Maximum-and-Concatenation Networks

Xingyu Xie [1] **Hao Kong** [1] **Jianlong Wu** [2] **Wayne Zhang** [3] **Guangcan Liu** [4] **Zhouchen Lin** [1]

Abstract

While successful in many fields, deep neural networks (DNNs) still suffer from some open problems such as bad local minima and unsatisfactory generalization performance. In this work, we propose a novel architecture called Maximum-and-Concatenation Networks (MCN) to try eliminating bad local minima and improving generalization ability as well. Remarkably, we prove that MCN has a very nice property; that is, *every local minimum of an $(l+1)$-layer MCN can be better than, at least as good as, the global minima of the network consisting of its first l layers*. In other words, by increasing the network depth, MCN can autonomously improve its local minima's goodness, what is more, *it is easy to plug MCN into an existing deep model to make it also have this property*. Finally, under mild conditions, we show that MCN can approximate certain continuous functions arbitrarily well with *high efficiency*; that is, the covering number of MCN is much smaller than most existing DNNs such as deep ReLU. Based on this, we further provide a tight generalization bound to guarantee the inference ability of MCN when dealing with testing samples.

1. Introduction

Deep neural networks (DNNs) have been showing superior performance in various fields such as computer vision, speech recognition, natural language processing, and so on. At the first glance, DNN learning is not an enigmatic technique, as its basic idea is quite simple and mostly about learning a possibly over-parameterized DNN from a huge number of training samples; namely,

$$\min_{\boldsymbol{\theta}} L(\boldsymbol{\theta}) := \frac{1}{n}\sum_{i=1}^{n} \ell(\mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x}_i), \mathbf{y}_i), \tag{1}$$

where $\mathbf{x}_i \in \mathbb{R}^{d_x}$ and $\mathbf{y}_i \in \mathbb{R}^{d_y}$ denote an input and a target, respectively, $\mathbf{f}_{\boldsymbol{\theta}}(\cdot)$ standards for a DNN with parameters $\boldsymbol{\theta}$, and $\ell(\cdot)$ is some loss function. Notice that, some kind of regularization schema has already been implanted into the network to constrain the parameter space, though there is no explicit regularizer imposed on $\boldsymbol{\theta}$ (Arora et al., 2019a). Despite its ordinary appearance, DNN learning is meanwhile quite complicated in many ways, and the current DNNs still suffer from several weaknesses, e.g., the training procedure may easily get stuck in *bad local minima* (i.e., the local minima with large training error), the learnt model may be prone to *over-fit the training data* (i.e., the testing error is large when small training error is obtained), etc. Overcoming these difficulties are crucial for DNNs to solve the real-world problems that are more challenging and significant, but they are still *open* problems.

To address the issue of bad local minima, many heuristic techniques have been proposed, e.g., batch normalization (Ioffe & Szegedy, 2015), group normalization (Wu & He, 2018), dropout (Srivastava et al., 2014), etc. These techniques would be useful under certain context, but may not be generally helpful and, even worse, it is hard to know when and which method should be used. In fact, the elimination of bad local minima, i.e., having small empirical training error at *all* local minima, is really important for DNN learning. Some recent theories (Zhang et al., 2017; Wei & Ma, 2019; Cao & Gu, 2019; Li & Liang, 2018; Allen-Zhu et al., 2018; Arora et al., 2019c) have revealed that, whenever the local minima produces only small training error, DNNs have probably good generalization performance at these local minima. That is to say, in some cases, good local minima mean good predictors which are the ultimate goal of supervised learning. With the hope of pursuing the property of *no bad local minima*, some learning theories (Kawaguchi, 2016; Arora et al., 2018; Hardt & Ma, 2016; Liang et al., 2018a;b) have been established to prove that, under certain conditions, any local minima of a certain DNN are also global minima. While impressive, existing studies are still unsatisfactory in some aspects:

- Most existing theories about "all local minima are

[1]Key Lab. of Machine Perception (MoE), School of EECS, Peking University [2]School of Computer Science and Technology, Shandong University [3]SenseTime Research [4]B-DAT and CICAEET, School of Automation, Nanjing University of Information Science and Technology. Correspondence to: Guangcan Liu <gcliu@nuist.edu.cn>, Zhouchen Lin <zlin@pku.edu.cn>.

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

global minima" are built upon on some unrealistic network architectures, e.g., without activation function, which means that they cannot be applied to common deep learning tasks. The work (Kawaguchi & Kaelbling, 2019) considers general architectures, but requires additional regularizer and is limited to shallow case. In addition, strictly speaking, the conclusion of "all local minima are global minima" cannot really ensure that "DNN has no bad local minima". This is because, whenever the adopted network itself is poorly designed, global minima can still lead to large training error. In one word, existing studies have not gained convenient schemes that can be easily used to reduce the training error of general DNNs.

- Though small training error may bring good generalization for some specially designed DNNs (Zhang et al., 2017; Wei & Ma, 2019; Cao & Gu, 2019; Li & Liang, 2018; Allen-Zhu et al., 2018; Arora et al., 2019c), a rigorous generalization bound is still important for general DNNs to produce superior performance in practice. There is sparse research in the direction of generalization analysis, e.g., deep ReLU (Yarotsky, 2017). However, the covering number in deep ReLU is very large, which means that the approximation ability of the network is rather weak.

- What is more, to our knowledge, there is no theoretical study that addresses the issues of local minima and generalization ability simultaneously. These two problems are closely related and should be investigated at the same time.

To relieve the issues highlighted above, we propose a novel multi-layer DNN termed **M**aximum-and-**C**oncatenation **N**etworks (MCN). In our MCN, one hidden layer is formed by concatenating together two parts, with one being a linear transformation of the output of the previous layer, and the other being a maximum of two piecewise smooth functions. The output of the final layer is further transformed by some linear operators, so as to stay in step with the configuration of the target output. In general, the concatenation operator is a good option during designing DNNs, and it is indeed a primary cause of the superiorities of MCN over existing architectures.

We prove that MCN naturally ensures the effectiveness of its learning process, i.e., the no bad local minima property. To be more precise, suppose that $\boldsymbol{\theta}'$ is a global minimum to (1) with $\mathbf{f}_{\boldsymbol{\theta}'}$ being an l-layer MCN (briefly, we say that $\boldsymbol{\theta}'$ is a global minimum of an l-layer MCN), and $\boldsymbol{\theta}$ is a local minimum of the $(l+1)$-layer MCN obtained by adding one layer to the former l-layer network. Then we have $L(\boldsymbol{\theta}) \leq L(\boldsymbol{\theta}')$, which means that *the global minima of an l-layer MCN may be outperformed, at least can be attained, by simply increasing the network depth*. More importantly, MCN can be easily appended to many existing network architectures, and we prove that, under mild conditions, *the modified DNN will get the nice properties of MCN*. This property is achieved mainly due to a *skip connection* with a proper activation function: With the help of this skip connection, the bad local minima are moved to infinity, while the implicit regularizer carried by the network itself may encourage the optimization procedure to seek for the remaining good local minima.

Notice that, piecewise liner functions can approximate any Lipschitz continuous function up to arbitrarily small error, and the maximum operator can model the piecewise linear function efficiently (Telgarsky, 2016). Based on these facts, we show that MCN with *sparse* connection can approximate a wide range of continuous functions arbitrarily well. Our analysis framework is new and quite different from the previous studies (Lu et al., 2020; Yarotsky, 2018; 2017), which rely on Taylor expansion and requires a parameter complexity of $\mathcal{O}(N^{d_x})$, where $N \gg 1$ is a quantity that controls the approximation accuracy [1]. By sharp contrast, we show that a complexity of only $\mathcal{O}(N(\ln N)^{d_x-2})$ is enough to approximate the target function.

Based on the approximation analysis, we further investigate the generalization ability of MCN to cope with testing samples, proving that MCN has much smaller covering number than deep ReLU. Interestingly, our results suggest that the width has less effects than the depth on the generalization bound. Our results also show that, whenever the training data are exactly fitted, *MCN achieves the statistically optimal rate in the minmiax sense*; this confirms the conjectures in (Wei & Ma, 2019; Arora et al., 2019c; Belkin et al., 2018b) that ultra-deep networks may generalize well on testing data [2]. To summarize, the contributions of this paper mainly include:

- We propose a novel architecture termed MCN and prove that MCN can help to overcome the issue of bad local minima. Namely, the global minima of an l-layer MCN can be always attained or even outperformed by simply increasing the network depth (Theorem 1). More importantly, we show that MCN is able to turn a possibly poorly-designed DNN into a good one, which also has the nice property of "no bad local minima" under certain conditions (Corollary 4.2 and Theorem 5). These results would be more significant than (Kawaguchi, 2016; Hardt & Ma, 2016), which only show that all local minima of a certain DNN with fixed depth are global minima, but provide no practical guidance for the users to seek better solutions to their

[1] $N^{-\beta}$ is the dominant term in approximation error, where $\beta > 0$ relates to the smoothness of the target function.

[2] Note here that we have no intention to suggest using infinitely deep networks, as the computational cost is also a matter and the required data amount in the extreme case could be huge.

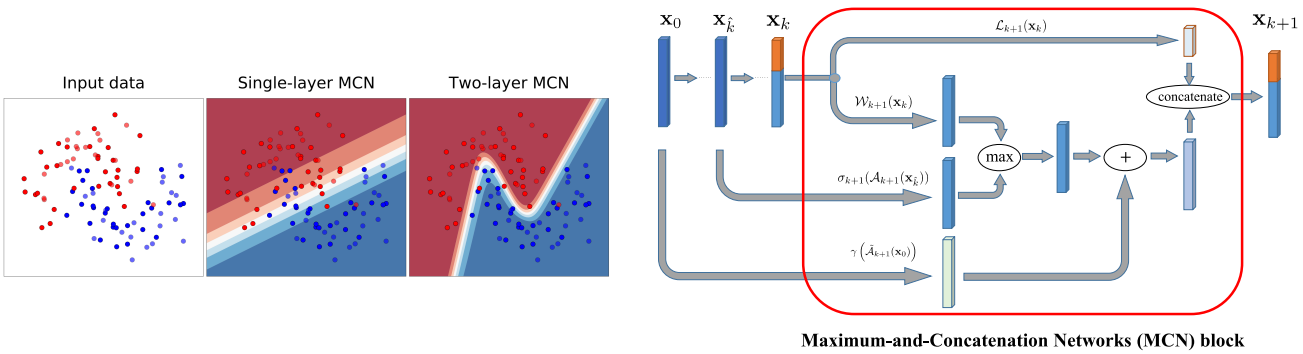

Figure 1. Left: Illustration of the motivations for inventing MCN, which is indeed a generalization of the piecewise smooth function. The composition of MCNs may increase the pieces exponentially. Right: One block of MCN, where each layer consists of four parts.

tasks—just finding the globally optimal solutions to some over-simplified optimization problems is essentially not enough.

- We devise a new framework to analyze the approximation ability of MCN, showing that MCN can approximate some classes of continuous functions arbitrarily well by only using a parameter complexity of $\mathcal{O}(N(\ln N)^{d_x-2})$ (Theorem 2). This is much lower than the $\mathcal{O}(N^{d_x})$ complexity obtained by the previous studies (Lu et al., 2020; Yarotsky, 2018; 2017).
- Unlike the previous analyses in (Liang et al., 2018a;b; Kawaguchi & Kaelbling, 2019), which focus on the elimination of local minima but ignore the generalization performance, we provide rigorous analysis to guarantee the generalization ability of MCN under certain conditions (Theorem 3 and Corollary 3.1). In particular, our results show that MCN has a much smaller covering number than deep ReLU, revealing that the depth is more important than the width for generalization; this supports the mechanism of deep learning.

2. Model and Setting

This section introduces the technical details of MCN, as well as the setup for establishing theoretical analysis.

2.1. Maximum-and-Concatenation Networks

The design of our MCN—a linearity and maximum concatenation network—is inspired by the following observations. Consider the task of shattering some points that are not linearly separable, which is shown in Figure 1. Intuitively, the maximum of two hyperplanes may produce smaller classification error than every single one of them. Therefore, we may reduce the classification error by replacing parts of the current classifier with some maximum-derived units. Such a replacement process can be repeated several times, learning progressively a refined classification surface that will be piecewise smooth. Moreover, considering the regression problem, we have a classical claim from the Stone-Weierstrass approximation theorem.

Claim 1. *Any Lipschitz continuous function can be approximated arbitrarily well by a piecewise linear function.*

By composing a series of maximum operators, we can easily construct a piecewise smooth function. Consider approximating the quadratic function $x \to x^2$. Define the operator $\mathcal{T}^m(x) := \max\{-x/2, x/2 - 2^{1-2m}\}$ and let $g^m(x) := \mathcal{T}^m \circ \mathcal{T}^{m-1} \circ \cdots \mathcal{T}^1(x)$. It is known that $x + \sum_{i=1}^{m} g^i(x)$ approximates x^2 exponentially fast in m (Telgarsky, 2016). In contrast, to approximate a twice differentiable non-piecewise linear function f, it would be awkward to use some existing DNNs that need to rescale the second order differences: $(f(t+2\delta x) - 2f(t+x\delta) + f(x\delta))/(\delta^2 f''(t)) \to x^2$ for $\delta \to 0$ with $f''(t) \neq 0$. Note that $\delta \to 0$ will cause the scale of network parameters to be very large.

Beneath it all, the model of an l-layer MCN, which is indeed a mapping from input $\mathbf{x}$ to output $\mathbf{y}$, is designed as follows, for $k = 0, \cdots, l-1$:

$$\mathbf{x}_{k+1} = \left[\mathcal{L}_{k+1}(\mathbf{x}_k);\ \gamma\left(\tilde{\mathcal{A}}_{k+1}(\mathbf{x}_0)\right) + \mathcal{M}_{k+1}(\mathbf{x}_k)\right], \tag{2}$$

where

$$\mathcal{M}_{k+1}(\mathbf{x}_k) = \max\left\{\mathcal{W}_{k+1}(\mathbf{x}_k), \sigma_{k+1}\left(\mathcal{A}_{k+1}\left(\mathbf{x}_{\hat{k}}\right)\right)\right\},$$

$0 \leq \hat{k} \leq k$ ($\mathbf{x}_{\hat{k}}$ is the output of any intermediate layer between $\mathbf{x}_k$ and $\mathbf{x}_0$), $\gamma(\cdot)$ and $\sigma_{k+1}(\cdot)$ are some element-wise activation functions, $\mathbf{x}_0 = \mathbf{x} \in \mathbb{R}^{d_x}$ is the input data vector, $\mathbf{x}_k \in \mathbb{R}^{d_k}$ is the output of the k-th layer, $[\,;\,]$ is the operator that vertically concatenates two vectors into a single one, $\mathcal{L}_{k+1} : \mathbb{R}^{d_k} \to \mathbb{R}^{d_{\mathcal{L}}}$ is a learnable linear operator[3],

[3]For convenience, we assume that the output of $\mathcal{L}_{k+1}$ has a

and $\mathcal{A}_{k+1}(\cdot)$, $\tilde{\mathcal{A}}_{k+1}(\cdot)$ and $\mathcal{W}_{k+1}(\cdot)$ are all learnable linear operators from $\mathbb{R}^{d_k}$ to $\mathbb{R}^{d_{k+1}-d_{\mathcal{L}}}$.

In fact, as mentioned in Figure 1, MCN is a generalization of piecewise smooth functions, and it can contain many existing DNNs as special cases, e.g., ResNet, Maxout Network (Goodfellow et al., 2013) and Input Convex Neural Networks (ICNN) (Amos et al., 2017). In MCN, there are layers that directly connect the input $\mathbf{x}_0$ to the hidden units in deeper layers. Such connections are unnecessary for traditional networks, but very important for achieving the nice property of "no bad local minimum" which we will introduce later. The highway with the operator $\mathcal{L}_k$ connects the training loss with the geometric projection residual in the proper setting (Section B in supplementary material), which helps MCN perform well when it goes deeper and wider.

2.2. Setting

To analyze MCN theoretically, we consider a typical task of regression (or classification). Denote by $\mathbf{x} \in \mathbb{R}^{d_x}$ and $\mathbf{y} \in \mathbb{R}^{d_y}$ an input vector and a target, respectively. Let $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ be a training set consisting of n samples, with $\{\mathbf{x}_i\}_{i=1}^n$ being *distinct* points in $\mathbb{R}^{d_x}$. Denote by $\mathbf{x}_{k,i}$ the output of the k-th layer on the i-th training sample $\mathbf{x}_i$. Notice that MCN is primarily designed to learn some extrinsic structures from the data $\mathbf{x}$, and its outputs may be inconsistent with the target $\mathbf{y}$, e.g., they might have different dimensions. Hence, an additional mapping $\Psi : \mathbb{R}^{d_l} \to \mathbb{R}^{d_y}$ is used to further transform the network outputs, resulting in the following objective function for training an l-layer MCN:

$$L(\boldsymbol{\theta}_l) := \frac{1}{n}\sum_{i=1}^{n} \ell(\Psi(\mathbf{x}_{l,i}), \mathbf{y}_i), \tag{3}$$

where $\ell : \mathbb{R}^{d_y} \times \mathbb{R}^{d_y} \to \mathbb{R}$ is an arbitrary lower-bounded loss function (without losing generality, we assume the lower bound is 0), and $\boldsymbol{\theta}_l = \{\boldsymbol{\theta}(\mathcal{L}_k, \mathcal{W}_k, \tilde{\mathcal{A}}_k, \mathcal{A}_k)\}_{k=1}^l$ is a collection of all learnable parameters with $\boldsymbol{\theta}(\mathcal{L}_k, \mathcal{W}_k, \tilde{\mathcal{A}}_k, \mathcal{A}_k)$ being the parameters of the operators $\mathcal{L}_k, \mathcal{W}_k, \tilde{\mathcal{A}}_k$ and $\mathcal{A}_k$ defined in (2). In our setup, the extra mapping $\Psi(\cdot)$ could be either learnt or fixed [4], while the activation functions $\sigma_k(\cdot)$ and $\gamma(\cdot)$ are always fixed.

To obtain rigorous conclusions, some technical conditions are required. But for the ease of presentation, we would like to present them along with the established theorems.

fixed dimension $d_{\mathcal{L}}$, $\forall k = 0, \cdots, l-1$. Actually, our methods and theories do not need this assumption.

[4] There is no much difference between these two variants, as fixing the last layer of a DNN may cause very little influence (Hoffer et al., 2018).

3. Main Results

This section presents the main results of this paper, including a couple of theories regarding the optimality, fitting ability and generalization performance. All the detailed proofs of these theorems are provided in the supplementary material.

3.1. Effects of Depth

First note that an $(l+1)$-layer MCN is obtained by adding one layer into the network consisting of its first l layers, i.e., $\boldsymbol{\theta}_{l+1} = \{\boldsymbol{\theta}_l, \boldsymbol{\theta}(\mathcal{L}_{l+1}, \mathcal{W}_{l+1}, \tilde{\mathcal{A}}_{l+1}, \mathcal{A}_{l+1})\}$. Under some mild technical conditions, we prove that the training objective (3) is non-increasing, or even monotonically decreasing, as the network goes deeper[5].

Theorem 1 (Effects of Depth)**.** *Let the activation function $\gamma(\cdot)$ be the element-wise $\exp(\cdot)$. Suppose that the loss function $\ell(\cdot)$ in (3) is differentiable and convex. Denote by $\boldsymbol{\theta}_{l+1}$ any local minimum of an $(l+1)$-layer MCN. If $d_{l+1} = d_l$, then the following holds for any fixed injection $\Psi(\cdot)$:*

$$L(\boldsymbol{\theta}_{l+1}) \leq \min_{\boldsymbol{\theta}_l'} L(\boldsymbol{\theta}_l'),$$

where $\boldsymbol{\theta}_l'$ is a global minimum of the l-layer MCN. Moreover, if $\ell(\cdot)$ is strongly convex and there exists i such that $\mathbf{x}_{l+1,i} \neq \mathbf{x}_{l,i}'$, then the inequality is strict, namely $L(\boldsymbol{\theta}_{l+1}) < \min_{\boldsymbol{\theta}_l'} L(\boldsymbol{\theta}_l')$.

The setting of fixing $\Psi(\cdot)$ is to ensure that an $(l+1)$-layer MCN and its l-layer part are comparable. According to the above theorem, the global minima of an l-layer MCN can be attained, or even outperformed, by simply increasing the network depth by one. So, given the context of MCN, increasing network depth can not only "eliminate" local minima, but also help seek good solutions that possess smaller training error, providing a theoretically interpretation for a well-known empirical observation—deeper networks usually lead to better training results.

Among the other things, provided that the loss function is differentiable and strongly convex, we can further prove that the training error is able to go to zero. But the proof needs a key theorem established in the next subsection.

Remark 1: One may worry that there exist decreasing paths to infinity, and the weight may need to diverge to improve the performance of local minima (Sohl-Dickstein & Kawaguchi, 2019). The previous work (Kawaguchi, 2016; Liang et al., 2018a;b) may suffer from this problem, mainly due to their explicit regularization, whose coefficient should decay to zero to ensure the consistency of optimization.

[5] This is not in conflict with the learning-based optimization theories (Xie et al., 2019; Liu et al., 2019), which show that their networks can converge fast and need only a smaller number of layers to solve optimization problems. In fact, empirically, MCN will converge when the network is deep enough.

Hence, it leads to the divergence of some parameters to ensure the scale of output. However, our results hold without requiring any parameter to approach zero or infinity. Furthermore, for the classification problem, this divergence problem can be solved by proper parameter regularization (Liang et al., 2019). But, for the general regression problem, regularization may not work. Fortunately, under the over-parameterized setting, algorithmic analysis (Allen-Zhu et al., 2019; Du et al., 2019a) can entirely avoid the divergence risk. We leave the algorithmic analysis of MCN as our future work.

3.2. Approximation Ability

In general, it is unlikely that all mathematical functions can be approximated by DNNs. The following defines a class of functions which can be well approximately by MCN.

Condition 1. *For $\beta \in \mathbb{N}$, we define a modified β-th Sobolev space on the hypercube $[-1,1]^{d_x}$*

$$\mathcal{H}^{\beta} \coloneqq \left\{ \mathbf{f} : \mathrm{D}^{\boldsymbol{\alpha}} \mathbf{f} \in \mathrm{L}^2\left([-1,1]^{d_x}\right), \forall \boldsymbol{\alpha} : |\boldsymbol{\alpha}|_{\infty} \leq \beta \right\},$$

*where $\boldsymbol{\alpha} = (\alpha_1, \cdots, \alpha_{d_x}) \in \mathbb{N}^{d_x}$ is a multi-index, $\mathrm{D}^{\boldsymbol{\alpha}}$ corresponds to the **weak** derivatives operator $\partial_{x_1}^{\alpha_1} \ldots \partial_{x_{d_x}}^{\alpha_d}$ of order $|\boldsymbol{\alpha}| = \alpha_1 + \cdots + \alpha_{d_x}$ and $|\boldsymbol{\alpha}|_{\infty} = \max\{\alpha_i\}$. It is assumed that the function $\mathbf{f} \in \mathcal{H}^{2\beta+2}$ obeys the homogeneous Neumann boundary conditions up to order β:*

$$\partial_{x_j}^{2r+1} \mathbf{f}\Big|_{\partial\Omega_j} = 0, \quad j = 1, \ldots, d_x, \quad r = 0, \ldots, \beta - 1,$$

where $\partial\Omega_j = \left\{ \mathbf{x} \in [-1,1]^d : x_j = \pm 1 \right\}$ is the boundary.

The above condition depicts a class of continuous functions $\mathbf{f} \in \mathrm{L}^2([-1,1]^{d_x})$ such that $\mathbf{f}$ and its weak derivatives up to a certain order have finite L_2 norm. Note that the Neumann boundary condition of $[-1,1]^{d_x}$ is not harsh, and we can always extend the target function by firstly using the Sine or Cosine functions to introduce the homogeneous Neumann property and then scaling it to the interval $[-1,1]^{d_x}$.

As pointed out by (Barron, 1993), a standard fully connected neural network with enough, possibly infinite, hidden units can approximate any continuous function in compact domain. For MCN, we have an explicit approximation bound to connect the width and depth in a finite fashion.

Theorem 2 (Approximation Ability)**.** *Let $\mathbf{f}$ be a vector-valued function that obeys Condition 1, and let $w \geq 0, p \geq 0, N \gg 1$ be given numbers. Define $N_d = N\left(\ln N\right)^{d_x - 2}$, and denote by $\mathbf{f}_{\boldsymbol{\theta}}$ the output of an MCN. Suppose either $\mathbf{f}_{\boldsymbol{\theta}}$ is of width $\mathcal{O}(N_d d_x w p \ln p)$ and depth $\mathcal{O}(l \ln p + N^2)$, or $\mathbf{f}_{\boldsymbol{\theta}}$ has $\mathcal{O}(d_x w p \ln p)$ width and $\mathcal{O}(N_d l \ln p + N^2 N_d)$ depth. Then $\mathbf{f}$ can be approximated by MCN with proper parameters, in a sense that:*

$$\|\mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x}) - \mathbf{f}(\mathbf{x})\|_{\infty} \leq \epsilon, \quad \forall \mathbf{x} \in [-1,1]^{d_x},$$

where

$$\epsilon = \mathcal{O}\left(d_x 2^{d_x} p^2 2^{-wl} + N^{-2\beta-2} \left(\ln N\right)^{d_x - 1} \right).$$

The number of non-zero parameters in $\boldsymbol{\theta}$ is in the order of $\mathcal{O}\left(N_d \left(d_x w^2 p \ln p + N^2\right)\right)$.

Proof Sketch. We first construct the shallow MCNs that approximate $\sin(n\pi x)$ and $\cos(n\pi x)$ for different $n \in \mathbb{N}$ exponentially fast. Then we can obtain a multivariate function $\phi_{\mathbf{n}} \coloneqq \prod_{i=1}^{d_1} \sin(n_i \pi x) \prod_{k=d_1+1}^{d_x} \cos(n_k \pi x)$ by an MCN of $\mathcal{O}(\ln d_x)$ depth, where $d_1 \leq d_x$ and $\mathbf{n} \in \mathbb{N}^{d_x}$. Since the set $\{\phi_{\mathbf{n}}, \mathbf{n} \in \mathbb{N}^{d_x}\}$ is a Fourier orthogonal basis for $\mathrm{L}^2([-1,1]^{d_x})$, we can prove that $N(\ln N)^{d_x - 2}$ sub-MCNs suffice to approximate the target function, where $N = \prod_i n_i$. More detailed proofs can be founded in the supplementary material. □

Remarkably, Theorem 2 shows that MCN requires only a parameter complexity of $\mathcal{O}\left(d_x N (\ln N)^{d_x - 2}\right)$ to approximate the target function, which is dramatically lower than the $\mathcal{O}(d_x^{d_x} N^{d_x})$ required by deep ReLU (Yarotsky, 2017). This is mainly benefited from our analysis techniques. Unlike the analyses in (Lu et al., 2020; Yarotsky, 2018), which split the input space into small hyper-cubes and use a local network to approximate the Taylor expansion on those hyper-cubes, our analysis is built upon high-dimensional Fourier expansions and can therefore obtain higher decay rate for the approximation residual. Besides, the special network architecture of MCN is another cause of the advantage of lower complexity. Namely, the maximum operator makes the power of the decay term for approximating underlying polynomial be in the order of width×depth. By contrast, the decay power is just proportional to the depth in deep ReLU.

In summary, Theorem 2 illustrates that MCN with highly sparse connectivity between neurons can produce good approximation performance. This forms good basis for establishing tight generalization bound and eliminating bad local minima, as will be shown soon.

3.3. Generalization Bound

Theorem 2 ensures the existence of a good predictor when MCN goes deeper and wider. Now, one natural question is: does the generalization bound also shrink as the network becomes deeper? To analyze the generalization ability of DNNs or any other learning methods, it is indeed necessary to make some assumptions about the data. In this subsection, we set $d_y = 1$ and assume that $\mathbf{x}_i \in [0,1]^{d_x}$ for $i = 1, \cdots, n$. We consider the nonparametric regression task, i.e., there exists a target oracle function f_0 such that

$$y_i = f_0(\mathbf{x}_i) + \varepsilon_i, \quad i = 1, \cdots, n, \tag{4}$$

where the noise terms ε_i's are assumed to be i.i.d. Gassuian and independent of $\mathbf{x}_i$.

Denote the function class of our MCN as

$$\mathcal{F}(\boldsymbol{\theta}, s) \coloneqq \left\{ f_{\boldsymbol{\theta}} : \mathrm{Supp}\left(\boldsymbol{\theta}\right) < s, \|\boldsymbol{\theta}_k\|_F^2 < \infty, \forall k \leq l \right\},$$

where $\|\boldsymbol{\theta}_k\|_F$ is the Frobenius norm of all the parameters at the k-th layer, and the operator $\mathrm{Supp}(\cdot)$ denotes the support of a set, i.e., $\mathrm{Supp}\left(\boldsymbol{\theta}\right)$ is the number of nonzero parameters in MCN. The boundness assumption of $\mathrm{Supp}\left(\boldsymbol{\theta}\right) < s$ is made on the basis of Theorem 2, which shows that MCN with sparse connections can possess strong approximation ability. For convenience, we consider the case where the structure of $\mathcal{F}(\boldsymbol{\theta}, s)$ is deterministic, i.e., the input layer of $\mathcal{A}_k(\cdot)$ is the same for all MCNs in $\mathcal{F}(\boldsymbol{\theta}, s)$. Denote by $\mathcal{N}\left(\delta, \mathcal{F}(\boldsymbol{\theta}, s), \|\cdot\|_1\right)$ the minimal number of ℓ_1-balls with radius δ that covers $\mathcal{F}(\boldsymbol{\theta}, s)$. The logarithm of $\mathcal{N}\left(\delta, \mathcal{F}(\boldsymbol{\theta}, s), \|\cdot\|_1\right)$ is also called the *covering number* for convenience. For an operator $\mathcal{A}$, $\|\mathcal{A}\|_1$ denotes its ℓ_1 norm induced by the vector ℓ_1 norm, namely $\|\mathcal{A}\|_1 = \max_{\mathbf{x}\neq 0} \frac{\|\mathcal{A}(\mathbf{x})\|_1}{\|\mathbf{x}\|_1}$. Then we have the following theorem to bound the covering number (i.e., $\ln \mathcal{N}\left(\delta, \mathcal{F}(\boldsymbol{\theta}, s), \|\cdot\|_1\right)$).

Theorem 3 (Covering Number of MCN)**.** *Assume that the activation function $\sigma_k(\cdot)$ is ρ_k-Lipschitz and $\rho_k \leq \rho$ for $k = 1, \cdots, l$. Then one block of MCN is κ_k-Lipschitz continuous w.r.t. the input layers and*

$$\kappa_k = \left(1 + \max\{\rho_k, 2\}\|\boldsymbol{\theta}_k\|_1\right),$$

where

$$\|\boldsymbol{\theta}_k\|_1 \coloneqq \max\{\|\tilde{\mathcal{A}}_k\|_1, \|\mathcal{A}_k\|_1, \|\mathcal{W}_k + \mathcal{L}_k\|_1\}.$$

Moreover, we have

$$\ln \mathcal{N}\left(\mathcal{F}(\boldsymbol{\theta}, s), \delta, \|\cdot\|_1\right) \leq \mathcal{O}\left(sl \ln\left(\frac{\rho w \prod_{k=1}^{l} \kappa_k}{\delta}\right)\right),$$

where w and l are the width and depth of MCN, respectively.

The above theorem shows that the covering number of MCN is $\mathcal{O}\left(sl^2 \ln\left(w/\delta\right)\right)$, where $s = \Theta\left(d_x N (\ln N)^{d_x - 2}\right)$. By contrast, to achieve the same approximation accuracy, deep ReLU needs a covering number of $\mathcal{O}\left(s' l \ln\left(s' w^2 l/\delta\right)\right)$, with $s' = \Theta(d_x^{d_x} N^{d_x})$. In the situation of high-dimensional data, i.e., d_x is large, it is clear that MCN has much smaller covering number than deep ReLU, which means that the model complexity of MCN is much lower. Due to this, MCN provably owns good generalization performance, as shown in the following.

Corollary 3.1 (Generalization Bound)**.** *Consider the regression problem in (4) and assume $\max_{\mathbf{x}\in[0,1]^{d_x}} f_0(\mathbf{x}) < \infty$. Let f_{M} be any MCN from $\mathcal{F}(\boldsymbol{\theta}, s)$, and define*

$$\ell_n(f) \colon = \frac{1}{n} \sum_{i=1}^{n} \left(y_i - f(\mathbf{x}_i)\right)^2,$$

$$\Delta_n \coloneqq \mathbb{E}_{f_0}\left[\ell_n(f_{\mathrm{M}}) - \inf_{f\in\mathcal{F}} \ell_n(f)\right],$$

where $\mathbb{E}_{f_0}$ is the expectation taken with respect to the samples generated from the regression model (4). Define

$$\mathrm{dis}(f_{\mathrm{M}}, f_0) \coloneqq \mathbb{E}_{f_0}\left[\left(f_{\mathrm{M}}(\mathbf{x}) - f_0(\mathbf{x})\right)^2\right].$$

Then, we have

$$\mathrm{dis}(f_{\mathrm{M}}, f_0) \leq \mathcal{O}\left(\Delta_n + \inf_{f\in\mathcal{F}} \mathrm{dis}(f, f_0) + \frac{sl^2 \ln\left(wn\right)}{n}\right).$$

This corollary is indeed a direct application of the general statics generalization inequality in (Lu et al., 2020; Yarotsky, 2017). As we can see, the generalization bound depends on three parts, intuitively described as $\varepsilon_1 + \varepsilon_2 + \varepsilon_3$, where ε_1 is the gap from the obtained training loss to the global minimal one, ε_2 is the approximation error, and ε_3 is the covering number. Notably, Theorem 1 provides a way to reduce ε_1, and Theorem 3 ensures that small ε_2 unnecessarily results in large ε_3.

For nonparametric regression with square loss, when the target function f_0 is β-smooth, it is well-known that the statistically optimal estimation rate in terms of data size is $n^{-\frac{2\beta}{2\beta+d_x}}$ (Giné & Nickl, 2016), also called as *minimax estimation rate*. Owning the minimax estimation rate means that the estimator performs the best in the worst case. Interestingly, when the training data is fitted exactly, MCN also owns this property.

Theorem 4 (Minimax Estimation Rate)**.** *Suppose that the density $p(\cdot)$ over some compact set $\mathcal{C}$ satisfies*

$$0 < p_{\min} \leq p(\mathbf{x}) \leq p_{\max}, \quad \forall \mathbf{x} \in \mathcal{C}.$$

Assume that the target function f_0 is β-smooth and let $\ell(\cdot)$ in (3) be the square loss. Denote the final output of our model as $f_{\boldsymbol{\theta}}(\mathbf{x}_i)$, where $\boldsymbol{\theta}$ is the learnable parameters of MCN. If $f_{\boldsymbol{\theta}}(\mathbf{x}_i) = y_i$ for $i = 1, \cdots, n$, then for any data sample $\mathbf{x} \in \mathbb{R}^{d_x}$ located in the support set of p, the output of MCN satisfies the following with high probability:

$$\mathbb{E}_{\mathcal{S}^n}\left[\mathbb{E}_{\varepsilon}\left[\|f_{\boldsymbol{\theta}}\left(\mathbf{x}\right) - f_0\left(\mathbf{x}\right)\|^2 \mid \mathcal{S}^n\right]\right] \leq C n^{-\frac{2\beta}{2\beta+d_x}},$$

where $\mathcal{S}^n = \{(\mathbf{x}_i, y_i))\}_{i=1}^n$ and $C > 0$ is a number depends only on the numerical range of the outputs of MCN.

In general, the above theorem confirms the phenomenon that over-parameterized DNNs may not necessarily cause over-fitting (Belkin et al., 2019; 2018b). For Theorem 4 to hold, the training error has to be reduced to zero. This can actually be accomplished by using the techniques in (Gasca & Sauer, 2000) to link together Theorem 1 and Theorem 2, as will be shown in next subsection.

Remark 2: One may worry about the "exact fitting" assumption may not be satisfied since the noise belongs to an unbounded distribution. The derivatives or weights of DNN may diverge to infinity as $n \to \infty$. However, this may not be a problem and exact fitting can easily happen under mild condition. On the one hand, Gaussian distribution has an exponential decay tail. Thus, we can approximately treat it as bounded. On the other hand, some recent results (Arora et al., 2019c; Du et al., 2019a; Allen-Zhu et al., 2018; Liang et al., 2020; E et al., 2019) show that the DNNs, having universal approximation ability, can easily fit the Gaussian noise without any weight diverging. Even more, exact fitting can happen near the initial state of DNN as long as it is wide or deep enough; the depth or width is in the polynomial order of n. For MCN, we already prove its approximation ability in Theorem 2. Following the same road-map, we can conclude that its parameters do not diverge in the exact fitting case.

Remark 3: We remark that the estimator $f_{\boldsymbol{\theta}}$ does not belong to the β-smooth function class (its smoothness depends on the architecture and activation function). In conclusion, even though $f_{\boldsymbol{\theta}}$ is not β-smooth and fits the data exactly, it attains optimal excess loss rates. We refer the readers to (Rakhlin et al., 2017) for further discussion of optimal rates in non-parametric estimation and statistical learning.

Remark 4: For any $\mathbf{x}_i \in \mathcal{S}^n$,

$$\mathbb{E}_{\epsilon_i}\left[\|f_{\boldsymbol{\theta}}(\mathbf{x}_i) - f_0(\mathbf{x}_i)\|^2 \mid \mathcal{S}^n\right] = 1.$$

However,

$$\mathbb{E}_{\mathcal{S}^n}\left[\mathbb{E}_{\varepsilon}\left[\|f_{\boldsymbol{\theta}}(\mathbf{x}_i) - f_0(\mathbf{x}_i)\|^2 \mid \mathcal{S}^n\right]\right] \to 0, \text{ as } n \to \infty,$$

due to the measure of a specific point is 0.

3.4. No Bad Local Minima

As aforementioned, under mild technical conditions, the training error produced by MCN can be arbitrarily small when the network is deep enough.

Corollary 4.1 (Optimal Training Error)**.** *Suppose that the loss function $\ell(\cdot)$ is differentiable and strongly convex. Denote by $\boldsymbol{\theta}_l$ any local minimum of an l-layer MCN. For any $\epsilon > 0$, there exists a $D \in \mathbb{N}$ such that $L(\boldsymbol{\theta}_l) \leq \epsilon$ holds for any $l > D$.*

The "no bad local minima" property of MCN replies on its special network design, and is unnecessarily true for the other DNNs. In the following, we shall introduce two ways to refine an existing DNN that is possibly poorly designed. The first one is straightforward and simply to treat the output of an existing DNN as the input $\mathbf{x}_0$ to MCN, and the parameters of the existing network are not involved in re-training. In this case, it is easy to obtain the following result:

Corollary 4.2 (Partial Training)**.** *For fixed injection $\Psi(\cdot)$ and an existing l_0-layer DNN with output $\mathbf{h}_0$, construct an l-layer MCN with $\gamma(\cdot)$ being element-wisely exponential and input $\mathbf{x}_0 = \mathbf{h}_0$. If $\mathbf{h}_0$ is an injective function w.r.t. the input $\mathbf{x}$ and the loss $\ell(\cdot)$ is differentiable and strongly convex, then for any $\epsilon > 0$, there exists a large enough $D \in \mathbb{N}$, such that $L(\boldsymbol{\theta}_l) \leq \epsilon$ holds for any local minimum $\boldsymbol{\theta}_l$ with $l \geq D$.*

In above corollary, the existing DNN is assumed to be fixed and MCN is simply applied to its output. Actually, it is also feasible to re-train all the parameters, including the parameters of both the existing network and the appended MCN blocks.

Theorem 5 (Full Training)**.** *For fixed injection $\Psi(\cdot)$ and an existing l_0-layer DNN with output $\mathbf{h}_0$, append an l-layer MCN at its end with $\gamma(\cdot)$ being element-wisely exponential, resulting in a new model $\mathbf{h}_l$. Suppose that the loss $\ell(\cdot)$ is differentiable and strongly convex, and there exist parameters that make $\mathbf{h}_0$ be injective. Then, for any $\epsilon > 0$, there exists a large enough $D \in \mathbb{N}$ such that*

$$\frac{1}{n}\sum_{i=1}^{n} \ell(\Phi(\mathbf{h}_l(\mathbf{x}_i)), \mathbf{y}_i) \leq \epsilon$$

holds at any local minimum $\mathbf{h}_l$ with $l \geq D$.

One may have noticed that monotonic decreasing property in Theorem 1 is not enough to guarantee global minimal training loss. In fact, as aforementioned, Theorem 2 also plays an important role in gaining the above results, and we need use the techniques in (Gasca & Sauer, 2000) to link Theorem 1 and Theorem 2 together.

Remarkably, the above results illustrate that MCN is not just an approach for seeking the global optimal solution to certain optimization problems, but instead a powerful tool for helping seek better solutions to the primary task behind the optimization problems.

3.5. Discussions

There is another interpretation for why MCN can eliminate bad local minimum. When adopting the square loss, we find that the loss in (3) at the local minimum equals to a projection residual obtained by projecting the training data onto a subspace. The subspace is expanded by parameters in the concatenation linear part $\mathcal{L}_k(\cdot)$ for $k = 1, \cdots, l$, which means that the subspace is larger when more independent parameters are contained in the linear branch $\mathcal{L}_k(\cdot)$. On the other hand, large space often brings small projection residual. Please see Section B in the supplementary material for more details.

To summarize, this section establishes a collection of theorems to cope with the problems of bad local minima and generalization issue. More precisely, first, Theorem 1 and Corollary 4.1 reveal the "no bad local minima" property of MCN, and Corollary 4.2 and Theorem 5 extend this property to the other DNNs. Second, Theorem 2 shows the

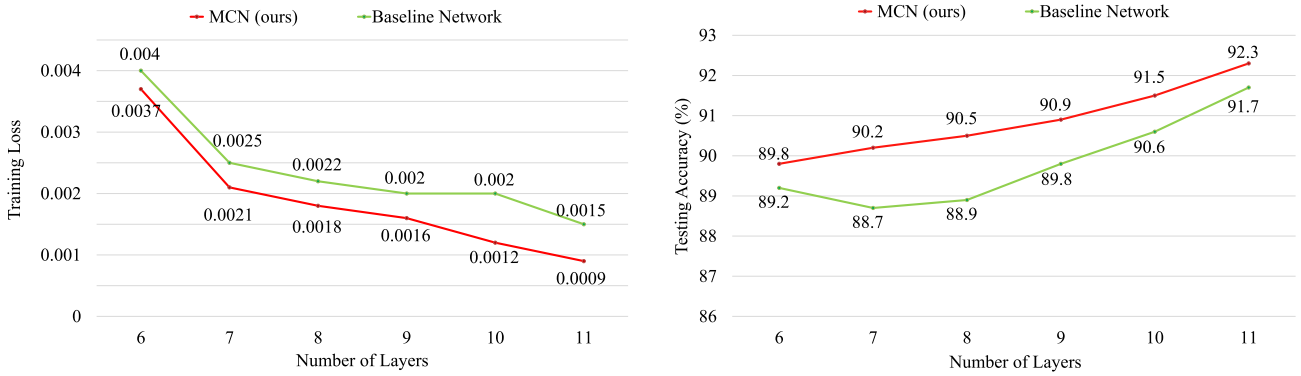

Figure 2. Left: Training loss of our MCN (red) and baseline network (green) with various number of layers. Right: Testing accuracy.

approximation ability of MCN, illustrating that MCN can obtain the same approximation error by using parameters much less than deep ReLU. The number of required parameters is far smaller than the network size, which implies that MCN allows to use some prevalent sparse patterns such as CNN structure and pruning tricks. The sparsity of network connections further leads to a small covering number for MCN in Theorem 3. Based on this, finally, we provide the generalization bound for MCN in Corollary 3.1.

4. Experiments

4.1. Theorems Verification

We conduct experiments on the commonly used CIFAR-10 dataset, with the purpose of validating our theorems as well as the effectiveness of MCN. We first construct a baseline network with 6 weighted layers, including five convolutional layers and one fully-connected layer. Then we add convolutional layers to make the network deeper. It contains five max pooling in total. For our MCN, we replace the convolutional layers after the third max pooling layer with our MCN block. To make a fair comparison, both networks have the same number of layers and parameters, and so for the random seed and learning rate. Also, batch normalization and ReLU are adopted by both networks. For detailed experimental settings and model configurations, please refer to the supplementary material.

Figure 2 shows the training loss and testing accuracy with different number of layers. According to the red line in the left part of Figure 2, the training loss of our MCN monotonically decreases with the increase of depth. This is consistent with our Theorems 1 and 2. From the red line in the right part of Figure 2, we can see that deeper MCN can achieve better testing accuracy, which demonstrates the generalization performance of MCN and confirms our Corollary 3.1 and Theorem 4. In addition, according to the green line in the right part of Figure 2, the testing accuracy of the baseline network does not monotonically increase as the network goes deeper. Therefore, the "no bad local minima" property should be a primary cause of the nice performance of MCN. In summary, compared with the baseline network, our MCN has much lower training loss as well as higher testing accuracy, revealing the superiority of MCN.

4.2. Appending MCN

To validate the merits of Corollary 4.2 and Theorem 5, we add two MCN blocks to VGG19 (Simonyan & Zisserman, 2014) and ResNet18 (He et al., 2016) as the treatment group. The original two architectures, VGG19 and ResNet18, are regarded as the first control group. To make a comparison, we also add two traditional convolutional layers to VGG19 and ResNet18, considered as the second control group. For the treatment group and the second control group, we consider two ways to train the appended VGG19 and ResNet18 (short as Res18). The first one is partial training which treats VGG19 and Res18 as the feature extractors whose parameters are not involved during training. The second one is full training which considers the appended networks as new models and train them from scratch.

Table 1 shows the comparison results among all the three groups, in terms of both training loss and testing accuracy. As we can see, the plugging of traditional convolution layers can decrease the training loss, however, the appending of MCN has more amount of improvement, which, again, show the benefits of the "no bad local minima" property. Interestingly, full training and partial training share comparable performance when appending MCN but not for convolution layers. Hence, both Corollary 4.2 and Theorem 5 are practical theories. Moreover, our MCN outperforms distinctly all the competing methods; this, again, confirms the superiority of our MCN architecture.

Table 1. The training error (Err.) and testing accuracy (Acc.) of different models on the CIFAR-10 dataset. We denote by C the added two convolutional layers and M the appended MCN blocks.

Models	VGG19	VGG19+ C(full)	VGG19+ C(part)	VGG19+ M(full)	VGG19+ M(part)	Res18	Res18+ C(full)	Res18+ C(part)	Res18+ M(full)	Res18+ M(part)
Err.	0.0016	0.0013	0.0015	**0.0010**	0.0011	0.0013	0.0011	0.0012	**0.0009**	**0.0009**
Acc.	92.0%	92.4%	92.1%	**92.8%**	92.6%	92.7%	93.5%	93.1%	93.7%	**93.8%**

Table 2. The training error (Err.) and testing accuracy (Acc.) of different models on the CIFAR-100 dataset. We denote by C the added two convolutional layers and M the appended MCN blocks.

Models	Res18	Res18+ C(full)	Res18+ C(part)	Res18+ M(full)	Res18+ M(part)	ResNeXt29	ResNeXt29+ C(full)	ResNeXt29+ C(part)	ResNeXt29+ M(full)	ResNeXt29+ M(part)
Err.	0.0020	0.0014	0.0015	**0.0009**	**0.0009**	0.0056	0.0051	0.0054	**0.0008**	0.0011
Acc.	76.15%	76.58%	76.48%	**76.95%**	76.87%	80.71%	80.78%	80.69%	**82.31%**	81.41%

4.3. Additional Experiments

To better demonstrate the representation ability of our MCN block, we further conduct some additional experiments on the more complex dataset CIFAR-100, and make comparisons with the SOTA of ResNeXt (Xie et al., 2017) (a more powerful network architecture).

Similar to the previous part, the original two architectures, Res18 and ResNeXt29, are regarded as the first control group. As for the second control group, we still add two traditional convolutional layers to Res18 and ResNeXt29. Besides, we append two MCN blocks to the end of both Res18 and ResNeXt29 as the treatment group.

For the treatment group and the second control group, the two ways to train the appended Res18 and ResNeXt29 remain the same as previous experiment. One is partial training which treats Res18 and ResNeXt29 as the feature extractors, while the other is full training which considers the appended DNNs as new models and train them from scratch.

We present the results of the partial training (i.e., fixing Res18 and ResNeXt29 when appending MCN blocks) in Table 2. It can be seen that, even in the case of handling complex data, our MCN achieves superior results. The treatment groups under two different training methods both outperform the control groups, which is consistent with Table 1. Moreover, by comparing Table 1 with Table 2, our MCN blocks have greatly improved the performance when handling more complex data. Please note that ordinarily appending CNNs cannot ensure the monotonicity of Err. and Acc. This phenomenon not only verifies Corollary 4.2 and Theorem 5 again, but also shows that our MCN has a stronger representation ability than general linear structure, which corresponds to Theorem 2.

5. Conclusion

In this paper, we propose a novel multi-layer DNN structure termed MCN, which can approximate some class of continuous functions arbitrarily well even with highly sparse connection. We prove that the global minima of an l-layer MCN may be outperformed, at least can be attained, by simply increasing the network depth. More importantly, MCN could be easily appended to any of the many existing DNN and the augmented DNN will share the same property of MCN. Finally, we analyze the generalization ability of MCN and reveal that depth is more important than width for generalization; this supports the mechanism of deep learning. In summary, this study does take a step towards the ultimate goal of deep learning theory—to understand why DNNs can work well in a wide variety of applications.

Acknowledgments

This work is supported in part by New Generation AI Major Project of Ministry of Science and Technology of China (grant no 2018AAA0102501), in part by NSF China (grant no.s 61625301 and 61731018), in part by Major Scientific Research Project of Zhejiang Lab (grant no.s 2019KB0AC01 and 2019KB0AB02), in part by Fundamental Research Funds of Shandong University, in part by Future Talents Research Funds of Shandong University, in part by Beijing Academy of Artificial Intelligence, in part by Qualcomm, and in part by SenseTime Research Fund.

References

Adcock, B. Multivariate modified fourier series and application to boundary value problems. *Numerische Mathematik*, 115(4):511–552, 2010.

Allen-Zhu, Z., Li, Y., and Liang, Y. Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918*, 2018.

Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252, 2019.

Amos, B., Xu, L., and Kolter, J. Z. Input convex neural networks. In *International Conference on Machine Learning*, pp. 146–155, 2017.

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

Arora, S., Cohen, N., Hu, W., and Luo, Y. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, pp. 7411–7422, 2019a.

Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R., and Wang, R. On exact computation with an infinitely wide neural net. *arXiv preprint arXiv:1904.11955*, 2019b.

Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019c.

Barron, A. R. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.

Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine learning and the bias-variance trade-off. *arXiv preprint arXiv:1812.11118*, 2018a.

Belkin, M., Hsu, D. J., and Mitra, P. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in Neural Information Processing Systems*, pp. 2300–2311, 2018b.

Belkin, M., Rakhlin, A., and Tsybakov, A. B. Does data interpolation contradict statistical optimality? In *International Conference on Artificial Intelligence and Statistics*, pp. 1611–1619, 2019.

Cao, Y. and Gu, Q. A generalization theory of gradient descent for learning over-parameterized deep relu networks. *arXiv preprint arXiv:1902.01384*, 2019.

Dou, X. and Liang, T. Training neural networks as learning data-adaptive kernels: Provable representation and approximation benefits. *arXiv preprint arXiv:1901.07114*, 2019.

Du, S. S., Wang, Y., Zhai, X., Balakrishnan, S., Salakhutdinov, R. R., and Singh, A. How many samples are needed to estimate a convolutional neural network? In *Advances in Neural Information Processing Systems*, pp. 373–383, 2018.

Du, S. S., Lee, J. D., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, 2019a.

Du, S. S., Zhai, X., Poczos, B., and Singh, A. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019b.

E, W., Ma, C., Wu, L., et al. On the generalization properties of minimum-norm solutions for over-parameterized neural network models. *arXiv preprint arXiv:1912.06987*, 2019.

Gasca, M. and Sauer, T. Polynomial interpolation in several variables. *Advances in Computational Mathematics*, 12 (4):377, 2000.

Giné, E. and Nickl, R. *Mathematical foundations of infinite-dimensional statistical models*, volume 40. Cambridge University Press, 2016.

Goodfellow, I., Warde-Farley, D., Mirza, M., Courville, A., and Bengio, Y. Maxout networks. In *International Conference on Machine Learning*, pp. 1319–1327, 2013.

Györfi, L., Kohler, M., Krzyzak, A., and Walk, H. *A distribution-free theory of nonparametric regression*. Springer Science & Business Media, 2006.

Hardt, M. and Ma, T. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Hoffer, E., Hubara, I., and Soudry, D. Fix your classifier: the marginal value of training the last weight layer. *arXiv preprint arXiv:1801.04540*, 2018.

Horn, R. A. and Johnson, C. R. Topics in matrix analysis, 1991. *Cambridge University Presss*, 37:39, 1991.

Huybrechs, D., Iserles, A., et al. From high oscillation to rapid approximation iv: Accelerating convergence. *IMA Journal of Numerical Analysis*, 31(2):442–468, 2011.

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

Kawaguchi, K. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, pp. 586–594, 2016.

Kawaguchi, K. and Kaelbling, L. P. Elimination of all bad local minima in deep learning. *arXiv preprint arXiv:1901.00279*, 2019.

Kawaguchi, K., Huang, J., and Kaelbling, L. P. Effect of depth and width on local minima in deep learning. *Neural Computation*, 2019.

Li, Y. and Liang, Y. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pp. 8157–8166, 2018.

Liang, S., Sun, R., Lee, J. D., and Srikant, R. Adding one neuron can eliminate all bad local minima. In *Advances in Neural Information Processing Systems*, pp. 4350–4360, 2018a.

Liang, S., Sun, R., Li, Y., and Srikant, R. Understanding the loss surface of neural networks for binary classification. *arXiv preprint arXiv:1803.00909*, 2018b.

Liang, S., Sun, R., and Srikant, R. Revisiting landscape analysis in deep neural networks: Eliminating decreasing paths to infinity. *arXiv preprint arXiv:1912.13472*, 2019.

Liang, T., Rakhlin, A., and Zhai, X. On the multiple descent of minimum-norm interpolants and restricted lower isometry of kernels. *arXiv preprint arXiv:1908.10292 [cs, math, stat]*, 2020.

Liu, J., Chen, X., Wang, Z., and Yin, W. ALISTA: Analytic weights are as good as learned weights in LISTA. In *International Conference on Learning Representations*, 2019.

Lu, J., Shen, Z., Yang, H., and Zhang, S. Deep network approximation for smooth functions. *arXiv preprint arXiv:2001.03040*, 2020.

Luxburg, U. v. and Bousquet, O. Distance-based classification with lipschitz functions. *Journal of Machine Learning Research*, 5(Jun):669–695, 2004.

Ma, C., Wu, L., et al. A priori estimates of the generalization error for two-layer neural networks. *arXiv preprint arXiv:1810.06397*, 2018.

Maillard, O. and Munos, R. Compressed least-squares regression. In *Advances in Neural Information Processing Systems*, pp. 1213–1221, 2009.

Olver, S. On the convergence rate of a modified fourier series. *Mathematics of Computation*, 78(267):1629–1645, 2009.

Rakhlin, A., Sridharan, K., Tsybakov, A. B., et al. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23 (2):789–824, 2017.

Rumerlhar, D. Learning representation by back-propagating errors. *Nature*, 323:533–536, 1986.

Schmidt-Hieber, J. Nonparametric regression using deep neural networks with relu activation function. *Annals of Statistics*, 2019.

Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: From theory to algorithms*. 2014.

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

Sohl-Dickstein, J. and Kawaguchi, K. Eliminating all bad local minima from loss landscapes without even adding an extra unit. *arXiv preprint arXiv:1901.03909*, 2019.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

Telgarsky, M. Benefits of depth in neural networks. In *Conference on Learning Theory*, pp. 1517–1539, 2016.

Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

Wei, C. and Ma, T. Data-dependent sample complexity of deep neural networks via lipschitz augmentation. *arXiv preprint arXiv:1905.03684*, 2019.

Wu, Y. and He, K. Group normalization. In *European Conference on Computer Vision*, pp. 3–19, 2018.

Xie, B., Liang, Y., and Song, L. Diverse neural network learns true target functions. In *International Conference on Artificial Intelligence and Statistics*, pp. 1216–1224, 2017.

Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5987–5995, 2017.

Xie, X., Wu, J., Zhong, Z., Liu, G., and Lin, Z. Differentiable linearized ADMM. In *International Conference on Machine Learning*, 2019.

Yarotsky, D. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.

Yarotsky, D. Optimal approximation of continuous functions by very deep relu networks. In *Conference on Learning Theory*, pp. 639–649, 2018.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Machine Learning*, 2017.

Zhang, X., Ling, C., and Qi, L. The best rank-1 approximation of a symmetric tensor and related spherical optimization problems. *SIAM Journal on Matrix Analysis and Applications*, 33(3):806–821, 2012.