

Gradient Descent Finds Global Minima for Generalizable Deep Neural Networks of Practical Sizes

Kenji Kawaguchi
MIT
Email: kawaguch@mit.edu

Jiaoyang Huang
Harvard University
Email: jiaoyang@math.harvard.edu

Abstract—In this paper, we theoretically prove that gradient descent can find a global minimum for nonlinear deep neural networks of sizes commonly encountered in practice. The theory developed in this paper requires only the number of trainable parameters to increase linearly as the number of training samples increases. This allows the size of the deep neural networks to be several orders of magnitude smaller than that required by the previous theories. Moreover, we prove that the linear increase of the size of the network is the optimal rate and that it cannot be improved, except by a logarithmic factor. Furthermore, deep neural networks with the trainability guarantee are shown to generalize well to unseen test samples with a natural dataset but not a random dataset.

I. INTRODUCTION

Deep neural networks have recently achieved significant empirical success in the fields of machine learning and its applications. Neural networks have been theoretically studied for a long time, dating back to the days of multilayer perceptron, with focus on the *expressivity* of shallow neural networks [1], [2], [3], [4], [5], [6]. More recently, the expressivity of neural networks was theoretically investigated for modern deep architectures with rectified linear units (ReLUs) [7], residual maps [8], and/or convolutional and max-pooling layers [9].

However, the expressivity of a neural network does not ensure its *trainability*. The expressivity of a neural network states that, given a training dataset, there exists an optimal parameter vector for the neural network to interpolate that given dataset. It does not guarantee that an algorithm will be able to find such an optimal vector, efficiently, during the training of neural networks. Indeed, finding the optimal vector for a neural network has been proven to be an NP-hard problem, in some cases [10], [11], [12].

Quite recently, it was proved in a series of papers that, if the size of a neural network is significantly larger than the size of the dataset, the (stochastic) gradient descent algorithm can find an optimal vector for shallow [13], [14], [15] and deep networks [16], [17], [18]. However, a considerable gap still exists between these trainability results and the expressivity theories; i.e., these trainability results require a significantly larger number of parameters, when compared to the expressivity theories. Table I summarizes the number of parameters required by each previous theory, in terms of

TABLE I: Number of parameters required to ensure the trainability, in terms of n , where n is the number of samples in a training dataset and H is the number of hidden layers.

Reference	# Parameters	Depth H	Trainability
[3], [4], [5]	$\tilde{\Omega}(n)$	1,2	No (expressivity only)
[8], [9], [7]	$\tilde{\Omega}(n)$	any H	No (expressivity only)
[13]	$\tilde{\Omega}(\text{poly}(n))$	1	Yes
[14]	$\tilde{\Omega}(n^6)$	1	Yes
[15]	$\tilde{\Omega}(n^2)$	1	Yes
[16], [18]	$\tilde{\Omega}(\text{poly}(n, H))$	any H	Yes
[17]	$\tilde{\Omega}(2^{O(H)}n^8)$	any H	Yes
this paper	$\tilde{\Omega}(n)$	any H	Yes

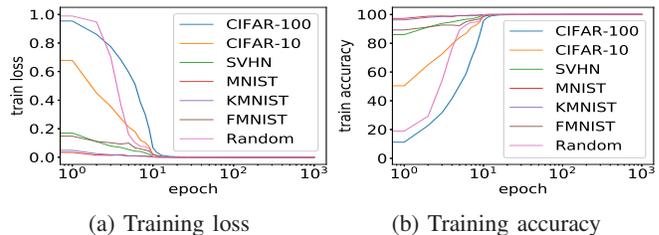


Fig. 1: Training loss and accuracy versus the number of epochs (in log scale) for pre-activation ResNet with 18 layers. Training accuracy reaches 100% (and training loss is approximately zero) for all datasets, even though the number of total parameters is several orders of magnitude smaller than that required by the previous theories.

the size n of the dataset, where the $\tilde{\Omega}(\cdot)$ notation ignores the logarithmic factors and the $\text{poly}(\cdot)$ notation hides the significantly large unknown polynomial dependencies: for example, $\text{poly}(n) \geq n^{60}$ in [16].

There is also a significant gap between the trainability theory and common practice. Typically, deep neural networks used in practical applications are trainable, and yet, much smaller than what the previous theories require to ensure trainability. Figure 1 illustrates this fact with various datasets and a pre-activation ResNet with 18 layers (PreActResNet18), which is widely used in practice. FMNIST represents the Fashion MNIST. RANDOM represents a randomly generated dataset of size 50000 (with the inputs being $3 \times 24 \times 24$ images of pixels drawn randomly from the standard normal distribution and the target being integer labels drawn uni-

Kenji Kawaguchi was supported in part by NSF grants (No.1523767 and 1723381), AFOSR grant (No. FA9550-17-1-0165), ONR grant (N00014-18-1-2847), Honda Research, and the MIT-Sensetime Alliance on AI.

formly from between 0 and 9). Here, the sizes of the training datasets vary from 50000 to 73257. For these datasets, the previous theories require at least $n^8 = (50000)^8$ parameters for the deep neural network to be trainable, which is several orders of magnitude larger than the number of parameters of PreActResNet18 (11169994 parameters) or even larger networks such as WideResNet18 (36479219 parameters).

In this paper, we aim to bridge these gaps by theoretically proving the upper and lower bounds for the number of parameters required to ensure trainability. In particular, we show that deep neural networks with $\tilde{\Omega}(n)$ parameters are efficiently trainable by using a gradient descent algorithm. That is, our theory only requires the number of total parameters to be in the order of n , which matches the practical observations. Moreover, we demonstrate that trainable deep neural networks of size $\tilde{\Omega}(n)$ are generalizable to unseen test points with a natural dataset, but not with a random dataset.

II. PRELIMINARIES

This paper studies feedforward neural networks with H hidden layers, where $H \geq 1$ is arbitrary. Given an input vector $x \in \mathbb{R}^{m_x}$ and a parameter vector θ , the output of the neural network is given by

$$f(x, \theta) = W^{(H+1)}x^{(H)} + b^{(H+1)} \in \mathbb{R}^{m_y}, \quad (1)$$

where $W^{(H+1)} \in \mathbb{R}^{m_y \times m_H}$ and $b^{(H+1)} \in \mathbb{R}^{m_y}$ are the weight matrix and bias term, respectively, of the output layer. The output of the last hidden layer $x^{(H)}$ is given by the set of recursive equations: $x^{(0)} = x$ and

$$x^{(l)} = \frac{1}{\sqrt{m_l}} \sigma(W^{(l)}x^{(l-1)} + b^{(l)}), \quad l = 1, 2, \dots, H, \quad (2)$$

where $W^{(l)} \in \mathbb{R}^{m_l \times m_{l-1}}$ is the weight matrix, $b^{(l)} \in \mathbb{R}^{m_l}$ is the bias term, and σ is the activation unit, which is applied coordinate-wise to its input. Here, $x^{(l)}$ is the output of the l -th layer, which has m_l neurons.

Then, the vector containing all trainable parameters is given by $\theta = (\text{vec}(\bar{W}^{(1)})^\top, \dots, \text{vec}(\bar{W}^{(H+1)})^\top)^\top$, where $\bar{W}^{(l)} = [W^{(l)}, b^{(l)}]$ and $\text{vec}(M)$ represents the standard vectorization of the matrix M . Thus, the total number of trainable parameters is

$$d = \sum_{l=0}^H (m_l m_{l+1} + m_{l+1}),$$

where $m_0 = m_x$ and $m_{H+1} = m_y$.

This paper analyzes the trainability in terms of the standard objective of empirical risk minimization:

$$J(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i, \theta), y_i),$$

where $\{(x_i, y_i)\}_{i=1}^n$ is a training dataset, y_i is the i -th target, and $\ell(\cdot, y_i)$ represents a loss criterion such as the squared loss or cross-entropy loss. The following assumptions are employed for the loss criterion $q \mapsto \ell(q, y_i)$ and activation unit $\sigma(x)$:

Assumption 1. (Use of common loss criteria) *For any $i \in \{1, \dots, n\}$, the function $\ell_i(q) = \ell(q, y_i) \in \mathbb{R}_{\geq 0}$ is*

differentiable and convex, and $\nabla \ell_i$ is ζ -Lipschitz (with the metric induced by the Euclidian norm $\|\cdot\|_2$).

Assumption 2. (Use of common activation units) *The activation function $\sigma(x)$ is real analytic, monotonically increasing, 1-Lipschitz, and the limits exist as: $\lim_{x \rightarrow -\infty} \sigma(x) = \sigma_- > -\infty$ and $\lim_{x \rightarrow +\infty} \sigma(x) = \sigma_+ \leq +\infty$.*

Assumption 1 is satisfied by simply using a common loss criterion such as the squared loss or cross-entropy loss. For example, $\zeta = 2$ for the squared loss, as $\|\nabla \ell_i(q) - \nabla \ell_i(q')\|_2 \leq 2\|q - q'\|_2$. The training objective function $J(\theta)$ is nonconvex in θ , even if the loss criterion $q \mapsto \ell(q, y_i)$ is convex in q .

Assumption 2 is satisfied by using common activation units such as sigmoid and hyperbolic tangents. Moreover, the softplus activation, which is defined as $\sigma_\alpha(x) = \ln(1 + \exp(\alpha x))/\alpha$, satisfies Assumption 2 with any hyperparameter $\alpha \in \mathbb{R}_{>0}$. The softplus activation can approximate the ReLU activation for any desired accuracy as

$$\sigma_\alpha(x) \rightarrow \text{relu}(x) \text{ as } \alpha \rightarrow \infty,$$

where relu represents the ReLU activation.

Throughout this paper, neural networks are initialized with random Gaussian weights, following the common initialization schemes used in practice. More precisely, the initial parameter vector θ^0 is randomly drawn as $(W_{ij}^{(l)})^0 \sim \mathcal{N}(0, c_w)$ and $(b^{(l)})^0 \sim \mathcal{N}(0, c_b)$, where c_w and c_b are constants and $(W_{ij}^{(l)})^0$ and $(b^{(l)})^0$ correspond to the initial vector θ^0 as $\theta^0 = (\text{vec}((\bar{W}^{(1)})^0)^\top, \dots, \text{vec}((\bar{W}^{(H+1)})^0)^\top)^\top$ with $(\bar{W}^{(l)})^0 = [(W^{(l)})^0, (b^{(l)})^0]$. With this random initialization scheme, the outputs are normalized properly as $\|x^{(l)}\|_2^2 = O(1)$ for $0 \leq l \leq H$, and $\|f(x, \theta)\|_2^2 = O(m_y)$ with high probability.

III. MAIN TRAINABILITY RESULTS

This section first introduces the formal definition of trainability, in terms of the number d of parameters, and then presents our main results for the trainability.

A. Problem formalization

The goal of this section is to formalize the question of trainability in terms of the number of parameters, d . Intuitively, given the dataset size n , depth H , and any $\delta > 0$, we define the *probable trainability* $\mathcal{P}_{n,H,\delta}$ as $\mathcal{P}_{n,H,\delta}(d) = \text{true}$ if having d parameters can ensure the trainability for all datasets with probability at least $1 - \delta$, and $\mathcal{P}_{n,H,\delta}(d) = \text{false}$ otherwise. We formalize this intuition as follows.

Let \mathcal{F}_d^H be the set of all neural network architectures $f(\cdot, \cdot)$ of the form in equation (1) with H hidden layers, at most d parameters, and activation units σ satisfying Assumption 2. Let \mathcal{S}_n be the set of all training datasets $S = \{(x_i, y_i)\}_{i=1}^n$ of size n such that the data points are normalized as $\|x_i\|_2^2 = 1$ and $y_i \in [-1, 1]^{m_y}$ for all $i \in \{1, \dots, n\}$. Let \mathcal{L}_S^ζ be a set of all loss functionals L such that for any $L \in \mathcal{L}_S^\zeta$, we have $L(g) = \frac{1}{n} \sum_{i=1}^n \ell(g(x_i), y_i)$ and $\text{argmin}_{g: \mathbb{R}^{m_x} \rightarrow \mathbb{R}^{m_y}} L(g) \neq \emptyset$, where $g: \mathbb{R}^{m_x} \rightarrow \mathbb{R}^{m_y}$ is a function, $S \in \mathcal{S}_n$ is a training dataset, and $q \mapsto \ell(q, y_i)$ is

a loss criterion satisfying Assumption 1. For any (θ, \bar{W}) , we define $\psi_i(\theta, \bar{W}) \in \mathbb{R}^d$ to be the parameter vector θ with the corresponding $\bar{W}^{(l)}$ entries replaced by \bar{W} . For example, $\psi_{H+1}(\theta, \bar{W}) = (\text{vec}(\bar{W}^{(1)})^\top, \dots, \text{vec}(\bar{W}^{(H)})^\top, \text{vec}(\bar{W})^\top)^\top$. We use the symbol \odot to represent the entrywise product (i.e., Hadamard product).

With these notations, we can now formalize the probable trainability $\mathcal{P}_{n,H,\delta}$, in terms of d , as follows:

Definition 1. $\mathcal{P}_{n,H,\delta} : \mathbb{N} \rightarrow \{\text{true}, \text{false}\}$ is a function such that $\mathcal{P}_{n,H,\delta}(d) = \text{true}$ if and only if the following statement holds true: $\forall \zeta > 0, \exists f \in \mathcal{F}_d^H, \exists \eta \in \mathbb{R}^d, \forall S \in \mathcal{S}_n, \forall L \in \mathcal{L}_S^\zeta, \exists c_\theta \in \mathbb{R}$, and $\forall \epsilon > 0$, with probability at least $1 - \delta$ (over randomly drawn initial weights θ^0), there exists $t = O(c_r \zeta / \epsilon)$ such that

$$J(\theta^t) = L(f(\cdot, \theta^t)) \leq L(f^*) + \epsilon, \quad (3)$$

and $\|\theta^t\|_2^2 \leq c_\theta$, where $f^* \in \text{argmin}_{g: \mathbb{R}^{m_x} \rightarrow \mathbb{R}^{m_y}} L(g)$ is a global minimum of the functional L , $(\theta^k)_{k \in \mathbb{N}}$ is the sequence generated by the gradient descent algorithm $\theta^{k+1} = \theta^k - \eta \odot \nabla J(\theta^k)$, and $c_r = \max_{l \in \{1, \dots, H+1\}} \inf_{\bar{W}^* \in \mathcal{W}_l^*} \|\bar{W}^* - (\bar{W}^{(l)})^0\|_F^2$ with $\mathcal{W}_l^* = \text{argmin}_{\bar{W}} L(f(\cdot, \psi_l(\theta^0, \bar{W})))$.

Here, $\mathcal{P}_{n,H,\delta}(d) = \text{true}$ implies that a gradient descent algorithm finds a global minimum of a deep neural network with d trainable parameters for any dataset (if a global minimum exists). To verify this, let $\tilde{\mathcal{P}}_{n,H,\delta}$ be equivalent to $\mathcal{P}_{n,H,\delta}$, except that inequality (3) is replaced by

$$L(f(\cdot, \theta^t)) \leq L(f(\cdot, \theta^*)) + \epsilon, \quad (4)$$

where $\theta^* \in \mathbb{R}^d$ is a global minimum of $J(\theta) = L(f(\cdot, \theta))$. As $L(f^*) \leq L(f(\cdot, \theta^*))$, $\mathcal{P}_{n,H,\delta}(d) = \text{true}$ implies that $\tilde{\mathcal{P}}_{n,H,\delta}(d) = \text{true}$, which is the desired statement.

The reason we use $\mathcal{P}_{n,H,\delta}(d)$ instead of $\tilde{\mathcal{P}}_{n,H,\delta}(d)$ is that $\tilde{\mathcal{P}}_{n,H,\delta}(d)$ admits trivial and unpreferred solutions in that a global minimum θ^* can have a large loss value $L(f(\cdot, \theta^*))$ when f is restricted. As an extreme example, one can set f to be a neural network with only one trainable parameter. Then, a bijection search can trivially find a global minimum with a large loss value. The use of $\mathcal{P}_{n,H,\delta}(d)$ instead of $\tilde{\mathcal{P}}_{n,H,\delta}(d)$ forces us to find nontrivial solutions with small loss values.

In the definition of $\mathcal{P}_{n,H,\delta}(d)$, the network architecture f and learning rate η must be fixed for all datasets. This forces the gradient descent algorithm to actually learn the predictor based on each dataset, instead of encoding too much information into the architecture and learning rate.

B. Analysis

The following theorem states that the probable trainability is ensured with the total parameter number d being linear in n :

Theorem 1. For any $n \in \mathbb{N}^+$, $H \geq 2$, and $\delta > 0$, it holds that $\mathcal{P}_{n,H,\delta}(d) = \text{true}$ for any

$$d \geq c \left(\left(n + m_x H^2 + H^5 \log \left(\frac{H n^2}{\delta} \right) \right) \log \left(\frac{H n^2}{\delta} \right) + n m_y \right),$$

where $c > 0$ is a universal constant.

Remark 1. In Theorem 1, we restrict ourselves to the case of $H \geq 2$. If $H = 1$, then by setting $m_{H-1} = m_x$ and $x_i^{(H-1)} = x_i$ in the proof of Theorem 1, it holds that $\mathcal{P}_{n,1,\delta}(d) = \text{true}$ for any $d \geq cn(m_x + m_y)$. In practice, m_x would be much larger than m_y , and if this is the case, the lower bound $\tilde{\Omega}(nm_x)$ for the case of $H = 1$ is worse than the lower bound $\tilde{\Omega}(nm_y)$ in Theorem 1.

In other words, Theorem 1 and Remark 1 state that there are trainable neural networks of size $\tilde{\Omega}(nm_y + m_x H^2 + H^5)$ if $H \geq 2$, and size $\tilde{\Omega}(n(m_x + m_y))$ if $H = 1$. This is significantly smaller than the sizes required by the previous studies. For deep neural networks, to the best of our knowledge, the state-of-the-art result, in terms of the size, is given in [17], where the neural networks are required to have size $\tilde{\Omega}(2^{O(H)} n^8 + n^4(m_x + m_y))$. For shallow networks, building on previous works [13], [14], it has been proven in [15] that, single-layer networks of size $\tilde{\Omega}(n^2(m_x + m_y))$ are trainable. Theorem 1 proves the probable trainability for considerably smaller networks when compared with the previous results.

Then, a natural question is whether we can further improve Theorem 1 by reducing d while keeping the probable trainability. The following theorem and its corollary state that Theorem 1 is already optimal and that it cannot be improved in terms of the order of the leading term nm_y :

Theorem 2. There exists a universal constant $\mathfrak{c} > 0$ such that the following holds: for any large $\beta > 0$, $\frac{nm_y}{d} - 1 \geq \frac{\mathfrak{c} \beta H \log n}{\log(1/\epsilon)}$, and deep neural network architecture $f \in \mathcal{F}_d^H$, there exists a dataset $S \in \mathcal{S}$ such that if

$$\sum_{i=1}^n \|f(x_i, \theta) - y_i\|_2^2 \leq \epsilon, \quad (5)$$

then $\|\theta\|_2^2 \geq n^\beta$.

Corollary 1. For any $n \in \mathbb{N}^+$, $H \geq 1$, and $\delta > 0$, it holds that $\mathcal{P}_{n,H,\delta}(d) = \text{false}$ for any $d < nm_y$.

Corollary 1 follows from Theorem 2 by taking the parameters $\beta = \sqrt{\log(1/\epsilon)}$ and $\epsilon \rightarrow 0$. If $n \gg H, m_x$, we have the lower bound $d = \tilde{\Omega}(nm_y)$ in Theorem 1, which matches the upper bound nm_y in Corollary 2, except for the logarithmic term and constant.

IV. PROOFS OF TRAINABILITY

This section presents the proofs of Theorems 1 and 2. Throughout this paper, we use c and C to represent various constants, which may be different from line to line.

A. Proof of Theorem 1

We first analyze the properties of the randomly initialized neural networks, and then, relate these properties to the trainability. The following lemma shows that if the input to a layer is normalized, then the outputs and their differences of the layer concentrate to the corresponding means with high probability:

Lemma 1. Consider two data points $x, x' \in \mathbb{R}^{m'}$ that satisfy $\|x\|_2^2 = O(1)$ and $\|x'\|_2^2 = O(1)$. Consider a random weight

matrix $W \in \mathbb{R}^{m \times m'}$ with $\mathcal{N}(0, c_w)$ entries and a random bias term $b \in \mathbb{R}^m$ with $\mathcal{N}(0, c_b)$ entries. Then, the following estimates hold:

$$\mathbb{P} \left(\left| \frac{\|\sigma(Wx + b)\|_2^2}{m} - \mathbb{E}[\sigma^2(g)] \right| \geq \frac{\beta}{\sqrt{m}} \right) \leq e^{-c\beta^2}, \quad (6)$$

$$\mathbb{P} \left(\left| \frac{\|\sigma(Wx + b) - \sigma(Wx' + b)\|_2^2}{m} - \mathbb{E}(\sigma(g) - \sigma(g'))^2 \right| \geq \frac{\beta}{\sqrt{m}} \right) \leq e^{-c\beta^2}, \quad (7)$$

where g, g' are joint Gaussian variables with zero mean and covariances $\mathbb{E}[g^2] = c_w \|x\|_2^2 + c_b$ and $\mathbb{E}[g'^2] = c_w \|x'\|_2^2 + c_b$, $\mathbb{E}[gg'] = c_w \langle x, x' \rangle + c_b$.

Proof of Lemma 1. Since W and b have independent Gaussian entries, $(Wx)_1 + b_1, (Wx)_2 + b_2, \dots, (Wx)_m + b_m$ are independent Gaussian variables with zero mean and variance $c_w \|x\|_2^2 + c_b$. We can rewrite the norm as

$$\frac{1}{m} \|\sigma(Wx + b)\|_2^2 = \frac{1}{m} \sum_{i=1}^m \sigma^2((Wx)_i + b_i). \quad (8)$$

By Assumption 2, the activation function σ is 1-Lipschitz. The random variables $\sigma^2((Wx)_i + b_i)$ are sub-exponential. Therefore, for $|\lambda| > 0$ sufficiently small, we have

$$\begin{aligned} & \mathbb{E} \left[e^{\lambda(\|\sigma(Wx+b)\|_2^2 - m\mathbb{E}[\sigma^2(g)])} \right] \\ &= \prod_{i=1}^m \mathbb{E} \left[e^{\lambda(\sigma^2(Wx_i + b_i) - \mathbb{E}[\sigma^2(g)])} \right] \leq e^{cm\lambda^2}. \end{aligned} \quad (9)$$

Inequality (6) follows from applying the Markov inequality to (9) and setting $\lambda = \pm\beta/(2c\sqrt{m})$. For (7), we can rewrite the norm of the difference as $\frac{1}{m} \|\sigma(Wx + b) - \sigma(Wx' + b)\|_2^2 = \frac{1}{m} \sum_{i=1}^m (\sigma((Wx)_i + b_i) - \sigma((Wx')_i + b_i))^2$. Moreover, the random variables $(\sigma((Wx)_i + b_i) - \sigma((Wx')_i + b_i))^2$ are sub-exponential. Thus, inequality (7) follows from the derivation of (6). \square

By repeatedly applying Lemma 1 to each layer, we obtain the following corollary, which approximates $\|x_i^{(l)}\|_2^2$ and $\|x_i^{(l)} - x_j^{(l)}\|_2^2$ using some constants $p^{(l)}$ and $p_{ij}^{(l)}$ with error terms $O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right)$:

Corollary 2. *For the randomly initialized neural network, the following holds: for any $\beta > 0$, with probability at least $1 - O(e^{-c\beta^2})$ over θ^0 ,*

$$\|x_i^{(l)}\|_2^2 = p^{(l)} + O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right), \quad (10)$$

$$\|x_i^{(l)} - x_j^{(l)}\|_2^2 = p_{ij}^{(l)} + O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right), \quad (11)$$

where $p^{(0)} = 1$, $p_{ij}^{(0)} = \|x_i - x_j\|_2^2 \geq \gamma$, and for $1 \leq l \leq H$, $p_i^{(l)} = \mathbb{E}[\sigma^2(g)]$, and $p_{ij}^{(l)} = \mathbb{E}(\sigma(g) - \sigma(g'))^2$. Here, g, g' are joint Gaussian variables with zero mean and covariances $\mathbb{E}[g^2] = \mathbb{E}[g'^2] = c_w p^{(l-1)} + c_b$ and $\mathbb{E}[gg'] = c_w (p^{(l-1)} - p_{ij}^{(l-1)}/2) + c_b$.

Proof of Corollary 2. We prove the statement by induction on l . The statements hold trivially for $l = 0$. In the following, we assume the statements for l , and prove them for $l + 1$. From Lemma 1, with probability at least $1 - O(e^{-c\beta^2})$,

$$\|x_i^{(l+1)}\|_2^2 = \mathbb{E}[\sigma^2(\tilde{g})] + O\left(\frac{\beta}{\sqrt{m_{l+1}}}\right), \quad (12)$$

$$\|x_i^{(l+1)} - x_j^{(l+1)}\|_2^2 = \mathbb{E}[(\sigma(\tilde{g}) - \sigma(\tilde{g}'))^2] + O\left(\frac{\beta}{\sqrt{m_{l+1}}}\right),$$

where \tilde{g}, \tilde{g}' are Gaussian variables with zero mean and covariances $\mathbb{E}[\tilde{g}^2] = c_w \|x_i^{(l)}\|_2^2 + c_b = c_w p^{(l)} + c_b + O(\sum_{i=1}^l \beta/\sqrt{m_i})$, $\mathbb{E}[\tilde{g}'^2] = c_w \|x_j^{(l)}\|_2^2 + c_b = c_w p^{(l)} + c_b + O(\sum_{i=1}^l \beta/\sqrt{m_i})$, $\mathbb{E}[\tilde{g}\tilde{g}'] = c_w \langle x_i^{(l)}, x_j^{(l)} \rangle + c_b = c_w (p^{(l)} - p_{ij}^{(l)}/2) + c_b + O(\sum_{i=1}^l \beta/\sqrt{m_i})$. We approximate \tilde{g}, \tilde{g}' by mean zero Gaussian variables g, g' such that $\mathbb{E}[g^2] = \mathbb{E}[\tilde{g}^2] = c_w p^{(l)} + c_b$, $\mathbb{E}[g'^2] = c_w p^{(l)} - p_{ij}^{(l)}/2 + c_b$. Since the activation function σ is 1-Lipschitz, we have

$$\begin{aligned} \mathbb{E}[\sigma^2(\tilde{g})] &= \mathbb{E}[\sigma^2(g)] + O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right) \\ &= p^{(l+1)} + O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right), \end{aligned} \quad (13)$$

and

$$\begin{aligned} \mathbb{E}[(\sigma(\tilde{g}) - \sigma(\tilde{g}'))^2] &= \mathbb{E}[(\sigma(g) - \sigma(g'))^2] + O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right) \\ &= p_{ij}^{(l+1)} + O\left(\sum_{i=1}^l \frac{\beta}{\sqrt{m_i}}\right). \end{aligned} \quad (14)$$

The statements for $l + 1$ follow from combining (12), (13), and (14). \square

Now that we have an understanding of the output $x_i^{(H)}$ for each i -th input, we analyze the set of outputs $\{x_i^{(H)}\}_{i=1}^n$ for all inputs. Let $\tilde{d} = m_H(m_{H-1} + 1)$, $\tilde{x} = x^{(H-1)}$, $\tilde{x}_i = x_i^{(H-1)}$, $\tilde{w} = [\tilde{w}_1^\top, \dots, \tilde{w}_{m_H}^\top]^\top = \text{vec}((W^{(H)})^\top)$, and $\tilde{b} = b^{(H)}$. Let $M(\tilde{w}, \tilde{b}) \in \mathbb{R}^{n \times (m_H + 1)}$ given by $M(\tilde{w}, \tilde{b})_{ij} = \sigma(\tilde{w}_j^\top \tilde{x}_i + \tilde{b}_j)/\sqrt{m_H}$ and $M(\tilde{w}, \tilde{b})_{i(m_H+1)} = 1$, for $1 \leq i \leq n$ and $1 \leq j \leq m_H$. The following lemma shows that if $x_i^{(H-1)}$ and $x_j^{(H-1)}$ are distinguishable and the last layer is wide, then the set of outputs $\{x_i^{(H)}\}_{i=1}^n$ is degenerate only when the weights are in a measure zero set:

Lemma 2. *If $\|x_i^{(H-1)}\|_2^2 - \langle x_i^{(H-1)}, x_j^{(H-1)} \rangle > c_\gamma$ for all $i \neq j$ and $m_H \geq n$, the Lebesgue measure of the set $\{(\tilde{w}, \tilde{b}) \in \mathbb{R}^{\tilde{d}} : \text{rank}(M(\tilde{w}, \tilde{b})) < n\}$ is zero.*

Proof of Lemma 2. Under our assumption, the function

$$\varphi(\tilde{w}, \tilde{b}) = \det(M(\tilde{w}, \tilde{b})M(\tilde{w}, \tilde{b})^\top)$$

is analytic since σ is analytic. With this function, we have that $\{(\tilde{w}, \tilde{b}) \in \mathbb{R}^{\tilde{d}} : \text{rank}(M(\tilde{w}, \tilde{b})) < n\} = \{(\tilde{w}, \tilde{b}) \in \mathbb{R}^{\tilde{d}} : \varphi(\tilde{w}, \tilde{b}) = 0\}$, which follows the fact that since $M(\tilde{w}, \tilde{b}) \in \mathbb{R}^{n \times (m_H + 1)}$, the rank of $M(\tilde{w}, \tilde{b})$ and the rank of the Gram matrix are equal.

Since φ is analytic, if φ is not identically zero ($\varphi \neq 0$), the Lebesgue measure of its zero set $\{(\tilde{w}, \tilde{b}) \in \mathbb{R}^d : \varphi(\tilde{w}, \tilde{b}) = 0\}$ is zero [19]. Therefore, it remains to show that $\varphi(\tilde{w}, \tilde{b}) \neq 0$ for some (\tilde{w}, \tilde{b}) .

We now construct a pair (\tilde{w}, \tilde{b}) such that $M(\tilde{w}, \tilde{b})$ is of rank n and $\varphi(\tilde{w}, \tilde{b}) \neq 0$. Set $\tilde{w}_j = \beta \tilde{x}_j$ and $\tilde{b}_j = \mathfrak{c}_\gamma \beta / 2 - \beta \|\tilde{x}_j\|_2^2$ for $j = 1, 2, \dots, n$. Then,

$$M(\tilde{w}, \tilde{b})_{ii} = \sigma(\mathfrak{c}_\gamma \beta / 2) / \sqrt{m_H}, \quad (15)$$

and for any $j \neq i$,

$$M(\tilde{w}, \tilde{b})_{ij} = \sigma(\mathfrak{c}_\gamma \beta / 2 + \tilde{w}_j^\top \tilde{x}_i - \beta \|\tilde{x}_j\|_2^2) / \sqrt{m_H} \quad (16)$$

$$\leq \sigma(-\mathfrak{c}_\gamma \beta / 2) / \sqrt{m_H},$$

which follows the assumption of $\|x_i^{(H-1)}\|_2^2 - \langle x_i^{(H-1)}, x_j^{(H-1)} \rangle > \mathfrak{c}_\gamma$, and the monotonicity of $\sigma(x)$. In (15) and (16), as $\beta \rightarrow \infty$, by our Assumption 2, $M(\tilde{w}, \tilde{b})_{ii} \rightarrow \sigma_+ / \sqrt{m_H}$, and $M(\tilde{w}, \tilde{b})_{ij} \rightarrow \sigma_- / \sqrt{m_H}$, for any $j \neq i$. Therefore, for β sufficiently large, and any $i \in \{1, \dots, n\}$,

$$|M(\tilde{w}, \tilde{b})_{ii} - \sigma_- / \sqrt{m_H}| > \sum_{j \neq i} |M(\tilde{w}, \tilde{b})_{ij} - \sigma_- / \sqrt{m_H}|.$$

This means that the matrix $\tilde{M} = [M(\tilde{w}, \tilde{b})_{ij} - (\sigma_- / \sqrt{m_H})]_{1 \leq i, j \leq n} \in \mathbb{R}^{n \times n}$ is strictly diagonally dominant and nonsingular; hence, its rank is n . This implies that $[\tilde{M}, 1] \in \mathbb{R}^{n \times (n+1)}$ has rank n , which then implies that $[\tilde{M}', 1] \in \mathbb{R}^{n \times (n+1)}$ has rank n , where $\tilde{M}' = [M(\tilde{w}, \tilde{b})_{ij}]_{1 \leq i, j \leq n}$, since elementary matrix operations preserve the matrix rank. Since $m_H \geq n$ and the columns of $M(\tilde{w}, \tilde{b})$ contain all columns of $[\tilde{M}', 1]$, this implies that $\text{rank}(M(\tilde{w}, \tilde{b})) = n$ and $\varphi(\tilde{w}, \tilde{b}) \neq 0$ for this constructed (\tilde{w}, \tilde{b}) , as desired. \square

We now derive an upper bound for the Lipschitz constant of the gradient of the objective function. Let $z_i = [(x_i^{(H)})^\top, 1]^\top$, $\bar{W} = \bar{W}^{(H+1)}$, and $\psi = \psi_{H+1}$. Let $\bar{W}^t = (\bar{W}^{(H+1)})^t$ correspond to θ^t as $\theta^t = (\text{vec}((\bar{W}^{(1)})^t)^\top, \dots, \text{vec}((\bar{W}^{(H+1)})^t)^\top)^\top$. Let $\bar{J}(w) = L(f(\cdot, \psi(\theta^0, \bar{W})))$, where $w = \text{vec}(\bar{W}) \in \mathbb{R}^d$. The following lemma bounds the Lipschitz constant.

Lemma 3. $\nabla \bar{J}$ is Lipschitz continuous with Lipschitz constant at most $\frac{\zeta}{n} \sum_{i=1}^n \|z_i\|_2^2$.

Proof of Lemma 3. Let $\theta_w = \psi(\theta^0, \bar{W})$ and $\theta_{w'} = \psi(\theta^0, \bar{W}')$ where $w' = \text{vec}((\bar{W}')^\top)$. Then,

$$\begin{aligned} & \|\nabla \bar{J}(w) - \nabla \bar{J}(w')\|_2 \\ &= \frac{1}{n} \left\| \sum_{i=1}^n \nabla_w (\ell_i \circ f_i)(\theta_w) - \nabla_{w'} (\ell_i \circ f_i)(\theta_{w'}) \right\|_2 \\ &\leq \frac{1}{n} \sum_{i=1}^n \|[z_i \otimes I_{m_y}]\|_2 \|\nabla \ell_i(f_i(\theta_w)) - \nabla \ell_i(f_i(\theta_{w'}))\|_2 \\ &\leq \frac{1}{n} \sum_{i=1}^n \zeta \|z_i\|_2 \|f(x_i, \theta_w) - f(x_i, \theta_{w'})\|_2 \\ &\leq \left(\frac{\zeta}{n} \sum_{i=1}^n \|z_i\|_2^2 \right) \|\bar{W} - \bar{W}'\|_2 \leq \left(\frac{\zeta}{n} \sum_{i=1}^n \|z_i\|_2^2 \right) \|w - w'\|_2, \end{aligned}$$

where the last line follows $\|\bar{W} - \bar{W}'\|_2 \leq \|\bar{W} - \bar{W}'\|_F = \|w - w'\|_2$. \square

Using these lemmas, we can complete the proof of Theorem 1. Let f be an arbitrary neural network architecture satisfying $m_1, m_2, \dots, m_{H-2} \geq O(C^2 H^2 \log(Hn^2/\delta))$, $m_{H-1} \geq O(C^2 \log(Hn^2/\delta))$, and $m_H \geq O(n)$, for some constant C . Since such an arbitrary network has a total number of parameters $d = \mathfrak{c}(m_x H^2 \log(\frac{Hn^2}{\delta}) + H^5 \log^2(\frac{Hn^2}{\delta}) + n(\log(\frac{Hn^2}{\delta}) + m_y))$ or higher, all we need to show now is that such an arbitrary architecture ensures the desired trainability.

By setting $\beta = \sqrt{\log(Hn^2/\delta)}/c$ and $l = H - 1$ in Corollary 2 and by taking a union bound, it holds that with probability at least $1 - \delta$, for any $1 \leq i \neq j \leq n$,

$$\begin{aligned} \|x_i^{(H-1)}\|_2^2 &= p^{(H-1)} + O(1/(cC)), \\ \|x_i^{(H-1)} - x_j^{(H-1)}\|_2^2 &= p_{ij}^{(H-1)} + O(1/(cC)). \end{aligned} \quad (17)$$

In particular, it follows by considering C sufficiently large that with probability at least $1 - \delta$, for any $1 \leq i \neq j \leq n$,

$$\begin{aligned} \|x_i^{(H-1)}\|_2 &= O(1), \\ \|x_i^{(H-1)}\|_2^2 - \langle x_i^{(H-1)}, x_j^{(H-1)} \rangle &= (p_{ij}^{(H-1)} + o(1))/2 > \mathfrak{c}_\gamma, \end{aligned} \quad (18)$$

where the constant \mathfrak{c}_γ depends only on γ .

Since the neural network is initialized by the Gaussian probability measure, which is absolutely continuous with respect to the Lebesgue measure, equations (17)–(18) and Lemma 2 imply that, with probability at least $1 - \delta$, $\text{rank}(M(\tilde{w}, \tilde{b})) = n$ and $\frac{1}{n} \sum_{i=1}^n \|z_i\|_2 \leq c_z$ for some constant c_z . Accordingly, we consider the case of $\text{rank}(M(\tilde{w}, \tilde{b})) = n$ and $\frac{1}{n} \sum_{i=1}^n \|z_i\|_2 \leq c_z$ in the following.

Since $\frac{1}{n} \sum_{i=1}^n \|z_i\|_2 \leq c_z$, from Lemma 3, $\nabla \bar{J}$ has Lipschitz constant at most $c_z \zeta$. Therefore, for any $w', w \in \mathbb{R}^d$,

$$\begin{aligned} \bar{J}(w') &= \bar{J}(w) + \int_0^1 \nabla \bar{J}(w + q(w' - w))^\top (w' - w) dq \\ &\leq \bar{J}(w) + \nabla \bar{J}(w)^\top (w' - w) + \frac{c_z \zeta}{2} \|w' - w\|_2^2. \end{aligned} \quad (19)$$

We set $\eta_i = \frac{1}{c_z \zeta}$ if $i > \sum_{l=0}^{H-1} m_l m_{l+1} + m_{l+1}$ and $\eta_i = 0$ otherwise. Using (19) with $w' = w^{k+1}$ and $w = w^k$, and the equation of $w^{k+1} = w^k - \frac{1}{c_z \zeta} \nabla \bar{J}(w^k)$, we obtain

$$\bar{J}(w^{k+1}) \leq \bar{J}(w^k) - \frac{c_z \zeta}{2} \|w^{k+1} - w^k\|_2^2 \leq \bar{J}(w^k). \quad (20)$$

Using (19) with $w' = w^{k+1}$ and $w = w^k$, we find that, for all $w \in \mathbb{R}^d$,

$$\begin{aligned} & \bar{J}(w^{k+1}) \\ &\leq \bar{J}(w^k) + \nabla \bar{J}(w^k)^\top (w^{k+1} - w^k) + \frac{c_z \zeta}{2} \|w^{k+1} - w^k\|_2^2 \\ &= \bar{J}(w^k) + \nabla \bar{J}(w^k)^\top (w - w^k) + \frac{c_z \zeta}{2} (\|w - w^k\|_2^2 - \|w - w^{k+1}\|_2^2) \\ &\leq \bar{J}(w) + \frac{c_z \zeta}{2} (\|w - w^k\|_2^2 - \|w - w^{k+1}\|_2^2), \end{aligned} \quad (21)$$

where the third line contains only arithmetic rearrangements using the equation of $\nabla \bar{J}(w^k) = c_z \zeta (w^k - w^{k+1})$, and the

last line follows the convexity of \bar{J} . Using (20) and (21), we have that, for any $w \in \mathbb{R}^d$,

$$\begin{aligned} t\bar{J}(w^t) &\leq \sum_{k=0}^{t-1} \bar{J}(w^{k+1}) \\ &\leq t\bar{J}(w) + \frac{c_z \zeta}{2} (\|w - w^0\|_2^2 - \|w - w^t\|_2^2). \end{aligned} \quad (22)$$

Let $f^*(X) = [f^*(x_1), \dots, f^*(x_n)] \in \mathbb{R}^{m_y \times n}$ and $f(X, \theta) = [f(x_1, \theta), \dots, f(x_n, \theta)] \in \mathbb{R}^{m_y \times n}$. If $\text{rank}(M(\tilde{w}, \tilde{b})) = n$, there exists a minimum norm solution $\bar{W}^* \in \mathbb{R}^{m_y \times (m_H+1)}$ such that $f(X, \psi(\theta^0, \bar{W}^*)) = \bar{W}^* M(\tilde{w}, \tilde{b})^\top = f^*(X)$, and hence $\bar{J}(w^*) = L(f(\cdot, \psi(\theta^0, \bar{W}^*))) = L(f^*)$, where $w^* = \text{vec}((W^*)^\top)$. Thus, using (22) and recalling the parameter c_r from Definition 1, we have $\bar{J}(w^t) \leq L(f^*) + \frac{c_z c_r \zeta}{2t}$, which implies that $\bar{J}(w^t) \leq \bar{J}(w^*) + \epsilon$, where $t = O(\frac{c_r \zeta}{\epsilon})$.

Therefore, recalling that we have $\text{rank}(M(\tilde{w}, \tilde{b})) = n$ and $\frac{1}{n} \sum_{i=1}^n \|z_i\|_2 \leq c_z$ with probability at least $1 - \delta$, it holds that with probability at least $1 - \delta$,

$$\bar{J}(w^t) \leq \bar{J}(w^*) + \epsilon$$

where $t = O(\frac{c_r \zeta}{\epsilon})$. Then, using (22), we have $\|w^* - w^t\|_2 \leq \|w^* - w^0\|_2$, which implies that $\|w^t\|_2^2 \leq (\|w^0\|_2 + 2\|w^*\|_2)^2 \leq c_\theta$ for some constant in $\epsilon > 0$. \square

B. Proof of Theorem 2

We consider the following map from the parameter space to the concatenation of the output of the model at x_1, x_2, \dots, x_n :

$$f_X : \theta \mapsto \text{vec}([f(x_1, \theta), f(x_2, \theta), \dots, f(x_n, \theta)]). \quad (23)$$

By Assumption 2, the map f_X is analytic in θ . We recall that the Jacobian of the map f_X is defined as

$$\text{Jac}(f_X)(\theta) = [\partial_k f(x_i, \theta)]_{1 \leq i \leq n, 1 \leq k \leq d} \in \mathbb{R}^{nm_y \times d}$$

In general, the image of the map f_X may not be a manifold. Sard's theorem asserts that the set of critical values, i.e., the image of the set of critical points $\{\theta : \text{rank Jac}(f_X)(\theta) < d\}$, has Lebesgue measure 0. For any noncritical point θ , i.e., $\text{rank Jac}(f_X)(\theta) = d$, there exists a small neighborhood $U(\theta)$ of θ , such that over $U(\theta)$, the rank of the Jacobian matrix of f_X is d . Then, the rank theorem states that, the image $f_X(U(\theta))$ is a manifold of dimension d . Therefore, the volume of the image of the map f_X is well defined, and we have the upper bound:

$$\begin{aligned} \text{vol} f_X(\{\theta : \|\theta\|_2^2 \leq R^2\}) &\leq \text{vol}(\mathbb{B}_d(R)) \sup_{\theta \in \mathbb{B}_d(R)} \det \text{Jac} f_X(\theta) \\ &= \frac{\pi^{d/2} R^d}{\Gamma(d/2 + 1)} \sup_{\theta \in \mathbb{B}_d(R)} \det \text{Jac} f_X(\theta), \end{aligned} \quad (24)$$

where $\mathbb{B}_d(R)$ is the radius- R ball in \mathbb{R}^d .

In the following, we show that if for any point $\text{vec}([y_1, y_2, \dots, y_n]) \in [-1, 1]^{nm_y}$, there exists some $\theta \in$

\mathbb{R}^d with $\sum_{i=1}^n \|f(x_i, \theta) - y_i\|_2^2 \leq \epsilon$, then there exists a large universal constant \mathfrak{c} such that $\frac{nm_y}{d} - 1 \leq \frac{\mathfrak{c} \beta H \log n}{\log(1/\epsilon)}$. If this is the case, then the $\sqrt{\epsilon}$ -neighborhood of the image set of the map f_X covers all possible labels $[-1, 1]^{nm_y}$. This fact, combined with (24), implies that

$$\epsilon^{(nm_y-d)/2} \frac{\pi^{d/2} R^d}{\Gamma(d/2 + 1)} \sup_{\theta \in \mathbb{B}_d(R)} \det \text{Jac} f_X(\theta) \geq 2^{nm_y}, \quad (25)$$

The following lemma provides an upper bound on $\det \text{Jac} f_X(\theta)$, which will be used to obtain the lower bound for the Euclidean norm of θ .

Lemma 4. *We have the following estimates for the determinant of the Jacobian of f_X :*

$$\begin{aligned} &\sup_{\theta \in \mathbb{B}_d(R)} \det \text{Jac} f_X(\theta) \\ &\leq \left(\frac{2(H+1)n}{d} \left(\frac{m_y^2 + H + R^2}{H+1} \right)^{H+1} \right)^{d/2} \end{aligned} \quad (26)$$

Proof of Lemma 4. For any θ , we denote the singular values of $\text{Jac} f_X(\theta)$ as s_1, s_2, \dots, s_d . Then,

$$\begin{aligned} \det \text{Jac} f_X(\theta) &= \prod_{i=1}^d s_i \leq \left(\frac{\sum_{i=1}^d s_i^2}{d} \right)^{d/2} \\ &= \left(\frac{\|\text{Jac} f_X(\theta)\|_F^2}{d} \right)^{d/2}. \end{aligned} \quad (27)$$

In the following, we derive an upper bound for the Frobenius norm of $\text{Jac} f_X(\theta)$. Then, inequality (27) gives an upper bound for the determinant of $\text{Jac} f_X(\theta)$.

By the definition of the Jacobian matrix,

$$\begin{aligned} \|\text{Jac} f_X(\theta)\|_F^2 &= \sum_{i=1}^n \|\partial_\theta f(x_i, \theta)\|_F^2 \\ &= \sum_{l=1}^{H+1} \sum_{i=1}^n \|\partial_{W^{(l)}} f(x_i, \theta)\|_F^2 + \|\partial_{b^{(l)}} f(x_i, \theta)\|_F^2. \end{aligned} \quad (28)$$

We have the following estimates for the derivatives for $1 \leq l \leq H$,

$$\begin{aligned} &\|\partial_{W^{(l)}} f(x_i, \theta)\|_F^2 + \|\partial_{b^{(l)}} f(x_i, \theta)\|_F^2 \\ &\leq \|W^{(H+1)}\|_F^2 (1 + \|x^{(l-1)}\|_2^2) \prod_{i=l+1}^H \|W^{(i)}\|_2^2, \end{aligned} \quad (29)$$

and for $l = H+1$

$$\begin{aligned} &\|\partial_{W^{(H+1)}} f(x_i, \theta)\|_F^2 + \|\partial_{b^{(H+1)}} f(x_i, \theta)\|_F^2 \\ &= m_y^2 (1 + \|x^{(H)}\|_2^2), \end{aligned} \quad (30)$$

since the activation function is 1-Lipschitz. From the defining relation of a feedforward neural network, and from the fact that the activation function is 1-Lipschitz, we obtain the following recursive bound for $x^{(l)}$,

$$\begin{aligned} \|x^{(l)}\|_2^2 &\leq \|W^{(l)} x^{(l-1)} + b^{(l)}\|_2^2 \\ &\leq (\|W^{(l)}\|_F^2 + \|b^{(l)}\|_F^2) (\|x^{(l-1)}\|_2^2 + 1). \end{aligned} \quad (31)$$

We can iterate inequality (31) and obtain the following bound for $1 + \|x^{(l)}\|_2^2$,

$$1 + \|x^{(l)}\|_2^2 \leq (1 + \|x\|_2^2) \prod_{i=1}^l (1 + \|W^{(i)}\|_F^2 + \|b^{(i)}\|_F^2). \quad (32)$$

Using (29), (30), and (32), we conclude the following estimate for the Euclidean norm of $\partial_\theta f(x_i, \theta)$,

$$\begin{aligned} & \|\partial_\theta f(x_i, \theta)\|_F^2 \\ &= \sum_{l=1}^{H+1} \|\partial_{W^{(l)}} f(x_i, \theta)\|_F^2 + \|\partial_{b^{(l)}} f(x_i, \theta)\|_F^2 \\ &\leq 2(H+1)(m_y^2 + \|W^{(H+1)}\|_F^2) \prod_{i=1}^H (1 + \|W^{(i)}\|_F^2 + \|b^{(i)}\|_F^2) \\ &\leq 2(H+1) \left(\frac{m_y^2 + H + \|\theta\|_2^2}{H+1} \right)^{H+1} \end{aligned} \quad (33)$$

where the last line follows the AM–GM inequality. Lemma 4 follows from combining (27), (28) and (33), and noticing $\theta \in \mathbb{B}_d(R)$. \square

Using Lemma 4, we can finish the proof of Theorem 2. By substituting (26) into (25), and raising both sides to the $1/d$ -th power, we obtain the following key estimate

$$\frac{CRn}{d^{3/2}} \frac{(m_y^2 + H + R^2)^{H+1}}{HH} \geq \left(\frac{2}{\sqrt{\epsilon}} \right)^{nm_y/d-1}, \quad (34)$$

where C is a universal constant. It follows that there exists a large universal constant c such that if $\frac{nm_y}{d} - 1 \geq \frac{c\beta H \log n}{\log(1/\epsilon)}$, then $R \geq n^\beta$. This finishes the proof of Theorem 2. \square

V. GENERALIZATION BOUND AND EXPERIMENTS

The previous sections presented the construction of deep neural network architectures of practical sizes, with the trainability guarantee. A major question remaining now is whether the constructed neural networks can generalize to unseen data points after training, which is the focus of this section.

This section considers multiclass classification with the one-hot vector $y \in \{0, 1\}^{m_y}$. Let $j(y) \in \{1, \dots, m_y\}$ be the index of the one-hot vector y having entry one as $y_{j(y)} = 1$. Let ℓ_{01} represent the 0–1 loss as $\ell(f(x, \theta), y) = \mathbb{1}\{\arg\max_j f(x, \theta)_j \neq j(y)\}$, with which we can write the expected test error $\mathbb{E}_{(x,y)}[\ell_{01}(f(x, \theta), y)]$. Let ℓ_ρ be a standard multiclass margin loss defined by $\ell_\rho(f(x, \theta), y) = \min(\max(1 - (f(x, \theta)_{j(y)} - \max_{j' \neq j(y)} f(x, \theta)_{j'})/\rho, 0), 1)$. We set f and η as constructed in the proof of Theorem 1 (i.e., $m_1, m_2, \dots, m_{H-2} = O(H^2 \log(Hn^2/\delta))$, $m_{H-1} = O(\log(Hn^2/\delta))$, and $m_H = O(n)$).

The following proposition provides a data-dependent generalization bound, which shows that the trainable deep neural networks can generalize to unseen data points if the weight norm turns out to be small after training:

Proposition 1. Fix $\rho > 0$ and $\varsigma \geq 1$. Then, for any $\delta' > 0$, with probability at least $1 - \delta - \delta'$ over θ^0 and i.i.d.

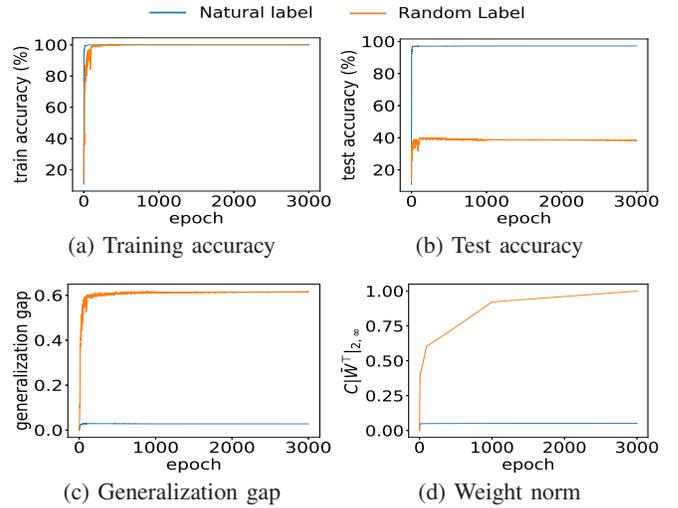


Fig. 2: Training accuracy, test accuracy, generalization gap, and weight norm for a neural network of practical size with the trainability guarantee, which is constructed in the proof of Theorem 1. Even though the trainable neural network has the capacity to memorize any dataset, it generalizes well with the natural label, but not with the random label. This behavior matches the growth of the weight norm as predicted by Proposition 1.

$((x_i, y_i))_{i=1}^n$, the following holds for any θ^t generated by the gradient descent (as $\theta^t = \theta^{t-1} - \eta \odot \nabla J(\theta^{t-1})$):

$$\begin{aligned} & \mathbb{E}_{(x,y)}[\ell_{01}(f(x, \theta^t), y)] - \frac{1}{n} \sum_{i=1}^n \ell_\rho(f(x_i, \theta^t), y_i) \\ &\leq \frac{cm_y^2 [\varsigma \|(\bar{W}^t)^\top\|_{2,\infty}]}{\rho\varsigma\sqrt{n}} + \sqrt{\frac{\ln \frac{\pi^2 [\varsigma \|(\bar{W}^t)^\top\|_{2,\infty}]^2}{\delta'}}{2n}}. \end{aligned}$$

for some constant $c = O(1)$.

Figure 2 shows the training accuracy, test accuracy, generalization gap, and weight norm of one of the neural networks trained with our trainability guarantee. In the figure, the trainable deep neural network generalizes well with a natural dataset, while it does not generalize well with a random dataset, as predicted by the values of the weight norm. Here, we use the softplus activation, $H = 2$, $m_1 = 16 \log(Hn^2/\delta)$, and $m_H = 4n$. We employ the MNIST dataset [20], which is a popular dataset for recognizing handwritten digits with $m_x = 784$ and $m_y = 10$. For the random-label experiment, the natural labels in the MNIST dataset are replaced by randomly generated labels. The generalization gap plotted in subfigure 2c is the value of (training accuracy - test accuracy)/100. The weight norm plotted in subfigure 2d is the value of $C\|\bar{W}^\top\|_{2,\infty}$, where C is the normalization constant.

VI. CONCLUSIONS

In this paper, we have proven that there are trainable and generalizable deep neural networks of sizes growing only linearly in the dataset size n . We have shown that this is already the optimal rate in terms of the dataset size n and that

it cannot be improved further, except by a logarithmic factor. In terms of the rate, these theoretical results are consistent with the practical observations and previous expressivity theories. Future work involves improvements in terms of constant and logarithmic factors.

Looking forward, the formalization of the probable trainability $\mathcal{P}_{n,H,\delta}$ would contribute to set a common language in the future studies on trainability. For example, one can consider data-dependent probable trainability by redefining \mathcal{S}_n and architecture-dependent probable trainability by reformulating \mathcal{F}_d^H , in the definition of $\mathcal{P}_{n,H,\delta}$. Our trainability results differ from recent results of practical guarantees on loss landscape with representation learning effects [21], [22].

APPENDIX

A. Proof of Proposition 1

Define $\Theta_k = \{\theta \in \mathbb{R}^d : (\exists \bar{W} \in \mathcal{W}_k)[\theta = \psi(\theta^0, \bar{W})]\}$ for all $k \in \mathbb{N}^+$, where $\mathcal{W}_k = \{\bar{W} \in \mathbb{R}^{m_y \times (m_H+1)} : k-1 \leq \varsigma \|\bar{W}^\top\|_{2,\infty} < k\}$. Let $\mathcal{T}(\Theta_k) = \{x \mapsto f(x, \theta)_j : \theta \in \Theta_k, j \in \mathcal{J}\}$ where $\mathcal{J} = \{1, \dots, m_y\}$. Then, the previous result [23] implies that for any $\delta'_k > 0$, with probability at least $1 - \delta'_k$, the following holds for all $\theta \in \Theta_k$: $\mathbb{E}_{(x,y)}[\ell_{01}(f(x, \theta), y)] - \frac{1}{n} \sum_{i=1}^n \ell_\rho(f(x_i, \theta), y_i) \leq \frac{2m_y^2}{\rho} \mathfrak{R}_n(\mathcal{T}(\Theta_k)) + \sqrt{\frac{\ln(1/\delta'_k)}{2n}}$, where $\mathfrak{R}_n(\mathcal{T}(\Theta_k))$ is the Rademacher complexity of the set $\mathcal{T}(\Theta_k)$, given by: $\mathfrak{R}_n(\mathcal{T}(\Theta_k)) = \mathbb{E}_{S,\xi}[\sup_{\theta \in \Theta_k, j \in \mathcal{J}} \frac{1}{n} \sum_{i=1}^n \xi_i f(x_i, \theta)_j]$. Here, ξ_1, \dots, ξ_n are independent uniform random variables taking values in $\{-1, 1\}$ (i.e., Rademacher variables).

Set $\delta'_k = \delta' \frac{6}{\pi^2 k^2}$, with which $\sum_{k=1}^{\infty} \delta'_k = \delta'$. By taking the union bound over $k \in \mathbb{N}^+$, for any $\delta' > 0$, with probability at least $1 - \delta'$, the following holds for all $k \in \mathbb{N}^+$ and all $\theta \in \Theta_k$:

$$\begin{aligned} & \mathbb{E}_{(x,y)}[\ell_{01}(f(x, \theta), y)] - \frac{1}{n} \sum_{i=1}^n \ell_\rho(f(x_i, \theta), y_i) \\ & \leq \frac{2m_y^2}{\rho} \mathfrak{R}_n(\mathcal{T}(\Theta_k)) + \sqrt{\frac{\ln \frac{\pi^2 [\varsigma \|\bar{W}^\top\|_{2,\infty}]^2}{\delta'}}{2n}}. \end{aligned} \quad (35)$$

By using the Cauchy-Schwarz inequality, $\mathfrak{R}_n(\mathcal{T}(\Theta_k)) \leq \frac{[\varsigma \|\bar{W}^\top\|_{2,\infty}]}{\varsigma n} \mathbb{E}_{S,\xi}[\|\sum_{i=1}^n \xi_i z_i\|_2]$. By using linearity of expectation and Jensen's inequality (since the square root is concave in its domain), $\mathbb{E}_{S,\xi}[\|\sum_{i=1}^n \xi_i z_i\|_2] \leq (\mathbb{E}_S \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}_\xi[\xi_i \xi_j] z_i^\top z_j)^{1/2} = (\sum_{i=1}^n \mathbb{E}_S[\|z_i\|_2^2])^{1/2} \leq (c/2)\sqrt{n}$, where we utilize the fact that, with probability at least $1 - \delta$, $\|z_i\|_2 \leq c/2$ for some constant $c = O(1)$, as shown in the proof of Theorem 1. Therefore, with probability at least $1 - \delta$,

$$\mathfrak{R}_n(\mathcal{T}(\Theta_k)) \leq \frac{(c/2) [\varsigma \|\bar{W}^\top\|_{2,\infty}]}{\varsigma \sqrt{n}}. \quad (36)$$

The desired statement follows by taking the union bound for the events of (35) and (36). \square

REFERENCES

- [1] E. B. Baum, "On the capabilities of multilayer perceptrons," *J. Complexity*, vol. 4, no. 3, pp. 193–215, 1988. [Online]. Available: [https://doi.org/10.1016/0885-064X\(88\)90020-9](https://doi.org/10.1016/0885-064X(88)90020-9)
- [2] T. M. Cover, "Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition," *IEEE Trans. Electronic Computers*, vol. 14, no. 3, pp. 326–334, 1965. [Online]. Available: <https://doi.org/10.1109/PGEC.1965.264137>
- [3] G. Huang, "Learning capability and storage capacity of two-hidden-layer feedforward networks," *IEEE Trans. Neural Networks*, vol. 14, no. 2, pp. 274–281, 2003. [Online]. Available: <https://doi.org/10.1109/TNN.2003.809401>
- [4] S. Huang and Y. Huang, "Bounds on the number of hidden neurons in multilayer perceptrons," *IEEE Trans. Neural Networks*, vol. 2, no. 1, pp. 47–55, 1991. [Online]. Available: <https://doi.org/10.1109/72.80290>
- [5] G. Huang and H. A. Babri, "Upper bounds on the number of hidden neurons in feedforward networks with arbitrary bounded nonlinear activation functions," *IEEE Trans. Neural Networks*, vol. 9, no. 1, pp. 224–229, 1998. [Online]. Available: <https://doi.org/10.1109/72.655045>
- [6] M. Yamasaki, "The lower bound of the capacity for a neural network with multiple hidden layers," in *International Conference on Artificial Neural Networks*. Springer, 1993, pp. 546–549.
- [7] C. Yun, S. Sar, and A. Jadbabaie, "Small relu networks are powerful memorizers: a tight analysis of memorization capacity," *arXiv preprint, arXiv:1810.07770*, 2018.
- [8] M. Hardt and T. Ma, "Identity matters in deep learning," in *International Conference on Learning Representations*, 2017.
- [9] Q. Nguyen and M. Hein, "Optimization landscape and expressivity of deep cnns," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, 2018, pp. 3727–3736. [Online]. Available: <http://proceedings.mlr.press/v80/nguyen18a.html>
- [10] P. Bartlett and S. Ben-David, "Hardness results for neural network approximation problems," in *European Conference on Computational Learning Theory*. Springer, 1999, pp. 50–62.
- [11] A. Blum and R. L. Rivest, "Training a 3-node neural network is np-complete," in *Advances in neural information processing systems*, 1989, pp. 494–501.
- [12] R. Livni, S. Shalev-Shwartz, and O. Shamir, "On the computational efficiency of training neural networks," in *Advances in Neural Information Processing Systems*, 2014, pp. 855–863.
- [13] Y. Li and Y. Liang, "Learning overparameterized neural networks via stochastic gradient descent on structured data," in *Advances in Neural Information Processing Systems*, 2018, pp. 8157–8166.
- [14] S. S. Du, X. Zhai, B. Póczos, and A. Singh, "Gradient descent provably optimizes over-parameterized neural networks," in *ICLR, arXiv:1810.02054*, 2018.
- [15] Z. Song and X. Yang, "Quadratic suffices for over-parametrization via matrix chernoff bound," *arXiv preprint arXiv:1906.03593*, 2019.
- [16] Z. Allen-Zhu, Y. Li, and Z. Song, "A convergence theory for deep learning via over-parameterization," in *ICML, arXiv:1811.03962*, 2018.
- [17] S. S. Du, J. D. Lee, H. Li, L. Wang, and X. Zhai, "Gradient descent finds global minima of deep neural networks," *ICML, arXiv:1811.03804*, 2018.
- [18] D. Zou, Y. Cao, D. Zhou, and Q. Gu, "Stochastic gradient descent optimizes over-parameterized deep relu networks," *arXiv preprint arXiv:1811.08888*, 2018.
- [19] B. Mityagin, "The zero set of a real analytic function," *arXiv preprint arXiv:1512.07276*, 2015.
- [20] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [21] K. Kawaguchi and Y. Bengio, "Depth with nonlinearity creates no bad local minima in ResNets," *Neural Networks*, vol. 118, pp. 167–174, 2019.
- [22] K. Kawaguchi, J. Huang, and L. P. Kaelbling, "Every local minimum value is the global minimum value of induced model in non-convex machine learning," *Neural Computation, to appear*, 2019.
- [23] V. Koltchinskii and D. Panchenko, "Empirical margin distributions and bounding the generalization error of combined classifiers," *Annals of Statistics*, pp. 1–50, 2002.