
Batch Normalization Orthogonalizes Representations in Deep Random Networks

Hadi Daneshmand
INRIA Paris
seyed.daneshmand@inria.fr

Amir Joudaki
ETH Zurich
amir.joudaki@inf.ethz.ch

Francis Bach
INRIA-ENS-PSL Paris
francis.bach@inria.fr

Abstract

This paper underlines a subtle property of batch-normalization (BN): Successive batch normalizations with random linear transformations make hidden representations increasingly orthogonal across layers of a deep neural network. We establish a non-asymptotic characterization of the interplay between depth, width, and the orthogonality of deep representations. More precisely, under a mild assumption, we prove that the deviation of the representations from orthogonality rapidly decays with depth up to a term inversely proportional to the network width. This result has two main implications: 1) Theoretically, as the depth grows, the distribution of the representation –after the linear layers– contracts to a Wasserstein-2 ball around an isotropic Gaussian distribution. Furthermore, the radius of this Wasserstein ball shrinks with the width of the network. 2) Practically, the orthogonality of the representations directly influences the performance of stochastic gradient descent (SGD). When representations are initially aligned, we observe SGD wastes many iterations to orthogonalize representations before the classification. Nevertheless, we experimentally show that starting optimization from orthogonal representations is sufficient to accelerate SGD, with no need for BN.

1 Introduction

Batch Normalization (BN) (Ioffe and Szegedy, 2015) enhances training across a wide range of deep network architectures and experimental setups (He et al., 2016; Huang et al., 2017; Silver et al., 2017). The practical success of BN has inspired research into the underlying mechanism of BN (Santurkar et al., 2018; Karakida et al., 2019; Arora et al., 2019b; Bjorck et al., 2018). BN influences first-order optimization methods by avoiding the rank collapse in deep representation (Daneshmand et al., 2020), direction-length decoupling of optimization (Kohler et al., 2018), influencing the convergence of the steepest descent (Arora et al., 2019b; Bjorck et al., 2018), and smoothing the optimization objective function (Santurkar et al., 2018; Karakida et al., 2019). However, the benefits of BN go beyond its critical role in optimization. For example, Frankle et al. (2021) shows that BN networks with random weights also achieve surprisingly high performance after only minor adjustments of their weights. This striking result motivates us to study the representational power of random networks with BN.

We study hidden representations across layers of a laboratory random BN with linear activations. Consider a batch of samples passing through consecutive BN and linear layers with Gaussian weights. The representations of these samples are perturbed by each random linear transformation, followed by a non-linear BN. At first glance, the deep representations appear unpredictable

after many stochastic and non-linear transformations. Yet, we show that these transformations orthogonalize the representations. To prove this statement, we introduce the notion of “orthogonality gap”, defined in Section 3, to quantify the deviation of representations from a perfectly orthogonal representation. Then, we prove that the orthogonality gap decays exponentially with the network depth and stabilizes around a term inversely related to the network width. More precisely, we prove

$$\mathbb{E} \left[\text{orthogonality gap} \right] = \mathcal{O} \left((1 - \alpha)^{\text{depth}} + \frac{\text{batchsize}}{\alpha \sqrt{\text{width}}} \right)$$

holds for $\alpha > 0$ that is an absolute constant under a mild assumption. In probability theoretic terms, we prove stochastic stability of the Markov chain of hidden representations (Kushner, 1967; Kushner and Yin, 2003; Khasminskii, 2011). The orthogonality of deep representations allows us to prove that the distribution of the representations after linear layers contracts to a Wasserstein-2 ball around isotropic Gaussian distribution as the network depth grows. Moreover, the radius of the ball is inversely proportional to the network width. Omitting details, we prove the following bound holds:

$$\text{Wasserstein}_2(\text{representations}, \text{Gaussian})^2 = \mathcal{O} \left((1 - \alpha)^{\text{depth}} (\text{batchsize}) + \frac{(\text{batchsize})^2}{\alpha \sqrt{\text{width}}} \right).$$

The above equation shows how depth, width, and batch size, interact with the Gaussian approximation of the representations. Since the established rate is exponential with depth, the distribution of the representations stays in a Wasserstein ball around isotropic Gaussian distribution after a few layers. Thus, BN not only stabilizes the distribution of the representations, which is its main promise (Ioffe and Szegedy, 2015), but also enforces Gaussian isotropic distribution in deep layers.

There is growing interest in bridging the gap between neural networks, as the most successful parametric methods for learning, and Gaussian processes and kernel methods, as well-understood classical models for learning (Jacot et al., 2018; Matthews et al., 2018; Lee et al., 2019; Bietti and Mairal, 2019; Huang et al., 2014). This link is inspired by studying random neural networks in the asymptotic regime of infinite width. The seminal work by Neal (1996) sparks that a single-layer network resembles a Gaussian process as its width goes to infinity. However, increasing the depth may significantly shift the distribution of the representations away from Gaussian (Ioffe and Szegedy, 2015). This distributional shift breaks the link between Gaussian processes and deep neural networks. To ensure Gaussian representations, Matthews et al. (2018) suggests increasing the width of the network proportional to the network depth. Here, we show that BN ensures Gaussian representations even for deep networks with *finite* width. This result bridges the link between deep neural networks and Gaussian processes in the regime of finite width. Many studies rely on deep Gaussian representations in an infinite width setting (Yang et al., 2019; Schoenholz et al., 2017; Pennington et al., 2018; Klambauer et al., 2017; De Palma et al., 2019). Our non-asymptotic Gaussian approximation can be incorporated into their analysis to extend these results to the regime of finite width.

Since training starts from random networks, representations in these networks directly influence training. Hence, recent theoretical studies has investigated the interplay between initial hidden representations and training (Daneshmand et al., 2020; Bjorck et al., 2018; Frankle et al., 2021; Schoenholz et al., 2017; Saxe et al., 2014; Bahri et al., 2020). In particular, it is known that hidden representations in random networks *without BN* become correlated as the network grows in depth, thereby drastically slowing training (Daneshmand et al., 2020; He et al., 2016; Bjorck et al., 2018; Saxe et al., 2014). On the contrary, we prove that deep representations in networks *with BN* are almost orthogonal. We experimentally validate that initial orthogonal representations can save training time that would otherwise be needed to orthogonalize them. By proposing a novel initialization scheme, we ensure the orthogonality of hidden representations. Such an initialization effectively avoids the training slowdown with depth for vanilla networks, with no need for BN. This observation further motivates studying the inner workings of BN to replace or improve it for deep learning. Theoretically, we made the following contributions:

1. For MLPs with batch normalization, linear activation, and Gaussian weights, we prove that representations across layers become increasingly orthogonal up to a constant inversely proportional to the network width.
2. Leveraging the orthogonality, we prove that the distribution of the representations contracts to a Wasserstein ball around a Gaussian distribution as the depth grows. Up to the best of

our knowledge, this is the first *non-asymptotic* Gaussian approximation for deep neural networks with finite width.

Experimentally, we made the following contribution¹:

3. Inspired by our theoretical understanding, we propose a novel weight initialization for standard neural networks that ensure orthogonal representations without BN. Experimentally, we show that this initialization effectively avoids training slowdown with depth in the absence of BN.

2 Preliminaries

2.1 Notations

Akin to Daneshmand et al. (2020), we focus on a Multi-Layer Perceptron (MLP) with batch normalization and linear activation. Theoretical studies of linear networks is a growing research area (Saxe et al., 2014; Daneshmand et al., 2020; Bartlett et al., 2019; Arora et al., 2019a). When weights are initialized randomly, linear and non-linear networks share similar properties such as the rank collapse issue studied in (Daneshmand et al., 2020). For ease of analysis, we assume activations are linear. Yet, we will argue that our findings extend to non-linear in Appendix G.

We use n to denote batch size, and d to denote the width across all layers, which we further assume is larger than the batch size $d \geq n$. Let $H_\ell \in \mathbb{R}^{d \times n}$ denote representations for n samples in layer ℓ , with $H_0 \in \mathbb{R}^{d \times n}$ corresponding to n input samples in the batch with d features. Successive representations are connected by Gaussian weight matrices $W_\ell \sim \mathcal{N}(0, I_d/d)$ as

$$H_{\ell+1} = \frac{1}{\sqrt{d}} \text{BN}(W_\ell H_\ell), \quad \text{BN}(M) = \text{diag}(MM^\top)^{-1/2} M, \quad (1)$$

where $\text{diag}(\cdot)$ zeros out off-diagonal elements of its input matrix, and the scaling factor $1/\sqrt{d}$ ensures that all matrices $\{H_k\}_k$ have unit Frobenius norm (see Appendix 2). The BN function in Eq. (1) differs slightly from the commonly used definition for BN as the mean correction is omitted. However, Daneshmand et al. (2020) shows this difference does not change the network properties qualitatively. Appendix G further discusses about the role of the mean correction term in our analysis. Readers can find all the notations in Appendix Table 2.

2.2 The linear independence of hidden representations

Daneshmand et al. (2020) observe that if inputs are linearly independent, then their hidden representations remain linearly independent in all layers as long as $d = \Omega(n^2)$. Under technical assumptions, Daneshmand et al. (2020) establishes a lower-bound on the average of the rank of hidden representations over infinite layers. Based on this study, we assume that the linear independence holds and build our analysis upon that. This avoids restating the technical assumptions of Daneshmand et al. (2020) and also further technical refinements of their theorems. The next assumption presents the formal statement of the linear independence property.

Assumption $\mathcal{A}_1(\alpha, \ell)$. *There exists an absolute positive constant α such that the minimum singular value of H_k is greater than (or equal to) α for all $k = 1, \dots, \ell$.*

The linear independence of the representations is a shared property across all layers. However, the representations constantly change when passing through random layers. In this paper, we mathematically characterize the dynamics of the representations across layers.

3 Orthogonality of deep representations

3.1 A warm-up observation

To illustrate the difference between BN and vanilla networks, we compare hidden representations of two input samples across the layers of these networks. Figure 1 plots the absolute value of

¹Implementations are available at <https://github.com/hadidaneshmand/batchnorm21.git>

cosine similarity of these samples across layers. This plot shows a stark contrast between vanilla and BN networks: While representations become increasingly orthogonal across layers of a BN network, they become increasingly aligned in a vanilla network. More specifically, we observe that BN is able to orthogonalize almost aligned representations; while, the vanilla network provides almost aligned representation of two orthogonal samples in deep layers. While this behaviour has been theoretically studied for vanilla networks (Daneshmand et al., 2020; Bjorck et al., 2018; Saxe et al., 2014), up to the best of our knowledge, it is not theoretically analyzed for BN networks for networks with finite width regime. In the following section, we formalize and prove this orthogonalizing property for BN networks with finite widths.




Figure 1: **Orthogonality: BN vs. vanilla networks.** The horizontal axis shows the number of layers, and the vertical axis shows the absolute value of cosine similarity between two samples across the layers ($d = 32$). Mean and 95% confidence intervals of 20 independent runs.

3.2 Theoretical analysis

The notion of orthogonality gap plays a central role in our analysis. Given the hidden representation $H \in \mathbb{R}^{d \times n}$, matrix $H^\top H$ constitutes inner products between representations of different samples. Note that $H^\top H \in \mathbb{R}^{n \times n}$ is different from the covariance matrix $HH^\top/n \in \mathbb{R}^{d \times d}$. The orthogonality gap of H is defined as the deviations of $H^\top H$ from identity matrix, after proper scaling. More precisely, define $V : \mathbb{R}^{d \times n} \setminus \mathbf{0} \rightarrow \mathbb{R}_+$ as

$$V(H) := \left\| \left(\frac{1}{\|H\|_F^2} \right) H^\top H - \left(\frac{1}{\|I_n\|_F^2} \right) I_n \right\|_F. \quad (2)$$

The following theorem establishes a bound on the orthogonality of representation in layer ℓ .

Theorem 1. *Under Assumption $\mathcal{A}_1(\alpha, \ell)$, the following holds:*

$$\mathbb{E}[V(H_{\ell+1})] \leq 2 \left(1 - \frac{2}{3} \alpha \right)^\ell + \frac{3n}{\alpha \sqrt{d}}. \quad (3)$$

Assumption \mathcal{A}_1 is studied by Daneshmand et al. (2020) who note that $\mathcal{A}_1(\alpha, \infty)$ holds as long as $d = \Omega(n^2)$. If \mathcal{A}_1 does not hold, one can still prove that there is a function of representations that decays with depth up to a constant (see Appendix C).

The above result implies that BN is an approximation algorithm for orthogonalizing the hidden representations. If we replace $\text{diag}(M)^{-1}$ by $(M)^{-1}$ in BN formula, in Eq. (1), then all the hidden representation will become exactly orthogonal. However, computing the inverse of a *non-diagonal* $d \times d$ matrix is computationally expensive, which must repeat for all layers throughout training, and differentiation must propagate back through this inversion. The diagonal approximation in BN significantly reduces the computational complexity of the matrix inversion. Since the orthogonality gap decays at an exponential rate with depth, the approximate orthogonality is met after only a few layers. Interestingly, this yields a desirable cost-accuracy trade-off: For a larger width, the orthogonality is more accurate, due to term $1/\sqrt{d}$ in the orthogonality gap, and also the computational gain is more significant.

From a different angle, Theorem 1 proves the stochastic stability of the Markov chain of hidden representations. In expectation, stochastic processes may obey an inherent Lyapunov-type of stability

(Kushner, 1967; Kushner and Yin, 2003; Khasminskii, 2011). One can analyze the mixing and hitting times of Markov chains based on the stochastic stability of Markov chains, (Kemeny and Snell, 1976; Eberle, 2009). In our analysis, the orthogonality gap is a Lyapunov function characterizing the stability of the chain of hidden representations. This stability opens the door to more theoretical analysis of this chain, such as studying mixing and hitting times. For example the stability can be used for mixing analysis of the chain $\{H_1, \dots, H_\ell, \dots\}$. Let π denote the stationary distribution of this chain. Under a particular stochastic stability condition called geometric drift condition,

$$|\mathbb{E}[\varphi(H_\ell)] - \mathbb{E}_{H \sim \pi}[\varphi(H)]| \leq \alpha^\ell |\mathbb{E}[\varphi(H_0)] - \mathbb{E}_{H \sim \pi}[\varphi(H)]|$$

holds for $\alpha \in (0, 1)$ and measurable function $\varphi : \mathbb{R}^{d \times n} \rightarrow \mathbb{R}$ (see Theorem 3.6 of (Hairer, 2010)). The drift condition holds if there exists a Lyapunov function $L : \mathbb{R}^{d \times n} \rightarrow [0, 1]$ and constant $K \geq 0$ such that the following holds:

$$E[L(H_{\ell+1})|H_\ell] \leq \gamma L(H_\ell) + K.$$

We prove that the above condition holds in Appendix Eq (12). However, such theoretical analyses may not be of interest to the machine learning community, hence we focus on the implications of these results for understanding the underpinnings of neural networks.

It is helpful to compare the orthogonality gap in BN networks to studies on vanilla networks (Bjorck et al., 2018; Daneshmand et al., 2020; Saxe et al., 2014). The function implemented by a vanilla linear network is a linear transformation as the product of random weight matrices. The spectral properties of the product of i.i.d. random matrices are the subject of extensive studies in probability theory (Bougerol et al., 2012). Invoking these results, one can check that the orthogonality gap of the hidden representations in vanilla networks rapidly increases since the rank of representations converges to a one matrix as the depth grows (Daneshmand et al., 2020). We refer readers to Appendix D for more details about vanilla networks.

Remarkably, Agarwal et al. (2021) proves if activation functions obey a self-normalization property, then a specific kernel of hidden representations becomes well-conditioned as the network depth grows. However, it is not clear how to impose the self-normalization property. Here, we establish the whitening for an explicitly defined normalization used in practice.

3.3 Experimental validations

Our experiments presented in Fig. 2a validate the exponential decay rate of V with depth. In this plot, we see that $\log(V_\ell)$ linearly decreases for $\ell = 1, \dots, 20$, then it wiggles around a small constant. Our experiments in Fig. 2b suggest that the $\mathcal{O}(1/\sqrt{d})$ dependency on width is almost tight. Since $V(H_\ell)$ rapidly converges to a ball, the average of $V(H_\ell)$ over layers estimates the radius of this ball. This plot shows how the average of $V(H_\ell)$, over 500 layers, changes with the network width, validating the $\mathcal{O}(1/\sqrt{d})$ dependency implied by Theorem 1.




Figure 2: **Orthogonality gap vs. depth and width.** Left: $\log(V(H_\ell))$ vertically versus ℓ horizontally. Right: $\log(\frac{1}{500} \sum_{\ell=100}^{600} V(H_\ell))$ vertically versus $\log(d)$ horizontally. The chain starts from a diagonal H_0 with one relatively large diagonal value. This structure imposes a large orthogonality gap for H_0 . Mean and 95% confidence interval of 20 independent runs.

Consider an input matrix $H_0 \in \mathbb{R}^{d \times n}$ containing n samples in \mathbb{R}^d . Assuming that elements of this matrix are i.i.d. zero-mean and unit variance random variables, the minimum singular value of H_0

is greater than $\sqrt{d} - \sqrt{n}$. In practical applications, the batchsize used for normalization is relatively smaller than the network width, hence $\mathcal{A}_1(\Theta(1), 0)$ holds. For the intermediate representations, we also observed that $\mathcal{A}_1(\alpha, \ell)$ holds –for an α independent from ℓ – as long as $d \gg n$. Let α_0 denote the minimum singular value of H_0 . For different values of d and n , we check whether $\mathcal{A}_1(0.1\alpha_0, 1000)$ holds. Fig. 3 illustrates that this assumption hold for $d = \Omega(n^2)$, which is also confirmed by Daneshmand et al. (2020).




Figure 3: **Validations for \mathcal{A}_1** Pixel (n, d) marks whether \mathcal{A}_1 holds in all 10 independent runs: The black color indicates $\mathcal{A}_1(0.1\alpha_0, 1000)$ failed in at least once (where α_0 is the minimum singular value of H_0). The blue curve marks $d = (n - 28)^{1.8}$ highlighting \mathcal{A}_1 holds for $d = \Omega(n^2)$.

4 Gaussian approximation

4.1 Orthogonality yields Gaussian approximation

The established result on the orthogonality gap allows us to show that the representations after linear layers are approximately Gaussian. If H_ℓ is orthogonal, the random matrix $W_\ell H_\ell$ is equal in distribution to a standard Gaussian matrix due to the invariance of standard Gaussian distribution under linear orthogonal transformations. We can formalize this notion of Gaussianity by bounding the Wasserstein-2 distance between the distribution of $W_\ell H_\ell$ and standard Gaussian distribution. Let $\mathcal{W}_2(R_1, R_2)$ denote the Wasserstein-2 distance between probability distributions of two random variables R_1 , and R_2 . The next lemma formally establishes the link between the orthogonality and the distribution of the representations.

Lemma 2. *Given $G \in \mathbb{R}^{d \times n}$ with i.i.d. zero-mean $1/d$ -variance Gaussian elements, the following Gaussian approximation holds:*

$$\mathcal{W}_2(W_\ell H_\ell, G/\sqrt{n})^2 \leq 2n\mathbb{E}[V(H_\ell)]. \quad (4)$$

Combining the above result with Theorem 1 yields the result presented in the next corollary.

Corollary 3. *For $G \in \mathbb{R}^{d \times n}$ with i.i.d. zero-mean $1/d$ -variance Gaussian elements,*

$$\mathcal{W}_2(W_\ell H_\ell, G/\sqrt{n})^2 \leq 4n \left(1 - \frac{2}{3}\alpha\right)^\ell + \frac{6n^2}{\alpha\sqrt{d}} \quad (5)$$

holds under Assumption $\mathcal{A}_1(\alpha, \ell)$.

In other words, the distribution of the representations contracts to a Wasserstein 2 ball around an isotropic Gaussian distribution as the depth grows. The radius of the Wasserstein 2 ball is at most $\mathcal{O}(1/\sqrt{\text{width}})$. As noted in the last section, \mathcal{A}_1 is extensively studied by Daneshmand et al. (2020) where it is shown that $\mathcal{A}_1(\alpha > 0, \infty)$ holds as long as $d = \Omega(n^2)$.

4.2 Deep neural networks as Gaussian processes

Leveraging BN, Corollary 3 establishes the first *non-asymptotic* Gaussian approximation for deep random neural networks. For vanilla networks, the Gaussianity is guaranteed only in the *asymptotic* regime of infinite width (Garriga-Alonso et al., 2019; Neal, 1996; Lee et al., 2019; Neal, 1996;

Hazan and Jaakkola, 2015). Particularly, Matthews et al. (2018) links vanilla networks to Gaussian processes when their width is infinite and grows in successive layers, while our Gaussian approximation holds for networks with finite width across layers. Table 1 briefly compares Gaussian approximation for vanilla and BN networks.

Network	Width	Depth	Distribution of Outputs
Vanilla MLP	infinite	finite	Converges to Gaussian as width $\rightarrow \infty$
Vanilla Convnet	infinite	finite	Converges to Gaussian as width $\rightarrow \infty$
BN MLP (Cor. 3)	(in)finite	(in)finite	In a $\mathcal{O}(\text{width}^{-1/4})$ - \mathcal{W}_2 ball around Gaussian

Table 1: **Distribution of representations in random vanilla and BN networks.** For the convolutional network, the width refers to the number of channels. Results for Vanilla MLPs and Vanilla convolution networks are established by Matthews et al. (2018), and Garriga-Alonso et al. (2019), respectively. Remarkably, Corollary 3 holds for MLP with linear activations.

The link between Gaussian processes and infinite-width neural networks has inspired several studies to rely on Gaussian representations in deep random networks (Klambauer et al., 2017; De Palma et al., 2019; Pennington et al., 2018; Schoenholz et al., 2017; Yang et al., 2019). Assuming the representations are Gaussian, Klambauer et al. (2017) designed novel activation functions that improve the optimization performance, De Palma et al. (2019) studies the sensitivity of random networks, Pennington et al. (2018) highlights the spectral universality in random networks, Schoenholz et al. (2017) studies information propagating through the network layers, and Yang et al. (2019) studies gradients propagation through the depth. Indeed, our analysis implies that including BN imposes the Gaussian representations required for these analyses. Although our result is established for linear activations, we conjecture a similar result holds for non-linear MLPs (see Appendix G).

5 The orthogonality and optimization

In the preceding sections, we elaborated on the theoretical properties of BN networks in controlled settings. In this section, we demonstrate the practical applications of our findings. In the first part, we focus on the relationship between depth and orthogonality. Increasing depth drastically slows the training of neural networks with BN. Furthermore, we observe that as depth grows, the training slowdown highly correlates with the orthogonality gap. This observation suggests that SGD needs to orthogonalize deep representations in order to start classification. This intuition leads us to the following question: If orthogonalization is a prerequisite for training, can we save optimization time by starting from orthogonal representations? To test this experimentally, we devised a weight initialization that guarantees orthogonality of representations. Surprisingly, even in a network without BN, our experiments showed that this initialization avoids the training slow down, affirmatively answering the question.

Throughout the experiments, we use vanilla MLP (without BN) with a width of 800 across all hidden layers, ReLU activation, and used Xavier’s method for weights initialization (Glorot and Bengio, 2010). We use SGD with stepsize 0.01 and batch size 500 and for training. The learning task is classification with cross entropy loss for CIFAR10 dataset (Krizhevsky et al., 2009, MIT license). We use PyTorch (Paszke et al., 2019, BSD license) and Google Colaboratory platform with a single Tesla-P100 GPU with 16GB memory in all the experiments. The reported orthogonality gap is the average of the orthogonality gap of representation in the last layer.

5.1 Orthogonality correlates with optimization performance

In the first experiment, we show that the orthogonality of representations at the initialization correlates with optimization speed. For networks with 15, 30, 45, 60, and 75 widths, we register training loss after 30 epochs and compare it with the initial orthogonality gap. Figure 4a shows the training loss (blue) and the initial orthogonality gap (red) as a function of depth. We observe that representations are more entangled, i.e., orthogonal, when we increase depth, coinciding with the training slowdown. Intuitively, the slowdown is due to the additional time SGD must spend to orthogonalize the representations before classification. In the second experiment, we validate this intuitive argument by tracking the orthogonality gap during training. Figure 4b plots the orthogonality gap of

output and training loss for a network with 20 layers. We observe that SGD updates are iteratively orthogonalizing representations, marked by the reduction in the orthogonality gap.




Figure 4: **Orthogonality and Optimization** Left: the orthogonality gap at initialization (red, left axis) and the training loss after 30 epochs (blue, right axis) with depth. Right: the orthogonality gap (red, left axis) and the training loss in each epoch (blue, right axis). Mean and 95% confidence interval of 4 independent runs.

5.2 Learning with initial orthogonal representations

We have seen that the slowdown in SGD for deeper networks correlates with the orthogonality gap before training. Here we show that by preemptively orthogonalizing representations, we avoid the slowdown with depth. While in MPL with linear activations, hidden representations remain orthogonal simply by taking orthogonal matrices as weights (Pennington et al., 2018; Saxe et al., 2014), the same does not hold for networks with non-linear activations, such as ReLU. To enforce the orthogonality in the absence of BN, we introduce a dependency between weights of successive layers that ensures deep representations remain orthogonal. More specifically, we incorporate the SVD decomposition of the hidden representation of each layer into the initialization of the subsequent layer. To emphasize this dependency between layers and to distinguish it from purely orthogonal weight initialization, we refer to this as *iterative orthogonalization*.

We take a large batch of samples $n \geq d$, as the input batch for initialization. Let us assume that weights are initialized up to layer $W_{\ell-1}$. To initialize W_ℓ , we compute SVD decomposition of the representations $H_\ell = U_\ell \Sigma_\ell V_\ell^\top$ where matrices $U_\ell \in \mathbb{R}^{d \times d}$ and $V_\ell \in \mathbb{R}^{n \times d}$ are orthogonal. Given this decomposition, we initialize W_ℓ by

$$W_\ell = \frac{1}{\|\Sigma_\ell^{1/2}\|_F} V'_\ell \Sigma_\ell^{-1/2} U_\ell^\top, \quad (6)$$

where $V'_\ell \in \mathbb{R}^{d \times d}$ is an orthogonal matrix obtained by slicing $V_\ell \in \mathbb{R}^{n \times d}$. Notably, the inverse in the above formula exists when n is sufficiently larger than d^2 . It is easy to check that $V(W_\ell H_\ell) < V(H_\ell)$ holds for the above initialization (see Appendix F), similar to BN. By enforcing the orthogonality, this initialization significantly alleviates the slow down of training with depth (see Fig. 5), with no need for BN. This initialization is not limited to MLPs. In Appendix H, we compare iterative orthogonalization with a BN-replacement method recently proposed by Brock et al. (2021). In Appendix I, we propose a similar SVD-based initialization for convolutional networks that effectively accelerates the training of deep convolutional networks.

6 Discussion

To recap, we proved the recurrence of random linear transformations and BN orthogonalizes samples. Our experiments underline practical applications of this theoretical finding: starting from orthogonal representations effectively avoids the training slowdown with depth for MLPs. Based on our experimental observations in Appendix I, we believe that this result extends to standard convolution networks used in practice. In other words, a proper initialization ensuring the orthogonality

²We may inductively assume that H_ℓ is almost orthogonal by the choice of $W_1, \dots, W_{\ell-1}$. Thus, Σ_ℓ is invertible.




Figure 5: **Iterative orthogonalization.** Horizontal axis: depth. Vertical axis: the training loss after 30 epochs for Xavier’s initialization (blue), our initialization (red). Mean and 95% confidence interval of 4 independent runs.

of hidden representations may replace BN in neural architectures. This future research direction has the potentials to boost the training of deep neural networks and change benchmarks in deep learning.

Although our theoretical bounds hold for MLPs with linear activations, our experiments confirm similar orthogonal stability for various neural architectures. Appendix Figure 6 demonstrates this stability for MLPs with ReLU activations and a convolutional networks. This plot compares the evolution of the orthogonality gap, through layers, for BN networks with vanilla networks. For vanilla networks, we observe that the gap increases with depth. On the contrary, the gap decreases by adding BN layers and stabilizes at a term that is constant with regard to depth. Based on our observations, we conjecture that hidden representations of modern BN networks obey similar stability. A more formal statement of the conjecture is presented in Appendix G.

Lubana et al. (2021) experimentally compares properties of different normalization techniques including layer normalization (LN) (Ba et al., 2016). According to this study, LN does not necessarily orthogonalize the outputs of deep neural networks. Hence, more theoretical studies are required to understand the essence of different normalization techniques in deep learning.

Acknowledgements

We thank Gideon Dresdner, Vincent Fortuin, and Ragnar Groot Koerkamp for their helpful comments and discussions. This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001(PRAIRIE 3IA Institute). We also acknowledge support from the European Research Council (grant SEQUOIA 724063).

References

- Agarwal, N., Awasthi, P., and Kale, S. (2021). A deep conditioning treatment of neural networks. In *Algorithmic Learning Theory*. PMLR.
- Arora, S., Cohen, N., Golowich, N., and Hu, W. (2019a). A convergence analysis of gradient descent for deep linear neural networks. *International Conference on Learning Representations*.
- Arora, S., Li, Z., and Lyu, K. (2019b). Theoretical analysis of auto rate-tuning by batch normalization. *International Conference on Learning Representations*.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. (2016). Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bahri, Y., Kadmon, J., Pennington, J., Schoenholz, S. S., Sohl-Dickstein, J., and Ganguli, S. (2020). Statistical mechanics of deep learning. *Annual Review of Condensed Matter Physics*.

- Bartlett, P. L., Helmbold, D. P., and Long, P. M. (2019). Gradient descent with identity initialization efficiently learns positive-definite linear transformations by deep residual networks. *Neural computation*.
- Bietti, A. and Mairal, J. (2019). On the inductive bias of neural tangent kernels. *Advances in Neural Information Processing Systems*.
- Bjorck, J., Gomes, C., Selman, B., and Weinberger, K. Q. (2018). Understanding batch normalization. *Advances in Neural Information Processing Systems*.
- Bougerol, P. et al. (2012). *Products of random matrices with applications to Schrödinger operators*. Springer Science & Business Media.
- Brock, A., De, S., Smith, S. L., and Simonyan, K. (2021). High-performance large-scale image recognition without normalization. *arXiv preprint arXiv:2102.06171*.
- Daneshmand, H., Kohler, J., Bach, F., Hofmann, T., and Lucchi, A. (2020). Batch normalization provably avoids rank collapse for randomly initialised deep networks. *Advances in Neural Information Processing Systems*.
- De Palma, G., Kiani, B. T., and Lloyd, S. (2019). Random deep neural networks are biased towards simple functions. *Advances in Neural Information Processing Systems*.
- Eberle, A. (2009). Markov processes. *Lecture Notes at University of Bonn*.
- Frankle, J., Schwab, D. J., and Morcos, A. S. (2021). Training batchnorm and only batchnorm: On the expressive power of random features in cnns. *ICLR*.
- Garriga-Alonso, A., Rasmussen, C. E., and Aitchison, L. (2019). Deep convolutional networks as shallow gaussian processes. *International Conference on Learning Representations*.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*.
- Hairer, M. (2010). Convergence of markov processes. *Lecture notes*.
- Hazan, T. and Jaakkola, T. (2015). Steps toward deep kernel methods from infinite neural networks. *arXiv preprint arXiv:1508.05133*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. (2017). Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Huang, P.-S., Avron, H., Sainath, T. N., Sindhvani, V., and Ramabhadran, B. (2014). Kernel methods match deep neural networks on timit. In *ICASSP*.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*.
- Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*.
- Karakida, R., Akaho, S., and Amari, S.-i. (2019). The normalization method for alleviating pathological sharpness in wide neural networks. In *Advances in Neural Information Processing Systems*.
- Kemeny, J. G. and Snell, J. L. (1976). *Markov chains*. Springer-Verlag, New York.
- Khasminskii, R. (2011). *Stochastic stability of differential equations*, volume 66. Springer Science & Business Media.
- Klambauer, G., Unterthiner, T., Mayr, A., and Hochreiter, S. (2017). Self-normalizing neural networks. *Advances in Neural Information Processing Systems*.

- Kohler, J., Daneshmand, H., Lucchi, A., Zhou, M., Neymeyr, K., and Hofmann, T. (2018). Exponential convergence rates for batch normalization: The power of length-direction decoupling in non-convex optimization. *arXiv preprint arXiv:1805.10694*.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- Kushner, H. and Yin, G. G. (2003). *Stochastic approximation and recursive algorithms and applications*. Springer Science & Business Media.
- Kushner, H. J. (1967). Stochastic stability and control. Technical report.
- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. (2019). Deep neural networks as gaussian processes. *International Conference on Learning Representations*.
- Lubana, E. S., Dick, R. P., and Tanaka, H. (2021). Beyond batchnorm: Towards a general understanding of normalization in deep learning. *NeurIPS*.
- Matthews, A. G. d. G., Rowland, M., Hron, J., Turner, R. E., and Ghahramani, Z. (2018). Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*.
- Neal, R. M. (1996). *Bayesian learning for neural networks*. Springer Science & Business Media.
- Newman, C. M. (1986). The distribution of lyapunov exponents: Exact results for random matrices. *Communications in mathematical physics*.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*.
- Pennington, J., Schoenholz, S., and Ganguli, S. (2018). The emergence of spectral universality in deep networks. In *International Conference on Artificial Intelligence and Statistics*.
- Santurkar, S., Tsipras, D., Ilyas, A., and Madry, A. (2018). How does batch normalization help optimization?(no, it is not about internal covariate shift). *Advances in Neural Information Processing Systems*.
- Sawa, T. (1972). Finite-sample properties of the k-class estimators. *Econometrica: Journal of the Econometric Society*, pages 653–680.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. (2014). Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *International Conference on Learning Representations*.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. (2017). Deep information propagation. *International Conference on Learning Representations*.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *nature*.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*.
- Yang, G., Pennington, J., Rao, V., Sohl-Dickstein, J., and Schoenholz, S. S. (2019). A mean field theory of batch normalization. *International Conference on Learning Representations*.

Appendix

A Preliminaries

A.1 Notations

Let $v, w \in \mathbb{R}^k$ then $v \odot w \in \mathbb{R}^k$ with coordinates

$$[v \odot w]_i = v_i w_i \quad (7)$$

Furthermore $v^{\otimes 2} \in \mathbb{R}^{k \times k}$ with entities

$$[v^{\otimes 2}]_{ij} = v_i v_j. \quad (8)$$

In Table 2, we summarize notations introduced previously.

Notation	Type	Definition
ℓ	integer	number of layers
n	integer	batch size
d	integer	width of the network
k	integer	output dimension
X	$\mathbb{R}^{d \times n}$	input matrix
H_ℓ	$\mathbb{R}^{d \times n}$	hidden representations at ℓ (obeying Eq. (4))
BN	$\mathbb{R}^{d \times n} \rightarrow \mathbb{R}^{d \times n}$	batch normalization layer (defined in Eq.(4))
$\text{Law}(X)$		the law of random matrix X
$\sigma_i(M)$	$\mathbb{R}^{k_1 \times k_2} \rightarrow \mathbb{R}_+$	the i th largest singular value of matrix M
I_k	$\mathbb{R}^{k \times k}$	Identity matrix of size k
$\mathbf{1}_k$	\mathbb{R}^k	all-ones vector

Table 2: Notations

A.2 The Markov chain of hidden representations

Recall the chain of the hidden representations, denoted by $\{H_\ell \in \mathbb{R}^{d \times n}\}$, obeying the following recurrence:

$$H_{\ell+1} = \frac{1}{\sqrt{d}} BN(W_\ell H_\ell), \quad BN(M) = (\text{diag}(MM^\top))^{-1/2} M, \quad (9)$$

where $W_\ell \in \mathbb{R}^{d \times n}$ are random weight matrices with i.i.d. zero-mean Gaussian elements. It is easy to check that the Frobenius norm of H_ℓ is one due to the row-wise normalization:

$$\begin{aligned} \text{Tr}(BN(H)BN(H)^\top) &= \text{Tr}\left(\text{diag}(HH^\top)^{-1/2} HH^\top \text{diag}(HH^\top)^{-1/2}\right) \\ &= \text{Tr}\left(\text{diag}(HH^\top)^{-1} HH^\top\right) \\ &= d. \end{aligned} \quad (10)$$

A.3 Lyapunov function characterizing the orthogonality

We introduce the following function serving as a Lyapunov function $\widehat{V} : \mathbb{R}^{d \times n}$ characterizing the orthogonality of the hidden representations:

$$\widehat{V}(H) = \frac{1}{n} - (\sigma_n(H))^2, \quad (11)$$

where $\sigma_n(H)$ is the minimum singular value of matrix H . Next lemma proves $\widehat{V}(H_\ell)$ bounds the orthogonality gap.

Lemma 4. *For all hidden representations H_ℓ , the following holds:*

$$V(H_\ell) \leq 2n\widehat{V}(H_\ell).$$

Proof. Let $\sigma_1, \dots, \sigma_n$ be the singular values of H_ℓ . Given these singular values, one can compute $V(H_\ell)$ as

$$(V(H_\ell))^2 = \sum_{i=1}^n \left(\sigma_i^2 - \frac{1}{n} \right)^2.$$

According to Eq. (10), $\sum_{i=1}^2 \sigma_i^2 = 1$ holds. The proof is an immediate consequence of this property.

$$\begin{aligned} V^2(H_\ell) &= \sum_{i=1}^n \sigma_i^4 - 2 \underbrace{\left(\sum_{i=1}^n \sigma_i^2 \right)}_{=1} \frac{1}{n} + \frac{1}{n} \\ &= \sum_{i=1}^n \sigma_i^4 - \frac{1}{n}. \end{aligned}$$

For a fixed σ_n , the maximum of $\sum_{i=1}^n \sigma_i^4 - \frac{1}{n}$ subject to $\sum_{i=1}^n \sigma_i^2 = 1$ is met when $\sigma_1^2 = 1 - (n-1)\sigma_n^2$ and $\sigma_2 = \dots = \sigma_n$. Using this optimal values, we get the following bound:

$$V^2(H_\ell) \leq 2(n-1)^2 \underbrace{\left(\frac{1}{n} - \sigma_n^2 \right)^2}_{=\widehat{V}^2(H_\ell(X))}.$$

Taking the square root of both sides concludes the proof. \square

B Proof of Theorem 1

The proof of Theorem 1 relies on the following Theorem that characterizes the change of \widehat{V} in consecutive layers.

Theorem 5. *The sequence $\{H_\ell\}$ obeys*

$$\mathbb{E} \left[\widehat{V}(H_{\ell+1}) | H_\ell \right] \leq \left(1 - \frac{2}{3} \left(\frac{1}{n} - \widehat{V}(H_\ell) \right) \right) \widehat{V}(H_\ell) + \frac{1}{\sqrt{d}}.$$

Notably, the above result does not rely on Assumption \mathcal{A}_1 . Assuming that Assumption \mathcal{A}_1 holds, we complete the proof of Theorem 1. Combining the last Theorem by this assumption, we get

$$\mathbb{E} \left[\widehat{V}(H_{\ell+1}) \right] \leq \left(1 - \frac{2}{3} \alpha \right) \mathbb{E} \left[\widehat{V}(H_\ell) \right] + \frac{1}{\sqrt{d}} \quad (12)$$

Induction over ℓ yields

$$\begin{aligned} \mathbb{E} \left[\widehat{V}(H_{\ell+1}(X)) \right] &\leq \left(1 - \frac{2}{3} \alpha \right)^\ell \mathbb{E} \left[\widehat{V}(H_1) \right] + \left(\sum_{k=1}^{\ell} \left(1 - \frac{2}{3} \alpha \right)^k \right) \frac{1}{\sqrt{d}} \\ &\leq \left(1 - \frac{2}{3} \alpha \right)^\ell \mathbb{E} \left[\widehat{V}(H_1) \right] + \frac{3}{2\alpha\sqrt{d}} \end{aligned}$$

An application of Lemma 4 completes the proof:

$$\begin{aligned} \mathbb{E} [V(H_{\ell+1})] &\leq 2n \mathbb{E} \left[\widehat{V}(H_{\ell+1}) \right] \\ &\leq 2 \left(1 - \frac{2}{3} \alpha \right)^\ell + \frac{3n}{2\alpha\sqrt{d}} \end{aligned}$$

B.1 Stability analysis without \mathcal{A}_1

Using Theorem 5, we can prove stability of the chain $\{H_\ell\}$ without Assumption \mathcal{A}_1 . After rearrangement of terms in Theorem 5, we get

$$\mathbb{E} \left[\widehat{V}(H_{\ell+1}) | H_\ell \right] - \widehat{V}(H_\ell) \leq -\frac{2}{3} \widehat{V}(H_\ell) \left(\frac{1}{n} - \widehat{V}(H_\ell) \right) + \frac{1}{\sqrt{d}}$$

Taking the expectation over H_ℓ and average over ℓ yields

$$\mathbb{E} \left[\frac{1}{\ell} \sum_{k=1}^{\ell} \widehat{V}(H_k) \left(\frac{1}{n} - \widehat{V}(H_k) \right) \right] \leq \left(\frac{3\mathbb{E}[\widehat{V}(H_0)]}{2\ell} \right) + \frac{3}{2\sqrt{d}}$$

C Proof of Theorem 5

C.1 Spectral decomposition.

Consider the SVD decomposition of H_ℓ as $H_\ell = U \text{diag}(\sigma) V^\top$ where U and V are orthogonal matrices. Given this decomposition, we get

$$W_\ell H_\ell = \underbrace{W_\ell U}_W \text{diag}(\sigma) V^\top \quad (13)$$

Since W_ℓ is Gaussian and U is an orthogonal matrix, entities of W are also i.i.d. standard normal. We will repeatedly use the above decomposition in our analyses.

C.2 Concentration

Consider matrix $C_{\ell+1} := H_{\ell+1}^\top H_{\ell+1}$ whose eigenvalues are $\sigma_1^2, \dots, \sigma_n^2$. The SVD decomposition of H_ℓ in Eq. (13) allows us to write $C_{\ell+1}$ as

$$C_{\ell+1} = \frac{1}{d} \sum_{i=1}^d \left(\frac{w_i \odot \sigma}{\|w_i \odot \sigma\|_2} \right)^{\otimes 2}$$

where $w_i \in \mathbb{R}^n$ is the i -th row of W , and $\sigma \in \mathbb{R}^n$ is the vector of singular values of H_ℓ . Thus, conditioned on σ , $C_{\ell+1}$ is an empirical average of i.i.d. random vectors. This allows us to prove that this empirical average is concentrated around its expected value. The next lemma states this concentration.

Lemma 6. *The following concentration always holds*

$$\mathbb{E}_{W_\ell} \|C_{\ell+1} - \mathbb{E}_{W_\ell} [C_{\ell+1}]\|^2 \leq 1/d$$

where

$$\mathbb{E}_{W_\ell} [C_{\ell+1}] = \text{diag}(p_1(\sigma), \dots, p_n(\sigma)), \quad p_i(\sigma) := \mathbb{E} \left[\frac{\sigma_n^2 w_n^2}{\sum_{k=1}^n \sigma_k^2 w_k^2} \right], \quad w_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1) \quad (14)$$

The concentration of $C_{\ell+1}$ allows us to prove that the Lyapunov function $\widehat{V}(H_{\ell+1})$ is concentrated around $1/n - p_n(\sigma)$.

Lemma 7. *The following holds*

$$\mathbb{E}_{W_\ell} \left[\left(\widehat{V}(H_{\ell+1}) - (1/n - p_n(\sigma)) \right)^2 \right] \leq \frac{1}{d}$$

The last lemma allows us to predict the value of random variable $\widehat{V}(H_{\ell+1})$ by deterministic term $1/n - p_n(\sigma)$.

C.3 Contraction.

The decay in $\widehat{V}(H_{\ell+1})$ with ℓ is due to term $1/n - p_n(\sigma)$ in the last lemma. This term is less than (or equal to) $V(H_\ell)$.

Lemma 8. *For $p_n(\sigma)$ defined in Eq. (14), the following holds:*

$$\left(\frac{1}{n} - p_n(\sigma) \right) \leq \left(1 - \frac{2}{3} \left(\frac{1}{n} - \widehat{V}(H_\ell) \right) \right) \widehat{V}(H_\ell).$$

Combining the last lemma by Lemma 7 concludes the proof of Theorem 5:

$$\mathbb{E}_{W_\ell} [\widehat{V}(H_{\ell+1})] \leq \left(1 - \frac{2}{3} \left(\frac{1}{n} - \widehat{V}(H_\ell) \right) \right) \widehat{V}(H_\ell) + \frac{1}{\sqrt{d}}. \quad (15)$$

To complete the proof, we prove Lemmas 6, 7, and 8.

C.4 Proof of Lemma 6

Given the spectral decomposition of H_ℓ in Eq. (13), we compute element ij of $C_{\ell+1}$, which is denoted by $[C_{\ell+1}]_{ij}$:

$$[C_{\ell+1}]_{ij} = [A^\top A]_{ij} = \frac{1}{d} \sum_{k=1}^d A_{ki} A_{kj}, \quad A_{ki} = W_{ki} \sigma_i / \sqrt{v_k}$$

where $v_k = \sum_{m=1}^n W_{km}^2 \sigma_m^2$. Since W_{km} are zero mean and unit variance, we get

$$\begin{aligned} \mathbb{E}[C_{\ell+1}]_{ij} &= 0 \\ \mathbb{E}[C_{\ell+1}]_{ij}^2 &= \frac{1}{d^2} \sum_{k=1}^d A_{ki}^2 A_{kj}^2 + \frac{1}{d^2} \sum_{k,k'}^d \underbrace{\mathbb{E}[A_{ki}^2 A_{kj}^2 A_{k'i}^2 A_{k'j}^2]}_{=0} \\ &= \frac{1}{d^2} \sum_{k=1}^d A_{ki}^2 A_{kj}^2 \end{aligned}$$

holds for $i \neq j$. By summing up over $i \neq j$, we get

$$\sum_{i \neq j} \mathbb{E}[C_{\ell+1}]_{ij} = \frac{1}{d^2} \left(\sum_k \left(\sum_i A_{ki}^2 \right)^2 - \sum_{ik} A_{ki}^4 \right)$$

For the diagonal elements, we get

$$\begin{aligned} \mathbb{E}[C_{\ell+1}]_{ii} &= \frac{1}{d} \sum_{k=1}^d \mathbb{E} A_{ki}^2 \\ &= \frac{1}{d} \sum_{k=1}^d \mathbb{E} A_{ki}^2 \\ &= \frac{1}{d} \sum_{k=1}^d \underbrace{\mathbb{E} \left[\frac{W_{ki}^2 \sigma_i^2}{\sum_{j=1}^n W_{kj}^2 \sigma_j^2} \right]}_{p_i(\sigma)} \end{aligned}$$

The variance of $[C_{\ell+1}]_{ii}$ is bounded as

$$\begin{aligned} \text{var}([C_{\ell+1}]_{ii}) &= \mathbb{E} \left(\frac{1}{d} \sum_{k=1}^d (A_{ki}^2 - p_i(\sigma)) \right)^2 \\ &= \frac{1}{d^2} \sum_{k=1}^d (A_{ki}^2 - p_i(\sigma))^2 \\ &\leq \frac{1}{d^2} \sum_{k=1}^d A_{ki}^4 \end{aligned}$$

Combining results for the diagonal and off-diagonal elements yields

$$\begin{aligned} \mathbb{E} \|C_{\ell+1} - \mathbb{E}_{W_\ell} [C_{\ell+1}]\|_F^2 &= \sum_{ij} \text{var}([C_{\ell+1}]_{ij}) \\ &\leq \frac{1}{d^2} \left(\sum_i A_{ki}^2 \right)^2 = 1/d \end{aligned}$$

C.5 Proof of Lemma 7

Notably, the eigenvalues of $C_{\ell+1}$ are squared singular values of $H_{\ell+1}$. Let $\lambda_n(C)$ denote the n th largest eigenvalue of matrix C .

$$\begin{aligned} \mathbb{E}_{W_\ell} \left[\left(\widehat{V}(C_{\ell+1}) - \left(\frac{1}{n} - p_n(\sigma) \right) \right)^2 \right] &\leq \mathbb{E}_{W_\ell} \left[(\lambda_n(C_{\ell+1}) - \lambda_n(\mathbb{E}_{W_\ell}[C_{\ell+1}]))^2 \right] \\ &= \mathbb{E}_{W_\ell} [\|C_{\ell+1} - \mathbb{E}_{W_\ell}[C_{\ell+1}]\|_F^2] \\ &\leq \frac{1}{d} \end{aligned}$$

where the last inequality relies on Lemma 6.

C.6 Proof of Lemma 8

The proof is based an application of moment generating function that allows us to compute expectations of ratios of random variables.

Lemma 9 (Lemma 1 in (Sawa, 1972)). *Let X_1 be a random variable that is positive with probability one and X_2 be an arbitrary random variable. Suppose that there exists a joint moment generating function of X_1 and X_2 :*

$$\phi(\theta_1, \theta_2) = \mathbb{E} [\exp(\theta_1 X_1 + \theta_2 X_2)]$$

for $\theta_1 \leq \epsilon$ and $|\theta_2| < \epsilon$ where ϵ is some positive constant. Then

$$\mathbb{E} \left[\frac{X_2}{X_1} \right] = \int_{-\infty}^0 \left[\frac{\partial \phi(\theta_1, \theta_2)}{\partial \theta_2} \right]_{\theta_2=0} d\theta_1$$

To estimate $p_n(\sigma)$, we set $X_2 := \sigma_i^2 w_i^2$ and $X_1 = \sum_j \sigma_j^2 w_j^2$, which obtains

$$\begin{aligned} \phi(\theta_1, \theta_2) &= \mathbb{E} [\exp(\theta_1 X_1 + \theta_2 X_2)] \\ &= (2\pi)^{-n/2} \int_{-\infty}^{\infty} \exp((\theta_1 + \theta_2) \sigma_i^2 w_i^2 + \sum_{j \neq i} \theta_1 \sigma_j^2 w_j^2) \exp(-\sum_k w_k^2/2) dw \\ &= (2\pi)^{-n/2} \int_{-\infty}^{\infty} \exp((-0.5 + (\theta_1 + \theta_2) \sigma_i^2) w_i^2) dw_i \left(\prod_{j \neq i} \int_{-\infty}^{\infty} \exp((-0.5 + \theta_1 \sigma_j^2) w_j^2) dw_j \right) \\ &= \frac{1}{\sqrt{1 - 2(\theta_1 + \theta_2) \sigma_i^2}} \left(\prod_{j \neq i} \frac{1}{\sqrt{1 - 2\theta_1 \sigma_j^2}} \right). \end{aligned}$$

Taking derivative with respect to θ_2 yields

$$\frac{\partial \phi}{\partial \theta_2}(\theta_1, 0) = \frac{\sigma_i^2}{(1 - 2\theta_1 \sigma_i^2)^{3/2}} \left(\prod_{j \neq i} \frac{1}{\sqrt{1 - 2\theta_1 \sigma_j^2}} \right)$$

Using the result of the last lemma, we get

$$p_i(\sigma) = \int_{-\infty}^0 \frac{\sigma_i^2}{(1 - 2\theta \sigma_i^2)} \left(\prod_j \frac{1}{\sqrt{1 - 2\theta \sigma_j^2}} \right) d\theta$$

Therefore,

$$p_n(\sigma) = \sigma_n^2 f_n(\sigma), \quad f_n(\sigma) := \int_{-\infty}^0 \frac{d\theta}{(1 - 2\theta \sigma_n^2)^{3/2} \prod_{j \neq n} (1 - 2\theta \sigma_j^2)^{1/2}} \quad (16)$$

Since $\sum_{i=1}^n \sigma_i^2 = 1$ holds (see Eq. (10)), $f_n(\sigma)$ in the above formulation is minimized when the σ_j^2 s are all equal for all $j \neq n$. Let $\sigma_n^2 := 1/n - \delta$ and $\sigma_j^2 := 1/n + \delta/(n-1)$ for all $j \neq i$. This allows us to establish a lowerbound on $f_n(\sigma)$ as

$$f_n(\sigma) \geq \int_0^\infty \underbrace{\left(1 + 2\theta \left(\frac{1}{n} - \delta\right)\right)^{-\frac{3}{2}} \left(1 + 2\theta \left(\frac{1}{n} + \frac{\delta}{n-1}\right)\right)^{-\frac{n-1}{2}}}_{g(\delta):=} d\theta \quad (17)$$

Next lemma proves that $g(\delta)$ is a convex function for $\delta \in [0, 1/n]$.

Lemma 10. *The function $g(\delta)$, which is defined in Eq. (17), is a convex function on domain $\delta \in [0, 1/n]$.*

The convexity of $g(\delta)$ yields

$$g''(\delta) \geq 0 \quad \forall \delta \implies g(\delta) \geq g(0) + \delta g'(0)$$

The above bound allow us to bound $f_n(\sigma)$ as

$$\begin{aligned} f_n(\sigma) &\geq \int_0^\infty g(0)d\theta + \delta \int_0^\infty 2\theta \left(1 + \frac{2\theta}{n}\right)^{-\frac{n}{2}-2} d\theta \\ &= 1 + \frac{2\delta n}{n+2} \end{aligned}$$

Note that we use integration by parts to compute the above integrals. Recall $\delta = \frac{1}{n} - \sigma_n^2$. Combining the above inequality by Eq. (16) concludes the proof of the Lemma 8:

$$\frac{1}{n} - p_n(\sigma) \leq \left(1 - \frac{2n}{(n+2)}\sigma_n^2\right) \left(\frac{1}{n} - \sigma_n^2\right)$$

C.7 Proof of Lemma 10

We can show convexity of $g(\delta)$ by showing $g''(\delta) \geq 0$ for all $\delta \in [0, 1/n]$. To this end, we define function $h_1(\delta)$ (for the compactness of notations) as

$$h_1(\delta) := \left(1 + 2\theta \left(\frac{1}{n} - \delta\right)\right)^{-5/2} \left(1 + 2\theta \left(\frac{1}{n} + \frac{\delta}{n-1}\right)\right)^{-\frac{n-1}{2}-1}.$$

Given h_1 , the derivative of g reads as

$$\begin{aligned} g'(\delta) &= h_1(\delta) \left(-\frac{3}{2}(-2\theta) \left(1 + 2\theta \left(\frac{1}{n} + \frac{\delta}{n-1}\right)\right) - \frac{n-1}{2} \left(\frac{2\theta}{n-1}\right) \left(1 + 2\theta \left(\frac{1}{n} - \delta\right)\right)\right) \\ &= \theta h_1(\delta) \left(3 + \theta \frac{6}{n} + \theta \delta \frac{6}{n-1} - 1 - \frac{2\theta}{n} + 2\theta \delta\right) \\ &= \theta h_1(\delta) \underbrace{\left(2 + \theta \frac{4}{n} + \theta \delta \frac{2n+4}{n-1}\right)}_{h_2(\delta):=} \\ &= \theta h_1(\delta) h_2(\delta). \end{aligned}$$

One can readily check that $h_2'(\delta) \geq 0$. Hence, $h_1'(\delta) \geq 0$ ensures the convexity of $g(\delta)$. Consider the following auxiliary function

$$h_3(\delta) := \left(1 + 2\theta \left(\frac{1}{n} - \delta\right)\right)^{-7/2} \left(1 + 2\theta \left(\frac{1}{n} + \frac{\delta}{n-1}\right)\right)^{-\frac{n-1}{2}-2}.$$

Given h_3 , we compute h'_1 as

$$\begin{aligned}
h'_1(\delta) &= h_3(\delta) \left(\left(-\frac{5}{2} \right) (-2\theta) \left(1 + 2\theta \left(\frac{1}{n} + \frac{\delta}{n-1} \right) \right) + \left(-\frac{n-3}{2} \right) \frac{2\theta}{n-1} \left(1 + 2\theta \left(\frac{1}{n} - \delta \right) \right) \right) \\
&= h_3(\delta) \theta \left(5 \left(1 + 2\theta/n + \frac{2\theta\delta}{n-1} \right) - \frac{n-3}{n-1} \left(1 + 2\theta/n - 2\theta\delta \right) \right) \\
&= h_3(\delta) \theta \underbrace{\left(\frac{4n-2}{n-1} + 2\theta \frac{4n-2}{n(n-1)} + 2\theta\delta \frac{n+2}{n-1} \right)}_{h_4(\delta) :=} \\
&= \theta h_3(\delta) h_4(\delta).
\end{aligned}$$

Clearly we have $h_3(\delta), h_4(\delta) \geq 0$ for $\delta \in [0, \frac{1}{n}]$. Therefore, the proof is complete.

D Results for vanilla networks

To compare networks with and without BN, next lemma formally establishes the contraction of orthogonality gap to a large value for networks without BN.

Lemma 11. *Let $S_\ell = W_\ell \dots W_1$. Then, there exists a positive constant δ such that the following holds*

$$\lim_{\ell \rightarrow \infty} \frac{1}{2\ell} \log \left(\left| V(S_\ell) - \sqrt{\frac{n-1}{n}} \right| \right) \leq -\delta. \quad (18)$$

In other words, the gap $V(S_\ell)$ converges to $\sqrt{(n-1)/n}$ with asymptotic rate $\exp(-\delta\ell)$ for the identity inputs. While, Thm. 1 proves the gap for BN networks converges to n/\sqrt{d} with an exponential rate. For a sufficiently large d , $n/\sqrt{d} \ll \sqrt{(n-1)/n}$ holds.

Proof. Let $\sigma_1(\ell) \geq \sigma_2(\ell) \geq \dots \geq \sigma_n(\ell)$ denote singular values of matrix S_ℓ , then it is known that

$$\lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log(\sigma_i^2(\ell)) = \frac{1}{2} \left(\log(2) + \Psi \left(\frac{d-i+1}{2} \right) \right) \quad (19)$$

holds where Ψ is digamma function (Newman, 1986), which a monotonically decreasing function. Therefore,

$$\lim_{\ell \rightarrow \infty} \frac{1}{\ell} (\log(\sigma_2^2(\ell)) - \log(\sigma_1^2(\ell))) = -\delta < 0 \quad (20)$$

holds for $\delta > 0$ that can be exactly computed using function Ψ . The above inequality implies that $\sigma_1^2(\ell)$ increases (or decreases) faster than $\sigma_2^2(\ell)$ with an exponential rate. Using this result, we get the following limit for $j \neq 1$:

$$\lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log \left(\frac{\sigma_j^2(\ell)}{\sum_i \sigma_i^2(\ell)} \right) \leq \lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log \left(\frac{\sigma_2^2(\ell)}{\sigma_1^2(\ell)} \right) = -\delta \quad (21)$$

Furthermore,

$$\lim_{\ell \rightarrow \infty} \frac{1}{2\ell} \log \left| \frac{\sigma_1^2(\ell)}{\sum_i \sigma_i^2(\ell)} - 1 \right| \leq \lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log \left| \frac{n\sigma_2^2(\ell)}{\sigma_1^2(\ell)} \right| \leq -\delta + \lim_{\ell \rightarrow \infty} \log(n)/\ell \leq -\delta \quad (22)$$

Let $\sigma(\ell) = (\sigma_1^2(\ell), \dots, \sigma_n^2(\ell)) \in \mathbb{R}^n$ and $\mathbf{1}_n$ is the all one vector in \mathbb{R}^n and $e_1 = (1, 0, \dots, 0) \in \mathbb{R}^n$. Using triangular inequality, we get

$$\left| V(S_\ell) - \sqrt{(n-1)/n} \right| = \left\| \frac{\sigma(\ell)}{\|\sigma(\ell)\|_1} - \frac{1}{n} \mathbf{1}_n \right\| - \left\| e_1 - \frac{1}{n} \mathbf{1}_n \right\| \leq \left\| \frac{\sigma(\ell)}{\|\sigma(\ell)\|_1} - e_1 \right\| \quad (23)$$

Therefore, we get

$$\begin{aligned}
& \lim_{\ell \rightarrow \infty} \frac{1}{\ell} \log \left(\left| V(S_\ell) - \sqrt{\frac{n-1}{n}} \right| \right) \\
& \leq \lim_{\ell \rightarrow \infty} \left(\frac{1}{4\ell} \log \left(2n \left(\frac{\sigma_2^2(\ell)}{\sum_j \sigma_j^2(\ell)} \right) \right) + \frac{1}{4\ell} \log \left(2 \left(\frac{\sigma_1^2(\ell)}{\sum_j \sigma_j^2(\ell)} - 1 \right)^2 \right) \right) \quad (24)
\end{aligned}$$

which is bounded by $-\delta$ according to the established bounds in Eq. (21), and Eq. (22). \square

E Proof of Lemma 2

The main idea is based on a particular coupling of random matrices $W_\ell H_\ell$ and G . Consider the truncated SVD decomposition of H_ℓ as $H_\ell = U \text{diag}(\sigma) V^\top$ where $U \in \mathbb{R}^{d \times n}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal matrices. Due to the orthogonality, the law of $W_\ell U$ is the same as those of GV . By coupling $W_\ell U = GV$, we get

$$\begin{aligned}
(\mathcal{W}_2(W_\ell H_\ell, G/\sqrt{n}))^2 &= \inf_{\text{all the couplings}} \mathbb{E} \|W_\ell U \text{diag}(\sigma) V^\top - G V V^\top / \sqrt{n}\|_F^2 \\
&\leq \mathbb{E} \|GV (\text{diag}(\sigma) - I/\sqrt{n}) V^\top\|_F^2 \\
&= \mathbb{E} \text{Tr}(GV (\text{diag}(\sigma) - I/\sqrt{n}) V^\top V (\text{diag}(\sigma) - I/\sqrt{n}) V^\top G^\top) \\
&= \text{Tr}(V (\text{diag}(\sigma) - I/\sqrt{n}) (\text{diag}(\sigma) - I/\sqrt{n}) V^\top \mathbb{E}[G^\top G]) \\
&= \mathbb{E} \text{Tr}((\text{diag}(\sigma) - I/\sqrt{n}) (\text{diag}(\sigma) - I/\sqrt{n}) V^\top V) \\
&= \mathbb{E} \|(\text{diag}(\sigma) - I/\sqrt{n})\|_F^2 \\
&= \mathbb{E} \left[\sum_{i=1}^n (\sigma_i - 1/\sqrt{n})^2 \right] \\
&= \mathbb{E} \left[\sum_{i=1}^n (\sigma_i^2 - 1/n)^2 / (\sigma_i + 1/\sqrt{n})^2 \right] \\
&\leq n \mathbb{E} \left[\sum_{i=1}^n (\sigma_i^2 - 1/n)^2 \right] \\
&\leq n \mathbb{E} [V^2(H_\ell)] \\
&\leq 2n \mathbb{E} [V(H_\ell)].
\end{aligned}$$

F Orthogonality gap for the iterative initialization

Recall the proposed initialization for weights based on SVD decomposition $H_\ell = U_\ell \Sigma_\ell V_\ell^\top$.

$$W_\ell = \frac{1}{\|\Sigma_\ell^{1/2}\|_F} V_\ell' \Sigma_\ell^{-1/2} U_\ell^\top.$$

Here, we show that

$$V(H_\ell) > V(W_\ell H_\ell) \quad (25)$$

holds as long as $V(H_\ell) \neq 0$ and $\mathcal{A}_1(\alpha, \ell)$ holds. Given singular values H_ℓ , we get

$$\begin{aligned}
V(H_\ell) &= \sum_{i=1}^n \left(\sigma_i^2 - \frac{1}{n} \right)^2 \\
&= \sum_i \sigma_i^4 - 1/n,
\end{aligned} \quad (26)$$

where we used $\sum_{i=1}^n \sigma_i^2 = 1$ (see Eq. (2)). Now, we compute $V(H_\ell)$ using the singular values.

$$W_\ell H_\ell = \frac{1}{\|\Sigma_\ell^{1/2}\|_F} V_\ell' \Sigma_\ell^{1/2} V_\ell$$

Hence, the following holds

$$\begin{aligned} V(W_\ell H_\ell) &= \sum_{i=1}^n \left(\frac{\sigma_i}{\sum_j \sigma_j} - \frac{1}{n} \right)^2 \\ &= \frac{1}{(\sum_{i=1}^n \sigma_i)^2} - 1/n \end{aligned}$$

Combining with Eq. (26), we get

$$V(H_\ell) - V(W_\ell H_\ell) = \sum_{i=1}^n \sigma_i^4 - \frac{1}{(\sum_{i=1}^n \sigma_i)^2}.$$

To show that the right side of the above equation is positive, we need to prove

$$\sum_i \sigma_i^4 \left(\sum_{i=1}^n \sigma_i \right)^2 > 1 = \left(\sum_i \sigma_i^2 \right)^4$$

holds. Using Cauchy-Schwarz inequality, we get

$$\begin{aligned} \left(\sum_i \sigma_i^2 \right)^4 &= \left(\sum_i \sigma_i^{3/2} \sigma_i^{1/2} \right)^4 \\ &\leq \left(\sum_i \sigma_i \right)^2 \left(\sum_i \sigma_i^2 \sigma_i \right)^2 \\ &\leq \left(\sum_i \sigma_i \right)^2 \left(\sum_i \sigma_i^4 \right), \end{aligned}$$

where the equality in the above inequality is met only for $\sigma_i^2 = \frac{1}{n}$ (so $V(H_\ell) = 0$) under \mathcal{A}_1 .

G Activation functions and the orthogonality

We have seen hidden representations become increasingly orthogonal through layers of BN MLPs with linear activations. We observed similar results for MLP with ReLU activations and also a simple convolutional network (see Figure 6). Here, we elaborate more on the role activation functions in our analysis.




Figure 6: Orthogonality gap of layers for BatchNorm vs vanilla network with ReLU, for MLP and a basic convolutional network on CIFAR10. Batch size was set to 8, the number of hidden layers for both networks was set to 30, with 100 width for MLP and 100 channels for ConvNet in all hidden layers. For the ConvNet, the kernel size was set to 3, across all hidden layers, with zero-padding to make convolutional feature maps equal sized. The standard deviations were computed for 50 different randomly sampled batches.

For an arbitrary activation function F , the hidden representations make the following Markov chain:

$$Q_{\ell+1} = BN'(W_\ell F(Q_\ell)).$$

where we assume that BN' makes rows of its input matrix zero-mean and unit variance (the mean correction is included).

G.1 Conjecture

Suppose that $\{Q_\ell\}$ is *ergodic* (Eberle, 2009), and admits a unique invariant distribution denoted by ν . We define function $L : \mathbb{R}^{d \times n} \rightarrow \mathbb{R}_+$ as

$$L(Q_\ell) = \|Q_\ell^\top Q_\ell - \mathbb{E}_{Q \sim \nu} [Q^\top Q]\|_F. \quad (27)$$

We conjecture that there exists an integer k such that for $\ell > k$ the following holds:

$$\mathbb{E}[L(Q_\ell)] = \mathcal{O}\left(\frac{1}{\sqrt{d}}\right). \quad (28)$$

Figure 7 experimentally validates that the above result holds for standard activations such as ReLU, sigmoid, and tanh.




Figure 7: **Validating the conjecture.** Horizontal axis: $\log(d)$; vertical axis: $\log(\frac{1}{1000} \sum_{k=1}^{1000} L(Q_k))$. The difference between the plots is negligible, hence not visible.

G.2 Odd activations

We observe that Theorem 1 holds for odd activations such as sin and tanh with no need to the mean correction in BN. For these activation result of Theorem 1.

$$\mathbb{E}[V(H_{\ell+1})] = \mathcal{O}\left((1 - \alpha)^\ell + \frac{1}{\sqrt{d}}\right)$$

holds under $\mathcal{A}_1(\alpha, \ell)$. Remarkably, this is consistent with our conjecture with $V = L$ for these activations.

H Comparisons with a BN-replacement

Here, we compare iterative orthogonalization with two baselines: (i) initialization with random orthogonal weights (Saxe et al., 2014), and (ii) adaptive gradient clipping (Brock et al., 2021).

- (i). Orthogonal weights achieve orthogonal representations in deep *linear* networks. Consider a MLP whose weight are initialized by Xavier’s scheme. Let the weight matrix at layer ℓ admit the SVD decomposition $W_\ell = U_\ell \text{diag}(\sigma_\ell) V_\ell^\top$. Then, we replace the weight matrix by the orthogonal matrix $(\text{mean}(\sigma_\ell)) U_\ell V_\ell^\top$. Since ReLU networks with orthogonal weights are prone to the alignment of representations in deep layers, we observed that such an initialization does not help with the slow down of training with depth (see Fig. 8a).
- (ii). Recently, Brock et al. (2021) propose an effective replacement for BN — based on gradient clipping. Let G_ℓ denotes the gradient of training loss with respect to W_ℓ . Given a clipping parameter λ , adaptive gradient clipping adjusts the norm of G_ℓ as

$$\hat{G}_\ell = \begin{cases} \left(\lambda \frac{\|W_\ell\|_F^*}{\|G_\ell\|_F}\right) G_\ell & \frac{\|G_\ell\|_F}{\|W_\ell\|_F^*} \geq \lambda \\ G_\ell & \text{otherwise} \end{cases} \quad (29)$$

where $\|W_\ell\|_F^* = \max\{\|W_\ell\|_F, 10^{-2}\}$.

Fig. 8a, and 8b presents results for two different choice of clipping parameters $\lambda = 0.1$ and $\lambda = 1$, respectively. These plots demonstrate adaptive gradient clipping effectively alleviates the training slow down with depth. Yet, we observe iterative orthogonalization achieves a better training loss after 30 epochs. It is not known how the clipping enhance the training. While, iterative orthogonalization is inspired by orthogonalization of representations with BN layers, which is theoretically established in this paper.




Figure 8: Iterative orthogonalization vs adaptive gradient clipping. Horizontal axis: the network depth. Vertical axis: training loss for each network after 30 SGD epochs. For more details, see Sec. 5. Mean and 95% confidence interval of 4 independent runs.

I Experiments for convolutional networks

The introduced initialization scheme in Section 5 extends to convolutional networks. In convolutional networks, hidden representation are in the tensor form $H_\ell \in \mathbb{R}^{d \times m \times m \times n}$ where d is the number of filters and k is the image dimensions. Let matrix $W_\ell^{d \times k^2}$ are weights of convolutions layer ℓ with kernel size k and d filters. A 2D convolution is a matrix multiplication combined by a projection as:

$$\text{conv2d}(W_\ell, H_\ell) = \text{fold} \left(W_\ell \underbrace{\text{unfold}(H_\ell)}_{H'_\ell} \right) \quad (30)$$

where the unfolding extract batches of $k \times k$ from images in tensor H_ℓ , and folding operation combines the computed convolution for the extracted batches. We use the SVD decomposition $H'_\ell = U_\ell \Sigma_\ell V_\ell$ to initialize weights W_ℓ , exactly the same as Eq. (6):

$$W_\ell = \frac{1}{\|\Sigma_\ell^{1/2}\|_F} V'_\ell \Sigma_\ell^{-1/2} U_\ell^\top,$$

where V'_ℓ is a slice of V_ℓ . We experimentally validate the performance of this initialization for two convolution networks with ReLU activations consisting of 20 and 80 convolutions. Table 3 outlines the details of these neural architectures.

The network with 20 layers	The network with 80 layers
Conv2d(3, 64)+ MaxPool2d(2) + ReLU	Conv2d(3, 64)+ MaxPool2d(2) + ReLU
Conv2d(64, 192) + MaxPool2d(2) + ReLU	Conv2d(64, 192) + MaxPool2d(2) + ReLU
Conv2d(192, 256) + Conv2d(256, 256) + ReLU	Conv2d(192, 256) Conv2d(256, 256) + ReLU
{Conv2d(256, 256)+ReLU} \times 16	{Conv2d(256, 256)+ReLU} \times 76
MaxPool2d(2) + ReLU	MaxPool2d(2) + ReLU
Dropout(0.5) + Linear(1024,4096) + ReLU	Dropout(0.5) + Linear(1024,4096) + ReLU
Dropout(0.5) + Linear(4096,4096) + ReLU	Dropout(0.5) + Linear(4096,4096) + ReLU
Linear(4096,10)	Linear(4096,10)

Table 3: Convolutional networks used in the experiment. These neural architectures are obtained by adding layers to AlexNet (Srivastava et al., 2014). The kernel size is 3 for all the convolutional layers.

Figure 9 shows the performance of SGD (with stepsize 0.001 and batchsize 30) for the proposed initialization. This initialization slows SGD for the shallow network with 20 layers compared to




Figure 9: **Training convolutional networks.** Convergence of SGD for Xavier’s initialization (Glorot and Bengio, 2010) (in blue) and also the proposed initialization method that ensures the orthogonality of hidden representations (in red). Mean and 95% confidence interval of 4 independent runs.




Figure 10: **Training convolutional networks.** Convergence of SGD for Xavier’s initialization (Glorot and Bengio, 2010) (in blue) and also the proposed initialization method that ensures the orthogonality of hidden representations (in red). Mean of 2 independent runs.

standard Xavier’s initialization (Glorot and Bengio, 2010). However, it outperforms Xavier’s initialization when the depth is significantly large. This result substantiates the role of orthogonality in training. We repeated this experiment for deeper networks and more epochs. Figure 10 presents our results for a convolutional network with 90 layers and 80 epochs demonstrating orthogonal initialization boosts SGD performance even after many SGD epochs.

Although the convolution network used in this experiment is not a conventional network, this experimental result motivates further investigations of the role of orthogonality in standard convolutional networks.