

Complexity bounds for approximations with deep ReLU neural networks in Sobolev norms

Ingo Gühring*

Gitta Kutyniok[†]

Philipp Petersen[‡]

April 8, 2019

Abstract

We analyze to what extent deep ReLU neural networks can efficiently approximate Sobolev regular functions if the approximation error is measured with respect to weaker Sobolev norms. In this context, we first establish upper approximation bounds by ReLU neural networks for Sobolev regular functions by explicitly constructing the approximating ReLU neural networks. Then, we establish lower approximation bounds for the same type of function classes. A trade-off between the regularity used in the approximation norm and the complexity of the neural network can be observed in upper and lower bounds. Our results extend recent advances in the approximation theory of ReLU networks to the regime that is most relevant for applications in the numerical analysis of partial differential equations.

1 Introduction

Powered by modern, highly parallelized hardware and the immense amount of accessible data, deep neural networks often substantially outperform both traditional modeling approaches, e.g. those based on differential equations, and classical machine learning methods in a wide range of applications. Prominent areas of application include image classification [34, 35, 54], speech recognition [17, 32, 59], and natural language processing [62].

Despite this overwhelming success in applications, a comprehensive mathematical explanation of this success has not yet been found. However, a deep theoretical understanding of these techniques is crucial to design more efficient architectures and essential in safety-critical applications, such as autonomous driving or medical diagnostics. One main concern is the understanding of the generalization error, which measures to which extent a neural network trained on a training data set is able to also correctly predict values for previously unseen data. From the point of view of statistical learning theory, the task of estimating the generalization error decomposes into a statistical problem depending on the samples and an approximation theoretical problem, independent of the samples (see [14]). The second question is typically coined *expressivity*, and aims at analyzing the approximation power of deep neural networks.

One of the applications of deep learning where a strong knowledge of the approximation capabilities of neural networks directly translates into quantifiable theoretical results, appears when solving *partial differential equations (PDEs)* using deep learning techniques. Some notable advances in this direction have been made in [6, 19, 20, 21, 28, 36, 37, 55]. In order to derive error estimates in appropriate norms, not only the approximation fidelity with respect to standard Lebesgue spaces is of high interest, but also that with respect to Sobolev-type norms. First results in this direction were reported in [52].

In this work, we analyze the expressive power of deep neural networks with respect to Sobolev-regular functions. In particular, we give a comprehensive analysis of the approximation rates of deep neural networks for Sobolev-regular functions with respect to (fractional) Sobolev norms.

*Institut für Mathematik, Technische Universität Berlin, E-mail: guehring@math.tu-berlin.de

[†]Institut für Mathematik, Technische Universität Berlin, E-mail: kutyniok@math.tu-berlin.de

[‡]Mathematical Institute, University of Oxford, E-mail: Philipp.Petersen@maths.ox.ac.uk

1.1 A mathematical view of deep neural networks

To provide a first insight into our results, we start by recalling the mathematical framework of deep neural networks. There exists a wide variety of neural network architectures, each adapted to specific tasks. One of the most commonly-used architectures is a *feedforward* architecture also known as *multi-layer perceptron* which implements a function as a sequence of affine-linear transformations followed by a componentwise application of a non-linear function, called *activation function*. The length L of the sequence is referred to as the number of *layers* of the network and the input of a layer consists of the output of the preceding layer. In the sequel, we denote by $N_0 \in \mathbb{N}$ the dimension of the input space and by $N_l \in \mathbb{N}$ the dimension of the output of the l -th layer for $l = 1, \dots, L$. Thus, if $\varrho : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function and $T_l : \mathbb{R}^{N_{l-1}} \rightarrow \mathbb{R}^{N_l}$ is the affine-linear transformation of the l -th layer, then the computation in that layer can be described as

$$f_l : \mathbb{R}^{N_{l-1}} \rightarrow \mathbb{R}^{N_l}, x \mapsto \varrho(T_l(x)) \quad \text{for } l = 1, \dots, L-1 \quad \text{and} \quad f_L : \mathbb{R}^{N_{L-1}} \rightarrow \mathbb{R}^{N_L}, x \mapsto T_L(x). \quad (1.1)$$

Note that ϱ acts componentwise in all but the last layer and in the last layer no activation function is applied.

The parameters defining the affine-linear transformations T_l are referred to as *weights* and denoted by w . Thus, given a set of weights w , a neural network \mathcal{N} implements a function $\mathcal{N}(\cdot|w) : \mathbb{R}^{N_0} \rightarrow \mathbb{R}^{N_L}$ defined by the scheme in (1.1). The sum $\sum_{k=0}^L N_k$ is called the number of *neurons* of the network since each output coordinate of f_l can be seen as a small computational unit, similar to neurons in the brain. If a network has three or more layers, then it is usually called *deep* and a two-layer network is called *shallow*. The *complexity of a neural network* is typically measured in the number of layers, weights, and neurons (see [3]). The term *deep learning* refers to the subset of machine learning methods associated with deep neural networks.

Recently, more general feedforward architectures which allow connections between non-neighboring layers, so-called *skip connections*, have been shown to yield state-of-the-art results in object recognition tasks [30, 34]. Here, the input of a layer consists of the output of *all* preceding layers. Note, that the case where no skip connections are allowed is a special case of this more general architecture.

One of today's most widely-used activation functions is the *Rectified Linear Unit (ReLU)* (see [38]), defined as

$$\varrho : \mathbb{R} \rightarrow \mathbb{R}, \quad x \mapsto \max\{0, x\}.$$

The popularity of the ReLU can be explained by a number of factors. It is cheap to compute, promotes sparsity in data representation [7], alleviates the problem of vanishing-gradients [24], and thus yields better optimization properties. Moreover, the ReLU is the only commonly-used activation function such that the associated function spaces are not necessary non-closed [47]. In this work, we will mostly focus on ReLU networks with a feedforward architecture allowing skip connections.

1.2 Expressivity of neural networks in classical norms

Many attempts at unraveling the extreme efficiency of deep neural networks have been made in the context of approximation theory. The universal approximation theorem [15, 33] establishes that every continuous function on a compact domain can be uniformly approximated by shallow neural networks. More refined results that also set into relation the size and the approximation fidelity of neural networks have been reported for smooth activation functions in, for example, [5, 9, 42, 53].

As previously mentioned, in most applications, the activation functions are not smooth but taken to be the piecewise linear ReLU function. For these networks, approximation bounds for classes of smooth functions have been established in [60] and for piecewise smooth functions in [48]. Connections to sparse grids and the associated approximation rates were established in [44] and connections to linear finite element approximation were reported in [29]. It was also discovered that the depth of neural networks, i.e., the number of layers, crucially influences the approximation capabilities of these networks in the sense that deeper networks are more efficient approximators [13, 43, 48, 51].

1.3 Neural networks and Sobolev spaces

Expressivity results have so far mainly focussed on L^p -type function classes, for $p \in [1, \infty]$. However, recent advances in applications of deep neural networks have shown the pressing need to derive a compre-

hensive understanding of the expressive power of deep neural networks with respect to Sobolev-regular functions.

Spaces of functions that admit generalized derivatives fulfilling suitable integrability properties, so-called *Sobolev spaces*, are a crucial concept in modern theory of PDEs (see e.g. [1, 22, 49]). Given some domain $\Omega \subset \mathbb{R}^d$, integrability order $1 \leq p \leq \infty$, and regularity $n \in \mathbb{N}$, the Sobolev space $W^{n,p}(\Omega)$ is defined as

$$W^{n,p}(\Omega) := \left\{ f : \Omega \rightarrow \mathbb{R} : \int_{\Omega} |D^{\alpha} f|^p dx < \infty \text{ for all } \alpha \in \mathbb{N}_0^d \text{ with } |\alpha| \leq n \right\},$$

and equipped with the norm

$$\|f\|_{W^{n,p}(\Omega)} := \left(\sum_{0 \leq |\alpha| \leq n} \int_{\Omega} |D^{\alpha} f|^p dx \right)^{1/p}.$$

To study properties of PDEs using functional analytic tools, the *variational formulation* of the PDE is derived. For this, a differential equation is reformulated via a differential operator mapping one function space to another. For a wide range of differential equations the appropriate spaces in this formulation are Sobolev spaces.

Motivated by the performance of deep learning-based solutions in classical machine learning tasks and, in particular, by their ability to overcome the curse of dimension, neural networks are now also applied for the approximative solution of PDEs (see [20, 28, 37, 55]). As a motivation for our results, we describe one of these methods in greater detail.

In [20] the authors present their meshfree *Deep Ritz method* for approximating solutions of potentially high-dimensional PDEs where the solution u can be expressed as the minimum $u = \arg \min_{v \in V} J(v)$ of a functional $J : V \rightarrow \mathbb{R}$ encoding the differential operator and external forces in the PDE. In this setting, V is typically a Sobolev space with regularity $n = 1$. Classical approaches compute an approximation u_h to the solution u by solving the minimization problem over a finite dimensional hypothesis space V_h (instead of V), where V_h is for example a finite element space with underlying mesh of fineness $h > 0$. Standard results in the literature relate the approximation error $\|u_h - u\|_V$ with the fineness h of the mesh. In case of the Deep Ritz method, the hypothesis space is a set of functions parameterized by the weights w of a neural network \mathcal{N}_{PDE} . This leads to the following optimization problem

$$w^* = \arg \min_w J(\mathcal{N}_{\text{PDE}}(w)),$$

where $\mathcal{N}_{\text{PDE}}(w^*)$ yields an approximation to the solution u . To derive error bounds relating the network complexity to the approximation error $\|\mathcal{N}_{\text{PDE}}(w^*) - u\|_V$, similar as in the classical setting, one needs to study the expressivity of neural networks with respect to Sobolev-regular norms.

The theoretical foundation for approximating a function with a neural network in Sobolev type norms was already given in a less known version of the universal approximation theorem by Hornik in [33, Thm. 3]. In particular, it was shown that if the activation function ϱ is k -times continuously differentiable, non-constant, and bounded, then any k -times continuously differentiable function f and its derivatives up to order k can be uniformly approximated by a shallow neural network on compact sets. In [16], it was shown that the theorem also holds for shallow ReLU networks if $k = 1$. Theorem 3 in [33] was also used in [55] to show the existence of a shallow network approximating solutions of the PDEs considered in that paper. Note though that in the above results only shallow networks are considered and it is not clear how the number of weights and neurons relates to the approximation error.

The relation of network complexity and approximation error for deep neural networks is extensively covered in the paper [60], where Yarotsky derived upper and lower complexity bounds for approximations in L^∞ norm of functions from the Sobolev space $W^{n,\infty}((0,1)^d)$ with neural networks. However, these results are restricted to the case $p = \infty$ and the derivatives are not included in the approximation norm as is necessary for the analysis of neural network based methods for the solution of PDEs.

1.4 Our contribution

The aim of this work is to study the simultaneous approximation of a function and its (fractional) derivative by a neural network. We extend the results of [60] to the case where the approximation is measured in the $W^{s,p}((0,1)^d)$ norm for $0 \leq s \leq 1$ and $1 \leq p \leq \infty$. Our contributions are three-fold:

- *Upper complexity bounds.* In Theorem 4.1, we derive upper complexity bounds for approximations of functions from the Sobolev space $W^{n,p}((0,1)^d)$ in weaker Sobolev norms with neural networks for continuous piecewise linear activation functions with a finite number of breakpoints for $1 \leq p \leq \infty$. For this, we explicitly construct ReLU neural networks realizing an ε -approximation error in a Sobolev norm with regularity s , where the complexity of the constructed network depends exponentially on $d/(n-s)$. Specifically, we show that, for any approximation accuracy $\varepsilon > 0$ and regularity $0 \leq s \leq 1$, there is a ReLU neural network \mathcal{N}_ε with at most $c \cdot \log_2(\varepsilon^{-n/(n-s)})$ layers and $c \cdot \varepsilon^{-d/(n-s)} \cdot \log_2(\varepsilon^{-n/(n-s)})$ weights and neurons such that for any $f \in W^{n,p}((0,1)^d)$ with $\|f\|_{W^{n,p}((0,1)^d)} \leq B$ there is a choice of weights w_f with

$$\|\mathcal{N}_\varepsilon(\cdot|w_f) - f(\cdot)\|_{W^{s,p}((0,1)^d)} \leq \varepsilon.$$

For $s = 1$ and $p = \infty$ the function $\mathcal{N}_\varepsilon(\cdot|w_f)$ and its weak gradient uniformly approximate f and the weak gradient of f , respectively. For non-integer s the function $\mathcal{N}_\varepsilon(\cdot|w_f)$ approximates the function f and its fractional derivative of order s approximates the fractional derivative of order s of f . The case $0 \leq s \leq 1$ and $p = \infty$ was already shown by one of the authors, I. Gühring, in his Master thesis [27].

- *Trade-off.* Similar as in classical approximation results using e.g. B-splines (see [46]), Theorem 4.1 shows a trade-off between the regularity s used in the approximation norm and the regularity used in the complexity bounds. When considering a function class with regularity n , measuring the approximation error ε in a norm with regularity $0 \leq s \leq 1$ results in a complexity penalty reciprocal to $n-s$ in the exponent of $1/\varepsilon$. In other words, a larger s necessitates a more complex network to satisfy the approximation task.
- *Lower complexity bounds.* We also establish lower bounds for the case $p = \infty$ and $s = 1$. In detail, we show in Theorem 4.3 that if a neural network \mathcal{N}_ε is able to realize an ε -approximation in $W^{1,\infty}$ norm for any $f \in W^{n,\infty}((0,1)^d)$ with $\|f\|_{W^{n,\infty}((0,1)^d)} \leq B$, then the number of weights of \mathcal{N}_ε is lower bounded by $c' \cdot \varepsilon^{-d/2(n-1)}$. This shows that the *curse of dimension*, i.e. the exponential dependency on d , cannot be avoided for uniform approximations in $W^{1,\infty}$ norm. A comparison with the approximation rates for the L^∞ norm derived in [60, Thm. 4 a)], shows that the same regularity-complexity trade-off can be observed for the lower bounds on an integer scale, i.e. $s \in \{0, 1\}$.

Our results can lead to complexity bounds for ReLU networks used in the numerical solution of PDEs by deep learning based approaches (see Section 1.3). In effect, we anticipate our results to be essential to improve the theoretical understanding of these methods.

1.5 Outline

In Section 2, we rigorously define a neural network and its architecture in mathematical terms and develop a network calculus. In Section 3, we briefly introduce (fractional) Sobolev spaces. In Section 4, we present our results which will be discussed in Section 5. To allow a concise presentation of our results most proofs are given in the appendix: Appendix A and B contain the necessary preparation for the proof of the results from Section 4.1 which is given in Appendix C. The results from Section 4.2 are proven in Appendix D.

1.6 Notation

In this subsection, we collect the notation used throughout this paper. Since most of the notation can be considered standard in the literature, we suggest that the reader should skip this subsection and only consider it if a symbol is unclear.

We set $\mathbb{N} := \{1, 2, \dots\}$ and $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. For $k \in \mathbb{N}_0$ the set $\mathbb{N}_{\geq k} := \{k, k+1, \dots\}$ consists of all natural numbers larger than or equal to k . For a set A we denote its cardinality by $|A| \in \mathbb{N} \cup \{\infty\}$. If $x \in \mathbb{R}$, then we write $\lceil x \rceil := \min\{k \in \mathbb{Z} : k \geq x\}$ where \mathbb{Z} is the set of integers.

If $d \in \mathbb{N}$ and $\|\cdot\|$ is a norm on \mathbb{R}^d , then we denote for $x \in \mathbb{R}^d$ and $r > 0$ by $B_{r,\|\cdot\|}(x)$ the open ball around x in \mathbb{R}^d with radius r , where the distance is measured in $\|\cdot\|$. By $|x|$ we denote the euclidean norm of x and by $\|x\|_{\ell^\infty}$ the maximum norm.

We endow \mathbb{R}^d with the standard topology and for $A \subset \mathbb{R}^d$ we denote by \bar{A} the closure of A and by ∂A the boundary of A . For the convex hull of A we write $\text{conv } A$. The diameter of a non-empty set $A \subset \mathbb{R}^d$ is always taken with respect to the euclidean distance, i.e. $\text{diam } A := \text{diam}_{|\cdot|} A := \sup_{x,y \in A} |x - y|$. If $A, B \subset \mathbb{R}^d$, then we write $A \subset\subset B$ if \bar{A} is compact in B .

For $d_1, d_2 \in \mathbb{N}$ and a matrix $A \in \mathbb{R}^{d_1, d_2}$ the number of nonzero entries of A is counted by $\|\cdot\|_{\ell^0}$, i.e.

$$\|A\|_{\ell^0} := |\{(i, j) : A_{i,j} \neq 0\}|.$$

If $d_1, d_2, d_3 \in \mathbb{N}$ and $A \in \mathbb{R}^{d_1, d_2}, B \in \mathbb{R}^{d_1, d_3}$, then we use the common block matrix notation and write for the horizontal concatenation of A and B

$$\left[\begin{array}{cc} A & B \end{array} \right] \in \mathbb{R}^{d_1, d_2 + d_3} \quad \text{or} \quad \left[\begin{array}{c|c} A & B \end{array} \right] \in \mathbb{R}^{d_1, d_2 + d_3}$$

where the second notation is used if a stronger delineation between different blocks is necessary. A similar notation is used for the vertical concatenation of $A \in \mathbb{R}^{d_1, d_2}$ and $B \in \mathbb{R}^{d_3, d_2}$.

Of course, the same notation also applies to (block) vectors.

For a function $f : X \rightarrow \mathbb{R}$ we denote by $\text{supp } f$ the support of f . If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are two functions, then we write $g \circ f : X \rightarrow Z$ for their composition.

We use the usual *multiindex* notation, i.e. for $\alpha \in \mathbb{N}_0^d$ we write $|\alpha| := \alpha_1 + \dots + \alpha_d$ and $\alpha! := \alpha_1! \dots \alpha_d!$. Moreover, if $x \in \mathbb{R}^d$, then we have

$$x^\alpha := \prod_{i=1}^d x_i^{\alpha_i}.$$

Let $\Omega \subset \mathbb{R}^d$ be open. For $n \in \mathbb{N}_0 \cup \{\infty\}$, we denote by $C^n(\Omega)$ the set of n times continuously differentiable functions on Ω . The space of *test functions* is

$$C_c^\infty(\Omega) := \{f \in C^\infty(\Omega) \mid \text{supp } f \subset\subset \Omega\}.$$

For $f \in C^n(\Omega)$ and $\alpha \in \mathbb{N}_0^d$ with $|\alpha| \leq n$ we write

$$D^\alpha f := \frac{\partial^{|\alpha|} f}{\partial x_1^{\alpha_1} \partial x_2^{\alpha_2} \dots \partial x_d^{\alpha_d}}.$$

We say that a function $f : \Omega \rightarrow \mathbb{R}^m$ for $\Omega \subset \mathbb{R}^d$ is *Lipschitz continuous* if there is a constant $L > 0$ such that

$$|f(x) - f(y)| \leq L|x - y|$$

for all $x, y \in \Omega$. If we want to emphasize the constant L , then we say that f is *L-Lipschitz*.

If X is a linear space and $\|\cdot\|_1$ and $\|\cdot\|_2$ are two norms on X , then we say that $\|\cdot\|_1$ and $\|\cdot\|_2$ are *equivalent* if there exist constants $C_1, C_2 > 0$ such that

$$C_1\|x\|_1 \leq \|x\|_2 \leq C_2\|x\|_1 \quad \text{for all } x \in X.$$

For two normed linear spaces X, Y we denote by $\mathcal{L}(X, Y)$ the set of bounded linear operators mapping X to Y and for $T \in \mathcal{L}(X, Y)$ the induced operator norm of T is denoted by

$$\|T\|_{\mathcal{L}(X, Y)} := \sup \{\|Tx\|_Y : x \in X, \|x\|_X \leq 1\}.$$

Let $a > 0$, then we say for two functions $f : (0, a) \rightarrow [0, \infty)$ and $g : (0, a) \rightarrow [0, \infty)$ that $f(\varepsilon)$ is in $\mathcal{O}(g(\varepsilon))$ if there exists $0 < \delta < a$ and $C > 0$ such that $f(\varepsilon) \leq Cg(\varepsilon)$ for all $\varepsilon \in (0, \delta)$.

2 Neural networks

We start by introducing the necessary notation for neural networks which we utilize in this paper. We will consider a general type of feedforward architecture that also allows for connections of neurons in non-neighboring layers. It can be seen, though, that any function realized by such a network can also be realized by a network with a more restrictive feedforward architecture where only neurons from neighboring layers can be connected (Lem. 2.13). As in [48], we draw a distinction between the neural network and the function that the network realizes. This gives us the possibility to develop a network calculus in the spirit of [48, Chap. 2]. This paper introduced the notion of network concatenation and parallelization, which we require in our proofs as well and therefore include here for the sake of completeness.

The following definition differs slightly from [48, Def. 2.1] in that we also allow connections between non-neighboring layers.

Definition 2.1. Let $d, L \in \mathbb{N}$.

- (1) A *neural network* Φ with *input dimension* d and L *layers* is a sequence of matrix-vector tuples

$$\Phi = ((A_1, b_1), (A_2, b_2), \dots, (A_L, b_L)),$$

where $N_0 = d$ and $N_1, \dots, N_L \in \mathbb{N}$, and where each A_l is an $N_l \times \sum_{k=0}^{l-1} N_k$ matrix, and $b_l \in \mathbb{R}^{N_l}$.

- (2) Let Φ be a neural network and $\varrho : \mathbb{R} \rightarrow \mathbb{R}$. Then we define the associated *realization of Φ with activation function ϱ* as the map $R_\varrho(\Phi) : \mathbb{R}^d \rightarrow \mathbb{R}^{N_L}$ such that

$$R_\varrho(\Phi)(x) = x_L,$$

where x_L results from the scheme given by

$$\begin{aligned} x_0 &:= x, \\ x_l &:= \varrho \left(A_l \left[\begin{array}{c|c} x_0^T & \dots & x_{l-1}^T \end{array} \right]^T + b_l \right), \quad \text{for } l = 1, \dots, L-1, \\ x_L &:= A_L \left[\begin{array}{c|c} x_0^T & \dots & x_{L-1}^T \end{array} \right]^T + b_L. \end{aligned}$$

Here ϱ acts componentwise, i.e., $\varrho(y) = [\varrho(y^1), \dots, \varrho(y^m)]$ for $y = [y^1, \dots, y^m] \in \mathbb{R}^m$. We sometimes write A_l in block-matrix form as

$$A_l = \left[\begin{array}{c|c} A_{l,x_0} & \dots & A_{l,x_{l-1}} \end{array} \right],$$

where A_{l,x_k} is an $N_l \times N_k$ matrix for $k = 0, \dots, l-1$ and $l = 1, \dots, L$. Then

$$\begin{aligned} x_l &= \varrho \left(A_{l,x_0} x_0 + \dots + A_{l,x_{l-1}} x_{l-1} + b_l \right), \quad \text{for } l = 1, \dots, L-1, \\ x_L &= A_{L,x_0} x_0 + \dots + A_{L,x_{L-1}} x_{L-1} + b_L. \end{aligned}$$

- (3) We call $N(\Phi) := d + \sum_{j=1}^L N_j$ the *number of neurons* of the network Φ , $L = L(\Phi)$ the *number of layers*, and finally $M(\Phi) := \sum_{j=1}^L (\|A_j\|_{\ell^0} + \|b_j\|_{\ell^0})$ denotes the number of nonzero entries of all A_l, b_l which we call the *number of weights* of Φ . Moreover, we refer to N_L as the *dimension of the output layer* of Φ .

We will now define the class of neural networks where only connections between neighboring layers are allowed. Networks without skip connections are a special case of the networks of Definition 2.1. Since they are more frequently used in the literature, we coin such networks *standard neural networks*.

Definition 2.2. If $\Phi = ((A_1, b_1), (A_2, b_2), \dots, (A_L, b_L))$ is a neural network as in Definition 2.1 and

$$A_l = \left[\begin{array}{c|c} A_{l,x_0} & \dots & A_{l,x_{l-1}} \end{array} \right],$$

satisfies $A_{l,x_i} = 0$ for $l = 1, \dots, L$ and $i = 0, \dots, l-2$, then we call Φ a *standard neural network*. In this case the computation scheme reduces to computing

$$\begin{aligned} x_0 &:= x, \\ x_l &:= \varrho \left(A_{l,x_{l-1}} x_{l-1} + b_l \right), \quad \text{for } l = 1, \dots, L-1, \\ x_L &:= A_{L,x_{L-1}} x_{L-1} + b_L. \end{aligned}$$

In practice, before training a neural network, i.e. adjusting the weights of the network, one has to decide which network architecture to use. The following definition will clarify the notion of a network architecture.

Definition 2.3. Let $d, L \in \mathbb{N}$.

- (1) A *neural network architecture* \mathcal{A} with input dimension d and L layers is a sequence of matrix-vector tuples

$$\mathcal{A} = ((A_1, b_1), (A_2, b_2), \dots, (A_L, b_L)),$$

where $N_0 = d$ and $N_1, \dots, N_L \in \mathbb{N}$, and where each A_l is an $N_l \times \sum_{k=0}^{l-1} N_k$ matrix with elements in $\{0, 1\}$, and b_l a vector of length N_l with elements in $\{0, 1\}$. Thus, a neural network architecture is a neural network with binary weights.

- (2) We call $N(\mathcal{A}) := d + \sum_{j=1}^L N_j$ the *number of neurons* of the architecture \mathcal{A} , $L = L(\mathcal{A})$ the *number of layers*, and finally $M(\mathcal{A}) := \sum_{j=1}^L (\|A_j\|_{\ell^0} + \|b_j\|_{\ell^0})$ denotes the number of nonzero entries of all A_l, b_l which we call the *number of weights* of \mathcal{A} . Moreover, we refer to N_L as the *dimension of the output layer* of \mathcal{A} .
- (3) We say that a *neural network* $\Phi = ((A'_1, b'_1), (A'_2, b'_2), \dots, (A'_L, b'_L))$ with input dimension d and L layers has architecture \mathcal{A} if
- (a) $N_l(\Phi) = N_l$ for all $l = 1, \dots, L$ and
 - (b) $[A_l]_{i,j} = 0$ implies $[A'_l]_{i,j} = 0$, and $[b_l]_i = 0$ implies $[b'_l]_i = 0$ for all $l = 1, \dots, L$ with $i = 1, \dots, N_l$ and $j = 1, \dots, \sum_{k=0}^{l-1} N_k$.
- (4) In the spirit of Definition 2.2 we say that \mathcal{A} is a *standard neural network architecture* if all weights connecting neurons from non-neighboring layers are zero.
- (5) By fixing an ordering of the nonzero entries in A_l for $l = 1, \dots, L$ we can uniquely identify a vector $w \in \mathbb{R}^{M(\mathcal{A})}$ with a (standard) neural network that has (standard) architecture \mathcal{A} . We denote this (standard) neural network by $\mathcal{A}(w)$.

Note that in general an architecture of a neural network does not need to be unique, i.e. a neural network Φ can have architecture \mathcal{A}_1 and \mathcal{A}_2 for $\mathcal{A}_1 \neq \mathcal{A}_2$.

Moreover, we define one of the in practice commonly used (see [38]) activation functions.

Definition 2.4. The function $\rho : \mathbb{R} \rightarrow \mathbb{R}$, $x \mapsto \max(0, x)$, is called *ReLU* (Rectified Linear Unit) *activation function*.

It can be shown, see [60, Prop. 1], that in terms of approximation and upper complexity bounds for networks with continuous and piecewise linear activation functions it suffices to focus on the ReLU.

2.1 Network concatenation

The concatenation of two networks allows us to realize compositions of network realizations, and will serve as a useful tool in some of our proofs. In detail, if Φ^1 and Φ^2 are two networks, then we are interested in a network $\Phi^1 \bullet \Phi^2$ such that $R_\rho(\Phi^1 \bullet \Phi^2)(x) = R_\rho(\Phi^1) \circ R_\rho(\Phi^2)(x)$ for all $x \in \mathbb{R}^d$. We will start by constructing a network $\Phi^1 \bullet \Phi^2$ for general activation functions ρ . In the case where the activation function is the ReLU, we use the fact that it is possible to find a network which realizes the identity function and construct a so-called *sparse concatenation* which gives us control over the complexity of the resulting network.

Definition 2.5. Let $L_1, L_2 \in \mathbb{N}$ and let

$$\Phi^1 = ((A_1^1, b_1^1), \dots, (A_{L_1}^1, b_{L_1}^1)), \quad \Phi^2 = ((A_1^2, b_1^2), \dots, (A_{L_2}^2, b_{L_2}^2))$$

be two neural networks such that the input layer of Φ^1 has the same dimension as the output layer of Φ^2 . Then, $\Phi^1 \bullet \Phi^2$ denotes the following $L_1 + L_2 - 1$ layer network:

$$\Phi^1 \bullet \Phi^2 := \left((A_1^2, b_1^2), \dots, (A_{L_2-1}^2, b_{L_2-1}^2), (\tilde{A}_1^1, \tilde{b}_1^1), \dots, (\tilde{A}_{L_1}^1, \tilde{b}_{L_1}^1) \right),$$

where

$$\tilde{A}_l^1 := [A_{l,x_0}^1 A_{L_2}^2 \mid A_{l,x_1}^1 \mid \dots \mid A_{l,x_{l-1}}^1] \quad \text{and} \quad \tilde{b}_l^1 := A_{l,x_0}^1 b_{L_2}^2 + b_l^1,$$

for $l = 1, \dots, L_1$. We call $\Phi^1 \bullet \Phi^2$ the *concatenation of Φ^1 and Φ^2* .

The following lemma can be verified with a direct computation.

Lemma 2.6. *Let Φ_1 and Φ_2 be two neural networks such that the input layer of Φ^1 has the same dimension as the output layer of Φ^2 . Then, $R_\varrho(\Phi^1 \bullet \Phi^2) = R_\varrho(\Phi^1) \circ R_\varrho(\Phi^2)$.*

Next, we adapt [48, Lem. 2.3] to the case where skip connections are allowed and define a network which realizes an identity function if the activation function is the ReLU.

Lemma 2.7. *Let $\varrho : \mathbb{R} \rightarrow \mathbb{R}$ be the ReLU and $d \in \mathbb{N}$. We define*

$$\Phi^{\text{Id}} := ((A_1, b_1), (A_2, b_2))$$

with

$$A_1 := \begin{bmatrix} \text{Id}_{\mathbb{R}^d} \\ -\text{Id}_{\mathbb{R}^d} \end{bmatrix}, \quad b_1 := 0, \quad A_2 := [0_{\mathbb{R}^d, d} \mid \text{Id}_{\mathbb{R}^d} \mid -\text{Id}_{\mathbb{R}^d}], \quad b_2 := 0.$$

Then $R_\varrho(\Phi^{\text{Id}}) = \text{Id}_{\mathbb{R}^d}$.

Using the identity network Φ^{Id} from the previous lemma we can now (see also [48, Def. 2.5]) define an alternative network concatenation which allows us to control the complexity of the resulting network for the case of the ReLU activation function.

Definition 2.8. Let $\varrho : \mathbb{R} \rightarrow \mathbb{R}$ be the ReLU, let $L_1, L_2 \in \mathbb{N}$, and let

$$\Phi^1 = ((A_1^1, b_1^1), \dots, (A_{L_1}^1, b_{L_1}^1)), \quad \Phi^2 = ((A_1^2, b_1^2), \dots, (A_{L_2}^2, b_{L_2}^2))$$

be two neural networks such that the input dimension of Φ^1 is equal to the output dimension of Φ^2 . Let Φ^{Id} be as in Lemma 2.7. Then the *sparse concatenation of Φ^1 and Φ^2* is defined as

$$\Phi^1 \odot \Phi^2 := \Phi^1 \bullet \Phi^{\text{Id}} \bullet \Phi^2.$$

The following remark shows that the sparse concatenation gives us the desired control over the complexity of the resulting network.

Remark 2.9. *We have*

$$\Phi^1 \odot \Phi^2 = \left((A_1^2, b_1^2), \dots, (A_{L_2-1}^2, b_{L_2-1}^2), \left(\begin{bmatrix} A_{L_2}^2 \\ -A_{L_2}^2 \end{bmatrix}, \begin{bmatrix} b_{L_2}^2 \\ -b_{L_2}^2 \end{bmatrix} \right), (\hat{A}_1^1, b_1^1), \dots, (\hat{A}_{L_1}^1, b_{L_1}^1) \right),$$

where

$$\hat{A}_l^1 := [0_{N_l^1, \sum_{k=0}^{L_2-1} N_k^2} \mid A_{l,x_0}^1 \mid -A_{l,x_0}^1 \mid A_{l,x_1}^1 \mid \dots \mid A_{l,x_{l-1}}^1],$$

for $l = 1, \dots, L_1$. Furthermore, it holds that $L(\Phi^1 \odot \Phi^2) = L_1 + L_2$, and $M(\Phi^1 \odot \Phi^2) \leq 2M(\Phi^1) + 2M(\Phi^2)$ and $N(\Phi^1 \odot \Phi^2) \leq 2N(\Phi^1) + 2N(\Phi^2)$.

2.2 Network parallelization

In some cases it is useful to combine several networks into one larger network. In particular, if Φ^1 and Φ^2 are two networks, then we are looking for a network $P(\Phi^1, \Phi^2)$ such that $R_\varrho(P(\Phi^1, \Phi^2))(x) = (R_\varrho(\Phi^1)(x), R_\varrho(\Phi^2)(x))$ for all $x \in \mathbb{R}^d$.

Definition 2.10. Let $L_1, L_2 \in \mathbb{N}$ and let

$$\Phi^1 = ((A_1^1, b_1^1), \dots, (A_{L_1}^1, b_{L_1}^1)), \quad \Phi^2 = ((A_1^2, b_1^2), \dots, (A_{L_2}^2, b_{L_2}^2))$$

be two neural networks with d -dimensional input. If $L_1 \leq L_2$, then we define

$$P(\Phi^1, \Phi^2) := \left(\left(\tilde{A}_1^1, \tilde{b}_1^1 \right), \dots, \left(\tilde{A}_{L_2}^1, \tilde{b}_{L_2}^1 \right) \right),$$

where

$$\tilde{A}_l := \left[\begin{array}{c|cc|ccc} A_{l,x_0}^1 & A_{l,x_1}^1 & 0 & \cdots & A_{l,x_{l-1}}^1 & 0 \\ \hline A_{l,x_0}^2 & 0 & A_{l,x_1}^2 & \cdots & 0 & A_{l,x_{l-1}}^2 \end{array} \right], \quad \tilde{b}_l := \begin{bmatrix} b_l^1 \\ b_l^2 \end{bmatrix} \quad \text{for } 1 \leq l < L_1,$$

and

$$\tilde{A}_l := \left[\begin{array}{c|c|ccc|ccc|ccc} A_{l,x_0}^2 & 0 & A_{l,x_1}^2 & \cdots & 0 & A_{l,x_{L_1-1}}^2 & A_{l,x_{L_1}}^2 & \cdots & A_{l,x_{l-1}}^2 \\ \hline A_{l,x_0}^1 & A_{l,x_1}^1 & 0 & \cdots & A_{l,x_{L_1-1}}^1 & 0 & 0 & \cdots & 0 \end{array} \right], \quad \tilde{b}_l := b_l^2$$

for $L_1 \leq l < L_2$ and

$$\tilde{A}_{L_2} := \left[\begin{array}{c|cc|ccc|ccc|ccc} A_{L_2,x_0}^1 & A_{L_2,x_1}^1 & 0 & \cdots & A_{L_2,x_{L_1-1}}^1 & 0 & 0 & \cdots & 0 \\ \hline A_{L_2,x_0}^2 & 0 & A_{L_2,x_1}^2 & \cdots & 0 & A_{L_2,x_{L_1-1}}^2 & A_{L_2,x_{L_1}}^2 & \cdots & A_{L_2,x_{L_2-1}}^2 \end{array} \right],$$

$$\tilde{b}_{L_2} := \begin{bmatrix} b_{L_2}^1 \\ b_{L_2}^2 \end{bmatrix}.$$

A similar construction is used for the case $L_1 > L_2$. Then $P(\Phi^1, \Phi^2)$ is a neural network with d -dimensional input and $\max\{L_1, L_2\}$ layers, which we call the *parallelization of Φ^1 and Φ^2* .

The following lemma is easy to see from the definition above.

Lemma 2.11. *Let Φ_1 and Φ_2 be two neural networks with d -dimensional input. Then $M(P(\Phi^1, \Phi^2)) = M(\Phi^1) + M(\Phi^2)$, $N(P(\Phi^1, \Phi^2)) = N(\Phi^1) + N(\Phi^2) - d$ and $R_\rho(P(\Phi^1, \Phi^2))(x) = (R_\rho(\Phi^1)(x), R_\rho(\Phi^2)(x))$.*

Remark 2.12. *The previous Definition 2.11 can be generalized to a parallelization of an arbitrary number of neural networks. Let $n \in \mathbb{N}$ and, for $i = 1, \dots, n$, let Φ^i be a neural network with d -dimensional input and L_i layers. The network*

$$P(\Phi^i : i = 1, \dots, n) := P(\Phi^1, \dots, \Phi^n) := P(\Phi^1, P(\Phi^2, P(\dots, P(\Phi^{n-1}, \Phi^n) \dots)))$$

with d -dimensional input and $L = \max\{L_1, \dots, L_n\}$ layers is called the parallelization of $\Phi^1, \Phi^2, \dots, \Phi^n$. We have

$$M(P(\Phi^1, \dots, \Phi^n)) = \sum_{i=1}^n M(\Phi^i) \quad \text{and} \quad N(P(\Phi^1, \dots, \Phi^n)) = \sum_{i=1}^n N(\Phi^i) - (n-1)d,$$

and $R_\rho(P(\Phi^1, \dots, \Phi^n))(x) = (R_\rho(\Phi^1)(x), \dots, R_\rho(\Phi^n)(x))$.

2.3 Standard neural networks

Next, we will show that if a function is a realization of a neural network, then it can also be realized by a standard neural network with similar complexity. The proof of this result is heavily based on the idea of the construction of an identity function with ReLU networks (Lem. 2.7).

Lemma 2.13. *Let ρ be the ReLU. Then there exist absolute constants $C_1, C_2 > 0$ such that for any neural network Φ with d -dimensional input and $L = L(\Phi)$ layers, $N = N(\Phi)$ neurons and $M = M(\Phi)$ nonzero weights, there is a standard neural network Φ_{st} with d -dimensional input, and L layers, at most $C_1 LN$ neurons and $C_2 \cdot (LN + M)$ nonzero weights such that*

$$R_\rho(\Phi)(x) = R_\rho(\Phi_{st})(x)$$

for all $x \in \mathbb{R}^d$.

Proof. If $L = 1$, then there is nothing to show. For $L > 1$ we start by defining the matrices

$$A_{1,x_0}^{\text{st}} := \begin{bmatrix} \text{Id}_{\mathbb{R}^{N_0}} \\ -\text{Id}_{\mathbb{R}^{N_0}} \\ A_{1,x_0} \end{bmatrix}, \quad b_0^{\text{st}} := \begin{bmatrix} 0 \\ 0 \\ b_0 \end{bmatrix}$$

and

$$A_{l,x_{l-1}}^{\text{st}} := \begin{bmatrix} \text{Id}_{\mathbb{R}^{2N_0}} & & & & 0 \\ & & & & \\ & 0 & & & \text{Id}_{\mathbb{R}^{\sum_{k=1}^{l-1} N_k}} \\ A_{l,x_0} & -A_{l,x_0} & A_{l,x_1} & \dots & A_{l,x_{l-1}} \end{bmatrix}, \quad b_l^{\text{st}} := \begin{bmatrix} 0 \\ 0 \\ b_l \end{bmatrix}$$

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

$$A_{L,x_{L-1}}^{\text{st}} := [A_{L,x_0} \quad -A_{L,x_0} \mid A_{L,x_1} \mid \dots \mid A_{L,x_{L-1}}], \quad b_L^{\text{st}} := b_L.$$

Now we set

$$A_l^{\text{st}} := [0 \mid \dots \mid 0 \mid A_{l,x_{l-1}}^{\text{st}}],$$

and

$$\Phi_{\text{st}} := ((A_1^{\text{st}}, b_1^{\text{st}}), \dots, (A_L^{\text{st}}, b_L^{\text{st}})).$$

Then Φ_{st} is a standard neural network with d -dimensional input and L layers and one can verify that $R_\varrho(\Phi)(x) = R_\varrho(\Phi_{\text{st}})(x)$ for all $x \in \mathbb{R}^d$. For the number of neurons in the l -th layer N_l^{st} we have $N_l^{\text{st}} = 2N_0 + \sum_{k=1}^l N_k \leq 2N$. Hence, $N(\Phi_{\text{st}}) \leq 2LN$. To estimate the number of weights, note that $\|A_l^{\text{st}}\|_{\ell^0} \leq 2N + 2\|A_l\|_{\ell^0}$ and $\|b_l^{\text{st}}\|_{\ell^0} = \|b_l\|_{\ell^0}$, from which we conclude that $N(\Phi_{\text{st}}) \leq 2(LN + M)$. \square

In the next remark we collect some properties of the class of functions that can be realized by a neural network.

Remark 2.14. *If Φ is a neural network with d -dimensional input and m -dimensional output and ϱ is the ReLU, then $R_\varrho(\Phi)$ is a Lipschitz-continuous, piecewise affine-linear function. This follows easily from Definition 2.2 and Lemma 2.13 which show that $R_\varrho(\Phi)$ can be expressed as the composition of Lipschitz-continuous, piecewise affine-linear transformations.*

In [4, Thm. 2.1] also the converse was shown, i.e. every piecewise affine-linear function can be realized by a neural network with the ReLU as activation function.

3 Sobolev spaces

A historical background of the development of Sobolev spaces in the context of PDEs can be found in [57, Lecture 1]. For a detailed treatment of the broad theory of Sobolev spaces we refer the reader to [1, 11, 22].

For this entire section, let $d \in \mathbb{N}$ and, if not stated otherwise, $\Omega \subset \mathbb{R}^d$ denote an open subset of \mathbb{R}^d . For $p \in [1, \infty]$, we denote by $L^p(\Omega)$ the standard Lebesgue spaces on Ω . The following definition can be found in [11, Chap. 4].

Definition 3.1 (Sobolev space). Let $n \in \mathbb{N}_0$ and $1 \leq p \leq \infty$. Then we define

$$W^{n,p}(\Omega) := \{f \in L^p(\Omega) : D^\alpha f \in L^p(\Omega) \text{ for all } \alpha \in \mathbb{N}_0^d \text{ with } |\alpha| \leq n\}.$$

Furthermore, for $f \in W^{n,p}(\Omega)$ and $1 \leq p < \infty$, we define the norm

$$\|f\|_{W^{n,p}(\Omega)} := \left(\sum_{0 \leq |\alpha| \leq n} \|D^\alpha f\|_{L^p(\Omega)}^p \right)^{1/p}$$

and

$$\|f\|_{W^{n,\infty}(\Omega)} := \max_{0 \leq |\alpha| \leq n} \|D^\alpha f\|_{L^\infty(\Omega)}.$$

We write $\nabla f := (D^1 f, \dots, D^d f)$. Moreover, the space $W^{n,p}(\Omega)$ equipped with the norm $\|\cdot\|_{W^{n,p}(\Omega)}$ is a Banach space. For $f \in W^{n,p}(\Omega)$ we shall use the notation

$$\|f\|_{W^{n,p}(\Omega)} = \|f\|_{W^{n,p}} = \|f(x)\|_{W^{n,p}(\Omega; dx)}$$

if the domain is clear from the context, or if we want to emphasize the variable x that the function f depends on, respectively. If $n = 0$, then the Sobolev space is just a Lebesgue space, i.e. $W^{0,p}(\Omega) = L^p(\Omega)$.

Many results about function spaces defined on a domain Ω require Ω to fulfill certain regularity conditions. Different geometrical conditions and the resulting properties have been intensively studied and can for example be found in [1]. For our purposes it is enough to focus on the condition introduced in the next definition which can be found in [22, Appendix C.1].

Definition 3.2 (Lipschitz-domain). We say that a bounded and open set $\Omega \subset \mathbb{R}^d$ is a *Lipschitz-domain* if for each $x_0 \in \partial\Omega$ there exists $r > 0$ and a Lipschitz continuous function $g : \mathbb{R}^{d-1} \rightarrow \mathbb{R}$ such that

$$\Omega \cap B_{r,|\cdot|}(x_0) = \{x \in B_{r,|\cdot|}(x) : x_d > g(x_1, \dots, x_{d-1})\},$$

after possibly relabeling and reorienting the coordinate axes.

In the sequel, we will only work with convex domains. Moreover, every open, bounded, and convex domain is a Lipschitz domain [25, Corollary 1.2.2.3].

3.1 Fractional Sobolev spaces

Fractional Sobolev spaces play an important role in the analysis of partial differential equations. In particular, they characterize the regularity of functions from Sobolev spaces defined on a domain Ω restricted to the boundary $\partial\Omega$ of the domain where the restriction is realized by a so-called trace operator (see e.g. [11]). Moreover, a detailed description of various areas of further application is listed in [12]. For a more in-depth analysis of these spaces the interested reader is referred to [39] and [57].

Next, we define fractional-order spaces in terms of an intrinsic norm.

Definition 3.3 (Sobolev-Slobodeckij spaces). We set for $0 < s < 1$ and $1 \leq p \leq \infty$

$$W^{s,p}(\Omega) := \left\{ f \in L^p(\Omega) : \|f\|_{W^{s,p}(\Omega)} < \infty \right\}$$

with

$$\|f\|_{W^{s,p}(\Omega)} := \left(\|f\|_{L^p(\Omega)}^p + \int_{\Omega} \int_{\Omega} \left(\frac{|f(x) - f(y)|}{|x - y|^{s+d/p}} \right)^p dx dy \right)^{1/p} \quad \text{for } 1 \leq p < \infty$$

and

$$\|f\|_{W^{s,\infty}(\Omega)} := \max \left\{ \|f\|_{L^\infty(\Omega)}, \operatorname{ess\,sup}_{x,y \in \Omega} \frac{|f(x) - f(y)|}{|x - y|^s} \right\}.$$

The space $W^{s,p}(\Omega)$ endowed with the norm $\|\cdot\|_{W^{s,p}(\Omega)}$ is a Banach space called *Sobolev-Slobodeckij space*.¹

4 Approximations with deep ReLU neural networks in Sobolev type norms

In this section, we present our main results which are upper and lower complexity bounds for approximations of functions from certain Sobolev spaces with deep ReLU networks. Our result includes [60, Thm. 1] (upper bounds) as a special case for $p = \infty$ and $s = 0$, and extends [60, Thm. 4 a)] (lower bounds) to the case where the approximation error is measured in a Sobolev-type norm.

From now on we assume that $\varrho : \mathbb{R} \rightarrow \mathbb{R}$ is the ReLU activation function (cf. Def. 2.4). We are interested in approximating functions in subsets of the Sobolev space $W^{n,p}((0,1)^d)$ with realizations of neural networks. For this we define the set:

$$\mathcal{F}_{n,d,p,B} := \left\{ f \in W^{n,p}((0,1)^d) : \|f\|_{W^{n,p}((0,1)^d)} \leq B \right\}.$$

¹For $p = \infty$ the space coincides with the space of bounded s -Hölder continuous functions. Precisely, each equivalence class contains a bounded s -Hölder continuous representative.

4.1 Upper complexity bounds

We measure the approximation with respect to a continuous scale of higher-order Sobolev-type norms for arbitrary $1 \leq p \leq \infty$ and additionally, reveal a trade-off between the regularity used in the norm in which the approximation error is measured and the regularity used in the complexity bounds. In particular, we show the following theorem:

Theorem 4.1. *Let $d \in \mathbb{N}$, $n \in \mathbb{N}_{\geq 2}$, $1 \leq p \leq \infty$, $B > 0$, and $0 \leq s \leq 1$. Then, there exists a constant $c = c(d, n, p, B, s) > 0$ with the following properties:*

For any $\varepsilon \in (0, 1/2)$, there is a neural network architecture $\mathcal{A}_\varepsilon = \mathcal{A}_\varepsilon(d, n, p, B, s, \varepsilon)$ with d -dimensional input and one-dimensional output such that for any $f \in \mathcal{F}_{n,d,p,B}$ there is a neural network Φ_ε^f that has architecture \mathcal{A}_ε such that

$$\|R_\varrho(\Phi_\varepsilon^f) - f\|_{W^{s,p}((0,1)^d)} \leq \varepsilon$$

and

- (i) $L(\mathcal{A}_\varepsilon) \leq c \cdot \log_2(\varepsilon^{-n/(n-s)});$
- (ii) $M(\mathcal{A}_\varepsilon) \leq c \cdot \varepsilon^{-d/(n-s)} \cdot \log_2(\varepsilon^{-n/(n-s)});$
- (iii) $N(\mathcal{A}_\varepsilon) \leq c \cdot \varepsilon^{-d/(n-s)} \cdot \log_2(\varepsilon^{-n/(n-s)}).$

Note that if Φ is a neural network and ϱ is the ReLU, then the restriction of $R_\varrho(\Phi)$ to $(0, 1)^d$ is bounded and Lipschitz continuous (cf. Rem. 2.14) and hence in view of [10, Chap. 1.3] an element of the Sobolev space $W^{1,\infty}((0, 1)^d)$. Therefore, the expressions in the previous theorem are well-defined.

Remark. *Theorem 4.1 holds for all activation functions that are similar to the ReLU in the following sense: If $\hat{\varrho}$ is any continuous piecewise linear activation function with M breakpoints, where $1 \leq M < \infty$, then it follows directly from [60, Prop. 1] that Theorem 4.1 holds.*

As a consequence of Theorem 4.1 and Lemma 2.13, we can easily derive a similar statement for standard neural networks.

Corollary 4.2. *Let $d \in \mathbb{N}$, $n \in \mathbb{N}_{\geq 2}$, $1 \leq p \leq \infty$, $B > 0$, and $0 \leq s \leq 1$. Then, there exists a constant $c = c(d, n, p, B, s) > 0$ with the following properties:*

For any $\varepsilon \in (0, 1/2)$, there is a standard neural network architecture $\mathcal{A}_{st,\varepsilon} = \mathcal{A}_{st,\varepsilon}(d, n, p, B, s, \varepsilon)$ with d -dimensional input and one-dimensional output such that for any $f \in \mathcal{F}_{n,d,p,B}$ there is a standard neural network $\Phi_{st,\varepsilon}^f$ that has architecture $\mathcal{A}_{st,\varepsilon}$ such that

$$\|R_\varrho(\Phi_{st,\varepsilon}^f) - f\|_{W^{s,p}((0,1)^d)} \leq \varepsilon$$

and

- (i) $L(\mathcal{A}_{st,\varepsilon}) \leq c \cdot \log_2(\varepsilon^{-n/(n-s)});$
- (ii) $M(\mathcal{A}_{st,\varepsilon}) \leq c \cdot \varepsilon^{-d/(n-s)} \cdot \log_2^2(\varepsilon^{-n/(n-s)});$
- (iii) $N(\mathcal{A}_{st,\varepsilon}) \leq c \cdot \varepsilon^{-d/(n-s)} \cdot \log_2^2(\varepsilon^{-n/(n-s)}).$

4.2 Lower complexity bounds

Next, we focus on lower bounds and show that the same regularity-complexity trade-off that can be observed for upper complexity bounds can be shown for lower bounds. In Theorem 4 a) in [60] Yarotsky proves that a network architecture capable of approximating any function $f \in \mathcal{F}_{n,d,\infty,1}$ up to a L^∞ -error ε has at least $c \cdot \varepsilon^{-d/(2n)}$ weights. We prove an extension of this result to approximations in $W^{1,\infty}$ norm. For the sake of readability, we combine the result from [60] with our result in the following theorem.

Theorem 4.3. *Let $d \in \mathbb{N}$, $n \in \mathbb{N}_{\geq 2}$, $B > 0$ and $k \in \{0, 1\}$. Then, there is a constant $c = c(d, n, B, k) > 0$ with the following property:*

If $\varepsilon \in (0, 1/2)$ and $\mathcal{A}_\varepsilon = \mathcal{A}_\varepsilon(d, n, B, k, \varepsilon)$ is an architecture such that for any $f \in \mathcal{F}_{n,d,\infty,B}$ there is a neural network Φ_ε^f that has architecture \mathcal{A}_ε and

$$\|R_\varrho(\Phi_\varepsilon^f) - f\|_{W^{k,\infty}((0,1)^d)} \leq \varepsilon,$$

then \mathcal{A}_ε has at least $M(\mathcal{A}_\varepsilon) \geq c\varepsilon^{-d/2(n-k)}$ weights.

Remark. Note that the case of standard neural network architectures is included in the theorem above and, thus, the same lower bounds hold true for this case.

Moreover, the proof can easily be adapted to piecewise polynomial activation functions (not necessarily continuous) with a finite number of breakpoints.

5 Discussion and future work

The motivation behind deriving theoretical results for deep neural networks is to improve our understanding of their empirical success in a broad area of applications. We now highlight some limitations and possible extensions of our results.

- *Curse of dimension.* One of the main reasons for the current interest in deep neural networks is their outstanding performance in high-dimensional problems. As an example, consider the popular ILSVR¹ challenge (see [50]), an image recognition task on the ImageNet database (see [18]) containing variable-resolution images which are typically downsampled before training to yield a roughly $(240 \times 240 \times 3)$ -dimensional input space (cf. e.g. [35, 54, 63]).

Our asymptotic upper complexity bounds for the number of weights and neurons are in $\mathcal{O}(\varepsilon^{-d/(n-s)})$, and thus depend strongly on the dimension d of the input space. Moreover, the constants in the upper bounds for the number of layers, weights, and neurons also increase exponentially with increasing d . Even if there is still a gap between our upper and lower bounds for the weights of a network architecture, Theorem 4.3 shows that one cannot expect to circumvent the curse of dimension in the setting considered in this paper.

Several ways have been proposed so far to tackle this common problem. One idea is to think of the data as residing on or near a manifold \mathcal{M} embedded in \mathbb{R}^d with dimension $d_{\mathcal{M}} \ll d$ (see [7]). Considering, for example, image classification problems this idea seems to be rather intuitive since most of the elements in $\mathbb{R}^{240 \times 240 \times 3}$ are not perceived as images. A similar approach is to narrow down the approximated function space by incorporating invariances (see [40]). If, for example, the approximated function f maps images to labels, then it makes sense to assume that f is translation and rotation invariant. The additional structure of the function space can then be exploited in the approximation (see [48, Section 5]). In particular, the last approach could also be suitable to tackle the curse of dimension in our result.

- *The power of depth.* For $s = 0$, the approximations obtained in Theorem 4.1 and [60, Thm. 1] coincide. In this case, Yarotsky showed in [60] that the constructed unbounded depth approximations for functions in $W^{n,\infty}((0,1)^d)$ (with $n > 2$) are asymptotically more efficient in the number of weights and neurons than approximations with fixed length L if

$$\frac{d}{n} < \frac{1}{2(L-1)}.$$

As a consequence, to be more efficient than a shallow network, i.e. a network with depth $L = 2$, one needs $n > 2d$ regularity. Even if this result does not completely explain the success of deep networks over shallow ones, since d is typically very large, it would be interesting to obtain similar results for higher-order Sobolev norms.

- *Unbounded complexity of weights.* When training a neural network on a computer, the weights have to be stored in memory. In practice, storing weights up to an arbitrary precision or even unbounded weights (in absolute value) is infeasible. In the construction of the neural network $R_{\varrho}(\Phi_{\varepsilon}^f)$ that approximates a function f up to an $W^{s,\infty}$ -error ε and realizes the upper complexity bound from Theorem 4.1 we used weights W_{ε} with $|W_{\varepsilon}| \rightarrow \infty$ as $\varepsilon \downarrow 0$ (see construction of $\Phi_{m,l}$ in the proof of Lemma C.3 (v)).

In [26, 48], neural networks are restricted to possess quantized weights (see [48, Def. 2.9]) which controls the complexity of the weights depending on the approximation error ε . It would be interesting to extend Theorem 4.1 to networks with quantized weights and to analyze how coarse

¹Typically abbreviated with *ILSRC* which stands for *ImageNet Large Scale Visual Recognition Challenge*.

a quantization is possible to maintain the associated upper bounds. This approach would also give rise to an alternative and possibly more efficient way to proving lower bounds as in [48]. Thereby the gap between the upper and lower bounds could potentially be closed.

Acknowledgements

I.G. would like to thank Mones Raslan for fruitful discussions on the topic and acknowledges support from the Research Training Group "Differential Equation- and Data-driven Models in Life Sciences and Fluid Dynamics: An Interdisciplinary Research Training Group (DAEDALUS)" (GRK 2433) funded by the German Research Foundation (DFG). G.K. acknowledges partial support by the Bundesministerium für Bildung und Forschung (BMBF) through the Berliner Zentrum for Machine Learning (BZML), Project AP4, by the Deutsche Forschungsgemeinschaft (DFG) through grants CRC 1114 "Scaling Cascades in Complex Systems", Project B07, CRC/TR 109 "Discretization in Geometry and Dynamics", Projects C02 and C03, RTG DAEDALUS (RTG 2433), Projects P1 and P3, RTG BIOQIC (RTG 2260), Projects P4 and P9, and SPP 1798 "Compressed Sensing in Information Processing", Coordination Project and Project Massive MIMO-I/II, by the Berlin Mathematics Research Center MATH+, Projects EF1-1 and EF1-4, and by the Einstein Foundation Berlin. P.P. is supported by a DFG Research Fellowship "Shearlet-based energy functionals for anisotropic phase-field methods".

A Interpolation spaces

In this section, we recall some essential notations and results from interpolation theory in Banach spaces. In particular, we present the K -method of real interpolation.

For two Banach spaces B_0 and B_1 interpolation techniques allow the definition and study of a family of spaces that in some sense bridge the gap between B_0 and B_1 . This rather abstract concept will be applied to Sobolev spaces in Section 3.1 to derive spaces of fractional regularity. Interpolation theory is a valuable tool in harmonic analysis and the study of partial differential equations. A small summary about the history of the development of the theory of interpolation spaces can be found in [57, Lecture 21]. Interested readers are recommended to read the monographs [8, 39, 57] and [58] for a detailed treatment in the context of partial differential equations.

The basic notions of functional analysis are assumed to be known and we refer to [61] for an introduction.

If B_0, B_1 are two Banach spaces such that B_1 is continuously embedded in B_0 , then we write $B_1 \subset B_0$ and call (B_0, B_1) an *interpolation couple*. One can also consider more general pairs of Banach spaces but this setting is sufficiently general for our purposes.

Definition A.1. Let (B_0, B_1) be an interpolation couple. For every $u \in B_1$ and $t > 0$ we define

$$K(t, u, B_0, B_1) := \inf_{v \in B_1} (\|u - v\|_{B_0} + t\|v\|_{B_1}).$$

The first term measures how well u can be approximated by elements from the smaller space B_1 in $\|\cdot\|_{B_0}$ and the second term is a penalty term weighted by t .

Now we can define interpolation spaces.

Definition A.2. Let (B_0, B_1) be an interpolation couple. Let $0 < \theta < 1$ and $1 \leq p \leq \infty$. We set

$$\|u\|_{(B_0, B_1)_{\theta, p}} := \begin{cases} \left(\int_0^\infty t^{-\theta p} K(t, u, B_0, B_1)^p \frac{dt}{t} \right)^{1/p}, & \text{for } 1 \leq p < \infty \\ \sup_{0 < t < \infty} t^{-\theta} K(t, u, B_0, B_1), & \text{for } p = \infty. \end{cases}$$

Moreover, we define the set

$$(B_0, B_1)_{\theta, p} := B_{\theta, p} := \left\{ u \in B_0 : \|u\|_{(B_0, B_1)_{\theta, p}} < \infty \right\}.$$

The norm $u \mapsto \|u\|_{(B_0, B_1)_{\theta, p}}$ turns $(B_0, B_1)_{\theta, p}$ into a Banach space (cf. [39, Prop. 1.5]) which is called a *real interpolation space*.

The next lemma shows that in the sense of a continuous embedding the spaces $B_{\theta, p}$ form a family of nested Banach spaces.

Lemma A.3. *Let $1 \leq p \leq \infty$, then the following holds:*

(i) *If (B_0, B_1) is an interpolation couple and $0 < \theta_1 < \theta_2 < 1$, then*

$$B_1 \subset B_{\theta_2, p} \subset B_{\theta_1, p} \subset B_0;$$

(ii) *if B is a Banach space and $0 < \theta < 1$, then $(B, B)_{\theta, p} = B$.*

Proof. For (i) see [39, p. 2] and for (ii) [39, Prop. 1.4]. □

One of the most important theorems in interpolation theory shows how the norm of an operator defined on interpolation couples relates to the operator norm with respect to the corresponding interpolation spaces. The theorem can be found in [39, Thm. 1.6].

Theorem A.4. *Let (A_0, A_1) and (B_0, B_1) be two interpolation couples. If $T \in \mathcal{L}(A_0, B_0) \cap \mathcal{L}(A_1, B_1)$, then $T \in \mathcal{L}(A_{\theta, p}, B_{\theta, p})$ for every $0 < \theta < 1$ and $1 \leq p \leq \infty$. Furthermore,*

$$\|T\|_{\mathcal{L}(A_{\theta, p}, B_{\theta, p})} \leq \|T\|_{\mathcal{L}(A_0, B_0)}^{1-\theta} \cdot \|T\|_{\mathcal{L}(A_1, B_1)}^{\theta}. \quad (\text{A.1})$$

As a corollary from the previous theorem the following useful result can be obtained (see [39, Corollary 1.7]).

Corollary A.5. *Let (B_0, B_1) be an interpolation couple. Moreover, let $0 < \theta < 1$ and $1 \leq p \leq \infty$, then there is a constant $c = c(\theta, p) > 0$ such that for all $u \in B_1$ we have*

$$\|u\|_{B_{\theta, p}} \leq c \|u\|_{B_0}^{1-\theta} \|u\|_{B_1}^{\theta}.$$

For the case $p = \infty$ we have $c(\theta, \infty) = 1$.

B Sobolev spaces

In this section, we give a brief introduction to Sobolev spaces and present related concepts which are important for the proof of Theorem 4.1. In detail, we show estimates for compositions and products of functions from certain Sobolev spaces in Subsection B.1 and B.2, respectively. In Subsection B.3, we introduce a generalization of Taylor polynomials which is one of the key tools in the proof of Theorem 4.1. In the last subsection, we apply Banach space interpolation theory to (integer) Sobolev spaces to define fractional Sobolev spaces. Again, let for this entire section, $d \in \mathbb{N}$ and, if not stated otherwise, $\Omega \subset \mathbb{R}^d$ denote an open subset of \mathbb{R}^d .

As a generalization of Definition 3.1, we now introduce Sobolev spaces of vector-valued functions (see [49, Chap. 1.2.3]).

Definition B.1. We set for $n \in \mathbb{N}_0$, $m \in \mathbb{N}$ and $1 \leq p \leq \infty$

$$W^{n, p}(\Omega; \mathbb{R}^m) := \{(f_1, \dots, f_m) : f_i \in W^{n, p}(\Omega)\}$$

with

$$\|f\|_{W^{n, p}(\Omega; \mathbb{R}^m)} := \left(\sum_{i=1}^m \|f_i\|_{W^{n, p}(\Omega)}^p \right)^{1/p} \quad \text{for } 1 \leq p < \infty$$

and

$$\|f\|_{W^{n, \infty}(\Omega; \mathbb{R}^m)} := \max_{i=1, \dots, m} \|f_i\|_{W^{n, \infty}(\Omega)}.$$

These spaces are again Banach spaces. To simplify the exposition, we introduce notation for a family of semi-norms.

Definition B.2 (Sobolev semi-norm). For $n, k \in \mathbb{N}_0$ with $k \leq n$, $m \in \mathbb{N}$ and $1 \leq p \leq \infty$ we define for $f \in W^{n,p}(\Omega; \mathbb{R}^m)$ the *Sobolev semi-norm*

$$|f|_{W^{k,p}(\Omega; \mathbb{R}^m)} := \left(\sum_{i=1, \dots, m, |\alpha|=k} \|D^\alpha f_i\|_{L^p(\Omega)}^p \right)^{1/p} \quad \text{for } 1 \leq p < \infty$$

and

$$|f|_{W^{k,\infty}(\Omega; \mathbb{R}^m)} := \max_{\substack{i=1, \dots, m, \\ |\alpha|=k}} \|D^\alpha f_i\|_{L^\infty(\Omega)}.$$

For $m = 1$ we only write $|\cdot|_{W^{k,p}(\Omega)}$. Note that $|\cdot|_{W^{0,p}(\Omega)}$ coincides with $\|\cdot\|_{L^p(\Omega)}$.

Rademacher's Theorem states that a Lipschitz continuous function $f : \Omega \rightarrow \mathbb{R}$ is differentiable (in the classical sense) almost everywhere in Ω (see e.g. [22, Thm. 5.8.6]). The following theorem yields an alternative characterization for the space $W^{1,\infty}$ using Lipschitz continuous functions and connects the notion of the classical and the weak derivative.

Theorem B.3. *Let $\Omega \subset \mathbb{R}^d$ be open and convex. If $f \in W^{1,\infty}(\Omega)$, then f is $\|\cdot\|_{L^\infty(\Omega; dx)}$ -Lipschitz. Conversely, if f is bounded and L -Lipschitz on Ω , then $f \in W^{1,\infty}(\Omega)$, the weak gradient of f agrees almost everywhere with the classical gradient, and*

$$\|\cdot\|_{L^\infty(\Omega; dx)} \leq L.$$

A proof can be found in [31, Thm. 4.1] together with [31, Rem. 4.2].

To use the bounds for the gradient obtained in the previous theorem in the Sobolev semi-norm setting from Definition B.2, the following observation turns out to be useful. If $f \in W^{1,\infty}(\Omega)$, then

$$|f|_{W^{1,\infty}(\Omega)} \leq \|\cdot\|_{L^\infty(\Omega)} \leq \sqrt{d} |f|_{W^{1,\infty}(\Omega)}. \quad (\text{B.1})$$

Moreover, it is not hard to see that if $\Omega \subset \mathbb{R}^d$ open and bounded, then $W^{n,\infty}(\Omega) \subset W^{n,p}(\Omega)$ for $1 \leq p < \infty$, and $n \in \mathbb{N}_0$.

B.1 Composition estimate

In this subsection, we derive an estimate for the semi-norm $|\cdot|_{W^{1,\infty}}$ of the composition $g \circ f$ of two functions $f \in W^{1,\infty}(\Omega; \mathbb{R}^m)$ and $g \in W^{1,\infty}(\mathbb{R}^m)$. Since in general this is not a well-defined notion we need to clarify what $g \circ f$ denotes in this case. To demonstrate the problem let f and g be as above and recall that f and g are equivalence classes of functions. Let now \hat{f} be a representative for the equivalence class f and \hat{g}_1, \hat{g}_2 be two representatives for g . Then $\hat{g}_1 = \hat{g}_2$ almost everywhere but in general we do not have $\hat{g}_1 \circ \hat{f} = \hat{g}_2 \circ \hat{f}$ almost everywhere. To circumvent this problem we denote by $g \circ f$ the equivalence class of $\hat{g} \circ \hat{f}$ where \hat{g} and \hat{f} are Lipschitz continuous representatives of g and f , respectively.

The next lemma recalls a well known fact about the composition of Lipschitz functions and is not hard to see.

Lemma B.4. *Let $d, m \in \mathbb{N}$ and $\Omega_1 \subset \mathbb{R}^d$, $\Omega_2 \subset \mathbb{R}^m$. Let $f : \Omega_1 \rightarrow \Omega_2$ be a L_1 -Lipschitz function and $g : \Omega_2 \rightarrow \mathbb{R}$ be a L_2 -Lipschitz function for some $L_1, L_2 > 0$. Then the composition $g \circ f : \Omega_1 \rightarrow \mathbb{R}$ is a $(L_1 \cdot L_2)$ -Lipschitz function.*

The following corollary establishes a chain rule estimate for $W^{1,\infty}$.

Corollary B.5. *Let $d, m \in \mathbb{N}$ and $\Omega_1 \subset \mathbb{R}^d$, $\Omega_2 \subset \mathbb{R}^m$ both be open, bounded, and convex. Then, there is a constant $C = C(d, m) > 0$ with the following property:*

If $f \in W^{1,\infty}(\Omega_1; \mathbb{R}^m)$ and $g \in W^{1,\infty}(\Omega_2)$ are Lipschitz continuous representatives such that $\text{ran } f \subset \Omega_2$, then for the composition $g \circ f$ it holds that $g \circ f \in W^{1,\infty}(\Omega_1)$ and we have

$$|g \circ f|_{W^{1,\infty}(\Omega_1)} \leq C |g|_{W^{1,\infty}(\Omega_2)} |f|_{W^{1,\infty}(\Omega_1; \mathbb{R}^m)}$$

and

$$\|g \circ f\|_{W^{1,\infty}(\Omega_1)} \leq C \max \left\{ \|g\|_{L^\infty(\Omega_2)}, |g|_{W^{1,\infty}(\Omega_2)} |f|_{W^{1,\infty}(\Omega_1; \mathbb{R}^m)} \right\}.$$

Proof. We set for $j = 1, \dots, m$

$$L_j := \|\nabla f_j\|_{L^\infty(\Omega_1)} \quad \text{and} \quad L_f := (L_1, \dots, L_m).$$

It then follows from Theorem B.3 that f_j is L_j -Lipschitz and hence, f is $|L_f|$ -Lipschitz. In the same manner we define

$$L_g := \|\nabla g\|_{L^\infty(\Omega_2)}$$

and have that g is L_g -Lipschitz. With Lemma B.4 we conclude that also the composition $g \circ f$ is $(|L_f| \cdot L_g)$ -Lipschitz. Since g is bounded and hence also $g \circ f$ applying Theorem B.3 again yields that $g \circ f \in W^{1,\infty}(\Omega_1)$ and furthermore

$$\|\nabla(g \circ f)\|_{L^\infty(\Omega_1)} \leq |L_f| \cdot L_g. \quad (\text{B.2})$$

We get

$$\begin{aligned} |g \circ f|_{W^{1,\infty}(\Omega_1)} &\leq \|\nabla(g \circ f)\|_{L^\infty(\Omega_1)} \\ &\leq |L_f| \cdot L_g \\ &\leq \sqrt{m} \|L_f\|_{\ell^\infty} \cdot \sqrt{m} |g|_{W^{1,\infty}(\Omega_2)} \\ &\leq \sqrt{dm} |f|_{W^{1,\infty}(\Omega_1; \mathbb{R}^m)} \cdot |g|_{W^{1,\infty}(\Omega_2)}, \end{aligned}$$

where we used Equation (B.1) in the first step and the previous estimate (B.2) in the second. The third step is due to estimating the ℓ^2 norm with the ℓ^∞ norm on \mathbb{R}^m , and applying (B.1) to the definition of L_g . The last step uses Equation (B.1) again. This shows the first estimate. The second estimate is now an easy consequence of the first. \square

B.2 Product estimate

In the following lemma, we show an estimate for the semi-norm of a product of weakly differentiable functions.

Lemma B.6. *Let $f \in W^{1,\infty}(\Omega)$ and $g \in W^{1,p}(\Omega)$ with $1 \leq p \leq \infty$, then $fg \in W^{1,p}(\Omega)$ and there exists a constant $C = C(d, p) > 0$ such that*

$$|fg|_{W^{1,p}(\Omega)} \leq C \left(|f|_{W^{1,\infty}(\Omega)} \|g\|_{L^p(\Omega)} + \|f\|_{L^\infty(\Omega)} |g|_{W^{1,p}(\Omega)} \right).$$

For $p = \infty$, we have $C = 1$.

Proof. Clearly, $fg, fD^i g + gD^i f \in L^1_{\text{loc}}(\Omega)$ so that the product formula in [23, Chap. 7.3] yields that for the weak derivatives of fg it holds

$$D^i(fg) = (D^i f)g + f(D^i g)$$

for $i = 1, 2, \dots, d$. Thus, we have

$$\begin{aligned} |fg|_{W^{1,p}(\Omega)} &= \left(\sum_{i=1}^d \|(D^i f)g + f(D^i g)\|_{L^p(\Omega)}^p \right)^{1/p} \\ &\leq \sum_{i=1}^d \|(D^i f)g + f(D^i g)\|_{L^p(\Omega)} \\ &\leq \sum_{i=1}^d \|(D^i f)g\|_{L^p(\Omega)} + \|f(D^i g)\|_{L^p(\Omega)}. \end{aligned}$$

Note that we have

$$\|(D^i f)g\|_{L^p(\Omega)} \leq \|D^i f\|_{L^\infty(\Omega)} \|g\|_{L^p(\Omega)} \leq |f|_{W^{1,\infty}(\Omega)} \|g\|_{L^p(\Omega)}$$

and

$$\|f(D^i g)\|_{L^p(\Omega)} \leq \|f\|_{L^\infty(\Omega)} \|D^i g\|_{L^p(\Omega)}.$$

Finally, we get

$$\begin{aligned} |fg|_{W^{1,p}(\Omega)} &\leq d|f|_{W^{1,\infty}(\Omega)} \|g\|_{L^p(\Omega)} + \|f\|_{L^\infty(\Omega)} \sum_{i=1}^d \|D^i g\|_{L^p(\Omega)} \\ &\leq C \left(|f|_{W^{1,\infty}(\Omega)} \|g\|_{L^p(\Omega)} + \|f\|_{L^\infty(\Omega)} |g|_{W^{1,p}(\Omega)} \right), \end{aligned}$$

where $C = C(d, p) > 0$ is a constant. Furthermore, it is clear that also $fg \in L^p(\Omega)$. \square

B.3 Averaged Taylor polynomial

In this subsection, we develop a polynomial approximation in the spirit of Taylor polynomials but appropriate for Sobolev spaces. A polynomial approximation P_f of a function $f \in \mathcal{F}_{n,d,p,B}$ is the first step towards an approximation of f realized by a neural network in the proof of Theorem 4.1.

A reference for this entire subsection is [10, Chap. 4.1].

Definition B.7 (averaged Taylor polynomial). Let $n \in \mathbb{N}$, $1 \leq p \leq \infty$ and $f \in W^{n-1,p}(\Omega)$, and let $x_0 \in \Omega$, $r > 0$ such that for the ball $B := B_{r,|\cdot|}(x_0)$ it holds that $B \subset\subset \Omega$. The corresponding *Taylor polynomial of order n of f averaged over B* is defined for $x \in \Omega$ as

$$Q^n f(x) := \int_B T_y^n f(x) \phi(y) dy, \quad (\text{B.3})$$

where

$$T_y^n f(x) := \sum_{|\alpha| \leq n-1} \frac{1}{\alpha!} D^\alpha f(y) (x-y)^\alpha \quad (\text{B.4})$$

and ϕ is an arbitrary cut-off function supported in \bar{B} , i.e.

$$\phi \in C_c^\infty(\mathbb{R}^d) \quad \text{with} \quad \phi(x) \geq 0 \text{ for all } x \in \mathbb{R}^d, \quad \text{supp } \phi = \bar{B} \quad \text{and} \quad \int_{\mathbb{R}^n} \phi(x) dx = 1.$$

A cut-off function as used in the previous definition always exists. A possible choice is

$$\psi(x) = \begin{cases} e^{-(1-(|x-x_0|/r)^2)^{-1}}, & \text{if } |x-x_0| < r \\ 0, & \text{else} \end{cases}$$

normalized by $\int_{\mathbb{R}^d} \psi(x) dx$. Next, we derive some properties of the averaged Taylor polynomial.

Remark B.8. *From the linearity of the weak derivative we can easily conclude that the averaged Taylor polynomial is linear in f .*

Recall that the averaged Taylor polynomial is defined via an integral and some cut-off function (cf. (B.3)) that perform an averaging of a polynomial expression (cf. (B.4)). Additionally, the following lemma shows that the averaged Taylor polynomial of order n is indeed a polynomial of degree less than n in x .

Lemma B.9. *Let $n \in \mathbb{N}$, $1 \leq p \leq \infty$ and $f \in W^{n-1,p}(\Omega)$, and let $x_0 \in \Omega$, $r > 0$, $R \geq 1$ such that for the ball $B := B_{r,|\cdot|}(x_0)$ it holds that $B \subset\subset \Omega$ and $B \subset B_{R,\|\cdot\|_{\ell^\infty}}(0)$. Then the Taylor polynomial of order n of f averaged over B can be written as*

$$Q^n f(x) = \sum_{|\alpha| \leq n-1} c_\alpha x^\alpha$$

for $x \in \Omega$.

Moreover, there exists a constant $c = c(n, d, R) > 0$ such that the coefficients c_α are bounded with $|c_\alpha| \leq cr^{-d/p} \|f\|_{W^{n-1,p}(\Omega)}$ for all α with $|\alpha| \leq n-1$.

Proof. The first part of the proof of this lemma follows closely the chain of arguments in [10, Eq. (4.1.5) - (4.1.8)]. We write for $\alpha \in \mathbb{N}_0^d$

$$(x - y)^\alpha = \prod_{i=1}^d (x_i - y_i)^{\alpha_i} = \sum_{\substack{\gamma, \beta \in \mathbb{N}_0^d, \\ \gamma + \beta = \alpha}} a_{(\gamma, \beta)} x^\gamma y^\beta,$$

where $a_{(\gamma, \beta)} \in \mathbb{R}$ are suitable constants with

$$|a_{(\gamma, \beta)}| \leq \binom{\gamma + \beta}{\gamma} = \frac{(\gamma + \beta)!}{\gamma! \beta!} \quad (\text{B.5})$$

in multi-index notation. Then, combining Equation (B.3) and (B.4) yields

$$\begin{aligned} Q^n f(x) &= \sum_{|\alpha| \leq n-1} \sum_{\gamma + \beta = \alpha} \frac{1}{\alpha!} a_{(\gamma, \beta)} x^\gamma \int_B D^\alpha f(y) y^\beta \phi(y) dy \\ &= \sum_{|\gamma| \leq n-1} x^\gamma \underbrace{\sum_{|\gamma + \beta| \leq n-1} \frac{1}{(\gamma + \beta)!} a_{(\gamma, \beta)} \int_B D^{\gamma + \beta} f(y) y^\beta \phi(y) dy}_{=: c_\gamma}. \end{aligned} \quad (\text{B.6})$$

For the second part, note that

$$\begin{aligned} \left| \int_B D^{\gamma + \beta} f(y) y^\beta \phi(y) dy \right| &\leq \int_B |D^{\gamma + \beta} f(y)| |y^\beta| |\phi(y)| dy \\ &\leq R^{|\beta|} \|f\|_{W^{n-1, p}(B)} \|\phi\|_{L^q(B)}, \end{aligned} \quad (\text{B.7})$$

where we used $B \subset B_{R, \|\cdot\|_{\ell^\infty}}(0)$ and the Hölder's inequality with $1/q = 1 - 1/p$. Next, since $\phi \in L^1(B) \cap L^\infty(B)$ and $\|\phi\|_{L^1} = 1$ we have (see [2, Chap. X.4 Exercise 4])

$$\|\phi\|_{L^q} \leq \|\phi\|_{L^1}^{1/q} \|\phi\|_{L^\infty}^{1-1/q} = \|\phi\|_{L^\infty}^{1/p}.$$

Combining the last estimate with Equation (B.7) yields

$$\begin{aligned} \left| \int_B D^{\gamma + \beta} f(y) y^\beta \phi(y) dy \right| &\leq R^{n-1} \|f\|_{W^{n-1, p}(\Omega)} \|\phi\|_{L^\infty(B)}^{1/p} \\ &\leq cR^{n-1} \|f\|_{W^{n-1, p}(\Omega)} r^{-d/p}, \end{aligned} \quad (\text{B.8})$$

where the second step follows from $\|\phi\|_{L^\infty} \leq cr^{-d}$ for some constant $c = c(d) > 0$ (see [10, Section 4.1]). To estimate the absolute value of the coefficients c_γ (defined in Equation B.6), we have

$$\begin{aligned} |c_\gamma| &\leq \sum_{|\gamma + \beta| \leq n-1} \frac{1}{(\gamma + \beta)!} |a_{(\gamma, \beta)}| \left| \int_B D^{\gamma + \beta} f(y) y^\beta \phi(y) dy \right| \\ &\leq cR^{n-1} \|f\|_{W^{n-1, p}(\Omega)} r^{-d/p} \sum_{|\gamma + \beta| \leq n-1} \frac{1}{\gamma! \beta!} = c' \|f\|_{W^{n-1, p}(\Omega)} r^{-d/p}. \end{aligned}$$

Here, the second step used Equation (B.8) together with Equation (B.5), and $c' = c'(n, d, R) > 0$ is a constant. \square

The next step is to derive approximation properties of the averaged Taylor polynomial. To this end, recall that for the (standard) Taylor expansion of some function f defined on a domain Ω in x_0 to yield an approximation at some point $x_0 + h$ the whole path $x_0 + th$ for $0 \leq t \leq 1$ has to be contained in Ω (see [41, Thm. 6.8.10]). In case of the averaged Taylor polynomial the expansion point x_0 is replaced by a ball B and we require that the path between each $x_0 \in B$ and each $x \in \Omega$ is contained in Ω . This geometrical condition is made precise in the following definition.

Definition B.10. Let $\Omega, B \subset \mathbb{R}^d$. Then Ω is called *star-shaped with respect to B* if,

$$\overline{\text{conv}}(\{x\} \cup B) \subset \Omega \quad \text{for all } x \in \Omega.$$

The next definition introduces a geometric notion which becomes important when given a family of subdivisions \mathcal{T}^h , $0 < h \leq 1$ of a domain Ω which becomes finer for smaller h . One typically needs to control the nondegeneracy of $(\mathcal{T}^h)_h$ which can be done e.g. with a uniformly bounded chunkiness parameter.

Definition B.11. Let $\Omega \subset \mathbb{R}^d$ be bounded. We define the set

$$\mathcal{R} := \left\{ r > 0 : \begin{array}{l} \text{there exists } x_0 \in \Omega \text{ such that } \Omega \text{ is} \\ \text{star-shaped with respect to } B_{r,|\cdot|}(x_0) \end{array} \right\}.$$

If $\mathcal{R} \neq \emptyset$, then we define

$$r_{\max}^* := \sup \mathcal{R} \quad \text{and call} \quad \gamma := \frac{\text{diam}(\Omega)}{r_{\max}^*}$$

the *chunkiness parameter* of Ω .

To emphasize the dependence on the set Ω , we will occasionally write $r_{\max}^*(\Omega)$ and $\gamma(\Omega)$.

Finally, the next lemma shows approximation properties of the averaged Taylor polynomial.

Lemma B.12 (Bramble-Hilbert). *Let $\Omega \subset \mathbb{R}^d$ be open and bounded, $x_0 \in \Omega$ and $r > 0$ such that Ω is star-shaped with respect to $B := B_{r,|\cdot|}(x_0)$, and $r > (1/2)r_{\max}^*$. Moreover, let $n \in \mathbb{N}$, $1 \leq p \leq \infty$ and denote by γ the chunkiness parameter of Ω . Then there exists a constant $C = C(n, d, \gamma) > 0$ such that for all $f \in W^{n,p}(\Omega)$*

$$|f - Q^n f|_{W^{k,p}(\Omega)} \leq Ch^{n-k} |f|_{W^{n,p}(\Omega)} \quad \text{for } k = 0, 1, \dots, n,$$

where $Q^n f$ denotes the Taylor polynomial of order n of f averaged over B and $h = \text{diam}(\Omega)$.

A proof can be found in [10, Lem. 4.3.8].

B.4 Fractional Sobolev spaces

In this subsection, we derive a generalization of Sobolev spaces characterized by fractional-order derivatives using two different approaches. First, we interpolate integer-valued Sobolev spaces, and secondly, we define an intrinsic norm. Both approaches are then shown to be equivalent under some regularity condition on the domain Ω .

We start by defining fractional-order Sobolev spaces $W^{s,p}$ for $0 < s < 1$ via Banach space interpolation (see Section A).

Definition B.13. Let $0 < s < 1$ and $1 \leq p \leq \infty$, then we set

$$\widetilde{W}^{s,p}(\Omega) := (L^p(\Omega), W^{1,p}(\Omega))_{s,p}.$$

The next theorem shows that, given suitable regularity conditions of the domain Ω , Definition B.13 and Definition 3.3 yield the same spaces with equivalent norms. The theorem can be found in [10, Thm. 14.2.3] for $1 \leq p < \infty$. The case $p = \infty$ is not included in [10, Thm. 14.2.3] but can be shown with the same technique.

Theorem B.14. *Let $\Omega \subset \mathbb{R}^d$ be a Lipschitz-domain. Then we have for $0 < s < 1$ and $1 \leq p \leq \infty$ that*

$$W^{s,p}(\Omega) = \widetilde{W}^{s,p}(\Omega)$$

with equivalence of the respective norms.

This motivates the following remark.

Remark. *We make use of the norm equivalence by only using the $W^{s,p}$ norm in the following section and transferring results obtained by interpolation theory for the $\widetilde{W}^{s,p}$ norm to the $W^{s,p}$ norm without mentioning it again.*

C Upper bounds for approximations

The strategy of the proof of Theorem 4.1 follows the idea of the proof of Theorem 1 in [60]: A polynomial approximation P_f of a function $f \in \mathcal{F}_{n,d,p,B}$ is approximated by a function realized by a neural network Φ_{P_f} . As in [60], we start by constructing a neural network that approximates the square function on the interval $(0, 1)$ up to an approximation error at most ε . In our result the error is measured in the $W^{1,\infty}$ norm (Prop. C.1). This result is then used to obtain an approximation of a multiplication operator (Prop. C.2). Next, we derive a partition of unity (Lem. C.3) as a product of univariate functions that can be realized by neural networks. Using this construction we then build a localized Taylor approximation for a function $f \in \mathcal{F}_{n,d,p,B}$ in the L^p and $W^{1,p}$ norm. An interpolation argument shows that same construction can be used for an approximation in the $W^{s,p}$ norm, where $0 \leq s \leq 1$ (Lem. C.4). In the next step, we lay the foundation for approximating products of the partition of unity and monomials, i.e. localized monomials, with ReLU networks (Lem. C.5) in the $W^{k,\infty}$ norm for $k = 0, 1$. Via an embedding this result is then used to approximate the localized polynomials in $W^{k,p}$ with ReLU networks (Lem. C.6).

All Sobolev spaces encountered in the following proofs are defined on open, bounded, and convex domains which are, in particular, Lipschitz-domains (see [25, Corollary 1.2.2.3]).

The results presented in the following two propositions can be found in a similar way in [52, Prop. 3.1]. However, since our results contain some minor extensions we decided to give the proof here for the sake of completeness. In detail, Proposition C.2 (ii) and (iii) are not shown in [52, Prop. 3.1].

In [60, Prop. 2], a network is constructed that approximates the square function $x \rightarrow x^2$ on the interval $(0, 1)$ in the L^∞ norm. Interestingly, the same construction can be used when measuring the approximation error in the $W^{1,\infty}$ norm. In particular, the depth and the number of weights and neurons of the network do not grow asymptotically faster to satisfy the approximation accuracy with respect to this stronger norm.

Proposition C.1. *There exist constants $c_1, c_2, c_3, c_4 > 0$, such that for all $\varepsilon \in (0, 1/2)$ there is a neural network Φ_ε^{sq} with at most $c_1 \cdot \log_2(1/\varepsilon)$ nonzero weights, at most $c_2 \cdot \log_2(1/\varepsilon)$ layers, at most $c_3 \cdot \log_2(1/\varepsilon)$ neurons, and with one-dimensional input and output such that*

$$\|R_\varrho(\Phi_\varepsilon^{sq})(x) - x^2\|_{W^{1,\infty}((0,1);dx)} \leq \varepsilon \quad (\text{C.1})$$

and $R_\varrho(\Phi_\varepsilon^{sq})(0) = 0$. Furthermore, it holds that

$$|R_\varrho(\Phi_\varepsilon^{sq})|_{W^{1,\infty}((0,1))} \leq c_4. \quad (\text{C.2})$$

Proof. In the proof of [60, Prop. 2] it is shown that there exist constants $c_1, c_2, c_3 > 0$, such that for each $m \in \mathbb{N}$ there is a neural network Φ_m with at most $c_1 \cdot m$ nonzero weights, at most $c_2 \cdot m$ layers, at most $c_3 \cdot m$ neurons the realization of which is a piecewise linear interpolation of $x \mapsto x^2$ on $(0, 1)$. In detail, it is shown there, that the network Φ_m satisfies for all $k \in \{0, \dots, 2^m - 1\}$ and all $x \in [\frac{k}{2^m}, \frac{k+1}{2^m}]$

$$R_\varrho(\Phi_m)(x) = \left(\frac{(k+1)^2}{2^m} - \frac{k^2}{2^m} \right) \left(x - \frac{k}{2^m} \right) + \left(\frac{k}{2^m} \right)^2. \quad (\text{C.3})$$

Thus, $R_\varrho(\Phi_m)$ is a piecewise linear interpolant of f with $2^m + 1$ uniformly distributed breakpoints $\frac{k}{2^m}, k = 0, \dots, 2^m$. In particular, $R_\varrho(\Phi_m)(0) = 0$. Furthermore, it is shown in the proof of [60, Prop. 2] that

$$\|R_\varrho(\Phi_m)(x) - x^2\|_{L^\infty((0,1);dx)} \leq 2^{-2-2m}. \quad (\text{C.4})$$

We will now show that the approximation error of the derivative can be bounded in a similar way. In particular, we show the estimate

$$|R_\varrho(\Phi_m) - x^2|_{W^{1,\infty}((0,1);dx)} \leq 2^{-m}. \quad (\text{C.5})$$

From Equation (C.3) we get for all $k = 0, \dots, 2^m - 1$

$$\begin{aligned} |R_\varrho(\Phi_m)(x) - x^2|_{W^{1,\infty}((k/2^m, (k+1)/2^m); dx)} &= \left\| \frac{(k+1)^2}{2^m} - \frac{k^2}{2^m} - 2x \right\|_{L^\infty((k/2^m, (k+1)/2^m); dx)} \\ &= \left\| \frac{2k+1}{2^m} - 2x \right\|_{L^\infty((k/2^m, (k+1)/2^m); dx)} \\ &= \max \left\{ \left| \frac{2k+1}{2^m} - 2 \frac{k}{2^m} \right|, \left| \frac{2k+1}{2^m} - 2 \frac{k+1}{2^m} \right| \right\} = 2^{-m}. \end{aligned}$$

Combining Equation (C.4) and (C.5) yields

$$\|R_\varrho(\Phi_m)(x) - x^2\|_{W^{1,\infty}((0,1); dx)} \leq \max \{2^{-2m-2}, 2^{-m}\} = 2^{-m}.$$

Clearly, the weak derivative of Φ_m is a piecewise constant function, which assumes its maximum on the last piece. Hence,

$$|R_\varrho(\Phi_m)|_{W^{1,\infty}((0,1))} \leq \frac{(2^m)^2 - (2^m - 1)^2}{2^m} = 2 - \frac{1}{2^m} \leq 2. \quad (\text{C.6})$$

Let now $\varepsilon \in (0, 1/2)$ and choose $m = \lceil \log_2(1/\varepsilon) \rceil$. Now, $\Phi_\varepsilon^{\text{sq}} := \Phi_m$ satisfies the approximation bound in Equation (C.1) and $R_\varrho(\Phi_\varepsilon^{\text{sq}})(0) = 0$. The estimate (C.2) holds because of Equation (C.6). The number of weights can be bounded by

$$M(\Phi_\varepsilon^{\text{sq}}) \leq c_1 \cdot m \leq c_1 \cdot (\log_2(1/\varepsilon) + 1) \leq 2 \cdot c_1 \cdot \log_2(1/\varepsilon).$$

In the same way, the number of neurons and layers can be bounded. This concludes the proof. \square

As in [60], we will now use Proposition C.1 and the polarization identity

$$xy = \frac{1}{2}((x+y)^2 - x^2 - y^2) \quad \text{for } x, y \in \mathbb{R}, \quad (\text{C.7})$$

to define an approximate multiplication where the approximation error is again (and in contrast to [60]) measured in the $W^{1,\infty}$ norm.

Proposition C.2. *For any $M \geq 1$, there exist constants $c, c_1 > 0$ and $c_2 = c_2(M) > 0$ such that for any $\varepsilon \in (0, 1/2)$, there is a neural network $\tilde{\times}$ with two-dimensional input and one-dimensional output that satisfies the following properties:*

- (i) $\|R_\varrho(\tilde{\times})(x, y) - xy\|_{W^{1,\infty}((-M, M)^2; dx dy)} \leq \varepsilon$;
- (ii) if $x = 0$ or $y = 0$, then $R_\varrho(\tilde{\times})(x, y) = 0$;
- (iii) $|R_\varrho(\tilde{\times}_\varepsilon)|_{W^{1,\infty}((-M, M)^2)} \leq cM$;
- (iv) the depth and the number of weights and neurons in $\tilde{\times}$ is at most $c_1 \log_2(1/\varepsilon) + c_2$.

Proof. Set $\delta := \varepsilon/(6M^2C)$, where C is the constant from Corollary B.5 for $n = 2$ and $m = 1$, and let Φ_δ^{sq} be the approximate squaring network from Proposition C.1 such that

$$\|R_\varrho(\Phi_\delta^{\text{sq}}) - x^2\|_{W^{1,\infty}((0,1); dx)} < \delta. \quad (\text{C.8})$$

As in the proof of [60, Prop. 3], we use the fact that $|x| = \varrho(x) + \varrho(-x)$ to see that a network $\tilde{\times}$ can be constructed with two-dimensional input and one-dimensional output that satisfies

$$R_\varrho(\tilde{\times})(x, y) = 2M^2 \left(R_\varrho(\Phi_\delta^{\text{sq}}) \left(\frac{|x+y|}{2M} \right) - R_\varrho(\Phi_\delta^{\text{sq}}) \left(\frac{|x|}{2M} \right) - R_\varrho(\Phi_\delta^{\text{sq}}) \left(\frac{|y|}{2M} \right) \right).$$

As a consequence of Proposition C.1, there exists a constant $c_0 > 0$ such that $\tilde{\times}$ has at most $c_0 \log_2(1/\varepsilon) + c_0 \log_2(6M^2) + 3 \leq c_0 \log_2(1/\varepsilon) + c_1$ layers, $3c_0 \log_2(1/\varepsilon) + 3c_0 \log_2(6M^2) + 9 \leq c' \log_2(1/\varepsilon) + c_2$ neurons and $3c_0 \log_2(1/\varepsilon) + 3c_0 \log_2(6M^2) + 17 \leq c'' \log_2(1/\varepsilon) + c_3$ nonzero weights. Here $c', c'' > 0$ and $c_1 = c_1(M), c_2 = c_2(M), c_3 = c_3(M) > 0$ are suitable constants and (iv) is hence satisfied.

Since $R_\rho(\Phi_\delta^{\text{sq}})(0) = 0$, (ii) easily follows.

Using the polarization identity (C.7), we can write

$$\begin{aligned} xy &= 4M^2 \frac{x}{2M} \frac{y}{2M} \\ &= 4M^2 \frac{1}{2} \left(\left(\frac{x}{2M} + \frac{y}{2M} \right)^2 - \left(\frac{x}{2M} \right)^2 - \left(\frac{y}{2M} \right)^2 \right) \\ &= 2M^2 \left(\left(\frac{|x+y|}{2M} \right)^2 - \left(\frac{|x|}{2M} \right)^2 - \left(\frac{|y|}{2M} \right)^2 \right). \end{aligned} \quad (\text{C.9})$$

To keep the following calculations simple, we introduce some notation. In particular, we define

$$u_{xy} : (-M, M)^2 \rightarrow (0, 1), \quad (x, y) \mapsto \frac{|x+y|}{2M},$$

$$u_x : (-M, M)^2 \rightarrow (0, 1), \quad (x, y) \mapsto \frac{|x|}{2M},$$

and

$$u_y : (-M, M)^2 \rightarrow (0, 1), \quad (x, y) \mapsto \frac{|y|}{2M}.$$

Setting $f : (0, 1) \rightarrow \mathbb{R}$, $x \mapsto x^2$ and using (C.9) we get

$$xy = 2M^2 \left(f \circ u_{xy}(x, y) - f \circ u_x(x, y) - f \circ u_y(x, y) \right). \quad (\text{C.10})$$

Now, we can estimate

$$\begin{aligned} &\|R_\rho(\tilde{\times})(x, y) - xy\|_{W^{1,\infty}((-M,M)^2; dx dy)} \\ &= \left\| 2M^2 (R_\rho(\Phi_\delta^{\text{sq}}) \circ u_{xy} - R_\rho(\Phi_\delta^{\text{sq}}) \circ u_x - R_\rho(\Phi_\delta^{\text{sq}}) \circ u_y) \right. \\ &\quad \left. - 2M^2 (f \circ u_{xy} - f \circ u_x - f \circ u_y) \right\|_{W^{1,\infty}((-M,M)^2)} \\ &\leq 2M^2 \left(\|(R_\rho(\Phi_\delta^{\text{sq}}) - f) \circ u_{xy}\|_{W^{1,\infty}((-M,M)^2)} + \|(R_\rho(\Phi_\delta^{\text{sq}}) - f) \circ u_x\|_{W^{1,\infty}((-M,M)^2)} \right. \\ &\quad \left. + \|(R_\rho(\Phi_\delta^{\text{sq}}) - f) \circ u_y\|_{W^{1,\infty}((-M,M)^2)} \right). \end{aligned} \quad (\text{C.11})$$

Note that for the inner functions in the compositions in Equation (C.11), it holds that

$$|u_{xy}|_{W^{1,\infty}((-M,M)^2)} = |u_x|_{W^{1,\infty}((-M,M)^2)} = |u_y|_{W^{1,\infty}((-M,M)^2)} = 1/(2M). \quad (\text{C.12})$$

Hence, to finally prove (i), we apply Corollary B.5 to (C.11) and get

$$\begin{aligned} &\|R_\rho(\tilde{\times})(x, y) - xy\|_{W^{1,\infty}((-M,M)^2; dx dy)} \\ &\leq 2M^2 3C \max \left\{ \|R_\rho(\Phi_\delta^{\text{sq}}) - f\|_{L^\infty((0,1))}, |R_\rho(\Phi_\delta^{\text{sq}}) - f|_{W^{1,\infty}((0,1))} \cdot 1/(2M) \right\} \\ &\leq 6M^2 C \delta = \varepsilon, \end{aligned}$$

where we used (C.8) in the second step.

To show (iii) we use the chain rule estimate from Corollary B.5 and get

$$\begin{aligned} &|R_\rho(\tilde{\times})|_{W^{1,\infty}((-M,M)^2)} \\ &\leq 2M^2 C \sum_{z \in \{x, y, xy\}} |R_\rho(\Phi_\delta^{\text{sq}})|_{W^{1,\infty}((0,1))} |u_z|_{W^{1,\infty}((-M,M)^2)} \\ &\leq 2M^2 C 3C' \frac{1}{2M} = 3CC'M, \end{aligned}$$

where we used Equation (C.12) in the third step and C' is the absolute constant from Equation (C.2) in Proposition C.1. \square

Next, we construct a partition of unity (in the same way as in [60, Thm. 1]) that can be defined as a product of piecewise linear functions, such that each factor of the product can be realized by a neural network. We will use this product structure in Lemma C.6 together with a generalized version of the approximate multiplication from Proposition C.2 to approximate localized polynomials with ReLU networks.

Lemma C.3. *For any $d, N \in \mathbb{N}$ there exists a collection of functions*

$$\Psi = \{\phi_m : m \in \{0, \dots, N\}^d\}$$

with $\phi_m : \mathbb{R}^d \rightarrow \mathbb{R}$ for all $m \in \{0, \dots, N\}^d$ with the following properties:

- (i) $0 \leq \phi_m(x) \leq 1$ for every $\phi_m \in \Psi$ and every $x \in \mathbb{R}^d$;
- (ii) $\sum_{\phi_m \in \Psi} \phi_m(x) = 1$ for every $x \in [0, 1]^d$;
- (iii) $\text{supp } \phi_m \subset B_{1/N, \|\cdot\|_{\ell^\infty}}(m/N)$ for every $\phi_m \in \Psi$;
- (iv) there exists a constant $c \geq 1$ such that $\|\phi_m\|_{W^{k, \infty}(\mathbb{R}^d)} \leq (c \cdot N)^k$ for $k \in \{0, 1\}$;
- (v) there exist absolute constants $C, c \geq 1$ such that for each $\phi_m \in \Psi$ there is a neural network Φ_m with d -dimensional input and d -dimensional output, with at most three layers, Cd nonzero weights and neurons, that satisfies

$$\prod_{l=1}^d [R_\varrho(\Phi_m)]_l = \phi_m$$

$$\text{and } \|[R_\varrho(\Phi_m)]_l\|_{W^{k, \infty}((0,1)^d)} \leq (cN)^k \text{ for all } l = 1, \dots, d \text{ and } k \in \{0, 1\}.$$

Proof. As in [60], we define the functions

$$\psi : \mathbb{R} \rightarrow \mathbb{R}, \quad \psi(x) := \begin{cases} 1, & |x| < 1, \\ 0, & 2 < |x|, \\ 2 - |x|, & 1 \leq |x| \leq 2, \end{cases}$$

and $\phi_m : \mathbb{R}^d \rightarrow \mathbb{R}$ as a product of scaled and shifted versions of ψ . Concretely, we set

$$\phi_m(x) := \prod_{l=1}^d \psi\left(3N\left(x_l - \frac{m_l}{N}\right)\right), \quad (\text{C.13})$$

for $m = (m_1, \dots, m_d) \in \{0, \dots, N\}^d$. Then, (i), (ii) and (iii) follow easily from the definition.

To show (iv), note that $\|\phi_m\|_{L^\infty} \leq 1$ follows already from (i). It now suffices to show the claim for the $W^{1, \infty}$ semi-norm. For this, let $l \in \{1, \dots, d\}$ and $x \in \mathbb{R}^d$, then

$$\begin{aligned} \left| \frac{\partial}{\partial x_l} \phi_m(x) \right| &= \left| \prod_{i=1, i \neq l}^d \psi\left(3N\left(x_i - \frac{m_i}{N}\right)\right) \right| \left| \psi'\left(3N\left(x_l - \frac{m_l}{N}\right)\right) \right| 3N \\ &\leq 3N. \end{aligned}$$

It follows that $|\phi_m|_{W^{1, \infty}} \leq 3N$.

To show (v), we start by constructing a network Φ_ψ that realizes the function ψ . For this we set

$$A_1 := \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \quad b_1 := \begin{bmatrix} 2 \\ 1 \\ -1 \\ -2 \end{bmatrix} \quad \text{and} \quad A_2 := [0 \mid 1 \quad -1 \quad -1 \quad 1], \quad b_2 := 0,$$

and $\Phi_\psi := ((A_1, b_1), (A_2, b_2))$. Then Φ_ψ is a two-layer network with one-dimensional input and one-dimensional output, with 12 nonzero weights and 6 neurons such that

$$R_\varrho(\Phi_\psi)(x) = \psi(x) \quad \text{for all } x \in \mathbb{R}.$$

We denote by $\Phi_{m,l}$ the one-layer network with d -dimensional input and one-dimensional output, with 2 nonzero weights and $d + 1$ neurons such that $R_\varrho(\Phi_{m,l}) = 3N(x_l - m_l/N)$ for all $x \in \mathbb{R}^d$ and $m \in \{0, \dots, N\}^d$, $l = 1, \dots, d$. Finally, we define

$$\Phi_m := P(\Phi_\psi \odot \Phi_{m,l} : l = 1, \dots, d),$$

which is a three-layer network with d -dimensional input and d -dimensional output, with at most $d \cdot 2 \cdot (12 + 2) = d \cdot 28$ nonzero weights and at most $9d$ neurons, and

$$[R_\varrho(\Phi_m)]_l(x) = R_\varrho(\Phi_\psi \odot \Phi_{m,l})(x) = \psi\left(3N\left(x_l - \frac{m_l}{N}\right)\right)$$

for $l = 1, \dots, d$ and $x \in \mathbb{R}^d$. Clearly, it follows that $\prod_{l=1}^d [R_\varrho(\Phi_m)]_l(x) = \phi_m(x)$ for all $x \in \mathbb{R}^d$. The last part of (v) can be shown similarly as (iv). \square

The following lemma uses the partition of unity from Lemma C.3 and the Bramble-Hilbert Lemma B.12 in a classical way to derive an approximation with localized polynomials in the L^p norm and the $W^{1,p}$ norm. Using an interpolation argument, we can generalize this result to the case where the approximation is performed with respect to the $W^{s,\infty}$ norm, where $0 \leq s \leq 1$.

Lemma C.4. *Let $d, N \in \mathbb{N}$, $n \in \mathbb{N}_{\geq 2}$, $1 \leq p \leq \infty$ and $\Psi = \Psi(d, N) = \{\phi_m : m \in \{0, \dots, N\}^d\}$ be the partition of unity from Lemma C.3. Then there is a constant $C = C(d, n, p) > 0$ such that for any $f \in W^{n,p}((0,1)^d)$, there exist polynomials $p_{f,m}(x) = \sum_{|\alpha| \leq n-1} c_{f,m,\alpha} x^\alpha$ for $m \in \{0, \dots, N\}^d$ with the following properties:*

Let $0 \leq s \leq 1$ and set $f_N := \sum_{m \in \{0, \dots, N\}^d} \phi_m p_{f,m}$, then the operator $T_s : W^{n,p}((0,1)^d) \rightarrow W^{s,p}((0,1)^d)$ with $T_s f = f - f_N$ is linear and bounded with

$$\|T_s f\|_{W^{s,p}((0,1)^d)} \leq C \left(\frac{1}{N}\right)^{n-s} \|f\|_{W^{n,p}((0,1)^d)}.$$

Furthermore, there is a constant $c = c(d, n) > 0$ such that for any $f \in W^{n,p}((0,1)^d)$ the coefficients of the polynomials $p_{f,m}$ satisfy

$$|c_{f,m,\alpha}| \leq c \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})} N^{d/p}$$

for all $\alpha \in \mathbb{N}_0^d$ with $|\alpha| \leq n - 1$ and $m \in \{0, \dots, N\}^d$, where $\Omega_{m,N} := B_{\frac{1}{N}, \|\cdot\|_{\ell^\infty}}\left(\frac{m}{N}\right)$ and $\tilde{f} \in W^{n,p}(\mathbb{R}^d)$ is an extension of f .

Proof. The idea of the proof is similar to the first part of the proof of [60, Thm. 1]. We use approximation properties of averaged Taylor polynomials (see Bramble-Hilbert Lemma B.12) to derive local estimates and then combine them using a partition of unity to obtain a global estimate. In order to use this strategy also near the boundary, we make use of an extension operator.

For this, let $E : W^{n,p}((0,1)^d) \rightarrow W^{n,p}(\mathbb{R}^d)$ be the extension operator from [56, Thm. VI.3.1.5] and set $\tilde{f} := Ef$. Note that for arbitrary $\Omega \subset \mathbb{R}^d$ and $1 \leq k \leq n$ it holds

$$\|\tilde{f}\|_{W^{k,p}(\Omega)} \leq \|\tilde{f}\|_{W^{n,p}(\mathbb{R}^d)} \leq C_E \|f\|_{W^{n,p}((0,1)^d)}, \quad (\text{C.14})$$

where $C_E = C_E(n, p, d)$ is the norm of the extension operator.

Step 1 (Averaged Taylor polynomials): For each $m \in \{0, \dots, N\}^d$ we set

$$\Omega_{m,N} := B_{\frac{1}{N}, \|\cdot\|_{\ell^\infty}}\left(\frac{m}{N}\right) \quad \text{and} \quad B_{m,N} := B_{\frac{3}{4N}, |\cdot|}\left(\frac{m}{N}\right),$$

and denote by $p_m = p_{f,m}$ the Taylor polynomial of order n of \tilde{f} averaged over $B_{m,N}$ (cf. Def. B.7). It follows from Proposition B.9 (for $\Omega = \Omega_{m,N}$, $B = B_{m,N}$ and $R = 2$) that we can write $p_m = \sum_{|\alpha| \leq n-1} c_{m,\alpha} x^\alpha$ and that there is a constant $c' = c'(n, d) > 0$ such that

$$|c_{m,\alpha}| \leq c' \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})} \left(\frac{3}{4N}\right)^{-d/p} \leq c'' \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})} N^{d/p}$$

for $m \in \{0, \dots, N\}^d$, where $c'' = c''(n, d, p) > 0$ is a suitable constant. It now suffices to show (i) and (ii).

Step 2 (Local estimates in $\|\cdot\|_{W^{k,p}}$, $k \in \{0, 1\}$): To check that the conditions of the Bramble-Hilbert Lemma B.12 are fulfilled, note that $B_{m,N} \subset\subset \Omega_{m,N}$. Furthermore, $B_{m,N}$ is a ball in $\Omega_{m,N}$ such that $\Omega_{m,N}$ is star-shaped with respect to $B_{m,N}$. We have $\text{diam}_{|\cdot|}(\Omega_{m,N}) = (2\sqrt{d})/N$ and $r_{\max}^*(\Omega_{m,N}) = 1/N$ and, thus,

$$r_{|\cdot|}(B_{m,N}) = \frac{3}{4N} > \frac{1}{2} \cdot \frac{1}{N} = \frac{1}{2} \cdot r_{\max}^*(\Omega_{m,N}).$$

Finally, we have for the chunkiness parameter of $\Omega_{m,N}$

$$\gamma(\Omega_{m,N}) = \text{diam}(\Omega_{m,N}) \cdot \frac{1}{r_{\max}^*(\Omega_{m,N})} = \frac{2\sqrt{d}}{N} \cdot N = 2\sqrt{d}. \quad (\text{C.15})$$

Applying the Bramble-Hilbert Lemma B.12 yields for each $m \in \{0, \dots, N\}^d$ the local estimate

$$\|\tilde{f} - p_m\|_{L^p(\Omega_{m,N})} \leq C_1 \left(\frac{2\sqrt{d}}{N}\right)^n |\tilde{f}|_{W^{n,p}(\Omega_{m,N})} \leq C_2 \left(\frac{1}{N}\right)^n \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})}. \quad (\text{C.16})$$

Here, $C_1 = C_1(n, d) > 0$ is the constant from Lemma B.12 which only depends on n and d , since the chunkiness parameter of $\Omega_{m,N}$ is a constant depending only on d (see (C.15)) and $C_2 = C_2(n, d) > 0$. In the same way, we get

$$|\tilde{f} - p_m|_{W^{1,p}(\Omega_{m,N})} \leq C_3 \left(\frac{1}{N}\right)^{n-1} \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})}, \quad (\text{C.17})$$

where $C_3 = C_3(n, d) > 0$ is a suitable constant.

The first step towards a global estimate is now to combine Equation (C.16) and (C.17) with the cut-off functions from the partition of unity. We have

$$\begin{aligned} \|\phi_m(\tilde{f} - p_m)\|_{L^p(\Omega_{m,N})} &\leq \|\phi_m\|_{L^\infty(\Omega_{m,N})} \cdot \|\tilde{f} - p_m\|_{L^p(\Omega_{m,N})} \\ &\stackrel{(\text{Lem. C.3 (iv), Eq. (C.16)})}{\leq} C_2 \left(\frac{1}{N}\right)^n \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})}. \end{aligned} \quad (\text{C.18})$$

Furthermore, using the product inequality for weak derivatives from Lemma B.6 we get that there is a constant $C' = C'(d, p) > 0$ such that

$$\begin{aligned} |\phi_m(\tilde{f} - p_m)|_{W^{1,p}(\Omega_{m,N})} &\leq C' |\phi_m|_{W^{1,\infty}(\Omega_{m,N})} \cdot \|\tilde{f} - p_m\|_{L^p(\Omega_{m,N})} \\ &\quad + C' \|\phi_m\|_{L^\infty(\Omega_{m,N})} \cdot |\tilde{f} - p_m|_{W^{1,p}(\Omega_{m,N})} \\ &\leq C' \cdot \tilde{c}N \cdot C_2 \left(\frac{1}{N}\right)^n \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})} + C' \cdot C_3 \left(\frac{1}{N}\right)^{n-1} \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})} \\ &= C_4 \left(\frac{1}{N}\right)^{n-1} \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})}, \end{aligned} \quad (\text{C.19})$$

where the first part of the second inequality follows from Lemma C.3 (iv) with $k = 1$ and an absolute constant $\tilde{c} \geq 1$, together with Equation (C.16) and the second part from Lemma C.3 (iv) with $k = 0$, together with Equation (C.17). Here, $C_4 = C_4(n, d, p) > 0$. Now it easily follows from (C.18) and (C.19) that

$$\|\phi_m(\tilde{f} - p_m)\|_{W^{1,p}(\Omega_{m,N})} \leq C_5 \left(\frac{1}{N}\right)^{n-1} \|\tilde{f}\|_{W^{n,p}(\Omega_{m,N})}, \quad (\text{C.20})$$

for some constant $C_5 = C_5(n, d, p) > 0$.

Step 3 (Global estimate in $\|\cdot\|_{W^{k,p}}$, $k \in \{0, 1\}$): To derive the global estimate, we start by noting that with property (ii) from Lemma C.3 we have

$$\tilde{f}(x) = \sum_{m \in \{0, \dots, N\}^d} \phi_m(x) \tilde{f}(x), \quad \text{for a.e. } x \in (0, 1)^d. \quad (\text{C.21})$$

Using that \tilde{f} is an extension of f on $(0, 1)^d$ we can write for $k \in \{0, 1\}$

$$\begin{aligned} \left\| f - \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m \right\|_{W^{k,p}((0,1)^d)}^p &= \left\| \tilde{f} - \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m \right\|_{W^{k,p}((0,1)^d)}^p \\ &\stackrel{\text{(Eq. (C.21))}}{=} \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m (\tilde{f} - p_m) \right\|_{W^{k,p}((0,1)^d)}^p \\ &\leq \sum_{\tilde{m} \in \{0, \dots, N\}^d} \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m (\tilde{f} - p_m) \right\|_{W^{k,p}(\Omega_{\tilde{m}, N})}^p, \end{aligned} \quad (\text{C.22})$$

where the last step follows from $(0, 1)^d \subset \bigcup_{\tilde{m} \in \{0, \dots, N\}^d} \Omega_{\tilde{m}, N}$. Now we obtain for each $\tilde{m} \in \{0, \dots, N\}^d$

$$\begin{aligned} \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m (\tilde{f} - p_m) \right\|_{W^{k,p}(\Omega_{\tilde{m}, N})} &\leq \sum_{\substack{m \in \{0, \dots, N\}^d, \\ \|m - \tilde{m}\|_{\ell^\infty} \leq 1}} \|\phi_m (\tilde{f} - p_m)\|_{W^{k,p}(\Omega_{\tilde{m}, N})} \\ &\leq \sum_{\substack{m \in \{0, \dots, N\}^d, \\ \|m - \tilde{m}\|_{\ell^\infty} \leq 1}} \|\phi_m (\tilde{f} - p_m)\|_{W^{k,p}(\Omega_{m, N})} \\ &\leq C_6 \left(\frac{1}{N}\right)^{n-k} \sum_{\substack{m \in \{0, \dots, N\}^d, \\ \|m - \tilde{m}\|_{\ell^\infty} \leq 1}} \|\tilde{f}\|_{W^{n,p}(\Omega_{m, N})} \end{aligned} \quad (\text{C.23})$$

where the triangle inequality together with the support property (iii) from Lemma C.3 is used in the first step. The second step follows again from Lemma C.3 (iii) and the third step follows from (C.18) for $k = 0$ and from (C.20) for $k = 1$. Here $C_6 = C_6(n, d, p) > 0$ can be chosen independent of k (e.g. $C_6 := \max\{C_2, C_5\}$).

Finally, the boundedness claim in (i) and (ii) follows from using the definition of f_N and combining Equation (C.22) with Equation (C.23):

$$\begin{aligned} \|f - f_N\|_{W^{k,p}((0,1)^d)}^p &\leq \sum_{\tilde{m} \in \{0, \dots, N\}^d} C_6^p \left(\frac{1}{N}\right)^{(n-k)p} \left(\sum_{\substack{m \in \{0, \dots, N\}^d, \\ \|m - \tilde{m}\|_{\ell^\infty} \leq 1}} \|\tilde{f}\|_{W^{n,p}(\Omega_{m, N})} \right)^p \\ &\stackrel{\text{(Hölder's inequality; } q := 1 - 1/p)}{\leq} C_6^p \left(\frac{1}{N}\right)^{(n-k)p} \sum_{\tilde{m} \in \{0, \dots, N\}^d} \sum_{\substack{m \in \{0, \dots, N\}^d, \\ \|m - \tilde{m}\|_{\ell^\infty} \leq 1}} \|\tilde{f}\|_{W^{n,p}(\Omega_{m, N})}^p (3^d)^{p/q} \\ &\leq C_6^p 3^{dp/q} \left(\frac{1}{N}\right)^{(n-k)p} 3^d \sum_{\tilde{m} \in \{0, \dots, N\}^d} \|\tilde{f}\|_{W^{n,p}(\Omega_{\tilde{m}, N})}^p \\ &\leq C_6^p 3^{dp/q} \left(\frac{1}{N}\right)^{(n-k)p} 3^d 2^d \|\tilde{f}\|_{W^{n,p}(\bigcup_{\tilde{m} \in \{0, \dots, N\}^d} \Omega_{\tilde{m}, N})}^p, \end{aligned}$$

where the last two steps follow from the definition of $\Omega_{\tilde{m}, N}$. Thus, we have

$$\|f - f_N\|_{W^{k,p}((0,1)^d)} \leq C_7 \left(\frac{1}{N}\right)^{n-k} \|\tilde{f}\|_{W^{n,p}(\mathbb{R}^d)} \leq C_8 \left(\frac{1}{N}\right)^{n-k} \|f\|_{W^{n,p}((0,1)^d)}$$

for $k \in \{0, 1\}$, where Equation (C.14) was used in the first and second step. Here $C_7 = C_7(n, d, p) > 0$ and $C_8 = C_8(n, d, p)$ are constants. The linearity of T_k , $k \in \{0, 1\}$ is a consequence of the linearity of the averaged Taylor polynomial (cf. Rem. B.8).

Step 4 (Interpolation): For $0 < s < 1$ we use a Banach space interpolation argument. Set

$$A_0 = A_1 = W^{n,p}((0,1)^d)$$

together with

$$B_0 = L^p((0,1)^d) \quad \text{and} \quad B_2 = W^{1,p}((0,1)^d).$$

Then, we can apply Theorem A.4 to the operator T with $Tf := f - f_N$ and get for a constant $C = C(n, d, p) > 0$ that

$$\begin{aligned} \|T\|_{\mathcal{L}(W^{n,p}, W^{s,p})} &\leq \|T\|_{\mathcal{L}(W^{n,p}, L^p)}^{1-s} \|T\|_{\mathcal{L}(W^{n,p}, W^{1,p})}^s \\ &\leq C \cdot \left(\frac{1}{N}\right)^{n(1-s)} \left(\frac{1}{N}\right)^{(n-1)s} = C \cdot \left(\frac{1}{N}\right)^{n-s}, \end{aligned}$$

where we used Lemma A.3 (ii) to see that $(A_0, A_1)_{s,p} = W^{n,p}((0,1)^d)$. \square

The following technical lemma lays the foundation for approximating localized monomials with neural networks. Using the notation and statement (v) from Lemma C.3, a localized monomial $\phi_m x^\alpha$ can be expressed as the product of the output components of a network $\Phi_{(m,\alpha)}$ with a suitable output dimension n , i.e.

$$\phi_m(x) x^\alpha = \prod_{l=1}^n [R_\varrho(\Phi_{(m,\alpha)})]_l(x).$$

Given a network Φ with n -dimensional output, we construct a network Ψ_Φ that approximates the product of the output components of Φ .

Lemma C.5. *Let $d, m, K \in \mathbb{N}$ and $N \geq 1$ be arbitrary. Then there is a constant $C = C(m) > 0$ such that the following holds:*

For any $\varepsilon \in (0, 1/2)$, and any neural network Φ with d -dimensional input and n -dimensional output where $n \leq m$, and with number of layers, neurons and weights all bounded by K , such that

$$\|[R_\varrho(\Phi)]_l\|_{W^{k,\infty}((0,1)^d)} \leq N^k \quad \text{for } k \in \{0, 1\} \text{ and } l = 1, \dots, n$$

there exists a neural network $\Psi_{\varepsilon,\Phi}$ with d -dimensional input and one-dimensional output, and with number of layers, neurons and weights all bounded by $K C \log_2(1/\varepsilon)$, such that

$$\left\| R_\varrho(\Psi_{\varepsilon,\Phi}) - \prod_{l=1}^n [R_\varrho(\Phi)]_l \right\|_{W^{k,\infty}((0,1)^d)} \leq c N^k \varepsilon \quad (\text{C.24})$$

for $k \in \{0, 1\}$ and some constant $c = c(d, m, k)$. Moreover,

$$R_\varrho(\Psi_{\varepsilon,\Phi})(x) = 0 \quad \text{if} \quad \prod_{l=1}^n [R_\varrho(\Phi)]_l(x) = 0 \quad (\text{C.25})$$

for $x \in (0, 1)^d$.

Proof. Let $d, K \in \mathbb{N}$ and $N \geq 1$. We show by induction over $m \in \mathbb{N}$ that the statement holds. To make the induction argument easier we will additionally show that $c = c(d, m, k) = m^{1-k} c_1^k$, where $c_1 = c_1(d, m) > 0$, and that the network $\Psi_{\varepsilon,\Phi}$ can be chosen such that the first $L(\Phi) - 1$ layers of $\Psi_{\varepsilon,\Phi}$ and Φ coincide and $\|R_\varrho(\Psi_{\varepsilon,\Phi})\|_{W^{1,\infty}((0,1)^d)} \leq C_1 N$ for a constant $C_1 = C_1(d, m) > 0$.

If $m = 1$, then we can choose $\Psi_{\varepsilon,\Phi} = \Phi$ for any $\varepsilon \in (0, 1/2)$ and the claim holds.

Now, assume that the claim holds for some $m \in \mathbb{N}$. We show that it also holds for $m + 1$. For this, let $\varepsilon \in (0, 1/2)$ and let $\Phi = ((A_1, b_1), (A_2, b_2), \dots, (A_L, b_L))$ be a neural network with d -dimensional input and n -dimensional output, where $n \leq m + 1$, and with number of layers, neurons and weights all bounded by K , where each A_l is an $N_l \times \sum_{k=0}^{l-1} N_k$ matrix, and $b_l \in \mathbb{R}^{N_l}$ for $l = 1, \dots, L$.

Case 1: If $n \leq m$, then we use the induction hypothesis and get that there is a constant $C_0 = C_0(m) > 0$ and a neural network $\Psi_{\varepsilon, \Phi}$ with d -dimensional input and one-dimensional output, and at most $KC_0 \log_2(1/\varepsilon)$ layers, neurons and weights such that

$$\left\| R_{\varrho}(\Psi_{\varepsilon, \Phi}) - \prod_{l=1}^n [R_{\varrho}(\Phi)]_l \right\|_{W^{k, \infty}((0,1)^d)} \leq m^{1-k} c_1^k N^k \varepsilon \leq (m+1)^{1-k} c_1^k N^k \varepsilon$$

for $k \in \{0, 1\}$ and $c_1 = c_1(d, m)$. Moreover,

$$R_{\varrho}(\Psi_{\varepsilon, \Phi})(x) = 0 \quad \text{if} \quad \prod_{l=1}^n [R_{\varrho}(\Phi)]_l(x) = 0,$$

for any $x \in (0, 1)^d$. Furthermore, we have $|R_{\varrho}(\Psi_{\varepsilon, \Phi})|_{W^{1, \infty}((0,1)^d)} \leq C_1 N$, for $C_1 = C_1(d, m)$.

Case 2: Now, we assume that $n = m + 1$ and show the claim for constants \tilde{C}_0, \tilde{c}_1 and \tilde{C}_1 depending on $m + 1$, possibly different from the constants C_0, c_1 and C_1 from Case 1, respectively. The maximum of each pair of constants fulfills then the claim for networks with output dimension $n \leq m + 1$.

We denote by Φ_m the neural network with d -dimensional input and m -dimensional output which results from Φ by removing the last output neuron and corresponding weights. In detail, we write

$$A_L = \begin{bmatrix} A_L^{(1,m)} \\ a_L^{(m+1)} \end{bmatrix} \quad \text{and} \quad b_L = \begin{bmatrix} b_L^{(1,m)} \\ b_L^{(m+1)} \end{bmatrix},$$

where $A_L^{(1,m)}$ is a $m \times \sum_{k=0}^{L-1} N_k$ matrix and $a_L^{(m+1)}$ is a $1 \times \sum_{k=0}^{L-1} N_k$ vector, and $b_L^{(1,m)} \in \mathbb{R}^m$ and $b_L^{(m+1)} \in \mathbb{R}^1$. Now we set

$$\Phi_m := \left((A_1, b_1), (A_2, b_2), \dots, (A_{L-1}, b_{L-1}), (A_L^{(1,m)}, b_L^{(1,m)}) \right).$$

Using the induction hypothesis and the constants C_0, c_1 and C_1 from Case 1, we get that there is a neural network $\Psi_{\varepsilon, \Phi_m} = ((A'_1, b'_1), (A'_2, b'_2), \dots, (A'_{L'}, b'_{L'}))$ with d -dimensional input and one-dimensional output, and at most $KC_0 \log_2(1/\varepsilon)$ layers, neurons and weights such that

$$\left\| R_{\varrho}(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_{\varrho}(\Phi_m)]_l \right\|_{W^{k, \infty}((0,1)^d)} \leq m^{1-k} c_1^k N^k \varepsilon \quad (\text{C.26})$$

for $k \in \{0, 1\}$. Moreover,

$$R_{\varrho}(\Psi_{\varepsilon, \Phi_m})(x) = 0 \quad \text{if} \quad \prod_{l=1}^m [R_{\varrho}(\Phi_m)]_l(x) = 0, \quad (\text{C.27})$$

for any $x \in (0, 1)^d$. Furthermore, we can assume that $|R_{\varrho}(\Psi_{\varepsilon, \Phi_m})|_{W^{1, \infty}((0,1)^d)} \leq C_1 N$, and that the first $L(\Phi) - 1$ layers of $\Psi_{\varepsilon, \Phi_m}$ and Φ_m coincide and, thus, also the first $L(\Phi) - 1$ layers of $\Psi_{\varepsilon, \Phi_m}$ and Φ , i.e. $A_l = A'_l$ for $l = 1, \dots, L(\Phi) - 1$.

Now, we add the formerly removed neuron with corresponding weights back to the last layer of $\Psi_{\varepsilon, \Phi_m}$. For the resulting network

$$\tilde{\Psi}_{\varepsilon, \Phi} := \left((A'_1, b'_1), (A'_2, b'_2), \dots, (A'_{L'-1}, b'_{L'-1}), \left(\begin{bmatrix} A'_{L'} & \\ a_L^{(m+1)} & 0_{\mathbb{R}^1, \sum_{k=L}^{L'} N'_k} \end{bmatrix}, \begin{bmatrix} b'_{L'} \\ b_L^{(m+1)} \end{bmatrix} \right) \right)$$

it holds that the first $L - 1$ layers of $\tilde{\Psi}_{\varepsilon, \Phi}$ and Φ coincide, and $\tilde{\Psi}_{\varepsilon, \Phi}$ is a neural network with two-dimensional output. Note that

$$\begin{aligned} & \left\| [R_{\varrho}(\tilde{\Psi}_{\varepsilon, \Phi})]_1 \right\|_{L^\infty((0,1)^d)} \\ &= \|R_{\varrho}(\Psi_{\varepsilon, \Phi_m})\|_{L^\infty((0,1)^d)} \\ &\leq \left\| R_{\varrho}(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_{\varrho}(\Phi_m)]_l \right\|_{L^\infty((0,1)^d)} + \left\| \prod_{l=1}^m [R_{\varrho}(\Phi_m)]_l \right\|_{L^\infty((0,1)^d)} \\ &\leq m\varepsilon + 1 < m + 1, \end{aligned}$$

where we used Equation (C.26) for $k = 0$. Additionally, we have

$$\left\| [R_{\varrho}(\tilde{\Psi}_{\varepsilon, \Phi})]_2 \right\|_{L^\infty((0,1)^d)} = \|[R_{\varrho}(\Phi)]_{m+1}\|_{L^\infty((0,1)^d)} \leq 1.$$

Now, we denote by $\tilde{\times}$ the network from Proposition C.2 with $M = m + 1$ and accuracy ε and define

$$\Psi_{\varepsilon, \Phi} := \tilde{\times} \odot \tilde{\Psi}_{\varepsilon, \Phi}.$$

Consequently, $\Psi_{\varepsilon, \Phi}$ has d -dimensional input, one-dimensional output and, combining the induction hypothesis with statement (iv) of Proposition C.2 and Remark 2.9, at most

$$2KC_0 \log_2(1/\varepsilon) + 2(c' \log_2(1/\varepsilon) + c'') \leq KC \log_2(1/\varepsilon)$$

layers, number of neurons and weights. Here c' and $c'' = c''(m+1)$ are the constants from Proposition C.2 (iv) for the choice $M = m + 1$ and $C = C(m + 1) > 0$ is a suitable constant. Clearly, the first $L - 1$ layers of $\Psi_{\varepsilon, \Phi}$ and Φ coincide and for the approximation properties it holds that

$$\begin{aligned} & \left| R_{\varrho}(\Psi_{\varepsilon, \Phi}) - \prod_{l=1}^{m+1} [R_{\varrho}(\Phi)]_l \right|_{W^{k, \infty}((0,1)^d)} \\ &= \left| R_{\varrho}(\tilde{\times}) \circ R_{\varrho}(\tilde{\Psi}_{\varepsilon, \Phi}) - [R_{\varrho}(\Phi)]_{m+1} \cdot \prod_{l=1}^m [R_{\varrho}(\Phi)]_l \right|_{W^{k, \infty}((0,1)^d)} \\ &\leq \left| R_{\varrho}(\tilde{\times}) \circ (R_{\varrho}(\Psi_{\varepsilon, \Phi_m}), [R_{\varrho}(\Phi)]_{m+1}) - R_{\varrho}(\Psi_{\varepsilon, \Phi_m}) \cdot [R_{\varrho}(\Phi)]_{m+1} \right|_{W^{k, \infty}((0,1)^d)} \\ &\quad + \left| [R_{\varrho}(\Phi)]_{m+1} \cdot \left(R_{\varrho}(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_{\varrho}(\Phi)]_l \right) \right|_{W^{k, \infty}((0,1)^d)}, \end{aligned} \tag{C.28}$$

for $k \in \{0, 1\}$. We continue by considering the first term of the Inequality (C.28) for $k = 0$ and obtain

$$\begin{aligned} & \left\| R_{\varrho}(\tilde{\times}) \circ (R_{\varrho}(\Psi_{\varepsilon, \Phi_m}), [R_{\varrho}(\Phi)]_{m+1}) - R_{\varrho}(\Psi_{\varepsilon, \Phi_m}) \cdot [R_{\varrho}(\Phi)]_{m+1} \right\|_{L^\infty((0,1)^d)} \\ &\leq \|R_{\varrho}(\tilde{\times})(x, y) - x \cdot y\|_{L^\infty((- (m+1), m+1)^2; dx dy)} \leq \varepsilon, \end{aligned} \tag{C.29}$$

where we used Proposition C.2 (i) for the last step. Next, we consider the same term for $k = 1$ and apply the chain rule from Corollary B.5. For this, let $\hat{C} = \hat{C}(d)$ be the constant from Corollary B.5 (for $n = d$ and $m = 2$). We get

$$\begin{aligned} & \left| R_{\varrho}(\tilde{\times}) \circ (R_{\varrho}(\Psi_{\varepsilon, \Phi_m}), [R_{\varrho}(\Phi)]_{m+1}) - R_{\varrho}(\Psi_{\varepsilon, \Phi_m}) \cdot [R_{\varrho}(\Phi)]_{m+1} \right|_{W^{1, \infty}((0,1)^d)} \\ &\leq \hat{C} \cdot |R_{\varrho}(\tilde{\times})(x, y) - x \cdot y|_{W^{1, \infty}((- (m+1), m+1)^2; dx dy)} \left| R_{\varrho}(\tilde{\Psi}_{\varepsilon, \Phi}) \right|_{W^{1, \infty}((0,1)^d; \mathbb{R}^2)} \\ &\leq \hat{C} \cdot \varepsilon \max\{C_1 N, N\} = C'_1 \varepsilon N, \end{aligned} \tag{C.30}$$

where we used the induction hypothesis together with $\|[R_{\varrho}(\Phi)]_{m+1}\|_{W^{1, \infty}((0,1)^d)} \leq N$ in the third step, and $C'_1 = C'_1(d, m + 1) > 0$ is a suitable constant.

To estimate the second term of (C.28) for $k = 0$ we use the induction hypothesis (for $k = 0$) and get

$$\begin{aligned} & \left\| [R_\varrho(\Phi)]_{m+1} \cdot \left(R_\varrho(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_\varrho(\Phi)]_l \right) \right\|_{L^\infty((0,1)^d)} \\ & \leq \| [R_\varrho(\Phi)]_{m+1} \|_{L^\infty((0,1)^d)} \cdot \left\| R_\varrho(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_\varrho(\Phi)]_l \right\|_{L^\infty((0,1)^d)} \leq 1 \cdot m \cdot \varepsilon. \end{aligned} \quad (\text{C.31})$$

For $k = 1$ we apply the product rule from Lemma B.6 together with $\| [R_\varrho(\Phi)]_{m+1} \|_{L^\infty} \leq 1$ and get

$$\begin{aligned} & \left| [R_\varrho(\Phi)]_{m+1} \cdot \left(R_\varrho(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_\varrho(\Phi)]_l \right) \right|_{W^{1,\infty}((0,1)^d)} \\ & \leq \| [R_\varrho(\Phi)]_{m+1} \|_{W^{1,\infty}((0,1)^d)} \cdot \left\| R_\varrho(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_\varrho(\Phi)]_l \right\|_{L^\infty((0,1)^d)} \\ & \quad + \| [R_\varrho(\Phi)]_{m+1} \|_{L^\infty((0,1)^d)} \cdot \left| R_\varrho(\Psi_{\varepsilon, \Phi_m}) - \prod_{l=1}^m [R_\varrho(\Phi)]_l \right|_{W^{1,\infty}((0,1)^d)} \\ & \leq N \cdot m\varepsilon + 1 \cdot c_1 N\varepsilon = c'_1 N\varepsilon, \end{aligned} \quad (\text{C.32})$$

where we used the induction hypothesis for $k = 1$, and $c'_1 = c'_1(d, m+1) > 0$.

Combining (C.28) with (C.29) and (C.31) we have

$$\left\| R_\varrho(\Psi_{\varepsilon, \Phi}) - \prod_{l=1}^{m+1} [R_\varrho(\Phi)]_l \right\|_{L^\infty((0,1)^d)} \leq \varepsilon + m \cdot \varepsilon = (m+1) \cdot \varepsilon, \quad (\text{C.33})$$

and in the same way a combination of (C.28) with (C.30) and (C.32) yields

$$\left| R_\varrho(\Psi_{\varepsilon, \Phi}) - \prod_{l=1}^{m+1} [R_\varrho(\Phi)]_l \right|_{W^{1,\infty}((0,1)^d)} \leq (C'_1 + c'_1) \cdot N \cdot \varepsilon = c''_1 N\varepsilon,$$

where $c''_1 = c''_1(d, m+1) > 0$. Putting together the two previous estimates yields

$$\left\| R_\varrho(\Psi_{\varepsilon, \Phi}) - \prod_{l=1}^{m+1} [R_\varrho(\Phi)]_l \right\|_{W^{1,\infty}((0,1)^d)} \leq c'''_1 N\varepsilon, \quad (\text{C.34})$$

for a suitable constant $c'''_1 = c'''_1(d, m+1) > 0$.

We now show Equation (C.25) for $m+1$. To this end, assume that $[R_\varrho(\Phi)]_l(x) = 0$ for some $l \in \{1, \dots, m+1\}$ and $x \in (0,1)^d$. If $l \leq m$, then Equation (C.27) implies that

$$[R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi})]_1(x) = R_\varrho(\Psi_{\varepsilon, \Phi_m})(x) = 0.$$

If $l = m+1$, then we have

$$[R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi})]_2(x) = [R_\varrho(\Phi)(x)]_{m+1}(x) = 0.$$

Hence, by application of Proposition C.2, we have

$$R_\varrho(\Psi_{\varepsilon, \Phi})(x) = R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi})(x) \left([R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi})]_1(x), [R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi})]_2(x) \right) = 0.$$

Finally, we need to show that there is a constant $C''_1 = C''_1(d, m+1) > 0$ such that

$$\| R_\varrho(\Psi_{\varepsilon, \Phi}) \|_{W^{1,\infty}((0,1)^d)} \leq C''_1 N.$$

Similarly as in (C.30) we have for a constant $C_1'' = C_1''(d, m+1) > 0$ that

$$\begin{aligned} |R_\varrho(\Psi_{\varepsilon, \Phi})|_{W^{1, \infty}((0,1)^d)} &= \left| R_\varrho(\tilde{\times}) \circ R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi}) \right|_{W^{1, \infty}((0,1)^d)} \\ &\leq \hat{C} \cdot |R_\varrho(\tilde{\times})|_{W^{1, \infty}((-m+1, m+1)^2)} \cdot \left| R_\varrho(\tilde{\Psi}_{\varepsilon, \Phi}) \right|_{W^{1, \infty}((0,1)^d; \mathbb{R}^2)} \\ &\leq \hat{C} \cdot \hat{c} \cdot (m+1) \cdot \max\{C_1 N, N\} = C_1'' \cdot N, \end{aligned}$$

where Corollary B.5 was used for the second step and \hat{c} is the constant from Proposition C.2 (iii) which together with an argument as in (C.30) implies the third step.

Taking the maximum of the each pair of constants derived in Case 1 and Case 2 concludes the proof. \square

In the next lemma, we approximate a sum of localized polynomials with a neural network. One of the difficulties is to control the derivative of the localizing functions from the partition of unity.

Lemma C.6. *Let $d, N \in \mathbb{N}$, $n \in \mathbb{N}_{\geq 2}$, $1 \leq p \leq \infty$, and $0 \leq s \leq 1$, and let $\Psi = \Psi(d, N) = \{\phi_m : m \in \{0, \dots, N\}^d\}$ be the partition of unity from Lemma C.3. Then, there are constants $C_1 = C_1(n, d, p, s) > 0$ and $C_2 = C_2(n, d), C_3 = C_3(n, d) > 0$ with the following properties:*

For any $\varepsilon \in (0, 1/2)$ there is a neural network architecture $\mathcal{A}_\varepsilon = \mathcal{A}_\varepsilon(d, n, N, \varepsilon)$ with d -dimensional input and one-dimensional output, with at most $C_2 \cdot \log_2(1/\varepsilon)$ layers and $C_3(N+1)^d \log_2(1/\varepsilon)$ weights and neurons, such that the following holds: Let $f \in W^{n, p}((0,1)^d)$ and $p_m(x) := p_{f, m}(x) = \sum_{|\alpha| \leq n-1} c_{m, \alpha} x^\alpha$ for $m \in \{0, \dots, N\}^d$ be the polynomials from Lemma C.4, then there is a neural network $\Phi_{P, \varepsilon}$ that has architecture \mathcal{A}_ε such that

$$\left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m - R_\varrho(\Phi_{P, \varepsilon}) \right\|_{W^{s, p}((0,1)^d)} \leq C_1 \|f\|_{W^{n, p}((0,1)^d)} N^s \varepsilon.$$

Proof. Step 1 (Approximating localized monomials $\phi_m x^\alpha$): Let $|\alpha| \leq n-1$ and $m \in \{0, \dots, N\}^d$. It is easy to see that there is a neural network Φ_α with d -dimensional input and $|\alpha|$ -dimensional output, with one layer and at most $n-1$ weights and $n-1+d$ neurons such that

$$x^\alpha = \prod_{l=1}^{|\alpha|} [R_\varrho(\Phi_\alpha)]_l(x) \quad \text{for all } x \in (0,1)^d$$

and

$$\|[R_\varrho(\Phi_\alpha)]_l\|_{W^{k, \infty}((0,1)^d)} \leq 1 \quad \text{for all } l = 1, \dots, |\alpha| \text{ and } k \in \{0, 1\}. \quad (\text{C.35})$$

Let now Φ_m be the neural network and $C, c \geq 1$ the constants from Lemma C.3 (v) and define the network

$$\Phi_{m, \alpha} := P(\Phi_m, \Phi_\alpha).$$

Then $\Phi_{m, \alpha}$ has at most $3 \leq K_0$ layers, $Cd + n - 1 \leq K_0$ nonzero weights, and $Cd + n - 1 + d \leq K_0$ neurons for a suitable constant $K_0 = K_0(n, d) \in \mathbb{N}$, and $\prod_{l=1}^{|\alpha|+d} [R_\varrho(\Phi_{m, \alpha})]_l(x) = \phi_m x^\alpha$ for all $x \in (0,1)^d$. Moreover, as a consequence of Lemma C.3 (v) together with Equation (C.35) we have

$$\|[R_\varrho(\Phi_{m, \alpha})]_l\|_{W^{k, \infty}((0,1)^d)} \leq (cN)^k \quad \text{for all } l = 1, \dots, |\alpha| + d \text{ and for } k \in \{0, 1\}.$$

To construct an approximation of the localized monomials $\phi_m x^\alpha$, let $\Psi_{\varepsilon, (m, \alpha)}$ be the neural network provided by Lemma C.5 (with $\Phi_{m, \alpha}$ instead of Φ , $m = n-1+d \in \mathbb{N}$, $K = K_0 \in \mathbb{N}$ and cN instead of N) for $m \in \{0, \dots, N\}^d$ and $\alpha \in \mathbb{N}_0^d$, $|\alpha| \leq n-1$. There exists a constant $C_1 = C_1(n, d) \geq 1$ such that $\Psi_{\varepsilon, (m, \alpha)}$ has at most $C_1 \log_2(1/\varepsilon)$ layers, number of neurons and weights. Moreover,

$$\|\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon, (m, \alpha)})(x)\|_{W^{k, \infty}((0,1)^d; dx)} \leq c' N^k \varepsilon \quad (\text{C.36})$$

for a constant $c' = c'(n, d, k) > 0$ and $k \in \{0, 1\}$, and

$$R_\varrho(\Psi_{\varepsilon, (m, \alpha)})(x) = 0 \quad \text{if } \phi_m(x) x^\alpha = 0 \quad \text{for all } x \in (0,1)^d. \quad (\text{C.37})$$

Step 2 (Constructing an architecture capable of approximating sums of localized polynomials): We set

$$M := |\{(m, \alpha) : m \in \{0, \dots, N\}^d, \alpha \in \mathbb{N}_0^d, |\alpha| \leq n-1\}|$$

and define the matrix $A_{\text{sum}} \in \mathbb{R}^{1 \times M}$ by $A_{\text{sum}} := [c_{m,\alpha} : m \in \{0, \dots, N\}^d, \alpha \in \mathbb{N}_0^d, |\alpha| \leq n-1]$ and the neural network $\Phi_{\text{sum}} := ((A_{\text{sum}}, 0))$. Finally, we set

$$\Phi_{P,\varepsilon} := \Phi_{\text{sum}} \odot P(\Psi_{\varepsilon,(m,\alpha)} : m \in \{0, \dots, N\}^d, \alpha \in \mathbb{N}_0^d, |\alpha| \leq n-1).$$

Then, there are constants $C_2 = C_2(n, d), C_3 = C_3(n, d) > 0$ such that $\Phi_{P,\varepsilon}$ is a neural network with d -dimensional input and one-dimensional output, with at most $1 + C_1 \log_2(1/\varepsilon) \leq C_2 \log_2(1/\varepsilon)$ layers,

$$2(M + MC_1 \log_2(1/\varepsilon)) \leq 4MC_1 \log_2(1/\varepsilon) \leq C_3(N+1)^d \log_2(1/\varepsilon)$$

nonzero weights and neurons, and

$$R_\varrho(\Phi_{P,\varepsilon}) = \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} c_{m,\alpha} \Psi_{\varepsilon,(m,\alpha)}.$$

Note that the network $\Phi_{P,\varepsilon}$ only depends on $p_{f,m}$ (and thus on f) via the coefficients $c_{m,\alpha}$. Now, it is easy to see that there exists a neural network architecture $\mathcal{A}_\varepsilon = \mathcal{A}_\varepsilon(d, n, N, \varepsilon)$ with $L(\mathcal{A}_\varepsilon) \leq C_2 \log_2(1/\varepsilon)$ layers and number of neurons and weights bounded by $C_3(N+1)^d \log_2(1/\varepsilon)$ such that $\Phi_{P,\varepsilon}$ has architecture \mathcal{A}_ε for every choice of coefficients $c_{m,\alpha}$, and hence for every choice of f .

Step 3 (Estimating the approximation error in $\|\cdot\|_{W^{k,p}}$, $k \in \{0, 1\}$): For each $m \in \{0, \dots, N\}^d$ we set

$$\Omega_{m,N} := B_{\frac{1}{N}, \|\cdot\|_\infty} \left(\frac{m}{N} \right)$$

and we get for $k \in \{0, 1\}$

$$\begin{aligned} & \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m(x) p_m(x) - R_\varrho(\Phi_{P,\varepsilon})(x) \right\|_{W^{k,p}((0,1)^d; dx)}^p \\ &= \left\| \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} c_{m,\alpha} \left(\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right) \right\|_{W^{k,p}((0,1)^d; dx)}^p \\ &\leq \sum_{\tilde{m} \in \{0, \dots, N\}^d} \left\| \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} c_{m,\alpha} \left(\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right) \right\|_{W^{k,p}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)}^p, \quad (\text{C.38}) \end{aligned}$$

where the last step is a consequence of $(0, 1)^d \subset \bigcup_{\tilde{m} \in \{0, \dots, N\}^d} \Omega_{\tilde{m},N}$. For $\tilde{m} \in \{0, \dots, N\}^d$ we have

$$\begin{aligned} & \left\| \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} c_{m,\alpha} \left(\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right) \right\|_{W^{k,p}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)} \\ &\leq \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} |c_{m,\alpha}| \left\| \phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right\|_{W^{k,p}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)} \\ &\leq c_1 \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} \|\tilde{f}\|_{W^{n-1,p}(\Omega_{m,N})} N^{d/p} \left\| \phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right\|_{W^{k,p}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)}, \quad (\text{C.39}) \end{aligned}$$

where we used the triangle inequality in the first step and Lemma C.4 in the second step. Here $c_1 = c_1(n, d) > 0$ is a constant. Next, note that

$$\begin{aligned} & \left\| \phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right\|_{W^{k,p}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)} \\ &\leq \lambda \left(\Omega_{\tilde{m},N} \cap (0, 1)^d \right)^{1/p} (d+1)^{1/p} \left\| \phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right\|_{W^{k,\infty}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)} \\ &\leq c_3 \left(\frac{1}{N} \right)^{d/p} \left\| \phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon,(m,\alpha)})(x) \right\|_{W^{k,\infty}(\Omega_{\tilde{m},N} \cap (0,1)^d; dx)}, \end{aligned}$$

where λ denotes the Lebesgue measure and $c_3 = c_3(d, p) > 0$ is a constant. Combining Equation (C.39) with the last estimate yields

$$\begin{aligned}
& \left\| \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} c_{m, \alpha} \left(\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon, (m, \alpha)})(x) \right) \right\|_{W^{k, p}(\Omega_{\tilde{m}, N} \cap (0, 1)^d; dx)} \\
& \leq c_4 \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})} \|\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon, (m, \alpha)})(x)\|_{W^{k, \infty}(\Omega_{\tilde{m}, N} \cap (0, 1)^d; dx)} \\
& \leq c_4 \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})} \|\phi_m(x) x^\alpha - R_\varrho(\Psi_{\varepsilon, (m, \alpha)})(x)\|_{W^{k, \infty}(\Omega_{\tilde{m}, N} \cap (0, 1)^d; dx)} \\
& \quad \|\tilde{m} - \tilde{m}\|_{\ell^\infty} \leq 1 \\
& \leq c_4 c' N^k \varepsilon \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})} \quad \|\tilde{m} - \tilde{m}\|_{\ell^\infty} \leq 1
\end{aligned} \tag{C.40}$$

where we used Lemma C.3 (iii) together with Equation (C.37) in the second step and Equation (C.36) in the last step. Here, $c_4 = c_4(n, d, p) > 0$ is a constant. Now, using that $|\{\alpha : \alpha \in \mathbb{N}_0^d, |\alpha| \leq n-1\}| \leq (n-1)^d$ and Hölder's inequality shows

$$\begin{aligned}
& \sum_{m \in \{0, \dots, N\}^d} \sum_{|\alpha| \leq n-1} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})} \leq (n-1)^d \sum_{m \in \{0, \dots, N\}^d} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})} \\
& \quad \|\tilde{m} - \tilde{m}\|_{\ell^\infty} \leq 1 \\
& \quad (\text{Hölder's inequality; } q := 1 - 1/p) \leq (n-1)^d 3^{d/q} \left(\sum_{m \in \{0, \dots, N\}^d} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})}^p \right)^{1/p} \quad \|\tilde{m} - \tilde{m}\|_{\ell^\infty} \leq 1
\end{aligned} \tag{C.41}$$

Combining Equation (C.40) with Equation (C.41) and plugging the result in Equation (C.38) finally yields

$$\begin{aligned}
& \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m(x) p_m(x) - R_\varrho(\Phi_{P, \varepsilon})(x) \right\|_{W^{k, p}((0, 1)^d; dx)}^p \\
& \leq \left(c_4 c' 3^{d/q} (n-1)^d N^k \varepsilon \right)^p \sum_{\tilde{m} \in \{0, \dots, N\}^d} \sum_{m \in \{0, \dots, N\}^d} \|\tilde{f}\|_{W^{n-1, p}(\Omega_{m, N})}^p \\
& \quad \|\tilde{m} - \tilde{m}\|_{\ell^\infty} \leq 1 \\
& \leq c_5 N^{kp} \varepsilon^p \|f\|_{W^{n, p}((0, 1)^d)}^p,
\end{aligned} \tag{C.42}$$

where the last step is the same as Step 3 of the proof of Lemma C.4 and $c_5 = c_5(n, d, p, k) > 0$ is a constant. Hence the case $s = 0$ and $s = 1$ is proven.

Step 4 (Interpolation): To show the general statement for $0 \leq s \leq 1$ we use the interpolation inequality from Corollary A.5 together with Equation (C.42) and directly get

$$\left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m(x) p_m(x) - R_\varrho(\Phi_{P, \varepsilon})(x) \right\|_{W^{s, p}((0, 1)^d; dx)} \leq c_6 N^s \varepsilon \|f\|_{W^{n, p}((0, 1)^d)}$$

for a constant $c_6 = c_6(n, p, d, s) > 0$. This concludes the proof of the lemma. \square

Finally, we are ready to proof the upper complexity bounds.

Proof of Theorem 4.1. The proof can be divided into two steps: First, we approximate the function f by a sum of localized polynomials and then approximate this sum by a network.

For the first step, we set

$$N := \left\lceil \left(\frac{\varepsilon}{2CB} \right)^{-1/(n-s)} \right\rceil, \quad (\text{C.43})$$

where $C = C(n, d, p) > 0$ is the constant from Corollary C.4. Without loss of generality we may assume that $CB \geq 1$. The same corollary yields that if $\Psi = \Psi(d, N) = \{\phi_m : m \in \{0, \dots, N\}^d\}$ is the partition of unity from Lemma C.3, then there exist polynomials $p_m(x) = \sum_{|\alpha| \leq n-1} c_{m,\alpha} x^\alpha$ for $m \in \{0, \dots, N\}^d$ such that

$$\begin{aligned} \left\| f - \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m \right\|_{W^{s,p}((0,1)^d)} &\leq CB \left(\frac{1}{N} \right)^{n-s} \\ &\leq CB \frac{\varepsilon}{2CB} = \frac{\varepsilon}{2}. \end{aligned} \quad (\text{C.44})$$

For the second step, let $C_1 = C_1(n, d, p, s) > 0$, $C_2 = C_2(n, d) > 0$ and $C_3 = C_3(n, d) > 0$ be the constants from Lemma C.6 and $\Phi_{P,\varepsilon}$ be the neural network provided by Lemma C.6 with $\varepsilon^{n/(n-s)}/(8CB^2C_1)$ instead of ε . Then $\Phi_{P,\varepsilon}$ has at most

$$C_2 \log_2 \left(8CB^2C_1 \varepsilon^{-n/(n-s)} \right) \leq C_2 \left(\log_2(8CB^2C_1) + \log_2 \left(\varepsilon^{-n/(n-s)} \right) \right) \leq C' \log_2 \left(\varepsilon^{-n/(n-s)} \right)$$

layers for a constant $C' = C'(n, d, p, B, s) > 0$ and at most

$$\begin{aligned} C_3 \left(\left(\frac{\varepsilon}{2CB} \right)^{-1/(n-s)} + 2 \right)^d \log_2 \left(8CB^2C_1 \varepsilon^{-n/(n-s)} \right) \\ \leq C_3 3^d \left(\frac{\varepsilon}{2CB} \right)^{-d/(n-s)} \log_2 \left(8CB^2C_1 \varepsilon^{-n/(n-s)} \right) \\ \leq C'' \varepsilon^{-d/(n-s)} \log_2 \left(\varepsilon^{-n/(n-s)} \right) \end{aligned}$$

nonzero weights and neurons. Here $C'' = C''(n, d, p, B, s)$ is a suitable constant and we used $(2CB)/\varepsilon \geq 1$ in the first step. It holds that the architecture of $\Phi_{P,\varepsilon}$ is independent of the function f . Furthermore, we have

$$\begin{aligned} \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m - R_\varrho(\Phi_{P,\varepsilon}) \right\|_{W^{s,p}((0,1)^d)} &\leq C_1 B N^s \frac{\varepsilon^{n/(n-s)}}{8CB^2C_1} \\ &\leq \left(\left(\frac{\varepsilon}{2CB} \right)^{-1/(n-s)} + 1 \right)^s \frac{\varepsilon^{n/(n-s)}}{8CB} \\ &\leq \left(\left(\frac{\varepsilon}{2CB} \right)^{-s/(n-s)} + 1 \right) \frac{\varepsilon^{n/(n-s)}}{8CB}, \end{aligned}$$

where we used the inequality $(x+y)^\sigma \leq x^\sigma + y^\sigma$ for $x, y \geq 0$ and $0 \leq \sigma \leq 1$ in the last step. We now continue the above computation using that $s/(n-s) \leq 1$ and $2CB \geq 1$ and therefore $(2CB)^{s/(n-s)} \leq 2CB$ and get

$$\begin{aligned} \left(\left(\frac{\varepsilon}{2CB} \right)^{-s/(n-s)} + 1 \right) \frac{\varepsilon^{n/(n-s)}}{8} &\leq \frac{1}{8CB} \left(2CB\varepsilon + \varepsilon^{n/(n-s)} \right) \\ &\leq \frac{1}{8CB} (2CB\varepsilon + 2CB\varepsilon) = \frac{\varepsilon}{2}. \end{aligned}$$

Combining the previous computations we get

$$\left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m - R_\varrho(\Phi_{P,\varepsilon}) \right\|_{W^{s,p}((0,1)^d)} \leq \frac{\varepsilon}{2}. \quad (\text{C.45})$$

Using the triangle inequality and Equations (C.44) and (C.45) we finally obtain

$$\begin{aligned} & \|f - R_\varrho(\Phi_{P,\varepsilon})\|_{W^{s,p}((0,1)^d)} \\ & \leq \left\| f - \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m \right\|_{W^{s,p}((0,1)^d)} + \left\| \sum_{m \in \{0, \dots, N\}^d} \phi_m p_m - R_\varrho(\Phi_{P,\varepsilon}) \right\|_{W^{s,p}((0,1)^d)} \\ & \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon, \end{aligned}$$

which concludes the proof. \square

D Lower bounds for approximations

Lower complexity bounds for approximations in L^∞ norm, which correspond to the case $k = 0$ in Theorem 4.3, have already been shown in Theorem 4 a) in [60]. For lower complexity bounds of $W^{1,\infty}$ approximations we modify the proof strategy outlined in [60, Thm. 4 a)].

We start by showing an auxiliary result, that is used in the proof of Proposition D.2.

Lemma D.1. *Let $d \in \mathbb{N}$ and Φ be a neural network with d -dimensional input and one-dimensional output. Moreover, let $x \in (0, 1)^d$ and $\nu \in \mathbb{R}^d$. Then, there exists an open set $T = T(x, \nu) \subset (0, 1)^d$ and $\delta = \delta(x, \nu, T) > 0$ with $x + \lambda \delta \nu \in \bar{T}$ for $0 \leq \lambda \leq 1$ and $R_\varrho(\Phi)$ is affine-linear on T .*

Proof. We start by defining the set $U := \{x \in \mathbb{R}^d : R_\varrho(\Phi) \text{ is affine-linear on a neighborhood of } x\}$. Standard results on the number of pieces of ReLU neural networks [45] yield that U has only finitely many polyhedral, connected components, $(V_i)_{i=1}^k$ for some $k \in \mathbb{N}$, with $U = \bigcup_{i=1}^k V_i$ and $\mathbb{R}^d = \bigcup_{i=1}^k \bar{V}_i$. Note that it follows from the definition of U that V_i is open for $i = 1, \dots, k$.

Now, set $x_n := x + (1/n)\nu$. By the pigeonhole principle, there exists $q \in \{1, \dots, k\}$, such that \bar{V}_q contains infinitely many x_n . It is not hard to see that if a closed polyhedron contains a converging sequence on a line, then it also contains a small section of the line including the limit point of the sequence. Thus, there exists $\delta > 0$ such that $\{x + \lambda \delta \nu : 0 \leq \lambda \leq 1\} \subset \bar{V}_q \cap (0, 1)^d \subset V_q \cap (0, 1)^d$. Then, setting $T := V_q \cap (0, 1)^d$ shows the claim. \square

As in [60], we make use of a combinatorial quantity measuring the expressiveness of a set of binary valued functions H defined on some set X , called *VC-dimension* (see e.g. [3, Chap. 3.3]). We define

$$\text{VCdim}(H) := \sup \left\{ m \in \mathbb{N} : \begin{array}{l} \text{there exist } x_1, \dots, x_m \in X \text{ such that for any } y \in \{0, 1\}^m \\ \text{there is a function } h \in H \text{ with } h(x_i) = y_i \text{ for } i = 1, \dots, m \end{array} \right\}.$$

The idea of the next proposition is to relate the approximation error ε with the number of weights $M(\mathcal{A}_\varepsilon)$ of an architecture \mathcal{A}_ε capable of realizing such an approximation. To this end, we construct a set of functions H parameterized by elements of $\mathbb{R}^{M(\mathcal{A}_\varepsilon)+1}$.

If $w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)}$ and $\delta > 0$ is chosen appropriately, then a directional derivative of the function realized by $\mathcal{A}_\varepsilon(w)$ is computed for the evaluation of $h((w, \delta), \cdot) \in H$. By exploiting the approximation capacity of derivatives of functions realized by $\mathcal{A}_\varepsilon(w)$ we can find a lower bound for $\text{VCdim}(H)$ depending on ε (Claim 1).

On the other hand, [3, Thm. 8.4] yields an upper bound of the VC-dimension of H in terms of the number of computations and the dimension of the parameterization of H which can be expressed as a function of $M(\mathcal{A}_\varepsilon)$ (Claim 2). Together this gives the desired relation.

Proposition D.2. *Let $d \in \mathbb{N}$, $n \in \mathbb{N}_{\geq 2}$ and $B > 0$. Then, there are constants $c = c(n, B) > 0$ and $C = C(d)$ with the following property:*

Let $N \in \mathbb{N}$, $0 < \varepsilon \leq cN^{-(n-1)}$ and $\mathcal{A}_\varepsilon = \mathcal{A}_\varepsilon(d, n, \varepsilon)$ be a neural network architecture with d -dimensional input and one-dimensional output such that for any $f \in \mathcal{F}_{n,d,\infty,B}$ there is a neural network Φ_ε^f that has architecture \mathcal{A}_ε and

$$\|R_\varrho(\Phi_\varepsilon^f) - f\|_{W^{1,\infty}((0,1)^d)} \leq \varepsilon, \quad (\text{D.1})$$

then

$$N^d \leq C \cdot M(\mathcal{A}_\varepsilon)^2.$$

Proof. We prove the proposition by showing that there exists a function $h : \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1} \times [0, 1]^d \rightarrow \{0, 1\}$ with

$$N^d \leq \text{VCdim} \left(\left\{ x \mapsto h(w, x) : w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1} \right\} \right) \leq C \cdot M(\mathcal{A}_\varepsilon)^2.$$

To simplify the notation, we set

$$H := \left\{ x \mapsto h(w, x) : w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1} \right\}.$$

Step 1 (Construction of h): Let now $0 < \varepsilon < c_1 N^{-(n-1)}/(3\sqrt{d})$ for some constant $c_1 = c_1(n) > 0$ to be chosen later, and $\mathcal{A}_\varepsilon = \mathcal{A}_\varepsilon(d, n, \varepsilon)$ be a neural network architecture as in the claim of the proposition.

For $x \in [0, 1]^d$ we define a direction $\nu(x) \in \mathbb{R}^d$ that points from x into $(0, 1)^d$ if $x \in [0, 1]^d \setminus (0, 1)^d$ and equals e_1 if $x \in (0, 1)^d$. We set

$$\tilde{\nu}(x) := \begin{cases} e_1, & \text{if } 0 < x_k < 1 \text{ for } k = 1, \dots, d \\ \left[\chi_{\{0\}}(x_k) - \chi_{\{1\}}(x_k) : k = 1, \dots, d \right], & \text{else,} \end{cases}$$

and define $\nu(x) := \tilde{\nu}(x)/|\tilde{\nu}(x)|$. Moreover, we set $\tilde{x} := x + \nu(x)/(4N)$ for $x \in [0, 1]^d$.

To construct h we start by defining a function $g : \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1} \times [0, 1]^d \rightarrow \mathbb{R}$ and then, to get a binary valued function, define h by thresholding g . In detail, we set

$$g((w, \delta), x) := \begin{cases} \frac{1}{\delta} \cdot \left(R_\varrho(\mathcal{A}_\varepsilon(w))(\tilde{x} - \delta\nu(x)) - R_\varrho(\mathcal{A}_\varepsilon(w))(\tilde{x}) \right), & \text{if } \delta \neq 0 \\ 0, & \text{if } \delta = 0 \end{cases}$$

for $w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)}$, $\delta \in \mathbb{R}$ and $x \in [0, 1]^d$. Now, we define $h : \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1} \times [0, 1]^d \rightarrow \{0, 1\}$ by

$$h((w, \delta), x) := \begin{cases} 1, & \text{if } g((w, \delta), x) > cN^{-(n-1)}/2 \\ 0, & \text{else} \end{cases}$$

for $w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)}$, $\delta \in \mathbb{R}$ and $x \in [0, 1]^d$.

Claim 1 ($N^d \leq \text{VCdim}(H)$): Let $x_1, \dots, x_{N^d} \in [0, 1]^d$ such that $|x_m - x_n| \geq 1/N$ for all $m, n = 1, \dots, N^d$ with $m \neq n$ and such that $\tilde{x}_m \in (0, 1)^d$ for $m = 1, \dots, N^d$. Moreover, let $y_1, \dots, y_{N^d} \in \{0, 1\}$ be arbitrary. We aim to construct $w_y \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)}$ and $\delta_y \in \mathbb{R}$ with

$$h((w_y, \delta_y), x_m) = y_m \quad \text{for } m = 1, \dots, N^d.$$

To this end, we first define a function $f_y \in \mathcal{F}_{n,d,\infty,B}$ with $f_y(x_m) = y_m \cdot a$ for some constant $a > 0$, and then make use of a neural network Φ^{f_y} that approximates f_y .

Step 2 (Construction of f_y): We start by defining a bump function $\psi \in C^\infty(\mathbb{R}^d)$ by

$$\psi(x) := \begin{cases} e^{-(1-4|x|^2)^{-1}+1}, & \text{if } |x| < 1/2, \\ 0, & \text{else,} \end{cases}$$

such that $\psi(0) = 1$ and $\text{supp } \psi \subset B_{1/2,| \cdot |}(0)$. For the derivative $D\psi$ of ψ it holds that there exists a function $\phi : (-1/2, 1/2) \rightarrow \mathbb{R}_{>0}$ such that¹

$$(D\psi)(x) = \phi(|x|) \cdot (-x) \quad \text{for all } x \text{ with } |x| < 1/2.$$

Thus, if $x \in \mathbb{R}^d$ with $|x| < 1/2$, then we have for the derivative of ψ in direction $-x/|x|$ at x that $(D_{-x/|x|}\psi)(x) = \phi(|x|)|x| > 0$ only depends on the norm of x .

Next, we define $f_y \in C^\infty(\mathbb{R}^d)$ by

$$f_y(x) := \sum_{m=1}^{N^d} y_m \frac{BN^{-n}}{\|\psi\|_{W^{n,\infty}((0,1)^d)}} \psi(N(x - x_m))$$

¹Precisely, $\phi(r) = e^{-(1-4r^2)^{-1}+1} \cdot \frac{8}{(1-4r^2)^2}$.

for $x \in \mathbb{R}^d$. We have $|f_y|_{W^{k,\infty}((0,1)^d)} \leq BN^{-n}N^k \leq B$ for $1 \leq k \leq n$ and, consequently, $f_y \in \mathcal{F}_{n,d,\infty,B}$. Furthermore, for $x \in \mathbb{R}^d$ with $|x| < 1/(2N)$ it holds that

$$(D_{-x/|x|}f_y)(x_m + x) = y_m \frac{BN^{-n}}{\|\psi\|_{W^{n,\infty}((0,1)^d)}} \phi(|Nx|)N^2|x|$$

and, in particular, if $|x| = 1/(4N)$, then

$$(D_{-x/|x|}f_y)(x_m + x) = y_m \phi(1/4) \cdot \frac{BN^{-(n-1)}}{4\|\psi\|_{W^{n,\infty}((0,1)^d)}} = y_m c_1 N^{-(n-1)}. \quad (\text{D.2})$$

Here, we defined the constant $c_1 = c_1(n, B) > 0$ which was left unspecified in the beginning of the proof by $c_1 := \frac{B\phi(1/4)}{4\|\psi\|_{W^{n,\infty}((0,1)^d)}}$.

Step 3 (Existence of w_y and δ_y): We can find a vector $w_y \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)}$ such that for the neural network $\Phi_\varepsilon^{f_y} := \mathcal{A}_\varepsilon(w_y)$ Equation (D.1) holds (with f_y instead of f). In particular, we have

$$|R_\rho(\Phi_\varepsilon^{f_y}) - f_y|_{W^{1,\infty}((0,1)^d)} \leq \varepsilon. \quad (\text{D.3})$$

Next, for $x \in (0, 1)^d$ and $\nu \in \mathbb{R}^d$ we get from Lemma D.1 that there exists an open set $T_{x,\nu} \subset \mathbb{R}^d$ and $\delta = \delta(x, \nu, T_{x,\nu}) > 0$ with $x + \lambda\delta\nu \in \overline{T_{x,\nu}}$ for $0 \leq \lambda \leq 1$ and $R_\rho(\Phi_\varepsilon^{f_y})$ is affine-linear on $T_{x,\nu}$. We define

$$\delta_y := \min_{m=1,\dots,N^d} \delta(\tilde{x}_m, -\nu(x_m), T_{\tilde{x}_m, -\nu(x_m)}) > 0.$$

Let $m \in \{1, \dots, N^d\}$ and let $F_m : \mathbb{R}^d \rightarrow \mathbb{R}$ be an affine-linear function such that $R_\rho(\Phi_\varepsilon^{f_y})(x) = F_m(x)$ for all $x \in T_{\tilde{x}_m, -\nu(x_m)}$. It then follows from the continuity of $R_\rho(\Phi_\varepsilon^{f_y})$ that $R_\rho(\Phi_\varepsilon^{f_y})(x) = F_m(x)$ for all $x \in \overline{T_{\tilde{x}_m, -\nu(x_m)}}$. This, together with the choice of δ_y implies that

$$g((w_y, \delta_y), x_m) = D_{-\nu(x_m)}F_m(\tilde{x}_m) = D_{-\nu(x_m)}F_m.$$

Recall that $\overline{T_{\tilde{x}_m, -\nu(x_m)}}$ is open and f_y and $R_\rho(\Phi_\varepsilon^{f_y})$ are continuously differentiable on $T_{\tilde{x}_m, -\nu(x_m)}$. Thus, the weak and strong derivative agree and Equation (D.3) implies

$$|D^i f_y(x) - D^i F_m| = |D^i f_y(x) - D^i R_\rho(\Phi_\varepsilon^{f_y})(x)| \leq \varepsilon$$

for all $x \in T_{\tilde{x}_m, -\nu(x_m)}$ and $i = 1, \dots, d$. Using the continuity of $D^i f_y$ we get that $|D^i f_y(\tilde{x}_m) - D^i F_m| \leq \varepsilon$ for $i = 1, \dots, d$ and hence

$$|D_\nu f_y(\tilde{x}_m) - D_\nu F_m| \leq \sqrt{d}\varepsilon|\nu| \quad \text{for } \nu \in \mathbb{R}^d. \quad (\text{D.4})$$

An addition of zero yields

$$g((w_y, \delta_y), x_m) = D_{-\nu(x_m)}F_m = D_{-\nu(x_m)}f_y(\tilde{x}_m) + D_{-\nu(x_m)}F_m - D_{-\nu(x_m)}f_y(\tilde{x}_m). \quad (\text{D.5})$$

We get for the case $y_m = 1$ that

$$g((w_y, \delta_y), x_m) \geq D_{-\nu(x_m)}f_y(\tilde{x}_m) - \sqrt{d}\varepsilon \geq c_1 N^{-(n-1)} - c_1 N^{-(n-1)}/3, \quad (\text{D.6})$$

where we used Equation (D.5) together with Equation (D.4) and $|\nu(x_m)| = 1$ for the first step and Equation (D.2) together with the upper bound for ε for the second step.

In a similar way, we get for the case $y_m = 0$ that

$$g((w_y, \delta_y), x_m) \leq \sqrt{d}\varepsilon \leq c_1 N^{-(n-1)}/3, \quad (\text{D.7})$$

where we used that $D_{-\nu(x_m)}f_y(\tilde{x}_m) = 0$.

Finally, combining Equation (D.6) and Equation (D.7) reads as

$$g((w_y, \delta_y), x_m) \begin{cases} > c_1 N^{-(n-1)}/2, & \text{if } y_m = 1, \\ < c_1 N^{-(n-1)}/2, & \text{if } y_m = 0, \end{cases}$$

which proves Claim 1.

Claim 2 ($\text{VCdim}(H) \leq C \cdot M(\mathcal{A}_\varepsilon)^2$): We start by showing that there exists a constant $C' = C'(d)$ such that $h((w, \delta), x)$ can be computed using $C' \cdot M(\mathcal{A}_\varepsilon)$ operations of the following types

- the arithmetic operations $+$, $-$, \times , and $/$ on real numbers,
- jumps conditioned on $>$, \geq , $<$, \leq , $=$, and \neq comparisons of real numbers

for all $w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1}$ and $x \in [0, 1]^d$.

There exists an absolute constant $C_1 > 0$ such that at most $C_1 \cdot d$ operations of the specified type are needed to compute $\nu(x)$. Hence, the same holds true for \tilde{x} and $\tilde{x} - \delta\nu(x)$.

Note that the number of neurons that are needed for the computation of $R_\varrho(\mathcal{A}_\varepsilon(w))$ can be bounded by $M(\mathcal{A}_\varepsilon)$. Thus, $R_\varrho(\mathcal{A}_\varepsilon(w))$ can be computed using at most $C_2 \cdot M(\mathcal{A}_\varepsilon)$ operations where $C_2 > 0$ is an absolute constant. Hence, there exists a constant $C_3 = C_3(d)$ such that for the number of operations of the specified type t needed for the computation of $h(x)$ where $x \in [0, 1]^d$ it holds that $t \leq C_3 M(\mathcal{A}_\varepsilon)$.

Finally, [3, Thm. 8.4] implies

$$\text{VCdim} \left(\left\{ x \mapsto h(w, x) : w \in \mathbb{R}^{M(\mathcal{A}_\varepsilon)+1} \right\} \right) \leq 4(M(\mathcal{A}_\varepsilon) + 1)(C_3 \cdot M(\mathcal{A}_\varepsilon) + 2) \leq C_4 M(\mathcal{A}_\varepsilon)^2,$$

where $C_4 = C_4(d) > 0$ is a suitable constant. □

The proof of the lower complexity bounds is now a simple consequence of Proposition D.2.

Proof of Theorem 4.3. The case $k = 0$ corresponds to [60, Thm. 4 a)].

For the case $k = 1$, let $c = c(n, B)$ and $C = C(d)$ be the constants from Proposition D.2 and set

$$N := \left\lfloor \left(\frac{c}{\varepsilon} \right)^{1/(n-1)} \right\rfloor.$$

Then, $N \leq (c/\varepsilon)^{1/(n-1)}$ and, thus, $0 < \varepsilon \leq cN^{-(n-1)}$. Now, Proposition D.2 implies that

$$N^d \leq CM(\mathcal{A}_\varepsilon)^2. \tag{D.8}$$

We also have $(c/\varepsilon)^{1/(n-1)} \leq 2N$ and hence $c_1 \varepsilon^{-d/(n-1)} \leq N^d$ for a suitable constant $c_1 = c_1(n, B) > 0$. Combining this estimate with Equation D.8 yields

$$c_1 \varepsilon^{-d/(n-1)} \leq N^d \leq CM(\mathcal{A}_\varepsilon)^2,$$

which finally results in

$$C' \varepsilon^{-d/2(n-1)} \leq M(\mathcal{A}_\varepsilon)$$

for a constant $C' = C'(d, n, B) > 0$. □

References

- [1] R. Adams. *Sobolev Spaces*. Academic Press, New York, 1975.
- [2] H. Amann and E. Joachim. *Analysis III*. Grundlehrbuch der Mathematik. Birkhäuser, Basel, second edition, 2008.
- [3] M. Anthony and P. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, first edition, 2009.
- [4] R. Arora, A. Basu, P. Mianjy, and A. Mukherjee. Understanding deep neural networks with rectified linear units. In *Proceedings of the 2018 International Conference on Learning Representations (ICLR)*, 2018.
- [5] A. Barron. Approximation and estimation bounds for artificial neural networks. *Machine Learning*, 14(1):115–133, 1994.
- [6] C. Beck, S. Becker, P. Grohs, N. Jaafari, and A. Jentzen. Solving stochastic differential equations and Kolmogorov equations by means of deep learning. *arXiv preprint arXiv:1806.00421*, 2018.
- [7] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.

- [8] J. Bergh and J. Löfström. *Interpolation Spaces*, volume 223 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1976.
- [9] H. Bölcskei, P. Grohs, G. Kutyniok, and P. Petersen. Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*. In Press.
- [10] S. Brenner and R. Scott. *The Mathematical Theory of Finite Element Methods*, volume 15 of *Texts in Applied Mathematics*. Springer Science+Business Media, New York, third edition, 2008.
- [11] H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer Science+Business Media, New York, 2011.
- [12] C. Bucur and E. Valdinoci. *Nonlocal diffusion and applications*, volume 20 of *Lecture Notes of the Unione Matematica Italiana*. Springer International Publishing, 2016.
- [13] N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor analysis. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 698–728, 2016.
- [14] F. Cucker and D. Zhou. *Learning Theory: An Approximation Theory Viewpoint*, volume 24 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2007.
- [15] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, 1989.
- [16] W. Czarnecki, S. Osindero, M. Jaderberg, G. Swirszcz, and R. Pascanu. Sobolev training for neural networks. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, pages 4278–4287. 2017.
- [17] G. E. Dahl, D. Yu, L. Deng, and A. Acero. Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(1):30–42, 2012.
- [18] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [19] W. E, J. Han, and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- [20] W. E and B. Yu. The Deep Ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
- [21] D. Elbrächter, P. Grohs, A. Jentzen, and C. Schwab. DNN expression rate analysis of high-dimensional PDEs: Application to option pricing. *arXiv preprint arXiv:1809.07669*, 2018.
- [22] L. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 1999.
- [23] D. Gilbarg and N. Trudinger. *Elliptic Partial Differential Equations of Second Order*, volume 224 of *A Series of Comprehensive Studies in Mathematics*. Springer-Verlag, Berlin, second edition, 1998.
- [24] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 315–323, 2011.
- [25] P. Grisvard. *Elliptic Problems in Nonsmooth Domains*, volume 69 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 1985.
- [26] P. Grohs, D. Perekrestenko, D. Elbrächter, and H. Bölcskei. Deep neural network approximation theory. *arXiv preprint arXiv:1901.02220*, 2019.

- [27] I. Gühring. Error bounds for approximations with deep ReLU neural networks in general norms. Master thesis. https://www.math.tu-berlin.de/fileadmin/i26_fg-kutyniok/G%C3%BChring/publications/master_thesis.pdf, 2018. [Online; accessed 20-February-2019].
- [28] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [29] J. He, L. Li, J. Xu, and C. Zheng. ReLU deep neural networks and linear finite elements. *arXiv preprint arXiv:1807.03973*, 2018.
- [30] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [31] J. Heinonen. Lectures on Lipschitz analysis. <http://www.math.jyu.fi/research/reports/rep100.pdf>, 2005. [Online; accessed 20-September-2018].
- [32] G. Hinton, L. Deng, D. Yu, G. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- [33] K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991.
- [34] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2017.
- [35] A. Krizhevsky, I. Sutskever, and G. Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25 (NeurIPS 2012)*, pages 1097–1105, 2012.
- [36] G. Kutyniok, P. Petersen, M. Raslan, and R. Schneider. A theoretical analysis of deep neural networks and parametric PDEs. *arXiv preprint arXiv:1904.00377*, 2019.
- [37] I. Lagaris, A. Likas, and D. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks and Learning Systems*, 9(5):987–1000, 1998.
- [38] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [39] A. Lunardi. *Interpolation Theory*. Scuola Normale Superiore, Pisa, third edition, 2018.
- [40] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- [41] J. Marsden. *Elementary Classical Analysis*. W. H. Freeman and Company, San Francisco, 1974.
- [42] H. Mhaskar. Neural networks for optimal approximation of smooth and analytic functions. *Neural Computation*, 8(1):164–177, 1996.
- [43] H. Mhaskar, Q. Liao, and T. Poggio. When and why are deep networks better than shallow ones? In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 2343–2349, 2017.
- [44] H. Montanelli and Q. Du. New error bounds for deep ReLU networks using sparse grids. *arXiv preprint arXiv:1712.08688*, 2018.
- [45] G. Montúfar, R. Pascanu, K. Cho, and Y. Bengio. On the number of linear regions of deep neural networks. In *Advances in Neural Information Processing Systems 27 (NeurIPS 2014)*, pages 2924–2932, 2014.
- [46] P. Oswald. On the degree of nonlinear spline approximation in Besov-Sobolev spaces. *Journal of Approximation Theory*, 61(2):131–157, 1990.

- [47] P. Petersen, M. Raslan, and F. Voigtlaender. Topological properties of the set of functions generated by neural networks of fixed size. *arXiv preprint arXiv:1806.08459*, 2018.
- [48] P. Petersen and F. Voigtländer. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *Neural Networks*, 108:296–330, 2018.
- [49] T. Roubíček. *Nonlinear Partial Differential Equations with Applications*, volume 153 of *International Series of Numerical Mathematics*. Springer Science+Business Media, Basel, second edition, 2013.
- [50] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. Berg, and L. Fei-Fei. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [51] I. Safran and O. Shamir. Depth-width tradeoffs in approximating natural functions with neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2979–2987, 2017.
- [52] C. Schwab and J. Zech. Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in UQ. *Analysis and Applications*, 17(1):19–55, 2019.
- [53] U. Shaham, A. Cloninger, and R. Coifman. Provable approximation properties for deep neural networks. *Applied and Computational Harmonic Analysis*, 44(3):537–557, 2018.
- [54] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the 2015 International Conference on Learning Representations (ICLR)*, 2015.
- [55] J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- [56] E. Stein. *Singular Integrals and Differentiability Properties of Functions*. Princeton University Press, Princeton, third edition, 1979.
- [57] L. Tartar. *An Introduction to Sobolev Spaces and Interpolation Spaces*, volume 3 of *Lecture Notes of the Unione Matematica Italiana*. Springer-Verlag, Berlin, 2007.
- [58] H. Triebel. *Interpolation Theory, Function Spaces, Differential Operators*. North-Holland Publishing Company, Amsterdam, 1978.
- [59] C. Wu, P. Karanasou, M. J. Gales, and K. C. Sim. Stimulated deep neural network for speech recognition. Technical report, University of Cambridge, 2016.
- [60] D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017.
- [61] K. Yosida. *Functional Analysis*. Springer-Verlag, Berlin, sixth edition, 1980.
- [62] T. Young, D. Hazarika, S. Poria, and E. Cambria. Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3):55–75, 2018.
- [63] M. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *Computer Vision – ECCV 2014*, pages 818–833. Springer International Publishing, 2014.