
The Loss Surface of Deep and Wide Neural Networks

Quynh Nguyen¹ Matthias Hein¹

Abstract

While the optimization problem behind deep neural networks is highly non-convex, it is frequently observed in practice that training deep networks seems possible without getting stuck in suboptimal points. It has been argued that this is the case as all local minima are close to being globally optimal. We show that this is (almost) true, in fact almost all local minima are globally optimal, for a fully connected network with squared loss and analytic activation function given that the number of hidden units of one layer of the network is larger than the number of training points and the network structure from this layer on is pyramidal.

1. Introduction

The application of deep learning (LeCun et al., 2015) has in recent years lead to a dramatic boost in performance in many areas such as computer vision, speech recognition or natural language processing. Despite this huge empirical success, the theoretical understanding of deep learning is still limited. In this paper we address the non-convex optimization problem of training a feedforward neural network. This problem turns out to be very difficult as there can be exponentially many distinct local minima (Auer et al., 1996; Safran & Shamir, 2016). It has been shown that the training of a network with a single neuron with a variety of activation functions turns out to be NP-hard (Sima, 2002).

In practice local search techniques like stochastic gradient descent or variants are used for training deep neural networks. Surprisingly, it has been observed (Dauphin et al., 2014; Goodfellow et al., 2015) that in the training of state-of-the-art feedforward neural networks with sparse connectivity like convolutional neural networks (LeCun et al., 1990; Krizhevsky et al., 2012) or fully connected ones one

does not encounter problems with suboptimal local minima. However, as the authors admit themselves in (Goodfellow et al., 2015), the reason for this might be that there is a connection between the fact that these networks have good performance and that they are easy to train.

On the theoretical side there have been several interesting developments recently, see e.g. (Brutzkus & Globerson, 2017; Lee et al., 2016; Poggio & Liao, 2017; Rister & Rubin, 2017; Soudry & Hoffer, 2017; Zhou & Feng, 2017). For some class of networks one can show that one can train them globally optimal efficiently. However, it turns out that these approaches are either not practical (Janzamin et al., 2016; Haeffele & Vidal, 2015; Soltanolkotabi, 2017) as they require e.g. knowledge about the data generating measure, or they modify the neural network structure and objective (Gautier et al., 2016). One class of networks which are simpler to analyze are deep linear networks for which it has been shown that every local minimum is a global minimum (Baldi & Hornik, 1988; Kawaguchi, 2016). While this is a highly non-trivial result as the optimization problem is non-convex, deep linear networks are not interesting in practice as one efficiently just learns a linear function. In order to characterize the loss surface for general networks, an interesting approach has been taken by (Choromanska et al., 2015a). By randomizing the nonlinear part of a feedforward network with ReLU activation function and making some additional simplifying assumptions, they can relate it to a certain spin glass model which one can analyze. In this model the objective of local minima is close to the global optimum and the number of bad local minima decreases quickly with the distance to the global optimum. This is a very interesting result but is based on a number of unrealistic assumptions (Choromanska et al., 2015b). It has recently been shown (Kawaguchi, 2016) that if some of these assumptions are dropped one basically recovers the result of the linear case, but the model is still unrealistic.

In this paper we analyze the case of overspecified neural networks, that is the network is larger than what is required to achieve minimum training error. Under overspecification (Safran & Shamir, 2016) have recently analyzed under which conditions it is possible to generate an initialization so that it is in principle possible to reach the global optimum with descent methods. However, they can only deal with one hidden layer networks and have to make strong

¹Department of Mathematics and Computer Science, Saarland University, Germany. Correspondence to: Quynh Nguyen <quynh@cs.uni-saarland.de>.

assumptions on the data such as linear independence or cluster structure. In this paper overspecification means that there exists a very wide layer, where the number of hidden units is larger than the number of training points. For this case, we can show that a large class of local minima is globally optimal. In fact, we will argue that almost every critical point is globally optimal. Our results generalize previous work of (Yu & Chen, 1995), who have analyzed a similar setting for one hidden layer networks, to networks of arbitrary depth. Moreover, it extends results of (Gori & Tesi, 1992; Frasconi et al., 1997) who have shown that for certain deep feedforward neural networks almost all local minima are globally optimal whenever the training data is linearly independent. While it is clear that our assumption on the number of hidden units is quite strong, there are several recent neural network structures which contain a quite wide hidden layer relative to the number of training points e.g. in (Lin et al., 2016) they have 50,000 training samples and the network has one hidden layer with 10,000 hidden units and (Ba & Caruana, 2014) have 1.1 million training samples and a layer with 400,000 hidden units. We refer to (Ciresan et al., 2010; Neyshabur et al., 2015; Vincent et al., 2010; Caruana et al., 2001) for other examples where the number of hidden units of one layer is on the order of the number of training samples. We conjecture that for these kind of wide networks it still holds that almost all local minima are globally optimal. The reason is that one can expect linear separability of the training data in the wide layer. We provide supporting evidence for this conjecture by showing that basically every critical point for which the training data is linearly separable in the wide layer is globally optimal. Moreover, we want to emphasize that all of our results hold for neural networks used in practice. There are no simplifying assumptions as in previous work.

2. Feedforward Neural Networks and Backpropagation

We are mainly concerned with multi-class problems but our results also apply to multivariate regression problems. Let N be the number of training samples and denote by $X = [x_1, \dots, x_N]^T \in \mathbb{R}^{N \times d}$, $Y = [y_1, \dots, y_N]^T \in \mathbb{R}^{N \times m}$ the input resp. output matrix for the training data $(x_i, y_i)_{i=1}^N$, where d is the input dimension and m the number of classes. We consider fully-connected feedforward networks with L layers, indexed from 0, 1, 2, \dots , L , which correspond to the input layer, 1st hidden layer, etc, and output layer. The network structure is determined by the weight matrices $(W_k)_{k=1}^L \in \mathcal{W} := \mathbb{R}^{d \times n_1} \times \dots \times \mathbb{R}^{n_{k-1} \times n_k} \times \dots \times \mathbb{R}^{n_{L-1} \times m}$; where n_k is the number of hidden units of layer k (for consistency, we set $n_0 = d, n_L = m$), and the bias vectors $(b_k)_{k=1}^L \in \mathcal{B} := \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_L}$. We denote by $\mathcal{P} = \mathcal{W} \times \mathcal{B}$ the space of all possible parameters of the network. In this paper, $[a]$ denotes the set

of integers $\{1, 2, \dots, a\}$ and $[a, b]$ the set of integers from a to b . The activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is assumed at least to be continuously differentiable, that is $\sigma \in C^1(\mathbb{R})$. In this paper, we assume that all the functions are applied componentwise. Let $f_k, g_k : \mathbb{R}^d \rightarrow \mathbb{R}^{n_k}$ be the mappings from the input space to the feature space at layer k , which are defined as

$$f_0(x) = x, f_k(x) = \sigma(g_k(x)), g_k(x) = W_k^T f_{k-1}(x) + b_k$$

for every $k \in [L], x \in \mathbb{R}^d$. In the following, let $F_k = [f_k(x_1), f_k(x_2), \dots, f_k(x_N)]^T \in \mathbb{R}^{N \times n_k}$ and $G_k = [g_k(x_1), g_k(x_2), \dots, g_k(x_N)]^T \in \mathbb{R}^{N \times n_k}$ be the matrices that store the feature vectors of layer k after and before applying the activation function. One can easily check that

$$F_1 = \sigma(XW_1 + \mathbf{1}_N b_1^T), \\ F_k = \sigma(F_{k-1}W_k + \mathbf{1}_N b_k^T), \text{ for } k \in [2, L].$$

In this paper we analyze the behavior of the loss of the network without any form of regularization, that is the final objective $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ of the network is defined as

$$\Phi\left((W_k, b_k)_{k=1}^L\right) = \sum_{i=1}^N \sum_{j=1}^m l(f_{Lj}(x_i) - y_{ij}) \quad (1)$$

where $l : \mathbb{R} \rightarrow \mathbb{R}$ is assumed to be a continuously differentiable loss function, that is $l \in C^1(\mathbb{R})$. The prototype loss which we consider in this paper is the squared loss, $l(\alpha) = \alpha^2$, which is one of the standard loss functions in the neural network literature. We assume throughout this paper that the minimum of (1) is attained.

The idea of backpropagation is the core of our theoretical analysis. Lemma 2.1 below shows well-known relations for feed-forward neural networks, which are used throughout the paper. The derivative of the loss w.r.t. the value of unit j at layer k evaluated at a single training sample x_i is denoted as $\delta_{kj}(x_i) = \frac{\partial \Phi}{\partial g_{kj}(x_i)}$. We arrange these vectors for all training samples into a single matrix Δ_k , defined as

$$\Delta_k = [\delta_{k:}(x_1), \dots, \delta_{k:}(x_N)]^T \in \mathbb{R}^{N \times n_k}.$$

In the following we use the Hadamard product \circ , which for $A, B \in \mathbb{R}^{m \times n}$ is defined as $A \circ B \in \mathbb{R}^{m \times n}$ with $(A \circ B)_{ij} = A_{ij}B_{ij}$.

Lemma 2.1 *Let $\sigma, l \in C^1(\mathbb{R})$. Then it holds*

1. $\Delta_k = \begin{cases} l'(F_L - Y) \circ \sigma'(G_L), & k = L \\ (\Delta_{k+1} W_{k+1}^T) \circ \sigma'(G_k), & k \in [L-1] \end{cases}$
2. $\nabla_{W_k} \Phi = \begin{cases} X^T \Delta_1, & k = 1 \\ F_{k-1}^T \Delta_k, & k \in [2, L] \end{cases}$

$$3. \nabla_{b_k} \Phi = \Delta_k^T \mathbf{1}_N \forall k \in [L]$$

Proof:

1. By definition, it holds for every $i \in [N], j \in [n_L]$ that

$$\begin{aligned} (\Delta_L)_{ij} &= \delta_{Lj}(x_i) \\ &= \frac{\partial \Phi}{\partial g_{Lj}(x_i)} \\ &= l'(f_{Lj}(x_i) - y_{ij}) \sigma'(g_{Lj}(x_i)) \\ &= l'((F_L)_{ij} - Y_{ij}) \sigma'((G_L)_{ij}) \end{aligned}$$

and hence, $\Delta_L = l'(F_L - Y) \circ \sigma'(G_L)$.

For every $k \in [L-1]$, the chain rule yields for every $i \in [N], j \in [n_k]$ that

$$\begin{aligned} (\Delta_k)_{ij} &= \delta_{kj}(x_i) \\ &= \frac{\partial \Phi}{\partial g_{kj}(x_i)} \\ &= \sum_{l=1}^{n_{k+1}} \frac{\partial \Phi}{\partial g_{(k+1)l}(x_i)} \frac{\partial g_{(k+1)l}(x_i)}{\partial g_{kj}(x_i)} \\ &= \sum_{l=1}^{n_{k+1}} \delta_{(k+1)l}(x_i) (W_{k+1})_{jl} \sigma'(g_{kj}(x_i)) \\ &= \sum_{l=1}^{n_{k+1}} (\Delta_{(k+1)})_{il} (W_{k+1})_{lj}^T \sigma'((G_k)_{ij}) \end{aligned}$$

and hence $\Delta_k = (\Delta_{k+1} W_{k+1}^T) \circ \sigma'(G_k)$.

2. For every $r \in [d], s \in [n_1]$ it holds

$$\begin{aligned} \frac{\partial \Phi}{\partial (W_1)_{rs}} &= \sum_{i=1}^N \frac{\partial \Phi}{\partial g_{1s}(x_i)} \frac{\partial g_{1s}(x_i)}{\partial (W_1)_{rs}} \\ &= \sum_{i=1}^N \delta_{1s}(x_i) x_{ir} = \sum_{i=1}^N (X^T)_{ri} (\Delta_1)_{is} \\ &= (X^T \Delta_1)_{rs} \end{aligned}$$

and hence $\nabla_{W_1} \Phi = X^T \Delta_1$.

For every $k \in [2, L], r \in [n_{k-1}], s \in [n_k]$, one obtains

$$\begin{aligned} \frac{\partial \Phi}{\partial (W_k)_{rs}} &= \sum_{i=1}^N \frac{\partial \Phi}{\partial g_{ks}(x_i)} \frac{\partial g_{ks}(x_i)}{\partial (W_k)_{rs}} \\ &= \sum_{i=1}^N \delta_{ks}(x_i) f_{(k-1)r}(x_i) = \sum_{i=1}^N (F_{k-1}^T)_{ri} (\Delta_k)_{is} \\ &= (F_{k-1}^T \Delta_k)_{rs} \end{aligned}$$

and hence $\nabla_{W_k} \Phi = F_{k-1}^T \Delta_k$.

3. For every $k \in [1, L], s \in [n_k]$ it holds

$$\begin{aligned} \frac{\partial \Phi}{\partial (b_k)_s} &= \sum_{i=1}^N \frac{\partial \Phi}{\partial g_{ks}(x_i)} \frac{\partial g_{ks}(x_i)}{\partial (b_k)_s} \\ &= \sum_{i=1}^N \delta_{ks}(x_i) = (\Delta_k^T \mathbf{1}_N)_s \end{aligned}$$

and hence $\nabla_{b_k} \Phi = \Delta_k^T \mathbf{1}_N$.

□

Note that Lemma 2.1 does not apply to non-differentiable activation functions like the ReLU function, $\sigma_{\text{ReLU}}(x) = \max\{0, x\}$. However, it is known that one can approximate this activation function arbitrarily well by a smooth function e.g. $\sigma_\alpha(x) = \frac{1}{\alpha} \log(1 + e^{\alpha x})$ (a.k.a. softplus) satisfies $\lim_{\alpha \rightarrow \infty} \sigma_\alpha(x) = \sigma_{\text{ReLU}}(x)$ for any $x \in \mathbb{R}$.

3. Main Result

We first discuss some prior work and present then our main result together with extensive discussion. For improved readability we postpone the proof of the main result to the next section which contains several intermediate results which are of independent interest.

3.1. Previous Work

Our work can be seen as a generalization of the work of (Gori & Tesi, 1992; Yu & Chen, 1995). While (Yu & Chen, 1995) has shown that for a one-hidden layer network, that if $n_1 = N - 1$, then every local minimum is a global minimum, the work of (Gori & Tesi, 1992) considered also multi-layer networks. For the convenience of the reader, we first restate Theorem 1 of (Gori & Tesi, 1992) using our previously introduced notation. The critical points of a continuously differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ are the points where the gradient vanishes, that is $\nabla f(x) = 0$. Note that this is a necessary condition for a local minimum.

Theorem 3.1 (Gori & Tesi, 1992) *Let $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ be defined as in (1) with least squares loss $l(a) = a^2$. Assume $\sigma : \mathbb{R} \rightarrow [d, \bar{d}]$ to be continuously differentiable with strictly positive derivative and*

$$\begin{aligned} \lim_{a \rightarrow \infty} \frac{\sigma'(a)}{d - \sigma(a)} &> 0, & \lim_{a \rightarrow \infty} \frac{-\sigma''(a)}{d - \sigma(a)} &> 0 \\ \lim_{a \rightarrow -\infty} \frac{\sigma'(a)}{\sigma(a) - d} &> 0, & \lim_{a \rightarrow -\infty} \frac{\sigma''(a)}{\sigma(a) - d} &> 0 \end{aligned}$$

Then every critical point $(W_l, b_l)_{l=1}^L$ of Φ which satisfies the conditions

1. $\text{rank}(W_l) = n_l$ for all $l \in [2, L]$,
2. $[X, \mathbf{1}_N]^T \Delta_1 = 0$ implies $\Delta_1 = 0$

is a global minimum.

While this result is already for general multi-layer networks, the condition “[$X, \mathbf{1}_N$] $^T \Delta_1 = 0$ implies $\Delta_1 = 0$ ” is the main caveat. It is already noted in (Gori & Tesi, 1992), that “it is quite hard to understand its practical meaning” as it requires prior knowledge of Δ_1 at every critical point. Note that this is almost impossible as Δ_1 depends on all the weights of the network. For a particular case, when the training samples (biases added) are linearly independent, i.e. $\text{rank}([X, \mathbf{1}_N]) = N$, the condition holds automatically. This case is discussed in the following Theorem 3.4, where we consider a more general class of loss and activation functions.

3.2. First Main Result and Discussion

A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is real analytic if the corresponding multivariate Taylor series converges to $f(x)$ on an open subset of \mathbb{R}^d (Krantz & Parks, 2002). All results in this section are proven under the following assumptions on the loss/activation function and training data.

Assumptions 3.2 1. There are no identical training samples, i.e. $x_i \neq x_j$ for all $i \neq j$,

2. σ is analytic on \mathbb{R} , strictly monotonically increasing and

(a) σ is bounded or

(b) there are positive $\rho_1, \rho_2, \rho_3, \rho_4$, s.t. $|\sigma(t)| \leq \rho_1 e^{\rho_2 t}$ for $t < 0$ and $|\sigma(t)| \leq \rho_3 t + \rho_4$ for $t \geq 0$

3. $l \in C^2(\mathbb{R})$ and if $l'(a) = 0$ then a is a global minimum

These conditions are not always necessary to prove some of the intermediate results presented below, but we decided to provide the proof under the above strong assumptions for better readability. For instance, all of our results also hold for strictly monotonically decreasing activation functions. Note that the above conditions are not restrictive as many standard activation functions satisfy them.

Lemma 3.3 The sigmoid activation function $\sigma_1(t) = \frac{1}{1+e^{-t}}$, the tangent hyperbolic $\sigma_2(t) = \tanh(t)$ and the softplus function $\sigma_3(t) = \frac{1}{\alpha} \log(1 + e^{\alpha t})$ for $\alpha > 0$ satisfy Assumption 3.2.

Proof: Note that $\sigma_2(t) = \frac{2}{1+e^{-2t}} - 1$. Moreover, it is well known that $\phi(t) = \frac{1}{1+t}$ is real-analytic on $\mathbb{R}_+ = \{t \in \mathbb{R} \mid t \geq 0\}$. The exponential function is analytic with values in $(0, \infty)$. As composition of real-analytic function is real-analytic (see Prop 1.4.2 in (Krantz & Parks, 2002)), we get that σ_1 and σ_2 are real-analytic. Similarly, since $\log(1+t)$ is real-analytic on $(-1, \infty)$ and the composition with the exponential function is real-analytic, we get that σ_3 is a real-analytic function.

Finally, we note that $\sigma_1, \sigma_2, \sigma_3$ are strictly monotonically increasing. Since σ_1, σ_2 are bounded, they both satisfy Assumption 3.2. For σ_3 , we note that $1 + e^{\alpha t} \leq 2e^{\alpha t}$ for $t \geq 0$, and thus it holds for every $t \geq 0$ that

$$\begin{aligned} 0 \leq \sigma_3(t) &= \frac{1}{\alpha} \log(1 + e^{\alpha t}) \\ &\leq \frac{1}{\alpha} \log(2e^{\alpha t}) \\ &= \frac{\log(2)}{\alpha} + t, \end{aligned}$$

and with $\log(1+x) \leq x$ for $x > -1$ it holds $\log(1+e^{\alpha t}) \leq e^{\alpha t}$ for every $t \in \mathbb{R}$. In particular

$$0 \leq \sigma_3(t) \leq \frac{e^{\alpha t}}{\alpha} \quad \forall t < 0$$

which implies that σ_3 satisfies Assumption 3.2 for $\rho_1 = 1/\alpha, \rho_2 = \alpha, \rho_3 = 1, \rho_4 = \log(2)/\alpha$. \square

The conditions on l are satisfied for any twice continuously differentiable convex loss function. A typical example is the squared loss $l(a) = a^2$ or the Pseudo-Huber loss (Hartley & Zisserman, 2004) given as $l_\delta(a) = 2\delta^2(\sqrt{1+a^2/\delta^2} - 1)$ which approximates a^2 for small a and is linear with slope 2δ for large a . But also non-convex loss functions satisfy this requirement, for instance:

1. Blake-Zisserman: $l(a) = -\log(\exp(-a^2) + \delta)$ for $\delta > 0$. For small a , this curve approximates a^2 , whereas for large a the asymptotic value is $-\log(\delta)$.

2. Corrupted-Gaussian:

$$l(a) = -\log(\alpha \exp(-a^2) + (1-\alpha) \exp(-a^2/w^2)/w)$$

for $\alpha \in [0, 1], w > 0$. This function computes the negative log-likelihood of a gaussian mixture model.

3. Cauchy: $l(a) = \delta^2 \log(1 + a^2/\delta^2)$ for $\delta \neq 0$. This curve approximates a^2 for small a and the value of δ determines for what range of a this approximation is close.

We refer to (Hartley & Zisserman, 2004) (p.617-p.619) for more examples and discussion on robust loss functions.

As a motivation for our main result, we first analyze the case when the training samples are linearly independent, which requires $N \leq d+1$. It can be seen as a generalization of Corollary 1 in (Gori & Tesi, 1992).

Theorem 3.4 Let $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ be defined as in (1) and let the Assumptions 3.2 hold. If the training samples are linearly independent, that is $\text{rank}([X, \mathbf{1}_N]) = N$, then every critical point $(W_l, b_l)_{l=1}^L$ of Φ for which the weight matrices $(W_l)_{l=2}^L$ have full column rank, that is $\text{rank}(W_l) = n_l$ for $l \in [2, L]$, is a global minimum.




Figure 1. Examples of convex and non-convex loss functions that satisfy Assumption 3.2.

Proof: The proof is based on induction. At a critical point it holds $\nabla_{W_1} \Phi = X^T \Delta_1 = 0$ and $\nabla_{b_1} \Phi = \Delta_1^T \mathbf{1}_N = 0$ thus $[X, \mathbf{1}_N]^T \Delta_1 = 0$. By assumption, the data matrix $[X, \mathbf{1}_N]^T \in \mathbb{R}^{(d+1) \times N}$ has full column rank, this implies $\Delta_1 = 0$. Using induction, let us assume that $\Delta_k = 0$ for some $1 \leq k \leq L - 1$, then by Lemma 2.1, we have $\Delta_k = (\Delta_{k+1} W_{k+1}^T) \circ \sigma'(G_k) = 0$. As by assumption σ' is strictly positive, this is equivalent to $\Delta_{k+1} W_{k+1}^T = 0$ resp. $W_{k+1} \Delta_{k+1}^T = 0$. As by assumption W_{k+1} has full column rank, it follows $\Delta_{k+1} = 0$. Finally, we get $\Delta_L = 0$. With Lemma 2.1 we thus get $l'(F_L - Y) \circ \sigma'(G_L) = 0$ which implies with the same argument as above $l'(F_L - Y) = 0$. From our Assumption 3.2, it holds that if $l'(a) = 0$ then a is a global minimum of l . Thus each individual entry of $(F_L - Y)$ must represent a global minimum of l . This combined with (1) implies that the critical point must be a global minimum of Φ . \square

Theorem 3.4 implies that the weight matrices of potential saddle points or suboptimal local minima need to have low rank for one particular layer. Note however that the set of low rank weight matrices in \mathcal{W} has measure zero. At the moment we cannot prove that suboptimal low rank local minima cannot exist. However, it seems implausible that such suboptimal low rank local minima exist as every neighborhood of such points contains full rank matrices which increase the expressiveness of the network. Thus it should be possible to use this degree of freedom to further reduce the loss, which contradicts the definition of a local minimum. Thus we conjecture that all local minima are indeed globally optimal.

The main restriction in the assumptions of Theorem 3.4 is the linear independence of the training samples as it requires $N \leq d + 1$, which is very restrictive in practice. We prove in this section a similar guarantee in our main Theorem 3.8 by implicitly transporting this condition to some higher layer. A similar guarantee has been proven by (Yu & Chen, 1995) for a single hidden layer network, whereas we consider general multi-layer networks. The main ingredient of the proof of our main result is the observation in

the following lemma.

Lemma 3.5 *Let $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ be defined as in (1) and let the Assumptions 3.2 hold. Let $(W_l, b_l)_{l=1}^L \in \mathcal{P}$ be given. Assume there is some $k \in [L - 1]$ s.t. the following holds*

1. $\text{rank}([F_k, \mathbf{1}_N]) = N$
2. $\text{rank}(W_l) = n_l, l \in [k + 2, L]$
3. $\nabla_{W_{k+1}} \Phi((W_l, b_l)_{l=1}^L) = 0$
 $\nabla_{b_{k+1}} \Phi((W_l, b_l)_{l=1}^L) = 0$

then $(W_l, b_l)_{l=1}^L$ is a global minimum.

Proof: By Lemma 2.1 it holds that

$$\nabla_{W_{k+1}} \Phi = F_k^T \Delta_{k+1} = 0, \quad \nabla_{b_{k+1}} \Phi = \Delta_{k+1}^T \mathbf{1}_N = 0,$$

which implies $[F_k, \mathbf{1}_N]^T \Delta_{k+1} = 0$. By our assumption, $\text{rank}([F_k, \mathbf{1}_N]) = N$ it holds that $\Delta_{k+1} = 0$. Since $\text{rank}(W_l) = n_l, l \in [k + 2, L]$, we can apply a similar induction argument as in the proof of Theorem 3.4, to arrive at $\Delta_L = 0$ and thus a global minimum. \square

The first condition of Lemma 3.5 can be seen as a generalization of the requirement of linearly independent training inputs in Theorem 3.4 to a condition of linear independence of the feature vectors at a hidden layer. Lemma 3.5 suggests that if we want to make statements about the global optimality of critical points, it is sufficient to know when and which critical points fulfill these conditions. The third condition is trivially satisfied by a critical point and the requirement of full column rank of the weight matrices is similar to Theorem 3.4. However, the first one may not be fulfilled since $\text{rank}([F_k, \mathbf{1}_N])$ is dependent not only on the weights but also on the architecture. The main difficulty of the proof of our following main theorem is to prove that this first condition holds under the rather simple requirement that $n_k \geq N - 1$ for a subset of all critical points.

But before we state the theorem we have to discuss a particular notion of non-degenerate critical point.

Definition 3.6 (Block Hessian) Let $f : D \rightarrow \mathbb{R}$ be a twice-continuously differentiable function defined on some open domain $D \subseteq \mathbb{R}^n$. The Hessian w.r.t. a subset of variables $S \subseteq \{x_1, \dots, x_n\}$ is denoted as $\nabla_S^2 f(x) \in \mathbb{R}^{|S| \times |S|}$. When $|S| = n$, we write $\nabla^2 f(x) \in \mathbb{R}^{n \times n}$ to denote the full Hessian matrix.

We use this to introduce a slightly more general notion of non-degenerate critical point.

Definition 3.7 (Non-degenerate critical point) Let $f : D \rightarrow \mathbb{R}$ be a twice-continuously differentiable function defined on some open domain $D \subseteq \mathbb{R}^n$. Let $x \in D$ be a critical point, i.e. $\nabla f(x) = 0$, then

- x is non-degenerate for a subset of variables $S \subseteq \{x_1, \dots, x_n\}$ if $\nabla_S^2 f(x)$ is non-singular.
- x is non-degenerate if $\nabla^2 f(x)$ is non-singular.

Note that a non-degenerate critical point might not be non-degenerate for a subset of variables, and vice versa, if it is non-degenerate on a subset of variables it does not necessarily imply non-degeneracy on the whole set. For instance,

$$\nabla^2 f(x) = \begin{array}{cc|cc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{array}, \quad \nabla^2 f(y) = \begin{array}{cc|cc} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ \hline 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{array}$$

Clearly, $\det \nabla^2 f(x) = 0$ but $\det \nabla_{\{x_1, x_2\}}^2 f(x) \neq 0$, and $\det \nabla^2 f(y) \neq 0$ but $\det \nabla_{\{y_3, y_4\}}^2 f(y) = 0$. The concept of non-degeneracy on a subset of variables is crucial for the following statement of our main result.

Theorem 3.8 Let $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ be defined as in (1) and let the Assumptions 3.2 hold. Suppose $n_k \geq N - 1$ for some $k \in [L - 1]$. Then every critical point $(W_l^*, b_l^*)_{l=1}^L$ of Φ which satisfies the following conditions

1. $(W_l^*, b_l^*)_{l=1}^L$ is non-degenerate on $\{(W_l, b_l) \mid l \in \mathcal{I}, \text{ for some subset } \mathcal{I} \subseteq \{k+1, \dots, L\} \text{ satisfying } \{k+1\} \in \mathcal{I},\}$
2. $(W_l^*)_{l=k+2}^L$ has full column rank, that is, $\text{rank}(W_l^*) = n_l$ for $l \in [k+2, L]$,

is a global minimum of Φ .

First of all we note that the full column rank condition of $(W_l^*)_{l=k+2}^L$ in Theorem 3.4, and 3.8 implicitly requires that $n_{k+1} \geq n_{k+2} \geq \dots \geq n_L$. This means the network needs to have a pyramidal structure from layer $k+2$ to L . It is interesting to note that most modern neural network architectures have a pyramidal structure from some layer, typically the first hidden layer, on. Thus this is not a restrictive

requirement. Indeed, one can even argue that Theorem 3.8 gives an implicit justification as it hints on the fact that such networks are easy to train if one layer is sufficiently wide.

Note that Theorem 3.8 does not require fully non-degenerate critical points but non-degeneracy is only needed for some subset of variables that includes layer $k+1$. As a consequence of Theorem 3.8, we get directly a stronger result for non-degenerate local minima.

Corollary 3.9 Let $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ be defined as in (1) and let the Assumptions 3.2 hold. Suppose $n_k \geq N - 1$ for some $k \in [L - 1]$. Then every non-degenerate local minimum $(W_l^*, b_l^*)_{l=1}^L$ of Φ for which $(W_l^*)_{l=k+2}^L$ has full column rank, that is $\text{rank}(W_l^*) = n_l$, is a global minimum of Φ .

Proof: The Hessian at a non-degenerate local minimum is positive definite and every principal submatrix of a positive definite matrix is again positive definite, in particular for the subset of variables $(W_l, b_l)_{l=k+1}^L$. Then application of Theorem 3.8 yields the result. \square

Let us discuss the implications of these results. First, note that Theorem 3.8 is slightly weaker than Theorem 3.4 as it requires also non-degeneracy wrt to a set of variables including layer $k+1$. Moreover, similar to Theorem 3.4 it does not exclude the possibility of suboptimal local minima of low rank in the layers ‘‘above’’ layer $k+1$. On the other hand it makes also very strong statements. In fact, if $n_k \geq N - 1$ for some $k \in [L - 1]$ then even degenerate saddle points/local maxima are excluded as long as they are non-degenerate with respect to any subset of parameters of upper layers that include layer $k+1$ and the rank condition holds. Thus given that the weight matrices of the upper layers have full column rank, there is not much room left for degenerate saddle points/local maxima. Moreover, for a one-hidden-layer network for which $n_1 \geq N - 1$, every non-degenerate critical point with respect to the output layer parameters is a global minimum, as the full rank condition is not active for one-hidden layer networks.

Concerning the non-degeneracy condition of main Theorem 3.8, one might ask how likely it is to encounter degenerate points of a smooth function. This is answered by an application of Sard’s/Morse theorem in (Milnor, 1965).

Theorem 3.10 (A. Morse, p.11) If $f : U \subset \mathbb{R}^d \rightarrow \mathbb{R}$ is twice continuously differentiable. Then for almost all $w \in \mathbb{R}^d$ with respect to the Lebesgue measure it holds that f' defined as $f'(x) = f(x) + \langle w, x \rangle$ has only non-degenerate critical points.

Note that the theorem would still hold if one would draw w uniformly at random from the set $\{z \in \mathbb{R}^d \mid \|z\|_2 \leq \epsilon\}$ for any $\epsilon > 0$. Thus almost every linear perturbation f' of

a function f will lead to the fact all of its critical points are non-degenerate. Thus, this result indicates that exact degenerate points might be rare. Note however that in practice the Hessian at critical points can be close to singular (at least up to numerical precision), which might affect the training of neural networks negatively (Sagun et al., 2016).

As we argued for Theorem 3.4 our main Theorem 3.8 does not exclude the possibility of suboptimal degenerate local minima or suboptimal local minima of low rank. However, we conjecture that the second case cannot happen as every neighborhood of the local minima contains full rank matrices which increase the expressiveness of the network and this additional flexibility can be used to reduce the loss which contradicts the definition of a local minimum.

As mentioned in the introduction the condition $n_k \geq N - 1$ looks at first sight very strong. However, as mentioned in the introduction, in practice often networks are used where one hidden layer is rather wide, that is n_k is on the order of N (typically it is the first layer of the network). As the condition of Theorem 3.8 is sufficient and not necessary, one can expect out of continuity reasons that the loss surface of networks where the condition is approximately true, is still rather well behaved, in the sense that still most local minima are indeed globally optimal and the suboptimal ones are not far away from the globally optimal ones.

4. Proof of Main Result

For better readability, we first prove our main Theorem 3.8 for a special case where \mathcal{I} is the whole set of upper layers, i.e. $\mathcal{I} = \{k + 1, \dots, L\}$, and then show how to extend the proof to the general case where $\mathcal{I} \subseteq \{k + 1, \dots, L\}$. Our proof strategy is as follows. We first show that the output of each layer are real analytic functions of network parameters. Then we prove that there exists a set of parameters such that $\text{rank}([F_k, \mathbf{1}_N]) = N$. Using properties of real analytic functions, we conclude that the set of parameters where $\text{rank}([F_k, \mathbf{1}_N]) < N$ has measure zero. Then with the non-degeneracy condition, we can apply the implicit-function theorem to conclude that even if $\text{rank}([F_k, \mathbf{1}_N]) = N$ is not true at a critical point, then still in any neighborhood of it there exists a point where the conditions of Lemma 3.5 are true and the loss is minimal. By continuity of Φ , this implies that the loss must also be minimal at the critical point.

We introduce some notation frequently used in the proofs. Let $B(x, r) = \{z \in \mathbb{R}^d \mid \|x - z\|_2 < r\}$ be the open ball in \mathbb{R}^d of radius r around x .

Lemma 4.1 *If the Assumptions 3.2 hold, then the output of each layer f_l for every $l \in [L]$ are real analytic functions of the network parameters on \mathcal{P} .*

Proof: Any linear function is real analytic and the set of real analytic functions is closed under addition, multiplication and composition, see e.g. Prop. 2.2.2 and Prop. 2.2.8 in (Krantz & Parks, 2002). As we assume that the activation function is real analytic, we get that all the output functions of the neural network f_k are real analytic functions of the parameters as compositions of real analytic functions. \square

The concept of real analytic functions is important in our proofs as these functions can never be “constant” in a set of the parameter space which has positive measure unless they are constant everywhere. This is captured by the following lemma.

Lemma 4.2 (Nguyen, 2015; Mityagin, 2015) *If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a real analytic function which is not identically zero then the set $\{x \in \mathbb{R}^n \mid f(x) = 0\}$ has Lebesgue measure zero.*

In the next lemma we show that there exist network parameters such that $\text{rank}([F_k, \mathbf{1}_N]) = N$ holds if $n_k \geq N - 1$. Note that this is only possible due to the fact that one uses non-linear activation functions. For deep linear networks, it is not possible for F_k to achieve maximum rank if the layers below it are not sufficiently wide. To see this, one considers $F_k = F_{k-1}W_k + \mathbf{1}_N b_k^T$ for a linear network, then $\text{rank}(F_k) \leq \min\{\text{rank}(F_{k-1}), \text{rank}(W_k)\} + 1$ since the addition of a rank-one term does not increase the rank of a matrix by more than one. By using induction, one gets $\text{rank}(F_k) \leq \text{rank}(W_l) + k - l + 1$ for every $l \in [k]$.

The existence of network parameters where $\text{rank}([F_k, \mathbf{1}_N]) = N$ together with the previous lemma will then be used to show that the set of network parameters where $\text{rank}([F_k, \mathbf{1}_N]) < N$ has measure zero.

Lemma 4.3 *If the Assumptions 3.2 hold and $n_k \geq N - 1$ for some $k \in [L - 1]$, then there exists at least one set of parameters $(W_l, b_l)_{l=1}^k$ such that $\text{rank}([F_k, \mathbf{1}_N]) = N$.*

Proof: We first show by induction that there always exists a set of parameters $(W_l, b_l)_{l=1}^{k-1}$ s.t. F_{k-1} has distinct rows. Indeed, we have $F_1 = \sigma(XW_1 + \mathbf{1}_N b_1^T)$. The set of (W_1, b_1) that makes F_1 to have distinct rows is characterized by

$$\sigma(W_1^T x_i + b_1) \neq \sigma(W_1^T x_j + b_1), \quad \forall i \neq j.$$

Note, that σ is strictly monotonic and thus bijective on its domain. Thus this is equivalent to

$$W_1^T(x_i - x_j) \neq 0, \quad \forall i \neq j.$$

Let us denote the first column of W_1 by a , then the exis-

tence of a for which

$$a^T(x_i - x_j) \neq 0, \quad \forall i \neq j, \quad (2)$$

would imply the result. Note that by assumption $x_i \neq x_j$ for all $i \neq j$. Then the set $\{a \in \mathbb{R}^d \mid a^T(x_i - x_j) = 0\}$ is a hyperplane, which has measure zero and thus the set where condition (2) fails corresponds to the union of $\frac{N(N-1)}{2}$ hyperplanes which again has measure zero. Thus there always exists a vector a such that condition (2) is satisfied and thus there exists (W_1, b_1) such that the rows of F_1 are distinct. Now, assume that F_{p-1} has distinct rows for some $p \geq 1$, then by the same argument as above we need to construct W_p such that

$$W_p^T(f_{p-1}(x_i) - f_{p-1}(x_j)) \neq 0, \quad \forall i \neq j.$$

By construction $f_{p-1}(x_i) \neq f_{p-1}(x_j)$ and thus with the same argument as above we can choose W_p such that this condition holds. As a result, there exists a set of parameters $(W_l, b_l)_{l=1}^{k-1}$ so that F_{k-1} has distinct rows.

Now, given that F_{k-1} has distinct rows, we show how to construct (W_k, b_k) in such a way that $[F_k, \mathbf{1}_N] \in \mathbb{R}^{N \times (n_k+1)}$ has full row rank. Since $n_k \geq N - 1$, it is sufficient to make the first $N - 1$ columns of F_k together with the all-ones vector become linearly independent. In particular, let $F_k = [A, B]$ where $A \in \mathbb{R}^{N \times (N-1)}$ and $B \in \mathbb{R}^{N \times (n_k - N + 1)}$ be the matrices containing outputs of the first $(N - 1)$ hidden units and last $(n_k - N + 1)$ hidden units of layer k respectively. Let $W_k = [w_1, \dots, w_{N-1}, w_N, \dots, w_{n_k}] \in \mathbb{R}^{n_{k-1} \times n_k}$ and $b_k = [v_1, \dots, v_{N-1}, v_N, \dots, v_{n_k}] \in \mathbb{R}^{n_k}$. Let $Z = F_{k-1} = [z_1, \dots, z_N]^T \in \mathbb{R}^{N \times n_{k-1}}$ with $z_i \neq z_j$ for every $i \neq j$. By definition of F_k , it holds $A_{ij} = \sigma(z_i^T w_j + v_j)$ for $i \in [N], j \in [N - 1]$. As mentioned above, we just need to show there exists $(w_j, v_j)_{j=1}^{N-1}$ so that $\text{rank}([\mathbf{1}_N, A]) = N$ because then it will follow immediately that $\text{rank}([F_k, \mathbf{1}_N]) = N$. Pick any $a \in \mathbb{R}^{n_{k-1}}$ satisfying potentially after reordering w.l.o.g. $\langle a, z_1 \rangle < \langle a, z_2 \rangle < \dots < \langle a, z_N \rangle$. By the discussion above such a vector always exists since the complementary set is contained in $\bigcup_{i \neq j} \{a \in \mathbb{R}^{n_{k-1}} \mid \langle z_i - z_j, a \rangle = 0\}$ which has measure zero.

We first prove the result for the case where σ is bounded. Since σ is bounded and strictly monotonically increasing, there exist two finite values $\gamma, \mu \in \mathbb{R}$ with $\mu < \gamma$ s.t.

$$\lim_{\alpha \rightarrow -\infty} \sigma(\alpha) = \mu \quad \text{and} \quad \lim_{\alpha \rightarrow +\infty} \sigma(\alpha) = \gamma.$$

Moreover, since σ is strictly monotonically increasing it holds for every $\beta \in \mathbb{R}$, $\sigma(\beta) > \mu$. Pick some $\beta \in \mathbb{R}$. For $\alpha \in \mathbb{R}$, we define $w_j = -\alpha a$, $v_j = \alpha z_j^T a + \beta$ for every $j \in [N - 1]$. Note that the matrix A changes as we vary α . Thus, we consider a family of matrices $A(\alpha)$ defined as

$A(\alpha)_{ij} = \sigma(z_i^T w_j + v_j) = \sigma(\alpha(z_j - z_i)^T a + \beta)$. Then it holds for every $i \in [N], j \in [N - 1]$

$$\lim_{\alpha \rightarrow +\infty} A(\alpha)_{ij} = \begin{cases} \gamma & j > i \\ \sigma(\beta) & j = i \\ \mu & j < i \end{cases}$$

Let $E(\alpha) = [\mathbf{1}_N, A(\alpha)]$ then it holds

$$\lim_{\alpha \rightarrow +\infty} E(\alpha)_{ij} = \begin{cases} 1 & j = 1 \\ \mu & j \in [2, N], i \geq j \\ \sigma(\beta) & j = i + 1 \\ \gamma & \text{else} \end{cases}$$

Let $\hat{E}(\alpha)$ be a modified matrix where one subtracts every row i by row $(i - 1)$ of $E(\alpha)$, in particular, let

$$\hat{E}(\alpha)_{ij} = \begin{cases} E(\alpha)_{ij} & i = 1, j \in [N] \\ E(\alpha)_{ij} - E(\alpha)_{i-1,j} & i > 1, j \in [N] \end{cases} \quad (3)$$

then it holds

$$\lim_{\alpha \rightarrow +\infty} \hat{E}(\alpha)_{ij} = \begin{cases} 1 & i = j = 1 \\ \mu - \sigma(\beta) < 0 & i = j > 1 \\ 0 & i > j \end{cases}$$

We do not show the values of other entries as what matters is that the limit, $\lim_{\alpha \rightarrow +\infty} \hat{E}(\alpha)$, is an upper triangular matrix.

Thus, the determinant is equal to the product of its diagonal entries which is non-zero. Note that the determinant of $\hat{E}(\alpha)$ is the same as that of $E(\alpha)$ as subtraction of some row from some other row does not change the determinant, and thus we get that $\lim_{\alpha \rightarrow +\infty} E(\alpha)$ has full rank N . As the determinant of $E(\alpha)$ is a polynomial of its entries and thus continuous in α , there exists $\alpha_0 \in \mathbb{R}$ s.t. for every $\alpha \geq \alpha_0$ it holds $\text{rank}(E(\alpha)) = \text{rank}([\mathbf{1}_N, A(\alpha)]) = N$. Moreover, since A is chosen as the first $(N - 1)$ columns of F_k , one can always choose the weights of the first $(N - 1)$ hidden units of layer k so that $\text{rank}([F_k, \mathbf{1}_N]) = N$.

In the case where the activation function fulfills $|\sigma(t)| \leq \rho_1 e^{\rho_2 t}$ for $t < 0$ and $|\sigma(t)| \leq \rho_3 t + \rho_4$ for $t \geq 0$ we consider directly the determinant of the matrix $E(\alpha)$. In particular, let us pick some $\beta \in \mathbb{R}$ such that $\sigma(\beta) \neq 0$. We consider the family of matrices $A(\alpha)$ defined as $A(\alpha)_{ij} = \sigma(z_i^T w_j + v_j) = \sigma(\alpha(z_j - z_i)^T a + \beta)$ where $w_j = -\alpha a$, $v_j = \alpha z_j^T a + \beta$ for every $j \in [N - 1]$. Let $E(\alpha) = [A(\alpha), \mathbf{1}_N]$. Note that the all-ones vector is now situated at the last column of $E(\alpha)$ instead of first column as before. This column re-ordering does not change the rank of $E(\alpha)$. By the Leibniz-formula one has

$$\det(E(\alpha)) = \sum_{\pi \in S_N} \text{sign}(\pi) \prod_{j=1}^{N-1} E(\alpha)_{\pi(j)j},$$

where S_N is the set of all $N!$ permutations of the set $\{1, \dots, N\}$ and we used the fact that the last column of $E(\alpha)$ is equal to the all ones vector. Define the permutation γ as $\gamma(j) = j$ for $j \in [N]$. Then we have

$$\begin{aligned} & \det(E(\alpha)) \\ &= \text{sign}(\gamma)\sigma(\beta)^{N-1} + \sum_{\pi \in S_N \setminus \{\gamma\}} \text{sign}(\pi) \prod_{j=1}^{N-1} E(\alpha)_{\pi(j)j}. \end{aligned}$$

The idea now is to show that $\prod_{j=1}^{N-1} E(\alpha)_{\pi(j)j}$ goes to zero for every permutation $\pi \neq \gamma$ as α goes to infinity. And since the whole summation goes to zero while $\sigma(\beta) \neq 0$, the determinant would be non-zero as desired. With that, we first note that for any permutation $\pi \neq \gamma$ there has to be at least one component $\pi(j)$ where $\pi(j) > j$, in which case, $\delta_j = (z_j - z_{\pi(j)})^T a < 0$ and thus for sufficiently large α , it holds $\alpha\delta_j + \beta < 0$. Thus

$$|E(\alpha)_{\pi(j)j}| = |\sigma(\alpha(z_j - z_{\pi(j)})^T a + \beta)| \leq \rho_1 e^{\rho_2 \beta} e^{-\alpha \rho_2 |\delta_j|}.$$

If $\pi(j) = j$ then $E(\alpha)_{\pi(j)j} = \sigma(\beta)$. In cases where $\pi(j) < j$ ($j \neq N$) it holds that $\delta_j = (z_j - z_{\pi(j)})^T a > 0$ and thus for sufficiently large α , it holds $\alpha\delta_j + \beta > 0$ and we have

$$|E(\alpha)_{\pi(j)j}| = |\sigma(\alpha(z_j - z_{\pi(j)})^T a + \beta)| \leq \rho_3 \delta_j \alpha + \rho_3 \beta + \rho_4.$$

So far, we have shown that $|E(\alpha)_{\pi(j)j}|$ can always be upper-bounded by an exponential function resp. affine function of α when $\pi(j) > j$ resp. $\pi(j) < j$ or it is just a constant when $\pi(j) = j$. The above observations imply that there exist positive constants P, Q, R, S, T such that it holds for every $\pi \in S_N \setminus \{\gamma\}$,

$$\left| \prod_{j=1}^{N-1} E(\alpha)_{\pi(j)j} \right| \leq R(P\alpha + Q)^S e^{-\alpha T}.$$

As $\alpha \rightarrow \infty$ the upper bound goes to zero. As there are only finitely many such terms, we get

$$\lim_{\alpha \rightarrow \infty} \det(E(\alpha)) = \text{sign}(\gamma)\sigma(\beta)^{N-1} \neq 0,$$

and thus with the same argument as before we can argue that there exists a finite α_0 for which $E(\alpha)$ has full rank. \square

Now we combine the previous lemma with Lemma 4.2 to conclude the following.

Lemma 4.4 *If the Assumptions 3.2 hold and $n_k \geq N - 1$ for some $k \in [L - 1]$ then the set $S := \left\{ (W_l, b_l)_{l=1}^k \mid \text{rank}([F_k, \mathbf{1}_N]) < N \right\}$ has Lebesgue measure zero.*

Proof: Let $E_k = [F_k, \mathbf{1}_N] \in \mathbb{R}^{N \times (n_k + 1)}$. Note that with Lemma 4.1 the output F_k of layer k is an analytic function of the network parameters on \mathcal{P} . The set of low rank

matrices E_k can be characterized by a system of equations such that the $\binom{n_k + 1}{N}$ determinants of all $N \times N$ submatrices of E_k are zero. As the determinant is a polynomial in the entries of the matrix and thus an analytic function of the entries and composition of analytic functions are again analytic, we conclude that each determinant is an analytic function of the network parameters of the first k layers. By Lemma 4.3 there exists at least one set of network parameters of the first k layers such that one of these determinant functions is not identically zero and thus by Lemma 4.2 the set of network parameters where this determinant is zero has measure zero. But as all submatrices need to have low rank in order that $\text{rank}([F_k, \mathbf{1}_N]) < N$, it follows that the set of network parameters where $\text{rank}([F_k, \mathbf{1}_N]) < N$ has measure zero. \square

We conclude that for $n_k \geq N - 1$ even if there are network parameters such that $\text{rank}([F_k, \mathbf{1}_N]) < N$, then every neighborhood of these parameters contains network parameters such that $\text{rank}([F_k, \mathbf{1}_N]) = N$.

Corollary 4.5 *If the Assumptions 3.2 hold and $n_k \geq N - 1$ for some $k \in [L - 1]$, then for any given $(W_l^0, b_l^0)_{l=1}^k$ and for every $\epsilon > 0$, there exists at least one $(W_l, b_l)_{l=1}^k \in B\left((W_l^0, b_l^0)_{l=1}^k, \epsilon\right)$ s.t. $\text{rank}([F_k, \mathbf{1}_N]) = N$.*

Proof: Let $S := \left\{ (W_l, b_l)_{l=1}^k \mid \text{rank}([F_k, \mathbf{1}_N]) < N \right\}$.

The ball $B\left((W_l, b_l)_{l=1}^k, \epsilon\right)$ has positive Lebesgue measure while S has measure zero due to Lemma 4.4. Thus, for every $(W_l, b_l)_{l=1}^k \in B\left((W_l^0, b_l^0)_{l=1}^k, \epsilon\right) \setminus S$ it holds $\text{rank}([F_k, \mathbf{1}_N]) = N$. \square

The final proof of our main Theorem 3.8 is heavily based on the implicit function theorem, see e.g. (Marsden, 1974).

Theorem 4.6 *Let $\Psi : \mathbb{R}^s \times \mathbb{R}^t \rightarrow \mathbb{R}^t$ be a continuously differentiable function. Suppose $(u_0, v_0) \in \mathbb{R}^s \times \mathbb{R}^t$ and $\Psi(u_0, v_0) = 0$. If the Jacobian matrix w.r.t. v ,*

$$J_v \Psi(u_0, v_0) = \begin{bmatrix} \frac{\partial \Psi_1}{\partial v_1} & \dots & \frac{\partial \Psi_1}{\partial v_t} \\ \vdots & & \vdots \\ \frac{\partial \Psi_t}{\partial v_1} & \dots & \frac{\partial \Psi_t}{\partial v_t} \end{bmatrix} \in \mathbb{R}^{t \times t}$$

is non-singular at (u_0, v_0) , then there is an open ball $B(u_0, \epsilon)$ for some $\epsilon > 0$ and a unique function $\alpha : B(u_0, \epsilon) \rightarrow \mathbb{R}^t$ such that $\Psi(u, \alpha(u)) = 0$ for all $u \in B(u_0, \epsilon)$. Furthermore, α is continuously differentiable.

With all the intermediate results proven above, we are finally ready for the proof of the main result.

Proof of Theorem 3.8 for case $\mathcal{I} = \{k + 1, \dots, L\}$

Let us divide the set of all parameters of the network

into two subsets where one corresponds to all parameters of all layers up to k , for that we denote $u = [\text{vec}(W_1)^T, b_1^T, \dots, \text{vec}(W_k)^T, b_k^T]^T$, and the other corresponds to the remaining parameters, for that we denote $v = [\text{vec}(W_{k+1})^T, b_{k+1}^T, \dots, \text{vec}(W_L)^T, b_L^T]^T$. By abuse of notation, we write $\Phi(u, v)$ to denote $\Phi((W_l, b_l)_{l=1}^L)$. Let $s = \dim(u)$, $t = \dim(v)$ and $(u^*, v^*) \in \mathbb{R}^s \times \mathbb{R}^t$ be the corresponding vectors for the critical point $(W_l^*, b_l^*)_{l=1}^L$. Let $\Psi : \mathbb{R}^s \times \mathbb{R}^t \rightarrow \mathbb{R}^t$ be a map defined as $\Psi(u, v) = \nabla_v \Phi(u, v) \in \mathbb{R}^t$, which is the gradient mapping of Φ w.r.t. all parameters of the upper layers from $(k+1)$ to L . Since the gradient vanishes at a critical point, it holds that $\Psi(u^*, v^*) = \nabla_v \Phi(u^*, v^*) = 0$. The Jacobian of Ψ w.r.t. v is the principal submatrix of the Hessian of Φ w.r.t. v , that is, $J_v \Psi(u, v) = \nabla_v^2 \Phi(u, v) \in \mathbb{R}^{t \times t}$. As the critical point is assumed to be non-degenerate with respect to v , it holds that $J_v \Psi(u^*, v^*) = \nabla_v^2 \Phi(u^*, v^*)$ is non-singular. Moreover, Ψ is continuously differentiable since $\Phi \in C^2(\mathcal{P})$ due to Assumption 3.2. Therefore, Ψ and (u^*, v^*) satisfy the conditions of the implicit function theorem 4.6. Thus there exists an open ball $B(u^*, \delta_1) \subset \mathbb{R}^s$ for some $\delta_1 > 0$ and a continuously differentiable function $\alpha : B(u^*, \delta_1) \rightarrow \mathbb{R}^t$ such that

$$\begin{cases} \Psi(u, \alpha(u)) = 0, & \forall u \in B(u^*, \delta_1) \\ \alpha(u^*) = v^* \end{cases}$$

By assumption we have $\text{rank}(W_l^*) = n_l, l \in [k+2, L]$, that is the weight matrices of the ‘‘upper’’ layers have full column rank. Note that $(W_l^*)_{l=k+2}^L$ corresponds to the weight matrix part of v^* where one leaves out W_{k+1}^* . Thus there exists a sufficiently small ϵ such that for any $v \in B(v^*, \epsilon)$, the weight matrix part $(W_l)_{l=k+2}^L$ of v has full column rank. In particular, this, combined with the continuity of α , implies that for a potentially smaller $0 < \delta_2 \leq \delta_1$, it holds for all $u \in B(u^*, \delta_2)$ that

$$\Psi(u, \alpha(u)) = 0, \quad \alpha(u^*) = v^*,$$

and that the weight matrix part $(W_l)_{l=k+2}^L$ of $\alpha(u) \in \mathbb{R}^t$ has full column rank.

Now, by Corollary 4.5 for any $0 < \delta_3 \leq \delta_2$ there exists a $\tilde{u} \in B(u^*, \delta_3)$ such that the generated output matrix \tilde{F}_k at layer k of the corresponding network parameters of \tilde{u} satisfies $\text{rank}([\tilde{F}_k, \mathbf{1}_N]) = N$. Moreover, it holds for $\tilde{v} = \alpha(\tilde{u})$ that $\Psi(\tilde{u}, \tilde{v}) = 0$ and the weight matrix part $(\tilde{W}_l)_{l=k+2}^L$ of \tilde{v} has full column rank. Assume (\tilde{u}, \tilde{v}) corresponds to the following representation

$$\begin{cases} \tilde{u} = [\text{vec}(\tilde{W}_1)^T, \tilde{b}_1^T, \dots, \text{vec}(\tilde{W}_k)^T, \tilde{b}_k^T]^T \in \mathbb{R}^s \\ \tilde{v} = [\text{vec}(\tilde{W}_{k+1})^T, \tilde{b}_{k+1}^T, \dots, \text{vec}(\tilde{W}_L)^T, \tilde{b}_L^T]^T \in \mathbb{R}^t \end{cases}$$

We obtain the following

$$\begin{cases} \Psi(\tilde{u}, \tilde{v}) = 0 \Rightarrow \nabla_{W_{k+1}} \Phi((\tilde{W}_l, \tilde{b}_l)_{l=1}^k) = 0 \\ \Psi(\tilde{u}, \tilde{v}) = 0 \Rightarrow \nabla_{b_{k+1}} \Phi((\tilde{W}_l, \tilde{b}_l)_{l=1}^k) = 0 \\ \text{rank}(\tilde{W}_l) = n_l, \forall l \in [k+2, L] \\ \text{rank}([\tilde{F}_k, \mathbf{1}_N]) = N \end{cases}$$

Thus, Lemma 3.5 implies that $(\tilde{W}_l, \tilde{b}_l)_{l=1}^L$ is a global minimum of Φ . Let $p^* = \Phi((\tilde{W}_l, \tilde{b}_l)_{l=1}^L) = \Phi(\tilde{u}, \tilde{v})$. Note that this construction can be done for any $\delta_3 \in (0, \delta_2]$. In particular, let $(\gamma_r)_{r=1}^\infty$ be a strictly monotonically decreasing sequence such that $\gamma_1 = \delta_3$ and $\lim_{r \rightarrow \infty} \gamma_r = 0$. By Corollary 4.5 and the previous argument, we can choose for any $\gamma_r > 0$ a point $\tilde{u}_r \in B(u^*, \gamma_r)$ such that $\tilde{v}_r = \alpha(\tilde{u}_r)$ has full rank and $\Phi(\tilde{u}_r, \tilde{v}_r) = p^*$. Moreover, as $\lim_{r \rightarrow \infty} \gamma_r = 0$, it follows that $\lim_{r \rightarrow \infty} \tilde{u}_r = u^*$ and as α is a continuous function, it holds with $\tilde{v}_r = \alpha(\tilde{u}_r)$ that $\lim_{r \rightarrow \infty} \tilde{v}_r = \lim_{r \rightarrow \infty} \alpha(\tilde{u}_r) = \alpha(\lim_{r \rightarrow \infty} \tilde{u}_r) = \alpha(u^*) = v^*$. Thus we get $\lim_{r \rightarrow \infty} (\tilde{u}_r, \tilde{v}_r) = (u^*, v^*)$ and as Φ is a continuous function it holds

$$\lim_{r \rightarrow \infty} \Phi((\tilde{u}_r, \tilde{v}_r)) = \Phi(u^*, v^*) = p^*,$$

as Φ attains the global minimum for the whole sequence $(\tilde{u}_r, \tilde{v}_r)$.

Proof of Theorem 3.8 for general case

In the general case $\mathcal{I} \subseteq \{k+1, \dots, L\}$, the previous proof can be easily adapted. The idea is that we fix all layers in $\{k+1, \dots, L\} \setminus \mathcal{I}$. In particular, let

$$\begin{cases} u = [\text{vec}(W_1)^T, b_1^T, \dots, \text{vec}(W_k)^T, b_k^T]^T \\ v = [\text{vec}(W_{\mathcal{I}(1)})^T, b_{\mathcal{I}(1)}^T, \dots, \text{vec}(W_{\mathcal{I}(|\mathcal{I}|)})^T, b_{\mathcal{I}(|\mathcal{I}|)}^T]^T. \end{cases}$$

Let $s = \dim(u)$, $t = \dim(v)$ and $(u^*, v^*) \in \mathbb{R}^s \times \mathbb{R}^t$ be the corresponding vectors at $(W_l^*, b_l^*)_{l=1}^L$. Let $\Psi : \mathbb{R}^s \times \mathbb{R}^t \rightarrow \mathbb{R}^t$ be a map defined as $\Psi(u, v) = \nabla_v \Phi((W_l, b_l)_{l=1}^L)$ with $\Psi(u^*, v^*) = \nabla_v \Phi((W_l^*, b_l^*)_{l=1}^L) = 0$.

The only difference is that all the layers from $\{k+1, \dots, L\} \setminus \mathcal{I}$ are hold fixed. They are not contained in the arguments of Ψ , thus will not be involved in our perturbation analysis. In this way, the full rank property of the weight matrices of these layers are preserved, which is needed to obtain the global minimum.

5. Relaxing the Condition on the Number of Hidden Units

We have seen that $n_k \geq N - 1$ is a sufficient condition which leads to a rather simple structure of the critical points, in the sense that all local minima which have full

rank in the layers $k + 2$ to L and for which the Hessian is non-degenerate on any subset of upper layers that includes layer $k + 1$ are automatically globally optimal. This suggests that suboptimal locally optimal points are either completely absent or relatively rare. We have motivated before that networks with a certain wide layer are used in practice, which shows that the condition $n_k \geq N - 1$ is not completely unrealistic. On the other hand we want to discuss in this section how it could be potentially relaxed. The following result will provide some intuition about the case $n_k < N - 1$, but will not be as strong as our main result 3.8 which makes statements about a large class of critical points. The main idea is that with the condition $n_k \geq N - 1$ the data is linearly separable at layer k . As modern neural networks are expressive enough to represent any function, see (Zhang et al., 2017) for an interesting discussion on this, one can expect that in some layer the training data becomes linearly separable. We prove that any critical point, for which the “learned” network outputs at any layer are linearly separable (see Definition 5.1) is a global minimum of the training error.

Definition 5.1 (Linearly separable vectors) *A set of vectors $(x_i)_{i=1}^N \in \mathbb{R}^d$ from m classes $(C_j)_{j=1}^m$ is called linearly separable if there exist m vectors $(a_j)_{j=1}^m \in \mathbb{R}^d$ and m scalars $(b_j)_{j=1}^m \in \mathbb{R}$ so that $a_j^T x_i + b_j > 0$ for $x_i \in C_j$ and $a_j^T x_i + b_j < 0$ for $x_i \notin C_j$ for every $i \in [N], j \in [m]$.*

In this section, we use a slightly different loss function than in the previous section. The reason is that the standard least squares loss is not necessarily small when the data is linearly separable. Let C_1, \dots, C_m denote m classes. We consider the objective function $\Phi : \mathcal{P} \rightarrow \mathbb{R}$ from (1)

$$\Phi\left((W_l, b_l)_{l=1}^L\right) = \sum_{i=1}^N \sum_{j=1}^m l(f_{Lj}(x_i) - y_{ij}) \quad (4)$$

where the loss function now takes the new form

$$l(f_{Lj}(x_i) - y_{ij}) = \begin{cases} l_1(f_{Lj}(x_i) - y_{ij}) & x_i \in C_j \\ l_2(f_{Lj}(x_i) - y_{ij}) & x_i \notin C_j \end{cases}$$

where l_1, l_2 penalize the deviation from the label encoding for the true class resp. wrong classes. We assume that the minimum of Φ is attained over \mathcal{P} . Note that Φ is bounded from below by zero as l_1 and l_2 are non-negative loss functions. The results of this section are made under the following assumptions on the activation and loss function.

Assumptions 5.2 1. $\sigma \in C^1(\mathbb{R})$ and strictly monotonically increasing.

2. $l_1 : \mathbb{R} \rightarrow \mathbb{R}_+, l_1 \in C^1, l_1(a) = 0 \Leftrightarrow a \geq 0, l_1'(a) = 0 \Leftrightarrow a \geq 0$ and $l_1'(a) < 0 \forall a < 0$




Figure 2. An example of l_1, l_2 .

3. $l_2 : \mathbb{R} \rightarrow \mathbb{R}_+, l_2 \in C^1, l_2(a) = 0 \Leftrightarrow a \leq 0, l_2'(a) = 0 \Leftrightarrow a \leq 0$ and $l_2'(a) > 0 \forall a > 0$

In classification tasks, this loss function encourages higher values for the true class and lower values for wrong classes. An example of the loss function that satisfies Assumption 5.2 is given as (see Figure 2):

$$l_1(a) = \begin{cases} a^2 & a \leq 0 \\ 0 & a \geq 0 \end{cases} \quad l_2(a) = \begin{cases} 0 & a \leq 0 \\ a^2 & a \geq 0 \end{cases}$$

Note that for a $\{+1, -1\}$ -label encoding, $+1$ for the true class and -1 for all wrong classes, one can rewrite (4) as

$$\Phi\left((W_l, b_l)_{l=1}^L\right) = \sum_{i=1}^N \sum_{j=1}^m \max\{0, 1 - y_{ij} f_{Lj}(x_i)\}^2,$$

which is similar to the truncated squared loss (also called squared hinge loss) used in the SVM for binary classification. Since σ and l are continuously differentiable, all the results from Lemma 2.1 still hold.

Our main result in this section is stated as follows.

Theorem 5.3 *Let $\Phi : \mathcal{P} \rightarrow \mathbb{R}_+$ be defined as in (4) and let the Assumptions 5.2 hold. Then it follows:*

1. *Every critical point of Φ for which the feature vectors contained in the rows of F_k are linearly separable and all the weight matrices $(W_l)_{l=k+2}^L$ have full column rank is a global minimum.*
2. *If the training inputs are linearly separable then every critical point of Φ for which all the weight matrices $(W_l)_{l=2}^L$ have full column rank is a global minimum.*

Proof:

1. Let $\tilde{F}_k = [F_k, \mathbf{1}_N]$. Since F_k contains linearly separable feature vectors, there exists m vectors $h_1, \dots, h_m \in \mathbb{R}^{n_k+1}$ s.t. $\langle h_j, (\tilde{F}_k)_{i \cdot} \rangle > 0$ for $x_i \in C_j$ and $\langle h_j, (\tilde{F}_k)_{i \cdot} \rangle < 0$ for $x_i \notin C_j$. Let $H = [h_1, \dots, h_m] \in \mathbb{R}^{(n_k+1) \times m}$, one obtains

$$(H^T \tilde{F}_k^T)_{ji} = \langle h_j, (\tilde{F}_k)_{i \cdot} \rangle \begin{cases} > 0 & x_i \in C_j \\ < 0 & x_i \notin C_j \end{cases}.$$

On the other hand,

$$\begin{aligned} (\Delta_L)_{ij} &= \delta_{Lj}(x_i) \\ &= \frac{\partial \Phi}{\partial g_{Lj}(x_i)} \\ &= \begin{cases} l'_1(f_{Lj}(x_i) - y_{ij})\sigma'(g_{Lj}(x_i)) & x_i \in C_j \\ l'_2(f_{Lj}(x_i) - y_{ij})\sigma'(g_{Lj}(x_i)) & x_i \notin C_j \end{cases} \end{aligned}$$

We show that $H^T \tilde{F}_k^T \Delta_L = 0$ if and only if $\Delta_L = 0$. Indeed, if $\Delta_L = 0$ the implication is trivial. For the other direction, assume that $H^T \tilde{F}_k^T \Delta_L = 0$. Then it holds for every $j \in [m]$ that $0 = (H^T \tilde{F}_k^T \Delta_L)_{jj} = \sum_{i=1}^N (H^T \tilde{F}_k^T)_{ji} (\Delta_L)_{ij}$. In particular,

$$\begin{aligned} & \sum_{i=1}^N (H^T \tilde{F}_k^T)_{ji} (\Delta_L)_{ij} \\ &= \sum_{x_i \in C_j} \langle (\tilde{F}_k)_{i,:}, h_j \rangle l'_1(f_{Lj}(x_i) - y_{ij})\sigma'(g_{Lj}(x_i)) \\ &+ \sum_{x_i \notin C_j} \langle (\tilde{F}_k)_{i,:}, h_j \rangle l'_2(f_{Lj}(x_i) - y_{ij})\sigma'(g_{Lj}(x_i)) \leq 0 \end{aligned}$$

Note that under the assumptions on the loss and activation function and since the features are separable, the terms in both sums are non-positive and thus the sum can only vanish if all terms vanish which implies

$$\begin{cases} l'_1(f_{Lj}(x_i) - y_{ij}) = 0 & x_i \in C_j \\ l'_2(f_{Lj}(x_i) - y_{ij}) = 0 & x_i \notin C_j \end{cases} \forall i \in [N], j \in [m] \quad (5)$$

which yields $\Delta_L = 0$.

Back to the main proof, the idea is to prove that $\Delta_L = 0$ at the given critical point. Let us assume for the sake of contradiction that $\Delta_L \neq 0$. For every $l \in [L]$, $i_l \in [n_l]$ define $\Sigma_{i_l}^l = \text{diag}(\sigma'(g_{li}(x_1)), \dots, \sigma'(g_{li}(x_N)))$. Since the cumulative product $\prod_{l=k+1}^{L-1} \Sigma_{i_l}^l$ is a $N \times N$ diagonal matrix which contains only positive entries in its diagonal, it does not change the sign pattern of Δ_L , and thus it holds with, $H^T \tilde{F}_k^T \Delta_L = 0$ if and only if $\Delta_L = 0$, for every $(i_{k+1}, \dots, i_{L-1}) \in [n_{k+1}] \times \dots \times [n_{L-1}]$ that

$$0 \neq H^T \tilde{F}_k^T \left(\prod_{l=k+1}^{L-1} \Sigma_{i_l}^l \right) \Delta_L \quad (6)$$

$$0 \neq H^T \tilde{F}_k^T \left(\prod_{l=k+1}^{L-1} \Sigma_{i_l}^l \right) \Delta_L W_L^T \quad (\text{rank}(W_L) = n_L),$$

where the last inequality is implied by (6) as W_L has full column rank $n_L = m$. Since the above product

of matrices is a non-zero matrix, there must exist a non-zero column, say $p \in [n_{L-1}]$, then

$$0 \neq H^T \tilde{F}_k^T \left(\prod_{l=k+1}^{L-1} \Sigma_{i_l}^l \right) (\Delta_L W_L^T)_{:p}$$

Since i_{L-1} is arbitrary, pick $i_{L-1} = p$ one obtains

$$0 \neq H^T \tilde{F}_k^T \left(\prod_{l=k+1}^{L-2} \Sigma_{i_l}^l \right) \underbrace{\Sigma_p^{L-1} (\Delta_L W_L^T)_{:p}}_{(\Delta_{L-1})_{:p}} \quad (7)$$

Moreover, it holds for every $i \in [N]$

$$\begin{aligned} & (\Sigma_p^{L-1} (\Delta_L W_L^T)_{:p})_i \\ &= \sigma'(g_{(L-1)p}(x_i)) \sum_{j=1}^{n_L} \delta_{Lj}(x_i) (W_L)_{pj} \\ &= \delta_{(L-1)p}(x_i) \\ &= (\Delta_{L-1})_{ip} \end{aligned}$$

and thus from (7),

$$\begin{aligned} 0 &\neq H^T \tilde{F}_k^T \left(\prod_{l=k+1}^{L-2} \Sigma_{i_l}^l \right) (\Delta_{L-1})_{:p} \\ &\Rightarrow 0 \neq H^T \tilde{F}_k^T \left(\prod_{l=k+1}^{L-2} \Sigma_{i_l}^l \right) \Delta_{L-1} \end{aligned}$$

Compared to (6), we have reduced the product from $\prod_{l=k+1}^{L-1}$ to $\prod_{l=k+1}^{L-2}$. By induction, one can easily show that

$$0 \neq H^T \tilde{F}_k^T \Sigma_{i_{k+1}}^{k+1} \Delta_{k+2}$$

and hence $0 \neq H^T \tilde{F}_k^T \Delta_{k+1}$, which implies $0 \neq \tilde{F}_k^T \Delta_{k+1} = [(\nabla_{W_{k+1}})^T \Phi, \nabla_{b_{k+1}} \Phi]^T$. However, this is a contradiction to the fact that we assumed that $(W_l, b_l)_{l=1}^L$ is a critical point. Thus it follows that it has to hold $\Delta_L = 0$. As $\Delta_L = 0$ it holds (5) which implies

$$\begin{cases} f_{Lj}(x_i) \geq y_{ij} & x_i \in C_j \\ f_{Lj}(x_i) \leq y_{ij} & x_i \notin C_j \end{cases} \forall i \in [N], j \in [m].$$

This in turn implies $\Phi((W_l, b_l)_{l=1}^L) = 0$. Thus the critical point $(W_l, b_l)_{l=1}^L$ is a global minimum.

2. This can be seen as a special case of the first statement. In particular, assume one has a zero-layer which coincides with the training inputs, namely $F_0 = X$, then the result follows immediately.

□

Note that the second statement of Theorem 5.3 can be considered as a special case of the first statement. In the case where $L = 2$ and training inputs are linearly separable, the second statement of our Theorem 5.3 recovers the similar result of (Gori & Tesi, 1992; Frasconi et al., 1997) for one-hidden layer networks.

Even though the assumptions of Theorem 3.4 and Theorem 5.3 are different in terms of class of activation and loss functions, their results are related. In fact, it is well known that if a set of vectors is linearly independent then they are linearly separable, see e.g. p.340 (Barber, 2012). Thus Theorem 5.3 can be seen as a direct generalization of Theorem 3.4. The caveat, which is also the main difference to Theorem 3.8, is that Theorem 5.3 makes only statements for all the critical points for which the problem has become separable at some layer, whereas there is no such condition in Theorem 3.8. However, we still think that the result is of practical relevance, as one can expect for a sufficiently large network that stochastic gradient descent will lead to a network structure where the data becomes separable at a particular layer. When this happens all the associated critical points are globally optimal. It is an interesting question for further research if one can show directly under some architecture condition that the network outputs become linearly separable at some layer for any local minimum and thus every local minimum is a global minimum.

6. Discussion

Our results show that the loss surface becomes well-behaved when there is a wide layer in the network. Implicitly, such a wide layer is often present in convolutional neural networks used in computer vision. It is thus an interesting future research question how and if our result can be generalized to neural networks with sparse connectivity. We think that the results presented in this paper are a significant addition to the recent understanding why deep learning works so efficiently. In particular, since in this paper we are directly working with the neural networks used in practice without any modifications or simplifications.

Acknowledgment

The authors acknowledge support by the ERC starting grant NOLEPRO 307793.

References

Auer, P., Herbster, M., and Warmuth, M. K. Exponentially many local minima for single neurons. In *NIPS*, 1996.

Ba, J. and Caruana, R. Do deep nets really need to be deep?

In *NIPS*, 2014.

Baldi, P. and Hornik, K. Neural networks and principle component analysis: Learning from examples without local minima. *Neural Networks*, 2:53–58, 1988.

Barber, D. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, Berlin, 2012.

Brutzkus, A. and Globerson, A. Globally optimal gradient descent for a convnet with gaussian inputs, 2017. arXiv:1702.07966.

Caruana, R., Lawrence, S., and Giles, L. Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping. In *NIPS*, 2001.

Choromanska, A., Hena, M., Mathieu, M., Arous, G. B., and LeCun, Y. The loss surfaces of multilayer networks. In *AISTATS*, 2015a.

Choromanska, A., LeCun, Y., and Arous, G. B. Open problem: The landscape of the loss surfaces of multilayer networks. *JMLR*, 2015b.

Ciresan, D. C., Meier, U., Gambardella, L. M., and Schmidhuber, J. Deep, big, simple neural nets for handwritten digit recognition. *Neural Computation*, 22: 3207–3220, 2010.

Dauphin, Y., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *NIPS*, 2014.

Frasconi, P., Gori, M., and Tesi, A. Successes and failures of backpropagation: A theoretical investigation. *Progress in Neural Networks: Architecture*, 1997.

Gautier, A., Nguyen, Q., and Hein, M. Globally optimal training of generalized polynomial neural networks with nonlinear spectral methods. In *NIPS*, 2016.

Goodfellow, I. J., Vinyals, O., and Saxe, A. M. Qualitatively characterizing neural network optimization problems. In *ICLR*, 2015.

Gori, M. and Tesi, A. On the problem of local minima in backpropagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:76–86, 1992.

Haeffele, B. D. and Vidal, R. Global optimality in tensor factorization, deep learning, and beyond, 2015. arXiv:1506.07540v1.

Hartley, R. and Zisserman, A. *Multiple view geometry in computer vision*. Cambridge University Press, New York, 2004.

- Janzamin, M., Sedghi, H., and Anandkumar, A. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv:1506.08473*, 2016.
- Kawaguchi, K. Deep learning without poor local minima. In *NIPS*, 2016.
- Krantz, S. G. and Parks, H. R. *A Primer of Real Analytic Functions*. Birkhäuser, Boston, second edition, 2002.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., and Jackel, L.D. Handwritten digit recognition with a back-propagation network. In *NIPS*, 1990.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*, 521, 2015.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. Gradient descent only converges to minimizers. In *COLT*, 2016.
- Lin, Z., Memisevic, R., and Konda, K. How far can we go without convolution: Improving fully-connected networks. preprint, *arXiv:1511.02580*, 2016.
- Marsden, J. E. *Elementary classical analysis*. W.H.Freeman and Company, 1974.
- Milnor, J. *Lectures on H-Cobordism Theorem*. Princeton University Press, Princeton, 1965.
- Mityagin, B. The zero set of a real analytic function, 2015. *arXiv:1512.07276*.
- Neyshabur, B., Salakhutdinov, R. R., and Srebro, N. Pathsgd: Path-normalized optimization in deep neural networks. In *NIPS*, 2015.
- Nguyen, V. D. Complex powers of analytic functions and meromorphic renormalization in qft, 2015. *arXiv:1503.00995*.
- Poggio, T. and Liao, Q. Theory ii: Landscape of the empirical risk in deep learning, 2017. *arXiv:1703.09833*.
- Rister, B. and Rubin, D. L. Piecewise convexity of artificial neural networks, 2017. *arXiv:1607.04917*.
- Safran, I. and Shamir, O. On the quality of the initial basin in overspecified networks. In *ICML*, 2016.
- Sagun, L., Bottou, L., and LeCun, Y. Singularity of the hessian in deep learning, 2016. *arXiv:1611.07476*.
- Sima, J. Training a single sigmoidal neuron is hard. *Neural Computation*, 14:2709–2728, 2002.
- Soltanolkotabi, M. Learning relus via gradient descent, 2017. *arXiv:1705.04591*.
- Soudry, D. and Hoffer, E. Exponentially vanishing sub-optimal local minima in multilayer neural networks, 2017. *arXiv:1702.05777*.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *JLMR*, 11:3371–3408, 2010.
- Yu, X. and Chen, G. On the local minima free condition of backpropagation learning. *IEEE Transaction on Neural Networks*, 6:1300–1303, 1995.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, Oriol. Understanding deep learning requires re-thinking generalization. In *ICLR*, 2017.
- Zhou, P. and Feng, J. The landscape of deep learning algorithms, 2017. *arXiv:1705.07038*.