

# An analysis of training and generalization errors in shallow and deep networks 

H. N. Mhaskar and T. Poggio


#### Abstract

This paper is motivated by an open problem around deep networks, namely, the apparent absence of overfitting despite large over-parametrization which allows perfect fitting of the training data. In this paper, we analyze this phenomenon in the case of regression problems when each unit evaluates a periodic activation function. We argue that the minimal expected value of the square loss is inappropriate to measure the generalization error in approximation of compositional functions in order to take full advantage of the compositional structure. Instead, we measure the generalization error in the sense of maximum loss, and sometimes, as a pointwise error. We give estimates on exactly how many parameters ensure both zero training error as well as a good generalization error. We prove that a solution of a regularization problem is guaranteed to yield a good training error as well as a good generalization error and estimate how much error to expect at which test data.


# An analysis of training and generalization errors in shallow and deep networks 

H. N. Mhaskar* and T. Poggio ${ }^{\dagger}$


#### Abstract

This paper is motivated by an open problem around deep networks, namely, the apparent absence of overfitting despite large over-parametrization which allows perfect fitting of the training data. In this paper, we analyze this phenomenon in the case of regression problems when each unit evaluates a periodic activation function. We argue that the minimal expected value of the square loss is inappropriate to measure the generalization error in approximation of compositional functions in order to take full advantage of the compositional structure. Instead, we measure the generalization error in the sense of maximum loss, and sometimes, as a pointwise error. We give estimates on exactly how many parameters ensure both zero training error as well as a good generalization error. We prove that a solution of a regularization problem is guaranteed to yield a good training error as well as a good generalization error and estimate how much error to expect at which test data.


Keywords: Deep learning, generalization error, interpolatory approximation

## 1 Introduction

The main problem of machine learning is the following. Given data ( $\mathbf{x}, y$ ) sampled from an unknown probability distribution $\mu$, the goal is to find a function $P$ that minimizes the generalization error $\mathbb{E}_{\mu}\left((y-P(\mathbf{x}))^{2}\right)$ among all functions in some function class that is thought to represent the prior information about the distribution. Since we do not know $\mu$, classical machine learning paradigm expresses the generalization error as a sum of three components: the variance, the approximation error, and the sampling error. The variance is a lower bound on the generalization error, and the estimation typically focuses on the other two errors. The sum of these two is given by $\mathbb{E}\left((f-P(\mathbf{x}))^{2}\right.$ ), where the expectation is with respect to the marginal distribution of $\mathbf{x}$ and the target function $f$ is the conditional expectation of $y$ given $\mathbf{x}$. If the marginal distribution of $\mathbf{x}$ is known, then the split between approximation and sampling errors is no longer necessary, and one can obtain estimates as well as constructions directly from characteristics of the training data (e.g., [11, 14, 18]). In the classical paradigm where this distribution is not known, the approximation error decreases as the number of parameters in $P$ increases to $\infty$. However, this makes the empirical risk minimization problem harder to solve; making it essential to choose the number of parameters in $P$ to balance the two errors. In turn, this explains a commonly observed phenomenon that if one achieves a zero empirical risk on the training data by over-parametrized model $P$, the test error is generally not good.

There are several recent papers that demonstrate that this phenomenon is often not observed (e.g., 10, 22, 25, [29, 2, 23]). There are a few recent papers that address this issue in the case of classification problems. Belkin, Hsu, and Mitra 1 analyze the "excess error" in least square fits by piecewise linear interpolants over that obtained by the optimal Bayes' classifier. In [23, 24], the question is analyzed from the perspective of the geometry of the error surface with respect to different loss functions near the local extrema. In particular, it is shown in [23] that substituting the ReLU activation function by a polynomial approximation exhibits the same behavior as the original network.

In this paper, we focus on regression problems. We wish to consider from the point of view of approximation theory the overfitting puzzle for universal approximation in a more intrinsic and theoretical manner, independent

[^0]of the training mechanism used. This paper does not seek to "explain" the overfitting phenomenon as observed. Such an explanation needs to involve not just the analysis of the approximation problem itself, but also an analysis of the training algorithms used for this purpose. It is obvious that a network (or any other model) using a number of parameters that is less than the number of training data points cannot in general produce a zero training error in the absence of some strong prior knowledge about the target function that generated the training data, no matter what training algorithm is used. Our goal is to study the fundamental problem of function approximation to examine what characteristics of the data and how much overparametrization will give theoretical guarantees that the generalization error can be controlled while achieving a zero training error. We will do this without any prior assumption about the target function apart from continuity or at most differentiability.

We wish to address the following issues about the phenomenon of zero/small training error and small test error for universal approximation:

1. What characteristics of the data govern a zero training error and a good generalization error?
2. How much over-parametrization will give a mathematical guarantee of the model to exhibit this behavior?
3. Propose a regularization scheme whose solution will guarantee a good (but not necessarily zero) training error while maintaining a good generalization error at the same time. The emphasis here is on an estimation of how much over-parametrization is necessary to get theoretical guarantees.
4. What bounds on the generalization error can be guaranteed by the solution of the (global) regularization scheme at each point in the test data (rather than a global error estimate), compared to the nearest neighbor in the training data?

In recent years, convolutional neural networks (CNNs) have achieved a revolution in machine learning. A good survey can be found in [12. From a practical point of view, the central features of CNNs are locality and weight sharing. From a strictly mathematical point of view, convolution is a very special operation that requires a group structure on the data. According to the book [8], the CNNs "are a specialized kind of neural network for processing data that has a known, grid-like topology. Examples include time-series data, which can be thought of as a 1D grid taking samples at regular time intervals, and image data, which can be thought of as a 2D grid of pixels." For this kind of data, it is customary to treat it either as data on the whole real line/plane with zero-padding, or otherwise use a symmetrical extension to treat it as data on a circle/torus so that the standard group operations on these spaces can be used to define the operation of convolution.

In this paper, we will focus on function approximation on the torus. The most fundamental class for this purpose is the class of all trigonometric polynomials. Accordingly, we will study the problem of the lack of overfitting in the context of approximation by trigonometric polynomials. We will explain in Section 2.2 the theoretical relationship between trigonometric polynomials and neural/RBF networks in the periodic setting. In Section 3, we will show how the theorems about trigonometric approximation translate into theorems about shallow networks with arbitrary periodic activation functions.

In the study of approximation error in deep learning, it is observed in 20] that the compositional structure of the target function can be utilized effectively via a property called good propagation of error to obtain substantially better error bounds allowing us to overcome the curse of dimensionality observed in shallow networks. This allows us to "lift" many results on approximation by shallow networks to those on deep networks. However, there are a few barriers that prevent a straightforward extension.

One is that we do not know the functions evaluated at each node in the intervening layers; just the input/output relations between the data and output of the ultimate layers. There are some recent efforts in the direction of designing deep networks in some special applications using domain knowledge (e.g., 9,19 ). It is conceivable that this problem does not appear in these contexts.

The other problem is even more fundamental, requiring a change in the traditional notion of generalization error. For example, suppose we wish to approximate a function $f^{*}=f\left(f_{1}\left(\mathbf{x}_{1}\right), f_{2}\left(\mathbf{x}_{2}\right)\right)$ by a network of the form $P^{*}=P\left(P_{1}\left(\mathbf{x}_{1}\right), P_{2}\left(\mathbf{x}_{2}\right)\right)$, where $\mathbf{x}_{1}, \mathbf{x}_{2} \in \mathbb{R}^{d}$. Without the compositional structure, $f^{*}$ has to be treated as a function of $2 d$ variables. The compositional structure allows us to treat the approximation problem as a set of three approximation problems: approximating functions $f_{1}, f_{2}$ as functions of $d$ variables each, and a function $f$ of 2 variables.

How do we define generalization error? Defining it in terms of the original distribution of $\left(\left(\mathbf{x}_{1}, \mathbf{x}_{2}\right), y\right)$ is not sensitive to the compositional structure. On the other hand, the input $\left(f_{1}, f_{2}\right)$ to the function $f$ is not the same (even may not have the same distribution) as the input $\left(P_{1}, P_{2}\right)$ to the approximation $P$.

Therefore, we measure in this paper the errors in the uniform norm rather than searching for appropriate $L^{2}$ norms suitable for the compositionality structure of the target function. Thus, we define the generalization error as
the maximal error between the target function and the model at all possible test points. One consequence is that the decomposition of the generalization error into three parts breaks down. Therefore, new ideas are required to achieve the approximation in terms of the training data alone. The approximation errors themselves are studied in [20, 15], but the techniques suggested there require the data points $\mathbf{x}_{j}$ to be sufficiently dense in the domain (Euclidean space, sphere, cube, etc.). When the data is not dense, it is clearly not expected to get a good approximation on the whole domain. However, if the data does become denser and denser on some subset of the domain, one can expect a good approximation at points close by to the training data. Thus, we will establish pointwise bounds for the generalization error obtained by a solution of a regularization scheme suggested for this purpose.

In Section 2, we develop some notation and provide some background information that has motivated our current paper. In Section 3, we state our theorems for the case of shallow networks. In Section 4, we state the analogues of the results in Section 3 in the case of deep networks. The proofs of the results in Sections 3 and 4 are given in Section 5

## 2 Background

The purpose of this section is to explain the connection between trigonometric polynomials and neural networks (Section 2.2) as well as to explain a classical theorem which provides a motivation for our theorems in this paper (Section 2.3). In order to do so, we need first to develop some notation. This is done in Section 2.1.

### 2.1 Notation

Let $q \geq 1$ be an integer, $\mathbb{T}^{q}$ be the $q$ dimensional torus $\left(=\mathbb{R}^{q} /(2 \pi \mathbb{Z})^{q}\right)$. For $\mathbf{x}, \mathbf{y} \in \mathbb{T}^{q}$,

$$
|\mathbf{x}-\mathbf{y}|=\max _{1 \leq i \leq q}\left|\left(x_{i}-y_{i}\right)(\bmod 2 \pi)\right|
$$

For a multi-integer $\mathbf{k},|\mathbf{k}|_{p}$ is the $\ell^{p}$ norm of $\mathbf{k}$.
The space of all continuous functions $f: \mathbb{T}^{q} \rightarrow \mathbb{R}$, equipped with the supremum norm will be denoted by $C^{*}$ (or $C^{*}\left(\mathbb{T}^{q}\right)$ if we wish to emphasize the input dimension to the functions). The norm on $C^{*}\left(\mathbb{T}^{q}\right)$ will be denoted by $\|\cdot\|$ or $\|\cdot\|_{q}$ if it is important to identify the dimension. For $n>0$, the space $\mathbb{H}_{n}^{q}$ of trignometric polynomials in $q$ variables with (spherical) degree $<n$ is defined by

$$
\mathbb{H}_{n}^{q}=\operatorname{span}\left\{\exp (i \mathbf{k} \cdot \circ):|\mathbf{k}|_{2}<n\right\} .
$$

The dimension of $\mathbb{H}_{n}^{q}$ is $\sim n^{q}$. If $f \in C^{*}\left(\mathbb{T}^{q}\right)$, then its Fourier coefficients are defined by

$$
\begin{equation*}
\hat{f}(\mathbf{k})=\frac{1}{(2 \pi)^{q}} \int_{\mathbb{T}^{q}} f(\mathbf{x}) \exp (-i \mathbf{k} \cdot \mathbf{x}) d \mathbf{x}, \quad \mathbf{k} \in \mathbb{Z}^{q} \tag{2.1}
\end{equation*}
$$

and its degree of approximation from $\mathbb{H}_{n}^{q}$ is defined by

$$
E_{n}(f)=E_{n}(q ; f):=\inf _{T \in \mathbb{H}_{n}^{q}}\|f-T\|
$$

When our models are trigonometric polynomials in $\mathbb{H}_{n}^{q}$, the quantity $E_{n}(f)$ denotes the ideal generalization error for the target function $f$.

Let $h: \mathbb{R} \rightarrow[0,1]$ be an infinitely differentiable, even function such that $h(t)=1$ if $|t| \leq 1 / 2, h(t)=0$ if $|t| \geq 1$. We define

$$
\begin{equation*}
\Phi_{N}(\mathbf{x})=\sum_{\mathbf{k} \in \mathbb{Z}^{q}} h\left(\frac{|\mathbf{k}|_{2}}{N}\right) \exp (i \mathbf{k} \cdot \mathbf{x}), \quad \mathbf{x} \in \mathbb{T}^{q}, N>0 \tag{2.2}
\end{equation*}
$$

For $f \in C^{*}$, we define

$$
\begin{equation*}
\sigma_{N}(f)(\mathbf{x})=\sum_{\mathbf{k} \in \mathbb{Z}^{q}} h\left(\frac{|\mathbf{k}|_{2}}{N}\right) \hat{f}(\mathbf{k}) \exp (i \mathbf{k} \cdot \mathbf{x}), \quad \mathbf{x} \in \mathbb{T}^{q}, N>0 \tag{2.3}
\end{equation*}
$$

We note that the sums in both $(2.2$ and 2.3 are finite sums, although they are written as infinite sums for convenience of notation.

For any finite subset $\mathcal{C} \subset \mathbb{T}^{q}$, we define its minimal separation by

$$
\begin{equation*}
\eta(\mathcal{C})=\min _{\mathbf{x}, \mathbf{y} \in \mathcal{C}, \mathbf{x} \neq \mathbf{y}}|\mathbf{x}-\mathbf{y}| \tag{2.4}
\end{equation*}
$$

We assume a training data of the form $\mathcal{D}=\left\{\left(\mathbf{x}_{j}, y_{j}\right)\right\}_{j=1}^{M}$, where $\mathcal{C}=\left\{\mathbf{x}_{j}\right\}_{j=1}^{M} \subset \mathbb{T}^{q}$, and $y_{j}=f\left(\mathbf{x}_{j}\right)+\epsilon_{j}$ for some $f \in C^{*}$. We denote

$$
\begin{equation*}
\epsilon=\max _{1 \leq j \leq M}\left|\epsilon_{j}\right| . \tag{2.5}
\end{equation*}
$$

The quantity $\epsilon$ plays the role of variance in our theory in this paper. In regression problems, where numerical accuracy is expected in the supremum norm, the errors $\epsilon_{j}$ all need to be small. Thus, even though the quantity $\epsilon$ seems to increase with $M$, it remains small as $M \rightarrow \infty$. One example is when the values of $f$ are not computed exactly (as they rarely are), but only through a numerical computation (for example, $f$ is the ideal solution of a differential equation, but the data is obtained by solving this equation numerically at the grid points). The quantity $\epsilon$ then represents the maximal error in this numerical computation.

## The constant convention

The symbols $c, c_{1}, \cdots$ will denote generic positive constants, depending on such fixed parameters of the problem as $q$, $h, \mathcal{G}$, and $S$ (to be introduced later), etc. and other quantities explicitly indicated, but their values may be different at different occurrences, even within a single formula. The notation $A \sim B$ means that $c_{1} A \leq B \leq c_{2} A$.

### 2.2 Neural networks and trigonometric polynomials

The material in this section is based on [16]. For reasons that will become clear shortly, the term activation function in this paper will mean $\phi \in C^{*}$ such that $\hat{\phi}(1) \neq 0$. We note that a trigonometric polynomial is itself a neural network with the activation function $t \mapsto \cos t$. There is a close connection between trigonometric polynomials and networks with other activation functions. Let $\phi \in C^{*}(\mathbb{T})$ and $\hat{\phi}(1) \neq 0$. Then for $\mathbf{k} \in \mathbb{Z}^{q}$ and $\mathbf{x} \in \mathbb{T}^{q}$, it is not difficult to verify that

$$
\exp (i \mathbf{k} \cdot \mathbf{x})=\frac{1}{2 \pi \hat{\phi}(1)} \int_{\mathbb{T}} \phi(t) \exp (i(\mathbf{k} \cdot \mathbf{x}-t)) d t=\frac{1}{2 \pi \hat{\phi}(1)} \int_{\mathbb{T}} \phi(\mathbf{k} \cdot \mathbf{x}-t) \exp (i t) d t
$$

Discretizing the integral expression above, it can be shown (cf. [13, Proposition 4.2.1]) that for any integer $N \geq 1$, $\mathbf{k} \in \mathbb{Z}^{q}$,

$$
\begin{equation*}
\left\|\exp (i \mathbf{k} \cdot(\circ))-\frac{1}{(2 N+1) \hat{\phi}(1)} \sum_{j=0}^{2 N} \exp \left(\frac{2 \pi i j}{2 N+1}\right) \phi\left(\mathbf{k} \cdot(\circ)-\frac{2 \pi j}{2 N+1}\right)\right\| \leq \frac{4}{|\hat{\phi}(1)|} E_{N}(1 ; \phi) \tag{2.6}
\end{equation*}
$$

In particular, if $T$ is a trigonometric polynomial, let

$$
\begin{equation*}
\mathbb{G}_{N}(\phi, T)(\mathbf{x})=\frac{1}{(2 N+1) \hat{\phi}(1)} \sum_{j=0}^{2 N} \exp \left(\frac{2 \pi i j}{2 N+1}\right)\left(\sum_{\mathbf{k} \in \mathbb{Z}^{q}} \hat{T}(\mathbf{k}) \phi\left(\mathbf{k} \cdot(\mathrm{o})-\frac{2 \pi j}{2 N+1}\right)\right) \tag{2.7}
\end{equation*}
$$

where it is understood that $\hat{T}(\mathbf{k})=0$ if $|\mathbf{k}|_{2}$ exceeds the degree $n$ of $T$. The number of neurons involved in the network $\mathbb{G}_{N}(\phi, T)$ is $\sim N n^{q}$. The estimate (2.6) leads to

$$
\begin{equation*}
\left\|T-\mathbb{G}_{N}(\phi, T)\right\| \leq \frac{4}{|\hat{\phi}(1)|} E_{N}(1 ; \phi) \sum_{\mathbf{k} \in \mathbb{Z}^{q}}|\hat{T}(\mathbf{k})| . \tag{2.8}
\end{equation*}
$$

Thus, a trigonometric polynomial can be approximated by a neural network with activation function $\phi$ and the error bounds can be obtained by keeping track of the degree of approximation of the target function by trigonometric polynomials, the magnitude of its Fourier coefficients and the bound in (2.8) (cf. [16] for some examples). This leads to the following proposition, which will be proved in Section 5.

Proposition 2.1 Let $\phi \in C^{*}(\mathbb{T}), \hat{\phi}(1) \neq 0, f \in C^{*}$. Then for $n, N \geq 1$, we have

$$
\begin{align*}
\left\|f-\mathbb{G}_{N}\left(\phi, \sigma_{n}(f)\right)\right\| & =\left\|f-\frac{1}{(2 N+1) \hat{\phi}(1)} \sum_{j=0}^{2 N} \exp \left(\frac{2 \pi i j}{2 N+1}\right)\left(\sum_{|\mathbf{k}|_{2}<n} h\left(\frac{|\mathbf{k}|_{2}}{n}\right) \hat{f}(\mathbf{k}) \phi\left(\mathbf{k} \cdot(\circ)-\frac{2 \pi j}{2 N+1}\right)\right)\right\| \\
& \leq c(\phi)\left\{E_{n / 2}(q ; f)+n^{q / 2} E_{N}(1 ; \phi)\|f\|\right\} \tag{2.9}
\end{align*}
$$

Example 2.1 For example, we consider the smooth ReLU function $t \mapsto \log \left(1+e^{t}\right)=t_{+}+\mathcal{O}\left(e^{-|t|}\right)$. Then the function $\psi(t)=\log \left(\frac{\left(1+e^{t+\pi}\right)\left(1+e^{t-\pi}\right)}{\left(1+e^{t}\right)^{2}}\right)$ is integrable on $\mathbb{R}$. The periodization

$$
\begin{equation*}
\phi(t)=\sum_{j \in \mathbb{Z}} \psi(t+2 \pi j), \quad t \in \mathbb{R} \tag{2.10}
\end{equation*}
$$

is an analytic function on $\mathbb{T}$. So, Bernstein approximation theorem [28, Theorem 5.4.2] shows that there exists $\rho_{1}<1$ with $E_{N}(1 ; \phi) \leq \rho_{1}^{N}$ for all $N \geq c(\phi)$. In Proposition 2.1. if $f$ satisfies $E_{n}(q ; f)=\mathcal{O}\left(n^{-\gamma}\right)$ for some $\gamma$, then we may choose $N \sim \log n$ to get a network with $\mathcal{O}\left(n^{q} \log n\right)$ neurons to obtain an estimate $\mathcal{O}\left(n^{-\gamma}\right)$ on the right hand side of 2.9 .

This idea is generalized to many other settings, and algorithms are known to find the approximation to the target function using the training data, without assuming any prior on the target function (see, e.g. [13] for an early construction). However, formulating the problem directly as a minimization of the supremum norm error between the function and the neural network model may not work. The theory implies certain relationships between the coefficients and the weights and thresholds.

Conversely, one can approximate $\phi$ by trigonometric polynomials. Therefore, if one can obtain or assume some bounds on the coefficients of a neural network with $\phi$ as the activation function, then these bounds can be translated to bounds on the degree of approximation by trigonometric polynomials. This part is hard to do on the torus with neural networks, but has been done in a far more general setting with kernel based approximation [14, where Mercer expansions satisfying certain technical conditions are known.

In view of this close relationship between general neural networks and trigonometric polynomials (i.e., networks with activation function $t \mapsto \cos t$ ), we will focus in this paper on trigonometric polynomials, and demonstrate in Section 3 how these results are translated to those with other neural networks.

### 2.3 Interpolatory approximation

In the language of classical approximation theory, the problem of achieving a zero training error is the problem of interpolation. In the context of trigonometric polynomials, for any data of the form $\left\{\left(\theta_{j}, y_{j}\right)\right\}_{j=0}^{2 n}, y_{j} \in \mathbb{R}, \theta_{j} \in \mathbb{T}$, and $\theta_{j} \neq \theta_{k}$ if $j \neq k$, there exists $T \in \mathbb{H}_{n}^{1}$ such that $T\left(\theta_{j}\right)=y_{j}$ for $j=0, \cdots, 2 n$ [30, Chapter X, Theorem 1.2]. Thus, it is easy to obtain a zero training error with a minimal number of free parameters. As we argued in the introduction, the test error in this context should be measured in terms of uniform approximation to the target function. A well known theorem attributed in 21 to Faber and Bernstein states that for any system of interpolation nodes, there exists a function $f \in C^{*}(\mathbb{T})$ for which the sequence of interpolatory trigonometric polynomials (with minimal degree as above) does not converge uniformly to $f$.

In 1943, Erdős [6] initiated (in the context of algebraic polynomials) a study of the question whether one can get a convergent sequence of interpolatory polynomials if one allows a polynomial of higher than minimal degree. A positive answer was given by Szabados in [27]. Although the answer is given in terms of algebraic polynomials, Szabados remarks that the same is clearly true for trigonometric polynomials as well. An explicit statement to this effect is the following [17, Theorem 3.1(a)].

Theorem 2.1 Let $\theta_{1}, \ldots, \theta_{N}$ be distinct points in $[-\pi, \pi], \theta_{N+1}:=\theta_{1}+2 \pi, \alpha>0$, and

$$
\min _{1 \leq k \leq N}\left|\theta_{k+1}-\theta_{k}\right|=: \eta .
$$

Then for $f \in C^{*}\left(\mathbb{T}^{1}\right)$, there exists a trigonometric polynomial $T$ of degree at most $(1+2 / \alpha)(\pi / \eta)$ such that $f\left(\theta_{j}\right)=T\left(\theta_{j}\right), 1 \leq j \leq N$, and

$$
\begin{equation*}
\|f-T\| \leq(2+\alpha) E_{m}(1 ; f) \tag{2.11}
\end{equation*}
$$

where $m=(1+2 / \alpha)(\pi / \eta)$.

Remark 2.1 A curious feature of Theorem 2.1 is that one obtains the bound 2.11) on the generalization on the entire torus $\mathbb{T}^{1}$ without requiring that the training data $\mathcal{C}$ be "dense" in $\mathbb{T}^{1}$.

## 3 Shallow networks

Our first theorem in this section shows the connection between the structural properties of the training data and the construction of a trigonometric polynomial $T_{N}^{\#}(\mathcal{D})$ that can interpolate the noisy data (i.e., achieve a zero training error), as well as achieve a good generalization error in the sense defined in Section 1 .

Before stating our main results, we first discuss two straightforward ideas. The first of these is to construct a trigonometric interpolatory polynomial $\mathcal{I}(\mathcal{D})$ of minimal degree. The other is to use a larger degree $N$, and solve the system of equations (cf. 2.2 )

$$
\begin{equation*}
\sum_{j=1}^{M} a_{j} \Phi_{N}\left(\mathbf{x}_{\ell}-\mathbf{x}_{j}\right)=y_{\ell}, \quad \ell=1, \cdots, M \tag{3.1}
\end{equation*}
$$

For a sufficiently large value of $N$, one can show that this sytem of equations has a unique and stable solution. We denote the corresponding polynomial by

$$
\begin{equation*}
\mathcal{L}_{N}(\mathcal{D})(\mathbf{x})=\sum_{j=1}^{M} a_{j} \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right) \tag{3.2}
\end{equation*}
$$

We examine these constructions using a univariate example. We consider $f(x)=|\cos x|$, and consider only the case when exact samples of $f$ are known at 128 points. We note that for $x \in \mathbb{T}$,

$$
f(x)=\frac{2}{\pi}\left\{1-\sum_{k=1}^{\infty} \frac{(-1)^{k}}{4 k^{2}-1} \cos (2 k x)\right\}
$$

Therefore, the sequence $\left\{\left(k^{2}+1\right)^{s / 2} \hat{f}(k)\right\}$ represents a sequence of Fourier coefficients of a continuous function for all $s<1$, but not if $s=1$ (and of course, not for any $s>1$ ). In particular, there is no "optimal" class with which to use the results of [3] for this function.

Example 3.1 We consider $M=128, x_{j}=-\pi+2 \pi j / M, j=0, \cdots, 127$, and denote the corresponding data by $\mathcal{D}_{1}$. We will refer to this choice as the dense case. Figure 1 shows the errors at 1024 points on $[-\pi, \pi]$ for the operators $\mathcal{I}\left(\mathcal{D}_{1}\right) \in \mathbb{H}_{64}^{1}$ and $\mathcal{L}_{N}\left(\mathcal{D}_{1}\right) \in \mathbb{H}_{128}^{1}$ with $N=128$. It is noted that the collocation matrix for computing $\mathcal{I}\left(\mathcal{D}_{1}\right)$ is very ill-conditioned, but the matrix for computing $\mathcal{L}_{N}\left(\mathcal{D}_{1}\right)$ is well-conditioned. Of course, in this very special example, one can compute $\mathcal{I}\left(\mathcal{D}_{1}\right)$ explicitly without having to solve a system of equations, but the intent of this example is to demonstrate the need to have methods more sophisticated than a straightforward interpolation.

The situation is quite different when the data is not dense on $[-\pi, \pi)$.
Example 3.2 We consider $M=128, x_{j}=\pi / 4+\pi j /(2 M), j=0, \cdots, 127$, and denote the corresponding data by $\mathcal{D}_{2}$. We will refer to this choice as the non-dense case. Figure 2 shows the errors at 1024 points on $[-\pi, \pi]$ for the operators $\mathcal{I}\left(\mathcal{D}_{2}\right)$ and $\mathcal{L}_{N}\left(\mathcal{D}_{2}\right)$ with $N=128$, and $N=256$. As before, the collocation matrix for computing $\mathcal{I}\left(\mathcal{D}_{2}\right)$ is very ill-conditioned, but the matrix for computing $\mathcal{L}_{N}\left(\mathcal{D}_{2}\right)$ is well-conditioned. However, a comparison between the middle and right figure of Figure 2 shows the critical importance of choosing a degree higher than the minimal possible, commensurate with the minimal separation among the interpolation nodes. Also, we note in the right figure that the polynomial $\mathcal{L}_{256}\left(\mathcal{D}_{2}\right)$ is highly localized, so that the errors on the interval $[\pi / 4, \pi / 2)$ are small, but those away from this interval are not; in fact, the polynomial is close to 0 away from this interval.

As remarked in Remark 2.1, it is not possible to achieve a bound analogous to 2.11) unless the training data is sufficiently dense on $\mathbb{T}^{q}$ or unless more information about the target function is used in addition to its values at the training data. Therefore, let us now assume that the Fourier coefficients $\hat{f}(\mathbf{k})$ are known for all $\mathbf{k}$ with $|\mathbf{k}|_{2}<n$


Figure 1: With $x_{j}=-\pi+2 \pi j / M, j=0, \cdots, 127$, the differences at 1024 equidistant points of $[-\pi, \pi)$ between the true values $|\cos x|$ and the interpolatory polynomial of minimal degree, $\mathcal{I}\left(\mathcal{D}_{1}\right)(x)$ (left), and interpolation obtained by a higher degree localized kernel $\mathcal{L}_{128}\left(\mathcal{D}_{1}\right)(x)$ (right). The errors are magnified 1000 times and the points on the $x$-axis are multiples of $\pi$. The errors are similar near $\pm \pi / 2$, but decrease rapidly to 0 away from these in the figure on the right.


Figure 2: With $x_{j}=\pi / 4+\pi j /(256), j=0, \cdots, 127$, the differences at 1024 equidistant points of $[-\pi, \pi)$ between the true values $|\cos x|$ and the interpolatory polynomial of minimal degree, $\mathcal{I}\left(\mathcal{D}_{2}\right)(x)$ (left), interpolation obtained by a higher degree localized kernel $\mathcal{L}_{128}\left(\mathcal{D}_{2}\right)(x)$ (middle), and $\mathcal{L}_{256}\left(\mathcal{D}_{2}\right)(x)$ (right). The errors are magnified 1000 times and the points on the $x$-axis are multiples of $\pi$.
for some $n>0$. With this information, we can construct $\sigma_{n}(f)$ using 2.3 with $n$ in place of $N$. For a sufficiently large value of $N$, we may solve a system of equations

$$
\begin{equation*}
\sum_{j=1}^{M} a_{j} \Phi_{N}\left(\mathbf{x}_{\ell}-\mathbf{x}_{j}\right)=y_{\ell}-\sigma_{n}(f)\left(\mathbf{x}_{\ell}\right), \quad \ell=1, \cdots, M \tag{3.3}
\end{equation*}
$$

We then define

$$
\begin{equation*}
\mathcal{T}_{n, N}^{\#}(\mathcal{D})(\mathbf{x})=\sigma_{n}(f)(\mathbf{x})+\sum_{j=1}^{M} a_{j} \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right) \tag{3.4}
\end{equation*}
$$

Example 3.3 We continue the examples with the target functions as in Examples 3.1 and 3.2 . We estimate $\hat{f}(k)$ using a 128 point discrete Fourier transform. In the dense case, this is done using only the training data. In the non-dense case, the coefficients is the additional information we need, apart from the training data itself. In the calculation of the operators $\mathcal{T}_{n, N}^{\#}$, we use $n=128, N=256$. The resulting errors are reported in Figure 3 . We note that in the non-dense case, the errors are uniformly small on the entire interval $[-\pi, \pi)$, in contrast to the errors in the right-most figure in Figure 2. Also, the errors near the interpolation nodes are smaller than near the other singularity of the target function at $-\pi / 2$. In the non-dense case, we might as well use a larger number of Fourier coefficients to make $n=N$. We did this using 1024 point discrete Fourier transform, and $n=N=256$. The results are shown in Figure 3 as well.

Theorem 3.1 below is a generalization of Theorem 2.1 , showing that the operators $\mathcal{T}_{N}^{\#}=\mathcal{T}_{N, N}^{\#}$ yield the requisite approximation. Our proof is much simpler than that of Theorem 2.1] in either [17] or [27].


Figure 3: The errors at points of $[-\pi, \pi)$ using $\mathcal{T}_{128,256}^{\#}\left(\mathcal{D}_{1}\right)$ on the upper left, $\mathcal{T}_{128,256}^{\#}\left(\mathcal{D}_{2}\right)$ in the upper middle. The upper right figure is the error for $\mathcal{T}_{128,256}^{\#}\left(\mathcal{D}_{2}\right)$ at the interpolation nodes, to verify that the operator does the interpolation correctly. The bottom left and bottom right figures are analogous to the upper middle and upper right figures with $\mathcal{T}_{256,256}^{\#}\left(\mathcal{D}_{2}\right)$. All errors are maginfied 1000 times, and the $x$ axis has multiples of $\pi$.

Theorem 3.1 There exists $B>0$ with the following property: for $f \in C^{*}$ and $N \geq B \eta(\mathcal{C})^{-1}$, the trigonometric polynomial $T_{N}^{\#}(\mathcal{D})=\mathcal{T}_{N, N}^{\#}(\mathcal{D}) \in \mathbb{H}_{N}^{q}$ in (3.4) is well defined, and satisfies

$$
\begin{equation*}
T_{N}^{\#}(\mathcal{D})\left(\mathbf{x}_{j}\right)=y_{j}, \quad j=1, \cdots, M, \quad \text { (Zero training error) } \tag{3.5}
\end{equation*}
$$

and

$$
\begin{equation*}
\left\|f-T_{N}^{\#}(\mathcal{D})\right\| \leq c\left\{\epsilon+E_{N / 2}(f)\right\} . \quad \text { (Good generalization error). } \tag{3.6}
\end{equation*}
$$

In view of the observations in Section 2.2 , the proof of the above theorem can be modified with neural network approximations at each stage to obtain the following version, where we overload the notation a bit for simplicity.

Suppose the following system of equations has a solution for some $\tilde{N}$ and $N$ (cf. 2.7p):

$$
\begin{equation*}
\sum_{j=1}^{M} a_{j} \mathbb{G}_{\tilde{N}}\left(\phi, \Phi_{N}\right)\left(\mathbf{x}_{\ell}-\mathbf{x}_{j}\right)=y_{\ell}-\mathbb{G}_{\tilde{N}}\left(\phi, \sigma_{N}(f)\right)\left(\mathbf{x}_{\ell}\right), \quad \ell=1, \cdots, M \tag{3.7}
\end{equation*}
$$

We then define

$$
\begin{equation*}
\mathbb{G}_{\tilde{N}, N}^{\#}(\mathcal{D})(\mathbf{x})=\mathbb{G}_{\tilde{N}}\left(\phi, \sigma_{N}(f)\right)(\mathbf{x})+\sum_{j=1}^{M} a_{j} \mathbb{G}_{\tilde{N}}\left(\phi, \Phi_{N}\right)\left(\mathbf{x}-\mathbf{x}_{j}\right) \tag{3.8}
\end{equation*}
$$

Theorem 3.2 There exist $B, \alpha^{*}>0$ with the following property. Let $f \in C^{*}, \phi \in C^{*}(\mathbb{T}), \hat{\phi}(1) \neq 0, N \geq B \eta(\mathcal{C})^{-1}$. Let $\tilde{N}$ be such that

$$
\begin{equation*}
\left\|\Phi_{N}-\mathbb{G}_{\tilde{N}}\left(\phi, \Phi_{N}\right)\right\| \leq \alpha^{*} \tag{3.9}
\end{equation*}
$$

Then the network $\mathbb{G}_{N}^{\#}(\mathcal{D})=\mathbb{G}_{\tilde{N}, N}^{\#}(\mathcal{D})$ in (3.8) is well defined and satisfies

$$
\begin{equation*}
\mathbb{G}_{N}^{\#}(\mathcal{D})\left(\mathbf{x}_{j}\right)=y_{j}, \quad j=1, \cdots, M, \quad \text { (Zero training error) } \tag{3.10}
\end{equation*}
$$

and

$$
\begin{equation*}
\left\|f-\mathbb{G}_{N}^{\#}(\mathcal{D})\right\| \leq c\left\{\epsilon+E_{N / 2}(f)+N^{q / 2} E_{\tilde{N}}(1 ; \phi)\|f\|\right\} . \quad \text { (Good generalization error). } \tag{3.11}
\end{equation*}
$$

Remark 3.1 A volume comparison argument shows that the number $M$ of data points satisfies $M \leq c \eta(\mathcal{C})^{-q}$. Thus, Theorems 3.1 and 3.2 show that for a right configuration of the training data, a good generalization error as well as zero training error can be achieved by choosing the number of parameters proportional to the number of data points. It is demonstrated in [2] that for RBF approximation, this phenomenon seems to hold in many applications with the number of parameters exactly equal to the number of data points.

Remark 3.2 In practice, the training data is high dimensional and sparse; i.e., $\eta(\mathcal{C})$ is large. The requirement that $N \geq B \eta(\mathcal{C})^{-1}$ is therefore satisfied with moderate degrees $N$.

As demonstrated before, one cannot construct $T_{N}^{\#}(\mathcal{D})$ based only on the training data $\mathcal{C}$, unless $\mathcal{C}$ is sufficiently dense on $\mathbb{T}^{q}$. We have already discussed an effort in the form of the operators $\mathcal{L}_{N}(\mathcal{D})$ in Example 3.2 An even simpler approach of just considering

$$
\frac{1}{\Phi_{N}(0)} \sum_{k=1}^{M} y_{k} \Phi_{N}\left(\circ-\mathbf{x}_{k}\right)
$$

yields similar bounds on the generalization error as those obtained by $\mathcal{L}_{N}(\mathcal{D})$. Of course, the training error for this simple construction is not 0 , but the localization properties of $\Phi_{N}$ ensure that it is small. In general, if we anticipate a scenario where the training data sets become increasingly dense on some compact subset $K \subset \mathbb{T}^{q}$, then we cannot expect convergence of trigonometric polynomials that interpolate a noisy data, where the noise level does not decrease as well.

Another approach, described in 3], is to minimize a high order Sobolev norm of the trigonometric polynomial subject to the interpolatory conditions. This approach has been used to great advantage for a numerical solution of some notoriously hard partial differential equations in 2 or 3 dimensions. However, the calculations are very ill-conditioned and require very carefully designed algorithms.

We describe a softer regularization scheme that does not require high order Sobolev norms, and yields both good training and generalization errors. The generalization error is given point-wise, and is commensurate with the estimates given in Theorem 3.1 when there is no noise.

The space $W^{*}=W^{*}\left(\mathbb{T}^{q}\right)$ consists of all continuously differentiable functions $f \in C^{*}$. We define

$$
\begin{equation*}
\|f\|_{W^{*}}=\|f\|_{W^{*}\left(\mathbb{T}^{q}\right)}=\sum_{j=1}^{q}\left\|D_{j} f\right\| \tag{3.12}
\end{equation*}
$$

For $n>0$ and $T \in \mathbb{H}_{n}^{q}$, let

$$
\begin{equation*}
R_{n}(T)=\max _{1 \leq j \leq M}\left|y_{j}-T\left(\mathbf{x}_{j}\right)\right|+\frac{1}{n}\|T\|_{W^{*}} \tag{3.13}
\end{equation*}
$$

Theorem 3.3 Let $f \in W^{*}, B$ be as in Theorem 3.1, $N \geq B \eta(\mathcal{C})^{-1}$, and $T^{*}(\mathcal{D})=\underset{T \in \mathbb{H}_{N}^{q}}{\arg \min } R_{N}(T)$. Then

$$
\begin{equation*}
\max _{1 \leq j \leq M}\left|y_{j}-\mathcal{T}^{*}(\mathcal{D})\left(\mathbf{x}_{j}\right)\right| \leq \min _{T \in \mathbb{H}_{N}^{q}} R_{N}(T) \leq c\left\{\epsilon+\frac{1}{N}\|f\|_{W^{*}}\right\} . \quad \text { (Good training error) } \tag{3.14}
\end{equation*}
$$

Let $\mathbf{x} \in \mathbb{T}^{q}$, and $\delta=\min _{1 \leq j \leq M}\left|\mathbf{x}-\mathbf{x}_{j}\right|$. Then

$$
\begin{equation*}
\left|f(\mathbf{x})-T^{*}(\mathcal{D})(\mathbf{x})\right| \leq c(1+N \delta)\left\{\epsilon+\frac{1}{N}\|f\|_{W^{*}}\right\} . \quad \text { (Good generalization error) } \tag{3.15}
\end{equation*}
$$

Remark 3.3 Theorem 3.3 makes no assumption on the target function except for differentiability. This is in contrast to usual machine learning theory, where one has to assume that the target function belongs to a reproducing kernel Hilbert space, with the kernel prescribed by the learning algorithm.

Remark 3.4 The estimate (3.15 shows that if $\mathbf{x}$ is very close to the training data so that $N \delta<1$, then the generalization error at $\mathbf{x}$ satisfies the same upper bound (3.14) that holds for $R_{N}\left(T^{*}\right)$. As remarked earlier in Remark 3.2 , we have greater liberty in choosing $N$ when the data is sparse. To take advantage of this fact, let $\mathbf{x}$ be such that $N \delta \geq 1$. Then 3.15 can be reformulated in the form

$$
\begin{equation*}
\left|f(\mathbf{x})-T^{*}(\mathcal{D})(\mathbf{x})\right| \leq c \delta\left\{N \epsilon+\|f\|_{W^{*}}\right\} \tag{3.16}
\end{equation*}
$$

The term $\delta\|f\|_{W^{*}}$ is clearly a customary bound from numerical analysis in view of the mean value theorem. If $\epsilon \leq \eta(\mathcal{C})\|f\|_{W^{*}}$, it is possible to choose $N \sim \eta(\mathcal{C})^{-1}$ so that error bound is $c \delta\|f\|_{W^{*}}$.

Remark 3.5 It is well known (cf. [30, Chapter X, Theorem 7.28] for the univariate case) that for any $T \in \mathbb{H}_{N}^{q}$,

$$
\begin{equation*}
\|T\| \sim \max _{|\mathfrak{m}| \infty \leq 3 N-1}\left|T\left(\frac{2 \pi \mathfrak{m}}{3 N}\right)\right| . \tag{3.17}
\end{equation*}
$$

Therefore,

$$
R_{N}(T) \sim \max _{1 \leq j \leq M}\left|y_{\ell}-\sum_{\mathbf{k}} \hat{T}(\mathbf{k}) \exp \left(i \mathbf{k} \cdot \mathbf{x}_{j}\right)\right|+\frac{1}{N} \sum_{j=1}^{q} \max _{|\mathfrak{m}|_{\infty} \leq 3 N-1}\left|\sum_{\mathbf{k}} k_{j} \hat{T}(\mathbf{k}) \exp \left(2 \pi i \frac{\mathbf{k} \cdot \mathfrak{m}}{3 N}\right)\right| .
$$

If we replace $R_{N}(T)$ by the expression on the right hand side of the above equation, we get an optimization problem directly in terms of the coefficients $\hat{T}(\mathbf{k})$. The theoretical results are not affected except for the actual values of the constants involved.

Remark 3.6 In this remark, let $T^{*}=\mathcal{T}^{*}(\mathcal{D})$. In view of (2.8), the estimates (3.14) and (3.15) imply for any activation function $\phi$,

$$
\max _{1 \leq j \leq M}\left|y_{j}-\mathbb{G}_{\tilde{N}}\left(\phi, T^{*}\right)\left(\mathbf{x}_{j}\right)\right| \leq c\left\{\epsilon+\frac{1}{N}\|f\|_{W^{*}}\right\}+\frac{4 E_{\tilde{N}}(1 ; \phi)}{\phi \hat{(1)}} \sum_{\mathbf{k}}\left|\widehat{T^{*}}(\mathbf{k})\right|,
$$

and

$$
\left|f(\mathbf{x})-\mathbb{G}_{\tilde{N}}\left(\phi, T^{*}\right)(\mathbf{x})\right| \leq c(1+N \delta)\left\{\epsilon+\frac{1}{N}\|f\|_{W^{*}}\right\}+\frac{4 E_{\tilde{N}}(1 ; \phi)}{\phi \hat{(1})} \sum_{\mathbf{k}}\left|\widehat{T^{*}}(\mathbf{k})\right|
$$

respectively. In particular, for the activation function obtained from the smooth ReLU function, the extra error terms decrease exponentially rapidly with $\tilde{N}$. It is therefore tempting to set up the regularization functional (3.13) directly with neural networks with free coefficients, weights, and thresholds to be determined by a suitable optimization technique. However, the estimate $(2.8)$ implies a strong connection between these parameters for the network approximating a trigonometric polynomial. Therefore, the solution to such a direct approach with neural networks is not guaranteed to give the right training and testing errors.

## 4 Deep networks

The following discussion regarding the terminology for deep networks, including Figure 4 is based on the discussion in [20], and elaborates upon the same.

A commonly used definition of a deep network is the following. Let $\phi: \mathbb{R} \rightarrow \mathbb{R}$ be an activation function; applied to a vector $\mathbf{x}=\left(x_{1}, \cdots, x_{q}\right), \phi(\mathbf{x})=\left(\phi\left(x_{1}\right), \cdots, \phi\left(x_{q}\right)\right)$. Let $L \geq 1$ be an integer, for $\ell=0, \cdots, L$, let $q_{\ell} \geq 1$ be an integer $\left(q_{0}=q\right), T_{\ell}: \mathbb{R}^{q_{\ell}} \rightarrow \mathbb{R}^{q_{\ell+1}}$ be an affine transform, where $q_{L+1}=1$. A deep network with $L$ hidden layers is defined as the compositional function

$$
\begin{equation*}
\mathbf{x} \mapsto T_{L}\left(\phi \left(T _ { L - 1 } \left(\phi\left(T_{L-2} \cdots \phi\left(T_{0}(\mathbf{x})\right) \cdots\right)\right.\right.\right. \tag{4.1}
\end{equation*}
$$

This definition has several shortcomings. First, a function may have more than one compositional representation, as we will demonstrate shortly, so that the affine transforms and $L$ are not determined uniquely by the function itself. Second, this notion does not capture the connection between the nature of the target function and its approximation. Third, the affine transforms $T_{\ell}$ define a special directed acyclic graph (DAG). It is cumbersome to describe notions of weight sharing, convolutions, sparsity, skipping of layers, etc. in terms of these transforms. Therefore, we have proposed in [20], to describe a deep network more generally as a directed acyclic graph (DAG) architecture.

Let $\mathcal{G}$ be a DAG, with the set of nodes $V \cup \mathbf{S}$, where $\mathbf{S}$ is the set of source nodes, and $V$ that of non-source nodes. A $\mathcal{G}$-function is defined as follows. Each of the in-edges to each node in $V \cup S$ represents an input real variable. For each node $v \in V \cup \mathbf{S}$, we denote its in-degree by $d(v)$. A node $v \in V \cup \mathbf{S}$ itself represents the evaluation of a real valued function $f_{v}$ of the $d(v)$ inputs. The out-edges fan out the result of this evaluation. Each of the source nodes obtains an input from some Euclidean space. Other nodes can also obtain such an input, but by introducing dummy nodes, it is convenient to assume that only the source nodes obtain an input from the Euclidean space.

The notion of the level (or height) of a node is defined as follows. The level of a source node is 0 ; it represents a shallow network. The level of $v \in V$ is the length of the longest path from the nodes in $\mathbf{S}$ to $v$. The level of $v$ will be denoted by $H(v)$.

In 20, we have argued that deep networks display better approximation properties than shallow networks because they can take advantage of a compositional structure in the target function which shallow networks cannot. However, compositionality is the property of an expression for the function; not an intrinsic property of a function itself. A simple example in the univariate case is the constant function $f(x) \equiv 2, x \in[0,1]$, that can also be expressed as a compositional function

$$
f(x)=(x+1) \cosh \left(\log \left(\frac{2+\sqrt{3-2 x-x^{2}}}{x+1}\right)\right), \quad x \in[0,1]
$$

It is not clear whether two different DAG structures can give rise to the same function. Even if we assume a certain DAG, it is not clear that the choice of the constituent functions is uniquely determined for a given function on $\mathbb{R}^{q}$. For example, one can write

$$
\cos ^{4} x=\left((\cos x)^{2}\right)^{2}=\left(\cos ^{2} x\right)^{2}=(1 / 4)(1+\cos (2 x))^{2}
$$

The second expression above has the structure $h_{1}\left(h_{2}\left(h_{3}(x)\right)\right)$, and the other two have the structures $g_{1}\left(g_{2}(x)\right)$ or $f_{1}\left(f_{2}(x)\right)$, both representing the same DAG but with different constituent functions. Thus, the question of whether a given multivariate function is in fact compositional cannot be answered. Of course, for a given DAG, it is possible to use inverse/implicit function theorem (in theory) in some cases to decide whether a family of functions are compositional according to the given DAG.

For our mathematical analysis, we therefore find it convenient to think of a $\mathcal{G}$-function as a set of functions $f=\left\{f_{v}: \mathbb{R}^{d(v)} \rightarrow \mathbb{R}\right\}_{v \in V \cup \mathbf{S}}$, rather than a single function on $\mathbb{R}^{q}$. For example, the DAG $\mathcal{G}$ in Figure $4([20])$ represents the compositional function

$$
\begin{align*}
& f^{*}\left(x_{1}, \cdots, x_{9}\right)=\quad h_{19}\left(h _ { 1 7 } \left(h_{13}\left(h_{10}\left(x_{1}, x_{2}, x_{3}\right), h_{11}\left(x_{4}, x_{5}\right)\right)\right.\right. \\
&\left.h_{14}\left(h_{10}, h_{11}\right), h_{16}\left(h_{1} 2\left(x 6, x_{7}, x_{8}, x_{9}\right)\right), h_{18}\left(h_{15}\left(h_{11}, h_{12}\right), h_{16}\right)\right) . \tag{4.2}
\end{align*}
$$

The $\mathcal{G}$-function is $\left\{h_{10}, \cdots, h_{19}=f^{*}\right\}$. The individual functions $f_{v}$ will be called constituent functions.
We assume that there is only one sink node, $v^{*}\left(\right.$ or $\left.v^{*}(\mathcal{G})\right)$ whose output is denoted by $f^{*}$ (the target function). Technically, there are two functions involved here: one is the final output as a function of all the inputs to all source nodes, the other is the final output as a function of the inputs to the node $v^{*}$. We will use the symbol $f^{*}$ to denote both with comments on which meaning is intended when we feel that it may not be clear from the context. A similar convention is followed with respect to each of the constituent functions as well. For example, in the DAG of Figure 4, the function $h_{14}$ can be thought of both as a function of two variables, namely the outputs of $h_{10}$ and $h_{11}$ as well as a function of five variables $x_{1}, \cdots, x_{5}$. In particular, if each constituent function is a neural network, $h_{14}$ is a shallow network receiving two inputs.

In this paper, we are interested only in the case where each of the inputs to each of the source nodes is in $\mathbb{T}$ rather than $\mathbb{R}$. Although this is no longer true for the non-source nodes, it is possible to accomplish this in the case when each of the constituent functions is continuous, as follows. We observe first that there is a one-to-one correspondence between functions on $[-1,1]^{d}$ and functions on $\mathbb{T}^{d}$ that are even in each variable, given by

$$
F^{\circ}\left(\theta_{1}, \cdots, \theta_{d}\right)=F\left(\cos \theta_{1}, \cdots, \cos \theta_{d}\right), \quad\left(\theta_{1}, \cdots, \theta_{d}\right) \in \mathbb{T}^{d}
$$

Let $f_{u}$ be one of the constituent functions. With a re-normalization, we may assume that the range of $f_{u}$ is a subset of $[-1,1]$. If $u_{1}, \cdots, u_{d(v)}$ are children of $v$ (i.e., if there is an edge from each $u_{1}, \cdots, u_{d(v)}$ to $v$ ), then $f_{v}$ can be seen as a function on $[-1,1]^{d}$, from which one can construct an even function $f_{v}^{\circ}$ on $\mathbb{T}^{d}$ as just explained : informally, think of $f_{v}$ as an even function of the points $\left( \pm \arccos \left(f_{u_{1}}\right), \cdots, \pm \arccos \left(f_{u_{d(v)}}\right)\right) \in \mathbb{T}^{d}$. Rather than complicating our notations, we will therefore assume in this paper that the domain of each constituent function is a torus, and the range is a subset of $\mathbb{T}$. In particular, we may assume that every constituent function $f_{v} \in C^{*}\left(\mathbb{T}^{d(v)}\right)$.

We adopt the convention that for any function class $\mathbb{X}\left(\mathbb{T}^{d}\right)$, the class $\mathcal{G}-\mathbb{X}$ denotes the set of $\mathcal{G}$-functions $f=$ $\left\{f_{v}\right\}_{v \in V}$, where each constituent function $f_{v} \in \mathbb{X}\left(\mathbb{T}^{d(v)}\right)$. We define

$$
\begin{equation*}
\|f\|_{\mathbb{X}, \mathcal{G}}=\sum_{v \in V \cup \mathbf{S}}\left\|f_{v}\right\|_{\mathbb{X}\left(\mathbb{T}^{d(v)}\right)} \tag{4.3}
\end{equation*}
$$

Thus, for example, $\mathcal{G}$ - $W^{*}$ is the set of all $\mathcal{G}$-functions $f=\left\{f_{v}\right\}_{v \in V}$, where each constituent function $f_{v} \in \mathbb{X}\left(\mathbb{T}^{d(v)}\right)$, and we write

$$
\|f\|_{W^{*}, \mathcal{G}}=\sum_{v \in V \cup \mathbf{S}}\left\|f_{v}\right\|_{W^{*}\left(\mathbb{T}^{d(v)}\right)}
$$



Figure 4: An example of a $\mathcal{G}$-function ( $f^{*}$ given in $(4.2)$ ). The vertices of the DAG $\mathcal{G}$ are denoted by red dots. The black dots represent the input to the various nodes as indicated by the in-edges of the red nodes, and the blue dot indicates the output value of the $\mathcal{G}$-function, $f^{*}$ in this example.

For a vector $\mathbf{N}=\left(N_{v}\right)$, the symbol $\mathcal{G}-\mathbb{H}_{\mathbf{N}}$ denotes the set of $\mathcal{G}$-functions $\left\{T_{v}\right\}_{v \in V}$, where each $T_{v} \in \mathbb{H}_{N_{v}}^{d(v)}$, and for a $\mathcal{G}$-function $f=\left\{f_{v}\right\}_{v \in V}$,

$$
\begin{equation*}
E_{\mathbf{N}, \mathcal{G}}(f)=\sum_{v \in V \cup \mathbf{S}} E_{N_{v}}\left(d(v) ; f_{v}\right) . \tag{4.4}
\end{equation*}
$$

To make precise the various inputs to the constituent functions, we introduce some conventions. Let $\mathbf{S}=$ $\left\{v_{1}, \cdots, v_{s}\right\}, q=\sum_{j=1}^{s} d\left(v_{j}\right)$. The input $\mathbf{x}$ can be viewed as a vector in $\mathbb{R}^{q}$, but also as an element of $\prod_{j=1}^{s} \mathbb{R}^{d\left(v_{j}\right)} ;$ i.e., $\mathbf{x}=\left((\mathbf{x})_{v_{1}}, \cdots,(\mathbf{x})_{v_{s}}\right)$ so that $(\mathbf{x})_{v_{j}} \in \mathbb{R}^{d\left(v_{j}\right)}$. We note that in making this statement, we have tacitly introduced some possible dummy variables here: for example, the function with the compositional structure

$$
F\left(x_{1}, x_{2}, x_{3}\right)=f\left(f_{1}\left(x_{1}, x_{2}\right), f_{2}\left(x_{2}, x_{3}\right)\right)
$$

is viewed as a function of 4 variables

$$
\tilde{F}\left(x_{1}, \cdots, x_{4}\right)=f\left(f_{1}\left(x_{1}, x_{2}\right), f_{2}\left(x_{4}, x_{3}\right)\right)
$$

although $F$ is the restriction of $\tilde{F}$ to the hyper-plane $x_{2}=x_{4}$. In our opinion, it is a very deep question to figure out what domains the functions in practice are defined on, and to some extent, manifold learning addresses this issue. Much of the approximation theory literature is however limited to approximation on cubes, spheres, and other known domains, torus in the present paper. From a purely mathematical point of view, this is justified by an appeal to Stein's extension theorem [26, Chapter VI]. Therefore, making this tacit introduction of dummy variables allows us to simplify notation and results without any theoretical loss of generality.

Correspondingly, we define the sets

$$
\mathcal{C}_{v_{j}}=\left\{(\mathbf{x})_{v_{j}}: \mathbf{x} \in \mathcal{C}\right\}, \quad j=1, \cdots, s
$$

Thus, $\mathcal{C}_{v_{j}}$ is the training data "seen" by the source node $v_{j}$. This notion is extended recursively to other nodes of $\mathcal{G}$. Let $v$ not be a source node, $u_{1}, \cdots, u_{d(v)}$ be the children of $v, \mathcal{C}_{u_{1}}, \cdots, \mathcal{C}_{u_{d(v)}}$ be the training data seen by these nodes. Thus, for any $\mathbf{x} \in \mathcal{C}$, the components given by $(\mathbf{x})_{u_{j}}$ are seen by $u_{j}$, and

$$
\mathcal{C}_{v}=\left\{\left(f_{u_{1}}\left((\mathbf{x})_{u_{1}}\right), \cdots, f_{u_{d(v)}}\left((\mathbf{x})_{u_{d(v)}}\right)\right): \mathbf{x} \in \mathcal{C}\right\}
$$

For example, in the DAG of Figure 4, the children of $h_{14}$ are $h_{10}$ and $h_{11}$. For each $\mathbf{x} \in \mathbb{R}^{9}, h_{10}$ sees the components $\left(x_{1}, x_{2}, x_{3}\right)$, while $h_{11}$ sees the components $\left(x_{4}, x_{5}\right)$. We have

$$
\begin{aligned}
\mathcal{C}_{h_{10}} & =\left\{\left((\mathbf{x})_{1},(\mathbf{x})_{2},(\mathbf{x})_{3}\right): \mathbf{x} \in \mathcal{C}\right\}, \quad \mathcal{C}_{h_{11}}=\left\{\left((\mathbf{x})_{4},(\mathbf{x})_{5}\right): \mathbf{x} \in \mathcal{C}\right\} \\
\mathcal{C}_{h_{14}} & =\left\{\left(\left(h_{10}\left((\mathbf{x})_{1},(\mathbf{x})_{2},(\mathbf{x})_{3}\right), h_{11}\left((\mathbf{x})_{4},(\mathbf{x})_{5}\right)\right): \mathbf{x} \in \mathcal{C}\right\}\right.
\end{aligned}
$$

In this section, we state our theorems only for networks with activation function $t \mapsto \cos t$; i.e., networks that evaluate trigonometric polynomials. The transition to networks with other activation functions is obtained by approximating these by trigonometric polynomials using 2.8 as in Section 3. This only adds to an additional complication in the notation without adding anything new conceptually.

The analogue of Theorem 3.1 is the following.
Theorem 4.1 Let $\mathcal{G}$ be a $D A G$ with sink node $v^{*}$. There exists $C=C(\mathcal{G})>0$ with the following property: Let $f=\left\{f_{v}\right\}_{v \in V \cup \mathbf{S}}$ be a $\mathcal{G}$-function such that each of the constituent functions $\left\{f_{v}\right\}_{v \in V}$ is Lipschitz continuous with Lipschitz constant $\leq L$. If $N_{v} \geq C(\mathcal{G}) \eta\left(\mathcal{C}_{v}\right)^{-1}$, for each $v \in V$, and $\mathbf{N}=\left(N_{v}\right)$, the $\mathcal{G}$-function $T_{\mathbf{N}}^{\#}(\mathcal{D}) \in \mathcal{G}-\mathbb{H}_{\mathbf{N}}$ defined in 5.36) satisfies

$$
\begin{equation*}
\left(T_{\mathbf{N}}^{\#}(\mathcal{D})\right)_{v^{*}}\left(\mathbf{x}_{j}\right)=y_{j}, \quad j=1, \cdots, M, \quad \text { (Zero training error) } \tag{4.5}
\end{equation*}
$$

and

$$
\begin{equation*}
\left\|f-T_{\mathbf{N}}^{\#}(\mathcal{D})\right\|_{C^{*}, \mathcal{G}} \leq c(\mathcal{G}, L)\left\{\epsilon+E_{\mathbf{N} / 2, \mathcal{G}}(f)\right\} . \quad \text { (Good generalization error) } \tag{4.6}
\end{equation*}
$$

We recall that the polynomials in Theorem 3.1 and hence, in Theorem 4.1 are constructed using the Fourier coefficients of the various functions involved. One could use the polynomials $\mathcal{L}_{N}(\mathcal{D})$ defined in 3.2 instead to get a construction based only on the values of the various functions computed using the training data. However, since we do not know the constituent functions, this construction is not as constructive as with shallow networks. Besides, as with shallow networks, such a construction does not yield a uniform error bound analogous to 4.6.

In contrast, Theorem 3.3 can be extended in a purely constructive manner, so as to yield a good training error and to keep the gradient of the resulting approximation under control so that pointwise generalization error bounds can be obtained. Even though one does not know the constituent functions, one can construct a DAG trigonometric polynomial knowing the DAG structure and the training data alone.

For $\mathbf{n}=\left(n_{v}\right)$ and $T \in \mathcal{G}-\mathbb{H}_{\mathbf{n}}$, let

$$
\begin{equation*}
R_{\mathbf{n}, \mathcal{G}}(T)=\max _{1 \leq j \leq M}\left|y_{j}-T_{v^{*}}\left(\mathbf{x}_{j}\right)\right|+\sum_{v \in V} \frac{1}{n_{v}}\left\|T_{v}\right\|_{W^{*}\left(\mathbb{T}^{d(v)}\right)} \tag{4.7}
\end{equation*}
$$

In this definition, it is understood that $T_{v^{*}}$ is thought of as a function of the $q$-dimensional vector $\mathbf{x}$, but is computed using the all the constituent functions in $T$ using DAG structure prescribed by $\mathcal{G}$.

Theorem 4.2 Let $\mathcal{G}$ be a DAG with sink node $v^{*}$. Let $f=\left\{f_{v}\right\}_{v \in V \cup \mathbf{S}} \in \mathcal{G}-W^{*}$, and $\max _{v \in V \cup \mathbf{S}}\left\|f_{v}\right\|_{W^{*}\left(\mathbb{T}^{(d(v)}\right)} \leq L$. There exists $C(\mathcal{G})>0$ with the following property: If $N_{v} \geq C(\mathcal{G}) \eta\left(\mathcal{C}_{v}\right)^{-1}$, for each $v \in V$, and $\mathbf{N}=\left(N_{v}\right)$,

$$
\begin{equation*}
\min _{T \in \mathcal{G}-\mathbb{H}_{\mathbf{N}}} R_{\mathbf{N}, \mathcal{G}}(T) \leq c(\mathcal{G}, L)\left\{\epsilon+\sum_{v \in V} \frac{1}{N_{v}}\left\|f_{v}\right\|_{W^{*}\left(\mathbb{T}^{d(v)}\right)}\right\} . \quad \text { (Good training error) } \tag{4.8}
\end{equation*}
$$

Let $\mathbf{x} \in \mathbb{T}^{q}, \delta=\min _{1 \leq j \leq M}\left|\mathbf{x}-\mathbf{x}_{j}\right|, N=\max _{v \in V} N_{v}$. Then with

$$
T^{*}(\mathcal{D})=\underset{\mathcal{G}-\mathbb{H}_{\mathbf{N}}}{\arg \min } R_{\mathbf{N}, \mathcal{G}}(T)
$$

we have
$\left|f_{v^{*}}(\mathbf{x})-\left(T^{*}(\mathcal{D})\right)_{v^{*}}(\mathbf{x})\right| \leq c(\mathcal{G}, L)\left\{\epsilon+\sum_{v \in V} \frac{1}{N_{v}}\left\|f_{v}\right\|_{W^{*}\left(\mathbb{T}^{d(v)}\right)}\right\}+c(\mathcal{G}, L)(1+N \epsilon)^{H\left(v^{*}\right)} \delta . \quad$ (Good generalization error).

Remark 4.1 The theorems in this section indicate that the superiority of deep learning over shallow learning comes from two factors. One is that the compositional structure of the target function allows us to study the problem in a cascade of low dimensional problems. The other is that this structure might allow us to sparsify the training data as we move up the cascade; noting that the minimal separation is defined only among the distinct points in a set.

## 5 Proofs

### 5.1 Preliminary results on trigonometric approximation

We recall first some essential properties of the kernels and operators defined in 2.2 and 2.3 . Since the proofs of these facts are scattered among several of our publications, we will sketch the proof in some detail for the sake of completion.

Proposition 5.1 Let $S>q$ be an integer. For $N \geq 1$,

$$
\begin{equation*}
\left|\Phi_{N}(\mathbf{x})\right| \leq \frac{c N^{q}}{\max \left(1,(N|\mathbf{x}|)^{S}\right)}, \quad \mathbf{x} \in \mathbb{T}^{q} \tag{5.1}
\end{equation*}
$$

and

$$
\begin{equation*}
\left|\Phi_{N}(\mathbf{0})\right| \geq c N^{q} \tag{5.2}
\end{equation*}
$$

Proof. Our proof summarizes that of [3, Theorem 6.1]. We consider the function $H(\mathbf{t})=h\left(|\mathbf{t}|_{2}\right), \mathbf{t} \in \mathbb{R}^{q}$. Since $0 \leq H(\mathbf{t}) \leq 1$ for all $\mathbf{t} \in \mathbb{R}^{q}$, and $H(\mathbf{t})=1$ for $|\mathbf{t}|_{2} \leq 1 / 2$, it is clear that

$$
\begin{equation*}
\left|\Phi_{N}(\mathbf{x})\right| \leq \Phi_{N}(\mathbf{0}) \sim N^{q} \tag{5.3}
\end{equation*}
$$

In particular, this proves 5.2 . Since $h$ is constant in a neighborhood of $\mathbf{0}$, it is easy to see that $H$ is $S$ times continuously differentiable as well, so that its Fourier transform satisfies

$$
\begin{equation*}
|\hat{H}(\mathbf{u})| \leq \frac{c(H)}{|\mathbf{u}|^{S}}, \quad \mathbf{u} \neq \mathbf{0} \tag{5.4}
\end{equation*}
$$

Hence, the Poisson summation formula yields

$$
\begin{equation*}
\Phi_{N}(\mathbf{x})=N^{q} \sum_{\mathbf{j} \in \mathbb{Z}^{q}} \hat{H}(N(\mathbf{x}+2 \mathbf{j} \pi)), \quad \mathbf{x} \in \mathbb{T}^{q} \tag{5.5}
\end{equation*}
$$

For $\mathbf{x} \neq 0, \mathbf{j} \neq \mathbf{0}$,

$$
|\mathbf{x}+2 \mathbf{j} \pi| \geq 2|\mathbf{j}|_{\infty} \pi-|\mathbf{x}| \geq\left(2|\mathbf{j}|_{\infty}-1\right)|\mathbf{x}|
$$

Therefore, 5.4 and 5.5 show that for $\mathbf{x} \neq \mathbf{0}$,

$$
\left|\Phi_{N}(\mathbf{x})\right| \leq N^{q}\left\{|\hat{H}(N \mathbf{x})|+\sum_{\mathbf{j} \in \mathbb{Z}^{q}, \mathbf{j} \neq \mathbf{0}}|\hat{H}(N(\mathbf{x}+2 \mathbf{j} \pi))|\right\} \leq c(H) N^{q}\left\{\frac{1}{(N|\mathbf{x}|)^{S}}+\sum_{\mathbf{j} \in \mathbb{Z}^{q}, \mathbf{j} \neq \mathbf{0}} \frac{1}{((2|\mathbf{j}| \infty-1) N|\mathbf{x}|)^{S}}\right\}
$$

Since $S>q$, the infinite series converges. Therefore,

$$
\begin{equation*}
\left|\Phi_{N}(\mathbf{x})\right| \leq c \frac{N^{q}}{(N|\mathbf{x}|)^{S}}, \quad \mathbf{x} \neq \mathbf{0} \tag{5.6}
\end{equation*}
$$

Together with (5.3), this leads to (5.1).
Proposition 5.2 Let $S>q$ be an integer.
(a) There exists a constant $B>0$ with the following property: If $\mathcal{C}=\left\{\mathbf{x}_{1}, \cdots, \mathbf{x}_{M}\right\}$, and $N \geq B \eta(\mathcal{C})^{-1}$ then for $\mathbf{x} \in \mathbb{T}^{q}$,

$$
\begin{equation*}
\sum_{j=1}^{M}\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right)\right| \leq c N^{q} \tag{5.7}
\end{equation*}
$$

and

$$
\begin{equation*}
\sum_{j:\left|\mathbf{x}_{j}-\mathbf{x}\right| \geq \eta(\mathcal{C})}\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right)\right| \leq(1 / 2) \Phi_{N}(\mathbf{0})=(1 / 2) \Phi_{N}(\mathbf{x}-\mathbf{x}) \leq c N^{q} . \tag{5.8}
\end{equation*}
$$

(b) If $T \in \mathbb{H}_{N / 2}^{q}$ then $\sigma_{N}(T)=T$. Further,

$$
\begin{equation*}
\left\|\sigma_{N}(f)\right\| \leq c\|f\|, \quad f \in C^{*} \tag{5.9}
\end{equation*}
$$

Consequently,

$$
\begin{equation*}
E_{N}(f) \leq\left\|f-\sigma_{N}(f)\right\| \leq c E_{N / 2}(f), \quad f \in C^{*} . \tag{5.10}
\end{equation*}
$$

The two parts of the proposition are proved along the same lines. It is convenient to prove the following lemma (cf. [14. Proposition 5.1]), which we will use once with $\nu$ being the measure that associates the mass 1 with each $\mathbf{x}_{j}$ for part (a), and once with $\nu$ being the Lebesgue measure $\mu^{*}$ on $\mathbb{T}^{q}$ for part (b).
Lemma 5.1 Let $\nu$ be a positive measure on $\mathbb{T}^{q}, d \geq 0$, and for some $A_{\nu}>0$,

$$
\begin{equation*}
\nu(\mathbb{B}(\mathbf{x}, r)) \leq A_{\nu}(r+d)^{q}, \quad r>0, \mathbf{x} \in \mathbb{T}^{q} . \tag{5.11}
\end{equation*}
$$

Then for $r \geq 1 / N$,

$$
\begin{equation*}
\int_{\mathbb{T}^{q} \backslash \mathbb{B}(\mathbf{x}, r)}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \nu(\mathbf{u}) \leq c A_{\nu}(N r)^{q-S}(1+d / r)^{q}, \tag{5.12}
\end{equation*}
$$

where $c$ is a positive constant depending only on $q, S$, and $h$ but not on $r$, $d$, or $\nu$.
Proof. In this proof, we assume by re-normalization that $A_{\nu}=1$. Also, we write $\mathbb{A}_{k}=\mathbb{B}\left(\mathbf{x}, 2^{k+1} r\right) \backslash \mathbb{B}\left(\mathbf{x}, 2^{k} r\right)$. Then (5.11) shows that $\nu\left(\mathbb{A}_{k}\right) \leq\left(2^{k+1} r+d\right)^{q} \leq 2^{q} 2^{k q}(r+d)^{q}$. Using 5.6), we conclude that

$$
\int_{\mathbb{T} \backslash \backslash \mathbb{B}(\mathbf{x}, r)}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \nu(\mathbf{u}) \leq \sum_{k=0}^{\infty} \int_{\mathbb{A}_{k}}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \nu(\mathbf{u}) \leq c 2^{q}(N r)^{q-S}(1+d / r)^{q} \sum_{k=0}^{\infty} 2^{(q-S) k} .
$$

Proof of Proposition 5.2.
To prove part (a), let $\nu$ be a measure that associates the mass 1 with each $\mathbf{x}_{j}, j=1, \cdots, M$, and let $\eta=\eta(\mathcal{C})$. We claim that $\nu$ satisfies (5.11) with $d=\eta$ and $A_{\nu}=c \eta^{-q}$ for some $c>0$ depending only on $q$. Fix $\mathbf{x}$ and $r$. If $r<\eta / 2$, then $\mathbb{B}(\mathbf{x}, r) \cap \mathcal{C}$ can contain at most 1 point. Therefore the claim is satisfied trivially. Let $r \geq \eta / 2$, and $\mathbb{B}(\mathbf{x}, r) \cap \mathcal{C}=\left\{\mathbf{x}_{1}, \cdots, \mathbf{x}_{J}\right\}$. We observe that $\mu^{*}(\mathbb{B}(\mathbf{y}, s))=c s^{q}$ for all $\mathbf{y} \in \mathbb{T}^{q}$ and $s \in(0, \pi)$. Since the balls $\mathbb{B}\left(\mathbf{x}_{j}, \eta / 3\right)$ are disjoint, we have

$$
\nu(\mathbb{B}(\mathbf{x}, r))=|J|=c \eta^{-q} \sum_{j=1}^{J} \mu^{*}\left(\mathbb{B}\left(\mathbf{x}_{j}, \eta / 3\right)\right)=c \eta^{-q} \mu^{*}\left(\cup_{j=1}^{J} \mathbb{B}\left(\mathbf{x}_{j}, \eta / 3\right)\right) \leq c \eta^{-q} \mu^{*}(\mathbb{B}(\mathbf{x}, r+\eta / 3)) \leq c_{1} \eta^{-q}(r+\eta)^{q} .
$$

This proves the claim. Thus, we may use with $r=d=\eta, A_{\nu}=c \eta^{-q}$ to obtain for $N \geq \eta^{-1}$ that

$$
\sum_{j:\left|\mathbf{x}-\mathbf{x}_{j}\right| \geq \eta}\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right)\right|=\int_{\mathbb{T}^{q} \backslash \mathbb{B}(\mathbf{x}, \eta)}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \nu(\mathbf{u}) \leq c \eta^{-q}(N \eta)^{q-S}=c N^{q}(N \eta)^{-S} \leq c \Phi_{N}(\mathbf{0})(N \eta)^{-S} .
$$

We choose $B>0$ such that if $N \eta \geq B, c(N \eta)^{-S} \leq 1 / 2$. This proves 5.8). Further, since any ball $\mathbb{B}(\mathbf{x}, \eta)$ contains at most $2^{q}$ points of $\mathcal{C}$, 5.8) and (5.3) lead to (5.7).

To prove part (b), we use Lemma 5.1 with $\mu^{*}$ in place of $\nu$. Clearly, 5.11) is satisfied with $d=0$, and $A_{\mu^{*}}=c$. Therefore, using (5.3) and (5.12) (with $r=1 / N$ ) leads to

$$
\int_{\mathbb{T}^{q}}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \mu^{*}(\mathbf{u})=\int_{\mathbb{B}(\mathbf{x}, 1 / N)}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \mu^{*}(\mathbf{u})+\int_{\mathbb{T}^{q} \backslash \mathbb{B}(\mathbf{x}, 1 / N)}\left|\Phi_{N}(\mathbf{x}-\mathbf{u})\right| d \mu^{*}(\mathbf{u}) \leq c .
$$

It is now easy to deduce 5.9. If $T \in \mathbb{H}_{N / 2}^{q}$, then $\hat{T}(\mathbf{k})=0$ if $|\mathbf{k}|_{2} \geq N / 2$, while $h(t)=1$ if $|t| \leq 1 / 2$. It follows from the definition 2.3p that $\sigma_{N}(T)=T$. For any $f \in C^{*}$ and $T \in \mathbb{H}_{N / 2}^{q}, 5.9$ leads to

$$
E_{N}(f) \leq\left\|f-\sigma_{N}(f)\right\|=\left\|f-T-\sigma_{N}(f-T)\right\| \leq c\|f-T\| .
$$

This implies 5.10.

Proposition 5.3 Let $\mathcal{C}$, $B$, and $N$ be as in Proposition 5.2, $\Psi \in C^{*}$, and $\mathbf{B}$ be the matrix $\left[\Psi\left(\mathbf{x}_{j}-\mathbf{x}_{k}\right)\right]_{j, k=1}^{M}$. There exists $\alpha^{*}>0$ such that if

$$
\begin{equation*}
\left\|\Phi_{N}-\Psi\right\| \leq \alpha^{*} \tag{5.13}
\end{equation*}
$$

then the matrix $\mathbf{B}$ is invertible, and $\left\|\mathbf{B}^{-1}\right\|_{\ell \infty \rightarrow \ell^{\infty}} \leq c N^{-q}$. In particular, let $\mathbf{b} \in \mathbb{R}^{M}$. Then there exists (uniquely) $\mathbf{a} \in \mathbb{R}^{M}$ such that

$$
\begin{equation*}
\sum_{j=1}^{M} a_{j} \Psi\left(\mathbf{x}_{\ell}-\mathbf{x}_{j}\right)=b_{\ell}, \quad \ell=1, \cdots, M \tag{5.14}
\end{equation*}
$$

and

$$
\begin{equation*}
\max _{1 \leq j \leq M}\left|a_{j}\right| \leq c N^{-q} \max _{1 \leq \ell \leq M}\left|b_{\ell}\right| \tag{5.15}
\end{equation*}
$$

Proof. When $\Psi=\Phi_{N}$, the proposition follows from (5.2), (5.8), and standard facts from linear algebra, (cf. 14, Proposition 6.1]). In this proof, the matrix norm $\|\cdot\|$ will refer to the norm $\|\cdot\|_{\ell \infty \rightarrow \ell^{\infty}}$, and we write $\eta=\eta(\mathcal{C})$. Denoting the matrix $\left[\Phi_{N}\left(\mathbf{x}_{j}-\mathbf{x}_{k}\right)\right]_{j, k=1}^{M}$ by $\mathbf{A}$, we have observed that $\left\|\mathbf{A}^{-1}\right\| \leq c N^{-q}$. Also, it is easy to see that $M \leq c \eta^{-q}$. Therefore, recalling that $N \eta \geq B$,

$$
\begin{aligned}
\|\mathbf{A}-\mathbf{B}\| & =\max _{j} \sum_{k=1}^{M}\left|\Phi_{N}\left(\mathbf{x}_{j}-\mathbf{x}_{k}\right)-\Psi\left(\mathbf{x}_{j}-\mathbf{x}_{k}\right)\right| \leq M\left\|\Phi_{N}-\Psi\right\| \leq c \eta^{-q}\left\|\Phi_{N}-\Psi\right\| \\
& \leq c B^{-q} N^{q}\left\|\Phi_{N}-\Psi\right\| \leq \frac{c B^{-q}}{\left\|\mathbf{A}^{-1}\right\|}\left\|\Phi_{N}-\Psi\right\|
\end{aligned}
$$

We now choose $\alpha^{*}$ so that 5.13 implies

$$
\|\mathbf{A}-\mathbf{B}\| \leq \frac{1}{2\left\|\mathbf{A}^{-1}\right\|}
$$

A perturbation theorem from linear algebra [7. Theorem 2.3.4] then shows that

$$
\left\|\mathbf{A}^{-1}-\mathbf{B}^{-1}\right\| \leq\left\|\mathbf{A}^{-1}\right\| \leq c N^{-q}
$$

This shows that $\left\|\mathbf{B}^{-1}\right\| \leq c N^{-q}$.
We also need some results about approximation of a function and its derivatives. We recall a theorem from [5, Theorem $1^{\circ}$. Per Section 2.1, the notation $\|\circ\|_{1}$ indicates the univariate supremum norm.

Theorem 5.1 Let $q=1, f \in W^{*}(\mathbb{T}), n \geq 1$ be an integer, $E>0$, and $T \in \mathbb{H}_{n}^{1}$ satisfy

$$
\begin{equation*}
\|f-T\|_{1} \leq E \tag{5.16}
\end{equation*}
$$

Then

$$
\begin{equation*}
\left\|f^{\prime}-T^{\prime}\right\|_{1} \leq c\left\{n E+E_{n}\left(f^{\prime}\right)\right\} \tag{5.17}
\end{equation*}
$$

In the multivariate case, we take the derivatives one variable at a time to deduce the following corollary of Theorem 5.1.

Corollary 5.1 Let $f \in W^{*}, N \geq 1$ be an integer, $E>0$, and $T \in \mathbb{H}_{N}^{q}$ satisfy

$$
\begin{equation*}
\|f-T\| \leq E \tag{5.18}
\end{equation*}
$$

Then

$$
\begin{equation*}
\|f-T\|_{W^{*}} \leq c\left\{N E+\sum_{j=1}^{q} E_{N}\left(D_{j} f\right)\right\} \tag{5.19}
\end{equation*}
$$

In particular,

$$
\begin{equation*}
\|T\|_{W^{*}} \leq c N\left\{E+\frac{1}{N}\|f\|_{W^{*}}\right\} \tag{5.20}
\end{equation*}
$$

We note also the direct theorem of trigonometric approximation [28, Section 5.3].
Proposition 5.4 If $f \in W^{*}$ then

$$
\begin{equation*}
E_{N}(f) \leq \frac{c}{N}\|f\|_{W^{*}} \tag{5.21}
\end{equation*}
$$

### 5.2 Proof of Proposition 2.1.

Using (2.8) and Schwarz inequality, we obtain that

$$
\begin{align*}
\left\|\sigma_{n}(f)-\mathbb{G}_{N}\left(\phi, \sigma_{n}(f)\right)\right\| & \leq c(\phi) E_{N}(1 ; \phi) \sum_{\mathbf{k} \in \mathbb{Z}^{q}}\left|h\left(\frac{|\mathbf{k}|_{2}}{n}\right) \hat{f}(\mathbf{k})\right| \\
& \leq c(\phi) E_{N}(1 ; \phi)\left\{\sum_{\mathbf{k} \in \mathbb{Z}^{q}}\left(h\left(\frac{|\mathbf{k}|_{2}}{n}\right)\right)^{2}\right\}^{1 / 2}\left\{\sum_{\mathbf{k} \in \mathbb{Z}^{q}}|\hat{f}(\mathbf{k})|^{2}\right\}^{1 / 2} \tag{5.22}
\end{align*}
$$

Since $h(t)=0$ if $t \geq 1$, and $0 \leq h(t) \leq 1$ for all $t$,

$$
\sum_{\mathbf{k} \in \mathbb{Z}^{q}}\left(h\left(\frac{|\mathbf{k}|_{2}}{n}\right)\right)^{2} \leq c n^{q}
$$

In view of Bessel inequality, 5.22 now implies

$$
\left\|\sigma_{n}(f)-\mathbb{G}_{N}\left(\phi, \sigma_{n}(f)\right)\right\| \leq c(\phi) E_{N}(1 ; \phi) n^{q / 2}\left(\int_{\mathbb{T}^{q}}|f(\mathbf{x})|^{2} d \mathbf{x}\right)^{1 / 2} \leq c(\phi) E_{N}(1 ; \phi) n^{q / 2}\|f\|
$$

Together with (5.10), this leads to (2.9).

### 5.3 Proof of the theorems in Section 3.

## Proof of Theorem 3.1.

In this proof, let $B$ be as in Proposition 5.2(a). $N \geq B \eta(\mathcal{C})^{-1}$, and for $\ell=1, \cdots, M, z_{\ell}=y_{\ell}-\sigma_{N}(f)\left(\mathbf{x}_{\ell}\right)$. Proposition 5.3 then guarantees that there exist $a_{j} \in \mathbb{R}$ such that

$$
\begin{equation*}
\sum_{j=1}^{M} a_{j} \Phi_{N}\left(\mathbf{x}_{\ell}-\mathbf{x}_{j}\right)=z_{\ell}=y_{\ell}-\sigma_{N}(f)\left(\mathbf{x}_{\ell}\right), \quad \ell=1, \cdots, M \tag{5.23}
\end{equation*}
$$

and (cf. 5.15) and 5.10)

$$
\begin{align*}
\max _{1 \leq j \leq M}\left|a_{j}\right| & \leq c N^{-q} \max _{1 \leq \ell \leq M}\left|z_{\ell}\right|=c N^{-q} \max _{1 \leq \ell \leq M}\left|y_{\ell}-\sigma_{N}(f)\left(\mathbf{x}_{\ell}\right)\right| \\
& \leq c N^{-q}\left\{\max _{1 \leq \ell \leq M}\left|y_{\ell}-f\left(\mathbf{x}_{\ell}\right)\right|+\max _{1 \leq \ell \leq M}\left|f\left(\mathbf{x}_{\ell}\right)-\sigma_{N}(f)\left(\mathbf{x}_{\ell}\right)\right|\right\} \\
& \leq c N^{-q}\left\{\epsilon+E_{N / 2}(f)\right\} \tag{5.24}
\end{align*}
$$

We now recall that

$$
\begin{equation*}
\mathcal{T}_{N}^{\#}(\mathcal{D})(\mathbf{x})=\sigma_{N}(f)(\mathbf{x})+\sum_{j=1}^{M} a_{j} \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right), \quad \mathbf{x} \in \mathbb{T}^{q} \tag{5.25}
\end{equation*}
$$

Clearly, $\mathcal{T}_{N}^{\#}(\mathcal{D}) \in \mathbb{H}_{N}^{q}$, and $\mathcal{T}_{N}^{\#}(\mathcal{D})\left(\mathbf{x}_{\ell}\right)=y_{\ell}, \ell=1, \cdots, M$. This proves 3.5.
Moreover, for every $\mathrm{x} \in \mathbb{T}^{q}$, 5.25 and (5.10 lead to

$$
\begin{align*}
\left|f(\mathbf{x})-\mathcal{T}_{N}^{\#}(\mathcal{D})(\mathbf{x})\right| & \leq\left|f(\mathbf{x})-\sigma_{N}(f)(\mathbf{x})\right|+\left\{\max _{1 \leq j \leq M}\left|a_{j}\right|\right\} \sum_{j=1}^{M}\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right)\right| \\
& \leq c\left\{E_{N / 2}(f)+c N^{-q}\left\{\epsilon+E_{N / 2}(f)\right\} \sum_{j=1}^{M}\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{j}\right)\right|\right\} \tag{5.26}
\end{align*}
$$

In view of (5.7), this leads to (3.6).

## Proof of Theorem 3.2.

This proof is very similar to that of Theorem 3.1. The condition 3.9 and Proposition 5.3 ensure that the system of equations (3.7) has a unique solution satisfying (5.15). The same argument as in (5.26) then works with the operators as in Theorem 3.2 replacing those in Theorem 3.1 we use $(2.9)$ in place of $(5.10)$.

## Proof of Theorem 3.3.

In this proof, let $\mathcal{T}^{\#}=T_{N}^{\#}(\mathcal{D})$ be as in Theorem 3.1. and

$$
\begin{equation*}
E=\epsilon+E_{N / 2}(f) \tag{5.27}
\end{equation*}
$$

In view of (3.6), we may use Corollary 5.1 to deduce that

$$
\begin{equation*}
\left\|f-\mathcal{T}^{\#}\right\|_{W^{*}} \leq c\left\{N E+\sum_{j=1}^{q} E_{N}\left(D_{j} f\right)\right\} \leq c N\left\{E+\frac{1}{N}\|f\|_{W^{*}}\right\} \tag{5.28}
\end{equation*}
$$

and in particular,

$$
\begin{equation*}
\left\|\mathcal{T}^{\#}\right\|_{W^{*}} \leq c N\left\{E+\frac{1}{N}\|f\|_{W^{*}}\right\} \tag{5.29}
\end{equation*}
$$

Using (3.5), 5.27) and 5.29, we obtain:

$$
\begin{align*}
\min _{T \in \mathbb{H}_{N}} R_{N}(T) & \leq R_{N}\left(\mathcal{T}^{\#}\right) \leq \max _{1 \leq j \leq M}\left|y_{j}-\mathcal{T}^{\#}\left(\mathbf{x}_{j}\right)\right|+\frac{1}{N}\left\|\mathcal{T}^{\#}\right\|_{W^{*}} \\
& =\frac{1}{N}\left\|\mathcal{T}^{\#}\right\|_{W^{*}} \leq c\left\{E+\frac{1}{N}\|f\|_{W^{*}}\right\} \tag{5.30}
\end{align*}
$$

In view of Proposition 5.4,

$$
E \leq c\left\{\epsilon+\frac{1}{N}\|f\|_{W^{*}}\right\}
$$

Together with (5.30), this proves (3.14).
Next, let $\mathbf{x} \in \mathbb{T}^{q}$, and $\delta=\min _{1 \leq j \leq M}\left|\mathbf{x}-\mathbf{x}_{j}\right|=\left|\mathbf{x}-\mathbf{x}_{\ell}\right|$. For brevity, we write

$$
\tilde{E}=\epsilon+\frac{1}{N}\|f\|_{W^{*}}
$$

Using (3.14) and Corollary 5.1. we deduce that

$$
\begin{align*}
\left|f(\mathbf{x})-T^{*}(\mathbf{x})\right| & \leq\left|f(\mathbf{x})-f\left(\mathbf{x}_{\ell}\right)\right|+\left|f\left(\mathbf{x}_{\ell}\right)-T^{*}\left(\mathbf{x}_{\ell}\right)\right|+\left|T^{*}\left(\mathbf{x}_{\ell}\right)-T^{*}(\mathbf{x})\right| \\
& \leq\left|\mathbf{x}-\mathbf{x}_{\ell}\right|\|f\|_{W^{*}}+c \tilde{E}+\left|\mathbf{x}-\mathbf{x}_{\ell}\right|\left\|T^{*}\right\|_{W^{*}} \leq N \delta \frac{1}{N}\|f\|_{W^{*}}+c \tilde{E}+N \delta \tilde{E} \\
& \leq c(1+N \delta) \tilde{E} \tag{5.31}
\end{align*}
$$

This proves 3.15.
Remark 5.1 If $\delta=\min _{1 \leq j \leq M}\left|\mathbf{x}-\mathbf{x}_{j}\right|=\left|\mathbf{x}-\mathbf{x}_{\ell}\right|$, then we have

$$
\begin{aligned}
& \left|f(\mathbf{x})-\frac{1}{\Phi_{N}(0)} \sum_{k=1}^{M} y_{k} \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right)\right| \\
& \quad \leq\left|f(\mathbf{x})-f\left(\mathbf{x}_{\ell}\right)\right|+\left|\frac{1}{\Phi_{N}(0)} \sum_{k=1}^{M}\left(y_{k}-f\left(\mathbf{x}_{k}\right)\right) \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right)\right|+\left|f\left(\mathbf{x}_{\ell}\right)-\frac{1}{\Phi_{N}(0)} \sum_{k=1}^{M} f\left(\mathbf{x}_{k}\right) \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right)\right| \\
& \quad \leq \delta\|f\|_{W^{*}}+\epsilon+\frac{1}{\Phi_{N}(0)}\left|f\left(\mathbf{x}_{\ell}\right)\left(\Phi_{N}\left(\mathbf{x}_{\ell}-\mathbf{x}_{\ell}\right)-\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{\ell}\right)\right)\right|+\frac{1}{\Phi_{N}(0)} \sum_{k \neq \ell}\left|f\left(\mathbf{x}_{k}\right)\right|\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right)\right| \\
& \quad \leq \delta\|f\|_{W^{*}}+c \epsilon+c N \delta\|f\|+\frac{1}{\Phi_{N}(0)} \sum_{k \neq \ell}\left|f\left(\mathbf{x}_{k}\right)\right|\left|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right)\right| .
\end{aligned}
$$

Hence, if $2 \pi \geq \delta \geq 1 / N$, Lemma 5.1 used with $\nu$ as in the proof of Proposition 5.2 (a), $d=\eta(\mathcal{C}), r=\delta$ shows that

$$
\frac{1}{\Phi_{N}(0)} \sum_{k \neq \ell}\left|f\left(\mathbf{x}_{k}\right)\left\|\Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right) \mid \leq c(N \delta)^{-S}(\delta+\eta(\mathcal{C}))^{q}\right\| f\left\|\leq c(N \delta)^{-S}\right\| f \|\right.
$$

Therefore, if $1 / N \leq \delta \leq 2 \pi$,

$$
\begin{equation*}
\left|f(\mathbf{x})-\frac{1}{\Phi_{N}(0)} \sum_{k=1}^{M} y_{k} \Phi_{N}\left(\mathbf{x}-\mathbf{x}_{k}\right)\right| \leq c\left\{\delta\|f\|_{W^{*}}+\epsilon+c N \delta\|f\|\right\} \tag{5.32}
\end{equation*}
$$

A similar bound can be proved with the polynomial $\mathcal{L}_{N}(\mathcal{D})$.

### 5.4 Proofs of the theorems in Section 4 .

The induction argument given in the proof of Theorem 4.1 allows us to "lift" results about shallow networks to deep networks. We will refer to this argument as good propagation of error.

## Proof of Theorem 4.1.

We observe that each node in $V \cup \mathbf{S}$ can be thought of as the sink node of an appropriate sub-DAG $\mathcal{G}_{v}$ of $\mathcal{G}$. Moreover, in the definitions such as (4.3), $\|f\|_{\mathbb{X}, \mathcal{G}} \sim \sum_{v \in V \cup \mathbf{S}}\left\|f_{v}\right\|_{\mathbb{X}, \mathcal{G}_{v}}$, and similarly for (4.4). So, the statement of this theorem should be true also for each such sub-DAG. Let $C=\max _{v \in V \cup \mathbf{S}} B(d(v))$ where $B(d(v))$ is the constant introduced in Proposition 5.2 applied with $d(v)$ in place of $q$.

Theorem 4.1 applied to a vertex in $\mathbf{S}$ is the same as Theorem 3.1 .
Suppose the theorem is proved for every node of level up to $\ell$ for some $\ell \geq 0, v$ be a node at level $\ell+1$, and $u_{1}, \cdots, u_{d(v)}$ be its children. The sub-DAG with the outputs of these children as the input to the sink node $v$ is of course a shallow network, to which we may apply Theorem 3.1. Since $N_{v} \geq C \eta\left(\mathcal{C}_{v}\right)^{-1}$, Theorem 3.1 implies that there is $T_{v}^{\#} \in \mathbb{H}_{N_{v}}^{d(v)}$ satisfying ${ }^{1}$

$$
\begin{equation*}
T_{v}^{\#}(\mathbf{x})=f_{v}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{C}_{v}, \quad\left\|f_{v}-T_{v}^{\#}\right\|_{d(v)} \leq c E_{N_{v} / 2}\left(d(v) ; f_{v}\right) \tag{5.35}
\end{equation*}
$$

Our induction hypothesis shows that the analogues of the estimates (5.35) and (5.33) are true for each of the children $u_{1}, \cdots, u_{d(v)}$. Necessarily, $v \in V$, and so, each $f_{v}$ is Lipschitz continuous with Lipschitz constant $\leq L$. Using triangle inequality, we deduce that

$$
\begin{aligned}
& \left|f_{v}\left(f_{u_{1}}\left(\mathbf{x}_{u_{1}}\right), \cdots, f_{u_{d(v)}}\left(\mathbf{x}_{u_{d(v)}}\right)\right)-T_{v}^{\#}\left(T_{u_{1}}^{\#}\left(\mathbf{x}_{u_{1}}\right), \cdots, T_{u_{d(v)}}^{\#}\left(\mathbf{x}_{u_{d(v)}}\right)\right)\right| \\
& \leq \quad\left|f_{v}\left(T_{u_{1}}^{\#}\left(\mathbf{x}_{u_{1}}\right), \cdots, T_{u_{d(v)}}^{\#}\left(\mathbf{x}_{u_{d(v)}}\right)\right)-T_{v}^{\#}\left(T_{u_{1}}^{\#}\left(\mathbf{x}_{u_{1}}\right), \cdots, T_{u_{d(v)}}^{\#}\left(\mathbf{x}_{u_{d(v)}}\right)\right)\right| \\
& \quad \quad+\left|f_{v}\left(f_{u_{1}}\left(\mathbf{x}_{u_{1}}\right), \cdots, f_{u_{d(v)}}\left(\mathbf{x}_{u_{d(v)}}\right)\right)-f_{v}\left(T_{u_{1}}^{\#}\left(\mathbf{x}_{u_{1}}\right), \cdots, T_{u_{d(v)}}^{\#}\left(\mathbf{x}_{u_{d(v)}}\right)\right)\right| \\
& \quad \leq \quad \sup _{\mathbf{y}_{\in \in \mathbb{T}^{d(v)}}}\left|f_{v}(\mathbf{y})-T_{v}^{\#}(\mathbf{y})\right|+L\left\|\left(f_{u_{1}}\left(\mathbf{x}_{u_{1}}\right), \cdots, f_{u_{d(v)}}\left(\mathbf{x}_{u_{d(v)}}\right)\right)-\left(T_{u_{1}}^{\#}\left(\mathbf{x}_{u_{1}}\right), \cdots, T_{u_{d(v)}}^{\#}\left(\mathbf{x}_{u_{d(v)}}\right)\right)\right\|_{1} \\
& \leq\left\|f_{v}-T_{v}^{\#}\right\|_{d(v)}+L \sum_{j=1}^{d(v)}\left\|f_{u_{j}}-T_{u_{j}}^{\#}\right\|_{C^{*}, \mathcal{G}_{u_{j}}} .
\end{aligned}
$$

In view of our dual interpretation of the constituent functions as functions of their immediate input as well as the input seen by these functions and the induction hypothesis, this shows that Theorem 4.1 is valid for the sub-DAG with $v$ as the sink node. We define

$$
\begin{equation*}
T_{\mathbf{N}}^{\#}(\mathcal{D})=\left\{T_{v}^{\#}\right\} \in \mathcal{G}-\mathbb{H}_{\mathbf{N}} \tag{5.36}
\end{equation*}
$$

[^1]The equation 5.33 is the same as 4.5 . The estimate 4.6 follows by induction as just explained.

## Proof of Theorem 4.2.

This proof is essentially the same as that of Theorem 3.3. We point out some considerations required in the details which are different. Since each $f_{v} \in W^{*}\left(\mathbb{T}^{d(v)}\right)$, the simultaneous approximation theorem Corollary 5.1 holds for each $f_{v}$. We may assume that the perturbation exists only at the sink node, and the rest of the data is exact. Finally, we use the good propagation of error to obtain 4.8. The proof of 4.9) involves an argument similar to 5.31. The middle term on the first line of that chain of inequalities gives the first term on the right hand side in 4.9). The other terms can be estimated as in 5.31), except that the chain rule of differentiation needs to be used several times to get to the level of the source nodes. This is facilitated by our extra assumption that $\max _{v \in V}\left\|f_{v}\right\|_{W^{*}\left(\mathbb{T}^{(d(v)}\right)} \leq L$.

## 6 Conclusions and open problems

We have explored the puzzle that deep networks (and sometimes also shallow ones) do not exhibit over-fitting even though the number of parameters is very large and the training error is reduced to zero. We have initiated a rigorous study of this phenomenon from the point of view of function approximation, giving estimates on how many parameters are needed to exhibit a zero or good training error, which is also compatible with the generalization error. Our estimates are given in terms of the data characteristics and the smoothness of the target function.

One obvious problem is to reduce the number of parameters to be trained in the regularization functional (3.13). Expressing the trigonometric polynomials as linear combinations of the translates $\Phi_{N}(\circ-2 \pi \mathbf{j} / N)$ ensures that the resulting solution will have coefficients that are small away from the training data. It is therefore reasonable to take only some of these translates that are close to the training data. However, it is not clear that the theory will work with the space defined by the span of only those translates which are kept. Also, there may be some numerical instability problems with this basis.

A deeper and wider area of theoretical investigation is the following. Theorems 3.3 and 4.2 suggest that the approximation error given by the trigonometric polynomials constructed there gives an estimate on the support of the marginal distribution of the training and test data. In [4], the approximation errors of the convergent bounded interpolatory polynomials constructed in [3] were used for texture detection and segmentation of images. What would be the analogues for understanding the nature of the data using the approximation errors obtained here?

## References

[1] M. Belkin, D. J. Hsu, and P. Mitra. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In Advances in Neural Information Processing Systems, pages 2300-2311, 2018.
[2] M. Belkin, S. Ma, and S. Mandal. To understand deep learning we need to understand kernel learning. In Proceedings of the 35th International Conference on Machine Learning (PMLR), pages 80:541-549, 2018.
[3] S. Chandrasekaran, K. R. Jayaraman, and H. N. Mhaskar. Minimum Sobolev norm interpolation with trigonometric polynomials on the torus. Journal of Computational Physics, 249:96-112, 2013.
[4] S. Chandrasekaran, K. R. Jayaraman, J. Moffitt, H. N. Mhaskar, and S. Pauli. Minimum Sobolev norm schemes and applications in image processing. In $I S \mathcal{G} T / S P I E$ Electronic Imaging, pages 753507-753507. International Society for Optics and Photonics, 2010.
[5] J. Czipszer and G. Freud. Sur l'approximation d'une fonction périodique et de ses dérivées successires par un polynome trigonométrique et par ses dérivées successives. Acta Math.,(Sweden), 99:33-51, 1958.
[6] P. Erdős. On some convergence properties of the interpolation polynomials. Annals of Mathematics, pages 330-337, 1943.
[7] G. H. Golub and C. F. Van Loan. Matrix computations, volume 3. JHU Press, 2012.
[8] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio. Deep learning, volume 1. MIT press Cambridge, 2016.
[9] Y. Han, L. Sunwoo, and J. C. Ye. k-space deep learning for accelerated MRI. IEEE transactions on medical imaging, 2019.
[10] M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. In ICML'16 Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, pages 1225-1234, 2016.
[11] Q. T. Le Gia and H. N. Mhaskar. Localized linear polynomial operators and quadrature formulas on the sphere. SIAM Journal on Numerical Analysis, 47(1):440-466, 2009.
[12] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. Nature, 521(7553):436-444, 2015.
[13] H. N. Mhaskar. Approximation theory and neural networks. In Wavelet Analysis and Applications, Proceedings of the international workshop in Delhi, pages 247-289, 1999.
[14] H. N. Mhaskar. Eignets for function approximation on manifolds. Applied and Computational Harmonic Analysis, 29(1):63-87, 2010.
[15] H. N. Mhaskar. Function approximation with zonal function networks with activation functions analogous to the rectified linear unit functions. Journal of Complexity, 51:1-19, April 2019.
[16] H. N. Mhaskar and C. A. Micchelli. Degree of approximation by neural and translation networks with a single hidden layer. Advances in Applied Mathematics, 16(2):151-183, 1995.
[17] H. N. Mhaskar, F. J. Narcowich, N. Sivakumar, and J. D. Ward. Approximation with interpolatory constraints. Proceedings of the American Mathematical Society, 130(5):1355-1364, 2002.
[18] H. N. Mhaskar, P. Nevai, and E. Shvarts. Applications of classical approximation theory to periodic basis function networks and computational harmonic analysis. Bulletin of Mathematical Sciences, 3(3):485-549, 2013.
[19] H. N. Mhaskar, S. V. Pereverzyev, and M. D. van der Walt. A deep learning approach to diabetic blood glucose prediction. Frontiers in Applied Mathematics and Statistics, 3:14, 2017.
[20] H. N. Mhaskar and T. Poggio. Deep vs. shallow networks: An approximation theory perspective. Analysis and Applications, 14(06):829-848, 2016.
[21] I. P. Natanson. Constructive function theory. Ungar, 1964.
[22] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro. Exploring generalization in deep learning. In Advances in Neural Information Processing Systems, pages 5949-5958, 2017.
[23] T. Poggio, K. Kawaguchi, Q. Liao, B. Miranda, L. Rosasco, X. Boix, J. Hidary, and H. Mhaskar. Theory of deep learning III: explaining the non-overfitting puzzle. arXiv preprint arXiv:1801.00173, 2017.
[24] T. Poggio, Q. Liao, B. Miranda, A. Banburski, X. Boix, and J. Hidary. Theory IIIb: Generalization in deep networks. arXiv preprint arXiv:1806.11379, 2018.
[25] J. Sokolić, R. Giryes, G. Sapiro, and M. R. Rodrigues. Robust large margin deep neural networks. IEEE Transactions on Signal Processing, 65(16):4265-4280, 2016.
[26] E. M. Stein. Singular integrals and differentiability properties of functions (PMS-30), volume 30. Princeton university press, 2016.
[27] J. Szabados. On some convergent interpolatory polynomials. Fourier Analysis and Approximation Theory, Coll. Math. Soc. János Bolyai, 19:805-815, 1978.
[28] A. F. Timan. Theory of Approximation of Functions of a Real Variable: International Series of Monographs on Pure and Applied Mathematics, volume 34. Elsevier, 2014.
[29] C. Zhang, Q. Liao, A. Rakhlin, B. Miranda, N. Golowich, and T. Poggio. Musings on deep learning: Properties of SGD. CBMM Memo, 67, 2017.
[30] A. Zygmund. Trigonometric series, volume 1. Cambridge University Press, 2002.


[^0]:    *Institute of Mathematical Sciences, Claremont Graduate University, Claremont, CA 91711. The research of this author is supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2018-18032000002. email: hrushikesh.mhaskar@cgu.edu
    ${ }^{\dagger}$ Center for Brains, Minds, and Machines, McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, 02139. The research of this author is supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216. tp@mit.edu

[^1]:    ${ }^{1}$ Here, if $v=v^{*}$ is the sink node of the original DAG, the statement needs to be modified as follows: there is $T_{v^{*}}^{\#} \in \mathbb{H}_{N}^{d\left(v^{*}\right)}$ that satisfies (thought of as a function of all the inputs to the network)

    $$
    \begin{equation*}
    T_{v^{*}}^{\#}\left(\mathbf{x}_{j}\right)=T_{v^{*}}^{\#}\left(\left(\mathbf{x}_{j}\right)_{v^{*}}\right)=y_{j}, \quad j=1, \cdots, M, \tag{5.33}
    \end{equation*}
    $$

    and (thought of as a function on $\mathbb{T}^{d\left(v^{*}\right)}$ )

    $$
    \begin{equation*}
    \left\|f_{v^{*}}-T_{v^{*}}^{\#}\right\|_{d\left(v^{*}\right)} \leq c\left\{\epsilon+E_{N_{v^{*}} / 2}\left(d\left(v^{*}\right) ; f_{v^{*}}\right)\right\} \tag{5.34}
    \end{equation*}
    $$

