The Impact of Neural Network Overparameterization on Gradient Confusion and Stochastic Gradient Descent

Karthik A. Sankararaman* [1 2] **Soham De*** [3] **Zheng Xu** [2] **W. Ronny Huang** [2] **Tom Goldstein** [2]

Abstract

This paper studies how neural network architecture affects the speed of training. We introduce a simple concept called *gradient confusion* to help formally analyze this. When gradient confusion is high, stochastic gradients produced by different data samples may be negatively correlated, slowing down convergence. But when gradient confusion is low, data samples interact harmoniously, and training proceeds quickly. Through theoretical and experimental results, we demonstrate how the neural network architecture affects gradient confusion, and thus the efficiency of training. Our results show that, for popular initialization techniques, increasing the *width* of neural networks leads to *lower* gradient confusion, and thus faster model training. On the other hand, increasing the *depth* of neural networks has the opposite effect. Our results indicate that alternate initialization techniques or networks using both batch normalization and skip connections help reduce the training burden of very deep networks.

1. Introduction

Stochastic gradient descent (SGD) (Robbins & Monro, 1951) and its variants with momentum have become the standard optimization routine for neural networks due to their fast convergence and good generalization properties (Wilson et al., 2017; Sutskever et al., 2013; Smith et al., 2020). Yet the convergence behavior of SGD on neural networks still eludes full theoretical understanding. Furthermore, it is not well understood how design choices on neural network architecture affect training performance. In this paper, we make progress on these open questions.

*Equal contribution [1]Facebook. [2]University of Maryland, College Park. [3]DeepMind, London.. Correspondence to: Karthik A. Sankararaman <karthikabinavs@gmail.com>, Soham De <sohamde@google.com>.

Proceedings of the 37^{th} International Conference on Machine Learning, Online, PMLR 119, 2020.

Classical stochastic optimization theory predicts that the learning rate of SGD needs to decrease over time for convergence to be guaranteed to the minimizer of a convex function (Shamir & Zhang, 2013; Bertsekas, 2011). For strongly convex functions for example, such results show that a decreasing learning rate schedule of $O(1/k)$ is required to guarantee convergence to within ϵ-accuracy of the minimizer in $O(1/\epsilon)$ iterations, where k denotes the iteration number. Such decay schemes, however, typically lead to poor performance on standard neural network problems.

Neural networks operate in a regime where the number of parameters is much larger than the number of training data. In this "over-parameterized" regime, SGD seems to converge quickly with constant learning rates. Most neural network practitioners use a constant learning rate for the majority of training (with exponential decay only towards the end of training) without seeing the method stall (Krizhevsky et al., 2012; Simonyan & Zisserman, 2014; He et al., 2016; Zagoruyko & Komodakis, 2016). With constant learning rates, theoretical guarantees show that SGD converges quickly to a neighborhood of the minimizer, but then reaches a *noise floor* beyond which it stops converging; this noise floor depends on the learning rate and the variance of the gradients (Moulines & Bach, 2011; Needell et al., 2014). Recent results show that convergence without a noise floor is possible without decaying the learning rate, provided the model is strongly convex and overfitting occurs (Schmidt & Roux, 2013; Ma et al., 2017; Vaswani et al., 2018).

While these results do give important insights, they do not fully explain the dynamics of SGD on neural networks, and how they relate to over-parameterization. Furthermore, training performance is strongly influenced by network architecture. It is common knowledge among practitioners that, under standard Gaussian initialization techniques (Glorot & Bengio, 2010; He et al., 2015), deeper networks train slower (Bengio et al., 1994; Saxe et al., 2013). This has led to several innovations over the years to get deeper nets to train more easily, such as careful initialization strategies (Xiao et al., 2018), residual connections (He et al., 2016), and normalization schemes like batch normalization (Ioffe & Szegedy, 2015). Furthermore, there is evidence to indicate that wider networks are faster to train (Zagoruyko & Komodakis, 2016; Nguyen & Hein, 2017), and recent the-

oretical results suggest that the dynamics of SGD simplify considerably for very wide networks (Jacot et al., 2018; Lee et al., 2019). In this paper, we make progress on theoretically understanding these empirical observations and unifying existing theoretical results. To this end, we identify and analyze a condition that enables us to establish direct relationships between layer width, network depth, problem dimensionality, initialization schemes, and trainability and SGD dynamics for over-parameterized networks.

Our contributions. Typical neural networks are *over-parameterized* (*i.e.,* the number of parameters exceed the number of training points). In this paper, we ask how this over-parameterization, and more specifically the network architecture, affects the trainability of neural networks and the dynamics of SGD. Through extensive theoretical and experimental studies, we show how layer width, network depth, initialization schemes, and other architecture choices affect the dynamics. The following are our main contributions.[1]

- We identify a condition, termed *gradient confusion*, that impacts the convergence properties of SGD on over-parameterized models. We prove that high gradient confusion may lead to slower convergence, while convergence is accelerated (and could be faster than predicted by existing theory) if confusion is low, indicating a regime where constant learning rates work well in practice (sections 2 and 3). We use the gradient confusion condition to study the effect of various architecture choices on trainability and convergence.
- We study the effect of neural network architecture on gradient confusion at standard Gaussian initialization schemes (section 4), and prove (a) gradient confusion increases as the network depth increases, and (b) wider networks have lower gradient confusion. These indicate that deeper networks are more difficult to train and wider networks can improve trainability of networks. Directly analyzing the gradient confusion bound enables us to derive results on the effect of depth and width, without requiring restrictive assumptions like large layer widths (Du et al., 2018; Allen-Zhu et al., 2018). Our results hold for a large class of neural networks with different non-linear activations and loss-functions. In section 5, we present a more general result on the effect of depth on the trainability of networks without assuming the network is at initialization.
- We prove that for linear neural networks, gradient confusion is *independent of depth* when using orthogonal initialization schemes (section 6) (Saxe et al., 2013; Schoenholz et al., 2016). This indicates a way forward in developing techniques for training deeper models.
- We test our theoretical predictions using extensive experiments on wide residual networks (WRNs) (Zagoruyko & Komodakis, 2016), convolutional networks (CNNs) and multi-layer perceptrons (MLPs) for image classification tasks on CIFAR-10, CIFAR-100 and MNIST (section 7 and appendix A). We find that our theoretical results consistently hold across all our experiments. We further show that the combination of batch normalization and skip connections in residual networks help lower gradient confusion, thus indicating why SGD can efficiently train deep neural networks that employ such techniques.

[1] To keep the main text of the paper concise, all proofs and several additional experimental results are delegated to the appendix.

2. Gradient confusion

Notations. We denote vectors in bold lower-case and matrices in bold upper-case. We use $(\mathbf{W})_{i,j}$ to indicate the (i,j) cell in matrix $\mathbf{W}$ and $(\mathbf{W})_i$ for the i^{th} row of matrix $\mathbf{W}$. $\|\mathbf{W}\|$ denotes the operator norm of $\mathbf{W}$. $[N]$ denotes $\{1, 2, \ldots, N\}$ and $[N]_0$ denotes $\{0, 1, \ldots, N\}$.

Preliminaries. Given N training points (specified by the corresponding loss functions $\{f_i\}_{i\in[N]}$), we use SGD to solve empirical risk minimization problems of the form,

$$\min_{\mathbf{w}\in\mathbb{R}^d} F(\mathbf{w}) := \min_{\mathbf{w}\in\mathbb{R}^d} \frac{1}{N}\sum_{i=1}^{N} f_i(\mathbf{w}), \quad (1)$$

using the following iterative update rule for T rounds:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \alpha\nabla\tilde{f}_k(\mathbf{w}_k). \quad (2)$$

Here α is the learning rate and $\tilde{f}_k$ is a function chosen uniformly at random from $\{f_i\}_{i\in[N]}$ at iteration $k \in [T]$. $\mathbf{w}^\star = \arg\min_{\mathbf{w}} F(\mathbf{w})$ denotes the optimal solution.

Gradient confusion. SGD works by iteratively selecting a random function $\tilde{f}_k$, and modifying the parameters to move in the direction of the negative gradient of $\tilde{f}_k$. It may happen that the selected gradient $\nabla\tilde{f}_k$ is negatively correlated with the gradient of another term ∇f_j. When the gradients of different mini-batches are negatively correlated, the objective terms disagree on which direction the parameters should move, and we say that there is *gradient confusion*.[2]

Definition 2.1. *A set of objective functions $\{f_i\}_{i\in[N]}$ has gradient confusion bound $\eta \geq 0$ if the pair-wise inner products between gradients satisfy, for a fixed $\mathbf{w} \in \mathbb{R}^d$,*

$$\langle\nabla f_i(\mathbf{w}), \nabla f_j(\mathbf{w})\rangle \geq -\eta, \ \forall i \neq j \in [N]. \quad (3)$$

Observations in simplified settings. SGD converges fast when gradient confusion is low along its path. To see why,

[2] Gradient confusion is related to both gradient variance and gradient diversity (Yin et al., 2017), but with important differences, which we discuss in section 9. We also discuss alternate definitions of the gradient confusion condition in section 8.

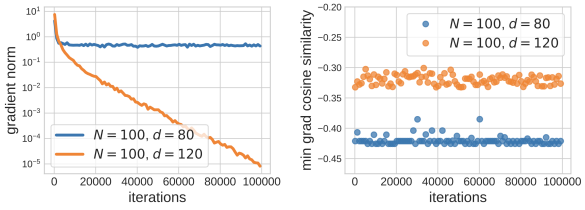

Figure 1. Linear regression on an over-parameterized ($d = 120$) and under-parameterized ($d = 80$) model with $N = 100$ samples generated randomly from a Gaussian, trained using SGD with mini-batch size 1. Plots are averaged over 3 independent runs. Gradient cosine similarities were calculated over all pairs of gradients.

consider the case of training a logistic regression model on a dataset with *orthogonal* vectors. We have $f_i(\mathbf{w}) = \mathcal{L}(y_i \mathbf{x}_i^\top \mathbf{w})$, where $\mathcal{L} : \mathbb{R} \to \mathbb{R}$ is the logistic loss, $\{\mathbf{x}_i\}_{i\in[N]}$ is a set of orthogonal training vectors, and $y_i \in \{-1, 1\}$ is the label for $\mathbf{x}_i$. We then have $\nabla f_i(\mathbf{w}) = \zeta_i \mathbf{x}_i$, where $\zeta_i = y_i \mathcal{L}'(y_i \cdot \mathbf{x}_i^\top \mathbf{w})$. Note that the gradient confusion is 0 since $\langle \nabla f_i(\mathbf{w}), \nabla f_j(\mathbf{w})\rangle = \zeta_i \zeta_j \langle \mathbf{x}_i, \mathbf{x}_j\rangle = 0, \forall i, j \in [N]$ and $i \neq j$. Thus, an update in the gradient direction f_i has *no* effect on the loss value of f_j for $i \neq j$. In this case, SGD decouples into (deterministic) gradient descent on each objective term separately, and we can expect to see the fast convergence rates attained by gradient descent.

Can we expect a problem to have low gradient confusion in practice? From the logistic regression problem, we have: $|\langle \nabla f_i(\mathbf{w}), \nabla f_j(\mathbf{w})\rangle| = |\langle \mathbf{x}_i, \mathbf{x}_j\rangle| \cdot |\zeta_i \zeta_j|$. This inner product is expected to be small for all $\mathbf{w}$; the logistic loss satisfies $|\zeta_i \zeta_j| < 1$, and for fixed N the quantity $\max_{ij} |\langle \mathbf{x}_j, \mathbf{x}_i\rangle|$ is $O(1/\sqrt{d})$ whenever $\{\mathbf{x}_i\}$ are randomly sampled from a sphere (see lemma B.1 for the formal statement).[3] Thus, we would expect a random linear model to have nearly orthogonal gradients, when the number of parameters is "large" and the number of training data is "small", i.e., when the model is over-parameterized. This is further evidenced by a toy example in figure 1, where we show a slightly over-parameterized linear regression model can have much faster convergence rates, as well as lower gradient confusion. One can prove a similar result for problems that have random and low-rank Hessians, which suggests that one might expect gradient to be small near the minimizer for many standard neural nets (see appendix C for more discussion).

The above arguments are a bit simplistic, considering toy scenarios and ignoring issues like the effect of network structure. In the following sections, we rigorously analyze the effect of gradient confusion on the speed of convergence on non-convex problems, and the effect of width and depth of the neural network architecture on the gradient confusion.

[3] Generally, this is true whenever $\mathbf{x}_i = \frac{1}{\sqrt{d}} \mathbf{y}_i$, where $\mathbf{y}_i$ is an isotropic random vector (Vershynin, 2018).

3. SGD is fast when gradient confusion is low

Several prior papers have analyzed the convergence rates of constant learning rate SGD (Nedić & Bertsekas, 2001; Moulines & Bach, 2011; Needell et al., 2014). These results show that for strongly convex and Lipschitz smooth functions, SGD with a constant learning rate α converges *linearly* to a neighborhood of the minimizer. The noise floor it converges to depends on the learning rate α and the variance of the gradients at the minimizer, i.e., $\mathbb{E}_i \|\nabla f_i(\mathbf{w}^\star)\|^2$. To guarantee convergence to ϵ-accuracy in such a setting, the learning rate needs to be small, i.e., $\alpha = O(\epsilon)$, and the method requires $T = O(1/\epsilon)$ iterations. Some more recent results show convergence of constant learning rate SGD without a noise floor and without small step sizes for models that can completely fit the data (Schmidt & Roux, 2013; Ma et al., 2017; Vaswani et al., 2018).

Gradient confusion is related to these results. Cauchy-Schwarz inequality implies that if $\mathbb{E}_i \|\nabla f_i(\mathbf{w}^\star)\|^2 = O(\epsilon)$, then $\mathbb{E}_{i,j} |\langle \nabla f_i(\mathbf{w}^\star), \nabla f_j(\mathbf{w}^\star)\rangle| = O(\epsilon)$, $\forall i, j$. Thus the gradient confusion at the minimizer is small when the variance of the gradients at the minimizer is small. Further note that when the variance of the gradients at the minimizer is $O(\epsilon)$, a direct application of the results in Moulines & Bach (2011) and Needell et al. (2014) shows that constant learning rate SGD has fast convergence to ϵ-accuracy in $T = O(\log(1/\epsilon))$ iterations, without the learning rate needing to be small. Generally however, bounded gradient confusion does not provide a bound on the variance of the gradients (see section 9). Thus, it is instructive to derive convergence bounds of SGD explicitly in terms of the gradient confusion to properly understand its effect.

We first consider functions satisfying the Polyak-Lojasiewicz (PL) inequality (Lojasiewicz, 1965), a condition related to, but weaker than, strong convexity, and used in recent work (Karimi et al., 2016; De et al., 2017). We provide bounds on the rate of convergence in terms of the optimality gap. We start with two standard assumptions.

(A1) $\{f_i\}_{i\in[N]}$ are *Lipschitz smooth*:
$f_i(\mathbf{w}') \leq f_i(\mathbf{w}) + \nabla f_i(\mathbf{w})^\top (\mathbf{w}' - \mathbf{w}) + \frac{L}{2} \|\mathbf{w}' - \mathbf{w}\|^2$.

(A2) $\{f_i\}_{i\in[N]}$ satisfy the *PL inequality*:
$\frac{1}{2} \|\nabla f_i(\mathbf{w})\|^2 \geq \mu (f_i(\mathbf{w}) - f_i^\star)$, $f_i^\star = \min_{\mathbf{w}} f_i(\mathbf{w})$.

We now state a convergence result of constant learning rate SGD in terms of the gradient confusion.

Theorem 3.1. *If the objective function satisfies (A1) and (A2), and has gradient confusion η, SGD converges linearly to a neighborhood of the minima of problem (1) as:*

$$\mathbb{E}[F(\mathbf{w}_T) - F^\star] \leq \rho^T (F(\mathbf{w}_0) - F^\star) + \frac{\alpha\eta}{1-\rho},$$

where $\alpha < \frac{2}{NL}$, $\rho = 1 - \frac{2\mu}{N}\left(\alpha - \frac{NL\alpha^2}{2}\right)$, $F^\star = \min_{\mathbf{w}} F(\mathbf{w})$ and $\mathbf{w}_0$ is the initialized weights.

This result shows that SGD converges *linearly* to a neighborhood of a minimizer, and the size of this neighborhood depends on the level of gradient confusion. When the gradient confusion is small, i.e., $\eta = O(\epsilon)$, SGD has fast convergence to $O(\epsilon)$-accuracy in $T = O(\log(1/\epsilon))$ iterations, without requiring the learning rate to be vanishingly small. We now extend this to general smooth functions.

Theorem 3.2. *If the objective satisfies (A1) and has gradient confusion η, then SGD converges to a neighborhood of a stationary point of problem (1) as:*

$$\min_{k=1,\dots,T} \mathbb{E}\|\nabla F(\mathbf{w}_k)\|^2 \leq \frac{\rho(F(\mathbf{w}_1)-F^\star)}{T} + \rho\eta,$$

for $\alpha < \frac{2}{NL}$, $\rho = \frac{2N}{2-NL\alpha}$, and $F^\star = \min_{\mathbf{w}} F(\mathbf{w})$.

Thus, as long as $\eta = O(1/T)$, SGD has fast $O(1/T)$ convergence on smooth non-convex functions. Theorems 3.1 and 3.2 predict an initial phase of optimization with fast convergence to the neighborhood of a minimizer or a stationary point. This behavior is often observed when optimizing neural nets (Darken & Moody, 1992; Sutskever et al., 2013), where a constant learning rate reaches a high level of accuracy on the model. As we show in subsequent sections, this is expected since for neural networks typically used, the gradient confusion is expected to be low. See section 9 for more discussion on the above results and how they relate to previous work. We stress that our goal is not to study convergence rates per se, nor is it to prove state-of-the-art rate bounds for this class of problems. Rather, we show the direct effect that the gradient confusion bound has on the convergence rate and the noise floor for constant learning rate SGD. As we show in the following sections, this new perspective in terms of the gradient confusion helps us more directly understand how neural network architecture design affects SGD dynamics and why.

4. Effect of neural network architecture at Gaussian initializations

To draw a connection between neural network architecture and training performance, we analyze gradient confusion for generic (i.e., random) model problems using methods from high-dimensional probability. In this section, we analyze the effect of neural network architecture at the beginning of training, when using standard Gaussian initialization techniques. Analyzing these models at initialization is important to understand which architectures are more easily trainable than others. Our results cover a wide range of scenarios compared to prior work, require minimal additional assumptions, and hold for a large family of neural networks with different non-linear activation functions and loss-functions. In particular, our results hold for fully connected networks (and can be extended to convolutional networks) with the square-loss and logistic-loss functions, and commonly used non-linear activations such as sigmoid, tanh and ReLU. We consider both the case where the input data is arbitrary but bounded (theorem 4.1, part 1), as well as where the input data is randomly drawn from the surface of a unit sphere (theorem 4.1, part 2).

Setting. We consider training data $\mathcal{D} = \{(\mathbf{x}_i, \mathcal{C}(\mathbf{x}_i))\}_{i\in[N]}$, with labeling function $\mathcal{C} : \mathbb{R}^d \to [-1, 1]$. For some of our results, we consider that the data points $\{\mathbf{x}_i\}$ are drawn uniformly at random from the surface of a d-dimensional unit sphere. The labeling function satisfies $|\mathcal{C}(\mathbf{x})| \leq 1$ and $\|\nabla_{\mathbf{x}}\mathcal{C}(\mathbf{x})\|_2 \leq 1$ for $\|\mathbf{x}\| \leq 1$. Note that this automatically holds for every model considered in this paper where the labeling function is *realizable* (i.e., where the model can express the labeling function using its parameters). More generally, this assumes a Lipschitz condition on the labels (i.e., the labels don't change too quickly with the inputs).

We consider two loss-functions: square-loss for regression and logistic loss for classification. The square-loss function is defined as $f_i(\mathbf{w}) = \frac{1}{2}(\mathcal{C}(\mathbf{x}_i) - g_{\mathbf{w}}(\mathbf{x}_i))^2$ and the logistic function is defined as $f_i(\mathbf{w}) = \log(1 + \exp(-\mathcal{C}(\mathbf{x}_i)g_{\mathbf{w}}(\mathbf{x}_i)))$. Here, $g_{\mathbf{w}} : \mathbb{R}^d \to \mathbb{R}$ denotes the parameterized function we fit to the training data and $f_i(\mathbf{w})$ denotes the loss-function of hypothesis $g_{\mathbf{w}}$ on data point $\mathbf{x}_i$.

Let $\mathbf{W}_0 \in \mathbb{R}^{\ell_1\times d}$ and $\{\mathbf{W}_p\}_{p\in[\beta]}$ where $\mathbf{W}_p \in \mathbb{R}^{\ell_p\times\ell_{p-1}}$ are weight matrices. Let $\mathbf{W}$ denote the tuple $(\mathbf{W}_p)_{p\in[\beta]_0}$. Define $\ell := \max_{p\in[\beta]} \ell_p$ to be the *width* and β to be the *depth* of the network. Then, the model $g_{\mathbf{W}}$ is defined as

$$g_{\mathbf{W}}(\mathbf{x}) := \sigma(\mathbf{W}_\beta\sigma(\mathbf{W}_{\beta-1}\dots\sigma(\mathbf{W}_1\sigma(\mathbf{W}_0\mathbf{x}))\dots)),$$

where σ denotes the non-linear activation function applied point-wise to its arguments. We assume that the activation is given by a function $\sigma(x)$ with the following properties.

- **(P1) Boundedness:** $|\sigma(x)| \leq 1$ for $x \in [-1, 1]$.
- **(P2) Bounded differentials:** Let $\sigma'(x)$ and $\sigma''(x)$ denote the first and second sub-differentials respectively. Then, $|\sigma'(x)| \leq 1$ and $|\sigma''(x)| \leq 1$ for all $x \in [-1, 1]$.

When $\|\mathbf{x}\| \leq 1$, activation functions such as *sigmoid*, *tanh*, *softmax* and *ReLU* satisfy these requirements.

Furthermore, in this section, we consider the following Gaussian weight initialization strategy.

Strategy 4.1. *$\mathbf{W}_0 \in \mathbb{R}^{\ell\times d}$ has independent $\mathcal{N}(0, \frac{1}{d})$ entries. For every $p \in [\beta]$, the weights $\mathbf{W}_p \in \mathbb{R}^{\ell_p\times\ell_{p-1}}$ have independent $\mathcal{N}\left(0, \frac{1}{\kappa\ell_{p-1}}\right)$ entries for some constant $\kappa > 0$.*

This initialization strategy with different settings of κ are used almost universally for neural networks (Glorot & Bengio, 2010; LeCun et al., 2012; He et al., 2015). For instance,

typically $\kappa = \frac{1}{2}$ when ReLU activations are used, and $\kappa = 1$ when tanh activations are used.

Main result. The following theorem shows how the width $\ell := \max_{p\in[\beta]} \ell_p$ and the depth β affect the gradient confusion condition at standard initializations. We show that *as width increases (for fixed depth) or depth decreases (for fixed width) the probability that the gradient confusion bound (equation 3) holds increases*. Thus, as the depth increases (with fixed width), training a model becomes harder, while as the width increases (with fixed depth), training a model becomes easier. Furthermore, note that this result also implies that training very deep *linear* neural networks (with identity activation functions) with standard Gaussian initializations is hard. Throughout the paper, we define the parameter $\zeta_0 := 2\sqrt{\beta}$. See the appendix (Lemma D.1) for a more careful definition of this quantity.

Theorem 4.1. *Let $\mathbf{W}_0, \mathbf{W}_1, \ldots, \mathbf{W}_\beta$ be weight matrices chosen according to strategy 4.1. There exists fixed constants $c_1, c_2 > 0$ such that we have the following.*

1. *Consider a fixed but arbitrary dataset $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N$ with $\|\mathbf{x}_i\| \leq 1$ for every $i \in [N]$. For $\eta > 4$, the gradient confusion bound in equation 3 holds with probability at least*

$$1 - \beta \exp\left(-c_1\kappa^2\ell^2\right) - N^2 \exp\left(\frac{-c\ell^2\beta(\eta-4)^2}{64\zeta_0^4(\beta+2)^4}\right).$$

2. *If the dataset $\{\mathbf{x}_i\}_{i\in[N]}$ is such that each $\mathbf{x}_i$ is an i.i.d. sample from the surface of d-dimensional unit sphere, then for every $\eta > 0$ the gradient confusion bound in equation 3 holds with probability at least*

$$1 - \beta \exp\left(-c_1\kappa^2\ell^2\right) - N^2 \exp\left(\frac{-c_2(\ell d+\ell^2\beta)\eta^2}{16\zeta_0^4(\beta+2)^4}\right).$$

Theorem 4.1 shows that under popular Gaussian initializations used, training becomes harder as networks get deeper. The result however also shows a way forward: layer width improves the trainability of deep networks. Other related work supports this showing that when the layers are infinitely wide, the learning dynamics of gradient descent simplifies considerably (Jacot et al., 2018; Lee et al., 2019). Hanin & Rolnick (2018) also suggest that the width should increase linearly with depth in a neural network to help dynamics at the beginning of training. In section 7 and appendix A, we show substantial empirical evidence that, given a sufficiently deep network, increasing the layer width often helps in lowering gradient confusion and speeding up convergence for a range of models.

5. A more general result on the effect of depth

While our results in section 4 hold at standard initialization schemes, in this section we derive a more general version of the result. In particular, we assume the setting where the data is drawn uniformly at random from a unit sphere and the weights lie in a ball around a local minimizer. Our results hold for both fully connected networks and convolutional networks with the square-loss and logistic-loss functions, and commonly-used non-linear activations such as sigmoid, tanh, softmax and ReLU.

We consider the same setup as in the previous section, and assume additionally that the data points $\{\mathbf{x}_i\}$ are drawn uniformly from the surface of a d-dimensional unit sphere. Additionally, instead of studying the network at initialization, we make the following assumption on the weights.

Assumption 1 (Small Weights)**.** *We assume that the operator norm of the weight matrices $\{\mathbf{W}_i\}_{i\in[\beta]_0}$ are bounded above by 1, i.e., for every $i \in [\beta]_0$ we have $\|\mathbf{W}_i\| \leq 1$.*

The operator norm of the weight matrices $\|\mathbf{W}\|$ being close to 1 is important for the trainability of neural networks, as it ensures that the input signal is passed through the network without exploding or shrinking across layers (Glorot & Bengio, 2010). Proving non-vacuous bounds in case of such blow-ups in magnitude of the signal or the gradient is not possible in general, and thus, we consider this restricted class of weights. Most standard neural networks are trained using *weight decay* regularizers of the form $\sum_i \|W_i\|_F^2$. This biases the weights to be small when training neural networks in practice. See appendix F for further discussion on the small weights assumption.

We now present a more general version of theorem 4.1.

Theorem 5.1. *Let $\mathbf{W}_0, \mathbf{W}_1, \ldots, \mathbf{W}_\beta$ satisfy assumption 1. For some fixed constant $c > 0$, the gradient confusion bound (equation 3) holds with probability at least*

$$1 - N^2 \exp\left(\frac{-cd\eta^2}{16\zeta_0^4(\beta+2)^4}\right).$$

Theorem 5.1 shows that (for fixed dimension d and number of samples N) when the depth β decreases, the probability that the gradient confusion bound in equation 3 holds increases, and vice versa. Thus, our results indicate that in the general case when the weights are small, increasing the network depth will typically lead to slower model training.

Note that on assuming $\|\mathbf{W}\| \leq 1$ for each weight matrix $\mathbf{W}$, the dependence of gradient confusion on the width goes away. To see why this, consider an example where each weight matrix in the neural network has exactly one non-zero element, which is set to 1. The operator norm of each such weight matrix is 1, but the forward or backward propagated signals would not depend on the width.

Note that the convergence rate results of SGD in section 3 assume that the gradient confusion bound holds at every point along the path of SGD. On the other hand, theorem

5.1 shows concentration bounds for the gradient confusion at a fixed weight $\mathbf{W}$. Thus, to make the above result more relevant for the convergence of SGD on neural networks, we now make the concentration bound in theorem 5.1 *uniform over all weights inside a ball $\mathcal{B}_r$ of radius r*.

Corollary 5.1. *Select a point $\mathbf{W} = (\mathbf{W}_0, \mathbf{W}_1, \ldots, \mathbf{W}_\beta)$, satisfying assumption 1. Consider a ball $\mathcal{B}_r$ centered at $\mathbf{W}$ of radius $r > 0$. If the data $\{\mathbf{x}_i\}_{i\in[N]}$ are sampled uniformly from a unit sphere, then the gradient confusion bound in equation 3 holds uniformly at all points $\mathbf{W}' \in \mathcal{B}_r$ with probability at least*

$$\begin{cases} 1 - N^2 \exp\left(-\frac{cd\eta^2}{64\zeta_0^4(\beta+2)^4}\right), & \text{if } r \leq \eta/4\zeta_0^2, \\ 1 - N^2 \exp\left(-\frac{cd\eta^2}{64\zeta_0^4(\beta+2)^4} + \frac{8d\zeta_0^2 r}{\eta}\right), & \text{otherwise.} \end{cases}$$

Corollary 5.1 shows that the probability that the gradient confusion bound holds decreases with increasing depth, for all weights in a ball around the minimizer.[4] This explains why, in the general case, training very deep models might always be hard. This raises the question why most deep neural networks used in practice are so efficiently trained using SGD. While careful Gaussian initialization strategies prevent vanishing or exploding gradients, these strategies still suffer from high gradient confusion for very deep networks unless the width is also increased with the depth, as we show in section 4. Practitioners over the years, however, have achieved state-of-the-art results by making networks deeper, without necessarily making networks wider. Thus, in section 7, we empirically study how popular techniques used in these models like skip connections and batch normalization affect gradient confusion. We find that these techniques drastically lower gradient confusion, making deep networks significantly easier to train. Furthermore, in the next section, we show how deep linear nets are trainable when used with orthogonal initialization techniques, indicating a way forward for training deeper models.

6. Gradient confusion is independent of depth for orthogonal initializations

In this section, we show that for deep linear neural networks, gradient confusion is independent of depth when the weight matrices are initialized as orthogonal matrices.[5] Consider the following linear neural network:

$$g_{\mathbf{W}}(\mathbf{x}) := \gamma \mathbf{W}_\beta \cdot \mathbf{W}_{\beta-1} \cdot \ldots \cdot \mathbf{W}_1 \cdot \mathbf{x}, \qquad (4)$$

where the rescaling parameter $\gamma = \frac{1}{\sqrt{2\beta}}$, and assume we use the squared loss function. Then we have the following.

Theorem 6.1. *Let $\{\mathbf{W}_i\}_{i\in[\beta]}$ be arbitrary orthogonal matrices that satisfy assumption 1. Let the dataset $\{\mathbf{x}_i\}_{i\in[N]}$ be such that each $\mathbf{x}_i$ is an i.i.d. sample from the surface of d-dimensional unit sphere. Consider the linear neural network in equation 4 that minimizes the empirical square loss function. For some fixed constant $c > 0$, the gradient confusion bound (equation 3) holds with probability at least*

$$1 - N^2 \exp\left(-cd\eta^2\right).$$

From Theorem 6.1, we see that the probability does not depend on the depth β or maximum width ℓ. Thus, trainability does not get worse with depth when using orthogonal initializations. This result matches previous theoretical and empirical results showing the efficiency of orthogonal initialization techniques for training very deep linear or tanh networks (Saxe et al., 2013; Schoenholz et al., 2016; Xiao et al., 2018). However, orthogonal initializations are not compatable with non-linear activation functions like sigmoids or ReLUs, which limit their use in practice. Nonetheless, this result suggests a promising direction in developing techniques for training deeper models.

7. Experimental results

To test our theoretical results and to probe why standard neural networks are efficiently trained with SGD, we now present experimental results showing the effect of the neural network architecture on the convergence of SGD and gradient confusion. It is worth noting that theorems 3.1 and 3.2 indicate that we would expect the effect of gradient confusion to be most prominent closer to the end of training.

We performed experiments on wide residual networks (WRNs) (Zagoruyko & Komodakis, 2016), convolutional networks (CNNs) and multi-layer perceptrons (MLPs) for image classification tasks on CIFAR-10, CIFAR-100 and MNIST. We present results for CNNs on CIFAR-10 in this section, and present all other results in appendix A. We use CNN-β-ℓ to denote WRNs that have no skip connections or batch normalization, with a depth β and width factor ℓ.[6] We turned off dropout and weight decay for all our experiments. We used SGD as the optimizer without any momentum. Following Zagoruyko & Komodakis (2016), we ran all experiments for 200 epochs with minibatches of size 128, and reduced the initial learning rate by a factor of 10 at epochs 80 and 160. We used the MSRA initializer (He et al., 2015) for the weights as is standard for this model, and used the same preprocessing steps for the CIFAR-10 images as described in Zagoruyko & Komodakis (2016). We ran each experiment 5 times, and we show the standard deviation across runs in our plots. We tuned the optimal initial

[4] The above results automatically hold for convolutional networks, since a convolution operation on $\mathbf{x}$ can be represented as a matrix multiplication $\mathbf{U}\mathbf{x}$ for an appropriate Toeplitz matrix $\mathbf{U}$.

[5] An orthogonal matrix $\mathbf{A}$ satisfies $\mathbf{A}^T \cdot \mathbf{A} = \mathbf{A} \cdot \mathbf{A}^T = \mathbf{I}$.

[6] The width factor denotes the number of filters relative to the original ResNet model (Zagoruyko & Komodakis, 2016).

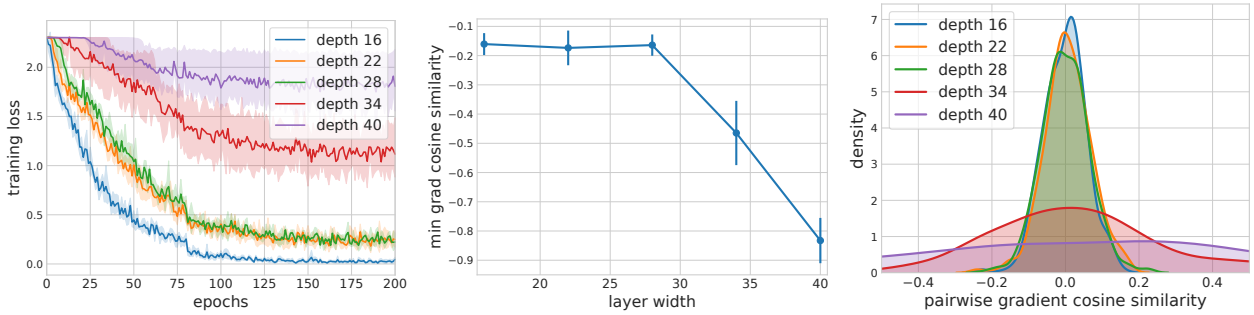

(a) (b) (c)

Figure 2. The effect of network depth with CNN-β-2 on CIFAR-10 for depths $\beta = 16$, 22, 28, 34 and 40. Plots show the (a) convergence curves for SGD, (b) minimum of pairwise gradient cosine similarities at the end of training, and the (c) kernel density estimate of the pairwise gradient cosine similarities at the end of training (over all independent runs).

learning rate for each model over a logarithmically-spaced grid and selected the run that achieved the lowest training loss value. To measure gradient confusion, at the end of every training epoch, we sampled 100 pairs of mini-batches each of size 128 (the same size as the training batch). We calculated gradients on each mini-batch, and then computed pairwise cosine similarities. See appendix A.2 for more details on the experimental setup and architectures used.

Effect of depth. To test our theoretical results, we consider CNNs with a fixed width factor of 2 and varying network depth. From figure 2, we see that our theoretical results are backed by the experiments: increasing depth slows down convergence, and increases gradient confusion. We also notice that with increasing depth, the density of pairwise gradient cosine similarities concentrates less sharply around 0, which makes the network harder to train.

Effect of width. We now consider CNNs with a fixed depth of 16 and varying width factors. From figure 3, we see that increasing width results in faster convergence and lower gradient confusion. We further see that gradient cosine similarities concentrate around 0 with growing width, indicating that SGD decouples across the training samples with growing width. Note that the smallest network considered (CNN-16-2) is still over-parameterized and achieves a high level of performance (see appendix A.3).

Effect of batch normalization and skip connections. Almost all state-of-the-art neural networks currently contain both skip connections and normalization layers. To help understand why such neural networks are so efficiently trained using SGD with constant learning rates, we test the effect of adding skip connections and batch normalization to CNNs of fixed width and varying depth. Figure 4 shows that adding skip connections or batch normalization individually help in training deeper models, but these models still suffer from worsening results and increasing gradient confusion as the network gets deeper. When these techniques are used together, the model has relatively low gradient confusion even for very deep networks, significantly improving trainability of deep models. Note that our observations are consistent with prior work (De & Smith, 2020; Yang et al., 2019).

8. Alternate definitions of gradient confusion

Note that the gradient confusion bound η in equation 3 is defined for the worst-case gradient inner product. However, all the results in this paper can be trivially extended to using a bound on the average gradient inner product of the form:

$$\sum_{i,j=1}^{N} \langle \nabla f_i(\mathbf{w}), \nabla f_j(\mathbf{w}) \rangle / N^2 \geq -\eta.$$

In this case, all theoretical results would remain the same up to constants. We can also define a normalized variant of the gradient confusion condition:

$$\langle \nabla f_i(\mathbf{w}), \nabla f_j(\mathbf{w}) \rangle / (\|\nabla f_i(\mathbf{w})\| \|\nabla f_j(\mathbf{w})\|) \geq -\eta.$$

This condition inherently makes an additional assumption that the norm of the stochastic gradients, $\|\nabla f_i(\mathbf{w})\|$, is bounded, and thus the gradient variance is also bounded (see discussion in section 9). Thus, while all our theoretical results would qualitatively remain the same under this condition, we can prove tighter versions of our current results.

Finally, note that gradient confusion condition in equation 3 is applicable even when the stochastic gradients are averaged over minibatches of size B. The variance of the gradient inner product scales down as $1/B^2$ in this case, and thus η is expected to decrease as B grows.

9. Related work

The gradient confusion bound and our theoretical results have interesting connections to prior work. In this section, we briefly discuss some of these connections.

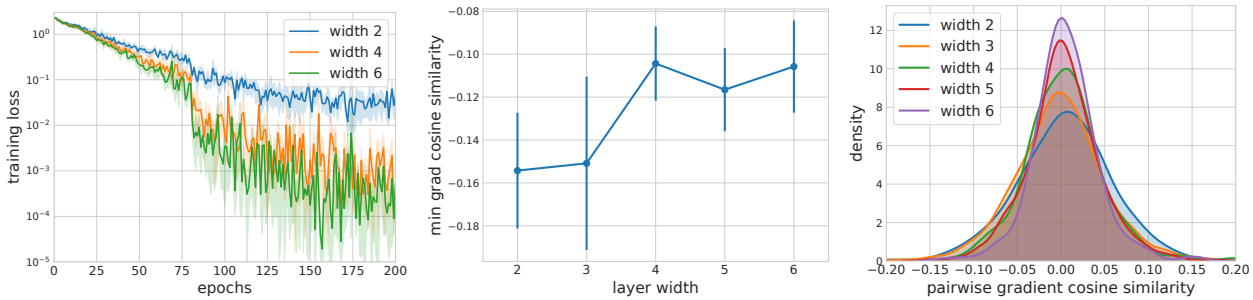

(a) (b) (c)

Figure 3. The effect of width with CNN-16-ℓ on CIFAR-10 for width factors $\ell = 2, 3, 4, 5$ and 6. Plots show the (a) convergence curves of SGD (for cleaner figures, we plot results for width factors 2, 4 and 6 here), (b) minimum of pairwise gradient cosine similarities at the end of training, and the (c) kernel density estimate of the pairwise gradient cosine similarities at the end of training (over all independent runs).

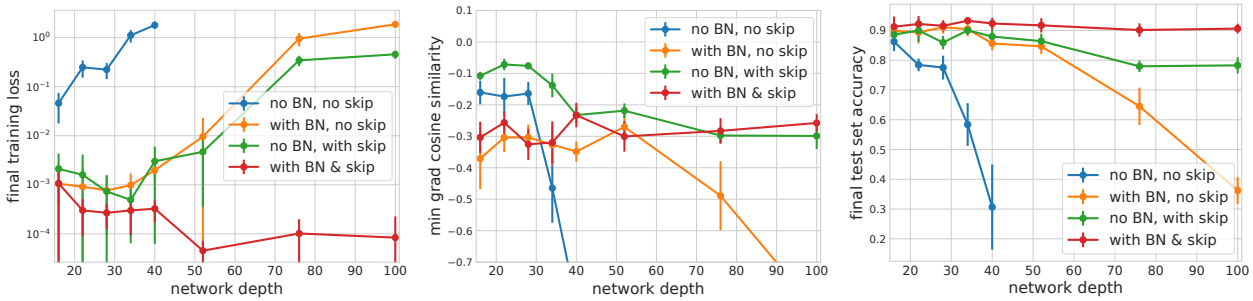

(a) (b) (c)

Figure 4. The effect of adding skip connections and batch normalization to CNN-β-2 on CIFAR-10 for depths $\beta = 16, 22, 28, 34, 40, 52, 76$ and 100. Plots show the (a) optimal training losses, (b) minimum pairwise gradient cosine similarities, and the (c) test set accuracies at the end of training.

Connections to the gradient variance: If we assume bounded gradient variance $\mathbb{E}_i\|\nabla f_i(\mathbf{w}) - \nabla F(\mathbf{w})\|^2 \leq \sigma^2$, we can bound the gradient confusion parameter η in terms of other quantities. For example, suppose the true gradient $\nabla F(\mathbf{w}) = \nabla f_1(\mathbf{w})/2 + \nabla f_2(\mathbf{w})/2$. Then we can write: $|\langle \nabla f_1(\mathbf{w}), \nabla f_2(\mathbf{w})\rangle| \leq \sigma^2 + \|\nabla F(\mathbf{w})\|^2$. However, in general one cannot bound the gradient variance in terms of the gradient confusion parameter. As a counter-example, consider a problem with the following distribution on the gradients: $\frac{1}{1-p}$ samples with gradient $\frac{1}{\epsilon}$ and $\frac{1}{p}$ samples with gradient ϵ, where $p = \epsilon \to 0$. In this case, the gradients are positive, so gradient confusion $\eta = 0$. The mean of the gradients is given by $1+\epsilon(1-\epsilon)$, which remains bounded as $\epsilon \to 0$. On the other hand, the variance (and thus the squared norm of the stochastic gradients) is unbounded ($O(1/\epsilon)$ as $\epsilon \to 0$). A consequence of this is that in theorems 3.1 and 3.2, the "noise term" (i.e., the second term in the RHS of the convergence bounds) does not depend on the learning rate in the general case. If gradients have unbounded variance, lowering the learning rate does not reduce the variance of the SGD updates, and thus does not reduce the noise term.

Connections to gradient diversity: Gradient diversity (Yin et al., 2017) also measures the degree to which individual gradients at different data samples are different from each other. However, the gradient diversity measure gets larger as the individual gradients become orthogonal to each other, and further increases as the gradients start pointing in opposite directions. On the other hand, gradient confusion between two individual gradients is zero unless the inner product between them is negative. As we show in this paper, this has important implications when we study the convergence of SGD in the over-parameterized setting: increased width makes gradients more orthogonal to each other improving trainability, while increased depth result in gradients pointing in opposite directions making networks harder to train. Thus, we view our papers to be complementary, providing insights about different issues (large batch distributed training vs. small batch convergence).

Related work on the impact of network architecture: Balduzzi et al. (2017) studied neural networks with ReLU activations at Gaussian initializations, and showed that gradients become increasingly negatively correlated with depth.

Hanin (2018) showed that the variance of gradients in fully connected networks with ReLU activations is exponential in the sum of the reciprocals of the hidden layer widths at Gaussian initializations. In a follow-up work, Hanin & Rolnick (2018) showed that this sum of the reciprocals of the hidden layer widths determines the variance of the sizes of the activations at each layer. When this sum of reciprocals is too large, early training dynamics are very slow, suggesting the difficulties of starting training on deeper networks, as well as the benefits of increased width.

Other work on SGD convergence: There has recently been a lot of interest in analyzing conditions under which SGD converges to global minimizers of over-parameterized linear and non-linear neural networks. Arora et al. (2018) shows SGD converges linearly to global minimizers for linear neural networks under certain conditions. Du et al. (2018); Allen-Zhu et al. (2018); Zou et al. (2018); Brutzkus et al. (2017) also show convergence to global minimizers of SGD for non-linear networks. This paper complements these recent results by studying how low gradient confusion contributes to SGD's success on over-parameterized neural networks used in practice.

10. Discussion

In this paper, we study how neural network architecture affects the trainability of networks and the dynamics of SGD. To rigorously analyze this, we introduce a concept called gradient confusion, and show that when gradient confusion is low, SGD has fast convergence. We show at standard Gaussian initializations, increasing layer width leads to lower gradient confusion, making the model easier to train. In contrast, increasing depth results in higher gradient confusion, making models harder to train. These results indicate that increasing the layer width with the network depth is important to maintain trainability of the neural network. This is supported by other recent work that suggest that the width should increase linearly with depth in a Gaussian-initialized neural network to help dynamics early in training (Hanin, 2018; Hanin & Rolnick, 2018).

Many previous results have shown how deeper models are more efficient at modeling higher complexity function classes than wider models, and thus depth is essential for the success of neural networks (Eldan & Shamir, 2016; Telgarsky, 2016). Indeed, practitioners over the years have achieved state-of-the-art results on various tasks by making networks deeper, without necessarily making networks wider. We thus study techniques that enable us to train deep models without requiring us to increase the width with depth. Most state-of-the-art neural networks currently contain both skip connections and normalization layers. We thus, empirically study the effect of introducing batch normalization and skip connections to a neural network. We show that the combination of batch normalization and skip connections lower gradient confusion and help train very deep models, explaining why many neural networks used in practice are so efficiently trained. Furthermore, we show how orthogonal initialization techniques provide a promising direction for improving the trainability of very deep networks.

Our results provide a number of important insights that can be used for neural network model design. We demonstrate that the gradient confusion condition could be useful as a measure of trainability of networks, and thus could potentially be used to develop algorithms for more efficient training. Additionally, the correlation between gradient confusion and the test set accuracies shown in appendix A suggest that an interesting topic for future work would be to investigate the connection between gradient confusion and generalization (Fort et al., 2019). Our results also suggest the importance of further work on orthogonal initialization schemes for neural networks with non-linear activations that make training very deep models possible.

Acknowledgements

The authors thank Brendan O'Donoghue, Aleksandar Botev, James Martens, Sudha Rao, and Samuel L. Smith for helpful discussions and for reviewing earlier versions of this manuscript. This paper was supported by the ONR MURI program, AFOSR MURI Program, and the National Science Foundation DMS directorate.

References

Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018.

Arora, S., Cohen, N., Golowich, N., and Hu, W. A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*, 2018.

Balduzzi, D., Frean, M., Leary, L., Lewis, J., Ma, K. W.-D., and McWilliams, B. The shattered gradients problem: If resnets are the answer, then what is the question? *arXiv preprint arXiv:1702.08591*, 2017.

Bengio, Y., Simard, P., and Frasconi, P. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.

Bertsekas, D. P. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Optimization for Machine Learning*, 2010(1-38):3, 2011.

Boucheron, S., Lugosi, G., and Massart, P. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

Brutzkus, A., Globerson, A., Malach, E., and Shalev-Shwartz, S. Sgd learns over-parameterized networks that provably generalize on linearly separable data. *arXiv preprint arXiv:1710.10174*, 2017.

Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. Entropy-sgd: Biasing gradient descent into wide valleys. *arXiv preprint arXiv:1611.01838*, 2016.

Cooper, Y. The loss landscape of overparameterized neural networks. *arXiv preprint arXiv:1804.10200*, 2018.

Darken, C. and Moody, J. Towards faster stochastic gradient search. In *Advances in neural information processing systems*, pp. 1009–1016, 1992.

De, S. and Smith, S. L. Batch normalization biases residual blocks towards the identity function in deep networks. *arXiv preprint arXiv:2002.10444*, 2020.

De, S., Yadav, A., Jacobs, D., and Goldstein, T. Automated inference with adaptive batches. In *Artificial Intelligence and Statistics*, pp. 1504–1513, 2017.

Du, S. S., Zhai, X., Poczos, B., and Singh, A. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.

Dziugaite, G. K. and Roy, D. M. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.

Eldan, R. and Shamir, O. The power of depth for feedforward neural networks. In *Conference on Learning Theory*, pp. 907–940, 2016.

Fort, S., Nowak, P. K., and Narayanan, S. Stiffness: A new perspective on generalization in neural networks. *arXiv preprint arXiv:1901.09491*, 2019.

Ghorbani, B., Krishnan, S., and Xiao, Y. An investigation into neural net optimization via hessian eigenvalue density. *arXiv preprint arXiv:1901.10159*, 2019.

Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256, 2010.

Hanin, B. Which neural net architectures give rise to exploding and vanishing gradients? In *Advances in Neural Information Processing Systems*, pp. 582–591, 2018.

Hanin, B. and Rolnick, D. How to start training: The effect of initialization and architecture. In *Advances in Neural Information Processing Systems*, pp. 571–581, 2018.

He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8580–8589, 2018.

Karimi, H., Nutini, J., and Schmidt, M. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

LeCun, Y. A., Bottou, L., Orr, G. B., and Müller, K.-R. Efficient backprop. In *Neural networks: Tricks of the trade*, pp. 9–48. Springer, 2012.

Lee, J., Xiao, L., Schoenholz, S. S., Bahri, Y., Sohl-Dickstein, J., and Pennington, J. Wide neural networks of any depth evolve as linear models under gradient descent. *arXiv preprint arXiv:1902.06720*, 2019.

Lojasiewicz, S. Ensembles semi-analytiques. *Lectures Notes IHES (Bures-sur-Yvette)*, 1965.

Ma, S., Bassily, R., and Belkin, M. The power of interpolation: Understanding the effectiveness of sgd in modern over-parametrized learning. *arXiv preprint arXiv:1712.06559*, 2017.

Martens, J. *Second-order optimization for neural networks*. University of Toronto (Canada), 2016.

Milman, V. D. and Schechtman, G. *Asymptotic Theory of Finite Dimensional Normed Spaces*. Springer-Verlag, Berlin, Heidelberg, 1986. ISBN 0-387-16769-2.

Moulines, E. and Bach, F. R. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pp. 451–459, 2011.

Nagarajan, V. and Kolter, J. Z. Generalization in deep networks: The role of distance from initialization. *arXiv preprint arXiv:1901.01672*, 2019.

Nedić, A. and Bertsekas, D. Convergence rate of incremental subgradient algorithms. In *Stochastic optimization: algorithms and applications*, pp. 223–264. Springer, 2001.

Needell, D., Ward, R., and Srebro, N. Stochastic gradient descent, weighted sampling, and the randomized kaczmarz algorithm. In *Advances in Neural Information Processing Systems*, pp. 1017–1025, 2014.

Nesterov, Y. *Lectures on convex optimization*, volume 137. Springer, 2018.

Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., and Srebro, N. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018.

Nguyen, Q. and Hein, M. The loss surface of deep and wide neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2603–2612. JMLR. org, 2017.

Oymak, S. and Soltanolkotabi, M. Overparameterized nonlinear learning: Gradient descent takes the shortest path? *arXiv preprint arXiv:1812.10004*, 2018.

Robbins, H. and Monro, S. A stochastic approximation method. *The annals of mathematical statistics*, pp. 400–407, 1951.

Sagun, L., Evci, U., Guney, V. U., Dauphin, Y., and Bottou, L. Empirical analysis of the hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017.

Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.

Schmidt, M. and Roux, N. L. Fast convergence of stochastic gradient descent under a strong growth condition. *arXiv preprint arXiv:1308.6370*, 2013.

Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. Deep information propagation. *arXiv preprint arXiv:1611.01232*, 2016.

Sedghi, H., Gupta, V., and Long, P. M. The singular values of convolutional layers. *arXiv preprint arXiv:1805.10408*, 2018.

Shamir, O. and Zhang, T. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning*, pp. 71–79, 2013.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Smith, S. L., Elsen, E., and De, S. On the generalization benefit of noise in stochastic gradient descent. *arXiv preprint arXiv:2006.15081*, 2020.

Sutskever, I., Martens, J., Dahl, G., and Hinton, G. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pp. 1139–1147, 2013.

Tao, T. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.

Telgarsky, M. Benefits of depth in neural networks. *arXiv preprint arXiv:1602.04485*, 2016.

Vaswani, S., Bach, F., and Schmidt, M. Fast and faster convergence of sgd for over-parameterized models and an accelerated perceptron. *arXiv preprint arXiv:1810.07288*, 2018.

Vershynin, R. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.

Wilson, A. C., Roelofs, R., Stern, M., Srebro, N., and Recht, B. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, pp. 4151–4161, 2017.

Wu, L., Zhu, Z., et al. Towards understanding generalization of deep learning: Perspective of loss landscapes. *arXiv preprint arXiv:1706.10239*, 2017.

Xiao, L., Bahri, Y., Sohl-Dickstein, J., Schoenholz, S. S., and Pennington, J. Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. *arXiv preprint arXiv:1806.05393*, 2018.

Yang, G., Pennington, J., Rao, V., Sohl-Dickstein, J., and Schoenholz, S. S. A mean field theory of batch normalization. *arXiv preprint arXiv:1902.08129*, 2019.

Yin, D., Pananjady, A., Lam, M., Papailiopoulos, D., Ramchandran, K., and Bartlett, P. Gradient diversity: a key ingredient for scalable distributed learning. *arXiv preprint arXiv:1706.05699*, 2017.

Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

Zou, D., Cao, Y., Zhou, D., and Gu, Q. Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*, 2018.