

On the Depth of Deep Neural Networks: A Theoretical View

Shizhao Sun,^{1,*} Wei Chen,² Liwei Wang,³ Xiaoguang Liu,¹ and Tie-Yan Liu²

¹College of Computer and Control Engineering, Nankai University, Tianjin, 300071, P. R. China

²Microsoft Research, Beijing, 100080, P. R. China

³Key Laboratory of Machine Perception (MOE), School of EECS, Peking University, Beijing, 100871, P. R. China

sunshizhao@mail.nankai.edu.cn, wche@microsoft.com, wanglw@cis.pku.edu.cn

liuxg@njb1.nankai.edu.cn, tyliu@microsoft.com

Abstract

People believe that depth plays an important role in success of deep neural networks (DNN). However, this belief lacks solid theoretical justifications as far as we know. We investigate role of depth from perspective of margin bound. In margin bound, expected error is upper bounded by empirical margin error plus Rademacher Average (RA) based capacity term. First, we derive an upper bound for RA of DNN, and show that it increases with increasing depth. This indicates negative impact of depth on test performance. Second, we show that deeper networks tend to have larger representation power (measured by Betti numbers based complexity) than shallower networks in multi-class setting, and thus can lead to smaller empirical margin error. This implies positive impact of depth. The combination of these two results shows that for DNN with restricted number of hidden units, increasing depth is not always good since there is a tradeoff between positive and negative impacts. These results inspire us to seek alternative ways to achieve positive impact of depth, e.g., imposing margin-based penalty terms to cross entropy loss so as to reduce empirical margin error without increasing depth. Our experiments show that in this way, we achieve significantly better test performance.

1 Introduction

Deep neural networks (DNN) have achieved great practical success in many machine learning tasks, such as speech recognition, image classification, and natural language processing (Hinton and Salakhutdinov 2006; Krizhevsky, Sutskever, and Hinton 2012; Hinton et al. 2012; Ciresan, Meier, and Schmidhuber 2012; Weston et al. 2012). Many people believe that the depth plays an important role in the success of DNN (Srivastava, Greff, and Schmidhuber 2015; Simonyan and Zisserman 2014; Lee et al. 2014; Romero et al. 2014; He et al. 2015; Szegedy et al. 2014). However, as far as we know, such belief is still lacking solid theoretical justification.

On one hand, some researchers have tried to understand the role of depth in DNN by investigating its generalization bound. For example, in (Bartlett, Maierov, and Meir 1998; Karpinski and Macintyre 1995; Goldberg and Jerrum 1995),

generalization bounds for multi-layer neural networks were derived based on Vapnik-Chervonenkis (VC) dimension. In (Bartlett 1998; Koltchinskii and Panchenko 2002), a margin bound was given to fully connected neural networks in the setting of binary classification. In (Neyshabur, Tomioka, and Srebro 2015), the capacity of different norm-constrained feed-forward networks was investigated. While these works shed some lights on the theoretical properties of DNN, they have limitations in helping us understand the role of depth, due to the following reasons. First, the number of parameters in many practical DNN models could be very large, sometimes even larger than the size of training data. This makes the VC dimension based generalization bound too loose to use. Second, practical DNN are usually used to perform multi-class classifications and often contains many convolutional layers, such as the model used in the tasks of ImageNet (Deng et al. 2009). However, most existing bounds are only regarding binary classification and fully connected networks. Therefore, the bounds cannot be used to explain the advantage of using deep neural networks.

On the other hand, in recent years, researchers have tried to explain the role of depth from other angles, e.g., deeper neural networks are able to represent more complex functions. In (Hastad 1986; Delalleau and Bengio 2011), authors showed that there exist families of functions that can be represented much more efficiently with a deep logic circuit or sum-product network than with a shallow one, i.e., with substantially fewer hidden units. In (Bianchini and Scarselli 2014; Montufar et al. 2014), it was demonstrated that deeper nets could represent more complex functions than shallower nets in terms of maximal number of linear regions and Betti numbers. However, these works are apart from the generalization of the learning process, and thus they cannot be used to explain the test performance improvement for DNN.

To reveal the role of depth in DNN, in this paper, we propose to investigate the margin bound of DNN. According to the margin bound, the expected 0-1 error of a DNN model is upper bounded by the empirical margin error plus a Rademacher Average (RA) based capacity term. Then we first derive an upper bound for the RA-based capacity term, for both fully-connected and convolutional neural networks in the multi-class setting. We find that with the increasing depth, this upper bound of RA will increase, which indicates that depth has its negative impact on the test perfor-

*This work was done when the author was visiting Microsoft Research Asia.

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

mance of DNN. Second, for the empirical margin error, we study the representation power of deeper networks, because if a deeper net can produce more complex classifiers, it will be able to fit the training data better w.r.t. any margin coefficient. Specifically, we measure the representation power of a DNN model using the Betti numbers based complexity (Bianchini and Scarselli 2014), and show that, in the multi-class setting, the Betti numbers based complexity of deeper nets are indeed much larger than that of shallower nets. This, on the other hand, implies the positive impact of depth on the test performance of DNN. By combining these two results, we can come to the conclusion that for DNN with restricted number of hidden units, arbitrarily increasing the depth is not always good since there is a clear tradeoff between its positive and negative impacts. In other words, with the increasing depth, the test error of DNN may first decrease, and then increase. This pattern of test error has been validated by our empirical observations on different datasets.

The above theoretical findings also inspire us to look for alternative ways to achieve the positive impact of depth, and avoid its negative impact. For example, it seems feasible to add a margin-based penalty term to the cross entropy loss of DNN so as to directly reduce the empirical margin error on the training data, without increasing the RA of the DNN model. For ease of reference, we call the algorithm minimizing the penalized cross entropy loss *large margin DNN* (LMDNN)¹. We have conducted extensive experiments on benchmark datasets to test the performance of LMDNN. The results show that LMDNN can achieve significantly better test performance than standard DNN. In addition, the models trained by LMDNN have smaller empirical margin error at almost all the margin coefficients, and thus their performance gains can be well explained by our derived theory.

The remaining part of this paper is organized as follows. In Section 2, we give some preliminaries for DNN. In Section 3, we investigate the roles of depth in RA and empirical margin error respectively. In Section 4, we propose the large margin DNN algorithms and conduct experiments to test their performances. In Section 5, we conclude the paper and discuss some future works.

2 Preliminaries

Given a multi-class classification problem, we denote $\mathcal{X} = \mathbb{R}^d$ as the input space, $\mathcal{Y} = \{1, \dots, K\}$ as the output space, and P as the joint distribution over $\mathcal{X} \times \mathcal{Y}$. Here d denotes the dimension of the input space, and K denotes the number of categories in the output space. We have a training set $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$, which is i.i.d. sampled from $\mathcal{X} \times \mathcal{Y}$ according to distribution P . The goal is to learn a prediction model $f \in \mathcal{F} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ from the training set, which produces an output vector $(f(x, k); k \in \mathcal{Y})$ for each instance $x \in \mathcal{X}$ indicating its likelihood of belonging to category k . Then the final classification is determined by $\arg \max_{k \in \mathcal{Y}} f(x, k)$. This naturally leads to the follow-

ing definition of the *margin* $\rho(f; x, y)$ of the model f at a labeled sample (x, y) :

$$\rho(f; x, y) = f(x, y) - \max_{k \neq y} f(x, k). \quad (1)$$

The classification accuracy of the prediction model f is measured by its expected 0-1 error, i.e.,

$$err_P(f) = \Pr_{(x,y) \sim P} \mathbb{I}_{[\arg \max_{k \in \mathcal{Y}} f(x, k) \neq y]} \quad (2)$$

$$= \Pr_{(x,y) \sim P} \mathbb{I}_{[\rho(f; x, y) < 0]}, \quad (3)$$

where $\mathbb{I}_{[\cdot]}$ is the indicator function.

We call the 0-1 error on the training set *training error* and that on the test set *test error*. Since the expected 0-1 error cannot be obtained due to the unknown distribution P , one usually uses the test error as its proxy when examining the classification accuracy.

Now, we consider using neural networks to fulfill the multi-class classification task. Suppose there are L layers in a neural network, including $L - 1$ hidden layers and an output layer. There are n_l units in layer l ($l = 1, \dots, L$). The number of units in the output layer is fixed by the classification problem, i.e., $n_L = K$. There are weights associated with the edges between units in adjacent layers of the neural network. To avoid over fitting, people usually constraint the size of the weights, e.g., impose a constraint A on the sum of the weights for each unit. We give a unified formulation for both fully connected and convolutional neural networks. Mathematically, we denote the function space of multi-layer neural networks with depth L , and weight constraint A as \mathcal{F}_A^L , i.e.,

$$\mathcal{F}_A^L = \left\{ (x, k) \rightarrow \sum_{i=1}^{n_{L-1}} w_i f_i(x); f_i \in \mathcal{F}_A^{L-1}, \sum_{i=1}^{n_{L-1}} |w_i| \leq A, w_i \in \mathbb{R} \right\}; \quad (4)$$

for $l = 1, \dots, L - 1$,

$$\mathcal{F}_A^l = \left\{ x \rightarrow \varphi \left(\phi(f_1(x)), \dots, \phi(f_{p_l}(x)) \right); f_1, \dots, f_{p_l} \in \bar{\mathcal{F}}_A^l \right\}, \quad (5)$$

$$\bar{\mathcal{F}}_A^l = \left\{ x \rightarrow \sum_{i=1}^{n_{l-1}} w_i f_i(x); f_i \in \mathcal{F}_A^{l-1}, \sum_{i=1}^{n_{l-1}} |w_i| \leq A, w_i \in \mathbb{R} \right\}; \quad (6)$$

and,

$$\mathcal{F}_A^0 = \left\{ x \rightarrow x_i; i \in \{1, \dots, d\} \right\}; \quad (7)$$

where w_i denotes the weight in the neural network, x_i is the i -th dimension of input x . The functions φ and ϕ are defined as follows:

(1) If the l -th layer is a convolutional layer, the outputs of the $(l - 1)$ -th layer are mapped to the l -th layer by means of filter, activation, and then pooling. That is, in Eqn (6), lots of weights equal 0, and n_l is determined by n_{l-1} as well as the number and domain size of the filters. In Eqn (5), p_l equals the size of the pooling region in the l -th layer, and function $\varphi : \mathbb{R}^{p_l} \rightarrow \mathbb{R}$ is called the *pooling function*. Widely-used pooling functions include

¹One related work is (Li et al. 2015), which combines the generative deep learning methods (e.g., RBM) with a margin-max posterior. In contrast, our approach aims to enlarge the margin of discriminative deep learning methods like DNN.

the max-pooling $\max(t_1, \dots, t_{p_l})$ and the average-pooling $(t_1 + \dots + t_{p_l})/p_l$. Function ϕ is increasing and usually called the *activation function*. Widely-used activation functions include the standard sigmoid function $\phi(t) = \frac{1}{1+e^{-t}}$, the tanh function $\phi(t) = \frac{e^t - e^{-t}}{e^t + e^{-t}}$, and the rectifier function $\phi(t) = \max(0, t)$. Please note that all these activation functions are 1-Lipschitz.

(2) If the l -th layer is a fully connected layer, the outputs of the $(l-1)$ -th layer are mapped to the l -th layer by linear combination and subsequently activation. That is, in Eqn (5) $p_l = 1$ and $\varphi(x) = x$.

Because distribution P is unknown and the 0-1 error is non-continuous, a common way of learning the weights in the neural network is to minimize the empirical (surrogate) loss function. A widely used loss function is the cross entropy loss, which is defined as follows,

$$C(f; x, y) = - \sum_{k=1}^K z_k \ln \sigma(x, k), \quad (8)$$

where $z_k = 1$ if $k = y$, and $z_k = 0$ otherwise. Here $\sigma(x, k) = \frac{\exp(f(x, k))}{\sum_{j=1}^K \exp(f(x, j))}$ is the softmax operation that normalizes the outputs of the neural network to a distribution.

Back-propagation algorithm is usually employed to minimize the loss functions, in which the weights are updated by means of stochastic gradient descent (SGD).

3 The Role of Depth in Deep Neural Networks

In this section, we analyze the role of depth in DNN, from the perspective of the margin bound. For this purpose, we first give the definitions of empirical margin error and Rademacher Average (RA), and then introduce the margin bound for multi-class classification.

Definition 1. Suppose $f \in \mathcal{F} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a multi-class prediction model. For $\forall \gamma > 0$, the empirical margin error of f at margin coefficient γ is defined as follows:

$$\text{err}_S^\gamma(f) = \frac{1}{m} \sum_{i=1}^m \mathbb{I}_{[\rho(f; x_i, y_i) \leq \gamma]}. \quad (9)$$

Definition 2. Suppose $\mathcal{F} : \mathcal{X} \rightarrow \mathbb{R}$ is a model space with a single dimensional output. The Rademacher average (RA) of \mathcal{F} is defined as follows:

$$R_m(\mathcal{F}) = \mathbf{E}_{\mathbf{x}, \sigma} \left[\sup_{f \in \mathcal{F}} \left| \frac{2}{m} \sum_{i=1}^m \sigma_i f(x_i) \right| \right], \quad (10)$$

where $\mathbf{x} = \{x_1, \dots, x_m\} \sim P_x^m$, and $\{\sigma_1, \dots, \sigma_m\}$ are i.i.d. sampled with $P(\sigma_i = 1) = 1/2, P(\sigma_i = -1) = 1/2$.

Theorem 1. (Koltchinskii and Panchenko 2002) Suppose $f \in \mathcal{F} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a multi-class prediction model. For $\forall \delta > 0$, with probability at least $1 - \delta$, we have, $\forall f \in \mathcal{F}$,

$$\begin{aligned} \text{err}_P(f) &\leq \inf_{\gamma > 0} \left\{ \text{err}_S^\gamma(f) + \frac{8K(2K-1)}{\gamma} R_m(\tilde{\mathcal{F}}) \right. \\ &\quad \left. + \sqrt{\frac{\log \log_2(2\gamma^{-1})}{m}} + \sqrt{\frac{\log(2\delta^{-1})}{2m}} \right\}. \end{aligned} \quad (11)$$

where $\tilde{\mathcal{F}} = \{x \rightarrow f(\cdot, k); k \in \mathcal{Y}, f \in \mathcal{F}\}$

According to the margin bound given in Theorem 1, the expected 0-1 error of a DNN model can be upper bounded by the sum of two terms, RA and the empirical margin error. In the next two subsections, we will make discussions on the role of depth in these two terms, respectively.

3.1 Rademacher Average

In this subsection, we study the role of depth in the RA-based capacity term.

In the following theorem, we derive an uniform upper bound of RA for both the fully-connected and convolutional neural networks.²

Theorem 2. Suppose input space $\mathcal{X} = [-M, M]^d$. In the deep neural networks, if activation function ϕ is L_ϕ -Lipschitz and non-negative, pooling function φ is max-pooling or average-pooling, and the size of pooling region in each layer is bounded, i.e., $p_l \leq p$, then we have,

$$R_m(\mathcal{F}_A^L) \leq cM \sqrt{\frac{\ln d}{m}} (pL_\phi A)^L. \quad (12)$$

where c is a constant.

Proof. According to the definition of \mathcal{F}_A^L and RA, we have,

$$\begin{aligned} R_m(\mathcal{F}_A^L) &= \mathbf{E}_{\mathbf{x}, \sigma} \left[\sup_{\|\mathbf{w}\|_1 \leq A, f_j \in \mathcal{F}_A^{L-1}} \left| \frac{2}{m} \sum_{i=1}^m \sigma_i \sum_{j=1}^{n_{L-1}} w_j f_j(x_i) \right| \right] \\ &= \mathbf{E}_{\mathbf{x}, \sigma} \left[\sup_{\|\mathbf{w}\|_1 \leq A, f_j \in \mathcal{F}_A^{L-1}} \left| \frac{2}{m} \sum_{j=1}^{n_{L-1}} w_j \sum_{i=1}^m \sigma_i f_j(x_i) \right| \right]. \end{aligned}$$

Supposing $\mathbf{w} = \{w_1, \dots, w_{n_{L-1}}\}$ and $\mathbf{h} = \{\sum_{i=1}^m \sigma_i f_1(x_i), \dots, \sum_{i=1}^m \sigma_i f_{n_{L-1}}(x_i)\}$, the inner product $\langle \mathbf{w}, \mathbf{h} \rangle$ is maximized when \mathbf{w} is at one of the extreme points of the l_1 ball, which implies:

$$\begin{aligned} R_m(\mathcal{F}_A^L) &\leq A \mathbf{E}_{\mathbf{x}, \sigma} \left[\sup_{f \in \mathcal{F}_A^{L-1}} \left| \frac{2}{m} \sum_{i=1}^m \sigma_i f(x_i) \right| \right] \\ &= AR_m(\mathcal{F}_A^{L-1}). \end{aligned} \quad (13)$$

For function class \mathcal{F}_A^{L-1} , if the $(L-1)$ -th layer is a fully connected layer, it is clear that $R_m(\mathcal{F}_A^{L-1}) \leq R_m(\phi \circ \tilde{\mathcal{F}}_A^{L-1})$ holds. If the $(L-1)$ -th layer is a convolutional layer with max-pooling or average-pooling, we have,

$$\begin{aligned} R_m(\mathcal{F}_A^{L-1}) &\leq \mathbf{E}_{\mathbf{x}, \sigma} \left[\sup_{f_1, \dots, f_{p_{L-1}} \in \tilde{\mathcal{F}}_A^{L-1}} \left| \frac{2}{m} \sum_{i=1}^m \sigma_i \sum_{j=1}^{p_{L-1}} \phi(f_j(x_i)) \right| \right] \\ &= p_{L-1} R_m(\phi \circ \tilde{\mathcal{F}}_A^{L-1}). \end{aligned} \quad (14)$$

The inequality (14) holds due to the fact that most widely used activation functions ϕ (e.g., standard sigmoid and rectifier) have non-negative outputs.

Therefore, for both fully connected layers and convolutional layers, $R_m(\mathcal{F}_A^{L-1}) \leq p_{L-1} R_m(\phi \circ \tilde{\mathcal{F}}_A^{L-1})$ uniformly

²To the best of our knowledge, an upper bound of RA for fully connected neural networks has been derived before (Bartlett and Mendelson 2003; Neyshabur, Tomioka, and Srebro 2015), but there is no result available for the convolutional neural networks.

holds. Further considering the Lipschitz property of ϕ , we have,

$$R_m(\mathcal{F}_A^{L-1}) \leq 2p_{L-1}L_\phi R_m(\bar{\mathcal{F}}_A^{L-1}). \quad (15)$$

Iteratively using maximization principle of inner product in (13), property of RA in (14) and Lipschitz property in (15), considering $p_l \leq p$, we can obtain the following inequality,

$$R_m(\mathcal{F}_A^L) \leq (2pL_\phi A)^{L-1} R_m(\bar{\mathcal{F}}_A^1). \quad (16)$$

According to (Bartlett and Mendelson 2003), $R_m(\bar{\mathcal{F}}_A^1)$ can be bounded by:

$$R_m(\bar{\mathcal{F}}_A^1) \leq cAM\sqrt{\frac{\ln d}{m}}, \quad (17)$$

where c is a constant.

Combining (16) and (17), we can obtain the upper bound on the RA of DNN. \square

From the above theorem, we can see that with the increasing depth, the upper bound of RA will increase, and thus the margin bound will become looser. This indicates that depth has its negative impact on the test performance of DNN.

3.2 Empirical Margin Error

In this subsection, we study the role of depth in empirical margin error.

To this end, we first discuss representation power of DNN models. In particular, we use the Betti numbers based complexity (Bianchini and Scarselli 2014) to measure the representation power. We generalize the definition of Betti numbers based complexity into multi-class setting as follows.

Definition 3. *The Betti numbers based complexity of functions implemented by multi-class neural networks \mathcal{F}_A^L is defined as $N(\mathcal{F}_A^L) = \sum_{i=1}^{K-1} B(S_i)$, where $B(S_i)$ is the sum of Betti numbers³ that measures the complexity of the set S_i . Here $S_i = \cap_{j=1, j \neq i}^K \{x \in \mathbb{R}^d \mid f(x, i) - f(x, j) \geq 0; f(x, \cdot) \in \mathcal{F}_A^L\}$, $i = 1, \dots, K-1$.*

As can be seen from the above definition, the Betti numbers based complexity considers classification output and merge those regions corresponding to the same classification output (thus is more accurate than the linear region number complexity (Montufar et al. 2014) in measuring the representation power). As far as we know, only for binary classification and fully connected networks, the bounds of the Betti numbers based complexity was derived (Bianchini and Scarselli 2014), and there is no result for the setting of multi-class classification and convolutional networks. In the following, we give our own theorem to fill in this gap.

Theorem 3. *For neural networks \mathcal{F}_A^L that has h hidden units. If activation function ϕ is a Pfaffian function with complexity (α, β, η) , pooling function φ is average-pooling and*

³For any subset $S \subset \mathbb{R}^d$, there exist d Betti numbers, denoted as $b_j(S)$, $0 \leq j \leq d-1$. Therefore, the sum of Betti numbers is denoted as $B(S) = \sum_{j=0}^{d-1} b_j(S)$. Intuitively, the first Betti number $b_0(S)$ is the number of connected components of the set S , while the j -th Betti number $b_j(S)$ counts the number of $(j+1)$ -dimension holes in S (Bianchini and Scarselli 2014).

$d \leq h\eta$, then

$$N(\mathcal{F}_A^L) \leq (K-1)^{d+1} 2^{h\eta(h\eta-1)/2} \times O\left((d((\alpha+\beta-1+\alpha\beta)(L-1)+\beta(\alpha+1)))^{d+h\eta}\right) \quad (18)$$

Proof. We first show that the functions $f(x, \cdot) \in \mathcal{F}_A^L$ are Pfaffian functions with complexity $((\alpha+\beta-1+\alpha\beta)(L-1)+\alpha\beta, \beta, h\eta)$, where \mathcal{F}_A^L can contain both fully-connected layers and convolutional layers. Assume the Pfaffian chain which defines activation function $\phi(t)$ is $(\phi_1(t), \dots, \phi_\eta(t))$, and then s^l is constructed by applying all ϕ_i , $1 \leq i \leq \eta$ on all the neurons up to layer $l-1$, i.e., $f^l \in \bar{\mathcal{F}}_A^l$, $l \in \{1, \dots, L-1\}$. As the first step, we need to get the degree of f^l in the chain s^l . Since $f^l = \frac{1}{p_{l-1}} \sum_{k=1}^{n_{l-1}} w_k(\phi(f_{k,1}^{l-1}) + \dots + \phi(f_{k,p_{l-1}}^{l-1}))$ and ϕ is a Pfaffian function, f^l is a polynomial of degree β in the chain s^l . Then, it remains to show that the derivative of each function in s^l , i.e., $\frac{\partial \phi_i(f^l)}{\partial x_i} = \frac{d\phi_i(f^l)}{df^l} \frac{\partial f^l}{\partial x_i}$, can be defined as a polynomial in the functions of the chain and the input. For average pooling, by iteratively using chain rule, we can obtain that the highest degree terms of $\frac{\partial f^l}{\partial x_i}$ are in the form of $\prod_{i=1}^{l-1} \frac{d\phi(f^i)}{df^i}$. Following the lemma 2 in (Bianchini and Scarselli 2014), we obtain the complexity of $f(x, \cdot) \in \mathcal{F}_A^L$.

Furthermore, the sum of two Pfaffian functions f_1 and f_2 defined by the same Pfaffian chain of length η with complexity $(\alpha_1, \beta_1, \eta)$ and $(\alpha_2, \beta_2, \eta)$ respectively is a Pfaffian function with complexity $(\max(\alpha_1, \alpha_2), \max(\beta_1, \beta_2), \eta)$ (Gabrielov and Vorobjov 2004). Therefore, $f(x, i) - f(x, j)$, $i \neq j$ is a Pfaffian function with complexity $((\alpha+\beta-1+\alpha\beta)(L-1)+\alpha\beta, \beta, h\eta)$.

According to (Zell 1999), since S_i is defined by $K-1$ sign conditions (inequalities or equalities) on Pfaffian functions, and all the functions defining S_i have complexity at most $((\alpha+\beta+\alpha\beta)(L-1)+\alpha\beta, \beta, h\eta)$, $B(S_i)$ can be upper bounded by $(K-1)^{d2^{h\eta(h\eta-1)/2}} \times O((d((\alpha+\beta-1+\alpha\beta)(L-1)+\beta(\alpha+1)))^{d+h\eta})$.

Summing over all $i \in \{1, \dots, K-1\}$, we get the results stated in Theorem 3. \square

Theorem 3 upper bounds the Betti numbers based complexity for general activation functions. For specific active functions, we can get the following results: when $\phi = \arctan(\cdot)$ and $d \leq 2h$, since \arctan is of complexity $(3, 1, 2)$, we have $N(\mathcal{F}_A^L) \leq (K-1)^{d+1} 2^{h(2h-1)} O((d(L-1)+d)^{d+2h})$; when $\phi = \tanh(\cdot)$ and $n \leq h$, since \tanh is of complexity $(2, 1, 1)$, we have $N(\mathcal{F}_A^L) \leq (K-1)^{d+1} 2^{h(h-1)/2} O((d(L-1)+d)^{d+h})$.

Basically, Theorem 3 indicates that in the multi-class setting, the Betti numbers based complexity grows with the increasing depth L . As a result, deeper nets will have larger representation power than shallower nets, which makes deeper nets fit better to the training data and achieve smaller empirical margin error. This indicates that depth has its positive impact on the test performance of DNN.

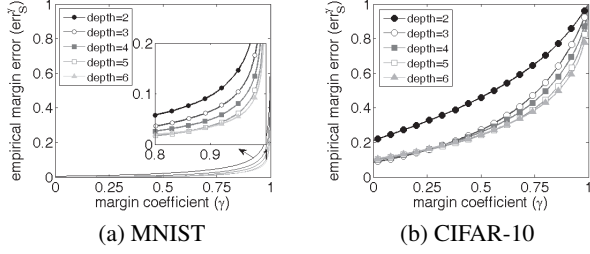


Figure 1: The influence of depth on empirical margin error.

Actually, above discussions about impact of depth on representation power are consistent with our empirical findings. We conducted experiments on two datasets, MNIST (LeCun et al. 1998) and CIFAR-10 (Krizhevsky 2009). To investigate the influence of network depth L , we trained fully-connected DNN with different depths and restricted number of hidden units. The experimental results are shown in Figure 1 and indicate that no matter on which dataset, deeper networks have smaller empirical margin errors than shallower networks for most of the margin coefficients.

3.3 Discussions

Based on discussions in previous two subsections, we can see that when the depth L of DNN increases, (1) the RA term in margin bound will increase (according to Theorem 2); (2) the empirical margin error in margin bound will decrease since deeper nets have larger representation power (according to Theorem 3). As a consequence, we can come to the conclusion that, for DNN with restricted number of hidden units, arbitrarily increasing depth is not always good since there is a clear tradeoff between its positive and negative impacts on test error. In other words, with the increasing depth, the test error of DNN may first decrease, and then increase.

Actually this theoretical pattern is consistent with our empirical observations on different datasets. We used the same experimental setting as that in the subsection 3.2 and repeated the training of DNN (with different random initializations) for 5 times. Figure 2 reports the average and minimum test error of 5 learned models. We can observe that as the depth increases, the test error first decreases (probably because increased representation power overwhelms increased RA capacity); and then increase (probably because RA capacity increases so quickly that representation power cannot compensate for negative impact of increased capacity).

4 Large Margin Deep Neural Networks

From the discussions in Section 3, we can see that one may have to pay the cost of larger RA capacity when trying to obtain better representation power by increasing the depth of DNN (not to mention that the effective training of very deep neural networks is highly non-trivial (Glorot and Bengio 2010; Srivastava, Greff, and Schmidhuber 2015)). Then a nature question is whether we can avoid this tradeoff, and achieve good test performance in an alternative way.

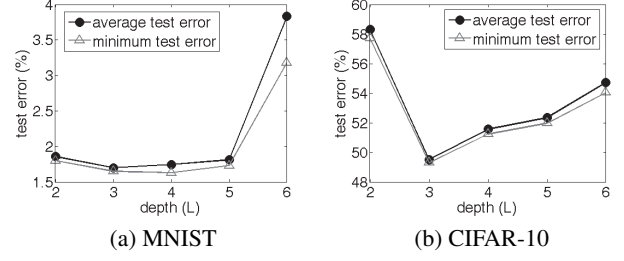


Figure 2: The influence of depth on test error.

To this end, let us revisit the positive impact of depth: it actually lies in that deeper neural networks tend to have larger representation power and thus smaller empirical margin error. Then the question is: can we directly minimize empirical margin error? Our answer to this question is yes, and our proposal is to add a margin-based penalty term to current loss function. In this way, we should be able to effectively tighten margin bound without manipulating the depth.

One may argue that widely used loss functions (e.g., cross entropy loss and hinge loss) in DNN are convex surrogates of margin error by themselves, and it might be unnecessary to introduce an additional margin-based penalty term. However, we would like to point out that unlike hinge loss for SVM or exponential loss for Adaboost, which have theoretical guarantee for convergence to margin maximizing separators as the regularization vanishes (Rosset, Zhu, and Hastie 2003), there is no optimization consistency guarantee for these losses used in DNN since neural networks are highly non-convex. Therefore, it makes sense to explicitly add a margin-based penalty term to loss function, in order to further reduce empirical margin error during training process.

4.1 Algorithm Description

We propose adding two kinds of margin-based penalty terms to the original cross entropy loss⁴. The first penalty term is the gap between the upper bound of margin (i.e., 1)⁵ and the margin of the sample (i.e., $\rho(f; x, y)$). The second one is the average gap between upper bound of margin and the difference between the predicted output for the true category and those for all the wrong categories. It can be easily verified that the second penalty term is an upper bound of the first penalty term. Mathematically, the penalized loss functions can be described as follows (for ease of reference, we call them C_1 and C_2 respectively): for model f , sample x, y ,

$$C_1(f; x, y) = C(f; x, y) + \lambda \left(1 - \rho(f; x, y)\right)^2,$$

$$C_2(f; x, y) = C(f; x, y) + \frac{\lambda}{K-1} \sum_{k \neq y} \left(1 - (f(x, y) - f(x, k))\right)^2.$$

⁴Although we take the most widely-used cross entropy loss as example, these margin-based penalty terms can also be added to other loss functions.

⁵Please note that, after softmax operation, the outputs are normalized to $[0, 1]$

	MNIST	CIFAR-10
DNN- C (%)	0.899 \pm 0.038	18.339 \pm 0.336
LMDNN- C_1 (%)	0.734 \pm 0.046	17.598 \pm 0.274
LMDNN- C_2 (%)	0.736 \pm 0.041	17.728 \pm 0.283

Table 1: Test error (%) of DNN- C and LMDNNs.

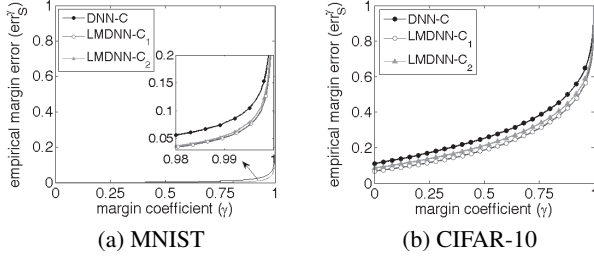


Figure 3: Empirical margin error of LMDNNs.

We call the algorithms that minimize the above new loss functions *large margin DNN algorithms* (LMDNN). For ease of reference, we denote LMDNN minimizing C_1 and C_2 as LMDNN- C_1 and LMDNN- C_2 respectively, and the standard DNN algorithms minimizing C as DNN- C . To train LMDNN, we also employ the back propagation method.

4.2 Experimental Results

Now we compare the performances of LMDNNs with DNN- C . We used well-tuned network structures in the Caffe (Jia et al. 2014) tutorial (i.e., LeNet⁶ for MNIST and AlexNet⁷ for CIFAR-10) and all the tuned hyper parameters on the validation set.

Each model was trained for 10 times with different initializations. Table 1 shows mean and standard deviation of test error over the 10 learned models for DNN- C and LMDNNs after tuning margin penalty coefficient λ . We can observe that, on both MNIST and CIFAR-10, LMDNNs achieve significant performance gains over DNN- C . In particular, LMDNN- C_1 reduce test error from 0.899% to 0.734% on MNIST and from 18.399% to 17.598% on CIFAR-10; LMDNN- C_2 reduce test error from 0.899% to 0.736% on MNIST and from 18.399% to 17.728% on CIFAR-10.

To further understand the effect of adding margin-based penalty terms, we plot empirical margin errors of DNN- C and LMDNNs in Figure 3. We can see that by introducing margin-based penalty terms, LMDNNs indeed achieve smaller empirical margin errors than DNN- C . Furthermore, the models with smaller empirical margin errors really have better test performances. For example, LMDNN- C_1 achieved both smaller empirical margin error and better test performance than LMDNN- C_2 . This is consistent with Theorem 3, and in return indicates reasonability of our theorem.

⁶<http://caffe.berkeleyvision.org/gathered/examples/mnist.html>

⁷<http://caffe.berkeleyvision.org/gathered/examples/cifar10.html>

html

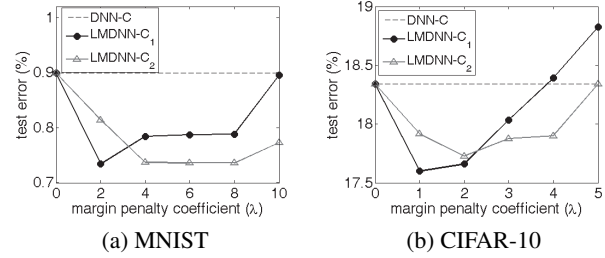


Figure 4: Test error of LMDNNs with different λ .

We also report mean test error of LMDNNs with different margin penalty coefficient λ (see Figure 4). In the figure, we use dashed line to represent mean test error of DNN- C (corresponding to $\lambda = 0$). From the figure, we can see that on both MNIST and CIFAR-10, (1) there is a range of λ where LMDNNs outperform DNN- C ; (2) although the best test performance of LMDNN- C_2 is not as good as that of LMDNN- C_1 , the former has a broader range of λ that can outperform DNN- C in terms of the test error. This indicates the value of using LMDNN- C_2 : it eases the tuning of hyper parameter λ ; (3) with increasing λ , test error of LMDNNs will first decrease, and then increase. When λ is in a reasonable range, LMDNNs can leverage both good the optimization property of cross entropy loss in training process and the effectiveness of margin-based penalty term, and thus achieve good test performance. When λ becomes too large, margin-based penalty term dominates cross entropy loss. Considering that margin-based penalty term may not have good optimization property as cross entropy loss in the training process, the drop of test error is understandable.

5 Conclusion and Future Work

In this work, we have investigated the role of depth in DNN from the perspective of margin bound. We find that while the RA term in margin bound is increasing w.r.t. depth, the empirical margin error is decreasing instead. Therefore, arbitrarily increasing the depth might not be always good, since there is a tradeoff between the positive and negative impacts of depth on test performance of DNN. Inspired by our theory, we propose two large margin DNN algorithms, which achieve significant performance gains over standard DNN algorithm. In the future, we plan to study how other factors influence the test performance of DNN, such as unit allocations across layers and regularization tricks. We will also work on the design of effective algorithms that can further boost the performance of DNN.

6 Acknowledgments

Liwei Wang was partially supported by National Basic Research Program of China (973 Program) (grant no. 2015CB352502), NSFC(61573026), and a grant from MOE-Microsoft Laboratory of Statistics of Peking University. Xiaoguang Liu was partially supported by NSF of China (61373018, 11301288, 11450110409) and Program for New Century Excellent Talents in University (NCET130301).

References

- Bartlett, P. L., and Mendelson, S. 2003. Rademacher and gaussian complexities: Risk bounds and structural results. *The Journal of Machine Learning Research* 3:463–482.
- Bartlett, P. L.; Maierov, V.; and Meir, R. 1998. Almost linear vc-dimension bounds for piecewise polynomial networks. *Neural computation* 10(8):2159–2173.
- Bartlett, P. L. 1998. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE Transactions on Information Theory* 44(2):525–536.
- Bianchini, M., and Scarselli, F. 2014. On the complexity of neural network classifiers: A comparison between shallow and deep architectures. *IEEE Transactions on Neural Networks*.
- Ciresan, D.; Meier, U.; and Schmidhuber, J. 2012. Multi-column deep neural networks for image classification. In *Computer Vision and Pattern Recognition*, 3642–3649.
- Delalleau, O., and Bengio, Y. 2011. Shallow vs. deep sum-product networks. In *Advances in Neural Information Processing Systems*, 666–674.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, 248–255. IEEE.
- Gabrielov, A., and Vorobjov, N. 2004. Complexity of computations with pfaffian and noetherian functions. *Normal forms, bifurcations and finiteness problems in differential equations* 211–250.
- Glorot, X., and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*, 249–256.
- Goldberg, P. W., and Jerrum, M. R. 1995. Bounding the vapnik-chervonenkis dimension of concept classes parameterized by real numbers. *Machine Learning* 18(2-3):131–148.
- Hastad, J. 1986. Almost optimal lower bounds for small depth circuits. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, 6–20. ACM.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *arXiv preprint arXiv:1502.01852*.
- Hinton, G. E., and Salakhutdinov, R. R. 2006. Reducing the dimensionality of data with neural networks. *Science* 313(5786):504–507.
- Hinton, G. E.; Deng, L.; Yu, D.; Dahl, G. E.; Mohamed, A.-r.; Jaitly, N.; Senior, A.; Vanhoucke, V.; Nguyen, P.; Sainath, T. N.; et al. 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine* 29(6):82–97.
- Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; and Darrell, T. 2014. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*.
- Karpinski, M., and Macintyre, A. 1995. Polynomial bounds for vc dimension of sigmoidal neural networks. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, 200–208. ACM.
- Koltchinskii, V., and Panchenko, D. 2002. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics* 1–50.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 1097–1105.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.
- Lee, C.-Y.; Xie, S.; Gallagher, P.; Zhang, Z.; and Tu, Z. 2014. Deeply-supervised nets. *arXiv preprint arXiv:1409.5185*.
- Li, C.; Zhu, J.; Shi, T.; and Zhang, B. 2015. Max-margin deep generative models. In *Advances in neural information processing systems*.
- Montufar, G. F.; Pascanu, R.; Cho, K.; and Bengio, Y. 2014. On the number of linear regions of deep neural networks. In *Advances in Neural Information Processing Systems*, 2924–2932.
- Neyshabur, B.; Tomioka, R.; and Srebro, N. 2015. Norm-based capacity control in neural networks. *arXiv preprint arXiv:1503.00036*.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Rosset, S.; Zhu, J.; and Hastie, T. J. 2003. Margin maximizing loss functions. In *Advances in neural information processing systems*, 1237–1244.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Srivastava, R. K.; Greff, K.; and Schmidhuber, J. 2015. Training very deep networks. In *Advances in neural information processing systems*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2014. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*.
- Weston, J.; Ratle, F.; Mobahi, H.; and Collobert, R. 2012. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*. Springer. 639–655.
- Zell, T. 1999. Betti numbers of semi-pfaffian sets. *Journal of Pure and Applied Algebra* 139(1):323–338.