

Constructive Deep ReLU Neural Network Approximation

L. Herrmann, J.A.A. Opschoor, C. Schwab

RICAM-Report 2021-03

Constructive Deep ReLU Neural Network Approximation

Lukas Herrmann[†], Joost A. A. Opschoor[‡], and Christoph Schwab[‡]

[†]Johann Radon Institute for Computational and Applied Mathematics, Austrian Academy of Sciences, Altenbergerstrasse 69, 4040 Linz, Austria.

lukas.herrmann@ricam.oeaw.ac.at

[‡]Seminar for Applied Mathematics, ETH Zürich, Rämistrasse 101, CH-8092 Zürich, Switzerland.
joost.opschoor@sam.math.ethz.ch, christoph.schwab@sam.math.ethz.ch

January 25, 2021

Abstract

We propose an efficient, *deterministic algorithm for constructing exponentially convergent deep neural network (DNN) approximations* of multivariate, analytic maps $f : [-1, 1]^K \rightarrow \mathbb{R}$. We address in particular networks with the rectified linear unit (ReLU) activation function. Similar results and proofs apply for many other popular activation functions. The algorithm is based on collocating f in deterministic families of grid points with small Lebesgue constants, and by a-priori (i.e., “offline”) emulation of a spectral basis with DNNs to prescribed fidelity.

Assuming availability of N function values of a possibly corrupted, numerical approximation \tilde{f} of f in $[-1, 1]^K$, a bound on $\|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$, and a given target accuracy $\varepsilon > 0$, we provide an explicit, computational construction of a deep ReLU NN which attains accuracy ε uniformly, with respect to the inputs. For analytic maps $f : [-1, 1]^K \rightarrow \mathbb{R}$, we prove exponential convergence of expression and generalization errors of the constructed ReLU DNNs with respect to ε , in the norm $W^{1, \infty}([-1, 1]^K; \mathbb{R})$, up to an additive data-corruption bound $\|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$ multiplied with a factor growing slowly with $1/\varepsilon$. The algorithmic construction of the ReLU DNNs which will realize the approximations, is explicit and deterministic in terms of the function values of \tilde{f} in tensorized Clenshaw–Curtis grids in $[-1, 1]^K$. We illustrate the proposed methodology by a constructive algorithm for (offline) computations of posterior expectations in Bayesian PDE inversion.

Key words: Deep ReLU neural networks, exponential convergence, neural network construction.

Subject Classification: 41A10, 41A50, 65D05, 65D15

1 Introduction

Approximation by deep neural networks (DNNs) receives increasing attention recently. DNNs realize mappings by a combination of affine mappings and a coordinatewise applied (generally) non-linear function, which is referred to as the activation function. Most rigorous analyses study the approximation properties of certain DNN *architectures* and the existence of DNN weights which guarantee a small error, cf. [35, 3, 27, 8]. A common theme among these references is that DNNs perform as well as the state-of-the-art numerical method in a variety of contexts.

In application the actual weights that represent the particular DNN are commonly computed by DNN training, effected with numerically minimizing a certain positive functional, the loss-function. Training of DNNs is generally challenging: the occurring optimization problem is highly non-convex. This has for example been approached in computational uncertainty quantification and also in image processing, cf. [28, 21, 2, 1].

On the other hand, there is empirical and mathematical evidence that this approach may not always be successful, cf. e.g. [12, 5]. In the present paper, we propose to explicitly construct values of the weights of DNNs via a deterministic algorithm thereby circumvent the need of computationally costly optimization routines.

1.1 Previous Work

The idea of constructive NN approximations via Chebyšev expansions and collocation in a tensor product Clenshaw-Curtis grid was already suggested in [23], in particular [23, Theorem 2.3] on the approximation of multivariate holomorphic functions by DNNs with smooth activations. There, robustness of the NN approximation rates under noisy data was described as “expected”, but not proved.

Constructive proofs for ReLU-based DNN emulation of polynomials were firstly seen in [19, 35]. Mainly in [35] a construction for DNNs with rectified linear unit (ReLU) activation that approximate the product of two scalars was found. The convergence of these constructions is exponential with respect to the size of the DNNs meaning the number of nonzero weights. Subsequent works used DNNs and polynomial approximations for smooth (or analytic) functions and established the existence of DNNs with exponential convergence, in the L^∞ - and some in the stronger $W^{1,\infty}$ -norm: For example, [31] provided $W^{1,\infty}$ -error bounds for the product network in [35], exponential convergence in $W^{1,\infty}$ for univariate analytic and Gevrey regular functions was shown in [25], [9] provided exponentially convergent DNN approximations of holomorphic maps on $[-1, 1]^d$, with respect to the L^∞ -norm. A more efficient approximation with respect to the $W^{1,\infty}$ -norm and under weaker smoothness assumptions was given in [26]. The proofs in these references are constructive, in principle, and algorithmic realizations could be based on the ReLU reapproximation of polynomials as linear combinations of monomials. It is well-known that such representations may have poor stability in finite-precision arithmetic (in particular, in the context of quantized DNN weights) due to exponential w.r. to the polynomial degree coefficient growth in monomial expansions. Furthermore, these constructions use e.g. Taylor, Legendre or Chebyšev coefficients of analytic functions, which cannot be determined exactly based on a finite number of function evaluations.

An alternative approach to exponential DNN approximation of multivariate holomorphic maps was used in [8], which used ReLU NN approximations of tensor products of univariate polynomials with real roots by using approximate ReLU NN multiplications of their linear factors. This did not lead to better asymptotic approximation rates than for the previously mentioned monomial-based polynomial approximations. The NNs inherit the well-known stability and conditioning issues from monomial representations of interpolation polynomials.

Constructive approximation of multivariate holomorphic parametric maps by tensorized Chebyšev polynomials with exponential convergence has been studied in the context of option pricing in [13]. The approximation of multivariate maps by DNNs based on Chebyšev polynomials has been considered in [32]. There, Chebyšev polynomials represented exactly by so called *RePU* DNNs serve as starting value for computing DNN weights in black box optimization routines. The presented DNNs are based on a well-known identity satisfied by the product of two Chebyšev polynomials. The initial error prior to NN optimization was not analyzed theoretically in [32].

Relevance of the present results is due to the fact that numerous applications in computational science and engineering aim at efficient numerical realization of input-output maps between suitable Banach spaces, such as for example data-to-solution maps for continuum models governed by partial differential equations. Equipping the input data space with suitable, affine-parametric representation systems such as (Riesz) bases or frames renders the maps of interest parametric. In many applications such maps are holomorphic, even for data spaces of inputs with possibly low spatial or temporal regularity, see for example [15]. In the present paper, following a general, algorithmic DNN construction and expression rate analysis, we develop one

example consisting in Data-to-Prediction maps for Bayesian Inverse Problems of elliptic PDEs which should be contrasted with other, standard (i.e. numerical minimization of loss functions) approaches (e.g. [34, 20]). Other applications include shape-to-solution maps of differential or integral equations (e.g. [17, 6, 14] and the references there).

1.2 Contributions

The principal contributions of this work are threefold: we propose and implement a constructive numerical approximation algorithm for DNN expression of maps $f : [-1, 1]^K \rightarrow \mathbb{R}$. Specifically, we propose and analyze an algorithm to build by explicit construction a DNN surrogate \tilde{f} to the map f with a) rigorous, sup-norm generalization error bounds and with b) good stability properties in finite precision arithmetic, and c) with low complexity.

We achieve this by adapting ideas from spectral collocation, for some finite index set $\Lambda \subset \mathbb{N}_0^K$ with associated polynomial space $\mathbb{P}_\Lambda = \text{span}\{\mathbf{x}^\nu : \nu \in \Lambda\}$ and corresponding data-sampling grid $\Gamma_\Lambda \subset [-1, 1]^K$ of unisolvent (for \mathbb{P}_Λ) interpolation points. We consider the (unique) polynomial interpolant $f_\Lambda \in \mathbb{P}_\Lambda$ which satisfies

$$\forall \mathbf{x} \in \Gamma_\Lambda : f(\mathbf{x}) = f_\Lambda(\mathbf{x}) .$$

Numerical stability of the polynomial interpolation process in finite precision arithmetic is required in order to preclude catastrophic amplification of, e.g., numerical errors in the function value queries $f(\mathbf{x})$, for $\mathbf{x} \in \Gamma_\Lambda$. It is well known to depend on two related issues: the choice of the sampling grid Γ_Λ and of the basis $\{p_\nu : \nu \in \Lambda\}$ spanning \mathbb{P}_Λ . Favorable numerical conditioning of the interpolation operator $I_\Lambda : C^0([-1, 1]^K) \rightarrow \mathbb{P}_\Lambda : f \mapsto f_\Lambda$ is known to be governed by the Lebesgue constant of Γ_Λ . Assuming the basis to have the *Kronecker property*, i.e. $p_\nu(\mathbf{x}^\mu) = \delta_{\nu, \mu}$ for all $\nu, \mu \in \Lambda$, we may represent I_Λ as

$$f_\Lambda(\mathbf{x}) := I_\Lambda[f](\mathbf{x}) = \sum_{\nu \in \Lambda} f(\mathbf{x}_\nu) p_\nu(\mathbf{x}), \quad \mathbf{x} \in [-1, 1]^K . \quad (1.1)$$

As $[-1, 1]^K$ is a cartesian product, the construction of unisolvent interpolation grids Γ_Λ with good stability properties and corresponding error bounds for $f - f_\Lambda$ in various norms in $[-1, 1]^K$ reduces to the univariate case.

The second step in the DNN construction is to replace in (1.1) the basis functions p_ν by suitable DNNs, denoted generically by $\mathbf{x} \mapsto \tilde{p}_\nu(\mathbf{x})$. Based on (1.1), the DNN is then given by (assuming at hand exact evaluations of f on Γ_Λ)

$$\mathbf{x} \mapsto \tilde{f}_\Lambda(\mathbf{x}) := \sum_{\nu \in \Lambda} f(\mathbf{x}_\nu) \tilde{p}_\nu(\mathbf{x}), \quad \mathbf{x} \in [-1, 1]^K .$$

The DNNs $\tilde{p}_\nu \in \mathbb{P}_\Lambda$ are independent of the data $f(\mathbf{x}_\nu)$ and are *constructed a-priori* (i.e., the DNN should “learn” the basis $\{p_\nu : \nu \in \Lambda\}$ of \mathbb{P}_Λ offline, prior to accounting for data $f(\mathbf{x}_\nu)$) to prescribed numerical accuracy $\varepsilon > 0$ being dictated, e.g., by the data fidelity. Our algorithm therefore mandates the construction of \tilde{p}_ν from an admissible class of DNNs. We consider here in particular the class of DNNs with ReLU activation, and we propose a DNN architecture of low complexity by exploiting certain algebraic properties of Chebyšev polynomials which are tensorized to build the p_ν . Based on the observation of [32], we obtain smaller architectures and better stability than with other bases p_ν (e.g. in [35, 26] either monomial bases were considered resulting in small DNNs with large Lebesgue constants or Legendre polynomials were considered which have better stability, but require larger DNNs for their accurate emulation).

For ease of exposition, we consider here only *isotropic tensor product interpolation*. I.e., for a prescribed polynomial degree $p \in \mathbb{N}$, $\Lambda := \{0 : p\}^K = \{\nu \in \mathbb{N}_0^K : \|\nu\|_\infty \leq p\}$. We hasten to add, however, that also more general polynomial spaces could be considered which are based

on, e.g., anisotropic, total degree or more general sparse grids $\Gamma_\Lambda \subset [-1, 1]^K$ which will be addressed elsewhere. The DNN expression and generalization error analysis will then proceed according to

$$\|f - \tilde{f}_\Lambda\| \leq \|f - f_\Lambda\| + \|f_\Lambda - \tilde{f}_\Lambda\|,$$

where the first error is an error of polynomial interpolation, and the second error term is, essentially, the DNN emulation error of the polynomial basis elements p_ν . The choice of the norm $\|\circ\|$ will be either the L^∞ or also the $W^{1,\infty}$ norm on the (scaled) input data domain $[-1, 1]^K$. The latter norm therefore also allows to control the generalization error of first order sensitivities of the DNNs.

We develop the error analysis and the emulation algorithm for DNNs with the rectified linear unit (ReLU) activation function. We prove the proposed construction algorithm (see ahead Algorithm 4.1) yields ReLU DNNs that are able to approximate multivariate analytic mappings that have a size of the DNN that depends only polylogarithmically on the reciprocal of the desired accuracy. This error bound is pointwise and holds for all possible inputs in the hypercube $[-1, 1]^K$. A so-called training routine, for example by the widely used stochastic gradient descent method, is not necessary.

More specifically, the algorithm combines mainly two building blocks. Firstly, the function to be approximated is collocated on a multivariate, tensor product Clenshaw–Curtis grid Γ which serves here as *training data*. Based on this data, exponentially convergent polynomial interpolants can be constructed as linear combination of tensor product Chebyšev polynomials (we recall the relevant details of polynomial interpolation in Section 2). Secondly, DNNs are constructed that emulate Chebyšev polynomials in a stable way. It shows that here the meticulous choice of training data points is crucial for the good performance of the algorithm, in particular for the exponential convergence of ReLU DNNs for analytic functions. This is in contrast to randomly chosen training points, which have been used in connection with the stochastic gradient method, see for example [21].

We provide a complete error analysis for the proposed DNN approximation in Proposition 3.5. Specifically, we prove exponential decay of three error contributions in terms of the number of nonzero NN weights by estimating (i) the error of approximating f by a polynomial, (ii) the additional error of polynomial interpolation, based on only finitely many function values of f , (iii) and the error of reapproximating the polynomial interpolant by a ReLU DNN. In the approximation literature, the first two are often referred to as *approximation error* and *generalization error*, respectively. In Proposition 3.5, it is also shown that if the function f can only be accessed through a possibly corrupted numerical approximation \tilde{f} , the proposed DNN approximations are stable with respect to the error $f - \tilde{f}$. In the worst case, the L^∞ -error $f - \tilde{f}$ is magnified by a factor $C(\log(\log(1/\varepsilon)))^K$, where ε denotes the accuracy of polynomial approximation (i.e. $\log(1/\varepsilon) \propto n$ in Proposition 3.5). Also the bound on the $W^{1,\infty}$ -error only depends on $f - \tilde{f}$ through $\|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$, multiplied with a factor $C(\log(1/\varepsilon))^2(\log(\log(1/\varepsilon)))^K$. In Remark 3.6 we quantify stability of the DNN architecture and DNN weights with respect to small changes in the approximated function f .

1.3 Notation

We will denote vectors and multiindices by bold characters. We denote $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}_0 = \{0, 1, 2, \dots\}$. For $k \in \mathbb{N}_0$ and a subset $S \subset \mathbb{N}_0$, we define $\mathbf{1}_S(k) := 1$ if $k \in S$, and $\mathbf{1}_S(k) := 0$ otherwise. For $K \in \mathbb{N}$ and $\mathbf{k} \in \mathbb{N}_0^K$, we define $\mathbf{1}_S(\mathbf{k}) := \sum_{j=1}^K \mathbf{1}_S(k_j)$ and denote by $|\mathbf{k}|_0 := \mathbf{1}_{\mathbb{N}}(\mathbf{k})$ the number of nonzero components of \mathbf{k} . For finite index sets $\Lambda \subset \mathbb{N}_0^K$, we denote the number of elements by $|\Lambda|$ and the maximum coordinatewise degree by $m_\infty(\Lambda) := \max_{\mathbf{k} \in \Lambda} \|\mathbf{k}\|_{\ell^\infty}$.

We denote by T_k , $k \in \mathbb{N}_0$, the univariate Chebyšev polynomials of the first kind, normalized such that $T_k(1) = 1$ for all $k \in \mathbb{N}_0$. For $K \in \mathbb{N}$ and $\mathbf{k} = (k_j)_{j=1}^K \in \mathbb{N}_0^K$, we denote tensor

product Chebyšev polynomials by $T_{\mathbf{k}}(\mathbf{x}) := \prod_{j=1}^K T_{k_j}(x_j)$, for $\mathbf{x} = (x_j)_{j=1}^K \in [-1, 1]^K$. We write $\cos(\boldsymbol{\theta}) := (\cos(\theta_1), \dots, \cos(\theta_K))$ for vectors $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K) \in \mathbb{R}^K$ and $\mathbf{x}^\nu := \prod_{j=1}^K x_j^{\nu_j}$ for $\mathbf{x} \in \mathbb{R}^K$ and $\boldsymbol{\nu} \in \mathbb{N}_0^K$.

We introduce the following notation for circles in the complex plane: For all $r > 0$ we define $\Gamma_r := \{z \in \mathbb{C} : |z| = r\}$. For all $r \geq 1$, the image of the closed disc of radius r in \mathbb{C} under the map $z \mapsto \frac{z+z^{-1}}{2}$ is a closed *Bernstein ellipse*, denoted by $\mathcal{E}_\rho := \left\{ \frac{z+z^{-1}}{2} \in \mathbb{C} : 1 \leq |z| \leq \rho \right\}$. For $p \geq 0$, the space of polynomials of degree at most p are denoted by \mathbb{P}_p . The space of polynomials in $K \in \mathbb{N}$ variables of coordinatewise degree at most p are denoted by $\mathbb{Q}_p := \text{span}\{\mathbf{x}^\nu : \boldsymbol{\nu} \in \mathbb{N}_0^K \text{ and } \|\boldsymbol{\nu}\|_{\ell^\infty} \leq p\}$.

Error estimates will be expressed in terms of the $W^{1,\infty}$ Sobolev norm, which is defined, for an open and bounded domain $\Omega \subset \mathbb{R}^K$, as $\|u\|_{W^{1,\infty}(\Omega)} = \max\{\|u\|_{L^\infty(\Omega)}, \max_{i=1}^K \|\frac{\partial}{\partial x_i} u\|_{L^\infty(\Omega)}\}$, $\frac{\partial}{\partial x_i}$ denoting weak derivatives.

1.4 Outline

In Section 2, we first recapitulate classical results on constructive polynomial approximation of multivariate, holomorphic functions on $[-1, 1]^K$. In Section 3, we recapitulate ReLU DNN approximation rates of multivariate maps $f : [-1, 1]^K \rightarrow \mathbb{R}$ from [26]. We also develop a constructive DNN approximation. It is based on standard spectral collocation approximation of f in a tensorized Clenshaw–Curtis grid, and on ReLU DNN emulation of tensorized Chebyšev polynomials. In Section 3.4, we comment on generalizations of our results to DNNs with activation functions other than ReLU. In Section 4, a construction algorithm is proposed and analyzed that converges exponentially for analytic functions. An application of the presented algorithm to a Bayesian inverse problem is discussed in Section 5 and numerical experiments confirming the theory are provided in Section 6.

2 Polynomial Approximation of Multivariate Holomorphic Functions

It is classical that univariate functions $f : [-1, 1] \rightarrow \mathbb{R}$ which are real-analytic in $[-1, 1]$ admit sequences of *polynomial approximations* $\{f_p\}_{p \geq 0}$ with $f_p \in \mathbb{P}_p$ which converge at an exponential rate. These univariate polynomial approximations can be tensorized to produce exponentially convergent, tensorized polynomial approximations to multivariate maps $f : [-1, 1]^K \rightarrow \mathbb{R}$. This argument was used in [26] to infer existence of tensorized truncated Legendre expansions of coordinatewise polynomial degree at most $p \in \mathbb{N}$ of f which could, in principle, serve as building block for the corresponding multivariate ReLU NNs which emulate the map f .

In this section, we present an alternative proof of such multivariate polynomial approximation of holomorphic maps $f : [-1, 1]^K \rightarrow \mathbb{R}$, which admit a holomorphic complex extension to the isotropic Bernstein polyellipse $\mathcal{E}_\rho = \mathcal{E}_\rho^K$ with polyradius $\boldsymbol{\rho} = (\rho, \dots, \rho) \in \mathbb{R}^K$ for some $\rho \in (1, \infty)$. The interpolation results are basically known, being based on tensorized Chebyšev polynomials. We detail them here, as they are the basis for the ensuing ReLU DNN approximations. Bounds on Lebesgue constants of Chebyšev points in $[-1, 1]$ are essential in quantifying numerical stability of the interpolation and, more importantly, of the DNN approximation process.

2.1 Chebyšev Expansion

We first recall the tensor product Chebyšev expansion, which is obtained by inductively taking the Chebyšev expansion with respect to each of the K coordinates. The following results are classical, we refer to [33, Theorem 3.1] for the arguments in the univariate case, which we

apply K times. Denoting the Chebyšev measure on $[-1, 1]$ by λ , which has Lebesgue density $(1 - x^2)^{-1/2}$ for $x \in (-1, 1)$, it holds for all $\mathbf{y} = (y_1, \dots, y_K) \in [-1, 1]^K$

$$\begin{aligned}
f(y_1, \dots, y_K) &= \sum_{k_1=0}^{\infty} 2^{\mathbb{1}_{\mathbb{N}}(k_1)} \pi^{-1} T_{k_1}(y_1) \int_{-1}^1 f(x_1, y_2, \dots, y_K) T_{k_1}(x_1) d\lambda(x_1) \\
&= \sum_{k_1=0}^{\infty} 2^{\mathbb{1}_{\mathbb{N}}(k_1)} \pi^{-1} T_{k_1}(y_1) \int_{-1}^1 \dots \\
&\quad \sum_{k_K=0}^{\infty} 2^{\mathbb{1}_{\mathbb{N}}(k_K)} \pi^{-1} T_{k_K}(y_K) \int_{-1}^1 f(x_1, \dots, x_K) T_{k_K}(x_K) d\lambda(x_K) \dots T_{k_1}(x_1) d\lambda(x_1) \\
&= \sum_{\mathbf{k} \in \mathbb{N}_0^K} T_{\mathbf{k}}(\mathbf{y}) 2^{|\mathbf{k}|_0} \pi^{-K} \int_{-1}^1 \dots \int_{-1}^1 f(\mathbf{x}) T_{\mathbf{k}}(\mathbf{x}) d\lambda(x_1) \dots d\lambda(x_K) \\
&=: \sum_{\mathbf{k} \in \mathbb{N}_0^K} T_{\mathbf{k}}(\mathbf{y}) f_{\mathbf{k}}. \tag{2.1}
\end{aligned}$$

In the third step, interchanging summation and integration is justified by the dominated convergence theorem. It implies that for all $j = 2, \dots, K$ and $\ell = 1, \dots, j-1$, all $(k_\ell, \dots, k_{j-1}) \in \mathbb{N}_0^{j-\ell}$, all $\mathbf{y} = (y_1, \dots, y_K) \in [-1, 1]^K$ and all $(x_1, \dots, x_{\ell-1}) \in [-1, 1]^{\ell-1}$

$$\begin{aligned}
&\int_{-1}^1 \sum_{k_j=0}^{\infty} \pi^{-(j-\ell)} \prod_{i=\ell+1}^j \left(2^{\mathbb{1}_{\mathbb{N}}(k_i)} T_{k_i}(y_i) \right) \\
&\quad \left[\int_{[-1, 1]^{j-\ell}} f(x_1, \dots, x_j, y_{j+1}, \dots, y_K) T_{k_j}(x_j) d\lambda(x_j) \dots T_{k_{\ell+1}}(x_{\ell+1}) d\lambda(x_{\ell+1}) \right] T_{k_\ell}(x_\ell) d\lambda(x_\ell) \\
&= \lim_{N \rightarrow \infty} \int_{-1}^1 \sum_{k_j=0}^N \pi^{-(j-\ell)} \prod_{i=\ell+1}^j \left(2^{\mathbb{1}_{\mathbb{N}}(k_i)} T_{k_i}(y_i) \right) \\
&\quad \left[\int_{[-1, 1]^{j-\ell}} f(x_1, \dots, x_j, y_{j+1}, \dots, y_K) T_{k_j}(x_j) d\lambda(x_j) \dots T_{k_{\ell+1}}(x_{\ell+1}) d\lambda(x_{\ell+1}) \right] T_{k_\ell}(x_\ell) d\lambda(x_\ell).
\end{aligned}$$

Use of the dominated convergence theorem is justified because for all $N \in \mathbb{N}_0$, Lemma 2.1 below can be applied to

$$g : [-1, 1]^{j-\ell} \rightarrow \mathbb{R} : (x_{\ell+1}, \dots, x_j) \mapsto f(x_1, \dots, x_j, y_{j+1}, \dots, y_K),$$

which implies that

$$\begin{aligned}
&\left| \sum_{k_j=0}^N \pi^{-(j-\ell)} \prod_{i=\ell+1}^j \left(2^{\mathbb{1}_{\mathbb{N}}(k_i)} T_{k_i}(y_i) \right) \right. \\
&\quad \left. \left[\int_{[-1, 1]^{j-\ell}} f(x_1, \dots, x_j, y_{j+1}, \dots, y_K) T_{k_j}(x_j) d\lambda(x_j) \dots T_{k_{\ell+1}}(x_{\ell+1}) d\lambda(x_{\ell+1}) \right] T_{k_\ell}(x_\ell) \right| \\
&\leq \sum_{k_j=0}^N \pi^{-(j-\ell)} \prod_{i=\ell+1}^j \left(2^{\mathbb{1}_{\mathbb{N}}(k_i)} |T_{k_i}(y_i)| \right) \\
&\quad \left| \int_{[-1, 1]^{j-\ell}} f(x_1, \dots, x_j, y_{j+1}, \dots, y_K) T_{k_j}(x_j) d\lambda(x_j) \dots T_{k_{\ell+1}}(x_{\ell+1}) d\lambda(x_{\ell+1}) \right| |T_{k_\ell}(x_\ell)| \\
&\leq \sum_{k_j=0}^N \pi^{-(j-\ell)} \prod_{i=\ell+1}^j 2^{\mathbb{1}_{\mathbb{N}}(k_i)}
\end{aligned}$$

$$\begin{aligned}
& \left| \int_{[-1,1]^{j-\ell}} f(x_1, \dots, x_j, y_{j+1}, \dots, y_K) T_{k_j}(x_j) \cdots T_{k_{\ell+1}}(x_{\ell+1}) \, d\lambda(x_j) \cdots d\lambda(x_{\ell+1}) \right| \\
&= \sum_{k_j=0}^N |g_{(k_{\ell+1}, \dots, k_j)}| \leq \prod_{i=\ell+1}^j 2^{\mathbb{1}_{\mathbb{N}}(k_i)} \max_{\mathbf{z} \in \mathcal{E}_\rho^{j-\ell}} |g(\mathbf{z})| \rho^{-(k_{\ell+1}, \dots, k_j)} < 2^{j-\ell} \frac{\rho}{\rho-1} \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})|,
\end{aligned}$$

where $g_{(k_{\ell+1}, \dots, k_j)}$ denotes a Chebyšev coefficient of g , analogous to $f_{\mathbf{k}}$ defined in Equation (2.1).

We next estimate the size of the coefficients, which we later need to bound the NN error.

Lemma 2.1 *Let $K \in \mathbb{N}$, and let $f : [-1, 1]^K \rightarrow \mathbb{R}$ be a map which admits a holomorphic complex extension to the isotropic Bernstein polyellipse $\mathcal{E}_\rho \subset \mathbb{C}^K$ with $\boldsymbol{\rho} = (\rho, \dots, \rho) \in (1, \infty)^K$ for some $\rho > 1$. Then, for every $\mathbf{k} \in \mathbb{N}_0^K$*

$$|f_{\mathbf{k}}| \leq 2^{|\mathbf{k}|_0} \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-\mathbf{k}}, \quad \sum_{\mathbf{k} \in \mathbb{N}_0^K} |f_{\mathbf{k}}| \leq \left(\frac{2\rho}{\rho-1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})|. \quad (2.2)$$

Proof. The Chebyšev coefficients of f satisfy (cf. e.g. [33, Equations (3.12) – (3.14)] and [29, Theorem 3.8])

$$\begin{aligned}
f_{\mathbf{k}} &= 2^{|\mathbf{k}|_0} \pi^{-K} \int_{-1}^1 \cdots \int_{-1}^1 f(\mathbf{x}) T_{\mathbf{k}}(\mathbf{x}) \, d\lambda(x_1) \cdots d\lambda(x_K) \\
&= 2^{|\mathbf{k}|_0} (2\pi)^{-K} \int_{-\pi}^{\pi} \cdots \int_{-\pi}^{\pi} f(\cos(\boldsymbol{\theta})) \cos(k_1 \theta_1) \, d\theta_1 \cdots \cos(k_K \theta_K) \, d\theta_K \\
&= 2^{|\mathbf{k}|_0} (2\pi i)^{-K} \int_{\Gamma_1} \cdots \int_{\Gamma_1} f\left(\frac{z_1+z_1^{-1}}{2}, \dots, \frac{z_K+z_K^{-1}}{2}\right) \left(\frac{z_1+z_1^{-1}}{2}\right) \frac{dz_1}{z_1} \cdots \left(\frac{z_K+z_K^{-1}}{2}\right) \frac{dz_K}{z_K} \\
&= 2^{|\mathbf{k}|_0} (2\pi i)^{-K} \int_{\Gamma_1} \cdots \int_{\Gamma_1} f\left(\frac{z_1+z_1^{-1}}{2}, \dots, \frac{z_K+z_K^{-1}}{2}\right) z_1^{-k_1} \frac{dz_1}{z_1} \cdots z_K^{-k_K} \frac{dz_K}{z_K},
\end{aligned} \quad (2.3)$$

where the last step follows by invariance of $\frac{z+z^{-1}}{2}$ under the transformation $z \mapsto z^{-1}$, for $z \in \mathbb{C} \setminus \{0\}$. By holomorphy of f , we can change the curve of integration from Γ_1 to Γ_ρ , so that

$$\begin{aligned}
|f_{\mathbf{k}}| &= \left| 2^{|\mathbf{k}|_0} (2\pi i)^{-K} \int_{\Gamma_\rho} \cdots \int_{\Gamma_\rho} f\left(\frac{z_1+z_1^{-1}}{2}, \dots, \frac{z_K+z_K^{-1}}{2}\right) \mathbf{z}^{-\mathbf{k}} \frac{dz_1}{z_1} \cdots \frac{dz_K}{z_K} \right| \\
&\leq 2^{|\mathbf{k}|_0} \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| (2\pi)^{-K} \left| \int_{\Gamma_\rho} \cdots \int_{\Gamma_\rho} \mathbf{z}^{-\mathbf{k}} \frac{dz_1}{z_1} \cdots \frac{dz_K}{z_K} \right| \\
&\leq 2^{|\mathbf{k}|_0} \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-\mathbf{k}}, \\
\sum_{\mathbf{k} \in \mathbb{N}_0^K} |f_{\mathbf{k}}| &\leq 2^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \left(\frac{\rho}{\rho-1} \right)^K = \left(\frac{2\rho}{\rho-1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})|.
\end{aligned}$$

□

2.2 Chebyšev Interpolation

Because in general the Chebyšev coefficients of f are not known explicitly, we next approximate the coefficients and thus consider polynomial interpolation of f : We approximate f by a polynomial whose coefficients with respect to the Chebyšev basis only depend on function values of f in the tensor product Clenshaw–Curtis grid. For $n \in \mathbb{N}$, we consider the approximation $\hat{f}_{\mathbf{k};n}$ of the Chebyšev coefficients, for $\mathbf{k} \in \{0, \dots, n\}^K$, obtained by applying an $(n+1)^K$ point

tensor product Gauss–Lobatto quadrature in the Clenshaw–Curtis points $\mathbf{x}_j^n := \cos(j\pi/n)$ for $\mathbf{j} \in \{0, \dots, n\}^K$ to the integrals in (2.3)

$$\begin{aligned} w_j &:= \left(\frac{\pi}{2n}\right)^K \prod_{\ell=1}^K 2^{\mathbb{1}_{\{1, \dots, n-1\}}(j_\ell)}, \quad \mathbf{j} \in \{0, \dots, n\}^K, \\ \hat{f}_{\mathbf{k};n} &:= 2^{\mathbb{1}_{S(n)}(\mathbf{k})} \pi^{-K} \sum_{\mathbf{j} \in \{0, \dots, n\}^K} w_j f(\cos(j\pi/n)) \cos(\mathbf{k}\mathbf{j}\pi/n) \\ &= 2^{\mathbb{1}_{S(n)}(\mathbf{k})} (2n)^{-K} \sum_{\mathbf{j} \in \{0, \dots, 2n-1\}^K} f(\cos(j\pi/n)) \cos(\mathbf{k}\mathbf{j}\pi/n) \\ &= 2^{\mathbb{1}_{S(n)}(\mathbf{k})} \text{IFFT} \left((f(\mathbf{x}_j^n))_{\mathbf{j} \in \{0, \dots, 2n-1\}^K} \right)_{\mathbf{k}}, \end{aligned} \quad (2.4)$$

where $S(n) := \{1, \dots, n-1\}$ and $\mathbf{k}\mathbf{j} = (k_1 j_1, \dots, k_K j_K) \in \mathbb{N}_0^K$ for $\mathbf{j} \in \{0, \dots, 2n-1\}^K$. In addition, IFFT denotes the inverse fast Fourier transform. Note that $\cos((2n-j)\pi/n) = \cos(j\pi/n)$ for $j = 1, \dots, n-1$, thus $(f(\mathbf{x}_j^n))_{\mathbf{j} \in \{0, \dots, n\}^K}$ are sufficient to determine $(f(\mathbf{x}_j^n))_{\mathbf{j} \in \{0, \dots, 2n-1\}^K}$. In $2^{\mathbb{1}_{S(n)}(\mathbf{k})}$, the factor 2 is omitted for each $k_j = n$, $j = 1, \dots, K$, because of aliasing in the coefficients of the Chebyšev interpolant (cf. [33, Chapter 4]).

The approximation

$$\hat{p}_{f,n} := I_n^{\text{Lob}}[f] := \sum_{\mathbf{k} \in \{0, \dots, n\}^K} \hat{f}_{\mathbf{k};n} T_{\mathbf{k}} \quad (2.5)$$

of f is in fact the Lagrange interpolant in the nodes $(\mathbf{x}_j^n)_{\mathbf{j} \in \{0, \dots, n\}^K}$ (cf. [29, Theorem 3.13] for $K = 1$, which can be applied inductively). It is known that the Lebesgue constant of tensor product Lagrange interpolation in these nodes satisfies

$$\|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} := \|I_n^{\text{Lob}}\|_{L^\infty([-1,1]^K), L^\infty([-1,1]^K)} \leq \left(\frac{2}{\pi} \log(n+1) + 1\right)^K$$

(cf. [29, Theorem 1.2] for $K = 1$ and the zeroes of Chebyšev polynomials as nodes, from which the result for $(\mathbf{x}_j^n)_{\mathbf{j} \in \{0, \dots, n\}^K}$ and $K = 1$ follows with [10, Theorem 4, Equation (4.6)]). Using Markov's inequality, a bound on the Lebesgue constant with respect to the $W^{1,\infty}$ -norm can be obtained as follows:

$$\begin{aligned} \|I_n^{\text{Lob}}\|_{W^{1,\infty}, W^{1,\infty}} &:= \|I_n^{\text{Lob}}\|_{W^{1,\infty}([-1,1]^K), W^{1,\infty}([-1,1]^K)} \\ &= \sup_{g \in W^{1,\infty}([-1,1]^K) \setminus \{0\}} \frac{\|I_n^{\text{Lob}}[g]\|_{W^{1,\infty}([-1,1]^K)}}{\|g\|_{W^{1,\infty}([-1,1]^K)}} \\ &\leq \sup_{g \in W^{1,\infty}([-1,1]^K) \setminus \{0\}} \frac{n^2 \|I_n^{\text{Lob}}[g]\|_{L^\infty([-1,1]^K)}}{\|g\|_{W^{1,\infty}([-1,1]^K)}} \\ &\leq \sup_{g \in W^{1,\infty}([-1,1]^K) \setminus \{0\}} \frac{n^2 \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|g\|_{L^\infty([-1,1]^K)}}{\|g\|_{W^{1,\infty}([-1,1]^K)}} \\ &\leq n^2 \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty}. \end{aligned}$$

In the third step, we used that for all $q \in \mathbb{Q}_n([-1,1]^K)$

$$\begin{aligned} &\max_{i=1}^K \left\| \frac{\partial}{\partial x_i} q \right\|_{L^\infty([-1,1]^K)} \\ &= \max_{i=1}^K \max_{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_K \in [-1,1]} \left\| \left(\frac{\partial}{\partial x_i} q \right) (x_1, \dots, x_{i-1}, \cdot, x_{i+1}, \dots, x_K) \right\|_{L^\infty([-1,1])} \\ &\leq \max_{i=1}^K \max_{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_K \in [-1,1]} n^2 \|q(x_1, \dots, x_{i-1}, \cdot, x_{i+1}, \dots, x_K)\|_{L^\infty([-1,1])} \\ &= n^2 \|q\|_{L^\infty([-1,1]^K)}, \end{aligned} \quad (2.6)$$

where we applied Markov's inequality to the univariate polynomial $z \mapsto q(x_1, \dots, x_{i-1}, z, x_{i+1}, \dots, x_K)$. This computation holds in particular for polynomials $q \in \mathbb{Q}_n([-1, 1]^K)$ which are themselves derivatives of other polynomials in $\mathbb{Q}_n([-1, 1]^K)$, hence the previous argument can be iterated to obtain for all $s \in \mathbb{N}$

$$\|I_n^{\text{Lob}}\|_{W^{s,\infty}, W^{s,\infty}} \leq n^{2s} \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty}.$$

The arguments used above also prove the following, slightly stronger result:

$$\begin{aligned} \|I_n^{\text{Lob}}\|_{L^\infty, W^{1,\infty}} &:= \|I_n^{\text{Lob}}\|_{L^\infty([-1,1]^K), W^{1,\infty}([-1,1]^K)} \\ &= \sup_{g \in W^{1,\infty}([-1,1]^K) \setminus \{0\}} \frac{\|I_n^{\text{Lob}}[g]\|_{W^{1,\infty}([-1,1]^K)}}{\|g\|_{L^\infty([-1,1]^K)}} \leq n^2 \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty}. \end{aligned} \quad (2.7)$$

Again, we get the analogous result for arbitrary $s \in \mathbb{N}$ by iterating (2.6) s times:

$$\|I_n^{\text{Lob}}\|_{L^\infty, W^{s,\infty}} \leq n^{2s} \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty}.$$

2.3 Chebyšev Interpolation Based On Approximate Function Values

In the case that the function f is only accessible through a (possibly corrupted) numerical approximation \tilde{f} in the Chebyšev points $(\mathbf{x}_j^n)_{j \in \{0, \dots, n\}^K}$, a further approximation of the coefficients $\hat{f}_{\mathbf{k};n}$ is made. The function values of f in Equation (2.4) are replaced by function values of \tilde{f} . For all $\mathbf{k} \in \{0, \dots, n\}^K$

$$\tilde{f}_{\mathbf{k};n} := 2^{\mathbb{1}_{S(n)}(\mathbf{k})} \pi^{-K} \sum_{j \in \{0, \dots, n\}^K} w_j \tilde{f}(\cos(j\pi/n)) \cos(\mathbf{k}j\pi/n) \quad (2.8)$$

$$\begin{aligned} &= 2^{\mathbb{1}_{S(n)}(\mathbf{k})} (2n)^{-K} \sum_{j \in \{0, \dots, 2n-1\}^K} \tilde{f}(\cos(j\pi/n)) \cos(\mathbf{k}j\pi/n) \\ &= 2^{\mathbb{1}_{S(n)}(\mathbf{k})} \text{IFFT} \left((\tilde{f}(\mathbf{x}_j^n))_{j \in \{0, \dots, 2n-1\}^K} \right)_{\mathbf{k}}. \end{aligned} \quad (2.9)$$

We define the corresponding interpolant $\tilde{p}_{f,n}$ as

$$\tilde{p}_{f,n} := \sum_{\mathbf{k} \in \{0, \dots, n\}^K} \tilde{f}_{\mathbf{k};n} T_{\mathbf{k}} = I_n^{\text{Lob}}[\tilde{f}]. \quad (2.10)$$

Lemma 2.2 *Let $K \in \mathbb{N}$, and let $f : [-1, 1]^K \rightarrow \mathbb{R}$ be a map which admits a holomorphic complex extension to the Bernstein polyellipse \mathcal{E}_ρ with $\rho = (\rho, \dots, \rho) \in (1, \infty)^K$ for some $\rho > 1$. We assume that an approximation \tilde{f} of f is available, and an upper bound on $\|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$. Then, for every $\rho' \in (1, \rho)$ and every $s \in \mathbb{N}$, there exists $C'(s, \rho, \rho') > 0$ such that*

$$\begin{aligned} \|f - \tilde{p}_{f,n}\|_{L^\infty([-1, 1]^K)} &\leq \left(1 + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty}\right) K \left(\frac{2\rho}{\rho-1}\right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-n-1} \\ &\quad + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}, \end{aligned} \quad (2.11)$$

$$\begin{aligned} \|f - \tilde{p}_{f,n}\|_{W^{s,\infty}([-1, 1]^K)} &\leq \left(1 + \|I_n^{\text{Lob}}\|_{W^{s,\infty}, W^{s,\infty}}\right) K \left(\frac{2C'\rho'}{\rho'-1}\right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho'^{-n-1} \\ &\quad + \|I_n^{\text{Lob}}\|_{L^\infty, W^{s,\infty}} \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}. \end{aligned} \quad (2.12)$$

Proof. We can estimate the error as follows. For all $q \in \mathbb{Q}_n$ we use $q = I_n^{\text{Lob}}[q]$ and obtain

$$\begin{aligned}
& \|f - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^\kappa)} \\
& \leq \|f - q\|_{L^\infty([-1,1]^\kappa)} + \|q - \hat{p}_{f,n}\|_{L^\infty([-1,1]^\kappa)} + \|\hat{p}_{f,n} - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^\kappa)} \\
& = \|f - q\|_{L^\infty([-1,1]^\kappa)} + \|I_n^{\text{Lob}}[q - f]\|_{L^\infty([-1,1]^\kappa)} + \|I_n^{\text{Lob}}[f - \tilde{f}]\|_{L^\infty([-1,1]^\kappa)} \\
& \leq \|f - q\|_{L^\infty([-1,1]^\kappa)} + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - q\|_{L^\infty([-1,1]^\kappa)} \\
& \quad + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^\kappa)}. \tag{2.13}
\end{aligned}$$

By the same argument it follows that

$$\begin{aligned}
\|f - \tilde{p}_{f,n}\|_{W^{s,\infty}([-1,1]^\kappa)} & \leq \|f - q\|_{W^{s,\infty}([-1,1]^\kappa)} + \|I_n^{\text{Lob}}\|_{W^{s,\infty}, W^{s,\infty}} \|f - q\|_{W^{s,\infty}([-1,1]^\kappa)} \\
& \quad + \|I_n^{\text{Lob}}\|_{L^\infty, W^{s,\infty}} \|f - \tilde{f}\|_{L^\infty([-1,1]^\kappa)}. \tag{2.14}
\end{aligned}$$

Now, we can take the infimum over $q \in \mathbb{Q}_n$ and replace $\|f - q\|$ by the error of best polynomial approximation. Next, we discuss two upper bounds on the error of best polynomial approximation, which are sufficient for our purposes.

In case f is not only holomorphic on \mathcal{E}_ρ , but also on the polydisk $B_{\rho'} := \times_{j=1}^K \{z_j \in \mathbb{C} : |z_j| \leq \rho'\}$, with $\rho' := (\rho', \dots, \rho') \in (1, \infty)^K$ for some $\rho' \in (1, \infty)$, an upper bound on the error of best approximation follows by taking $q \in \mathbb{Q}_n$ to be the Taylor approximation of f in 0 and using a bound on the error of the Taylor approximation, e.g. [15, Theorem 5.1]. It then follows that $\|f - q\|_{L^\infty([-1,1]^\kappa)} \leq C(\rho', f)\rho'^{-n}$. Using that $\frac{\partial}{\partial x_j} q$ is the Taylor approximation of $\frac{\partial}{\partial x_j} f$ for $j = 1, \dots, K$, the same bound on the error of the Taylor approximation gives that $\|f - q\|_{W^{1,\infty}([-1,1]^\kappa)} \leq C(\rho', f)\rho'^{-n+1} \leq C(\rho', f)\rho'^{-n}$, where the constant was increased in the last step. Applying this argument s times gives a bound of the $W^{s,\infty}$ -error.

If such a stronger condition of holomorphy on a polydisk does not hold, we can take q to be a truncation of the Chebyšev expansion: $q := \sum_{\mathbf{k} \in \{0, \dots, n\}^\kappa} f_{\mathbf{k}} T_{\mathbf{k}}$. We obtain a bound on the truncation error using the bound (2.2) on the Chebyšev coefficients. We use as notation $\Lambda_n^c := \mathbb{N}_0^K \setminus \{0, \dots, n\}^\kappa$ and $\mathbf{k}^{(j)} := (0, \dots, 0, n+1, 0, \dots, 0) \in \mathbb{N}_0^K$ for $j = 1, \dots, K$, with the nonzero component in the j 'th position. Using that $\Lambda_n^c = \bigcup_{j=1}^K \{\mathbf{k}^{(j)} + \mathbf{k} : \mathbf{k} \in \mathbb{N}_0^K\}$, and the fact that $\|T_{\mathbf{k}}\|_{L^\infty([-1,1]^\kappa)} = 1$ for all $\mathbf{k} \in \mathbb{N}_0^K$, it follows that

$$\begin{aligned}
\|f - q\|_{L^\infty([-1,1]^\kappa)} & \leq \sum_{\Lambda_n^c} |f_{\mathbf{k}}| \|T_{\mathbf{k}}\|_{L^\infty([-1,1]^\kappa)} \leq \sum_{j=1}^K \sum_{\mathbf{k} \in \mathbb{N}_0^K} |f_{\mathbf{k}^{(j)} + \mathbf{k}}| \\
& \leq \sum_{j=1}^K \sum_{\mathbf{k} \in \mathbb{N}_0^K} 2^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-\mathbf{k}^{(j)} - \mathbf{k}} \leq K \left(\frac{2\rho}{\rho-1}\right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-n-1}.
\end{aligned}$$

For all $s \in \mathbb{N}$, bounds on the $W^{s,\infty}$ -error follow similarly, using that for all $\rho' \in (1, \rho)$ there exists $C'(s, \rho, \rho') > 0$ such that $\rho^{-k} k^{2s} \leq C' \rho'^{-k}$ for all $k \in \mathbb{N}_0$. By the Markov inequality, the definition of the $\|\circ\|_{W^{s,\infty}}$ -norm (cf. Sec. 1.3), and the fact that $\|T_{\mathbf{k}}\|_{L^\infty([-1,1]^\kappa)} = 1$ for all $\mathbf{k} \in \mathbb{N}_0^K$, we get

$$\begin{aligned}
\|T_{\mathbf{k}}\|_{W^{s,\infty}([-1,1]^\kappa)} & = \max_{\mathbf{i} \in \mathbb{N}_0^K : |\mathbf{i}| \leq s} \|\partial_{\mathbf{x}}^{\mathbf{i}} \prod_{\ell=1}^K T_{k_\ell}(x_\ell)\|_{L^\infty([-1,1]^\kappa)} \\
& = \max_{\mathbf{i} \in \mathbb{N}_0^K : |\mathbf{i}| \leq s} \prod_{\ell=1}^K \|\partial_{x_\ell}^{i_\ell} T_{k_\ell}(\cdot)\|_{L^\infty([-1,1])}
\end{aligned}$$

$$\leq \max_{i \in \mathbb{N}_0^K: |i| \leq s} \mathbf{k}^{2i} \leq \prod_{\ell=1}^K \max\{1, k_\ell^{2s}\}.$$

We obtain that for all $\rho' \in (1, \rho)$

$$\begin{aligned} \|f - q\|_{W^{s,\infty}([-1,1]^K)} &\leq \sum_{\Lambda_n^c} |f_{\mathbf{k}}| \|T_{\mathbf{k}}\|_{W^{s,\infty}([-1,1]^K)} \leq \sum_{j=1}^K \sum_{\mathbf{k} \in \mathbb{N}_0^K} |f_{\mathbf{k}^{(j)} + \mathbf{k}}| \prod_{\ell=1}^K \max\{1, (\mathbf{k}^{(j)} + \mathbf{k})_\ell^{2s}\} \\ &\leq \sum_{j=1}^K \sum_{\mathbf{k} \in \mathbb{N}_0^K} 2^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-\mathbf{k}^{(j)} - \mathbf{k}} \prod_{\ell=1}^K \max\{1, (\mathbf{k}^{(j)} + \mathbf{k})_\ell^{2s}\} \\ &\leq K 2^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \left(\sum_{k=0}^{\infty} \max\{1, k^{2s}\} \rho^{-k} \right)^{K-1} \left(\sum_{k=n+1}^{\infty} \max\{1, k^{2s}\} \rho^{-k} \right) \\ &\leq K 2^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \left(\sum_{k=0}^{\infty} C' \rho'^{-k} \right)^{K-1} \left(\sum_{k=n+1}^{\infty} C' \rho'^{-k} \right) \\ &\leq K \left(\frac{2C' \rho'}{\rho' - 1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho'^{-n-1}. \end{aligned}$$

With (2.13) and (2.14) we obtain

$$\begin{aligned} \|f - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^K)} &\leq \|f - q\|_{L^\infty([-1,1]^K)} + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - q\|_{L^\infty([-1,1]^K)} \\ &\quad + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \\ &\leq \left(1 + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty}\right) K \left(\frac{2\rho}{\rho-1}\right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-n-1} \\ &\quad + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)}, \\ \|f - \tilde{p}_{f,n}\|_{W^{s,\infty}([-1,1]^K)} &\leq \|f - q\|_{W^{s,\infty}([-1,1]^K)} + \|I_n^{\text{Lob}}\|_{W^{s,\infty}, W^{s,\infty}} \|f - q\|_{W^{s,\infty}([-1,1]^K)} \\ &\quad + \|I_n^{\text{Lob}}\|_{L^\infty, W^{s,\infty}} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \\ &\leq \left(1 + \|I_n^{\text{Lob}}\|_{W^{s,\infty}, W^{s,\infty}}\right) K \left(\frac{2C' \rho'}{\rho' - 1}\right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho'^{-n-1} \\ &\quad + \|I_n^{\text{Lob}}\|_{L^\infty, W^{s,\infty}} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)}. \end{aligned}$$

□

Remark 2.3 *The estimates in Equations (2.13) and (2.14) apply more generally also for functions of finite regularity, using Jackson-type estimates on the error of best polynomial approximation. Error estimates for interpolation of functions of finite regularity on a Clenshaw–Curtis grid could also be based on other arguments. For example, see the proof of [23, Theorem 2.1], which uses such a result to show a lower bound on the convergence rate of NNs with a smooth activation function.*

The previous error bound leads to the following bound on the approximate Chebyšev coefficients:

Corollary 2.4 *Let $K \in \mathbb{N}$, and let $f : [-1, 1]^K \rightarrow \mathbb{R}$ be a map which admits a holomorphic complex extension to the Bernstein polyellipse \mathcal{E}_ρ with $\rho = (\rho, \dots, \rho) \in (1, \infty)^K$ for some $\rho > 1$.*

We assume that an approximation \tilde{f} of f is available, and an upper bound on $\|f - \tilde{f}\|_{L^\infty([-1,1]^K)}$. Then,

$$\sum_{\mathbf{k} \in \{0, \dots, n\}^K} \left| \tilde{f}_{\mathbf{k};n} \right| \leq C(K, \rho) \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + (n+1)^K \pi^K \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)}. \quad (2.15)$$

Proof. We compute

$$\begin{aligned} \left| f_{\mathbf{k}} - \tilde{f}_{\mathbf{k};n} \right| &\leq \left| \int_{-1}^1 \cdots \int_{-1}^1 (f - \tilde{p}_{f,n}) T_{\mathbf{k}} \, d\lambda(x_1) \cdots d\lambda(x_K) \right| \\ &\leq \|f - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^K)} \|T_{\mathbf{k}}\|_{L^\infty([-1,1]^K)} \left| \int_{-1}^1 \cdots \int_{-1}^1 d\lambda(x_1) \cdots d\lambda(x_K) \right| \\ &= \pi^K \|f - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^K)}, \\ \sum_{\mathbf{k} \in \{0, \dots, n\}^K} \left| \tilde{f}_{\mathbf{k};n} \right| &\leq \sum_{\mathbf{k} \in \mathbb{N}_0^K} |f_{\mathbf{k}}| + \sum_{\mathbf{k} \in \{0, \dots, n\}^K} \left| f_{\mathbf{k}} - \tilde{f}_{\mathbf{k};n} \right| \\ &\leq \left(\frac{2\rho}{\rho-1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + (n+1)^K \pi^K \|f - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^K)} \\ &\leq \left(\frac{2\rho}{\rho-1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \left[1 + K(n+1)^K \pi^K \left(1 + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \right) \rho^{-n-1} \right] \\ &\quad + (n+1)^K \pi^K \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)}. \end{aligned}$$

Here, $\left(\frac{2\rho}{\rho-1} \right)^K \left[1 + K(n+1)^K \pi^K \left(1 + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \right) \rho^{-n-1} \right]$ is bounded by a constant $C(K, \rho) > 0$ independent of n . \square

Remark 2.5 For all $\tilde{f}, \tilde{g} \in \mathbb{C}^0([-1,1]^K)$, it follows from (2.8) and the positivity of the quadrature weights w_j , which are defined just above (2.4) and sum up to π^K , that for all $n \in \mathbb{N}$ and $\mathbf{k} \in \{0, \dots, n\}^K$

$$\left| \tilde{f}_{\mathbf{k};n} - \tilde{g}_{\mathbf{k};n} \right| \leq 2^{\mathbf{1}_{S(n)}(\mathbf{k})} \left\| \tilde{f} - \tilde{g} \right\|_{L^\infty([-1,1]^K)} \leq 2^K \left\| \tilde{f} - \tilde{g} \right\|_{L^\infty([-1,1]^K)}.$$

3 Constructive Deep Neural Network Approximation of Multivariate Holomorphic Functions

As we showed in [26] by (straightforward) tensorization of the univariate results, multivariate holomorphic maps $f : [-1,1]^K \rightarrow \mathbb{R}$ admit DNN surrogates which approximate f to any accuracy $\varepsilon > 0$ with NN size scaling as $O(|\log(\varepsilon)|^{K+1})$ (with the constant implied in $O(\cdot)$ still depending on the input dimension K).

In the present section, we shall recapitulate this result, with a new proof via multivariate, tensorized Chebyšev expansions.

3.1 Definitions and Architecture of Deep ReLU NNs

We consider *feed-forward deep neural networks*. These DNNs result from repeated application of affine mappings and a specific non-linear map. This nonlinearity is specified via the so-called *activation function* $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ of the DNN. Here, we take $\sigma(x) = \max\{0, x\}$, $x \in \mathbb{R}$, to be the *rectified linear unit (ReLU)* activation function. The architecture of the DNN comprises a fixed number of hidden layers $L \in \mathbb{N}$, numbers $N_\ell \in \mathbb{N}$ of computation nodes in layer $\ell \in$

$\{1, \dots, L+1\}$, the map $\Phi : \mathbb{R}^{N_0} \rightarrow \mathbb{R}^{N_{L+1}}$ is said to be realized by a feedforward neural network, if for certain *weights* $A_{i,j}^\ell \in \mathbb{R}$, and *biases* $b_j^\ell \in \mathbb{R}$ it holds for all $\mathbf{x} = (x_i)_{i=1}^{N_0}$

$$w_j^1 = \sigma \left(\sum_{i=1}^{N_0} A_{i,j}^1 x_i + b_j^1 \right), \quad j \in \{1, \dots, N_1\},$$

and

$$w_j^{\ell+1} = \sigma \left(\sum_{i=1}^{N_\ell} A_{i,j}^{\ell+1} w_i^\ell + b_j^{\ell+1} \right), \quad \ell \in \{1, \dots, L-1\}, \quad j \in \{1, \dots, N_{\ell+1}\},$$

and finally

$$\Phi(\mathbf{x}) = (w_j^{L+1})_{j=1}^{N_{L+1}} = \left(\sum_{i=1}^{N_L} A_{i,j}^{L+1} w_i^L + b_j^{L+1} \right)_{j=1}^{N_{L+1}}.$$

In this case N_0 is the dimension of the DNN input, and N_{L+1} is the dimension of the output. The acronym NN for neural network shall also be used. The number of hidden layers L of a DNN is referred to as its *depth*, denoted by $\text{depth}(\Phi)$. Also, we define the total number of nonzero weights and biases as the *size* of the DNN, i.e. $\text{size}(\Phi) := |\{(i, j, \ell) : A_{i,j}^\ell \neq 0\}| + |\{(j, \ell) : b_j^\ell \neq 0\}|$.

We do emphasize that different DNN architectures, weights and biases can realize the same function. In this paper, we focus on a constructive procedure to find weights and biases such that the resulting DNN as a function approximates a given holomorphic map to a specified accuracy.

To construct such networks from smaller subnetworks, we use NN concatenation from [27, Remark 2.6], parallelization from [27, Definition 2.7] and [11, Setting 5.2], and networks emulating the identity from [27, Lemma 2.3], which we all recall below.

Let f and g be two NNs with the same depth $L \in \mathbb{N}_0$ and the same input dimension $n \in \mathbb{N}$. Denote by m_f the output dimension of f and by m_g the output dimension of g . There exists a neural network (f, g) , called *parallelization* of f and g , which in parallel emulates f and g , i.e.

$$(f, g) : \mathbb{R}^n \rightarrow \mathbb{R}^{m_f} \times \mathbb{R}^{m_g} : \mathbf{x} \mapsto (f(\mathbf{x}), g(\mathbf{x})),$$

and it satisfies $\text{depth}((f, g)) = L$ and $\text{size}((f, g)) = \text{size}(f) + \text{size}(g)$ ([27, Definition 2.7]).

Let f and g be ReLU NNs, such that the number of nodes in the output layer of g equals the number of nodes in the input layer of f . Denote by n the number of nodes in the input layer of g , and by m the number of nodes in the output layer of f . There exists a NN $f \circ g$, called *sparse concatenation of the NNs f and g* , which we will refer to simply as *concatenation of f and g* , which realizes the composition of f and g , i.e.

$$f \circ g : \mathbb{R}^n \rightarrow \mathbb{R}^m : \mathbf{x} \mapsto f(g(\mathbf{x})). \quad (3.1)$$

It satisfies $\text{depth}(f \circ g) = \text{depth}(f) + 1 + \text{depth}(g)$ and $\text{size}(f \circ g) \leq 2 \text{size}(f) + 2 \text{size}(g)$ ([27, Remark 2.6]).

Next, let f and g be two NNs with the same depth $L \in \mathbb{N}_0$, whose input dimensions n_f and n_g may be different, and whose output dimensions we will denote by m_f and m_g , respectively. There exists a NN $(f, g)_d$, called *full parallelization of networks with distinct inputs* of f and g , which in parallel emulates f and g , i.e.

$$(f, g)_d : \mathbb{R}^{n_f} \times \mathbb{R}^{n_g} \rightarrow \mathbb{R}^{m_f} \times \mathbb{R}^{m_g} : (\mathbf{x}, \tilde{\mathbf{x}}) \mapsto (f(\mathbf{x}), g(\tilde{\mathbf{x}})).$$

It satisfies $\text{depth}((f, g)_d) = L$ and $\text{size}((f, g)_d) = \text{size}(f) + \text{size}(g)$ ([11, Setting 5.2]).

Finally, for all $n \in \mathbb{N}$ and $L \in \mathbb{N}_0$ there exists an *identity network* $\Phi_{n,L}^{\text{Id}}$ of depth L which emulates the identity map $\text{Id}_{\mathbb{R}^n} : \mathbb{R}^n \rightarrow \mathbb{R}^n : \mathbf{x} \mapsto \mathbf{x}$. It satisfies $\text{size}(\Phi_{n,L}^{\text{Id}}) \leq 2n(L+1)$ ([27, Lemma 2.3]).

3.2 ReLU NN Approximations of Chebyšev Polynomials

We construct the ReLU NN approximation of f as a reapproximation of $\tilde{p}_{f,n}$ defined in (2.10). The coefficients of $\tilde{p}_{f,n}$ can be computed with (2.9). The polynomial approximation $\tilde{p}_{f,n}$ is a tensor product of univariate Chebyšev polynomials. We use the ReLU NN approximation of univariate Chebyšev polynomials from [24], and for the ReLU NN approximation of products with multiple arguments, we recall [26, Proposition 2.6], based on [31, Proposition 3.1].

Lemma 3.1 ([24]) *There exists $C > 0$ such that for all $n \in \mathbb{N}$ there exist ReLU NNs $\left\{ \Phi_\delta^{\text{Cheb},n} \right\}_{\delta \in (0,1)}$ with input dimension one and output dimension n which satisfy*

$$\begin{aligned} \left\| T_\ell - \left(\Phi_\delta^{\text{Cheb},n} \right)_\ell \right\|_{W^{1,\infty}([-1,1])} &\leq \delta, \quad \ell = 1, \dots, n, \\ \text{depth} \left(\Phi_\delta^{\text{Cheb},n} \right) &\leq C(1 + \log(n)) \log(1/\delta) + C(1 + \log(n))^3, \\ \text{size} \left(\Phi_\delta^{\text{Cheb},n} \right) &\leq Cn \log(1/\delta) + Cn(1 + \log(n)). \end{aligned}$$

The construction of these networks is described in Appendix A.

Proposition 3.2 ([26, Proposition 2.6]) *For any $\delta \in (0, 1)$, $n \in \mathbb{N}$ and $M \geq 1$ there exists a ReLU NN $\tilde{\Pi}_{\delta,M}^n : [-M, M]^n \rightarrow \mathbb{R}$ such that*

$$\sup_{(x_i)_{i=1}^n \in [-M, M]^n} \left| \prod_{j=1}^n x_j - \tilde{\Pi}_{\delta,M}^n(x_1, \dots, x_n) \right| \leq \delta, \quad (3.2)$$

$$\text{ess sup}_{(x_i)_{i=1}^n \in [-M, M]^n} \sup_{i=1, \dots, n} \left| \frac{\partial}{\partial x_i} \prod_{j=1}^n x_j - \frac{\partial}{\partial x_i} \tilde{\Pi}_{\delta,M}^n(x_1, \dots, x_n) \right| \leq \delta. \quad (3.3)$$

There exists a constant C independent of $\delta \in (0, 1)$, $n \in \mathbb{N}$ and $M \geq 1$ such that

$$\text{size} \left(\tilde{\Pi}_{\delta,M}^n \right) \leq C(1 + n \log(nM^n/\delta)) \quad \text{and} \quad \text{depth} \left(\tilde{\Pi}_{\delta,M}^n \right) \leq C(1 + \log(n) \log(nM^n/\delta)). \quad (3.4)$$

Similar to [26, Proposition 2.13], there holds the following result on deep ReLU NN emulation rates for Chebyšev polynomials, which is of independent interest. As compared to the results in [26] and as observed in [32], the functional structure of Chebyšev polynomials affords gains in DNN depth and size bounds.

Proposition 3.3 *There exists a constant $C > 0$, such that for every $K \in \mathbb{N}$, every finite subset $\Lambda \subset \mathbb{N}_0^K$ and every $\delta \in (0, 1)$ there exists a ReLU NN $\Phi_{\Lambda,\delta}$ with input dimension K and output dimension $|\Lambda|$, such that the outputs of $\Phi_{\Lambda,\delta}$, which we denote by $\{\tilde{T}_{\mathbf{k},\delta}\}_{\mathbf{k} \in \Lambda}$, satisfy*

$$\forall \mathbf{k} \in \Lambda : \quad \left\| T_{\mathbf{k}} - \tilde{T}_{\mathbf{k},\delta} \right\|_{W^{1,\infty}([-1,1]^K)} \leq \delta,$$

$$\begin{aligned} \text{depth}(\Phi_{\Lambda,\delta}) &\leq C(1 + \log m_\infty(\Lambda))^3 + C(1 + \log(K) + \log m_\infty(\Lambda)) \log(1/\delta) + CK \log(m_\infty(\Lambda)) \\ &\quad + CK \log K, \end{aligned}$$

$$\begin{aligned} \text{size}(\Phi_{\Lambda,\delta}) &\leq CK|\Lambda| \log(m_\infty(\Lambda)) + CK|\Lambda| \log(1/\delta) + CK^2|\Lambda| \\ &\quad + CKm_\infty(\Lambda) \log(m_\infty(\Lambda)) + CKm_\infty(\Lambda) \log(1/\delta) + CK^2m_\infty(\Lambda). \end{aligned}$$

Proof. This proof consists of three steps. In the first step, we define the network. In the second step, we estimate the error. In the third step, we give bounds on the network depth and size.

Step 1. We construct the network $\Phi_{\Lambda,\delta}$ as the concatenation of two subnetworks:

$$\Phi_{\Lambda,\delta} := \Phi_{\Lambda,\delta}^{(1)} \circ \Phi_{\Lambda,\delta}^{(2)}.$$

See figure 3.1 for a sketch of the network structure.

For $\beta = \frac{1}{2}\delta(1+\delta)^{-K}K^{-1}(m_\infty(\Lambda)^2+1)^{-1} \leq \delta$, the network $\Phi_{\Lambda,\delta}^{(2)}$ in parallel approximates univariate Chebyšev polynomials up to degree $m_\infty(\Lambda)$, in all K input variables, and is based on Lemma 3.1: Denoting the input of $\Phi_{\Lambda,\delta}$ by $\mathbf{y} = (y_1, \dots, y_K) \in [-1, 1]^K$,

$$\Phi_{\Lambda,\delta}^{(2)} := \left(\Phi_\beta^{\text{Cheb}, m_\infty(\Lambda)}, \dots, \Phi_\beta^{\text{Cheb}, m_\infty(\Lambda)} \right)_d,$$

which contains K copies of $\Phi_\beta^{\text{Cheb}, m_\infty(\Lambda)}$, the j 'th of which receives y_j as input, for $j = 1, \dots, K$.

For $\beta' = \frac{1}{2}\delta(m_\infty(\Lambda)^2+1)^{-1}$, the network $\Phi_{\Lambda,\delta}^{(1)}$ computes tensor products of univariate Chebyšev polynomials using networks from Proposition 3.2. Denoting the output of $\Phi_{\Lambda,\delta}^{(2)}$ by $\left\{ \tilde{T}_{k,\beta}(y_j) \right\}_{\substack{j=1,\dots,K \\ k=1,\dots,m_\infty(\Lambda)}}$, it holds that

$$\Phi_{\Lambda,\delta}^{(1)} \circ \Phi_{\Lambda,\delta}^{(2)}(y_1, \dots, y_K) = \left(\left\{ \prod_{\beta', 1+\delta}^K \left(\tilde{T}_{k_1,\beta}(y_1), \dots, \tilde{T}_{k_K,\beta}(y_K) \right) \right\}_{\mathbf{k} \in \Lambda} \right),$$

where, in case $k_j = 0$, the factor $\tilde{T}_{0,\beta} \equiv 1$ is implemented by a bias in the first layer of $\prod_{\beta', 1+\delta}^K$. We used that

$$\left\| \tilde{T}_{k_j,\beta} \right\|_{L^\infty([-1,1])} \leq \left\| T_{k_j} \right\|_{L^\infty([-1,1])} + \left\| T_{k_j} - \tilde{T}_{k_j,\beta} \right\|_{L^\infty([-1,1])} \leq 1 + \beta \leq 1 + \delta \leq 2,$$

so that $\left\{ \tilde{T}_{k_j,\beta}(y_j) \right\}_{j=1}^K$ can be used as inputs of $\prod_{\beta', 1+\delta}^K$.

Step 2. The error can be estimated as follows:

$$\begin{aligned} & \sup_{\mathbf{y} \in [-1,1]^K} \left| T_{\mathbf{k}}(\mathbf{y}) - \tilde{T}_{\mathbf{k},\delta}(\mathbf{y}) \right| \\ & \leq \sup_{\mathbf{y} \in [-1,1]^K} \left| T_{\mathbf{k}}(\mathbf{y}) - \prod_{j=1}^K \tilde{T}_{k_j,\beta}(y_j) \right| \\ & \quad + \sup_{\mathbf{y} \in [-1,1]^K} \left| \prod_{j=1}^K \tilde{T}_{k_j,\beta}(y_j) - \prod_{\beta', 1+\delta}^K \left(\left\{ \tilde{T}_{k_j,\beta}(y_j) \right\}_{j=1}^K \right) \right| \\ & \leq \sup_{\mathbf{y} \in [-1,1]^K} \sum_{i=1}^K \left| \prod_{j=1, \dots, i-1} \tilde{T}_{k_j,\beta}(y_j) \right| \cdot \left| T_{k_i}(y_i) - \tilde{T}_{k_i,\beta}(y_i) \right| \cdot \left| \prod_{j=i+1, \dots, K} T_{k_j}(y_j) \right| + \beta' \\ & \leq \sum_{i=1}^K (1+\delta)^{i-1} \beta + \beta' \leq K(1+\delta)^K \beta + \beta' \leq \delta. \end{aligned}$$

For the estimate on the error in the derivative, we consider only the derivative with respect to y_1 . The derivatives with respect to the other inputs satisfy analogous bounds. Using that

$$\left| \tilde{T}_{k_j,\beta} \right|_{W^{1,\infty}([-1,1])} \leq \left| T_{k_j} \right|_{W^{1,\infty}([-1,1])} + \left| T_{k_j} - \tilde{T}_{k_j,\beta} \right|_{W^{1,\infty}([-1,1])} \leq m_\infty(\Lambda)^2 + \beta \leq m_\infty(\Lambda)^2 + 1,$$

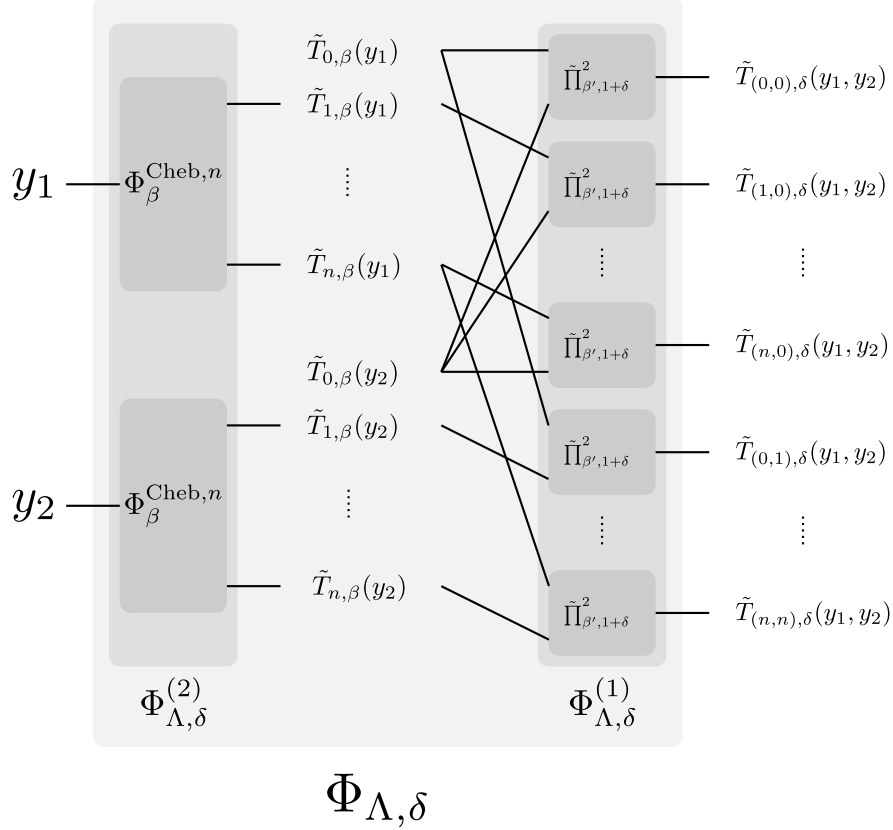


Figure 3.1: Sketch of $\Phi_{\Lambda, \delta}$ for $\Lambda = \{0, \dots, n\}^K$, with $n \in \mathbb{N}$, $K = 2$, and $\delta \in (0, 1)$. The subnetwork $\Phi_{\beta}^{\text{Cheb}, n}$ from Lemma 3.1 approximates univariate Chebyshev polynomials of degrees $1, \dots, n$, which are then multiplied by the networks $\tilde{\Pi}_{\beta', 1+\delta}^2$ from Proposition 3.2 to obtain approximations of tensor product Chebyshev polynomials. The Chebyshev polynomial of degree 0 is identically equal to 1 and implemented by a bias in the first layer of the product networks $\tilde{\Pi}_{\beta', 1+\delta}^2$.

we obtain

$$\begin{aligned}
& \operatorname{ess\,sup}_{\mathbf{y} \in [-1, 1]^K} \left| \frac{\partial}{\partial y_1} T_{\mathbf{k}}(\mathbf{y}) - \frac{\partial}{\partial y_1} \tilde{T}_{\mathbf{k}, \delta}(\mathbf{y}) \right| \\
& \leq \operatorname{ess\,sup}_{\mathbf{y} \in [-1, 1]^K} \left| \frac{\partial}{\partial y_1} T_{\mathbf{k}}(\mathbf{y}) - \frac{\partial}{\partial y_1} \prod_{j=1}^K \tilde{T}_{k_j, \beta}(y_j) \right| \\
& \quad + \operatorname{ess\,sup}_{\mathbf{y} \in [-1, 1]^K} \left| \frac{\partial}{\partial y_1} \prod_{j=1}^K \tilde{T}_{k_j, \beta}(y_j) - \frac{\partial}{\partial y_1} \tilde{\Pi}_{\beta', 1+\delta}^K \left(\left\{ \tilde{T}_{k_j, \beta}(y_j) \right\}_{j=1}^K \right) \right| \\
& \leq \operatorname{ess\,sup}_{\mathbf{y} \in [-1, 1]^K} \left| \frac{\partial}{\partial y_1} T_{k_1}(y_1) - \frac{\partial}{\partial y_1} \tilde{T}_{k_1, \beta}(y_1) \right| \cdot \left| \prod_{j=2, \dots, K} T_{k_j}(y_j) \right|
\end{aligned}$$

$$\begin{aligned}
& + \operatorname{ess\,sup}_{\mathbf{y} \in [-1,1]^K} \sum_{i=2,\dots,K} \left| \frac{\partial}{\partial y_1} \tilde{T}_{k_1,\beta}(y_1) \right| \cdot \left| \prod_{j=2,\dots,i-1} \tilde{T}_{k_j,\beta}(y_j) \right| \cdot |T_{k_i}(y_i) - \tilde{T}_{k_i,\beta}(y_i)| \cdot \left| \prod_{j=i+1,\dots,K} T_{k_j}(y_j) \right| \\
& + \operatorname{ess\,sup}_{\mathbf{y} \in [-1,1]^K} \left| \prod_{j=2,\dots,K} \tilde{T}_{k_j,\beta}(y_j) - \left(\frac{\partial}{\partial x_1} \tilde{\prod}_{\beta',1+\delta}^K \right) \left(\left\{ \tilde{T}_{k_j,\beta}(y_j) \right\}_{j=1}^K \right) \right| \cdot \left| \frac{\partial}{\partial y_1} \tilde{T}_{k_1,\beta}(y_1) \right| \\
& \leq \beta + \sum_{i=2}^K (m_\infty(\Lambda)^2 + 1)(1 + \delta)^{i-2} \beta + \beta' (m_\infty(\Lambda)^2 + 1) \\
& \leq K(1 + \delta)^K \beta (m_\infty(\Lambda)^2 + 1) + \beta' (m_\infty(\Lambda)^2 + 1) \\
& \leq \frac{\delta}{2} + \frac{\delta}{2} \leq \delta,
\end{aligned}$$

where $\frac{\partial}{\partial x_1} \tilde{\prod}_{\beta',1+\delta}^K$ denotes the (weak) derivative of $\tilde{\prod}_{\beta',1+\delta}^K : [-1 - \delta, 1 + \delta]^K \rightarrow \mathbb{R}$ with respect to its first argument, cf. Proposition 3.2.

Step 3. As a bound on the network depth and size we obtain, using $\log(1 + \delta) \leq 1$ in the last step,

$$\begin{aligned}
\operatorname{depth}(\Phi_{\Lambda,\delta}) & \leq \operatorname{depth}(\Phi_{\Lambda,\delta}^{(1)}) + 1 + \operatorname{depth}(\Phi_{\Lambda,\delta}^{(2)}) \\
& \leq \operatorname{depth}\left(\tilde{\prod}_{\beta',1+\delta}^K\right) + 1 + \operatorname{depth}\left(\Phi_\beta^{\operatorname{Cheb},m_\infty(\Lambda)}\right) \\
& \leq C(1 + \log(K) \log(K(1 + \delta)^K / \beta')) + 1 \\
& \quad + C(1 + \log m_\infty(\Lambda)) \log(1/\beta) + C(1 + \log m_\infty(\Lambda))^3 \\
& \leq C(\log(K)^2 + \log(K)K \log(1 + \delta) + \log(K) \log(m_\infty(\Lambda)) + \log(K) \log(1/\delta)) \\
& \quad + C((1 + \log m_\infty(\Lambda))^2 + (\log(K) + K \log(1 + \delta))(1 + \log m_\infty(\Lambda)) \\
& \quad + (1 + \log m_\infty(\Lambda)) \log(1/\delta)) + C(1 + \log m_\infty(\Lambda))^3 \\
& \leq C(1 + \log m_\infty(\Lambda))^3 + C(1 + \log(K) + \log m_\infty(\Lambda)) \log(1/\delta) + CK \log(m_\infty(\Lambda)) \\
& \quad + CK \log K, \\
\operatorname{size}(\Phi_{\Lambda,\delta}) & \leq C \left[\operatorname{size}(\Phi_{\Lambda,\delta}^{(1)}) + \operatorname{size}(\Phi_{\Lambda,\delta}^{(2)}) \right] \\
& \leq C \left[|\Lambda| \operatorname{size}\left(\tilde{\prod}_{\beta',1+\delta}^K\right) + K \operatorname{size}\left(\Phi_\beta^{\operatorname{Cheb},m_\infty(\Lambda)}\right) \right] \\
& \leq C \left[|\Lambda| (1 + K \log(K(1 + \delta)^K / \beta')) + K (m_\infty(\Lambda) \log(1/\beta) \right. \\
& \quad \left. + m_\infty(\Lambda)(1 + \log m_\infty(\Lambda))) \right] \\
& \leq C \left[|\Lambda| (K \log(K) + K^2 \log(1 + \delta) + K \log(m_\infty(\Lambda)) + K \log(1/\delta)) \right. \\
& \quad \left. + K (m_\infty(\Lambda) \log(m_\infty(\Lambda)) + m_\infty(\Lambda) \log(1/\delta) + (\log(K) + K \log(1 + \delta)) m_\infty(\Lambda) \right. \\
& \quad \left. + m_\infty(\Lambda)(1 + \log m_\infty(\Lambda))) \right] \\
& \leq C \left[K|\Lambda| \log(m_\infty(\Lambda)) + K|\Lambda| \log(1/\delta) + K^2|\Lambda| + Km_\infty(\Lambda) \log(m_\infty(\Lambda)) \right. \\
& \quad \left. + Km_\infty(\Lambda) \log(1/\delta) + K^2 m_\infty(\Lambda) \right].
\end{aligned}$$

This finishes the proof. \square

Remark 3.4 For $\delta \leq \exp(1/K) - 1$ it holds $K \log(1 + \delta) \leq 1$, so that the network depth is of the order $\log(K)^2$ and the size of the order $K \log(K)$.

3.3 Construction of ReLU NN Approximations of Multivariate Holomorphic Functions

Proposition 3.5 *Let $K \in \mathbb{N}$, and let $f : [-1, 1]^K \rightarrow \mathbb{R}$ be a map which admits a holomorphic complex extension to the Bernstein polyellipse \mathcal{E}_ρ with $\rho = (\rho, \dots, \rho) \in (1, \infty)^K$ for some $\rho > 1$. We assume that an approximation \tilde{f} of f is available, and an upper bound on $\|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$.*

Then, for all $n \in \mathbb{N}$ there exists a ReLU NN $\Phi_n^{\tilde{f}}$ with input dimension K and output dimension one such that for all $\rho' \in (1, \rho)$, there exists a constant $C > 0$ depending on K , ρ and ρ' and a constant $c > 0$ depending on ρ , but not on K and ρ' , such that

$$\begin{aligned} \|f - \Phi_n^{\tilde{f}}\|_{L^\infty([-1, 1]^K)} &\leq C \left(\max_{z \in \mathcal{E}_\rho} |f(z)| + \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)} \right) \rho'^{-n} \\ &\quad + \left(\frac{2}{\pi} \log(n+1) + 1 \right)^K \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}, \\ \|f - \Phi_n^{\tilde{f}}\|_{W^{1, \infty}([-1, 1]^K)} &\leq C \left(\max_{z \in \mathcal{E}_\rho} |f(z)| + \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)} \right) \rho'^{-n} \\ &\quad + n^2 \left(\frac{2}{\pi} \log(n+1) + 1 \right)^K \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}, \\ \text{depth} \left(\Phi_n^{\tilde{f}} \right) &\leq cKn(1 + \log(Kn)), \quad \text{size} \left(\Phi_n^{\tilde{f}} \right) \leq cK^2(n+1)^{K+1}. \end{aligned}$$

The NN weights and biases only depend on f through $\{\tilde{f}(\mathbf{x}_j^n)\}_{j \in \{0, \dots, n\}^K}$, where $\mathbf{x}_j^n = \cos(j\pi/n)$.

This means that for C as in Proposition 3.5 and for some $b > 0$ depending on ρ and ρ' , but not on K , that for all $n \in \mathbb{N}$ holds

$$\begin{aligned} \|f - \Phi_n^{\tilde{f}}\|_{W^{1, \infty}([-1, 1]^K)} &\leq C \left(\max_{z \in \mathcal{E}_\rho} |f(z)| + \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)} \right) \exp \left(-b \left(\text{size} \left(\Phi_n^{\tilde{f}} \right) \right)^{1/(K+1)} \right) \\ &\quad + n^2 \left(\frac{2}{\pi} \log(n+1) + 1 \right)^K \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}. \end{aligned} \quad (3.5)$$

The L^∞ -error satisfies the same estimate, even without the factor n^2 in the second term. Equation (3.5) follows from the error estimate in Proposition 3.5 and from the following estimate, where we assume, without loss of generality, that $c \geq 1$:

$$\begin{aligned} \log(\rho')n &\geq \log(\rho')c^{-1/(K+1)}K^{-2/(K+1)} \left(\text{size} \left(\Phi_n^{\tilde{f}} \right) \right)^{1/(K+1)} \\ &\geq \log(\rho')c^{-1}c' \left(\text{size} \left(\Phi_n^{\tilde{f}} \right) \right)^{1/(K+1)} =: b \left(\text{size} \left(\Phi_n^{\tilde{f}} \right) \right)^{1/(K+1)}. \end{aligned}$$

We also used the existence of $c' > 0$ such that $x^{-2/(x+1)} = \exp(-2 \log(x)/(x+1)) \geq c'$ for all $x \in [1, \infty)$.

Proof of Proposition 3.5. We consider the approximation $\tilde{p}_{f,n}$ of f from (2.10). For the ReLU NN approximation of tensor product Chebyshev polynomials, we apply Proposition 3.3 with $\Lambda = \{0, \dots, n\}^K$, $|\Lambda| = (n+1)^K$, $m_\infty(\Lambda) = n$ and $\delta = \rho^{-n}(n+1)^{-K}$.

We write $\mathbf{k}^{(\ell)}$ for the element of Λ satisfying $(\Phi_{\Lambda, \delta})_\ell = \tilde{T}_{\mathbf{k}^{(\ell)}, \delta}$ (the latter is defined in Proposition 3.3). With that notation, we define $A \in \mathbb{R}^{1 \times |\Lambda|}$ by $A_{1, \ell} = \tilde{f}_{\mathbf{k}^{(\ell)}, n}$ for $\ell = 1, \dots, |\Lambda|$. With $b := [0] \in \mathbb{R}^1$, we define

$$\Phi_n^{\tilde{f}} := ((A, b)) \circ \Phi_{\Lambda, \delta}$$

and observe that $\text{depth}(((A, b))) = 0$ and $\text{size}(((A, b))) \leq |\Lambda|$. Its realization satisfies for some $C > 0$ depending on K and ρ

$$\begin{aligned}
\Phi_n^{\tilde{f}} &= \sum_{\mathbf{k} \in \Lambda} \tilde{f}_{\mathbf{k};n} \tilde{T}_{\mathbf{k},\delta}, \\
\left\| \tilde{p}_{f,n} - \Phi_n^{\tilde{f}} \right\|_{W^{1,\infty}([-1,1]^K)} &\leq \sum_{\mathbf{k} \in \Lambda} |\tilde{f}_{\mathbf{k};n}| \left\| T_{\mathbf{k}} - \tilde{T}_{\mathbf{k},\delta} \right\|_{W^{1,\infty}([-1,1]^K)} \\
&\leq \delta \sum_{\mathbf{k} \in \Lambda} |\tilde{f}_{\mathbf{k};n}| \\
&\leq \delta C \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + \delta(n+1)^K \pi^K \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \\
&\leq C \rho^{-n} \left(\max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \right), \quad (3.6)
\end{aligned}$$

where we used the bound (2.15) on the sum of the coefficients. With Lemma 2.2, for some $C > 0$ depending on K and ρ , it further follows that

$$\begin{aligned}
\left\| f - \Phi_n^{\tilde{f}} \right\|_{L^\infty([-1,1]^K)} &\leq \|f - \tilde{p}_{f,n}\|_{L^\infty([-1,1]^K)} + \left\| \tilde{p}_{f,n} - \Phi_n^{\tilde{f}} \right\|_{L^\infty([-1,1]^K)} \\
&\leq \left(1 + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \right) K \left(\frac{2\rho}{\rho-1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho^{-n-1} \\
&\quad + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \\
&\quad + C \rho^{-n} \left(\max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \right) \\
&\leq C \left(\max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \right) \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \rho^{-n} \\
&\quad + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)}.
\end{aligned}$$

For all $\rho' \in (1, \rho)$ there exists a constant C depending on ρ, ρ' and K such that $\|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \rho^{-n} \leq \left(\frac{2}{\pi} \log(n+1) + 1\right)^K \rho^{-n} \leq C \rho'^{-n}$ for all $n \in \mathbb{N}$, resulting in the error bound stated in the proposition. Similarly, for every $\rho' \in (1, \rho)$ and with $C'(\rho, \rho') > 0$ as in Lemma 2.2, with Equation (3.6) it follows that there exists $C(K, \rho, \rho') > 0$ such that

$$\begin{aligned}
&\left\| f - \Phi_n^{\tilde{f}} \right\|_{W^{1,\infty}([-1,1]^K)} \\
&\leq \|f - \tilde{p}_{f,n}\|_{W^{1,\infty}([-1,1]^K)} + \left\| \tilde{p}_{f,n} - \Phi_n^{\tilde{f}} \right\|_{W^{1,\infty}([-1,1]^K)} \\
&\leq \left(1 + \|I_n^{\text{Lob}}\|_{W^{1,\infty}, W^{1,\infty}} \right) K \left(\frac{2C'\rho'}{\rho'-1} \right)^K \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| \rho'^{-n-1} \\
&\quad + \|I_n^{\text{Lob}}\|_{L^\infty, W^{1,\infty}} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \\
&\quad + C \rho^{-n} \left(\max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \right) \\
&\leq C \rho'^{-n} \|I_n^{\text{Lob}}\|_{W^{1,\infty}, W^{1,\infty}} \max_{\mathbf{z} \in \mathcal{E}_\rho} |f(\mathbf{z})| + C \rho^{-n} \|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)} \\
&\quad + \|I_n^{\text{Lob}}\|_{L^\infty, W^{1,\infty}} \|f - \tilde{f}\|_{L^\infty([-1,1]^K)}.
\end{aligned}$$

Reasoning as for the bound on the L^∞ -error, for all $\rho'' \in (1, \rho')$ there exists a constant $C > 0$ depending on K, ρ, ρ' and ρ'' such that $\|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \rho^{-n} \leq C \rho''^{-n}$ and $\|I_n^{\text{Lob}}\|_{W^{1,\infty}, W^{1,\infty}} \rho'^{-n} \leq$

$n^2(\frac{2}{\pi} \log(n) + 1)^K \rho'^{-n} \leq C \rho''^{-n}$ for all $n \in \mathbb{N}$. Using these estimates and Equation (2.7), the bound in the proposition is obtained when we write ρ' instead of ρ'' . To bound NN depth and size, we observe that there is $c > 0$ depending on ρ , but not on K , such that for every $n \in \mathbb{N}$,

$$\begin{aligned} \text{depth}(\Phi_n^{\tilde{f}}) &\leq \text{depth}(((A, b))) + 1 + \text{depth}(\Phi_{\Lambda, \delta}) \\ &\leq 1 + c(1 + \log m_\infty(\Lambda))^3 + c(1 + \log(K) + \log m_\infty(\Lambda)) \log(1/\delta) \\ &\quad + cK \log(m_\infty(\Lambda)) + cK \log(K) \\ &\leq cKn(1 + \log(Kn)), \end{aligned} \tag{3.7}$$

$$\begin{aligned} \text{size}(\Phi_n^{\tilde{f}}) &\leq c \text{size}(((A, b))) + c \text{size}(\Phi_{\Lambda, \delta}) \\ &\leq c|\Lambda| + c|\Lambda|K \log(m_\infty(\Lambda)) + c|\Lambda|K \log(1/\delta) + c|\Lambda|K^2 + cKm_\infty(\Lambda) \log(m_\infty(\Lambda)) \\ &\quad + cKm_\infty(\Lambda) \log(1/\delta) + cK^2 m_\infty(\Lambda) \\ &\leq cK^2(n+1)^{K+1}. \end{aligned} \tag{3.8}$$

□

Remark 3.6 For $K \in \mathbb{N}$ assume that $f, g : [-1, 1]^K \rightarrow \mathbb{R}$ satisfy the requirements of Proposition 3.5 for the same $\rho \in (1, \infty)$, with available approximations \tilde{f} and \tilde{g} . For all $n \in \mathbb{N}$, the networks $\Phi_n^{\tilde{f}}$ and $\Phi_n^{\tilde{g}}$ have the same architecture. In fact, all layers but the output layer are identical, and equal the DNN $\Phi_{\Lambda, \delta}$ from Proposition 3.3, for $\Lambda = \{0, \dots, n\}^K$ and $\delta = \rho^{-n}(n+1)^{-K}$. The weights in the output layer equal $\{\tilde{f}_{\mathbf{k};n}\}_{\mathbf{k} \in \Lambda}$ and $\{\tilde{g}_{\mathbf{k};n}\}_{\mathbf{k} \in \Lambda}$, respectively. These depend linearly on \tilde{f}, \tilde{g} , thus in particular continuously. For example, with Remark 2.5 it follows that

$$\max_{\mathbf{k} \in \Lambda} |\tilde{f}_{\mathbf{k};n} - \tilde{g}_{\mathbf{k};n}| \leq 2^K \|\tilde{f} - \tilde{g}\|_{L^\infty([-1, 1]^K)}.$$

3.4 More General Activations

The results derived here should extend to DNNs with activation functions other than ReLU. Specifically, the DNNs constructed in our proofs essentially consist of concatenations and parallelizations of identity networks and DNN approximations of products. The proposed ReLU DNN constructions therefore extend to DNNs with other activation functions, as long as they allow an exact emulation or efficient approximation of identity networks and products. We discuss several particular cases.

Poly-ReLU: Exact emulation of both is possible with DNNs with the *rectified power unit* (RePU) activation function $\sigma_r : x \mapsto \max\{0, x\}^r$, for $r \in \mathbb{N}$, $r \geq 2$ ([18]). Polynomial emulation with such networks was studied e.g. in [22, Proof of Theorem 3.3], [18]. For DNNs with the RePU activation function, efficient polynomial emulation based on the recursion in Equation (A.1) was already obtained in [32].

Rational Activation: Exact emulation of the identity and of multiplications is also possible by DNNs with rational activation. In [4], networks were studied which have in each computational node $\sigma = p/q$ as activation, for polynomials p, q of prescribed degrees, but whose coefficients can be trained and may be different for each node. Such networks can emulate both the identity and products exactly if the prescribed degrees satisfy $\deg(p) \geq 2$ and $\deg(q) \in \mathbb{N}_0$ (cf. [4, Proposition 10] and its proof). Because Lemma 2.2 holds for all $s \in \mathbb{N}$, networks whose activation function allows for the exact emulation of polynomials obtain exponential convergence with respect to the $W^{s, \infty}$ -norm for a larger range of $s \in \mathbb{N}$, with only the constant C in the error bound depending on s .

Sigmoidal Activation: The present results for RePU networks directly generalize to sigmoidal activation functions of order $k \in \mathbb{N}$, $k \geq 2$. A function $\tau : \mathbb{R} \rightarrow \mathbb{R}$ is *sigmoidal of order* $k \in \mathbb{N}_0$, if it satisfies $\lim_{x \rightarrow -\infty} \tau(x)/x^k = 0$, $\lim_{x \rightarrow \infty} \tau(x)/x^k = 1$, and $|\tau(x)| \leq C(1 + |x|^k)$,

$x \in \mathbb{R}$, for a constant $C > 0$ that does not depend on x , see for example [22, Equations (2.3) and (2.4)]. In [22, Lemma 3.6], it was shown that for every continuous sigmoidal activation function τ of order $k \geq 2$ and arbitrary $A > 0$, the RePU σ_k can be approximated on the interval $[-A, A]$ with arbitrarily small error ε with respect to the $L^\infty([-A, A])$ -norm by a τ -NN which has depth 1 and fixed network size independent of A and ε . When τ is uniformly continuous on \mathbb{R} , this result remains true w.r.t. the $L^\infty(\mathbb{R})$ -norm (instead of $L^\infty([-A, A])$), as remarked directly below [22, Lemma 3.6]. There it was also noted that a similar statement holds for the approximation of the ReLU by a NN with arbitrary sigmoidal activation function of the order $k = 1$. In the proof of [22, Lemma 3.6], it was observed that for every $k \in \mathbb{N}$ and every continuous, sigmoidal τ of order k , the τ -NN that approximates σ_k is uniformly continuous on $[-A, A]$. From this, it follows that σ_k -NNs can be approximated up to arbitrarily small $L^\infty([-A, A])$ -error ε by a τ -NN with network size independent of A and ε . If τ is uniformly continuous on \mathbb{R} , the same holds w.r.t. the $L^\infty(\mathbb{R})$ -norm.

Genuinely Nonlinear Activation: The existence of NNs that approximate a given RePU network with arbitrarily small $L^\infty([-A, A])$ -error ε and network size independent of A and ε , holds also for NNs with “genuinely nonlinear” activation function $\tau \in C^2(\mathbb{R})$, i.e. for which there exists $x \in \mathbb{R}$ where $\tau''(x) \neq 0$. In [30, Theorem 3.4] it was shown that for all $A, \varepsilon > 0$ there exists a τ -NN which approximates the product of two numbers from $[-A, A]$ up to $L^\infty([-A, A]^2)$ -error at most ε and whose network size is bounded independent of A and ε . See also [31, Section 3.3].

4 DNN Emulation Algorithms

In this section, we develop an algorithm which, assuming at hand a procedure to approximately evaluate a map $f : [-1, 1]^K \rightarrow \mathbb{R}$, will construct, for a given accuracy threshold $\varepsilon > 0$, an approximating NN whose estimated error is at most ε and whose depth and size satisfy the theoretical bounds from Section 3. Importantly, the proposed algorithm is deterministic, and does not resort to “black-box” loss-function minimization algorithms.

As before, we assume that f admits a holomorphic extension to a closed Bernstein polyellipse \mathcal{E}_ρ , with $\boldsymbol{\rho} = (\rho, \dots, \rho) \in \mathbb{R}^K$ for some $\rho > 1$. The algorithm is based on the observation that the proof of Proposition 3.5 is constructive, i.e., for $n \in \mathbb{N}$, $\Lambda = \{0, \dots, n\}^K$ and $\delta = (n+1)^{-K} \rho^{-n}$, the network $\Phi_n^{\tilde{f}} := ((A, b) \circ \Phi_{\Lambda, \delta})$ has been described explicitly. In the proof of Proposition 3.5, the accuracy $\delta = (n+1)^{-K} \rho^{-n}$ of the network $\Phi_{\Lambda, \delta}$ approximating tensor product Chebyšev polynomials is sufficient, but may be much smaller than needed, and because a priori no good estimate on ρ may be known, part of the algorithm is dedicated to determining a sufficient value of δ .

Thus, assuming access to an approximation \tilde{f} of f , the following algorithm determines values $n_* \in \mathbb{N}$ and $\delta_* \in (0, 1)$ and constructs a ReLU DNN $\Phi_{n_*, \delta_*}^{\tilde{f}}$ approximating f . When determining n_* , we note that in (2.13) the interpolation error $\|f - \tilde{p}_{f, n}\|_{L^\infty([-1, 1]^K)}$ is bounded by $\|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$, growing polylogarithmically in n , plus a term that decays exponentially with n . Because we cannot access f directly, the algorithm below cannot ensure smallness of $\|I_n^{\text{Lob}}\|_{L^\infty, L^\infty} \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$. It could be taken into account by setting an upper bound on n , we will not discuss this in detail. The NN error $\|\tilde{p}_{f, n} - \Phi_n^{\tilde{f}}\|_{L^\infty([-1, 1]^K)}$ is bounded by $\delta \sum_{\mathbf{k} \in \Lambda} |\tilde{f}_{\mathbf{k}; n}|$. In the proof of Proposition 3.5, we used that in theory, when neglecting rounding errors in the NN’s affine transformations, this error can be made small with respect to other error terms by choosing δ sufficiently small, namely $\delta = (n+1)^{-K} \rho^{-n}$. In practice, it holds that choosing a smaller δ cannot increase the error (as this holds for the networks in Proposition 3.2, and therefore also for those in Proposition 3.3). However, because the accuracy of affine transformations is limited by the machine precision, δ is in practice limited from below. Therefore, as the upper bound (2.15) on $\sum_{\mathbf{k} \in \Lambda} |\tilde{f}_{\mathbf{k}; n}|$ increases with n , the NN

error $\|\tilde{p}_{f,n} - \Phi_n^{\tilde{f}}\|_{L^\infty([-1,1]^\kappa)}$ may increase with n for large n , regardless of our choice of δ . Therefore, we use a greedy algorithm to determine moderate n_* , δ_* for which the estimated error w.r.t. \tilde{f} is smaller than a prescribed tolerance.

The construction in Algorithm 4.1 uses some elementary NNs, for which explicit constructions are available in the literature [35, 31]. In [35, Propositions 2 and 3], a NN is proposed that approximates the square denoted by $\Phi_{\ell,M}^{\text{sq}}$ and the product of two scalars denoted by $\tilde{\times}_{\ell,M}$ in a bounded interval $[-M, M]$ with ℓ layers (see also [31, Proposition 3.1] for a different construction). In Step 1 of the proof of [31, Proposition 3.3], an explicit NN is proposed that approximates the product of K scalars (already previously) denoted by $\tilde{\prod}_{\delta,M}^K$ with accuracy δ on the hypercube $[-M, M]^K$. Based on the scalar product NN $\tilde{\times}_{\ell,M}$, the NN $\Phi_\beta^{\text{Cheb},n}$ is constructed in Appendix A. Its properties are given by Lemma 3.1. Recall that this DNN has input dimension one and output dimension n such that the ℓ 'th component of the output layer approximates the Chebyšev polynomial of degree ℓ with accuracy β over the interval $[-1, 1]$. A DNN that approximates tensorized Chebyšev polynomials is constructed by a concatenation of K copies of the NN $\Phi_\beta^{\text{Cheb},n}$ with copies of the NN $\tilde{\prod}_{\delta,M}^K$ and is denoted by $\Phi_{\Lambda,\delta}$, where $\Lambda \subset \mathbb{N}_0^K$ is the index set of polynomial degrees of tensorized Chebyšev polynomials. The construction of $\Phi_{\Lambda,\delta}$ is described in detail in Step 1 of the proof of Proposition 3.3. These explicitly constructed DNNs with certified accuracy over all inputs in an a priori known hypercube¹ are used in the following algorithm. Finally, as in the proof of Proposition 3.5, let $b := [0] \in \mathbb{R}^1$ and $A \in \mathbb{R}^{1 \times |\Lambda|}$ be defined by $A_{1,\ell} = \tilde{f}_{\mathbf{k}^{(\ell)},n}$, where $\{\mathbf{k}^{(\ell)}\}_{\ell=1,\dots,|\Lambda|} \subset \Lambda$ is an enumeration of Λ such that $(\Phi_{\Lambda,\delta})_\ell = \tilde{T}_{\mathbf{k}^{(\ell)},\delta}$ (the latter is defined in Proposition 3.3). With this notation, we define $\Phi_{n,\delta}^{\tilde{f}} := ((A, b)) \circ \Phi_{\Lambda,\delta}$ for all $n \in \mathbb{N}$ and $\delta \in (0, 1)$.

Regarding line 5 of Algorithm 4.1, recall that $\cos((2n-j)\pi/n) = \cos(j\pi/n)$ for $j = 1, \dots, n-1$, thus $(f(\mathbf{x}_j^n))_{j \in \{0, \dots, n\}^\kappa}$ are sufficient to determine $(f(\mathbf{x}_j^n))_{j \in \{0, \dots, 2n-1\}^\kappa}$.

Next, we compute the complexity of the algorithm.

Proposition 4.1 *In total, under the assumption that the complexity of constructing a NN from its weights and biases and evaluating a NN are both proportional to the number of nonzero network weights and biases, the computational complexity of Algorithm 4.1 is bounded by $C [K^2(n_* + 1)^{2K+1} \log(1/\delta_*) + K^2(n_* + 1)^{2K} \log(1/\delta_*^2)]$, in addition to at most $(n_* + 1)^{K+1}$ evaluations of \tilde{f} . Here, the constant $C > 0$ does not depend on n_* , δ_* , K .*

The constructed network $\Phi_{n_*,\delta_*}^{\tilde{f}}$ satisfies

$$\begin{aligned} \text{depth}(\Phi_{n_*,\delta_*}^{\tilde{f}}) &\leq C [1 + \log(n_*)^3 + (\log(K) + \log(n_*)) \log(1/\delta_*) + K \log(n_*) + K \log(K)], \\ \text{size}(\Phi_{n_*,\delta_*}^{\tilde{f}}) &\leq C [1 + K^2(n_* + 1)^K \log(1/\delta_*) + n_* \log(n_*)]. \end{aligned}$$

Proof. Let $C > 0$ denote some generic constant. The while loop of lines 2–17 is executed at most $n_* - 2 + \log_2(\delta_0/\delta_*)$ times: after the initial execution of the while loop, it is executed $n_* - 3$ times after increasing n and $\log_2(\delta_0/\delta_*)$ times after decreasing δ .

Per execution of the while loop: Line 4 takes $(n+1)^K$ evaluations of \tilde{f} , and is executed for $n = 3, \dots, n_*$, i.e. only when state = true. Line 5 computes the IFFT of a K -dimensional array of size $(n+1)^K$. The complexity of each execution is $CK(n+1)^K \log(n) = C(n+1)^K \log((n+1)^K)$. Line 6 contains the construction of $\Phi_{\Lambda,\delta}$, which has computational cost $CK^2(n+1)^K \log_2(1/\delta)$ (see Step 3 in the proof of Proposition 3.3). Line 8 requires n^K evaluations of \tilde{f} and of $\Phi_{n,\delta}^{\tilde{f}}$, where each evaluation of $\Phi_{n,\delta}^{\tilde{f}}$ has complexity $CK^2(n+1)^K \log_2(1/\delta)$. The evaluations of \tilde{f} are the same as in line 4 (except for $n = 3$), the evaluation of $\Phi_{n,\delta}^{\tilde{f}}$ contributes $CK^2(n+1)^{2K} \log_2(1/\delta)$ for each execution of line 8.

¹The error bounds have been derived under the assumption that the affine transformations in the NNs are evaluated in exact arithmetic, without rounding.

Algorithm 4.1 Construction of a ReLU NN that approximates a function f , based on pointwise evaluations of an approximation \tilde{f} of f .

Input: $K \in \mathbb{N}$, a routine to evaluate $\tilde{f} : [-1, 1]^K \rightarrow \mathbb{R}$, a goal accuracy $\varepsilon \in (0, 1)$, $\alpha \in (0, 1)$, $\delta_0 \in (0, 1)$

Output: $\Phi_{n_*, \delta_*}^{\tilde{f}}$, for $n_* \in \mathbb{N}$ and $\delta_* \in (0, 1)$ to be determined.

```

1: Set  $n \leftarrow 3$ ,  $\delta \leftarrow \delta_0$ ,  $\text{err} \leftarrow 1$ , and  $\text{state} \leftarrow \text{true}$ 
2: while  $\text{err} > \varepsilon$  do
3:   Set  $\text{err}_{-1} \leftarrow \text{err}$ .
4:   Compute  $\{\tilde{f}(\mathbf{x}_j^n)\}_{j \in \{0, \dots, n\}^K}$ .  $\triangleright \mathbf{x}_j^n = \cos(j\pi/n)$ 
5:   Compute  $\{\tilde{f}_{\mathbf{k}; n}\}_{\mathbf{k} \in \{0, \dots, n\}^K}$  with (2.9). Store as  $((A, b))$ , as in the proof of Proposition 3.5.
6:   Construct  $\Phi_{\Lambda, \delta}$ , depending on  $n$  and  $\delta$  with  $\Lambda = \{0, \dots, n\}^K$ .
7:   Construct  $\Phi_{n, \delta}^{\tilde{f}} = ((A, b)) \circ \Phi_{\Lambda, \delta}$  with  $\Lambda = \{0, \dots, n\}^K$ .
8:   Update  $\text{err} \leftarrow \max_{x \in \{\mathbf{x}_j^{n-1} : j \in \{0, \dots, n-1\}^K\}} |\tilde{f}(x) - \Phi_{n, \delta}^{\tilde{f}}(x)|$ .  $\triangleright \mathbf{x}_j^{n-1} = \cos(j\pi/(n-1))$ 
9:   if  $\text{err} > \alpha \text{err}_{-1}$  then
10:      $\text{state} \leftarrow \text{complement}(\text{state})$ .
11:   end if
12:   if  $\text{state} = \text{true}$  then
13:      $n \leftarrow n + 1$ .
14:   else
15:      $\delta \leftarrow \delta/2$ .
16:   end if
17: end while
18: if  $\text{state} = \text{true}$  then
19:    $n_* \leftarrow n - 1$  and  $\delta_* \leftarrow \delta$ .
20: else
21:    $n_* \leftarrow n$  and  $\delta_* \leftarrow 2\delta$ .
22: end if
23: return  $\Phi_{n_*, \delta_*}^{\tilde{f}}$ .

```

In total, the computational complexity is upper bounded by $(n_* + 1)^{K+1}$ evaluations of \tilde{f} in addition to $CK^2(n_* + 1)^{2K+1} \log(1/\delta_*) + CK^2(n_* + 1)^{2K} \log(1/\delta_*)^2$, with C depending on δ_0 .

The bounds on the network depth and size follow from Proposition 3.3 and

$$\begin{aligned} \text{depth}(\Phi_{n_*, \delta_*}^{\tilde{f}}) &\leq \text{depth}(((A, b))) + 1 + \text{depth}(\Phi_{\Lambda, \delta}), \\ \text{size}(\Phi_{n_*, \delta_*}^{\tilde{f}}) &\leq C \text{size}(((A, b))) + C \text{size}(\Phi_{\Lambda, \delta}). \end{aligned}$$

□

If the first order partial derivatives of \tilde{f} can be evaluated as well, replacing line 8 by

$$\text{Update err} \leftarrow \max_{x \in \{\mathbf{x}_j^{n-1} : j \in \{0, \dots, n-1\}^K\}} \max\{|\tilde{f}(x) - \Phi_{n, \delta}^f(x)|, \max_{i=1}^K |\frac{\partial}{\partial x_i} \tilde{f}(x) - \frac{\partial}{\partial x_i} \Phi_{n, \delta}^f(x)|\}$$

will give a similar result, now with $W^{1, \infty}$ -error at most ε .

5 DNN Data-to-QoI Maps in Bayesian Inversion

In the present section, we illustrate the foregoing results, including the exponential convergence, see Proposition 3.5, in the context of PDE constrained Bayesian estimation.

Let X be a separable Banach space and π_0 be a *prior* probability measure on X . The expectation with respect to π_0 is denoted by \mathbb{E}^{π_0} . Let \mathcal{X} be a separable Banach space and let the *forward map* $\mathcal{S} : X \rightarrow \mathcal{X}$ be continuous. For given $u \in X$, we suppose that we can measure the *response* $\mathcal{S}(u)$ with observation functionals $O_i \in \mathcal{X}^*$, $i = 1, \dots, K$, where \mathcal{X}^* denotes the dual space of \mathcal{X} . We aim at a reconstruction of u from these measurements $O_i(\mathcal{S}(u))$, $i = 1, \dots, K$. In general this problem is ill-posed (e.g. [7, 16]). We consider additive Gaussian observation noise which regularizes the Bayesian inverse problem [7], i.e., we assume noisy observation data δ_i , $i = 1, \dots, K$, such that

$$\delta = O(\mathcal{S}(u)) + \eta,$$

where η is assumed to be an additive, centered, independent K -dimensional multivariate Gaussian random variable with nondegenerate covariance matrix Σ . We use the notation $\delta = (\delta_1, \dots, \delta_K)$ and $O = (O_1, \dots, O_K)$. The Bayesian estimate of a quantity of interest $\phi : X \rightarrow \mathbb{R}$, which is assumed to be continuous, is given by $\mathbb{E}^{\pi_\delta}[\phi(u)]$. Here, π_δ denotes the posterior probability conditional on given data δ .

In practice the computation of an approximation to this Bayesian estimate could be costly, in particular when the forward map may involve a solver of a partial differential equation. For the repeated efficient approximate numerical evaluation of the data-to-QoI map $\delta \mapsto \mathbb{E}^{\pi_\delta}[\phi(u)]$, surrogate maps of DNN type can lead to large computational savings. These maps can be evaluated rapidly in comparison to the pointwise, approximate evaluation of this map using, e.g., MCMC algorithms requiring one PDE solve for each proposal.

With the presently developed, explicit, DNN emulators, we are able to explicitly construct a surrogate for the Data-to-QoI map by means of a deep ReLU NN. The exponential error bound on the generalization error will imply few *parallel evaluations of the forward map, at synthetic data points* are required. This is due to the additive Gaussian noise η having a strong regularizing effect. This has been pointed out and was exploited in [15]. There in [15, Corollary 4.7] it is shown that for any $r > 0$, the map

$$[-r, r]^K \ni \delta \mapsto f(\delta) := \mathbb{E}^{\pi_\delta}[\phi(u)] \tag{5.1}$$

is analytic. Bounds on its modulus depend (exponentially) on the observation precision, i.e., on the inverse observation noise covariance. This allows quite general forward maps \mathcal{S} , which are

not required to be smooth, see [15, Section 3] for a few examples. Thus, the map $[-r, r]^K \ni \delta \mapsto f(\delta)$ satisfies the assumptions of Proposition 3.5 and Algorithm 4.1 constructs a deep ReLU NN with size and depth as in Proposition 4.1, if we are able to approximate the map f by some \tilde{f} sufficiently well.

6 Numerical Experiments

We consider the solution q to an elliptic partial differential equation given a realization of an uncertain diffusion coefficient function u as forward operator. Specifically, on a Lipschitz domain $D \subset \mathbb{R}^d$ let

$$-\nabla \cdot (u \nabla q) = g \quad q|_{\partial D} = 0, \quad (6.1)$$

where $g \in L^2(D)$ is assumed known and where $u \in \{\tilde{u} \in L^\infty(D) : \text{ess inf}_{x \in D} \tilde{u}(x) > 0\}$. For every such coefficient function u , the solution $q \in H_0^1(D) =: \mathcal{X}$ exists and is unique by the Lax–Milgram lemma. The forward data-to-solution map \mathcal{S} is defined by $\mathcal{S}(u) := q$. The observation operators $O_i, i = 1, \dots, K$, are taken as linear functionals in $H^{-1}(D) := (H_0^1(D))^*$. We construct an uncertain diffusion coefficient

$$u(\mathbf{y}) = \bar{u} + \sum_{j=1}^s y_j \psi_j,$$

where $s \in \mathbb{N}$, $\mathbf{y} \in [-1/2, 1/2]^s$ and $\text{ess inf}_{x \in D} \{\bar{u}(x)\} > \sum_{j=1}^s \|\psi_j\|_{L^\infty(D)}/2$. We assume that $\psi_j \in L^\infty(D)$, $j = 1, \dots, s$. We model the Bayesian prior as product probability measure μ on $[-1/2, 1/2]^s$ by $\mu(d\mathbf{y}) = \bigotimes_{j=1}^s dy_j$. The prior π_0 is chosen as the image measure of the random coefficient $\mathbf{y} \mapsto u(\mathbf{y})$ on $X = \{\bar{u} + \sum_{i=1}^s y_i \psi_i : \mathbf{y} \in [-1/2, 1/2]^s\}$, i.e., $\pi_0(A) := \mu(u^{-1}(A))$ for every measurable set $A \subset X$, where $u^{-1}(A)$ denotes the pre-image under the mapping $u : [-1/2, 1/2]^s \rightarrow X$. Let $\varphi \in H^{-1}(D)$. Then the QoI ϕ is given by the composition of φ and \mathcal{S} , i.e., $\phi = \varphi \circ \mathcal{S}$. According to [7, Theorem 14], the posterior expectation in (5.1) is given by

$$f(\delta) = \frac{1}{Z} \int_X \phi(u) \exp\left(-\frac{(\delta - O(\mathcal{S}(u)))^\top \Sigma^{-1} (\delta - O(\mathcal{S}(u)))}{2}\right) \pi_0(du) \quad (6.2)$$

and

$$Z = \int_X \exp\left(-\frac{(\delta - O(\mathcal{S}(u)))^\top \Sigma^{-1} (\delta - O(\mathcal{S}(u)))}{2}\right) \pi_0(du). \quad (6.3)$$

In our computations, we consider the univariate case and one or two parameters, i.e., $d = 1$, $D = (-1, 1)$, and $s = 1, 2$. Also the data space dimension is chosen to be one, i.e., $K = 1$. For given $u(y)$, the solution $q(y) = \mathcal{S}(u(y))$ is approximated by the Finite Element Method (FEM) with 65 degrees of freedom using a hat functions basis and the integrals in (6.2) and in (6.3) are efficiently approximated by a product Gauss–Legendre quadrature with 40 quadrature nodes in each parameter direction. This allows our numerical tests to focus on the algorithmic realization of deep ReLU NNs to approximate the data-to-QoI map $[-1/2, 1/2]^s \ni \delta \mapsto f(\delta)$.

In our setup, we choose $\bar{u} = 1.5$, $\psi_j(x) = 2j^{-2} \cos(j\pi x)$. The observation functional is taken to be $\mathcal{O}(v) := \int_{-1}^0 v(x) dx$ and the goal functional is taken as $\varphi(v) := \int_0^1 v(x) dx$ for every $v \in H_0^1(D)$. The right-hand side in (6.1) is taken as $g = 1$. We ran Algorithm 4.1 for given error tolerances and observed the resulting size of the constructed NN. Specifically, we ran an implementation of Algorithm 4.1 for the data-to-QoI map $\delta \rightarrow f(\delta)$ given in (6.2).

In Figures 6.1 and 6.2, we visualize the output of Algorithm 4.1. The size of the NN that is constructed in Algorithm 4.1, the depth, and the achieved accuracy are plotted in a semi-logarithmic scale. In the figures the error should impact the size with the square of the logarithm according to Proposition 3.5, which is why the square root of the size of the NN is plotted. As described in Algorithm 4.1 line 8, the error is estimated by the maximum of the absolute value

of the difference of the DNN approximation and the data-to-QoI map evaluated in Clenshaw–Curtis points of one degree less than used as training points. Note that Clenshaw–Curtis points are not nested. Because we cannot evaluate f exactly, the error was computed with respect to the previously described, quadrature based numerical approximation, corresponding to \tilde{f} in the notation of the algorithm. We tested this setting for additive Gaussian noise with variance $\sigma^2 = 0.04$, see (5.1). The step-like behavior of the size and depth of the constructed DNN that is visible in Figures 6.1 and 6.2 is a consequence of the employed construction of DNNs approximating Chebyšev polynomials, see Appendix A. There, for given polynomial degree $p \in \mathbb{N}$ and $k \in \mathbb{N}$ such that $2^{k-1} < p \leq 2^k$, DNNs approximating univariate Chebyšev polynomials up to degree 2^k are constructed, see also Remark A.1.

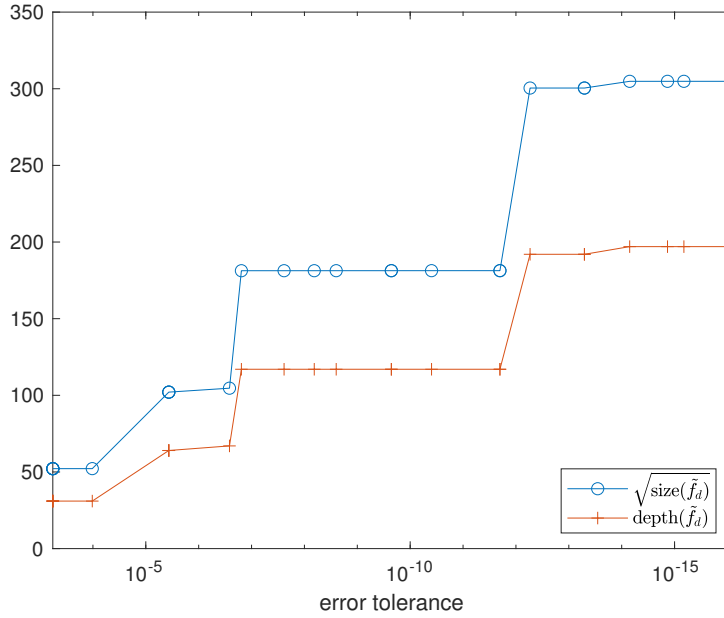


Figure 6.1: Accuracy vs. size and depth of the DNNs constructed by Algorithm 4.1. The data-to-QoI map $\delta \mapsto f(\delta)$ from (6.2) is approximated by constructed DNNs for $s = 1$ and $\sigma^2 = 0.04$.

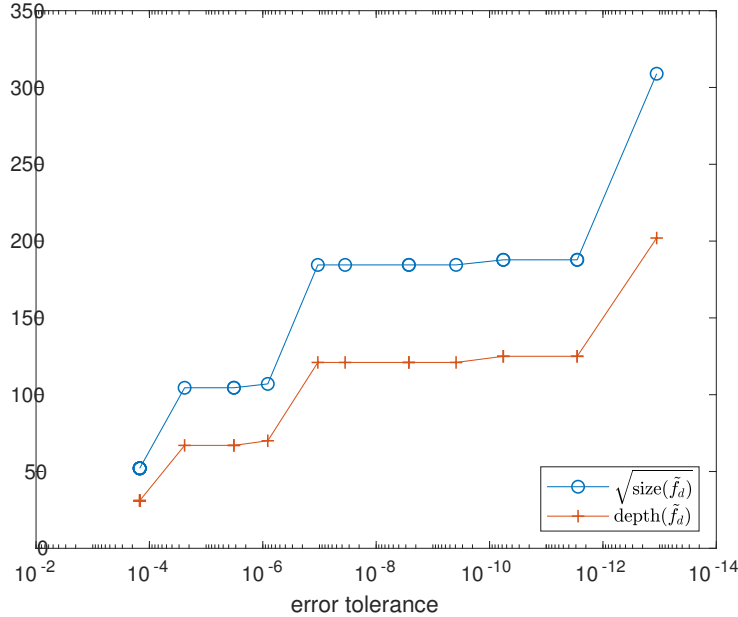


Figure 6.2: Accuracy vs. size and depth of the DNNs constructed by Algorithm 4.1. The data-to-QoI map $\delta \mapsto f(\delta)$ from (6.2) is approximated by constructed DNNs for $s = 2$ and $\sigma^2 = 0.04$.

7 Conclusions

We proved approximation results resp. expressive power bounds for DNNs of multivariate, analytic functions, and introduced a novel, deterministic, constructive approach for generating ReLU DNN surrogates. Unlike the (constructive) proofs of DNN emulation in [26], the presently proposed argument and algorithm depends in an essential way on the use of multivariate Chebyshev interpolants of the map f which are to be emulated by DNNs. On the one hand, while retaining exponential convergence for multivariate holomorphic functions, Chebyshev polynomials have favorable algorithmic and stability properties. In particular, through a close link with discrete Fourier transformation, Chebyshev interpolants of multivariate functions $f : [-1, 1]^K \rightarrow \mathbb{R}$ can be conveniently computed through function *collocation on $(n + 1)^K$ -point tensorized Clenshaw–Curtis grids*. The good (logarithmic w.r. to the polynomial degree n) stability of interpolation in Clenshaw–Curtis points also allowed us to perform an error analysis for interpolation based on “noisy”, *corrupted function evaluations*. We developed the corresponding error bounds for the DNN expression error. The error is bounded by $C \exp(-b(M^{1/(K+1)}))$ in terms of the network size M , for some $b > 0$ independent of K , if exact function evaluations of f in Clenshaw–Curtis points are used. The constructed DNN is stable in the sense that if instead the function values of f in Clenshaw–Curtis points are approximated with a possibly corrupted, numerical approximation \tilde{f} of f , there is an additional term in the L^∞ -error bounded by $C(\log(\log(1/\varepsilon)))^K \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$, where ε denotes the error of polynomial interpolation (for which it holds $n \propto \log(1/\varepsilon)$ in Proposition 3.5). The additional error in $W^{1,\infty}$ -norm is bounded by $C(\log(1/\varepsilon))^2 (\log(\log(1/\varepsilon)))^K \|f - \tilde{f}\|_{L^\infty([-1, 1]^K)}$.

In this paper, we applied the constructive algorithm to an example in the context of Bayesian PDE inversion, where under the assumption of additive Gaussian observation noise the data-to-quantity of interest map is analytic, cf. [15]. The algorithm is applicable much more generally, e.g. to any Lipschitz continuous function.

A Constructive ReLU DNN Approximation of T_n

We present an emulation of univariate Chebyšev polynomials $T_n(x)$ of arbitrary degrees by ReLU NNs, which will be developed in [24] and which follows closely the construction in [25] for univariate monomials. Specifically, we construct a DNN that approximates the Chebyšev polynomials of the first kind, denoted by $\{T_\ell\}_{\ell \in \mathbb{N}_0}$. As was derived for NNs with the RePU activation function $\sigma_r(x) = (\max\{x, 0\})^r$ for $r \in \mathbb{N}$ satisfying $r \geq 2$ in [32], they can be approximated efficiently also by ReLU NNs by exploiting the three term recursion

$$\forall m, n \in \mathbb{N}_0 : \quad T_{m+n} = 2T_m T_n - T_{|m-n|}, \quad T_0(x) = 1, \quad T_1(x) = x, \quad \forall x \in \mathbb{R}. \quad (\text{A.1})$$

This recurrence is specific to the T_n and follows from the addition rule for cosines.

To construct the DNNs that approximate all Chebyšev polynomials of degree $1, \dots, n$ on $\hat{I} := (-1, 1)$, we first construct inductively, for all $k \in \mathbb{N}$, NNs $\{\Psi_\delta^k\}_{\delta \in (0,1)}$ with input dimension one and output dimension $2^{k-1} + 2$ with the following properties: denoting all components of the output, except for the first one, by $\tilde{T}_{\ell,\delta} := \mathbb{R}(\Psi_\delta^k)_{2+\ell-2^{k-1}}$ for $\ell \in \{2^{k-1}, \dots, 2^k\}$, it holds that

$$\begin{aligned} \Psi_\delta^k(x) &= (x, \tilde{T}_{2^{k-1},\delta}(x), \dots, \tilde{T}_{2^k,\delta}(x)), \quad x \in \hat{I}, \\ \left\| T_\ell(x) - \tilde{T}_{\ell,\delta}(x) \right\|_{L^\infty(\hat{I})} &\leq \delta, \quad \ell \in \{2^{k-1}, \dots, 2^k\}. \end{aligned} \quad (\text{A.2})$$

We only provide the NN constructions, for proofs of the error bound and the estimates on the network depth and size in Lemma 3.1 we refer to [24].

Induction basis. Let $\delta \in (0, 1)$ be arbitrary and define $L_1 := \text{depth}(\tilde{\prod}_{\delta/4,1}^2)$. Also, define the matrix $A := [1, 1]^\top \in \mathbb{R}^{2 \times 1}$ and the vector $b' := [-1] \in \mathbb{R}^1$, and let $A_i, b_i, i = 1, \dots, L_1$ denote the weights and biases of $\tilde{\prod}_{\delta/4,1}^2$ as in Proposition 3.2. Then we define

$$\Psi_\delta^1 := (\Phi_{1,L_1}^{\text{Id}}, \Phi_{1,L_1}^{\text{Id}}, \Phi),$$

where the weights and biases of Φ are $A_1 A, A_2, \dots, A_{L_1-1}, 2A_{L_1}$ resp. $b_1, \dots, b_{L_1-1}, 2b_{L_1} + b'$. It follows that $(\Psi_\delta^1(x))_1 = x$, $\tilde{T}_{1,\delta}(x) := (\Psi_\delta^1(x))_2 = x = T_1(x)$ and $\tilde{T}_{2,\delta}(x) := (\Psi_\delta^1(x))_3 = 2(\tilde{\prod}_{\delta/4,1}^2)(x, x) - 1$ for all $x \in \hat{I} := (-1, 1)$.

Induction hypothesis (IH). For all $\delta \in (0, 1)$ and $k \in \mathbb{N}$, let $\theta := 2^{-2k-4}\delta$, and assume that there exists a NN Ψ_θ^k which satisfies Equation (A.2) with θ instead of δ .

Induction step. For δ and k as in (IH), we show that (A.2) holds with δ as in (IH) and with $k+1$ instead of k . We define, for $\Phi^{1,k}$ and $\Phi_\delta^{2,k}$ introduced below,

$$\Psi_\delta^{k+1} := \Phi_\delta^{2,k} \circ \Phi^{1,k} \circ \Psi_\theta^k. \quad (\text{A.3})$$

For a sketch of the network structure, see Figure A.1. The NN $\Phi^{1,k}$ of depth 0 implements the linear map

$$\mathbb{R}^{2^{k-1}+2} \rightarrow \mathbb{R}^{2^{k+1}+2} : (z_1, \dots, z_{2^{k-1}+2}) \mapsto (z_1, z_{2^{k-1}+2}, z_2, z_3, z_3, z_3, z_3, z_4, z_4, z_4, z_4, z_5, z_5, z_5, \dots, z_{2^{k-1}+1}, z_{2^{k-1}+2}, z_{2^{k-1}+2}, z_{2^{k-1}+2}).$$

Denoting its weights and biases by $A^{1,k}, b^{1,k}$, it holds that $b^{1,k} = 0$ and

$$(A^{1,k})_{m,i} = \begin{cases} 1 & \text{if } m = 1, i = 1, \\ 1 & \text{if } m = 2, i = 2^{k-1} + 2, \\ 1 & \text{if } m \in \{3, \dots, 2^{k+1} + 2\}, i = \lceil \frac{m+5}{4} \rceil, \\ 0 & \text{else.} \end{cases}$$

With $L_\theta := \text{depth}(\tilde{\prod}_{\theta,2}^2)$ we define

$$\Phi_\delta^{2,k} := \Phi \circ \left(\Phi_{2,L_\theta}^{\text{Id}}, \tilde{\prod}_{\theta,2}^2, \dots, \tilde{\prod}_{\theta,2}^2 \right)_d,$$

containing $2^k \tilde{\prod}_{\theta,2}^2$ -networks, with Φ denoting the depth 0 network with weights and biases $A^{2,k} \in \mathbb{R}^{(2^k+2) \times (2^k+2)}$ and $b^{2,k} \in \mathbb{R}^{2^k+2}$ defined as

$$(A^{2,k})_{m,i} := \begin{cases} 1 & \text{if } m = i \leq 2, \\ 2 & \text{if } m = i \geq 3, \\ -1 & \text{if } m \geq 3 \text{ is odd, } i = 1, \\ 0 & \text{else,} \end{cases} \quad (b^{2,k})_m = \begin{cases} -1 & \text{if } m \geq 3 \text{ is even,} \\ 0 & \text{else.} \end{cases}$$

The network Ψ_δ^{k+1} defined in Equation (A.3) realizes

$$(\Psi_\delta^{k+1}(x))_1 = x, \quad \text{for } x \in \hat{I}, \quad (\text{A.4})$$

$$(\Psi_\delta^{k+1}(x))_2 = \tilde{T}_{2^k, \theta}(x), \quad \text{for } x \in \hat{I}, \quad (\text{A.5})$$

$$(\Psi_\delta^{k+1}(x))_{\ell+2-2^k} = 2 \tilde{\prod}_{\theta,2}^2 \left(\tilde{T}_{\lceil \ell/2 \rceil, \theta}(x), \tilde{T}_{\lfloor \ell/2 \rfloor, \theta}(x) \right) - x^{\lceil \ell/2 \rceil - \lfloor \ell/2 \rfloor}, \quad \text{for } x \in \hat{I} \text{ and} \quad (\text{A.6})$$

$$\ell \in \{2^k + 1, \dots, 2^{k+1}\},$$

where $x^{\lceil \ell/2 \rceil - \lfloor \ell/2 \rfloor} = x = T_1(x)$ for odd ℓ and $x^{\lceil \ell/2 \rceil - \lfloor \ell/2 \rfloor} = 1 = T_0(x)$ for even ℓ . For $\ell \in \{2^k + 1, \dots, 2^{k+1}\}$ and $x \in \hat{I}$, the right-hand side of (A.6) will be denoted by

$$\tilde{T}_{\ell, \delta}(x) := (\Psi_\delta^{k+1}(x))_{\ell+2-2^k}.$$

This finishes the construction of Ψ_δ^{k+1} .

Next, we construct $\Phi_\delta^{\text{Cheb}, n}$ as in Lemma 3.1.

If $n = 1$, for all $\delta \in (0, 1)$ we define $\Phi_\delta^{\text{Cheb}, n} := ((A, b))$, where $A := [1] \in \mathbb{R}^{1 \times 1}$ and $b := [0] \in \mathbb{R}^1$.

If $n \geq 2$, let $k := \lceil \log_2(n) \rceil$. See Figure A.2 for a sketch of the NN construction below. We use the networks $\{\Psi_\delta^j\}_{\delta \in (0,1), j \in \{1, \dots, k\}}$ constructed above and take $\{\ell_j\}_{j=1}^k \subset \mathbb{N}$ such that $\text{depth}(\Psi_\delta^k) + 1 = \text{depth}(\Psi_\delta^j) + \ell_j$ for $j = 1, \dots, k$, and thus $\ell_j \leq \max_{j=1}^k \text{depth}(\Psi_\delta^j) = \text{depth}(\Psi_\delta^k)$. We define

$$\Phi_\delta^{\text{Cheb}, n} := \Phi^{3,n} \circ (\Psi_\delta^1 \circ \Phi_{1, \ell_1}^{\text{Id}}, \dots, \Psi_\delta^k \circ \Phi_{1, \ell_k}^{\text{Id}}),$$

where the NN $\Phi^{3,n}$ of depth 0 emulates the linear map $\mathbb{R}^{2^k+2k-1} \rightarrow \mathbb{R}^n$ satisfying

$$\Phi^{3,n}(z_1, \dots, z_{2^k+2k-1})_1 = z_2, \quad \Phi^{3,n}(z_1, \dots, z_{2^k+2k-1})_2 = z_3,$$

$$\text{and for all } \ell = 3, \dots, n, \text{ with } j := \lceil \log_2(\ell) \rceil : \quad \Phi^{3,n}(z_1, \dots, z_{2^k+2k-1})_\ell = z_{\ell+2j-1}.$$

The realization satisfies for all $\ell = 1, \dots, n$

$$(\Phi_\delta^{\text{Cheb}, n}(x))_\ell = \tilde{T}_{\ell, \delta}(x), \quad x \in \hat{I}, \ell \in \{1, \dots, n\}.$$

Remark A.1 *The subnetwork Ψ_δ^k of $\Phi_\delta^{\text{Cheb}, n}$ approximates all univariate Chebyšev polynomials of degree up to 2^k , also when $n < 2^k$. This causes the “step-like” behavior in Figures 6.1 and 6.2. This step-wise growth of the network size can easily be prevented by removing from Ψ_δ^k the product networks $\tilde{\prod}_{\theta,2}^2$ that compute $\tilde{T}_{\ell, \delta}(x)$ for $\ell > n$, and modifying $\Phi^{3,n}$ accordingly.*

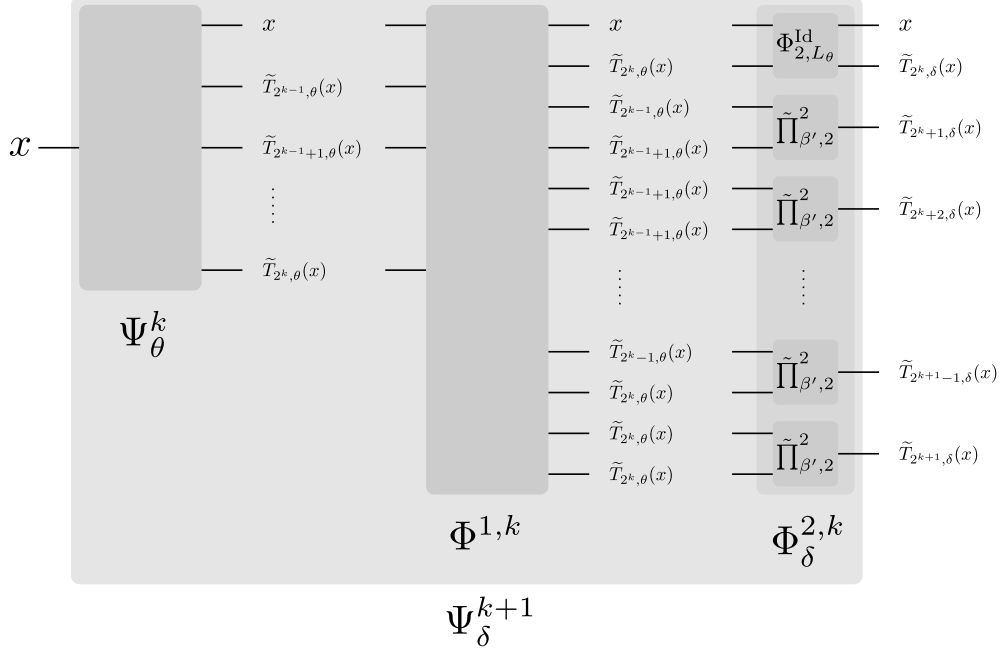


Figure A.1: Sketch of Ψ_δ^{k+1} for some $k \in \mathbb{N}$ and $\delta \in (0, 1)$, inductively constructed from Ψ_θ^k with $\theta = 2^{-2k-4}\delta$. The subnetwork $\Phi^{1,k}$ realizes a linear map, correctly coupling the output of Ψ_θ^k to the input of $\Phi_\delta^{2,k}$. The subnetwork $\Phi_\delta^{2,k}$ acts as the identity on the first two inputs, and as an approximate multiplication from Proposition 3.2 on pairs of the remaining inputs. Output are the input x and approximations of Chebyshev polynomials of degree $2^k, \dots, 2^{k+1}$, with accuracy δ .

References

- [1] B. Adcock and N. Dexter. The gap between theory and practice in function approximation with deep neural networks, 2020. ArXiv: 2001.07523.
- [2] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb. Solving inverse problems using data-driven models. *Acta Numer.*, 28:1–174, 2019.
- [3] H. Bölcskei, P. Grohs, G. Kutyniok, and P. Petersen. Optimal approximation with sparsely connected deep neural networks. *SIAM J. Math. Data Sci.*, 1(1):8–45, 2019.
- [4] N. Boull, Y. Nakatsukasa, and A. Townsend. Rational neural networks, 2020. Accepted for publication in *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, Vancouver, Canada.
- [5] P. Cheridito, A. Jentzen, and F. Rossmannek. Non-convergence of stochastic gradient descent in the training of deep neural networks, 2020. ArXiv: 2006.07075.
- [6] A. Cohen, C. Schwab, and J. Zech. Shape holomorphy of the stationary Navier-Stokes Equations. *SIAM J. Math. Analysis*, 50(2):1720–1752, 2018.
- [7] M. Dashti and A. M. Stuart. The Bayesian approach to inverse problems. In *Handbook of uncertainty quantification. Vol. 1, 2, 3*, pages 311–428. Springer, Cham, 2017.
- [8] J. Daws and C. Webster. Analysis of deep neural networks with quasi-optimal polynomial approximation rates, 2019. ArXiv: 1912.02302.
- [9] W. E and Q. Wang. Exponential convergence of the deep neural network approximation for analytic functions. *Sci. China Math.*, 61(10):1733–1740, 2018.
- [10] H. Ehlich and K. Zeller. Auswertung der Normen von Interpolationsoperatoren. *Math. Ann.*, 164:105–112, 1966.

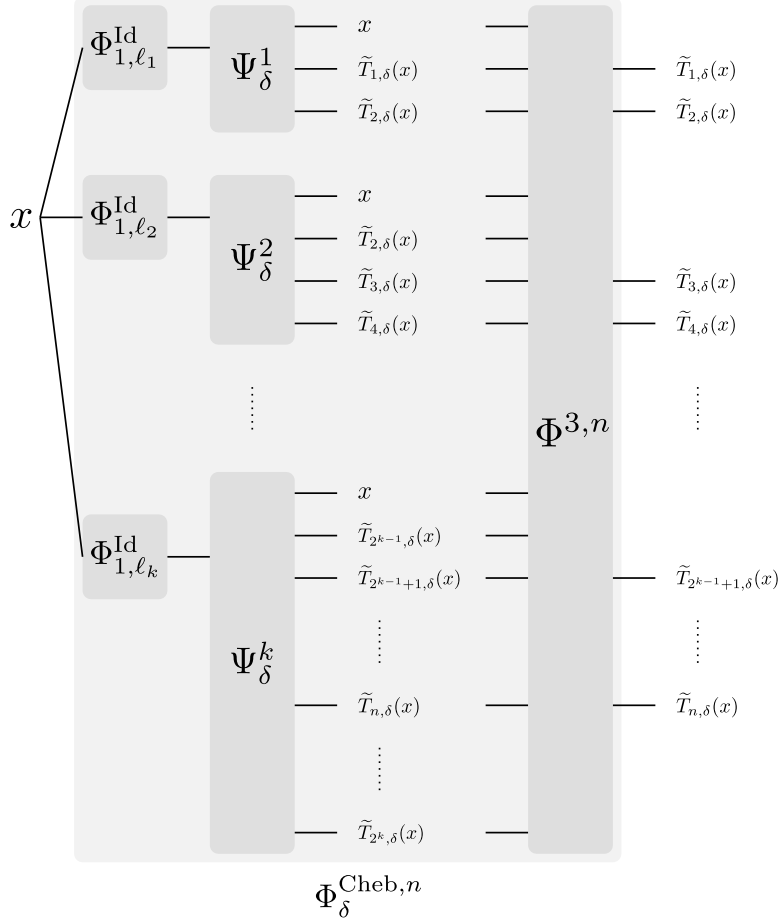


Figure A.2: Sketch of $\Phi_\delta^{\text{Cheb},n}$ for some $n \in \mathbb{N}$ and $\delta \in (0, 1)$, constructed from identity networks and the previously described $\Psi_\delta^1, \dots, \Psi_\delta^k$, with $k := \lceil \log_2(n) \rceil$. The subnetwork $\Phi^{3,n}$ realizes a linear map that selects the desired approximations of univariate Chebyshev polynomials from the outputs of the preceding layer.

- [11] D. Elbrächter, P. Grohs, A. Jentzen, and C. Schwab. DNN Expression Rate Analysis of High-dimensional PDEs: Application to Option Pricing, 2020. ArXiv:1809.07669, accepted for publication in *Constructive Approximation*.
- [12] D. Fokina and I. Oseledets. Growing axons: greedy learning of neural networks with application to function approximation, 2019. ArXiv: 1910.12686.
- [13] M. Gaß, K. Glau, M. Mahlstedt, and M. Mair. Chebyshev interpolation for parametric option pricing. *Finance and Stochastics*, 22(3):701–731, 2018.
- [14] F. Henriquez and C. Schwab. Shape Holomorphy of the Calderon Projector for the Laplacean in \mathbb{R}^2 . Technical Report 2019-43, Seminar for Applied Mathematics, ETH Zürich, 2019. (to appear in *Journ. Int. Eqns. Op. Theory* (2021)).
- [15] L. Herrmann, C. Schwab, and J. Zech. Deep neural network expression of posterior expectations in Bayesian PDE inversion. *Inverse Problems*, 36(12):125011, 2020.
- [16] B. Hosseini and N. Nigam. Well-posed Bayesian inverse problems: priors with exponential tails. *SIAM/ASA J. Uncertain. Quantif.*, 5(1):436–465, 2017.
- [17] C. Jerez-Hanckes, C. Schwab, and J. Zech. Electromagnetic Wave Scattering by Random Surfaces: Shape Holomorphy. *Math. Mod. Meth. Appl. Sci.*, 27(12):2229–2259, 2017.

- [18] B. Li, S. Tang, and H. Yu. Better approximations of high dimensional smooth functions by deep neural networks with rectified power units. *Communications in Computational Physics*, 27(2):379–411, 2019.
- [19] S. Liang and R. Srikant. Why deep neural networks for function approximation? In *Proc. of ICLR 2017*, pages 1 – 17, 2017. ArXiv:1610.04161.
- [20] L. Lu, P. Jin, and G. E. Karniadakis. DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators, 2020.
- [21] K. O. Lye, S. Mishra, and D. Ray. Deep learning observables in computational fluid dynamics. *J. Comput. Phys.*, 410:109339, 26, 2020.
- [22] H. N. Mhaskar. Approximation properties of a multilayered feedforward artificial neural network. *Advances in Computational Mathematics*, 1(1):61–80, Feb 1993.
- [23] H. N. Mhaskar. Neural networks for optimal approximation of smooth and analytic functions. *Neural Computation*, 8(1):164–177, 1996.
- [24] J. A. A. Opschoor. *In preparation*. PhD thesis, Dissertation, ETH Zürich, 202x.
- [25] J. A. A. Opschoor, P. C. Petersen, and C. Schwab. Deep ReLU networks and high-order finite element methods. *Analysis and Applications*, 18(05):715–770, 2020.
- [26] J. A. A. Opschoor, C. Schwab, and J. Zech. Exponential ReLU DNN expression of holomorphic maps in high dimension. Technical Report 2019-35_rev1, Seminar for Applied Mathematics, ETH Zürich, 2020. Accepted for publication in *Constructive Approximation*.
- [27] P. Petersen and F. Voigtlaender. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *Neural Netw.*, 108:296 – 330, 2018.
- [28] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, 378:686–707, 2019.
- [29] T. J. Rivlin. *The Chebyshev polynomials*. Wiley-Interscience [John Wiley & Sons], New York-London-Sydney, 1974. Pure and Applied Mathematics.
- [30] D. Rolnick and M. Tegmark. The power of deeper networks for expressing natural functions. In *International Conference on Learning Representations*, 2018.
- [31] C. Schwab and J. Zech. Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in UQ. *Analysis and Applications, Singapore*, 17(1):19–55, 2019.
- [32] S. Tang, B. Li, and H. Yu. ChebNet: Efficient and stable constructions of deep neural networks with rectified power units using Chebyshev approximations. Technical report, 2019. ArXiv: 1911.05467.
- [33] L. N. Trefethen. *Approximation theory and approximation practice*. Society for Industrial and Applied Mathematics, Philadelphia, extended edition, 2019.
- [34] L. Yang, X. Meng, and G. E. Karniadakis. B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *Journal of Computational Physics*, 425:109913, Jan 2021.
- [35] D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Netw.*, 94:103–114, 2017.