QR decomposition used in our training experiments does not lead to a rectangular matrix with i.i.d. Gaussian entries, because the rows or columns are orthogonal. Figure 14 compares the correlators in a rectangular network with $n = 50$ to those from random inputs with size 784 (analogous to the MNIST dataset), with first-layer weights in both cases given by the QR algorithm. The differences are largest in the first layer, as expected, but both initializations asymptote to the same values at large depth.

Appendix B. Variations on training experiments

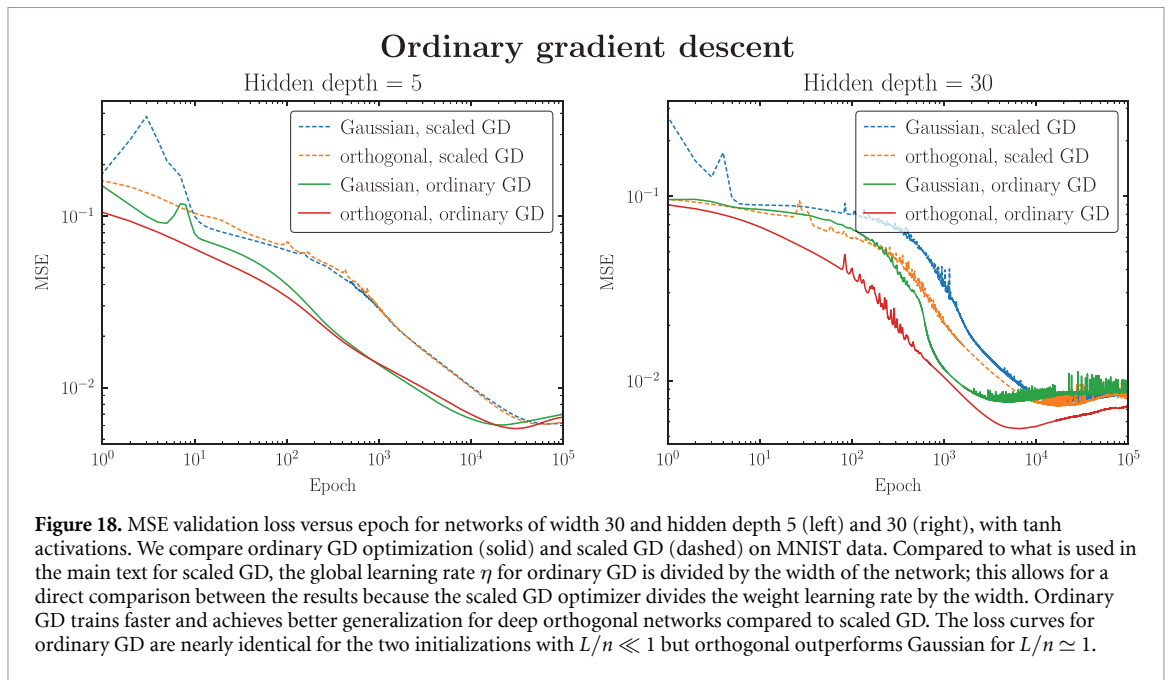
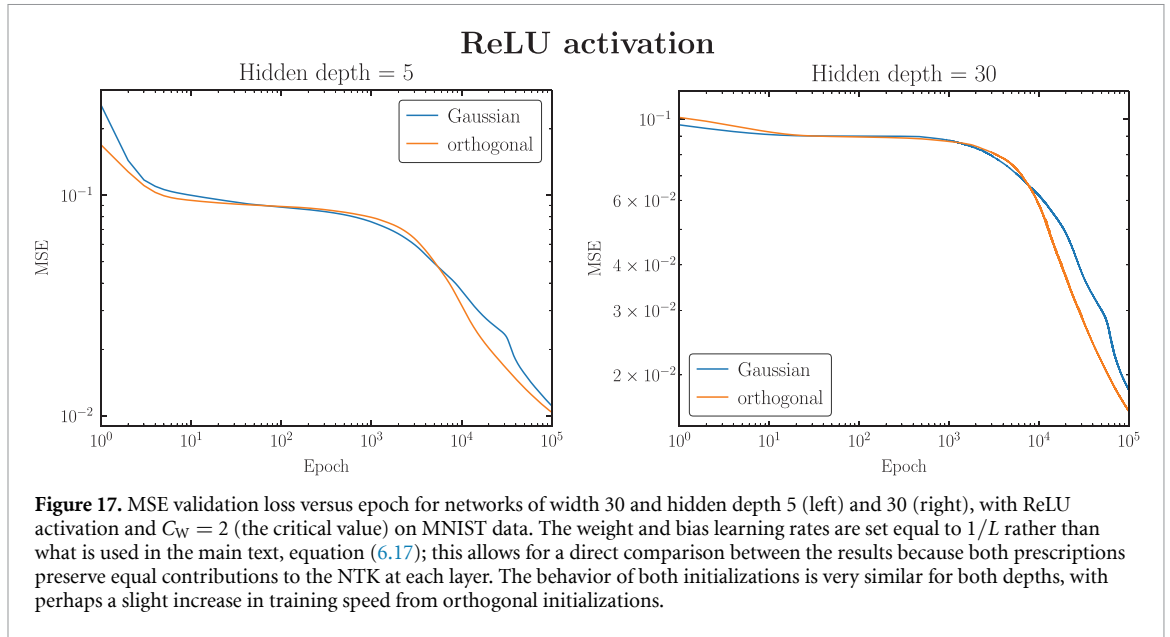
In this appendix we make various perturbations to our choices of hyperparameters to confirm that our conclusions about the training and generalization advantages of orthogonal networks are robust. All networks are trained on MNIST with hidden width 30.

First, figures 15 and 16 show the effect of choosing C_W smaller or larger than the critical value of $C_W = 1$. As expected by the analysis of [7], $C_W < 1$ (figure 15) leads to a vanishing gradient problem where networks either begin training much later (left, $L/n \ll 1$) or fail to train altogether (right, $L/n \simeq 1$). Similarly, $C_W > 1$



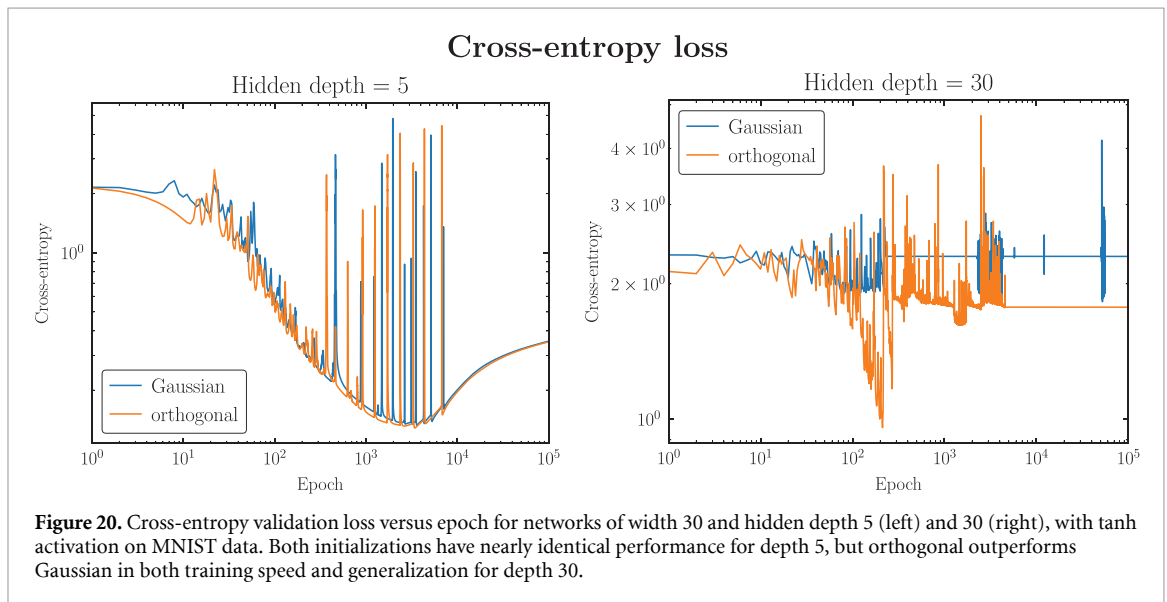

(figure 16) leads to an exploding gradient problem where an order-one fraction of runs give a diverging loss. Crucially, orthogonal initializations suffer from exactly the same issues as Gaussian initializations, confirming that the benefits of orthogonal initializations are only seen at networks tuned to criticality. Similarly, figure 17 shows the results of changing the activation function to ReLU, now taking $C_W = 2$ to maintain criticality. In addition, we make the associated change in the weight and bias learning rates from [7] to preserve equal contributions to the NTK at each layer: specifically, the bias learning rates are now independent of depth, but both weight and bias learning rates have an overall scaling with the total depth $1/L$. Curiously, both shallow and deep networks take considerably longer to train than for tanh activations (overfitting has not occurred by epoch 10^5), but orthogonal initializations offer only a mild improvement over Gaussian, if any. This is consistent with the analysis in the main text, where orthogonal ReLU networks still have correlators that grow linearly with depth, though at a slightly reduced slope compared to Gaussian ReLU networks.

Next, we consider varying the optimizer. Figure 18 shows the result of performing ordinary full-batch GD compared to the tensorial scaled GD defined in the main text. The global learning rate for ordinary GD was taken to be $10/n$ so that the effective weight learning rate is the same as for scaled GD, c.f. equation (6.3). Comparing Gaussian and orthogonal initializations for the same network architectures, the performance is nearly identical for $L/n \ll 1$ but orthogonal is clearly superior for $L/n \simeq 1$. On the other hand, comparing the two optimizers, there seems to be a clear advantage of ordinary GD compared to scaled GD for



orthogonal initializations away from the infinite-width limit, while for $L/n \ll 1$, ordinary GD trains faster but does not achieve a marked improvement in best validation loss for either initialization. It is possible that the depth-independence of orthogonal NTK correlators reduces the importance of equal contributions to the NTK at each layer, which motivated the scaled GD prescription in [7] for Gaussian initializations. Relatedly, it is possible that having *some* correlators growing with depth is in fact advantageous for feature learning, which may affect the optimal choice of η for each optimizer and initialization. In future work, where we will compute the NTK correlators for orthogonal initializations with both GD choices, we will investigate this phenomenon further. Figure 19 shows the result of using the Adam optimizer, which is much more common in practical applications than full-batch GD. To achieve reasonable performance, we changed the learning rate to $\eta = 10^{-3}$ (the default for Adam), and we caution that performance should not be directly compared with scaled GD. However, the comparison of the two initialization choices for the same optimizer shows once again that orthogonal outperforms Gaussian away from the infinite-width limit.

Finally, we consider the effect of changing the loss function to the more-typical cross-entropy loss paired with sigmoid activations on the output layer. Figure 20 (left) shows that Gaussian and orthogonal networks



still have essentially identical performance in the $L/n \ll 1$ limit, though the GD oscillations are much more pronounced (as also noted in [43]). Interestingly, the oscillations seem to damp out shortly after overfitting sets in, a phenomenon we intend to explore in future work. However, for $L/n \simeq 1$ (right panel), the Gaussian network completely fails to train, with a loss that only briefly surpasses that of a random 10-class classifier, $-\log(0.1) \approx 2.3$. The orthogonal network does train but overfitting occurs extremely fast, at epoch 200. Even so, the orthogonal network loss plateaus at a value below that of the Gaussian network, suggesting that somewhat better performance is still achieved with cross-entropy loss.

ORCID iDs

Hannah Day  0000-0002-4496-5600

Yonatan Kahn  0000-0002-9379-1838

Daniel A Roberts  0000-0002-5755-2274

References

- [1] Poole B, Lahiri S, Raghu M, Sohl-Dickstein J and Ganguli S 2016 Exponential expressivity in deep neural networks through transient chaos *Advances in Neural Information Processing Systems* vol 29
- [2] Schoenholz S S, Gilmer J, Ganguli S and Sohl-Dickstein J 2017 Deep information propagation *ICLR* (arXiv:1611.01232)

- [3] Hanin B 2018 Which neural net architectures give rise to exploding and vanishing gradients? *Advances in Neural Information Processing Systems* vol 31
- [4] Hanin B and Rolnick D 2018 How to start training: the effect of initialization and architecture *Advances in Neural Information Processing Systems* vol 31
- [5] Hanin B and Nica M 2020 Finite depth and width corrections to the neural tangent kernel *ICLR* (arXiv:1909.05989)
- [6] Yaida S 2020 Non-gaussian processes and neural networks at finite widths *Mathematical and Scientific Machine Learning* (PMLR) pp 165–92
- [7] Roberts D A, Yaida S and Hanin B 2022 *The Principles of Deep Learning Theory* (Cambridge University Press) p 5
- [8] Park D, Sohl-Dickstein J, Le Q and Smith S 2019 The effect of network width on stochastic gradient descent and generalization: an empirical study *Int. Conf. on Machine Learning* (PMLR) pp 5042–51
- [9] Yang G and Hu E J 2021 Tensor programs IV: feature learning in infinite-width neural networks *Int. Conf. on Machine Learning* (PMLR) pp 11727–37
- [10] Yaida S 2022 Meta-principled family of hyperparameter scaling strategies (arXiv:2210.04909)
- [11] Kingma D P and Ba J 2015 Adam: a method for stochastic optimization *ICLR* (arXiv:1412.6980)
- [12] Loshchilov I and Hutter F 2019 Decoupled weight decay regularization *ICLR* (arXiv:1711.05101)
- [13] Ba J L, Kiros J R and Hinton G E 2016 Layer normalization *NIPS Deep Learning Symposium* (arXiv:1607.06450)
- [14] Doshi D, He T and Gromov A 2023 Critical initialization of wide and deep neural networks using partial jacobians: general theory and applications *Advances in Neural Information Processing Systems* vol 36 40054–95
- [15] He T, Doshi D and Gromov A 2022 Autoinit: automatic initialization via jacobian tuning (arXiv:2206.13568)
- [16] Jacot A, Gabriel F and Hongler C 2018 Neural tangent kernel: convergence and generalization in neural networks *Advances in Neural Information Processing Systems* vol 31
- [17] Arora S, Du S S, Hu W, Li Z, Salakhutdinov R R and Wang R 2019 On exact computation with an infinitely wide neural net *Advances in Neural Information Processing Systems* vol 32
- [18] Lee J et al 2019 Wide neural networks of any depth evolve as linear models under gradient descent *Advances in Neural Information Processing Systems* vol 32
- [19] Saxe A M, McClelland J L and Ganguli S 2013 Exact solutions to the nonlinear dynamics of learning in deep linear neural networks (arXiv:1312.6120)
- [20] Pennington J, Schoenholz S and Ganguli S 2017 Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice *Advances in Neural Information Processing Systems* vol 30
- [21] Pennington J, Schoenholz S and Ganguli S 2018 The emergence of spectral universality in deep networks *Int. Conf. on Artificial Intelligence and Statistics* (PMLR) pp 1924–32
- [22] Xiao L, Bahri Y, Sohl-Dickstein J, Schoenholz S and Pennington J 2018 Dynamical isometry and a mean field theory of CNNs: how to train 10,000-layer vanilla convolutional neural networks *Int. Conf. on Machine Learning* (PMLR) pp 5393–402
- [23] Hu W, Xiao L and Pennington J 2020 Provable benefit of orthogonal initialization in optimizing deep linear networks *ICLR* (arXiv:2001.05992)
- [24] Huang W, Du W and Xu R 2021 On the neural tangent kernel of deep networks with orthogonal initialization *Proc. 13th Int. Joint Conf. on Artificial Intelligence, Int. Joint Conf. on Artificial Intelligence Organization*
- [25] Ringel Z, Rubin N, Mor E, Helias M and Seroussi I 2025 Applications of statistical field theory in deep learning
- [26] Krizhevsky A, Sutskever I and Hinton G E 2012 Imagenet classification with deep convolutional neural networks *Advances in Neural Information Processing Systems* vol 25
- [27] Vaswani A et al 2017 Attention is all you need *Advances in Neural Information Processing Systems* vol 30
- [28] Dinan E, Yaida S and Zhang S 2023 Effective theory of transformers at initialization (arXiv:2304.02034)
- [29] Nguyen A and Reddy G 2024 Differential learning kinetics govern the transition from memorization to generalization during in-context learning *ICLR* (arXiv:2412.00104)
- [30] Sirunyan A M et al shape CMS collaboration 2020 Inclusive search for highly boosted Higgs bosons decaying to bottom quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV *J. High Energy Phys. JHEP12(2020)085*
- [31] Necib L, Ostdiek B, Lisanti M, Cohen T, Freytsis M, Garrison-Kimmel S, Hopkins P F, Wetzel A and Sanderson R 2020 Evidence for a vast prograde stellar stream in the solar vicinity *Nat. Astron.* **4** 1078–83
- [32] Weingarten D 1978 Asymptotic behavior of group integrals in the limit of infinite rank *J. Math. Phys.* **19** 999
- [33] Collins B and Śniady P 2006 Integration with respect to the Haar measure on unitary, orthogonal and symplectic group *Commun. Math. Phys.* **264** 773–95
- [34] Collins B and Matsumoto S 2009 On some properties of orthogonal Weingarten functions *J. Math. Phys.* **50** 113516
- [35] Y Gu 2013 Moments of random matrices and Weingarten functions *Master's Thesis* Queen's University
- [36] Collins B, Matsumoto S and Novak J 2022 The Weingarten calculus *Not. Am. Math. Soc.* **69** 734–45
- [37] Dyer E and Gur-Ari G 2019 Asymptotics of wide networks from Feynman diagrams *ICLR* (arXiv:1909.11304)
- [38] Banta I, Cai T, Craig N and Zhang Z 2024 Structures of neural network effective theories *Phys. Rev. D* **109** 105007
- [39] Zavattone-Veth J and Pehlevan C 2021 Exact marginal prior distributions of finite bayesian neural networks *Advances in Neural Information Processing Systems* vol 34 pp 3364–75
- [40] Chollet F et al 2015 Keras (available at: <https://keras.io>)
- [41] Mezzadri F 2007 How to generate random matrices from the classical compact groups *Not. AMS* **54** 592–604
- [42] Hui L and Belkin M 2021 Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks *ICLR* (arXiv:2006.07322)
- [43] Cohen J, Kaur S, Li Y, Kolter J Z and Talwalkar A 2021 Gradient descent on neural networks typically occurs at the edge of stability *ICLR* (arXiv:2103.00065)
- [44] Nowak A J, Gniecki Ł, Szatkowski F and Tabor J 2025 Sparser, better, deeper, stronger: Improving sparse training with exact orthogonal initialization *ICML* (arXiv:2406.01755)
- [45] Chen T Y, Dey B, Ghosh A, Kagan M, Nord B and Ramachandra N, 2022 Interpretable uncertainty quantification in AI for HEP, Snowmass (arXiv:2208.03284)
- [46] Kindratenko V, Mu D, Zhan Y, Maloney J, Hashemi S H and Rabe B 2020 Hal: computer system for scalable deep learning *Practice and Experience in Advanced Research Computing, Pearc '20* (Association for Computing Machinery) pp 41–48