

# Chapter 1

## Learning smooth functions in high dimensions

From sparse polynomials to deep neural networks

Ben Adcock<sup>a,\*</sup>, Simone Brugiapaglia<sup>b</sup>, Nick Dexter<sup>c</sup>, and Sebastian Moraga<sup>a</sup>

<sup>a</sup>Simon Fraser University, Burnaby, BC, Canada, <sup>b</sup>Concordia University, Montreal, QC, Canada,

<sup>c</sup>Florida State University, Tallahassee, FL, United States

\*Corresponding author: e-mail address: [ben\\_adcock@sfu.ca](mailto:ben_adcock@sfu.ca)

### Contents

<b>1 Introduction</b>	<b>2</b>	5.2 Lower bounds for adaptive $m$ -widths	23
1.1 Motivations and challenges	3	5.3 Towards methods	24
1.2 Overview	3	<b>6 Learning sparse polynomial approximations from data</b>	<b>24</b>
1.3 Further literature	9	6.1 Setup	25
<b>2 Problem statement and notation</b>	<b>11</b>	6.2 Sampling discretizations for multivariate polynomials	26
<b>3 Holomorphic functions of infinitely many variables</b>	<b>12</b>	6.3 Resolving Challenge 1: lower and anchored sets	28
3.1 $(b, \varepsilon)$ -holomorphic functions	12	6.4 Resolving Challenge 2: weighted $k$ -term approximation	29
3.2 Holomorphy and parametric DEs	14	6.5 Weighted $\ell^1$ -minimization	31
3.3 Known and unknown anisotropy	16	6.6 Theoretical guarantee	32
3.4 $\ell^p$ -summability and the $\mathcal{H}(p)$ and $\mathcal{H}(p, M)$ classes	17	6.7 Discussion and extensions	34
<b>4 Best <math>s</math>-term polynomial approximation</b>	<b>17</b>	<b>7 DNN existence theory</b>	<b>34</b>
4.1 Orthogonal polynomials	17	7.1 Review	34
4.2 Orthogonal polynomial expansions	18	7.2 Neural network architectures	36
4.3 Best $s$ -term polynomial approximation	19	7.3 Emulating polynomials with DNNs: typical result	36
4.4 Rates of best $s$ -term polynomial approximation	19	7.4 Elements of the proof of Theorem 7.1	37
4.5 How high is high dimensional?	20	7.5 Existence theorem for $(b, \varepsilon)$ -holomorphic functions	38
<b>5 Limits of learnability from data</b>	<b>22</b>	<b>8 Practical existence theory: near-optimal DL</b>	<b>38</b>
5.1 Adaptive $m$ -widths	22	8.1 Setup	39

## 2 Numerical Analysis Meets Machine Learning

8.2	Practical existence theorem	39	9.3	Theorem 8.1 does not eliminate the theory-practice gap	43
8.3	The mechanism of practical existence theorems	40	9.4	Practical insights	43
9	<b>Epilogue</b>	<b>42</b>	9.5	Eliminating the gap: beating the Monte Carlo rate is key	44
9.1	Scientific computing and data scarcity	42	9.6	Conclusion	45
9.2	Potential benefits to DNNs over sparse polynomials	43	<b>References</b>		<b>45</b>

---

### Abstract

Learning approximations to smooth target functions of many variables from finite collections of pointwise samples is an important task in scientific computing and its many applications in computational science and engineering. Despite well over half a century of research on high-dimensional approximation, this remains a challenging problem. Yet, significant advances have been made in the last decade towards efficient methods for doing this, commencing with so-called *sparse polynomial approximation* methods and continuing most recently with methods based on *Deep Neural Networks (DNNs)* and *Deep Learning (DL)*. In tandem, there have been substantial advances in the relevant approximation theory and analysis of these techniques. In this work, we survey this recent progress. We describe the contemporary motivations for this problem, which stem from parametric models and computational uncertainty quantification; the relevant function classes, namely, classes of infinite-dimensional, Banach-valued, holomorphic functions; fundamental limits of learnability from finite data for these classes; and finally, sparse polynomial and DNN methods for efficiently learning such functions from finite data. In the case of the latter, there is currently a significant gap between the approximation theory of DNNs and the practical performance of DL. Aiming to narrow this gap, we develop the topic of *practical existence theory*, which asserts the existence of dimension-independent DNN architectures and DL training strategies that achieve provably near-optimal generalization errors in terms of the amount of training data.

### Keywords

High-dimensional approximation, Holomorphy, Scarce data, Polynomial approximation, Deep neural networks, Deep learning, Practical existence theory

### MSC Codes

65D40, 41A10, 41A63, 41A25, 68T07

---

## 1 Introduction

This article reviews the problem of learning infinite-dimensional functions from finite data. In this section, we begin by describing motivations for this problem and its challenges (§1.1). We then present an overview of the remainder of the article (§1.2). Finally, we discuss relevant literature (§1.3).

## 1.1 Motivations and challenges

Approximating functions of many variables is a classical topic in approximation theory and numerical analysis, which has been studied intensively for more than seventy years. The contemporary resurgence in this problem stems from applications to parametric models, parametric *Differential Equations (DEs)* and computational *Uncertainty Quantification (UQ)*. In such models, the target function represents some quantity-of-interest of a given physical system and its variables the parameters of the system. Complex physical models involve many parameters, which naturally leads to functions of many variables. It is also increasingly common to consider countably infinite parameterizations, a standard example being the use of a Karhunen–Loève expansion to represent a random field in, for example, parametric models of porous media (see, e.g., Le Maître and Knio, 2010, §2.1, or Sullivan, 2015, §11.1). In this case, the target function depends on countably many variables.

Therefore, the first challenge in approximating such functions is that the dimension is high or countably infinite. The second is that the function is usually expensive to evaluate. Typically, each evaluation involves either an expensive numerical simulation of the DE model or a costly physical experiment. Hence, one strives to learn an accurate approximation from a limited set of samples. This approximation – or *surrogate model* as it is often termed (see Smith, 2013, Chpt. 13, or Sullivan, 2015, Chpt. 13) – can then be used in place of the true function for various downstream tasks, such as parameter optimization, inverse parametric problems or uncertainty quantification. Nonetheless, since the target function is expensive to evaluate, the amount of data available to learn this approximation is usually highly limited. Put another way, the learning problem is highly *data-starved*.

A final challenge is that the target function is often not scalar-valued. In the case of parametric DEs, the output of the target function is the solution of the DE with the given input parameters. Hence this function takes values in an infinite-dimensional Banach or Hilbert space. This adds complications, both theoretically and practically, stemming from the need to discretize this output space in order to perform computations.

## 1.2 Overview

We now give a detailed overview of this article.

### *Problem statement and notation (§2)*

We consider an unknown target function  $f$  depending on countably many variables that range between finite maximum and minimum values and take values in a Banach space  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$ . After normalizing, we may assume that

$$f : \mathcal{U} \rightarrow \mathcal{V}, \quad \text{where } \mathcal{U} = [-1, 1]^{\mathbb{N}}.$$

We equip  $\mathcal{U}$  with the uniform probability measure  $\varrho$  and consider training data

$$\{(\mathbf{y}_i, f(\mathbf{y}_i) + e_i)\}_{i=1}^m \subseteq \mathcal{U} \times \mathcal{V}, \quad (1.1)$$

where  $\mathbf{y}_1, \dots, \mathbf{y}_m \sim_{\text{i.i.d.}} \varrho$  and  $e_i \in \mathcal{V}$  represents measurement error. The objective is to learn an approximation  $\hat{f}$  to  $f$  from the data (2.2).

### *Holomorphic functions of infinitely many variables (§3)*

An important property of many parametric model problems is that the parametric map  $f$  is a smooth function of its variables. There is a large body of literature, which we review further in §3, that establishes that solution maps for various parametric DEs are *holomorphic* functions – in other words, they admit holomorphic extensions to certain complex regions  $\mathcal{U} \subset \mathcal{O} \subseteq \mathbb{C}^{\mathbb{N}}$ .

In one dimension, it is common to consider complex regions  $[-1, 1] \subset \mathcal{O} \subset \mathbb{C}$  defined by *Bernstein ellipses* (see, e.g., Trefethen, 2013). The convergence of the  $s$ -term polynomial expansion is then *exponential* in  $s$ , with the rate being characterized by the largest Bernstein ellipse to which the function can be extended (see, e.g., Trefethen, 2013, Chpt. 8). For functions of  $d \geq 1$  variables, it is natural to look at regions  $[-1, 1]^d \subset \mathcal{O} \subset \mathbb{C}^d$  defined by *Bernstein polyellipses*, i.e., tensor-products of one-dimensional Bernstein ellipses. One can then show exponential convergence in  $s^{1/d}$  of a suitable  $s$ -term polynomial expansion, with rate depending once more on the largest polyellipse to which  $f$  admits an extension (see, e.g., Adcock et al., 2022, §3.5-3.6).

*( $\mathbf{b}, \varepsilon$ )-holomorphic functions.* The situation changes in infinite dimensions. Holomorphy in an arbitrary Bernstein polyellipse is no longer sufficient to guarantee convergence of a (in fact, any)  $s$ -term polynomial expansion. At the very least, one needs *anisotropy*, namely, increasing regularity with the variables  $y_i$  as  $i \rightarrow \infty$ . Moreover, for parametric DE applications, one needs to consider not just a single Bernstein ellipse, but complex regions  $\mathcal{R}(\mathbf{b}, \varepsilon) \subset \mathbb{C}^{\mathbb{N}}$  defined by certain *unions* of Bernstein polyellipses. Here  $\mathbf{b} = (b_j)_{j \in \mathbb{N}} \in [0, \infty)^{\mathbb{N}}$  with  $\mathbf{b} \in \ell^1(\mathbb{N})$  is a sequence and  $\varepsilon > 0$  is a scalar that parametrize the corresponding region. In Definition 3.3 we formalize the corresponding class of so-called *( $\mathbf{b}, \varepsilon$ )-holomorphic functions* (see, e.g., Chkifa et al., 2015b; Schwab and Zech, 2019). This class was first introduced in the context of parametric DEs (Cohen and DeVore, 2015), and has since become a standard setting in which to consider the approximation of holomorphic functions in infinite dimensions (Adcock et al., 2022, Chpt. 3). Throughout the majority of this work, we set  $\varepsilon = 1$ ; a property which can always be guaranteed by rescaling  $\mathbf{b}$ . For convenience, we also define the set

$$\mathcal{H}(\mathbf{b}) = \left\{ f : \mathcal{U} \rightarrow \mathcal{V} \text{ ( $\mathbf{b}, 1$ )-holomorphic} : \sup_{z \in \mathcal{R}(\mathbf{b}, 1)} \|f(z)\|_{\mathcal{V}} \leq 1 \right\}. \quad (1.2)$$

This is the main class of functions considered in the remainder of this work. In §3 we explain the relevance of this class to parametric DEs in more detail.

*Unknown anisotropy.* Having done so, we then explain how the sequence  $\mathbf{b}$  controls the *anisotropy* of  $f \in \mathcal{H}(\mathbf{b})$ . Larger  $b_i$  implies that  $f$  is less smooth with respect to its  $i$ th variable  $y_i$ , while smaller  $b_i$  implies more smoothness. In some specific situations, one may have knowledge of a suitable  $\mathbf{b}$  – information which could then be used when designing a learning algorithm. However, in practice, it is more often the case that  $\mathbf{b}$  is unknown. The focus of this work is the more realistic *unknown anisotropy* setting. In short, the learning method must be independent of  $\mathbf{b}$ , with the assumption  $f \in \mathcal{H}(\mathbf{b})$  being used only to derive bounds for the resulting generalization error.

### Orthogonal polynomials and best $s$ -term polynomial approximation (§4)

As previously observed, holomorphy is intimately related to polynomial approximation. In §4 we introduce multivariate orthogonal polynomials, orthogonal polynomial expansions of infinite-dimensional, Banach-valued functions and the concept of *best  $s$ -term polynomial approximation*. As we argue, best  $s$ -term polynomial approximation is an important theoretical benchmark against which to compare methods for learning holomorphic functions from data.

A signature result we recap in this section is the following.

**Informal Result 1** (Algebraic convergence of the best  $s$ -term approximation, Theorem 4.1). *The best  $s$ -term polynomial approximation of  $f \in \mathcal{H}(\mathbf{b})$  converges algebraically fast in  $s$ . Specifically, if  $\mathbf{b} \in \ell^p(\mathbb{N})$  for some  $0 < p < 1$ , then the  $L^2_{\mathcal{Q}}$ -norm error decreases like  $O(s^{1/2-1/p})$  as  $s \rightarrow \infty$ .*

This result implies that best  $s$ -term polynomial approximation in  $\mathcal{H}(\mathbf{b})$  is free from the *curse of dimensionality*, whereas the aforementioned exponential rates witnessed in  $d < \infty$  dimensions (which are exponential in  $s^{1/d}$ ) quickly succumb to this curse. Although the focus in this work is infinite-dimensional functions, in §4 we demonstrate that such algebraic rates also typically describe the true convergence behavior seen for functions of finitely many variables, whenever  $d$  is large enough (e.g.,  $d \geq 16$ ).

### Limits of learnability from data (§5)

We next turn our attention to learning holomorphic functions from the data (1.1). In §5, we study lower bounds using concepts from *information-based complexity*, in particular, *adaptive  $m$ -widths* (Definition 5.1). In order to consider the unknown anisotropy setting, we introduce the function classes

$$\mathcal{H}(p) = \bigcup \left\{ \mathcal{H}(\mathbf{b}) : \mathbf{b} \in \ell^p(\mathbb{N}), \mathbf{b} \in [0, \infty)^{\mathbb{N}}, \|\mathbf{b}\|_p \leq 1 \right\}, \quad 0 < p < 1,$$

where  $\mathcal{H}(\mathbf{b})$  is as in (1.2), and

$$\mathcal{H}(p, M) = \bigcup \left\{ \mathcal{H}(\mathbf{b}) : \mathbf{b} \in \ell^p_M(\mathbb{N}), \mathbf{b} \in [0, \infty)^{\mathbb{N}}, \|\mathbf{b}\|_{p, M} \leq 1 \right\}, \quad 0 < p < 1.$$

Here  $\ell_M^p(\mathbb{N})$  is the *monotone*  $\ell^p$ -space (see §2). We then have the following.

**Informal Result 2** (Limits of learnability, Theorem 5.2). *It is impossible to learn functions from  $\mathcal{H}(p)$ . Specifically, there does not exist a method for learning functions from  $m$  (adaptive) linear samples for which the  $L_Q^2$ -norm error decreases as  $m \rightarrow \infty$  uniformly for functions in  $\mathcal{H}(p)$ . Moreover, when restricting to the space  $\mathcal{H}(p, \mathbf{M})$ , or even  $\mathcal{H}(\mathbf{b})$ , the error cannot decrease faster than  $c \cdot m^{1/2-1/p}$  for some universal constant  $c > 0$ , even if the method is allowed to depend on the anisotropy parameter  $\mathbf{b} \in \ell_M^p(\mathbb{N})$ .*

This theorem illustrates a fundamental gap between approximation theory and learning from data. The best  $s$ -term approximation converges like  $s^{1/2-1/p}$  for any  $f \in \mathcal{H}(p)$ , yet no method can learn such functions from data. Note that this result holds not just for i.i.d. pointwise samples, as in (1.1), but *any* linear measurements, which may also be *adaptive*.

This gap is narrowed by restricting to the class  $\mathcal{H}(p, \mathbf{M})$ , wherein a function’s variables are “on average” ordered in terms of importance. Notably, this lower bound also implies that there is no benefit to knowing  $\mathbf{b}$ , at least in terms of sample complexity. The same lower bound  $c \cdot m^{1/2-1/p}$  holds, even if the method is allowed to depend on  $\mathbf{b}$ .

With this in mind, the remainder of this work is devoted to describing learning methods that achieve (close to) the rate  $m^{1/2-1/p}$  uniformly for functions in  $\mathcal{H}(p, \mathbf{M})$  for *any*  $0 < p < 1$ . Specifically, we focus on sparse polynomial and *Deep Neural Network (DNN)* methods.

### *Learning sparse polynomial approximations from data (§6)*

Inspired by the benchmark results on best  $s$ -term approximation, we first consider methods that learn polynomial approximations. This approach is heavily based on techniques from *compressed sensing* (Foucart and Rauhut, 2013; Vidyasagar, 2019). Superficially, this seems straightforward. Fast decay of the best  $s$ -term polynomial approximation means that the polynomial coefficients are approximately sparse. Thus, a natural first idea would be to formulate an  $\ell^1$ -minimization problem for the coefficients and leverage standard compressed sensing theory to derive recovery guarantees.

Unfortunately, this approach fails to deliver optimal rates. The specific reason is that there are  $s$ -term polynomials that require approximately  $m = O(s^2)$  i.i.d. samples to be stably recovered. The higher level reason is that the assumed *low-complexity* model (functions in  $\mathcal{H}(p)$  have approximately sparse coefficients) is simply too crude. To lower the sample complexity requirement, we refine the model. We do this via *weighted sparsity* (Rauhut and Ward, 2016), which encodes the additional information that coefficients are both approximately sparse and *decaying*.

**Informal Result 3** (Near-optimal learning via polynomials, Theorem 6.8). *There is a method (based on weighted  $\ell^1$ -minimization) for learning functions*

from the data (1.1) that achieves an  $L^2_{\mathcal{Q}}$ -norm error that decays like

$$O((m/\log^4(m))^{1/2-1/p}), \quad m \rightarrow \infty, \quad (1.3)$$

with high probability, for functions in  $\mathcal{H}(p, M)$  and any  $0 < p < 1$ .

Comparing with the lower bound of  $m^{1/2-1/p}$ , we see that this procedure is optimal up to the logarithmic term. Note that in the full theorem, we also account for other errors in learning process: namely, *measurement error* (i.e., the effect of the terms  $e_i$  in (1.1)), *physical discretization error* (i.e., the effect of discretizing the space  $\mathcal{V}$  in order to perform computations) and *optimization error* (i.e., inexact solution of the optimization problem).

### DNN existence theory (§7)

As we discuss further in §1.3, the last five years have seen a growing interest in the application of *Deep Learning (DL)* to challenging parametric model problems. DL has the potential to achieve significant performance gains. Yet, it is poorly understood from a theoretical perspective, especially in terms of the sample complexity, i.e., the amount of training data need to learn good DNN approximations for specific classes of functions.

We commence §7 by briefly reviewing the approximation theory of DNNs, which has been an area of significant research in the last few years. Broadly, this area aims to establish *existence theorems*. Building on the classical *universal approximation theory* of shallow *Neural Networks (NNs)*, these results assert the existence of DNNs of a given complexity – in terms of the width and depth, number of nonzero weights and biases, or other pertinent metrics – that approximate functions in a given class to within a prescribed accuracy.

Such results are typically proved through *emulation*. One first shows that DNNs can emulate a standard approximation scheme (e.g., polynomials, piecewise polynomials, splines, wavelets, and so forth), either exactly or up to some error that can be made arbitrarily small. Then one leverages known bounds for the standard scheme to derive the corresponding existence theorem.

Having briefly reviewed this literature, we then describe an exemplar existence theorem for holomorphic functions, which is obtained by emulating a certain  $s$ -term polynomial approximation.

**Informal Result 4** (Existence theorem for DNNs, Theorem 7.2). *Let  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}}) = (\mathbb{R}, |\cdot|)$ . Then for any  $\mathbf{b} \in [0, \infty)^{\mathbb{N}}$  with  $\mathbf{b} \in \ell^p(\mathbb{N})$  for some  $0 < p < 1$  and  $s \in \mathbb{N}$ , there is a family  $\mathcal{N}$  of DNNs of width  $O(s^2)$  and depth  $O(\log(s))$  with the following property. For any  $f \in \mathcal{H}(\mathbf{b})$  there is an element of  $\mathcal{N}$  that achieves an  $L^2_{\mathcal{Q}}$ -norm error of  $O(s^{1/2-1/p})$ .*

### Practical existence theory: near-optimal DL (§8)

While important, existence theorems such as this say little about the performance of DL, i.e., computing a DNN from training data by minimizing a loss

function. As noted, analysis of this practical scenario is currently lacking, certainly in terms of sample complexity. It has also been well documented that there is a gap between existence theory and the practical performance of DNNs when trained from finite numbers of samples (Abdeljawad and Grohs, 2023; Adcock and Dexter, 2021; Grohs and Voigtlaender, 2023).

In this section, we describe a theoretical framework, termed *practical existence theory* (Adcock et al., 2021, 2023a; Franco and Brugiapaglia, 2024), that seeks to narrow this gap. Its aim is to show that there not only exists a family of DNNs, but also a training strategy akin to standard practice – i.e., minimizing a loss function – that attains the near-optimal rates.

**Informal Result 5** (Practical existence theorem for DNNs, Theorem 8.1). *There is a class of DNNs  $\mathcal{N}$  (of width and depth depending on  $m$ ), a regularized  $\ell^2$ -loss function and a choice regularization parameter (depending on  $m$  only) such that any DNN obtained by solving the resulting training problem (minimizing the loss function over  $\mathcal{N}$ ) achieves the error decay rate (1.3).*

Like existence theorems, this result is proved by emulating the sparse polynomial method with DNNs. The class  $\mathcal{N}$  is ‘handcrafted’ in the sense that the weights and biases in the hidden layers are fixed, and chosen specifically to emulate certain polynomials. Only the weights in the final layer are trained.

### *Epilogue: the benefits of practical existence theory and the gap between theory and practice (§9)*

Practical existence theory does not explain the performance of standard DL strategies based on *fully trained models* (i.e., those where all layers of a DNN are trained). Nonetheless, it provides a number of key insights and benefits, which we now list (see §9 for details).

1. It narrows the gap between theory and practice, by showing the existence of good training procedures, which are near-optimal for certain function classes.
2. It thereby emphasizes the potential to achieve even better performance in practice with suitably chosen architectures and training strategies.
3. It also stresses the sample complexity aspect of DL for high-dimensional approximation. As argued, this is vital to parametric modelling applications, where data scarcity is of primary concern.
4. The architectures in Theorem 8.1 are much wider than they are deep. Smoother activation functions also lead to narrower and shallow DNNs. Both insights agree with the empirical performance of fully trained models.
5. Once trained, the DNN in Theorem 8.1 can be sparsified via *pruning*. This lends credence to the idea that sparse DNNs can perform well in practice.
6. The methodology of practical existence theory is flexible, and can be applied to other problems, such as reduced order models based on deep autoencoders (Franco and Brugiapaglia, 2024) and Physics-Informed Neural Networks (PINNs) for PDEs (Brugiapaglia et al., 2024).



While narrowed, the gap between theory and practice still persists. With an eye towards future research, we end §9 with some further discussion. We discuss how current techniques for analyzing fully trained models may at best achieve a rate of  $O(m^{-1/2})$ , regardless of  $p$ , which is strictly slower than those asserted by practical existence theorems. However, the far greater expressivity of DNNs – in particular, their ability to approximate continuous and discontinuous functions alike – sets them apart from more standard scientific computing tools such as polynomials. While the latter are near-optimal for holomorphic functions, the flexibility of former is a significant potential advantage.

### 1.3 Further literature

See Ghanem et al. (2017); Le Maître and Knio (2010); Smith (2013); Sullivan (2015) for general introductions to computational UQ and Adcock et al. (2022); Cohen and DeVore (2015) for more on parametric models and parametric DEs. Note that in the context of UQ, orthogonal polynomial expansions are often termed polynomial *chaos* expansions (see, e.g., Ghanem and Spanos, 2003). Best  $s$ -term polynomial approximation is a type of nonlinear approximation (DeVore, 1998). It was developed in a series of works in the context of parametric DEs (Beck et al., 2014, 2012; Bieri et al., 2010; Bonito et al., 2021; Chkifa et al., 2015b; Cohen et al., 2010, 2011; Hansen and Schwab, 2013a; Todor and Schwab, 2007; Tran et al., 2017), focusing on Taylor, Legendre or Chebyshev polynomial expansions in infinite dimensions. See, e.g., Cohen and DeVore (2015) and Adcock et al. (2022, Chpt. 3) for in-depth reviews. See §3.2 for a discussion holomorphic regularity of solutions of various parametric DEs, including the  $(\mathbf{b}, \varepsilon)$ -holomorphic class and relevant references. Our results on limits of learnability are from Adcock et al. (2024b) and, as noted, use ideas from information-based complexity. For overviews of this topic, see Novak (1988); Novak and Woźniakowski (2008, 2010); Traub et al. (1988).

In tandem with best  $s$ -term polynomial approximation theory, there has been a focus on learning polynomial approximations from data. Some early approaches in the context of parametric DEs included (adaptive) interpolation using sparse grids (Bäck et al., 2011; Babuška et al., 2007; Chkifa et al., 2014; Ganapathysubramanian and Zabaras, 2007; Gunzburger et al., 2014b; Ma and Zabaras, 2009; Mathelin et al., 2005; Nobile et al., 2008a,b; Xiu and Hesthaven, 2005). Another significant line of investigation has been on least-squares methods (Chkifa et al., 2015a; Cohen et al., 2013; Migliorati, 2013; Migliorati et al., 2014; Berveiller et al., 2006; Migliorati et al., 2013). See Guo et al. (2020); Hadigol and Doostan (2018) and Adcock et al. (2022, Chpt. 5) for reviews. While simpler than compressed sensing methods, these methods require *a priori* knowledge of a good multiindex set in which to construct the polynomial approximation, and are therefore best suited to the simpler *known anisotropy* setting. *Adaptive* least-squares methods (Migliorati, 2015, 2019) can address unknown anisotropy, but they currently lack theoretical guarantees. Compressed

sensing for learning polynomial approximations began with the works of Blatman and Sudret (2011); Doostan and Owhadi (2011); Mathelin and Gallivan (2012); Rauhut and Ward (2012); Yan et al. (2012); Yang and Karniadakis (2013), and has seen much subsequent development. See Hampton and Doostan (2017) and Adcock et al. (2022, Chpt. 7) and references therein. The focus in this work on *weighted* sparsity originated in Rauhut and Ward (2016). See also Adcock (2017); Adcock et al. (2019); Peng et al. (2014) and references therein. The results described in this work are based on Adcock et al. (2024a).

Early works on the application of DL to parametric DEs include Adcock et al. (2021); Cyr et al. (2020); Dal Santo et al. (2020); Geist et al. (2021); Khoo et al. (2021); Laakmann and Petersen (2021). See also Becker et al. (2023); Ciaci et al. (2022); Heiß et al. (2021, 2023); Khara et al. (2021); Lei et al. (2022); Scarabosio (2022) and references therein for more recent developments. Closely related to this topic is that of *operator learning* with DNNs – an area of significant current interest (Bhattacharya et al., 2021; Boullé and Townsend, 2023; Kovachki et al., 2023, 2024; Li et al., 2021; L. Lu et al., 2021). One can view holomorphic, Banach-valued functions as an example of holomorphic operators with specific parameterizations of the input space (Herrmann et al., 2022; Schwab et al., 2023). These operators constitute an important example in the field of operator learning for which there are strong guarantees on both parametric complexity (i.e., the number of DNN parameters required) and sample complexity. The relevant theory for these examples is based on the theory of holomorphic, Banach-valued functions described in this work. See Kovachki et al. (2024, §5.2) or Lanthaler (2023, §3.4) for further discussion.

See Cybenko (1989); Hornik et al. (1989) and Pinkus (1999) for the classical *universal approximation theory* of (shallow) NNs. The modern study of existence theory was initiated in Yarotsky (2017), and has since seen a wealth of developments. We review this literature in §7. The polynomial emulation results used in this work are based on Daws and C. Webster (2019); De Ryck et al. (2021); Opschoor et al. (2022).

The theory-to-practice gap in DL was studied empirically (Adcock and Dexter, 2021) and theoretically in Abedeljawad and Grohs (2023); Grohs and Voigtlaender (2023). Practical existence theorems were first established in Adcock et al. (2021); Adcock and Dexter (2021). Our results are based on Adcock et al. (2023a). This approach is not unique to high-dimensional regression problems. Similar ideas have been used to assert that it is possible to compute stable, accurate and efficient DNNs for inverse problems in imaging (Colbrook et al., 2022; Neyra-Nesterenko and Adcock, 2023). Here, one starts with a standard regularization problem, such as TV minimization, then exploits the fact that  $n$  steps of a first-order optimization method for solving this problem can be reinterpreted as a DNN with  $O(n)$  layers – a process known as *unrolling* (Monga et al., 2021) (see also Adcock and Hansen, 2021, Chpts. 19-21). Interestingly, this also leads to principled ways to design DNN architectures for DL in inverse

problems (Monga et al., 2021), although the trained networks may not be robust (Antun et al., 2023, 2020).

## 2 Problem statement and notation

In this section, we formalize the problem studied in this work and establish important notation.

Throughout this work,  $\mathcal{U} = [-1, 1]^{\mathbb{N}}$  and  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$  is a Banach space over  $\mathbb{R}$ . We write  $\mathbf{y} = (y_i)_{i \in \mathbb{N}}$  for the independent variable in  $\mathcal{U}$ . We equip  $\mathcal{U}$  with the uniform probability measure  $\varrho$  and, for  $1 \leq p \leq \infty$ , write  $L_{\varrho}^p(\mathcal{U}; \mathcal{V})$  for the Lebesgue–Bochner space of (equivalence classes of) strongly  $\varrho$ -measurable functions  $f : \mathcal{U} \rightarrow \mathcal{V}$  for which  $\|f\|_{L_{\varrho}^p(\mathcal{U}; \mathcal{V})} < \infty$ , where

$$\|f\|_{L_{\varrho}^p(\mathcal{U}; \mathcal{V})} := \begin{cases} \left( \int_{\mathcal{U}} \|f(\mathbf{y})\|_{\mathcal{V}}^p d\varrho(\mathbf{y}) \right)^{1/p} & 1 \leq p < \infty, \\ \text{ess sup}_{\mathbf{y} \in \mathcal{U}} \|f(\mathbf{y})\|_{\mathcal{V}} & p = \infty. \end{cases} \quad (2.1)$$

When  $\mathcal{V} = (\mathbb{R}, |\cdot|)$ , we just write  $L_{\varrho}^p(\mathcal{U})$  for the corresponding Lebesgue space of real-valued functions.

Let  $f \in L_{\varrho}^2(\mathcal{U}; \mathcal{V})$  be the unknown target function and consider sample points  $\mathbf{y}_1, \dots, \mathbf{y}_m \sim \text{i.i.d. } \varrho$ . Then we consider training data

$$\{(\mathbf{y}_i, f(\mathbf{y}_i) + e_i)\}_{i=1}^m \subseteq \mathcal{U} \times \mathcal{V}, \quad (2.2)$$

where  $e_i \in \mathcal{V}$  represents measurement error. In this work, we assume the noise is adversarial and of small norm, i.e.,  $\sum_{i=1}^m \|e_i\|_{\mathcal{V}}^2 \ll 1$ . Statistical models can also be considered – see, e.g., Migliorati et al. (2015). The problem is then to learn an approximation to  $f$  based on the data (2.2).

We require some notation for sequences indexed via possibly multiindices. Let  $d = \mathbb{N} \cup \{\infty\}$ . We write  $\mathbf{v} = (v_k)_{k=1}^d$  for an arbitrary multiindex in  $\mathbb{N}_0^d$ . Let  $\Lambda \subseteq \mathbb{N}_0^d$  be a finite or countable set of multiindices. For  $0 < p \leq \infty$ , we write  $\ell^p(\Lambda; \mathcal{V})$  for the space of all  $\mathcal{V}$ -valued sequences  $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$  for which  $\|\mathbf{c}\|_{p; \mathcal{V}} < \infty$ , where

$$\|\mathbf{c}\|_{p; \mathcal{V}} = \begin{cases} \left( \sum_{\mathbf{v} \in \Lambda} \|c_{\mathbf{v}}\|_{\mathcal{V}}^p \right)^{\frac{1}{p}} & 0 < p < \infty, \\ \sup_{\mathbf{v} \in \Lambda} \|c_{\mathbf{v}}\|_{\mathcal{V}} & p = \infty. \end{cases}$$

When  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}}) = (\mathbb{R}, |\cdot|)$ , we just write  $\ell^p(\Lambda)$  and  $\|\cdot\|_p$ . Given a sequence  $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$ , we define its support

$$\text{supp}(\mathbf{c}) = \{\mathbf{v} \in \Lambda : c_{\mathbf{v}} \neq 0\} \subseteq \Lambda.$$

Now let  $\Lambda = \mathbb{N}$  and  $0 < p \leq \infty$ . Given a real-valued sequence  $\mathbf{z} = (z_i)_{i \in \mathbb{N}} \in \ell^\infty(\mathbb{N})$ , we define its minimal monotone majorant as

$$\tilde{\mathbf{z}} = (\tilde{z}_i)_{i \in \mathbb{N}}, \quad \text{where } \tilde{z}_i = \sup_{j \geq i} |z_j|, \quad \forall i \in \mathbb{N} \quad (2.3)$$

and the *monotone*  $\ell^p$ -space  $\ell_M^p(\mathbb{N})$  as

$$\ell_M^p(\mathbb{N}) = \{\mathbf{z} \in \ell^\infty(\mathbb{N}) : \|\mathbf{z}\|_{p, \mathcal{M}} := \|\tilde{\mathbf{z}}\|_p < \infty\}. \quad (2.4)$$

Let  $d \in \mathbb{N} \cup \{\infty\}$ . We write  $\mathbf{e}_j$ ,  $j = 1, \dots, d$ , for the canonical basis vectors in  $\mathbb{R}^d$ . Given multiindices  $\mathbf{v} = (v_k)_{k=1}^d$  and  $\boldsymbol{\mu} = (\mu_k)_{k=1}^d$ , the inequality  $\mathbf{v} \geq \boldsymbol{\mu}$  is interpreted componentwise, i.e.,  $v_k \geq \mu_k, \forall k$ . We define  $\mathbf{v} > \boldsymbol{\mu}$  analogously. We write  $\mathbf{0}$  for the multiindex of zeros and  $\mathbf{1}$  for the multiindex of ones.

**Remark 2.1** (Other measures and domains). We work with the uniform probability measure on  $\mathcal{U} = [-1, 1]^{\mathbb{N}}$  for convenience. This means that we construct polynomial approximations using the corresponding multivariate Legendre polynomials. It is possible to work more generally, for instance by considering ultraspherical measures and the resulting Jacobi polynomials. See Adcock et al. (2024b). The cases where  $\mathcal{U} = \mathbb{R}^{\mathbb{N}}$  with the Gaussian measure (corresponding to Hermite polynomials) or  $\mathcal{U} = [0, \infty)^{\mathbb{N}}$  with the gamma measure (corresponding to Laguerre) have also been studied, although the theory is less complete. Results on best  $s$ -term polynomial approximation are known in this case, as are DNN existence theorems (Düng et al., 2023; Schwab and Zech, 2023). However, results akin to those we present in this work on learning sparse polynomial or DNN approximations from data are lacking.

**Remark 2.2** (Error metric). Throughout this work, we measure the error of the various approximations in the  $L_\rho^2(\mathcal{U}; \mathcal{V})$ -norm. The majority of the results we present extend to the stronger  $L_\rho^\infty(\mathcal{U}; \mathcal{V})$ -norm. In particular, the various upper bounds that establish algebraic rates with index  $1/2 - 1/p$  can be modified to show algebraic rates with index  $1 - 1/p$  (see, e.g., Adcock et al., 2024a). For succinctness, we do not present this modification.

### 3 Holomorphic functions of infinitely many variables

In this section, we first formally define the classes of holomorphic functions considered in this work and then discuss their relevance to parametric DEs.

#### 3.1 $(b, \varepsilon)$ -holomorphic functions

We commence with the definition of holomorphy.

**Definition 3.1** (Holomorphy). Let  $\mathcal{O} \subseteq \mathbb{C}^{\mathbb{N}}$  be an open set and  $\mathcal{V}$  be a Banach space. A function  $f : \mathcal{O} \rightarrow \mathcal{V}$  is *holomorphic in*  $\mathcal{O}$  if it is holomorphic with

respect to each variable in  $\mathcal{O}$ . That is to say, for any  $\mathbf{z} \in \mathcal{O}$  and any  $j \in \mathbb{N}$ , the following limit exists in  $\mathcal{V}$ :

$$\lim_{\substack{h \in \mathbb{C} \\ h \rightarrow 0}} \frac{f(\mathbf{z} + h\mathbf{e}_j) - f(\mathbf{z})}{h} \in \mathcal{V}.$$

**Definition 3.2** (Holomorphic extension). Let  $\mathcal{U} \subset \mathcal{O} \subseteq \mathbb{C}^{\mathbb{N}}$  be an open set. A function  $f : \mathcal{U} \rightarrow \mathcal{V}$  has a *holomorphic extension to  $\mathcal{O}$*  (or simply, *is holomorphic in  $\mathcal{O}$* ) if there is a  $\tilde{f} : \mathcal{O} \rightarrow \mathcal{V}$  that is holomorphic in  $\mathcal{O}$  for which  $\tilde{f}|_{\mathcal{U}} = f$ . In this case, we also define  $\|f\|_{L^\infty(\mathcal{O}; \mathcal{V})} := \|\tilde{f}\|_{L^\infty(\mathcal{O}; \mathcal{V})} = \sup_{\mathbf{z} \in \mathcal{O}} \|\tilde{f}(\mathbf{z})\|_{\mathcal{V}}$  or, when  $\mathcal{V} = \mathbb{C}$ , simply  $\|f\|_{L^\infty(\mathcal{O})}$ . If  $\mathcal{O}$  is a closed set, then we say that  $f$  is holomorphic in  $\mathcal{O}$  if it has a holomorphic extension to some open neighborhood of  $\mathcal{O}$ .

In this work, we consider functions that possess holomorphic extensions to (unions of) tensor-products of Bernstein ellipses. The *Bernstein ellipse* with parameter  $\rho > 1$  is the set

$$\mathcal{E}(\rho) = \left\{ (z + z^{-1})/2 : z \in \mathbb{C}, 1 \leq |z| \leq \rho \right\} \subset \mathbb{C}.$$

See, e.g., Trefethen (2013, Chpt. 8). This is an axis-aligned ellipse with foci at  $\pm 1$  and with major and minor semiaxis lengths given by  $(\rho + 1/\rho)/2$  and  $(\rho - 1/\rho)/2$ , respectively. By convention, we let  $\mathcal{E}(\rho) = [-1, 1]$  when  $\rho = 1$ . Next, we define the *Bernstein polyellipse* with parameter  $\boldsymbol{\rho} = (\rho_i)_{i \in \mathbb{N}} > \mathbf{1}$  by

$$\mathcal{E}(\boldsymbol{\rho}) = \mathcal{E}(\rho_1) \times \mathcal{E}(\rho_2) \times \cdots \subset \mathbb{C}^{\mathbb{N}}.$$

**Definition 3.3** ( $(\mathbf{b}, \varepsilon)$ -holomorphic functions). Let  $\mathbf{b} = (b_j)_{j \in \mathbb{N}} \in [0, \infty)^{\mathbb{N}}$  with  $\mathbf{b} \in \ell^1(\mathbb{N})$  and  $\varepsilon > 0$ . A function  $f : \mathcal{U} \rightarrow \mathcal{V}$  is  $(\mathbf{b}, \varepsilon)$ -holomorphic if it has a holomorphic extension to the region

$$\mathcal{R}(\mathbf{b}, \varepsilon) = \bigcup \left\{ \mathcal{E}(\boldsymbol{\rho}) : \boldsymbol{\rho} \geq \mathbf{1}, \sum_{j=1}^{\infty} \left( \frac{\rho_j + \rho_j^{-1}}{2} - 1 \right) b_j \leq \varepsilon \right\} \subset \mathbb{C}^{\mathbb{N}}. \quad (3.1)$$

See, e.g., Chkifa et al. (2015b); Schwab and Zech (2019). We discuss the motivations for this definition next. Beforehand, it is worth noting that the parameter  $\varepsilon$  is technically redundant, since we can always rescale  $\mathbf{b}$ . In the remainder of this work, we assume that  $\varepsilon = 1$ . For later use, we now introduce the class of  $(\mathbf{b}, 1)$ -holomorphic functions with norm at most one, i.e.,

$$\mathcal{H}(\mathbf{b}) = \left\{ f : \mathcal{U} \rightarrow \mathcal{V} \text{ } (\mathbf{b}, 1)\text{-holomorphic} : \|f\|_{L^\infty(\mathcal{R}(\mathbf{b}, 1); \mathcal{V})} \leq 1 \right\}.$$

**Remark 3.4** (Functions of finitely many variables). In finite dimensions (in particular,  $d = 1$ ), one normally considers functions that are holomorphic in a single Bernstein polyellipse. For such functions, it is well known that one

can find polynomial approximations that converge exponentially fast. As we see later, in infinite dimensions we use the assumption of holomorphy in the region (3.1) in order to obtain algebraic rates of convergence.

It is, however, worth noting that finite-dimensional holomorphy in a Bernstein polyellipse implies  $(\mathbf{b}, \varepsilon)$ -holomorphy. Let  $f : [-1, 1]^d \rightarrow \mathbb{R}$  be a function of finitely many variables that is holomorphic in  $\mathcal{E}(\bar{\rho}_1) \times \cdots \times \mathcal{E}(\bar{\rho}_d) \subset \mathbb{C}^d$ . We can extend  $f$  to a function  $\bar{f} : \mathcal{U} \rightarrow \mathbb{C}$  in the standard way as  $\bar{f}(y_1, y_2, \dots) = f(y_1, \dots, y_d)$ ,  $\forall \mathbf{y} = (y_i)_{i \in \mathbb{N}} \in \mathcal{U}$ . Now define

$$b_i = \varepsilon((\bar{\rho}_i + \bar{\rho}_i^{-1})/2 - 1)^{-1}, \quad i \in [d], \quad b_i = 0, \quad i \in \mathbb{N} \setminus [d].$$

Then  $\bar{f}$  is  $(\mathbf{b}, \varepsilon)$ -holomorphic. Therefore, all the results that follow on approximation of  $(\mathbf{b}, \varepsilon)$ -holomorphic functions also apply *mutatis mutandis* to functions of  $d$  variables that are holomorphic in a single Bernstein ellipse.

### 3.2 Holomorphy and parametric DEs

As mentioned,  $(\mathbf{b}, \varepsilon)$ -holomorphic functions were first developed in the context of parametric DEs. There is now a wealth of literature that establishes that solution maps of many parametric DEs possess this regularity.

The classical example of a parametric PDEs is the parametric stationary diffusion equation with parametrized diffusion coefficient

$$-\nabla_{\mathbf{x}} \cdot (a(\mathbf{x}, \mathbf{y}) \nabla_{\mathbf{x}} u(\mathbf{x}, \mathbf{y})) = F(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad u(\mathbf{x}, \mathbf{y}) = 0, \quad \mathbf{x} \in \partial\Omega. \quad (3.2)$$

Here  $\mathbf{x} \in \Omega$  is the spatial variable,  $\Omega \subset \mathbb{R}^k$  is the domain, which is assumed to have Lipschitz boundary, and  $\nabla_{\mathbf{x}}$  is the gradient with respect to  $\mathbf{x}$ . The function  $F$  is the forcing term, and is assumed to be nonparametric. We assume that (3.2) satisfies a *uniform ellipticity condition*

$$\operatorname{ess\,inf}_{\mathbf{x} \in \Omega} a(\mathbf{x}, \mathbf{y}) \geq r, \quad \forall \mathbf{y} \in \mathcal{U}, \quad (3.3)$$

for some  $r > 0$ . Let  $\mathcal{V} = H_0^1(\Omega)$  be the standard Sobolev space of functions with weak first-order derivatives in  $L^2(\Omega)$  and traces vanishing on  $\partial\Omega$ . Then the problem (3.2) has a well-defined *parametric solution map*

$$u : \mathcal{U} \rightarrow \mathcal{V}, \quad \mathbf{y} \mapsto u(\cdot, \mathbf{y}). \quad (3.4)$$

Now suppose that the diffusion coefficient has the affine parametrization

$$a(\mathbf{x}, \mathbf{y}) = a_0(\mathbf{x}) + \sum_{j=1}^{\infty} y_j \psi_j(\mathbf{x}), \quad (3.5)$$

where the functions  $a_0, \psi_1, \psi_2, \dots \in L^\infty(\Omega)$ . Notice that uniform ellipticity (3.3) for this problem is equivalent to the condition

$$\sum_{j=1}^{\infty} |\psi_j(\mathbf{x})| \leq a_0(\mathbf{x}) - r, \quad \forall \mathbf{x} \in \Omega.$$

Under this assumption, one can show (see, e.g., Adcock et al., 2022, Prop. 4.9) that the solution map (3.4) is well defined and  $(\mathbf{b}, \varepsilon)$ -holomorphic for  $0 < \varepsilon < r$  with

$$\mathbf{b} = (b_i)_{i=1}^{\infty}, \quad \text{with } b_i = \|\psi_i\|_{L^\infty(\Omega)}, \quad \forall i \in \mathbb{N}. \quad (3.6)$$

This example was first considered in Bieri et al. (2010); Cohen et al. (2010, 2011). Subsequently, these results were generalized to many other classes of parametric DEs and PDEs. One such extension (see, e.g., Chkifa et al., 2015b) considers the problem (3.2) with various nonaffine parametric diffusion coefficients, such as

$$a(\mathbf{x}, \mathbf{y}) = a_0(\mathbf{x}) + \left( \sum_{j=1}^{\infty} y_j \psi_j(\mathbf{x}) \right)^2 \quad \text{or} \quad a(\mathbf{x}, \mathbf{y}) = \exp \left( \sum_{j=1}^{\infty} y_j \psi_j(\mathbf{x}) \right).$$

This is part of a general framework for showing holomorphy for parametric weak problems in Hilbert spaces. See Chkifa et al. (2015b, Thm. 4.1), Cohen and DeVore (2015, Cor. 2.4) or Adcock et al. (2022, §4.3.1). See also Kunoth and Schwab (2013); Rauhut and Schwab (2017). Another related general framework considers parametric implicit operator equations (see Chkifa et al., 2015b, Thm. 4.3, Cohen and DeVore, 2015, Thm. 2.5, or Adcock et al., 2022, §4.3.1).

These frameworks can be used to establish holomorphy results for various problems beyond (3.2). This includes parabolic problems (see Chkifa et al., 2015b, §5.1, or Cohen and DeVore, 2015, §2.2) and various types of nonlinear, elliptic PDEs (see Chkifa et al., 2015b, §5.2, or Cohen and DeVore, 2015, §2.3). Another prominent example includes PDEs such as (3.2) over parametrized domains  $\Omega = \Omega_{\mathbf{y}}$ . Holomorphy of the solution map is often referred to as *shape holomorphy* in this context. See Chkifa et al. (2015b, §5.3), Cohen and DeVore (2015, §2.2), Castrillón-Candás et al. (2016); Cohen et al. (2018) and references therein. Other examples include parametric hyperbolic problems (Hoang and Schwab, 2012) and certain classes of parametric control problems (Kunoth and Schwab, 2013).

See Cohen and DeVore (2015) or Adcock et al. (2022, Chpt. 4) for further reviews. Note that the above problems all involve PDEs. Classes of parametric ODEs also admit holomorphic regularity. See Adcock et al. (2022, §4.1) and Hansen and Schwab (2013b). Recent work in this direction has also shown such regularity for parametric ODEs arising from diffusion on graphs (Ajavon, 2024).

### 3.3 Known and unknown anisotropy

Functions that are  $(\mathbf{b}, \varepsilon)$ -holomorphic are *anisotropic*: they depend more smoothly on some variables rather than others. This behavior is dictated by the parameter  $\mathbf{b}$ , with a larger (smaller) value of an entry  $b_i$  corresponding to less (more) smoothness with respect to the variable  $y_i$ . This can be seen by noting that the condition

$$\sum_{i=1}^{\infty} \left( \frac{\rho_i + \rho_i^{-1}}{2} - 1 \right) b_i \leq 1,$$

in (3.1) holds only for smaller (larger) values of  $\rho_i$  when  $b_i$  is large (small), meaning that  $f$  only admits an extension in the  $y_i$  variable to a relatively small (large) Bernstein ellipse.

In practice, *a priori* analysis of a given problem may establish that the target function is  $(\mathbf{b}, \varepsilon)$ -holomorphic for some  $\mathbf{b}$ . However, an optimal value for  $\mathbf{b}$  may be unknown. This situation arises in parametric DEs. For problems such as the stationary diffusion equation (3.2) with affine diffusion, one can find a sufficient value (3.6) of  $\mathbf{b}$  for which the parametric solution map is  $(\mathbf{b}, \varepsilon)$ -holomorphic, although this value may not be sharp. In general, since  $\mathcal{R}(\mathbf{b}, \varepsilon) \subseteq \mathcal{R}(\mathbf{b}', \varepsilon)$  whenever  $\mathbf{b}' \leq \mathbf{b}$ , it is difficult to know whether a value of  $\mathbf{b}$  obtained from some analysis is optimal. Further, for more complicated problems, such as some of those described above, one can derive theoretical guarantees that assert holomorphy, but without an estimate for the region itself (Cohen and DeVore, 2015, Rem. 2.6).

For this reason, the primary focus of this work is the *unknown anisotropy* setting. We may assume the target function  $f$  is  $(\mathbf{b}, \varepsilon)$ -holomorphic for some  $\mathbf{b}$ , but we do not have access to  $\mathbf{b}$  itself. In particular, this means that we cannot use  $\mathbf{b}$  to design a method for learning  $f$  from data: the holomorphy assumption can only be used to provide bounds for the resulting approximation error.

**Remark 3.5.** To emphasize this consideration further, consider the function

$$f(y_1, y_2) = \sin(1000y_2)/(1.1 - y_1).$$

This function is entire with respect to the variable  $y_2$ , but has a pole at  $y_1 = 1.1$ . By Remark 3.4, it is  $(\mathbf{b}, 1)$ -holomorphic for  $\mathbf{b} = (b_1, b_2, \dots) = (10, 0, 0, \dots)$ . When building a polynomial approximation to  $f$ , knowledge of  $\mathbf{b}$  may lead one to include only low-degree terms in  $y_2$  and more higher-degree terms in  $y_1$ . Yet this is completely the opposite of what one should do in practice. The function  $\sin(1000y_2)$ , while entire, is highly oscillatory, and can only be resolved by using high-degree polynomials. On the other hand, the function

$$g(y_1, y_2) = 1/(1.1 - y_1)$$

is also  $(\mathbf{b}, 1)$ -holomorphic with the same  $\mathbf{b}$ . Yet, it requires no nonconstant terms in  $y_2$  in order to approximate it accurately with a polynomial.



### 3.4 $\ell^p$ -summability and the $\mathcal{H}(p)$ and $\mathcal{H}(p, M)$ classes

As we will see in §4, it is generally impossible to approximate a  $(\mathbf{b}, \varepsilon)$ -holomorphic function without a further assumption on the sequence  $\mathbf{b}$ . In particular, the terms of  $\mathbf{b}$  need to decay sufficiently fast. Henceforth, we will assume that  $\mathbf{b}$  is not just in  $\ell^1(\mathbb{N})$ , but that, for some  $0 < p < 1$ , either

$$\mathbf{b} \in \ell^p(\mathbb{N}) \text{ or } \mathbf{b} \in \ell_M^p(\mathbb{N}).$$

Here we recall that  $\ell_M^p(\mathbb{N})$  is the monotone  $\ell^p$ -space (2.4). Since we consider the unknown anisotropy setting, we also define the function classes

$$\mathcal{H}(p) = \bigcup \left\{ \mathcal{H}(\mathbf{b}) : \mathbf{b} \in \ell^p(\mathbb{N}), \mathbf{b} \in [0, \infty)^\mathbb{N}, \|\mathbf{b}\|_p \leq 1 \right\} \quad (3.7)$$

and

$$\mathcal{H}(p, M) = \bigcup \left\{ \mathcal{H}(\mathbf{b}) : \mathbf{b} \in \ell_M^p(\mathbb{N}), \mathbf{b} \in [0, \infty)^\mathbb{N}, \|\mathbf{b}\|_{p, M} \leq 1 \right\}. \quad (3.8)$$

In what follows we will ask for a method to provide a uniform error bound (depending on  $m$ , the amount of data, and  $p$ ) for *any* function in these classes.

## 4 Best $s$ -term polynomial approximation

In this section, we make a first foray into polynomial approximation of holomorphic functions. We introduce orthogonal polynomials and orthogonal polynomial expansions in  $L_\rho^2(\mathcal{U}; \mathcal{V})$ . We then introduce best  $s$ -term polynomial approximation and provide a key theorem on algebraic convergence for it. See, e.g., Adcock et al. (2022, Chpts. 2-3) or Chkifa et al. (2015b); Cohen and DeVore (2015), for further information on this material.

### 4.1 Orthogonal polynomials

Let  $P_0, P_1, \dots$  denote the classical one-dimensional Legendre polynomials with the normalization

$$\|P_\nu\|_{L^\infty([-1, 1])} = P_\nu(1) = 1. \quad (4.1)$$

We consider the orthonormalized version of these polynomials. Since

$$\int_{-1}^1 |P_\nu(y)|^2 dy = \frac{2}{2\nu + 1}, \quad \forall \nu \in \mathbb{N}_0,$$

we define these as

$$\psi_\nu(y) = \sqrt{2\nu + 1} P_\nu(y), \quad \forall \nu \in \mathbb{N}_0. \quad (4.2)$$

The set  $\{\psi_\nu\}_{\nu \in \mathbb{N}_0} \subset L_\rho^2([-1, 1])$  forms an orthonormal basis, where  $\rho$  is the one-dimensional uniform measure.

We construct an orthonormal basis of  $L^2_\varrho(\mathcal{U})$  by tensorization. Let

$$\mathcal{F} = \left\{ \mathbf{v} = (v_k)_{k=1}^\infty \in \mathbb{N}_0^\mathbb{N} : |\{k : v_k \neq 0\}| < \infty \right\}$$

be the set of multiindices with finitely many nonzero terms and set

$$\Psi_{\mathbf{v}}(\mathbf{y}) = \prod_{i \in \mathbb{N}} \psi_{v_i}(y_i), \quad \forall \mathbf{y} \in \mathcal{U}, \mathbf{v} \in \mathcal{F}.$$

Note that  $\psi_0 = 1$  by construction. Therefore, this is equivalent to a product over finitely many terms, i.e.,

$$\Psi_{\mathbf{v}}(\mathbf{y}) = \prod_{i: v_i \neq 0} \psi_{v_i}(y_i). \quad (4.3)$$

Given these functions, it can now be shown that the set

$$\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathcal{F}} \subset L^2_\varrho(\mathcal{U}) \quad (4.4)$$

constitutes an orthonormal basis for  $L^2_\varrho(\mathcal{U})$  (Cohen and DeVore, 2015, §3).

Using (4.1)–(4.3), we also deduce that

$$\|\Psi_{\mathbf{v}}\|_{L^\infty_\varrho(\mathcal{U}; \mathcal{V})} = |\Psi_{\mathbf{v}}(\mathbf{1})| = \prod_{k \in \mathbb{N}} \sqrt{2v_k + 1} =: u_{\mathbf{v}}. \quad (4.5)$$

The values  $u_{\mathbf{v}}$  will be of use later. For convenience, we also define

$$\mathbf{u} = (u_{\mathbf{v}})_{\mathbf{v} \in \mathcal{F}}. \quad (4.6)$$

## 4.2 Orthogonal polynomial expansions

Let  $f \in L^2_\varrho(\mathcal{U}; \mathcal{V})$ . Then we have the convergent expansion (in  $L^2_\varrho(\mathcal{U}; \mathcal{V})$ )

$$f = \sum_{\mathbf{v} \in \mathcal{F}} c_{\mathbf{v}} \Psi_{\mathbf{v}}, \quad \text{where } c_{\mathbf{v}} = \int_{\mathcal{U}} f(\mathbf{y}) \Psi_{\mathbf{v}}(\mathbf{y}) \, d\varrho(\mathbf{y}) \in \mathcal{V}. \quad (4.7)$$

Note that the *coefficients*  $c_{\mathbf{v}}$  are elements of  $\mathcal{V}$  and defined by Bochner integrals. For convenience, we denote the infinite sequence of coefficients of  $f$  by

$$\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \mathcal{F}}. \quad (4.8)$$

Parseval's identity gives that  $\mathbf{c} \in \ell^2(\mathcal{F}; \mathcal{V})$  with  $\|\mathbf{c}\|_{2; \mathcal{V}} = \|f\|_{L^2_\varrho(\mathcal{U}; \mathcal{V})}$ .

### 4.3 Best $s$ -term polynomial approximation

Let  $s \in \mathbb{N}$ . An  $s$ -term polynomial approximation to  $f$  is an approximation

$$f \approx f_S = \sum_{\mathbf{v} \in S} c_{\mathbf{v}} \Psi_{\mathbf{v}} \quad (4.9)$$

for some multiindex set  $S \subset \mathcal{F}$ ,  $|S| = s$ . This raises the question: which index set  $S$  should one choose? One answer is to select a set  $S$  which provides the best approximation, an approach known as *best  $s$ -term approximation* and itself a type of nonlinear approximation (DeVore, 1998). Formally, a *best  $s$ -term approximation*  $f_s$  of  $f$  (with respect to the  $L^2_{\mathcal{Q}}$ -norm) is defined as

$$f_s = f_{S^*}, \quad \text{where } S^* \in \operatorname{argmin}\{\|f - f_S\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})} : S \subset \mathcal{F}, |S| = s\}. \quad (4.10)$$

This approximation can also be characterized in terms of the coefficients of  $f$ . Due to Parseval's identity, the error of the approximation  $f_s$  is precisely

$$\|f - f_s\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})}^2 = \sum_{\mathbf{v} \in \mathcal{F} \setminus S} \|c_{\mathbf{v}}\|_{\mathcal{V}}^2. \quad (4.11)$$

Therefore, any set  $S^*$  that yields a best  $s$ -term approximation consists of multi-indices corresponding to the largest  $s$  coefficient norms  $\|c_{\mathbf{v}}\|_{\mathcal{V}}$ . Specifically,

$$S^* = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_s\},$$

where  $\mathbf{v}_1, \mathbf{v}_2, \dots$  is an ordering of the multiindex set  $\mathcal{F}$  such that  $\|c_{\mathbf{v}_1}\|_{\mathcal{V}} \geq \|c_{\mathbf{v}_2}\|_{\mathcal{V}} \geq \dots$ . It follows immediately from this and (4.11) that the error

$$\|f - f_s\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})}^2 = \sum_{i > s} \|c_{\mathbf{v}_i}\|_{\mathcal{V}}^2 \quad (4.12)$$

is precisely the  $\ell^2(\mathcal{F}; \mathcal{V})$ -norm of the sequence of coefficients (4.8), excluding those coefficients with indices in  $S^*$ . Note that  $S^*$  (and therefore  $f_s$ ) is generally nonunique. However, this fact causes no difficulties in what follows.

### 4.4 Rates of best $s$ -term polynomial approximation

The best  $s$ -term approximation describes the best possible approximation obtainable with an  $s$ -term polynomial expansion. It is therefore important to provide bounds for this *benchmark* approximation in the case of holomorphic functions. Besides providing insight into limits of approximation, these bounds are also useful when we come to learn sparse polynomial approximations from data.

As we discussed above, Parseval's identity relates questions of best  $s$ -term approximation to  $f$  (in the  $L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})$  norm) to approximation of its coefficients

$\mathbf{c}$  (in the  $\ell^2(\mathcal{F}; \mathcal{V})$ -norm). This motivates one to study best  $s$ -term approximation of sequences, rather than functions. To this end, we now introduce some additional notation. Let  $\Lambda \subseteq \mathcal{F}$ ,  $0 < p \leq \infty$ ,  $\mathbf{c} \in \ell^p(\Lambda; \mathcal{V})$  and  $s \in \mathbb{N}_0$  with  $s \leq |\Lambda|$ . The  $\ell^p$ -norm best  $s$ -term approximation error of the sequence  $\mathbf{c}$  is defined as

$$\sigma_s(\mathbf{c})_{p; \mathcal{V}} = \min \left\{ \|\mathbf{c} - \mathbf{z}\|_{p; \mathcal{V}} : \mathbf{z} \in \ell^p(\Lambda; \mathcal{V}), |\text{supp}(\mathbf{z})| \leq s \right\}. \quad (4.13)$$

Note that when  $p = 2$  we have  $\sigma_s(\mathbf{c})_{2; \mathcal{V}} = \|f - f_s\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})}$ , where  $f_s$  is a best  $s$ -term approximation to  $f$  (4.10).

We now state a well known result regarding the best  $s$ -term approximation (of coefficients) of holomorphic functions (see, e.g., Adcock et al., 2022, Thm. 3.28, or Cohen and DeVore, 2015, §3.2).

**Theorem 4.1** (Algebraic convergence of the best  $s$ -term approximation). *Let  $\mathbf{b} \in [0, \infty)^{\mathbb{N}}$  be such that  $\mathbf{b} \in \ell^p(\mathbb{N})$  for some  $0 < p < 1$ . Then for any  $s \in \mathbb{N}$  and  $p \leq q \leq 2$ , there exists a set  $S \subset \mathcal{F}$  with  $|S| \leq s$  such that*

$$\sigma_s(\mathbf{c})_{p; \mathcal{V}} \leq \|\mathbf{c} - \mathbf{c}_S\|_{q; \mathcal{V}} \leq C(\mathbf{b}, p) \cdot s^{\frac{1}{q} - \frac{1}{p}},$$

for all  $f \in \mathcal{H}(\mathbf{b})$  with coefficients  $\mathbf{c}$  as in (4.8).

**Remark 4.2** (Sharpness of the algebraic rate). Theorem 4.1 provides an upper bound for the best  $s$ -term approximation error. However, it is also possible to provide a lower bound. As shown in Adcock and Monte (2023, Thm. 5.6), the are choices of  $\mathbf{b} \in \ell^p(\mathbb{N})$  and functions  $f \in \mathcal{H}(\mathbf{b})$  for which

$$\limsup_{s \rightarrow \infty} \left\{ s^{\frac{1}{r} - \frac{1}{2}} \cdot \|f - f_s\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})} \right\} = +\infty$$


for any  $0 < r < p$ . Thus, the algebraic exponent  $1/2 - 1/p$  is sharp.

## 4.5 How high is high dimensional?

As in Remark 3.4, let  $f : [-1, 1]^d \rightarrow \mathbb{R}$  be holomorphic in a single Bernstein ellipse  $\mathcal{E}(\bar{\rho}) = \mathcal{E}(\bar{\rho}_1) \times \cdots \times \mathcal{E}(\bar{\rho}_d) \subset \mathbb{C}^d$ . Then its best  $s$ -term polynomial approximation converges with exponential rate in  $s^{1/d}$ : specifically

$$\|f - f_s\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})} \leq \|f\|_{L^\infty(\mathcal{R}; \mathcal{V})} C(\epsilon) \sqrt{s} \exp \left( - \left( \frac{sd! \prod_{j=1}^d \log(\bar{\rho}_j)}{1 + \epsilon} \right)^{1/d} \right),$$

for any  $0 < \epsilon < 1$  and all  $s \geq \bar{s}$ , where  $\bar{s}$  is a constant depending on  $d$ ,  $\epsilon$  and  $\bar{\rho}$  only. See Tran et al. (2017) or Adcock et al. (2022, Thm. 3.21). In low dimensions, this rate accurately predicts the convergence of the best  $s$ -term approximation. However, for larger  $d$  this rate is generally not witnessed for



**FIGURE 1** Best  $s$ -term approximation error in the  $L^2_{\rho}(\mathcal{U})$ -norm for (4.14) with  $\delta_i = i^{3/2}$ . This figure also shows the exponential rate “exp. rate”, defined as  $C_{\text{exp}} \cdot \exp\left(-\left(sd! \prod_{i=1}^d \log(\rho_i)\right)^{1/d}\right)$ , where  $\rho_i$  is such that  $(\rho_i + 1/\rho_i)/2 = 1 + \delta_i$ , and the algebraic rate “alg. rate”, defined as  $C_{\text{alg}} \cdot s^{-1}$ . The constants  $C_{\text{exp}}$  and  $C_{\text{alg}}$  are chosen empirically to aid visualization.

any finite  $s$ , due to the exceedingly slow growth of the term  $s^{1/d}$ . For example,  $s^{1/d} \leq 3.2$  in  $d = 8$  dimensions for all  $1 \leq s \leq 10,000$  (the range of  $s$  that is quite reasonable in practice).

For larger  $d$ , the algebraic rates of Theorem 4.1 better describe the convergence behavior. An example of this effect is shown in Fig. 1 for

$$f(\mathbf{y}) = \prod_{i=1}^d \frac{(2\delta_i + \delta_i^2)^{1/2}}{y_i + 1 + \delta_i}, \quad \forall \mathbf{y} = (y_i)_{i=1}^d \in [-1, 1]^d. \quad (4.14)$$

This function is holomorphic in  $\mathcal{E}(\boldsymbol{\rho})$  for any  $\rho_i$  satisfying  $(\rho_i + 1/\rho_i)/2 < 1 + \delta_i$ . Thus, by Remark 3.4, it is also  $(\mathbf{b}, \varepsilon)$ -holomorphic for any  $b_i > \varepsilon/\delta_i$ . In this figure, we consider  $\delta_i = i^{3/2}$ , meaning that  $\mathbf{b}$  can be chosen so that  $\mathbf{b} \in \ell^p(\mathbb{N})$  for any  $2/3 < p < 1$ . Thus, Theorem 4.1 predicts dimension-independent algebraic convergence with order that is arbitrarily close to  $1/2 - 1/(2/3) = -1$ . As we see, in  $d = 4$  dimensions the error follows the exponential rate. However, once  $d = 16$  or  $d = 32$ , the algebraic rate better predicts the true error.

This discussion and example motivates the study of algebraic rates. Simply put, they typically better describe the convergence behavior, even in finite dimensions as long as  $d$  is not too small (e.g.,  $d \geq 16$ ).

## 5 Limits of learnability from data

We now consider learning  $(b, \varepsilon)$ -holomorphic functions from data. In this section, we present lower bounds for the amount of data that is necessary to learn such functions. We do this by appealing to techniques from information-based complexity (Novak, 1988; Novak and Woźniakowski, 2008, 2010; Traub et al., 1988) – in particular, *(adaptive) m-widths*.

Since our present focus is on lower bounds, in this section we do not assume that the training data takes the form of pointwise evaluations of the target function, as in (2.2). In fact, we can allow for arbitrary (adaptive) linear sampling operators. We will, for simplicity, consider only scalar-valued function approximation, i.e.,  $(\mathcal{V}, \|\cdot\|_{\mathcal{V}}) = (\mathbb{R}, |\cdot|)$ . See Remark 5.3 for discussion.

### 5.1 Adaptive $m$ -widths

Let  $C(\mathcal{U})$  be the Banach space of continuous functions  $f : \mathcal{U} \rightarrow \mathbb{R}$ . We define an *adaptive sampling operator* as a map of the form

$$\mathcal{S} : C(\mathcal{U}) \rightarrow \mathbb{R}^m, \quad \mathcal{S}(f) = \begin{bmatrix} S_1(f) \\ S_2(f; S_1(f)) \\ \vdots \\ S_m(f; S_1(f), \dots, S_{m-1}(f)) \end{bmatrix},$$

where  $S_1 : C(\mathcal{U}) \rightarrow \mathbb{R}$  is a bounded linear functional and, for  $i = 2, \dots, m$ ,  $S_i : C(\mathcal{U}) \times \mathbb{R}^{i-1} \rightarrow \mathbb{R}$  is bounded and linear in its first component.

This definition includes standard bounded linear operators as a special case. However, it also allows for situations where the  $i$ th sample is selected adaptively based on the existing measurements. In machine learning, this is commonly referred to as *active learning* (Settles, 2012). In information-based complexity, it is commonly termed *adaptive information* (Novak and Woźniakowski, 2008, Sec. 4.1.1). As noted, our primary concern in this work is where each  $S_i$  is a pointwise evaluation operator (so-called *standard information* (Novak and Woźniakowski, 2008, Sec. 4.1.1)). In this case,  $\mathcal{S}(f) = (f(y_i))_{i=1}^m \in \mathbb{R}^m$ , where  $y_i \in \mathcal{U}$  is the  $i$ th sample point, which is potentially chosen adaptively based on the previous measurements  $f(y_1), \dots, f(y_{i-1})$ .

**Definition 5.1** (Adaptive  $m$ -width). The *(adaptive)  $m$ -width* of a subset  $\mathcal{K} \subseteq C(\mathcal{U})$  is

$$\Theta_m(\mathcal{K}) = \inf_{\mathcal{S}, \mathcal{R}} \sup_{f \in \mathcal{K}} \|f - \mathcal{R}(\mathcal{S}(f))\|_{L^2_{\varrho}(\mathcal{U})}, \quad (5.1)$$

where the infimum is taken over all adaptive sampling operators  $\mathcal{S} : C(\mathcal{U}) \rightarrow \mathbb{R}^m$  and reconstruction maps  $\mathcal{R} : \mathbb{R}^m \rightarrow C(\mathcal{U})$ .

The adaptive  $m$ -width is related to the concept of *information complexity* (Novak and Woźniakowski, 2008, §4.1.4). It measures how well one can approximate functions from  $\mathcal{K}$  using an arbitrary (adaptive) linear sampling operator  $\mathcal{S}$  and an arbitrary (potentially nonlinear) reconstruction map  $\mathcal{R}$ . Due to the inner supremum, the approximation is measured in a worst-case (or uniform) sense; in other words, the sampling-recovery pair  $(\mathcal{S}, \mathcal{R})$  is required to provide a guaranteed bound simultaneously for all functions in the class.

## 5.2 Lower bounds for adaptive $m$ -widths

We now present the main results of this section. Following §3.4, we let

$$\theta_m(p) = \Theta_m(\mathcal{H}(p)), \quad \theta_m(p, M) = \Theta_m(\mathcal{H}(p, M)).$$

We also consider slightly weaker notions, which are defined as follows.

$$\begin{aligned} \overline{\theta}_m(p) &= \sup \left\{ \Theta_m(\mathcal{H}(\mathbf{b})) : \mathbf{b} \in \ell^p(\mathbb{N}), \mathbf{b} \in [0, \infty)^\mathbb{N}, \|\mathbf{b}\|_p \leq 1 \right\}, \\ \overline{\theta}_m(p, M) &= \sup \left\{ \Theta_m(\mathcal{H}(\mathbf{b})) : \mathbf{b} \in \ell_M^p(\mathbb{N}), \mathbf{b} \in [0, \infty)^\mathbb{N}, \|\mathbf{b}\|_{p, M} \leq 1 \right\}. \end{aligned}$$

The widths  $\theta_m(p)$  and  $\theta_m(p, M)$  pertain to the unknown anisotropy setting. A sampling-recovery pair  $(\mathcal{S}, \mathcal{R})$  does not have access to the anisotropy parameter  $\mathbf{b}$ , and must work uniformly for all holomorphic functions with  $\ell^p$  or  $\ell_M^p$ -summable  $\mathbf{b}$  (of norm at most one). Conversely,  $\overline{\theta}_m(p)$  and  $\overline{\theta}_m(p, M)$  are weaker and pertain to the known anisotropy setting, since they allow the sampling-recovery pair  $(\mathcal{S}, \mathcal{R})$  to depend on  $\mathbf{b}$ . In particular, we have

$$\theta_m(p) \geq \overline{\theta}_m(p) \text{ and } \theta_m(p, M) \geq \overline{\theta}_m(p, M).$$

**Theorem 5.2** (Limits of learnability). *Let  $m \geq 1$  and  $0 < p < 1$ . Then*

- the  $m$ -widths  $\overline{\theta}_m(p)$  and  $\overline{\theta}_m(p, M)$  satisfy

$$\overline{\theta}_m(p) \geq \overline{\theta}_m(p, M) \geq c \cdot 2^{-\frac{1}{p}} \cdot m^{\frac{1}{2} - \frac{1}{p}}, \quad (5.2)$$

- the  $m$ -width  $\theta_m(p)$  satisfies

$$\theta_m(p) \geq c \cdot 2^{\frac{1}{2} - \frac{2}{p}}, \quad (5.3)$$

- and the  $m$ -width  $\theta_m(p, M)$  satisfies

$$\theta_m(p, M) \geq \overline{\theta}_m(p, M) \geq c \cdot 2^{-\frac{1}{p}} \cdot m^{\frac{1}{2} - \frac{1}{p}}. \quad (5.4)$$

Here  $c > 0$  is a universal constant.

See Adcock et al. (2023c, Thms. 4.4 & 4.5). This theorem has several consequences. First, (5.3) shows that it is impossible to uniformly learn functions from the class  $\mathcal{H}(p)$ . In other words, mere  $\ell^p$ -summability of the anisotropy parameter  $\mathbf{b}$  is insufficient. By contrast, (5.4) shows that it may be possible to learn functions in  $\mathcal{H}(p, M)$ , where  $\mathbf{b}$  is now  $\ell_M^p$ -summable, but only at a rate that is algebraically decaying in  $m$ . We will see in the next section that these rates can in fact be attained (up to log terms) by practical methods. Finally, (5.2) asserts that knowledge of the anisotropy parameter  $\mathbf{b}$  does not help. The lower bound for the nonuniform width  $\overline{\theta}_m(p, M)$  is the same as the lower bound for the uniform width  $\theta_m(p, M)$ . Therefore, in terms of sample complexity, knowledge of  $\mathbf{b}$  conveys no benefit – we may as well deal exclusively with the unknown anisotropy setting.

It is worth relating this result back to Theorem 4.1. Theorem 5.2 shows that we can at best achieve a rate of  $m^{1/2-1/p}$  when learning functions in  $\mathcal{H}(\mathbf{b})$  from data, whenever  $\mathbf{b} \in \ell_M^p(\mathbb{N})$ . Theorem 4.1 (with  $p = 2$ ) shows that the best  $s$ -term polynomial approximation can achieve a rate of  $s^{1/2-1/p}$  subject to the weaker assumption  $\mathbf{b} \in \ell^p(\mathbb{N})$ . Hence, the discrepancy between  $\ell^p(\mathbb{N})$  (where learning is impossible) and  $\ell_M^p(\mathbb{N})$  (where learning is possible) must stem from the limited amount of data, not intrinsic properties of the function class itself.

**Remark 5.3** (Extensions). For simplicity, we have restricted this discussion to scalar-valued functions with  $\varrho$  being the uniform measure. The latter assumption can be easily relaxed to an arbitrary tensor-product probability measure. One may also consider Banach-valued functions, although some care is needed when defining an adaptive sampling operator in this case. See Adcock et al. (2023c) for details.

### 5.3 Towards methods

In summary, we have now seen that functions in  $\mathcal{H}(p, M)$  can be learned from training sets of size  $m$  with rates that are at best  $\mathcal{O}(m^{1/2-1/p})$ , regardless of the training data and learning method. In the remainder of this review, we focus on describing methods that achieve (close to) these rates when the training data consists of i.i.d. pointwise samples, as in (2.2). We consider, firstly, polynomial-based methods (§6) and, secondly, methods based on DNNs (§7-8).

## 6 Learning sparse polynomial approximations from data

We know from Theorem 4.1 that the best  $s$ -term approximation converges at the desired rate  $s^{1/2-1/p}$ . Therefore, our goal is to design methods that can compute best (or quasibest)  $s$ -term approximations using roughly  $s$  samples (up to constant and log factors). We shall do this using tools from compressed sensing (Foucart and Rauhut, 2013; Vidyasagar, 2019), which recast the problem of learning sparse polynomial approximations as recovering approximately sparse



vectors. Unfortunately, as we describe in this section, this task is easily laid out, but not so easily completed.

## 6.1 Setup

Throughout this section, we assume that  $\mathcal{V}$  is a Hilbert space (see §6.7 for some discussion on the Banach case). Moreover, in a practical scenario, we often cannot work directly with the space  $\mathcal{V}$ , since it may be infinite dimensional. Therefore, we now consider a finite-dimensional discretization  $\mathcal{V}_h \subseteq \mathcal{V}$ , where  $h > 0$  denotes a discretization parameter. For parametric PDEs, this is typically the mesh size when a Finite Element Method (FEM) is used. We also assume that the training data (2.2) belongs to  $\mathcal{U} \times \mathcal{V}_h$ , rather than  $\mathcal{U} \times \mathcal{V}$ , i.e.,

$$\{(\mathbf{y}_i, f(\mathbf{y}_i) + e_i)\}_{i=1}^m \subseteq \mathcal{U} \times \mathcal{V}_h. \quad (6.1)$$

This encapsulates the notion that  $f$  is evaluated using some numerical simulation (e.g., a FEM) that outputs values in  $\mathcal{V}_h$ . Notice that the error term  $e_i$  incorporates any error incurred by this computation. For convenience, we also define the orthogonal projection

$$\mathcal{P}_h : \mathcal{V} \rightarrow \mathcal{V}_h.$$

Moreover, if  $f \in L^2_\rho(\mathcal{U}; \mathcal{V})$ , then we let  $\mathcal{P}_h f \in L^2_\rho(\mathcal{U}; \mathcal{V}_h)$  be the function defined almost everywhere as  $(\mathcal{P}_h f)(\mathbf{y}) = \mathcal{P}_h(f(\mathbf{y}))$ ,  $\mathbf{y} \in \mathcal{U}$ .

Now consider a function  $f \in L^2_\rho(\mathcal{U}; \mathcal{V})$  with expansion (4.7) and coefficients (4.8). Standard compressed sensing involves the recovery of finite vectors. Therefore, we first need to truncate the expansion (4.7). Let  $\Lambda \subset \mathcal{F}$  be finite with  $|\Lambda| = N$  and an enumeration  $\Lambda = \{\mathbf{v}_1, \dots, \mathbf{v}_N\}$ . Then

$$f(\mathbf{y}_i) + e_i = f_\Lambda(\mathbf{y}_i) + (f - f_\Lambda)(\mathbf{y}_i) + e_i = \sum_{\mathbf{v} \in \Lambda} c_{\mathbf{v}} \Psi_{\mathbf{v}}(\mathbf{y}_i) + (f - f_\Lambda)(\mathbf{y}_i) + e_i,$$

where  $f_\Lambda = \sum_{\mathbf{v} \in \Lambda} c_{\mathbf{v}} \Psi_{\mathbf{v}}$ . Now let  $\mathbf{c}_\Lambda = (c_{\mathbf{v}_i})_{i=1}^N \in \mathcal{V}^N$ . Then we have

$$\mathbf{A} \mathbf{c}_\Lambda + \mathbf{e} + \mathbf{e}' = \mathbf{f}, \quad (6.2)$$

where

$$\mathbf{f} = \frac{1}{\sqrt{m}} (f(\mathbf{y}_i) + e_i)_{i=1}^m, \quad \mathbf{A} = \frac{1}{\sqrt{m}} (\Psi_{\mathbf{v}_j}(\mathbf{y}_i) / \sqrt{m})_{i,j=1}^{m,N} \in \mathbb{R}^{m \times N}, \quad (6.3)$$

and

$$\mathbf{e} = \frac{1}{\sqrt{m}} (e_i)_{i=1}^m, \quad \mathbf{e}' = \frac{1}{\sqrt{m}} ((f - f_\Lambda)(\mathbf{y}_i))_{i=1}^m \in \mathbb{R}^m.$$

The vector  $\mathbf{c}_\Lambda$  is approximately  $s$ -sparse. Therefore, we have recast the problem into a standard compressed sensing form: namely, the recovery of an approximately sparse vector from noisy linear measurements (6.2).

Of course, this problem is not quite standard, since the vector  $\mathbf{c}_\Lambda$  has entries taking values in  $\mathcal{V}$ , rather than  $\mathbb{R}$  or  $\mathbb{C}$ . Fortunately, this challenge can be dealt with by ‘lifting’ the appropriate theoretical tools from  $\mathbb{R}$  or  $\mathbb{C}$  to Hilbert spaces (see Remark 6.1 below). A more delicate challenge is the following.

**Challenge 1.** How should the truncation set  $\Lambda$  be chosen.

In general,  $f$  has an infinite expansion. Since we recover only those coefficients with indices in  $\Lambda$  (i.e., the vector  $\mathbf{c}_\Lambda$ ), the error in recovering  $f$  will always involve the expansion tail  $f - f_\Lambda$ . On the one hand,  $\Lambda$  should not be too large, since, as we see later in (6.21),  $\log(N)$  will enter into the sample complexity bound. Moreover, larger  $N$  will also increase the computational cost. On the other hand, we need  $f - f_\Lambda$  to be sufficiently small so as to obtain optimal learning rates. To do this, we need to ensure no large coefficients of  $f$  lie outside  $\Lambda$ . And here is where the problem lies. The result on best  $s$ -term approximation, Theorem 4.1, which has served as our rationale up to now for using compressed sensing, asserts approximate sparsity of the coefficients, but gives no guarantees whatsoever on where the largest  $s$  coefficients should lie within the set  $\mathcal{F}$ .

## 6.2 Sampling discretizations for multivariate polynomials

We will return to Challenge 1 in a moment. But, first, we also need to introduce a second challenge. Let us imagine some oracle gives us a ‘good’ set  $S \subset \mathcal{F}$  of size  $|S| = s$ . For example, this could even be a set  $S = S^*$  corresponding to the best  $s$ -term approximation (4.10). Then one would naturally learn an approximation to  $f$  via the empirical least-squares fit

$$\hat{f} \in \operatorname{argmin}_{p \in \mathbb{P}_{S; \mathcal{V}}} \frac{1}{m} \sum_{i=1}^m \|f(\mathbf{y}_i) + e_i - p(\mathbf{y}_i)\|_{\mathcal{V}}^2, \quad (6.4)$$

where  $\mathbb{P}_{S; \mathcal{V}}$  is the  $s$ -dimensional subspace

$$\mathbb{P}_{S; \mathcal{V}} = \left\{ \sum_{\mathbf{v} \in S} c_{\mathbf{v}} \Psi_{\mathbf{v}} : c_{\mathbf{v}} \in \mathcal{V} \right\} \subset L_{\mathcal{Q}}^2(\mathcal{U}; \mathcal{V})$$

of Hilbert-valued polynomials with nonzero coefficients in  $S$ .

The behavior of the estimator (6.4) is intimately related to the existence of a *sampling discretization* of the  $L_{\rho}^2(\mathcal{U})$ -norm for the scalar-valued analogue of this space, i.e.,  $\mathbb{P}_S = \mathbb{P}_{S; \mathbb{R}} = \operatorname{span}\{\Psi_{\mathbf{v}} : \mathbf{v} \in S\}$ . A sampling discretization (Kashin et al., 2022) (also known as a *Marcinkiewicz–Zygmund inequality* (Temlyakov, 2018)) is an inequality of the form

$$\alpha \|p\|_{L_{\mathcal{Q}}^2(\mathcal{U})}^2 \leq \frac{1}{m} \sum_{i=1}^m |p(\mathbf{y}_i)|^2 \leq \beta \|p\|_{L_{\rho}^2(\mathcal{U})}^2, \quad \forall p \in \mathbb{P}_S, \quad (6.5)$$

for constants  $0 < \alpha \leq \beta < \infty$ . It is known that (6.5) ensures both accuracy of the estimator (6.4) in the scalar-valued case and robustness to measurement error (see, e.g., Adcock et al., 2022, Thm. 5.3).

**Remark 6.1** (Lifting to Hilbert spaces). This is an instance of the aforementioned ‘lifting’ concept. The sampling discretization (6.5) is formulated for the space  $\mathbb{P}_S = \mathbb{P}_{S;\mathbb{R}}$ . It is a short argument (see, e.g., Adcock et al., 2024a, Lem. 7.5) to show that it is, in fact, equivalent to a sampling discretization for the space  $\mathbb{P}_{S;\mathcal{V}}$ , i.e.,

$$\alpha \|P\|_{L^2_{\varrho}(\mathcal{U};\mathcal{V})}^2 \leq \frac{1}{m} \sum_{i=1}^m \|P(\mathbf{y}_i)\|_{\mathcal{V}}^2 \leq \beta \|P\|_{L^2_{\varrho}(\mathcal{U};\mathcal{V})}^2, \quad \forall P \in \mathbb{P}_{S;\mathcal{V}}.$$

In turn, this implies accuracy and robustness of the estimator (6.4) in the Hilbert-valued case.

Sufficient conditions for sampling discretizations in linear subspaces with i.i.d. samples can be derived using standard matrix concentration inequalities. See, e.g., Cohen et al. (2013) or Adcock et al. (2022, Chpt. 5). These conditions involve the quantity

$$\kappa(\mathbb{P}_S) = \|\mathcal{K}(\mathbb{P}_S)\|_{L^{\infty}(\mathcal{U})}, \quad (6.6)$$

where  $\mathcal{K}(\mathbb{P}_S)$  is the (reciprocal) *Christoffel function* of the subspace  $\mathbb{P}_S$ :

$$\mathcal{K}(\mathbb{P}_S)(\mathbf{y}) = \sum_{\mathbf{v} \in S} |\Psi_{\mathbf{v}}(\mathbf{y})|^2, \quad \forall \mathbf{y} \in \mathcal{U}. \quad (6.7)$$

Specifically, one can show that if

$$m \geq c \cdot \kappa(\mathbb{P}_S) \cdot \log(2s/\epsilon), \quad (6.8)$$

where  $s = |S|$  and  $c > 0$  is a universal constant, then (6.5) holds with constants  $\beta \leq 2$  and  $\alpha \geq 1/2$  (these values are arbitrary) with probability at least  $1 - \epsilon$  on the draw of the sample points  $\mathbf{y}_1, \dots, \mathbf{y}_m \sim_{\text{i.i.d.}} \varrho$  (see, e.g., Adcock et al., 2022, Thm. 5.12).

Therefore, the *sample complexity* of the ‘oracle’ estimator (6.4) is governed by the maximal behavior of the Christoffel function (6.6). Note that  $\kappa(\mathbb{P}_S) \geq s$  for any set  $S$  with  $|S| = s$  (Adcock et al., 2022, §5.3). To use (6.8) to achieve the optimal rates described in Theorem 5.2 for this oracle estimator, we require that  $\kappa(\mathbb{P}_S) \lesssim s$  as well. Unfortunately, this is not the case:  $\kappa(\mathbb{P}_S)$  can be arbitrarily large in comparison to  $s$ . To see why, we recall (4.1). This implies that

$$\kappa(\mathbb{P}_S) = \mathcal{K}(P_S)(\mathbf{1}) = \sum_{\mathbf{v} \in S} u_{\mathbf{v}}^2, \quad (6.9)$$

where the  $u_{\mathbf{v}}$  are as in (4.5). Since  $u_{\mathbf{v}} \rightarrow \infty$  as  $\mathbf{v} \rightarrow \infty$ , we deduce the claim.

**Challenge 2.** Even if it is known, the index set  $S$  of the best  $s$ -term approximation may lead to a sample complexity estimate (6.8) that is arbitrarily large.

Similar to Challenge 1, this difficulty arises because Theorem 4.1 gives no guarantees about the set  $S = S^*$  which attains the best  $s$ -term approximation. In particular, it says nothing about the term  $\kappa(\mathbb{P}_{S^*})$ . The paths to resolving each challenge are therefore similar. We need to show that *near-best*  $s$ -term approximations can be obtained using suitably *structured* index sets.

**Remark 6.2** (Necessity of sampling discretizations). The sampling discretization (6.5) is sufficient condition for robustness of the estimator  $\hat{f}$ . However, the lower inequality is also essentially necessary. Indeed, let  $\mathcal{S} : C(\mathcal{U}) \rightarrow \mathbb{R}^m$ ,  $f \mapsto (f(y_i))_{i=1}^m / \sqrt{m}$  and  $\mathcal{R} : \mathbb{R}^m \rightarrow C(\mathcal{U})$  be any reconstruction map that is  $\delta$ -accurate over  $\mathbb{P}_S$ , i.e.,

$$\|p - \mathcal{R}(\mathcal{S}(p))\|_{L^2_{\rho}(\mathcal{U})} \leq \delta \|p\|_{L^2_{\rho}(\mathcal{U})}, \quad \forall p \in \mathbb{P}_S, \quad (6.10)$$

for some  $\delta > 0$ . Then it is a short argument to show that the  $\epsilon$ -Lipschitz constant

$$L_{\epsilon} = \sup_{f \in C(\mathcal{U})} \sup_{0 < \|e\|_2 \leq \epsilon} \|\mathcal{R}(\mathcal{S}(f) + e) - \mathcal{R}(\mathcal{S}(f))\|_{L^2_{\rho}(\mathcal{U})} / \|e\|_2$$

satisfies  $L_{\epsilon} \geq (1 - \delta) / \sqrt{\alpha}$ . Thus, when  $\alpha$  is small, a reconstruction map cannot be simultaneously accurate (in the sense of (6.10)) and stable.

### 6.3 Resolving Challenge 1: lower and anchored sets

In order to address Challenge 1, we now introduce the following concept.

**Definition 6.3** (Lower and anchored sets). A multiindex set  $S \subseteq \mathcal{F}$  is *lower* if the following holds for every  $\mathbf{v}, \boldsymbol{\mu} \in S$ :

$$(\mathbf{v} \in S \text{ and } \boldsymbol{\mu} \leq \mathbf{v}) \Rightarrow \boldsymbol{\mu} \in \Lambda.$$

A multiindex set  $S \subseteq \mathcal{F}$  is *anchored* if it is lower and if the following holds for every  $j \in \mathbb{N}$ :

$$\mathbf{e}_j \in S \Rightarrow \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_j\} \subseteq S.$$

Lower sets are classical objects in multivariate approximation theory. Anchored sets were introduced in the context of infinite-dimensional approximations. See, e.g., Adcock et al. (2022, §2.3.3 & 2.5.3) and references therein.

We now state a result that shows that the  $s^{1/q-1/p}$  rate asserted in Theorem 4.1 can also be obtained using a lower or (under an additional assumption) anchored set. See, e.g., Adcock et al. (2022, Thm. 3.33) or Cohen and DeVore (2015, §3.8).

**Theorem 6.4** (Algebraic convergence in lower or anchored sets). *Let  $\mathbf{b} \in [0, \infty)^\mathbb{N}$  be such that  $\mathbf{b} \in \ell^p(\mathbb{N})$  for some  $0 < p < 1$ . Then for any  $s \in \mathbb{N}$  and  $p \leq q \leq 2$ , there exists an lower set  $S \subset \mathcal{F}$  with  $|S| \leq s$  such that*

$$\|\mathbf{c} - \mathbf{c}_S\|_{q, \mathcal{V}} \leq C(\mathbf{b}, p) \cdot s^{\frac{1}{q} - \frac{1}{p}}$$

for all  $f \in \mathcal{H}(\mathbf{b})$  with coefficients  $\mathbf{c}$  as in (4.8). If  $\mathbf{b} \in \ell^p_{\mathcal{M}}(\mathbb{N})$  then  $S$  can also be chosen as an anchored set.

In particular, when  $q = 2$ , this theorem and Parseval's identity imply the existence of a lower or anchored set  $S$  with  $|S| \leq s$  such that

$$\|f - f_S\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})} \leq C(\mathbf{b}, p) \cdot s^{\frac{1}{2} - \frac{1}{p}}, \quad \forall f \in \mathcal{H}(\mathbf{b}). \quad (6.11)$$

Now recall that our aim is to learn functions in the space  $\mathcal{H}(p, \mathcal{M})$  defined by (3.8). Any  $f \in \mathcal{H}(p, \mathcal{M})$  satisfies  $f \in \mathcal{H}(\mathbf{b})$  for some  $\mathbf{b} \in \ell^p_{\mathcal{M}}(\mathbb{N})$ . Therefore, for any such function, we know there is an anchored set  $S$ ,  $|S| \leq s$ , that achieves (6.11). How does this allow us to overcome Challenge 1? The reason is because anchored sets of a fixed size lie within a *finite* subset of  $\mathcal{F}$ . Indeed, one can show (see, e.g., Cohen et al., 2017, or Adcock et al., 2022, Prop. 2.18) that  $S \subset \Lambda_s^{\text{HCl}}$  for all  $S \subset \mathcal{F}$  with  $|S| \leq s$ , where  $\Lambda_s^{\text{HCl}}$  is the finite set

$$\Lambda_s^{\text{HCl}} = \left\{ \mathbf{v} = (v_k)_{k=1}^{\infty} \in \mathcal{F} : \prod_{k=1}^{s-1} (v_k + 1) \leq s, v_k = 0, \forall k \geq s \right\}. \quad (6.12)$$

This set is in fact isomorphic to the  $(s - 1)$ -dimensional *hyperbolic cross* set of order  $s - 1$ , a very well-known object in high-dimensional approximation (Düng et al., 2018). With this knowledge, we set  $\Lambda = \Lambda_s^{\text{HCl}}$  and then apply Theorem 6.4:

$$\|f - f_{\Lambda}\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})} = \|\mathbf{c} - \mathbf{c}_{\Lambda}\|_{2, \mathcal{V}} \leq \|\mathbf{c} - \mathbf{c}_S\|_{2, \mathcal{V}} \leq C(\mathbf{b}, p) \cdot s^{\frac{1}{2} - \frac{1}{p}},$$

for all  $f \in \mathcal{H}(\mathbf{b})$  and  $\mathbf{b} \in \ell^p_{\mathcal{M}}(\mathbb{N})$ . This resolves Challenge 1.

**Remark 6.5.** Theorem 5.2 states that no method can learn functions in  $\mathcal{H}(p)$  from finite data. Theorem 6.4 says that we can always find a lower set that yields the desired rate (6.11) when  $\mathbf{b} \in \ell^p(\mathbb{N})$ . However, the union of all lower sets in infinite dimensions is not a finite set. This is precisely what prohibits one from learning functions in  $\mathcal{H}(p)$  using compressed sensing (which, of course, must not be possible in view of Theorem 5.2): namely, there is no way to construct a finite set  $\Lambda$  that ensures a uniformly small truncation error  $f - f_{\Lambda}$ .

## 6.4 Resolving Challenge 2: weighted $k$ -term approximation

We now return to Challenge 2. In the previous subsection, we identified anchored sets as structured sets which achieve near-best  $s$ -term approximation

rates. Unfortunately, while anchored sets do ameliorate the sample complexity issue, they do not fully resolve it. On the one hand, it is possible to show that  $\kappa(\mathbb{P}_S) \leq s^2$  whenever  $S$  is anchored (see, e.g., Adcock et al., 2022, Prop. 5.17), where  $\kappa(\mathbb{P}_S)$  is as in (6.6). Unfortunately, this bound is sharp. For example, the anchored set

$$S = \{j\mathbf{e}_1 : j = 0, \dots, s-1\}, \quad \text{where } \mathbf{e}_1 = (1, 0, 0, \dots),$$

satisfies  $\kappa(\mathbb{P}_S) = s^2$ . This follows from (6.9) and the fact that  $u_{j\mathbf{e}_1} = u_j = \sqrt{2j+1}$ . Therefore, the sample complexity bound (6.8) for learning with anchored sets is generically log-quadratic in  $s$ .

To overcome Challenge 2 we need a different concept, *weighted  $k$ -term approximation*. Motivated by (6.9), we now define the *weighted cardinality* of set  $S \subset \mathcal{F}$  with respect to *weights*  $\mathbf{u}$  as

$$|S|_{\mathbf{u}} := \sum_{\mathbf{v} \in S} u_{\mathbf{v}}^2.$$

Thus we may reinterpret (6.8) as follows: the sample complexity is determined not by the cardinality of  $S$ , but by its weighted cardinality.

Fortunately, as we next show, there exist sets of a given weighted cardinality that achieve near-optimal approximation rates. For this, we need some additional notation. Given  $1 \leq p \leq 2$ , we define the weighted  $\ell_{\mathbf{u}}^p(\Lambda; \mathcal{V})$  space as the space of  $\mathcal{V}$ -valued sequences  $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$  for which  $\|\mathbf{c}\|_{p, \mathbf{u}; \mathcal{V}} < \infty$ , where

$$\|\mathbf{c}\|_{p, \mathbf{u}; \mathcal{V}} = \left( \sum_{\mathbf{v} \in \Lambda} u_{\mathbf{v}}^{2-p} \|c_{\mathbf{v}}\|_{\mathcal{V}}^p \right)^{\frac{1}{p}}.$$

Notice that  $\|\cdot\|_{2, \mathbf{u}; \mathcal{V}}$  coincides with the unweighted norm  $\|\cdot\|_{2; \mathcal{V}}$ .

**Theorem 6.6** (Algebraic convergence of the weighted best  $k$ -term approximation). *Let  $\mathbf{b} \in [0, \infty)^{\mathbb{N}}$  be such that  $\mathbf{b} \in \ell^p(\mathbb{N})$  for some  $0 < p < 1$ . Then for any  $k > 0$  and  $p \leq q \leq 2$  there exists a set  $S \subset \mathcal{F}$  with  $|S|_{\mathbf{u}} \leq k$  such that*

$$\|\mathbf{c} - \mathbf{c}_S\|_{q, \mathbf{u}; \mathcal{V}} \leq C(\mathbf{b}, p) \cdot k^{\frac{1}{q} - \frac{1}{p}}$$

for all  $f \in \mathcal{H}(\mathbf{b})$  with coefficients  $\mathbf{c}$  as in (4.8).

This theorem shows that we can construct near-best approximations using set of weighted cardinality at most  $k$ . In particular, when  $q = 2$ , this theorem and Parseval's identity give that there exists a set  $S$  with  $|S|_{\mathbf{u}} \leq k$  such that

$$\|f - f_S\|_{L^2_{\mathbf{u}}(\mathcal{V})} \leq C(\mathbf{b}, p) \cdot k^{\frac{1}{2} - \frac{1}{p}}, \quad \forall f \in \mathcal{H}(\mathbf{b}).$$

Note that we do not require  $\mathbf{b} \in \ell^p_{\mathbb{M}}(\mathbb{N})$  for this result. This additional regularity is only needed to resolve Challenge 1.

## 6.5 Weighted $\ell^1$ -minimization

Challenge 2 is now resolved, at least insofar as the ‘oracle’ least-squares estimator goes. Of course, in practice we do not have access to such an oracle. However, Theorem 6.6 implies that the sequence  $\mathbf{c}$ , and therefore the finite vector  $\mathbf{c}_\Lambda$ , is approximately *weighted sparse*. In other words, it is well approximated by  $\mathbf{c}_S$  for some index set  $S \subseteq \Lambda$  of weighted cardinality  $|S|_u \leq k$ .

Weighted sparsity is a well-understood concept in compressed sensing (see Adcock, 2017; Rauhut and Ward, 2016 and references therein). Much like how sparsity can be promoted by solving a minimization problem involving the  $\ell^1$ -norm, weighted sparsity can be solving a minimization problem involving the weighted  $\ell^1_u$ -norm. A large weight penalizes the corresponding coefficient, much like how a large weight increases the weighted sparsity of any vector that has a nonzero coefficient at the corresponding index. There are a number of ways to formulate the weighted  $\ell^1$ -minimization problem, but, following Adcock et al. (2019, 2024a), we will consider the Hilbert-valued, *weighted square-root LASSO* program

$$\min_{z \in \mathcal{V}_h^N} \mathcal{G}(z), \quad \text{where } \mathcal{G}(z) := \lambda \|z\|_{1,u;\mathcal{V}} + \|\mathbf{A}z - \mathbf{f}\|_{2;\mathcal{V}}. \quad (6.13)$$

Here  $\|z\|_{1,u;\mathcal{V}} = \sum_{v \in \Lambda} u_v \|z_v\|_{\mathcal{V}}$  is the  $\ell^1_{u;\mathcal{V}}$ -norm and  $\lambda > 0$  is a parameter. Notice that this problem is posed over  $\mathcal{V}_h^N$ , which means it can be numerically solved. See Remark 6.7 below. In order to account for inexact solution of (6.13), given  $\gamma \geq 0$  we say that  $\hat{\mathbf{c}} = (\hat{c}_v)_{v \in \Lambda}$  is a  $\gamma$ -*minimizer* of (6.13) if

$$\mathcal{G}(\hat{\mathbf{c}}) \leq \min_{z \in \mathcal{V}_h^N} \mathcal{G}(z) + \gamma.$$

For such  $\hat{\mathbf{c}}$ , we define the corresponding sparse polynomial approximation

$$\hat{f} = \sum_{v \in \Lambda} \hat{c}_v \Psi_v. \quad (6.14)$$

**Remark 6.7** (Numerical solution of (6.13)). In practice, (6.13) is solved by first introducing a basis  $\{\varphi_i\}_{i=1}^K$  for the space  $\mathcal{V}_h$ , where  $K = \dim(\mathcal{V}_h)$  (in particular,  $K = 1$  in the scalar-valued case  $\mathcal{V} = (\mathbb{R}, |\cdot|)$ ). Instead of searching for a  $\mathcal{V}_h$ -valued vector of coefficients  $\hat{\mathbf{c}} \in \mathcal{V}_h^N$ , one now searches for an equivalent matrix of coefficients  $\hat{\mathbf{C}} \in \mathbb{R}^{N \times K}$ . It is a short exercise to show that (6.13) is equivalent to the problem

$$\min_{\mathbf{Z} \in \mathbb{R}^{N \times K}} \lambda \|\mathbf{Z}\|_{2,1,u} + \|(\mathbf{A}\mathbf{Z} - \mathbf{F})\mathbf{G}^{1/2}\|_F, \quad (6.15)$$

where  $\|\cdot\|_F$  is the Frobenius norm,  $\mathbf{G} = ((\varphi_i, \varphi_j))_{i,j=1}^K \in \mathbb{R}^{K \times K}$  and

$$\|\mathbf{Z}\|_{2,1,\mathbf{u}} = \sum_{i=1}^N u_i \sqrt{\sum_{j=1}^K |z_{ij}|^2}.$$

As discussed in Adcock et al. (2024a), the convex optimization problem (6.15) can be solved efficiently using Chambolle & Pock's primal-dual iteration (Chambolle and Pock, 2011, 2016) in combination with a restart scheme (Adcock et al., 2023b; Roulet and d'Aspremont, 2020). Specifically, one can obtain a  $\gamma$ -minimizer in at most  $\log(1/\gamma)$  iterations, where the cost-per-iteration is bounded by

$$c \cdot (m \cdot N \cdot K + (m + N) \cdot (F(\mathbf{G}) + K)). \quad (6.16)$$

Here  $F(\mathbf{G}) \leq K^2$  is the cost of performing the matrix-vector multiplication  $\mathbf{x} \mapsto \mathbf{G}\mathbf{x}$ . See Adcock et al. (2024a, Thm. 3.9 & Lem. 4.3).

## 6.6 Theoretical guarantee

We are now ready to present a theoretical guarantee for this estimator. For convenience, we define

$$L = L(m, \epsilon) = \log(m) \cdot (\log^3(m) + \log(\epsilon^{-1})). \quad (6.17)$$

**Theorem 6.8** (Near-optimal learning via polynomials). *Let  $m \geq 3$ ,  $0 < \epsilon < 1$  and  $n = \lceil m/L \rceil$ , where  $L = L(m, \epsilon)$  is as in (6.17). Let  $0 < p < 1$ ,  $\mathbf{b} \in \ell_M^p(\mathbb{N})$ ,  $f \in \mathcal{H}(\mathbf{b})$ ,  $\mathbf{y}_1, \dots, \mathbf{y}_m \sim_{\text{i.i.d.}} \mathcal{Q}$  and consider the training data (6.1). Then with probability at least  $1 - \epsilon$  the approximation  $\hat{f}$  defined by (6.14) for any  $\gamma$ -minimizer  $\hat{\mathbf{c}}$  of (6.13),  $\gamma \geq 0$ , with  $\Lambda = \Lambda_n^{\text{HCl}}$  satisfies*

$$\|f - \hat{f}\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})} \leq c \cdot \zeta, \quad (6.18)$$

where  $c \geq 1$  is a universal constant and

$$\zeta := C \cdot (m/L)^{1/2-1/p} + \|\mathbf{e}\|_{2;\mathcal{V}}/\sqrt{m} + \|f - \mathcal{P}_h(f)\|_{L^\infty(\mathcal{U}; \mathcal{V})} + \gamma. \quad (6.19)$$

See Adcock et al. (2024a, Thm. 3.7). The above theorem is slightly more general than Adcock et al. (2024a, Thm. 3.7) since it does not require  $\mathbf{b}$  to be monotonically decreasing. This generalization follows from techniques given in the proof of Theorem 7.2 in Adcock and Monte (2023).

This result demonstrates that functions in  $\mathcal{H}(p, M)$  can be learned from i.i.d. samples with a rate  $O((m/\log^4(m))^{1/2-1/p})$  which, in view of Theorem 5.2, is near-optimal. Moreover, the estimator is robust to other sources of error in the problem. Indeed, the term (6.19) is a linear combination of the following.



- (a) An *approximation error*  $C \cdot (m/L)^{1/2-1p}$ , as discussed.
- (b) A *measurement error*  $\|e\|_{2;\mathcal{V}}/\sqrt{m}$ , which accounts for any errors in computing the sample values  $f(\mathbf{y}_i)$ .
- (c) A *physical discretization error*  $\|f - \mathcal{P}_h(f)\|_{L^\infty(\mathcal{U};\mathcal{V})}$ , which accounts for the error induced when working over the finite-dimensional space  $\mathcal{V}_h$  instead of  $\mathcal{V}$  and depends on the orthogonal projection  $\mathcal{P}_h(f)$ .
- (d) An *optimization error*  $\gamma$ , which depends on the optimality gap of the computed solution  $\hat{\mathbf{c}}$ . See Remark 6.7 for further discussion on this term.

The proof of this theorem relies heavily on tools from compressed sensing. To relate it back to the discussion in §6.2, we remark that a key step involves establishing a sampling discretization of the form

$$\alpha \|p\|_{L^2_\varrho(\mathcal{U})}^2 \leq \frac{1}{m} \sum_{i=1}^m |p(\mathbf{y}_i)|^2 \leq \beta \|p\|_{L^2_\varrho(\mathcal{U})}^2, \quad \forall p \in \mathbb{P}_S, \quad S \subseteq \Lambda, \quad |S|_u \leq k. \quad (6.20)$$

This is stronger than (6.5), since it is required to hold simultaneously for all sets  $S \subseteq \Lambda$  of weighted cardinality at most  $k$ . It is an example of a *universal sampling discretization* (Dai and Temlyakov, 2023). The proof first shows that it is equivalent to a certain *weighted Restricted Isometry Property* (wRIP), then uses known results for the wRIP (see Brugiapaglia et al., 2021, as well as Rauhut and Ward, 2016). As proved therein, a sufficient condition for (6.20) to hold with probability at least  $1 - \epsilon$  and constants  $\alpha \geq 1 - \delta$  and  $\beta \leq 1 + \delta$  for some  $0 < \delta < 1$  is

$$m \geq c \cdot \delta^{-2} \cdot k \cdot (\log(eN) \cdot \log^2(k/\delta) + \log(2/\epsilon)). \quad (6.21)$$

Crucially, up to the log terms, there is now a linear scaling between  $m$  and  $k$ . This is what leads to the near-optimal approximation error term.

**Remark 6.9** (Sparse polynomial approximations). Generally, the vector  $\hat{\mathbf{c}}$  computed by solving (6.13) will not be sparse. Therefore,  $\hat{f}$  will not be a *sparse polynomial* approximation per se. Fortunately, one can always postprocess the solution to obtain such an approximation. This was shown in Adcock and Monte (2023, Thm. 7.2) in the scalar-valued case  $\mathcal{V} = (\mathbb{R}, |\cdot|)$ , but the technique extends straightforwardly to the Hilbert-valued case. In the case of Theorem 6.8, given a  $\gamma$ -minimizer  $\hat{\mathbf{c}} = (\hat{c}_\nu)_{\nu \in \Lambda} \in \mathcal{V}_h^N$  one first computes the index set  $S \subseteq \Lambda$  of the largest  $n = \lceil m/L \rceil$  entries of the vector  $(\|\hat{c}_\nu\|_{\mathcal{V}})_{\nu \in \Lambda}$ , and then replaces  $\hat{f}$  with the  $n$ -sparse polynomial approximation  $\check{f} = \sum_{\nu \in S} \hat{c}_\nu \Psi_\nu$ . As shown in Adcock and Monte (2023, Thm. 7.2),  $\check{f}$  attains the same error bounds up to possible changes in the constants. The computational cost of this postprocessing step, roughly  $\mathcal{O}(N \log(N) + N \cdot F(\mathbf{G}))$  operations, is generally negligible in comparison to the cost of computing the initial  $\gamma$ -minimizer  $\hat{\mathbf{c}}$ .

## 6.7 Discussion and extensions

We conclude this section a short discussion. First, Theorem 6.8 asserts that i.i.d. pointwise samples constitute *near-optimal information*. This is an interesting facet of infinite-dimensional holomorphic function approximation, which contrasts starkly with the finite-, and specifically, low-dimensional case, where samples drawn i.i.d. from the uniform measure are distinctly suboptimal. See Adcock and Monte (2023) for a detailed analysis of this phenomenon.

On the other hand, Theorem 6.8 assumes that  $\mathcal{V}$  is a Hilbert space, whereas the lower bounds in Theorem 5.2 allow for Banach spaces. Theorem 6.8 can be extended to Banach spaces, but with the suboptimal rate  $(m/L)^{1/2(1/2-1/p)}$  (Adcock et al., 2023a, Thm. 4.1). It is currently unknown whether this rate can be improved.

Moreover, the computational cost of this procedure can be prohibitive. As discussed in Adcock et al. (2024a), combining (6.16) with a standard estimate for  $N = |\Lambda_n^{\text{HCl}}|$  leads to a per-iteration costs that scales like  $m^{3+\log_2(m)}$  – i.e., subexponential, but superpolynomial in  $m$ . The reason for this is the need to construct the matrix  $\mathbf{A}$  corresponding to all polynomials in the large index set  $\Lambda_n^{\text{HCl}}$ . Whether functions in  $\mathcal{H}(p, M)$  can be learned to algebraic accuracy in  $m$  with polynomial computational cost in  $m$  is currently an open problem. Sublinear time algorithms (Choi et al., 2021a,b) may yield a solution to this problem.

## 7 DNN existence theory

We now turn our attention to DNNs. In this section, we discuss *existence theory*. Theorems of this ilk describe the *expressivity* of NNs: namely, their ability to approximate functions from specific classes to a desired accuracy. We commence with a review, before showing how to establish an existence theorem for  $(\mathbf{b}, \varepsilon)$ -holomorphic functions via the technique of polynomial emulation.

### 7.1 Review

Arguably, the first results on existence theory are the various *universal approximation theorems* for NNs. See Cybenko (1989); Hornik et al. (1989) and, in particular, Pinkus (1999). These show that shallow NNs with one hidden layer can approximate any continuous function.

Unfortunately, these classical results only apply to shallow networks, and do not always give quantitative bounds on the complexity (i.e., the width) of the corresponding networks (some notable exceptions to this include Mhaskar, 1996, and references therein). These issues have been investigated intensively over the last five years. Some of the first results were obtained in Yarotsky (2017). Here, explicit width and depth bounds were derived for DNNs with the *Rectified Linear Unit (ReLU)* activation function for approximating Lipschitz continuous functions in the  $L^\infty$ -norm over compact sets. This inspired many subsequent works. Other activations have been studied, including *hyperbolic*

*tangents (tanh)*, *Rectified Quadratic Units (ReQU)* and *Rectified Polynomial Units (RePU)*, and various others. And quantitative bounds have been shown for various different function classes. A partial list includes: functions in Sobolev spaces with ReLU (Yarotsky, 2017; Gühring et al., 2020), RePU (Li et al., 2020), tanh (De Ryck et al., 2021) or rational (Boullé et al., 2020) activations; piecewise smooth functions with ReLU (Petersen and Voigtlaender, 2018);  $C^k$  and Hölder smooth functions with ReLU (J. Lu et al., 2021; Schmidt-Hieber, 2020) or general (Ohn and Kim, 2019) activations; uniformly continuous functions with ReLU activations (Yarotsky, 2018); functions in Besov spaces with ReLU (Opschoor et al., 2020; Suzuki, 2019); spaces of mixed smoothness with ReLU (Blanchard and Bennouna, 2020; Montanelli and Du, 2019; Dǔng and Nguyen, 2021; Suzuki, 2019), RePU (Li et al., 2020) or smooth activations (Blanchard and Bennouna, 2020); Gevrey functions with ReLU (Opschoor et al., 2020), RePU (Opschoor et al., 2022) or tanh (De Ryck et al., 2021) activations; bandlimited functions with ReLU activations (Montanelli et al., 2021); functions in Barron spaces (E et al., 2021) with ReLU and non-ReLU activations; compositional functions (Poggio et al., 2017; Liang and Srikant, 2017; Schmidt-Hieber, 2020); smooth functions on manifolds (Chen et al., 2022; Shaham et al., 2018); finite-dimensional, analytic functions with smooth (Mhaskar, 1996), ReLU (E and Wang, 2018) or RePU (Opschoor et al., 2022) activation functions; and, most relevantly to this work, infinite-dimensional holomorphic functions with ReLU (Schwab and Zech, 2019, 2023; Dǔng et al., 2023), tanh (De Ryck et al., 2021) or RePU (Schwab and Zech, 2023) activations. See also DeVore et al. (2021); Elbrächter et al. (2021) for reviews.

These existence theorems are based on emulating suitable classical approximation schemes with DNNs. Emulation results include (localized) Taylor polynomials (De Ryck et al., 2021; Gühring and Raslan, 2021; Liang and Srikant, 2017; Li et al., 2020; J. Lu et al., 2021; Yarotsky, 2017; Schwab and Zech, 2019), orthogonal polynomials (Adcock et al., 2023a; Daws and C. Webster, 2019; Mhaskar, 1996; Opschoor et al., 2022; Tang et al., 2019; Dǔng et al., 2023; Opschoor and Schwab, 2023; Schwab and Zech, 2023), rational functions (Boullé et al., 2020; Telgarsky, 2017), wavelets (Shaham et al., 2018) and general affine systems (Bölcskei et al., 2019), B-splines (Mhaskar, 1993), free-knot splines (Opschoor et al., 2020), finite elements (Longo et al., 2023; Opschoor et al., 2020; Opschoor and Schwab, 2023) and sparse grids (Blanchard and Bennouna, 2020; Montanelli and Du, 2019; Dǔng and Nguyen, 2021; Suzuki, 2019).

To highlight some of these themes, we next describe an existence theorem for  $(\mathbf{b}, \varepsilon)$ -holomorphic functions. This is achieved by first emulating Legendre polynomials with DNNs. For ease of presentation, we consider tanh DNNs. But, as we discuss, other activation functions can also be used.

## 7.2 Neural network architectures

In this and subsequent sections, we consider standard feedforward DNN architectures of the form

$$\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^k, \mathbf{z} \mapsto \Phi(\mathbf{z}) = \mathcal{A}_{D+1}(\sigma(\mathcal{A}_D(\sigma(\cdots \sigma(\mathcal{A}_0(\mathbf{z})) \cdots))), \quad (7.1)$$

where  $\mathcal{A}_l : \mathbb{R}^{N_l} \rightarrow \mathbb{R}^{N_{l+1}}$ ,  $l = 0, \dots, D + 1$  are affine maps and  $\sigma$  is the activation function, which we assume acts componentwise. The values  $\{N_l\}_{l=1}^{D+1}$  are the widths of the hidden layers, and for convenience, we write  $N_0 = n$  and  $N_{D+2} = k$ . Given such a DNN  $\Phi$ , we write

$$\text{width}(\Phi) = \max\{N_1, \dots, N_{D+1}\}, \quad \text{depth}(\Phi) = D.$$

We denote a class of DNNs of the form (7.1) with a fixed architecture (i.e., fixed activation function, depth and widths) as  $\mathcal{N}$ , and define

$$\text{width}(\mathcal{N}) = \max\{N_1, \dots, N_{D+1}\}, \quad \text{depth}(\mathcal{N}) = D.$$

Finally, since the DNNs (7.1) take a finite input, yet the functions considered in this work take inputs in  $\mathcal{U} \subseteq \mathbb{R}^{\mathbb{N}}$ , we also need to introduce a restriction operator. Let  $\Theta \subset \mathbb{N}$  with  $|\Theta| = n$ . Then we define the *variable restriction operator* as

$$\mathcal{T}_\Theta : \mathbb{R}^{\mathbb{N}} \rightarrow \mathbb{R}^n, \mathbf{y} = (y_i)_{i \in \mathbb{N}} \rightarrow (y_i)_{i \in \Theta}. \quad (7.2)$$

If  $\Theta = \{1, \dots, n\}$  then we simply write  $\mathcal{T}_\Theta = \mathcal{T}_n$ .

## 7.3 Emulating polynomials with DNNs: typical result

We now present a result on emulating a finite set of multivariate Legendre polynomials with a tanh DNN.

**Theorem 7.1** (Emulating multivariate Legendre polynomials). *Let  $\Lambda \subset \mathcal{F}$  be a finite multiindex set and  $\Theta \subset \mathbb{N}$ ,  $|\Theta| = n$ , be such that  $\text{supp}(\mathbf{v}) \subseteq \Theta$ ,  $\forall \mathbf{v} \in \Lambda$ . Then for every  $0 < \delta < 1$  there exists a tanh DNN  $\Phi_{\Lambda, \delta} : \mathbb{R}^n \rightarrow \mathbb{R}^{|\Lambda|}$ , such that, if  $\Phi_{\Lambda, \delta}(\mathbf{z}) = (\Phi_{\mathbf{v}, \delta}(\mathbf{z}))_{\mathbf{v} \in \Lambda}$ ,  $\mathbf{z} = (z_j)_{j \in \Theta} \in \mathbb{R}^n$  and  $\mathcal{T}_\Theta$  is as in (7.2), then*

$$\|\Psi_{\mathbf{v}} - \Phi_{\mathbf{v}, \delta} \circ \mathcal{T}_\Theta\|_{L^\infty(\mathcal{U})} \leq \delta, \quad \forall \mathbf{v} \in \Lambda,$$

where  $\Psi_{\mathbf{v}}$  is the corresponding multivariate Legendre polynomial. The width and depth of this network satisfy

$$\text{width}(\Phi_{\Lambda, \delta}) \leq c_1 \cdot |\Lambda| \cdot m(\Lambda), \quad \text{depth}(\Phi_{\Lambda, \delta}) \leq c_2 \cdot \log_2(m(\Lambda)),$$

for universal constants  $c_1, c_2 > 0$ , where  $m(\Lambda) = \max_{\mathbf{v} \in \Lambda} \|\mathbf{v}\|_1$ .

This result was shown in Adcock et al. (2023a, Thm. 7.4), and is based on techniques of Daws and C. Webster (2019); De Ryck et al. (2021); Opschoor et al. (2022). To avoid unnecessary complications, we have stated it for tanh DNNs. It also applies without changes to DNNs with the sigmoid activation, since this is obtained from tanh via shifting and scaling. The result (Adcock et al., 2023a, Thm. 7.4) also considers the ReLU and RePU activations. As discussed therein, the proof readily extends to more general classes of activation functions. See also De Ryck et al. (2021, Rem. 3.9) and Opschoor et al. (2022, §2.4) for related discussion.

## 7.4 Elements of the proof of Theorem 7.1

The proof of Theorem 7.1 involves three key ingredients:

- (i) The map  $(x, y) \mapsto xy$  can be approximated by a shallow tanh NN.
- (ii) Using (i), the multiplication of  $d$  numbers  $(x_1, \dots, x_d) \mapsto x_1 \cdots x_d$  can be approximated by a tanh DNN of depth  $c \lceil \log_2(d) \rceil$  for some constant  $c > 0$ .
- (iii) By the fundamental theorem of algebra, the multivariate Legendre polynomial  $\Psi_{\mathbf{v}}$  can be expressed as a product of  $\|\mathbf{v}\|_1$  terms.

The constructions that lead to the proofs of (i) and (ii) have become standard. In the case of tanh DNNs, they can be found in De Ryck et al. (2021, Lem. 3.8). For (i), the basic idea is to use the relation  $xy = ((x + y)^2 - (x - y)^2)/4$  in combination with a tanh DNN for approximating the map  $x \mapsto x^2$ . The latter is constructed via finite differences (see De Ryck et al., 2021, Lem. 3.1, and Gühring and Raslan, 2021). Having shown (i), the next ingredient (ii) follows from arguments given (Schwab and Zech, 2019, Prop. 3.3), which formulates a DNN for multiplying  $d$  numbers as a binary tree of depth  $\lceil \log_2(d) \rceil$ . The final ingredient (iii) follows Daws and C. Webster (2019) and involves writing

$$\Psi_{\mathbf{v}}(\mathbf{y}) = \prod_{i \in \text{supp}(\mathbf{v})} \prod_{j=1}^{v_i} a_{ij}(y_i - r_{ij}), \quad \forall \mathbf{y} \in \mathcal{U}, \mathbf{v} \in \mathcal{F}, \quad (7.3)$$

for scalars  $a_{ij}, r_{ij} \in \mathbb{R}$ . We deduce that  $\Psi_{\mathbf{v}}(\mathbf{y})$  can be approximated by the composition of the affine map  $\mathbf{y} \mapsto (a_{ij}(y_i - r_{ij}))_{ij}$  and the tanh DNN that approximately computes the product in (7.3).

The precise bounds on the architecture given in Theorem 7.1 are now evident. First, each polynomial  $\Psi_{\mathbf{v}}$  involves a product of  $\|\mathbf{v}\|_1 \leq m(\Lambda)$  numbers. The depth bound now follows immediately. The width bound follows from the fact that the DNNs emulating each polynomial are stacked vertically to form the overall network  $\Phi_{\Lambda, \delta}$  which simultaneously emulates all  $|\Lambda|$  polynomials.

It is also evident that this proof can be adapted to any other activation function for which the multiplication map in (i) can be approximated with a DNN of a quantifiable width and depth. ReLU DNNs were also considered in Adcock

et al. (2023a, Thm. 7.4), with ingredients (i) and (ii) being based on Opschoor et al. (2022, Prop. 2.6). The resulting depth and width bounds are worse than the tanh case. Conversely, a RePU network of depth one and constant width can *exactly* represent the multiplication of two numbers (see, e.g., Li et al., 2020, Lem. 2.1). Therefore, a RePU network with the same width and depth bounds as in Theorem 7.1 can *exactly* emulate the Legendre polynomials, as opposed to approximately in the tanh case. The same holds for DNNs based on *rational* activation functions (Boullé et al., 2020).

Finally, we note that the proof can be easily adapted to other polynomial systems, such as Chebyshev or, more generally, Jacobi polynomials. In the RePU case, as before, this emulation is exact.

## 7.5 Existence theorem for $(\mathbf{b}, \varepsilon)$ -holomorphic functions

We now present an existence theorem for  $(\mathbf{b}, \varepsilon)$ -holomorphic functions.

**Theorem 7.2** (Existence theorem for DNNs). *Let  $\mathcal{V} = (\mathbb{R}, |\cdot|)$  and  $\mathbf{b} \in [0, \infty)^{\mathbb{N}}$  be such that  $\mathbf{b} \in \ell^p(\mathbb{N})$  for some  $0 < p < 1$ . Then, for every  $s \in \mathbb{N}$  there exists a class of tanh DNNs  $\Phi : \mathbb{R}^s \rightarrow \mathbb{R}$  with*

$$\text{width}(\mathcal{N}) \leq c_1 \cdot s^2, \quad \text{depth}(\mathcal{N}) \leq c_2 \cdot \log(s),$$

for universal constants  $c_1, c_2 > 0$  such that the following holds. For every  $f \in \mathcal{H}(\mathbf{b})$ , there is a  $\Phi \in \mathcal{N}$  such that

$$\|f - \Phi \circ \mathcal{T}_s\|_{L^2_{\mathbf{b}}(\mathcal{U})} \leq C(\mathbf{b}, p) \cdot s^{\frac{1}{2} - \frac{1}{p}}.$$

This result follows immediately from Theorems 6.4 and 7.1. The former asserts the existence of a lower set  $S$  of size  $|S| \leq s$  and depending on  $\mathbf{b}$  and  $p$  only which attains the desired rate of  $s^{1/2-1/p}$ . Then the latter asserts the existence of the DNN  $\Phi_{S,\delta}$ , with  $\delta = s^{-1/p}$ , of the desired size. Note that  $m(S) \leq s$  for any lower set. Indeed, if  $\mathbf{v} \in S$  then so does the set  $\{\boldsymbol{\mu} \in \mathcal{F} : \boldsymbol{\mu} \leq \mathbf{v}\}$ . Therefore

$$\|\mathbf{v}\|_1 \leq \prod_{i \in \mathbb{N}} (v_i + 1) = |\{\boldsymbol{\mu} \in \mathcal{F} : \boldsymbol{\mu} \leq \mathbf{v}\}| \leq |S| \leq s.$$

Theorem 7.2 only deals with scalar-valued functions. We will tackle Hilbert-valued functions in the next section when we consider learning via DNNs.

## 8 Practical existence theory: near-optimal DL

While existence theory provides crucial insight into the *expressivity* of DNNs, it says nothing about whether networks with similar approximation guarantees can be trained in practice and, in particular, how much training data suffices to do so. To narrow this gap between theory and practice, in this section we develop the topic of *practical existence theory*. The goal of this endeavor is to show that

there exists both an architecture *and* a training strategy that is similar to what is used in practice (i.e., minimizing a loss function) from which one provably learns near-optimal DNN approximations from the training data (6.1).

## 8.1 Setup

As in §6, we assume that  $\mathcal{V}$  is a Hilbert space with a finite-dimensional discretization  $\mathcal{V}_h$ . Following in Remark 6.7, we let  $\{\varphi_i\}_{i=1}^K$  be a basis for  $\mathcal{V}_h$ . If  $f : \mathcal{U} \rightarrow \mathcal{V}$  is the function to recover, then we write

$$f(\mathbf{y}) \approx \sum_{i=1}^K d_i(\mathbf{y})\varphi_i,$$

where the coefficients  $d_i : \mathcal{U} \rightarrow \mathbb{R}$  are scalar-valued functions. We now seek to approximate these functions using a DNN with  $K$  neurons on the output layer, with the  $i$ th neuron corresponding to the approximation of the function  $d_i$ . Let  $\mathcal{N}$  be a class of DNNs of the form  $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^K$  for some  $n \in \mathbb{N}$ . Then our aim is to use the training data (6.1) to compute a suitable  $\Phi \in \mathcal{N}$  that yields an approximation  $f_{\hat{\Phi}} \approx f$  defined by

$$f_{\Phi}(\mathbf{y}) := \sum_{i=1}^K (\Phi \circ \mathcal{T}_n(\mathbf{y}))_i \varphi_i, \quad \forall \mathbf{y} \in \mathcal{U}. \quad (8.1)$$

## 8.2 Practical existence theorem

The first practical existence theorems were shown in Adcock and Dexter (2021) (scalar-valued case) and Adcock et al. (2021) (Hilbert-valued case) for holomorphic function approximation in finite dimensions with ReLU DNNs. This was extended in Adcock et al. (2023a) to infinite-dimensional, holomorphic functions with ReLU, RePU or tanh activation functions. The following result is based on Adcock et al. (2023a, Thm. 4.4).

**Theorem 8.1** (Practical existence theorem for DNNs). *There are universal constants  $c_1, c_2, c_3 \geq 1$  such that the following holds. Let  $m \geq 3$ ,  $0 < \epsilon < 1$  and  $n = \lceil m/L \rceil$ , where  $L = L(m, \epsilon)$  is as in (6.17). There exists*

1. a class  $\mathcal{N}$  of DNNs  $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^K$  with tanh activation function satisfying

$$\text{width}(\mathcal{N}) \leq c_1 \cdot m^{3+\log_2(m)}, \quad \text{depth}(\mathcal{N}) \leq c_2 \cdot \log(m);$$

2. a regularization function  $\mathcal{J} : \mathcal{N} \rightarrow [0, \infty)$  equivalent to a certain norm of the trainable parameters;

3. and a choice of regularization parameter  $\lambda$  involving only  $m$  and  $\epsilon$ ;

such that following holds for every  $0 < p < 1$  and  $\mathbf{b} \in \ell_M^p(\mathbb{N})$ . Let  $f \in \mathcal{H}(\mathbf{b})$ ,  $\mathbf{y}_1, \dots, \mathbf{y}_m \sim_{\text{i.i.d.}} \mathcal{Q}$  and consider the training data (2.2). Then, with probability

at least  $1 - \epsilon$ , every  $\gamma$ -minimizer  $\hat{\Phi}$ ,  $\gamma \geq 0$ , of the training problem

$$\min_{\Phi \in \mathcal{N}} \mathcal{G}(\Phi), \quad \text{where } \mathcal{G}(\Phi) = \sqrt{\frac{1}{m} \sum_{i=1}^m \|f_{\Phi}(y_i) - d_i\|_{\mathcal{V}}^2} + \lambda \mathcal{J}(\Phi), \quad (8.2)$$

where  $f_{\Phi}$  is as in (8.1), satisfies, with  $\zeta$  as in (6.19),

$$\|f - f_{\hat{\Phi}}\|_{L^2_{\rho}(\mathcal{U}; \mathcal{V})} \leq c_3 \cdot \zeta. \quad (8.3)$$

In this theorem, as before, we term  $\hat{\Phi}$  a  $\gamma$ -minimizer of (8.2) if  $\mathcal{G}(\hat{\Phi}) \leq \min_{\Phi \in \mathcal{N}} \mathcal{G}(\Phi) + \gamma$ . Comparing with Theorem 6.8, we conclude that there is a DNN architecture and training procedure such that the resulting learned approximations achieve the same error bounds (up to possible constants) as sparse polynomial approximation based on weighted  $\ell^1$ -minimization. In particular, this procedure also achieves the near-optimal approximation rate  $(m/L)^{\frac{1}{2} - \frac{1}{p}}$ .

We term Theorem 8.1 a *practical existence theorem*. It not only asserts the existence of a DNN with a given architecture that achieves some desired rate of approximation, but demonstrates how to construct it from training data and gives generalization bounds that are explicit in the amount of training data  $m$ . Moreover, the training procedure is similar to standard DL procedures, in that it involves minimizing a (regularized) least-squares loss function.

While Theorem 8.1 only considers tanh DNNs, it can be readily adapted to other activations (Adcock et al., 2023a). As discussed next, all one requires to do this are variants of Theorem 7.1 for other activations, a topic we discussed previously in §7.4.

### 8.3 The mechanism of practical existence theorems

We now give some insight into the proof of Theorem 8.1, since this provides a general recipe for establishing practical existence theorems.

The overall mechanism involves not just emulating polynomials with DNNs, but also emulating the weighted  $\ell^1$ -minimization problem of §6.1 as a DNN training problem. More precisely, following §6.1 we proceed as follows.

- (i) Approximate the Legendre polynomials  $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \Lambda}$  using DNNs.
- (ii) Replace the polynomials in the matrix (6.3) by these DNNs, leading to a matrix  $A' \approx A$ . Then replace  $A$  by  $A'$  in (6.15).
- (iii) Re-cast (6.13) with  $A'$  as a training problem of the form (8.2), where the unknowns  $Z \in \mathbb{R}^{N \times K}$  correspond to the trainable parameters of the DNNs.
- (iv) Emulate the proof steps of Theorem 6.8, showing, as needed, that each step remains valid for the perturbed matrix  $A'$ .

Step (i) is accomplished straightforwardly by using the emulation result, Theorem 7.1. Since  $\Lambda$  is taken to be  $\Lambda = \Lambda_n^{\text{HCl}}$ , we have  $\text{supp}(\mathbf{v}) \subseteq \{1, \dots, n\}$



for  $\mathbf{v} \in \Lambda$ . Hence we set  $\Theta = \{1, \dots, n\}$  and consider a suitable parameter  $\delta$  that is chosen later in the proof to balance the ensuing error terms. Let  $\Phi_{\Lambda, \delta} : \mathbb{R}^n \rightarrow \mathbb{R}^{|\Lambda|}$  be the resulting DNN.

Step (ii) warrants no further discussion. Now consider Step (iii). Let  $N = |\Lambda|$  as before, and define the class of DNNs

$$\mathcal{N} = \left\{ \Phi = \mathbf{Z}^\top \Phi_{\Lambda, \delta} : \mathbf{Z} \in \mathbb{R}^{N \times K} \right\}.$$

This is a class of DNNs, where only the weight matrix on the output layer is trainable. The remaining layers are nontrainable, and are handcrafted to emulate the Legendre polynomials  $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \Lambda}$ .

Now let  $\mathbf{z} = (z_{\mathbf{v}_j})_{j=1}^N \in \mathcal{V}_h^N$ . Then we can associate  $\mathbf{z}$  with its matrix of coefficients  $\mathbf{Z} = (Z_{ij}) \in \mathbb{R}_{i,j=1}^{N \times K}$  via the relation

$$z_{\mathbf{v}_i} = \sum_{j=1}^K Z_{ij} \varphi_j, \quad \forall i = 1, \dots, N, \quad (8.4)$$

and consequently with the DNN  $\Phi = \mathbf{Z}^\top \Phi_{\Lambda, \delta} \in \mathcal{N}$ . Using this, it is a short argument to show that

$$f_\Phi(\mathbf{y}) = \sum_{\mathbf{v} \in \Lambda} z_{\mathbf{v}} \Phi_{\mathbf{v}, \delta} \circ \mathcal{T}_\Theta(\mathbf{y}).$$

Therefore,

$$\|\mathbf{A}' \mathbf{z} - \mathbf{f}\|_{2; \mathcal{V}}^2 = \frac{1}{m} \sum_{i=1}^m \left\| \sum_{\mathbf{v} \in \Lambda} z_{\mathbf{v}} \Phi_{\mathbf{v}, \delta}(\mathbf{y}_i) - d_i \right\|_{\mathcal{V}}^2 = \frac{1}{m} \sum_{i=1}^m \|f_\Phi(\mathbf{y}_i) - d_i\|_{\mathcal{V}}^2.$$

Now let  $\mathcal{J} : \mathcal{N} \rightarrow [0, \infty)$  be the regularization functional defined by

$$\mathcal{J}(\Phi) = \sum_{j=1}^N u_{\mathbf{v}_j} \left\| \sum_{k=1}^K Z_{jk} \varphi_k \right\|_{\mathcal{V}} = \sum_{j=1}^N u_{\mathbf{v}_j} \|z_{\mathbf{v}_j}\|_{\mathcal{V}} = \|\mathbf{z}\|_{1, \mathbf{u}; \mathcal{V}},$$

for  $\Phi = \mathbf{Z}^\top \Phi_{\Lambda, \delta} \in \mathcal{N}$ , where  $\mathbf{z} = (z_{\mathbf{v}_i})_{i=1}^N$  is as defined by (8.4). We readily see that  $\mathcal{J}$  is a norm over the trainable parameters. Using this and the previous expression we see that (6.13) with the matrix  $\mathbf{A}'$  in place of  $\mathbf{A}$  can be re-cast as a training problem of the form (8.2).

Finally, consider step (iv). This step is facilitated by the perturbation bound  $\|\mathbf{A} - \mathbf{A}'\|_2 \leq \sqrt{N} \delta$ , which in turn follows from a short argument via the Cauchy-Schwarz inequality:

$$\begin{aligned} \|(A - A')z\|_2^2 &= \frac{1}{m} \sum_{i=1}^m \left| \sum_{\mathbf{v} \in \Lambda} (\Psi_{\mathbf{v}}(y_i) - \Phi_{\mathbf{v}, \delta}(y_i)) z_{\mathbf{v}} \right|^2 \leq \left( \sum_{\mathbf{v} \in \Lambda} \delta |z_{\mathbf{v}}| \right)^2 \\ &\leq \delta^2 N \|z\|_2^2. \end{aligned}$$

The remainder of this step involves a careful modification of the proof of Theorem 6.8 to take into account this perturbation. Theorem 6.8 involves compressed sensing techniques, and relies on first asserting that the matrix  $A$  has a certain *weighted robust Null Space Property (rNSP)* (Rauhut and Ward, 2016) (see also Adcock et al., 2022, Chpt. 6). This is a weighted variant of the classical (unweighted) rNSP, and like the latter, is a slightly weaker condition than the better known (weighted) RIP (Foucart and Rauhut, 2013). Crucially, it can be shown that the (weighted) rNSP is preserved under sufficiently small perturbations (Adcock and Hansen, 2021, Lem. 8.5).

We note in passing that the width and depth bounds in Theorem 8.1 follow quite directly from Theorem 7.1. Recall that  $\Lambda = \Lambda_n^{\text{HCl}}$  is as in (6.12). A standard estimate (see Kühn et al., 2015, Thm. 4.9) gives that

$$|\Lambda| \leq en^{2+\log_2(n)}, \quad \forall n \in \mathbb{N}. \quad (8.5)$$

Moreover, by definition, any  $\mathbf{v} \in \Lambda$  satisfies  $\|\mathbf{v}\|_1 \leq \prod_{k=1}^{n-1} (v_k + 1) \leq n$ . Hence

$$m(\Lambda) \leq n. \quad (8.6)$$

The desired width and depth bounds now follow immediately from Theorem 7.1 and these estimates, along with the (somewhat loose) bound  $n \leq m$ .

## 9 Epilogue

This work has been about the approximation of smooth, infinite-dimensional functions from limited data. We close with a discussion on the benefits and consequences of practical existence theory and the gap between the *handcrafted* models on which it is based and the *fully trained* models used in practice.

### 9.1 Scientific computing and data scarcity

As discussed in §1.1, parametric DE problems are often *data scarce*. This is the case for many problems in scientific computing in which machine learning and, specifically, DL, is currently being applied. It also stands in stark contrast to more classical DL applications such as image classification, where datasets usually contain tens of millions of images or more. Therefore, understanding the sample complexity is crucial. Practical existence theorems show that DNNs can be learned from data in a sample-efficient manner.

## 9.2 Potential benefits to DNNs over sparse polynomials

Having said this, we emphasize that practical existence theorems are, at least in this work, intended primarily as *theoretical* contributions. Since the strategy in Theorem 8.1 involves emulating the sparse polynomial approximation scheme constructed in §6, there is no benefit to implementing it over the latter.

However, in related work in inverse problems in imaging (recall the discussion in §1.3), unrolling is both used to establish practical existence theorems *and* as a principled way to design DNN architectures which can then be trained as part of a DL strategy (Monga et al., 2021), potentially using the theoretical weights and biases as initialization. It remains to see whether similar ideas could be effective in the parametric DE setting. Some initial work in this direction can be found in Daws and C.G. Webster (2019).

In addition, existence theory establishes that DNNs have the capacity to approximate broad classes of functions efficiently. This is *not* the case for polynomials, which fail dramatically on, for instance, discontinuous functions. This work has focused on classes of holomorphic functions, where polynomials are well suited (in fact, near optimal). DNN-based schemes have the potential to succeed on quite different function classes, something that distinguishes them from traditional methods of scientific computing. We remark in passing that discontinuous or sharp transitions arise frequently in parametric model problems, and are difficult to treat with standard methods (Elman and Miller, 2012; Gunzburger et al., 2014a,b; Jakeman et al., 2011; Ma and Zabarar, 2009; Zhang et al., 2016).

## 9.3 Theorem 8.1 does not eliminate the theory-practice gap

DNNs in Theorem 8.1 are handcrafted to emulate Legendre polynomials, with only the final layer being trained. Standard DL methods use fully trained models, where all layers are trained. Theorem 8.1 says nothing directly about this practice. However, it does lead to some insights, as we now discuss.

## 9.4 Practical insights

### *Width and depth bounds*

The aim of practical existence theory is to express the error in terms of the sample complexity  $m$ . This differs from standard existence theory in which express the error is expressed in terms of the complexity of the network, e.g., its width and depth, or its *size* (number of nonzero weights and biases). Nonetheless, it is worth discussing the network complexity in Theorem 8.1.

This theorem describes architectures that are much wider than they are deep. In fact, the depth grows very slowly with the number of samples  $m$ , like  $\log(m)$ . This broadly agrees with empirical insights from the application of DL in scientific computing, where it is often observed that relatively shallow networks

perform well. See De Ryck et al. (2021) and references therein, as well as Adcock et al. (2021); Adcock and Dexter (2021).

However, while the depth bound in Theorem 8.1 is somewhat reasonable, the width bound of  $O(m^{3+\log_2(m)})$  grows extremely rapidly with  $m$  (albeit subexponentially). As discussed in §8.3, the large estimate for the width arises, in great part, from having to emulate all polynomials in the hyperbolic cross  $\Lambda_n^{\text{HCl}}$ , whose size, as shown in (8.5), behaves like  $O(n^{2+\log_2(n)})$ . This has nothing to do with DNN approximation itself. It stems from the sparse polynomial approximation scheme and, as discussed in §6.3, the need to build a finite search set  $\Lambda$  outside of which the error is ensured to be small. It is unclear whether the near-optimal approximation rates of Theorem 6.8 can be obtained with a smaller search space without further assumptions on the functions being approximated. This is an interesting prospect for future work.

### *Post-training pruning and sparsification*

On the other hand, it is always possible to sparsify the DNN  $\hat{\Phi}$  learned in Theorem 8.1 after training – a process known as *pruning* (see Frankle and Carbin, 2019, and references therein). This exploits the fact that the networks in  $\mathcal{N}$ , while very wide, are sparsely connected. This is done in much the same way as the in polynomial case (Remark 6.9). If  $\hat{\mathbf{C}} \in \mathbb{R}^{N \times K}$  are the weights of  $\hat{\Phi}$ , then one first forms  $\hat{\mathbf{c}} = (\hat{c}_v)_{v \in \Lambda} \in \mathcal{V}_h^N$  using (8.4), then computes the index set  $S \subseteq \Lambda$  of the largest  $n$  entries of  $(\|\hat{c}_v\|_{\mathcal{V}})_{v \in \Lambda}$ , and finally replaces  $\hat{\Phi}$  with  $\check{\Phi} = \mathbf{Z}_S^\top \Phi_{S,\delta}$ , where  $\mathbf{Z}_S \in \mathbb{R}^{n \times K}$  is formed from the rows of  $\mathbf{Z}$  with indices in  $S$ . By Theorem 7.1 and the bound  $m(S) \leq m(\Lambda) \leq n$  (recall (8.6)), we have  $\text{width}(\check{\Phi}) \leq c_1 m^2$  and  $\text{depth}(\check{\Phi}) \leq c_2 \log(m)$ . In particular,  $\check{\Phi}$  is significantly narrower than  $\hat{\Phi}$ .

This suggests that methods to promote sparsity in training (see Hoefler et al., 2021, and references therein) may be beneficial in practice when training fully-connected models for scientific computing problems. This requires further investigation.

## 9.5 Eliminating the gap: beating the Monte Carlo rate is key

Despite these insights, the gap between theory and practice persists. It is worth noting that standard approaches to estimating the generalization error for fully-trained models based on statistical learning theory and the *bias-variance decomposition* (see, e.g., Beck et al., 2022; Chen et al., 2022; Ohn and Kim, 2019; Schmidt-Hieber, 2020; Suzuki, 2019 and references therein) are not immediately applicable to this setting. These approaches use estimates for the covering number or Rademacher complexity of the relevant DNN classes. Unfortunately, they typically lead to rates that decay at best like  $m^{-1/2}$ . These rates are strictly slower than near-optimal rates  $m^{1/2-1/p}$ , up to log factors, asserted in Theorem 8.1. A major theme in parametric DEs and computational UQ is building methods that beat the *Monte Carlo* rate  $m^{-1/2}$  (Adcock et al., 2022, Chpt. 1).

Whether these approaches could be meaningfully combined with practical existence theorems is an interesting question for future work.

## 9.6 Conclusion

To summarize, practical existence theory is a promising way to study DNN approximation which has the potential to give new insights into the promises and challenges of data scarce applications in computational science and engineering. In addition to those outlined above, several interesting future avenues include the design of improved training methodologies and novel architectures and activation functions. Studying these areas is key for pushing the boundaries of what DNNs can achieve in scientific computing, particularly in the face of scarce data and complex computational tasks. These efforts may not only narrow the theory-to-practice gap but also unlock new DL approaches, making it more efficient, accurate, and applicable across a broader spectrum of scientific challenges and guiding the way towards more sophisticated and capable DNN models.

## References

- Abedeljawad, A., Grohs, P., 2023. Sampling complexity of deep approximation spaces. arXiv:2312.1337.
- Adcock, B., 2017. Infinite-dimensional  $\ell^1$  minimization and function approximation from pointwise data. *Constr. Approx.* 45 (3), 343–390.
- Adcock, B., Bao, A., Brugiapaglia, S., 2019. Correcting for unknown errors in sparse high-dimensional function approximation. *Numer. Math.* 142 (3), 667–711.
- Adcock, B., Brugiapaglia, S., 2023. Monte Carlo is a good sampling strategy for polynomial approximation in high dimensions. arXiv:2208.09045.
- Adcock, B., Brugiapaglia, S., Dexter, N., Moraga, S., 2021. Deep neural networks are effective at learning high-dimensional Hilbert-valued functions from limited data. In: Bruna, J., Hesthaven, J.S., Zdeborová, L. (Eds.), *Proceedings of the Second Annual Conference on Mathematical and Scientific Machine Learning*. In: *Proc. Mach. Learn. Res.* (PMLR), vol. 145. PMLR, pp. 1–36.
- Adcock, B., Brugiapaglia, S., Dexter, N., Moraga, S., 2023a. Near-optimal learning of Banach-valued, high-dimensional functions via deep neural networks. arXiv:2211.12633.
- Adcock, B., Brugiapaglia, S., Dexter, N., Moraga, S., 2024a. On Efficient Algorithms for Computing Near-Best Polynomial Approximations to High-Dimensional, Hilbert-Valued Functions from Limited Samples. *Mem. Eur. Math. Soc.* <https://dx.doi.org/10.4171/MEMS/13>.
- Adcock, B., Brugiapaglia, S., Webster, C.G., 2022. Sparse Polynomial Approximation of High-Dimensional Functions. *Comput. Sci. Eng. Society for Industrial and Applied Mathematics, Philadelphia, PA.*
- Adcock, B., Colbrook, M.J., Neyra-Nesterenko, M., 2023b. Restarts subject to approximate sharpness: a parameter-free and optimal scheme for first-order methods. arXiv:2301.02268.
- Adcock, B., Dexter, N., 2021. The gap between theory and practice in function approximation with deep neural networks. *SIAM J. Math. Data Sci.* 3 (2), 624–655.
- Adcock, B., Dexter, N., Moraga, S., 2023c. Optimal approximation of infinite-dimensional holomorphic functions II: recovery from i.i.d. pointwise samples. arXiv:2310.16940.
- Adcock, B., Dexter, N., Moraga, S., 2024b. Optimal approximation of infinite-dimensional holomorphic functions. *Calcolo* 61 (12).
- Adcock, B., Hansen, A.C., 2021. *Compressive Imaging: Structure, Sampling, Learning*. Cambridge University Press, Cambridge, UK.

- Ajavon, K., 2024. Surrogate models for diffusion on graphs: a high-dimensional polynomial approach. Master's thesis. Concordia University.
- Antun, V., Gottschling, N.M., Hansen, A.C., Adcock, B., 2023. Am (A)I hallucinating? Non-robustness, hallucinations and unpredictable performance of AI for MR image reconstruction. Preprint.
- Antun, V., Renna, F., Poon, C., Adcock, B., Hansen, A.C., 2020. On instabilities of deep learning in image reconstruction and the potential costs of AI. *Proc. Natl. Acad. Sci. USA* 117 (48), 30088–30095.
- Babuška, I.M., Nobile, F., Tempone, R., 2007. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.* 43 (3), 1005–1034.
- Bäck, J., Nobile, F., Tamellini, L., Tempone, R., 2011. Stochastic spectral Galerkin and collocation methods for PDEs with random coefficients: a numerical comparison. In: Hesthaven, Jan S., Rønquist, Einar M. (Eds.), *Spectral and High Order Methods for Partial Differential Equations*. In: *Lect. Notes Comput. Sci. Eng.*, vol. 76. Springer, Berlin, Heidelberg, Germany, pp. 43–62.
- Beck, C., Jentzen, A., Kuckuck, B., 2022. Full error analysis for the training of deep neural networks. *Infin. Dimens. Anal. Quantum Probab. Relat. Top.* 25 (2), 2150020.
- Beck, J., Nobile, F., Tamellini, L., Tempone, R., 2014. Convergence of quasi-optimal stochastic Galerkin methods for a class of PDEs with random coefficients. *Comput. Math. Appl.* 67 (4), 732–751.
- Beck, J., Tempone, R., Nobile, F., Tamellini, L., 2012. On the optimal polynomial approximation of stochastic PDEs by Galerkin and collocation methods. *Math. Models Methods Appl. Sci.* 22 (9), 1250023.
- Becker, S., Jentzen, A., Müller, M.S., von Wurstemberger, P., 2023. Learning the random variables in Monte Carlo simulations with stochastic gradient descent: machine learning for parametric PDEs and financial derivative pricing. *Math. Finance* 34 (1), 90–150.
- Berveiller, M., Sudret, B., Lemaire, M., 2006. Stochastic finite element: a non intrusive approach by regression. *Eur. J. Comput. Mech.* 15 (1–3), 81–92.
- Bhattacharya, K., Hosseini, N., Kovachki, B., Stuart, A., 2021. Model reduction and neural networks for parametric PDEs. *J. Comput. Math.* 7, 121–157.
- Bieri, M., Andreev, R., Schwab, C., 2010. Sparse tensor discretization of elliptic SPDEs. *SIAM J. Sci. Comput.* 31 (6), 4281–4304.
- Blanchard, M., Bennouna, M.A., 2020. The representation power of neural networks: breaking the curse of dimensionality. arXiv:2012.05451.
- Blatman, G., Sudret, B., 2011. Adaptive sparse polynomial chaos expansion based on least angle regression. *J. Comput. Phys.* 230, 2345–2367.
- Bölcskei, H., Grohs, P., Kutyniok, G., Petersen, P., 2019. Optimal approximation with sparsely connected deep neural networks. *SIAM J. Math. Data Sci.* 1 (1), 8–45.
- Bonito, A., DeVore, R., Guignard, D., Jantsch, P., Petrova, G., 2021. Polynomial approximation of anisotropic analytic functions of several variables. *Constr. Approx.* 53, 319–348.
- Boullé, N., Nakatsukasa, Y., Townsend, A., 2020. Rational neural networks. In: *Advances in Neural Information Processing Systems*, pp. 14243–14253.
- Boullé, N., Townsend, A., 2023. A mathematical guide to operator learning. arXiv:2312.14688.
- Brugiapaglia, S., Dexter, N., Karam, S., Wang, W., 2024. Physics-informed deep learning and compressive collocation for high-dimensional diffusion-reaction equations: practical existence theory and numerics. Preprint.
- Brugiapaglia, S., Dirksen, S., Jung, H.C., Rauhut, H., 2021. Sparse recovery in bounded Riesz systems with applications to numerical methods for PDEs. *Appl. Comput. Harmon. Anal.* 53, 231–269.
- Castrillón-Candás, J.E., Nobile, F., Tempone, R., 2016. Analytic regularity and collocation approximation for elliptic PDEs with random domain deformations. *Comput. Math. Appl.* 71 (6), 1173–1197.
- Chambolle, A., Pock, T., 2011. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.* 40 (1), 120–145.

- Chambolle, A., Pock, T., 2016. On the ergodic convergence rates of a first-order primal-dual algorithm. *Math. Program.* 159 (1–2), 253–287.
- Chen, M., Jiang, H., Liao, W., Zhao, T., 2022. Nonparametric regression on low-dimensional manifolds using deep ReLU networks: function approximation and statistical recovery. *Inf. Inference* 11 (4), 1203–1253.
- Chkifa, A., Cohen, A., Migliorati, G., Nobile, F., Tempone, R., 2015a. Discrete least squares polynomial approximation with random evaluations - application to parametric and stochastic elliptic PDEs. *ESAIM: Math. Model. Numer. Anal.* 49 (3), 815–837.
- Chkifa, A., Cohen, A., Schwab, C., 2014. High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs. *Found. Comput. Math.* 14 (4), 601–633.
- Chkifa, A., Cohen, A., Schwab, C., 2015b. Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs. *J. Math. Pures Appl.* 103 (2), 400–428.
- Choi, B., Iwen, M.A., Krahmer, F., 2021a. Sparse harmonic transforms: a new class of sublinear-time algorithms for learning functions of many variables. *Found. Comput. Math.* 21 (2), 275–329.
- Choi, B., Iwen, M.A., Volkmer, T., 2021b. Sparse harmonic transforms II: best  $s$ -term approximation guarantees for bounded orthonormal product bases in sublinear-time. *Numer. Math.* 148 (2), 293–362.
- Cicci, L., Fresca, S., Manzoni, A., 2022. Deep-HyROMnet: a deep learning-based operator approximation for hyper-reduction of nonlinear parametrized PDEs. *J. Sci. Comput.* 93, 57.
- Cohen, A., Davenport, M.A., Leviatan, D., 2013. On the stability and accuracy of least squares approximations. *Found. Comput. Math.* 13, 819–834.
- Cohen, A., DeVore, R.A., 2015. Approximation of high-dimensional parametric PDEs. *Acta Numer.* 24, 1–159.
- Cohen, A., DeVore, R.A., Schwab, C., 2010. Convergence rates of best  $N$ -term Galerkin approximations for a class of elliptic sPDEs. *Found. Comput. Math.* 10, 615–646.
- Cohen, A., DeVore, R.A., Schwab, C., 2011. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's. *Anal. Appl. (Singap.)* 9 (1), 11–47.
- Cohen, A., Migliorati, G., Nobile, F., 2017. Discrete least-squares approximations over optimized downward closed polynomial spaces in arbitrary dimension. *Constr. Approx.* 45, 497–519.
- Cohen, A., Schwab, C., Zech, J., 2018. Shape holomorphy of the stationary Navier–Stokes equations. *SIAM J. Math. Anal.* 50 (2), 1720–1752.
- Colbrook, M.J., Antun, V., Hansen, A.C., 2022. The difficulty of computing stable and accurate neural networks: on the barriers of deep learning and Smale's 18th problem. *Proc. Natl. Acad. Sci. USA* 119 (12), e2107151119.
- Cybenko, G., 1989. Approximation by superpositions of a sigmoidal function. *Math. Control Signals Syst.* 2 (4), 303–314.
- Cyr, E.C., Gulian, M.A., Patel, R.G., Perego, M., Trask, N.A., 2020. Robust training and initialization of deep neural networks: an adaptive basis viewpoint. In: Lu, Jianfeng, Ward, Rachel (Eds.), *Proceedings of the First Mathematical and Scientific Machine Learning Conference*. In: *Proceedings of Machine Learning Research*, vol. 107. Princeton University/PMLR, Princeton, NJ, USA, pp. 512–536.
- Dai, F., Temlyakov, V., 2023. Universal sampling discretization. *Constr. Approx.* 58, 589–613.
- Dal Santo, N., Deparis, S., Pegolotti, L., 2020. Data driven approximation of parametrized PDEs by reduced basis and neural networks. *J. Comput. Phys.* 416, 109550.
- Daws, J., Webster, C., 2019. Analysis of deep neural networks with quasi-optimal polynomial approximation rates. *arXiv:1912.02302*.
- Daws, J., Webster, C.G., 2019. A polynomial-based approach for architectural design and learning with deep neural networks. *arXiv:1905.10457*.
- De Ryck, T., Lanthaler, S., Mishra, S., 2021. On the approximation of functions by tanh neural networks. *Neural Netw.* 143, 732–750.
- DeVore, R., Hanin, B., Petrova, G., 2021. Neural network approximation. *Acta Numer.* 30, 327–444.
- DeVore, R.A., 1998. Nonlinear approximation. *Acta Numer.* 7, 51–150.

- Doostan, A., Owhadi, H., 2011. A non-adapted sparse approximation of PDEs with stochastic inputs. *J. Comput. Phys.* 230 (8), 3015–3034.
- Düing, D., Nguyen, V.K., 2021. Deep ReLU neural networks in high-dimensional approximation. *Neural Netw.* 142, 619–635.
- Düing, D., Nguyen, V.K., Pham, D.T., 2023. Deep ReLU neural network approximation in Bochner spaces and applications to parametric PDEs. *J. Complex.* 79, 101779.
- Düing, D., Temlyakov, V., Ullrich, T., 2018. *Hyperbolic Cross Approximation*. Adv. Courses Math. CRM Barcelona. Birkhäuser, Basel, Switzerland.
- E, W., Ma, C., Wu, L., 2021. The Barron space and the flow-induced function spaces for neural network models. *Constr. Approx.* 55, 369–406.
- E, W., Wang, Q., 2018. Exponential convergence of the deep neural network approximation for analytic functions. *Sci. China Math.* 61 (10), 1733–1740.
- Elbrächter, D., Perekrestenko, D., Grohs, P., Bölskei, H., 2021. Deep neural network approximation theory. *IEEE Trans. Inf. Theory* 67 (6), 2581–2623.
- Elman, H.C., Miller, C.W., 2012. Stochastic collocation with kernel density estimation. *Comput. Methods Appl. Mech. Eng.* 245–246, 36–46.
- Foucart, S., Rauhut, H., 2013. *A Mathematical Introduction to Compressive Sensing*. Appl. Numer. Harmon. Anal. Birkhäuser, New York, NY.
- Franco, N.R., Brugiapaglia, S., 2024. A practical existence theorem for reduced order models based on convolutional autoencoders. arXiv:2402.00435.
- Frankle, J., Carbin, M., 2019. The lottery ticket hypothesis: finding sparse, trainable neural networks. In: *ICLR*.
- Ganapathysubramanian, B., Zabaras, N., 2007. Sparse grid collocation schemes for stochastic natural convection problems. *J. Comput. Phys.* 225 (1), 652–685.
- Geist, M., Petersen, P., Raslan, M., Schneider, R., Kutyniok, G., 2021. Numerical solution of the parametric diffusion equation by deep neural networks. *J. Sci. Comput.* 88 (22).
- Ghanem, R.G., Spanos, P.D., 2003. *Stochastic Finite Elements: A Spectral Approach*, revised edition. Dover Publications, Inc., Mineola, NY.
- Ghanem, Roger, Higdon, David, Owhadi, Houman (Eds.), 2017. *Handbook of Uncertainty Quantification*. Springer, Switzerland.
- Grohs, P., Voigtlaender, F., 2023. Proof of the theory-to-practice gap in deep learning via sampling complexity bounds for neural network approximation spaces. *Found. Comput. Math.* <https://doi.org/10.1007/s10208-023-09607-w>.
- Gühring, I., Kutyniok, G., Petersen, P., 2020. Error bounds for approximations with deep ReLU neural networks in  $W^{s,p}$  norms. *Anal. Appl. (Singap.)* 18 (05), 803–859.
- Gühring, I., Raslan, M., 2021. Approximation rates for neural networks with encodable weights in smoothness spaces. *Neural Netw.* 134, 107–130.
- Gunzburger, M., Webster, C.G., Zhang, G., 2014a. An adaptive wavelet stochastic collocation method for irregular solutions of partial differential equations with random input data. In: Garcke, Jochen, Pflüger, Dirk (Eds.), *Sparse Grids and Applications – Munich 2012*. In: *Lect. Notes Comput. Sci. Eng.*, vol. 97. Springer, Cham, Switzerland, pp. 137–170.
- Gunzburger, M., Webster, C.G., Zhang, G., 2014b. Stochastic finite element methods for partial differential equations with random input data. *Acta Numer.* 23, 521–650.
- Guo, L., Narayan, A., Zhou, T., 2020. Constructing least-squares polynomial approximations. *SIAM Rev.* 62 (2), 483–508.
- Hadigol, M., Doostan, A., 2018. Least squares polynomial chaos expansion: a review of sampling strategies. *Comput. Methods Appl. Mech. Eng.* 332, 382–407.
- Hampton, J., Doostan, A., 2017. Compressive sampling methods for sparse polynomial chaos expansions. In: Ghanem, Roger, Higdon, David, Owhadi, Houman (Eds.), *Handbook of Uncertainty Quantification*. Springer, Cham, Switzerland, pp. 827–855.
- Hansen, M., Schwab, C., 2013a. Analytic regularity and nonlinear approximation of a class of parametric semilinear elliptic PDEs. *Math. Nachr.* 286 (8–9), 832–860.



- Hansen, M., Schwab, C., 2013b. Sparse adaptive approximation of high dimensional parametric initial value problems. *Vietnam J. Math.* 41 (2), 181–215.
- Hei, C., Ghring, I., Eigel, M., 2021. A neural multilevel method for high-dimensional parametric PDEs. In: *Advances in Neural Information Processing Systems*.
- Hei, C., Ghring, I., Eigel, M., 2023. Multilevel CNNs for parametric PDEs. *J. Mach. Learn. Res.* 24, 1–42.
- Herrmann, L., Schwab, C., Zech, J., 2022. Neural and spectral operator surrogates: unified construction and expression rate bounds. [arXiv:2207.04950v1](https://arxiv.org/abs/2207.04950v1).
- Hoang, V.H., Schwab, C., 2012. Regularity and generalized polynomial chaos approximation of parametric and random second-order hyperbolic partial differential equations. *Anal. Appl. (Singap.)* 10 (3), 295–326.
- Hoefler, T., Alistarh, D., Ben-Nun, T., Dryden, N., Peste, A., 2021. Sparsity in deep learning: pruning and growth for efficient inference and training in neural networks. *J. Mach. Learn. Res.* 23, 1–124.
- Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. *Neural Netw.* 2 (5), 359–366.
- Jakeman, J.D., Archibald, R., Xiu, D., 2011. Characterization of discontinuities in high-dimensional stochastic problems on adaptive sparse grids. *J. Comput. Phys.* 230 (10), 3977–3997.
- Kashin, B., Kosov, E., Limonova, I., Temlyakov, V., 2022. Sampling discretization and related problems. *J. Complex.* 71, 101653.
- Khara, B., Balu, A., Joshi, A., Sarkar, S., Hegde, C., Krishnamurthy, A., Ganapathysubramanian, B., 2021. NeuFENet: neural finite element solutions with theoretical bounds for parametric PDEs. [arXiv:2110.01601](https://arxiv.org/abs/2110.01601).
- Khoo, Y., Lu, J., Ying, L., 2021. Solving parametric PDE problems with artificial neural networks. *Eur. J. Appl. Math.* 32 (3), 421–435.
- Kovachki, N., Li, Z., Liu, B., Azizzadnesheli, K., Bhattacharya, K., Stuart, A., Anandkumar, A., 2023. Neural operator: learning maps between function spaces with applications to PDEs. *J. Mach. Learn. Res.* 24, 1–97.
- Kovachki, N.B., Lanthaler, S., Stuart, A.M., 2024. Operator learning: algorithms and analysis. [arXiv:2402.15715](https://arxiv.org/abs/2402.15715).
- Khn, T., Sickel, W., Ullrich, T., 2015. Approximation of mixed order Sobolev functions on the  $d$ -torus: asymptotics, preasymptotics, and  $d$ -dependence. *Constr. Approx.* 42, 353–398.
- Kunoth, A., Schwab, C., 2013. Analytic regularity and GPC approximation for control problems constrained by linear parametric elliptic and parabolic PDEs. *SIAM J. Control Optim.* 51 (3), 2442–2471.
- Laakmann, F., Petersen, P., 2021. Efficient approximation of solutions of parametric linear transport equations by ReLU DNNs. *Adv. Comput. Math.* 47 (11).
- Lanthaler, S., 2023. Operator learning with PCA-Net: upper and lower complexity bounds. [arXiv:2303.16317](https://arxiv.org/abs/2303.16317).
- Le Matre, O., Knio, O.M., 2010. *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Sci. Comput. Springer, Dordrecht, Netherlands.
- Lei, Z., Shi, L., Zeng, C., 2022. Solving parametric partial differential equations with deep rectified quadratic unit neural networks. *J. Sci. Comput.* 93, 80.
- Li, B., Tang, S., Yu, H., 2020. Better approximations of high dimensional smooth functions by deep neural networks with rectified power units. *Commun. Comput. Phys.* 27, 379–411.
- Li, Z., Kovachki, N., Azizzadnesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A., 2021. Fourier neural operator for parametric partial differential equations. In: *ICLR*.
- Liang, S., Srikant, R., 2017. Why deep neural networks for function approximation? In: *ICLR*.
- Longo, M., Opschoor, J.A.A., Disch, N., Schwab, C., Zech, J., 2023. De Rham compatible deep neural network FEM. *Neural Netw.* 165, 721–739.
- Lu, J., Shen, Z., Yang, H., Zhang, S., 2021. Deep network approximation for smooth functions. *SIAM J. Math. Anal.* 53 (5), 5465–5506.

- Lu, L., Jin, P., Pang, Z., Zhang, G., Karniadakis, G.E., 2021. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat. Mach. Intell.* 3, 218–229.
- Ma, X., Zabaras, N., 2009. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *J. Comput. Phys.* 228 (8), 3084–3113.
- Mathelin, L., Gollivan, K.A., 2012. A compressed sensing approach for partial differential equations with random input data. *Commun. Comput. Phys.* 12 (4), 919–954.
- Mathelin, L., Hussaini, M.Y., Zang, T.A., 2005. Stochastic approaches to uncertainty quantification in CFD simulations. *Numer. Algorithms* 38 (1–3), 209–236.
- Mhaskar, H., 1993. Approximation properties of a multilayered feedforward artificial neural network. *Adv. Comput. Math.* 1, 61–80.
- Mhaskar, H., 1996. Neural networks for optimal approximation of smooth and analytic functions. *Neural Comput.* 8 (1), 164–177.
- Migliorati, G., 2013. Polynomial approximation by means of the random discrete  $L^2$  projection and application to inverse problems for PDEs with stochastic data. PhD thesis. Politecnico di Milano.
- Migliorati, G., 2015. Adaptive polynomial approximation by means of random discrete least squares. In: Abdulle, Assyr, Deparis, Simone, Kressner, Daniel, Nobile, Fabio, Picasso, Marco (Eds.), *Numerical Mathematics and Advanced Applications – ENUMATH 2013*. Springer, Cham, Switzerland, pp. 547–554.
- Migliorati, G., 2019. Adaptive approximation by optimal weighted least squares methods. *SIAM J. Numer. Anal.* 57 (5), 2217–2245.
- Migliorati, G., Nobile, F., Tempone, R., 2015. Convergence estimates in probability and in expectation for discrete least squares with noisy evaluations at random points. *J. Multivar. Anal.* 142, 167–182.
- Migliorati, G., Nobile, F., von Schwerin, E., Tempone, R., 2013. Approximation of quantities of interest in stochastic PDEs by the random discrete  $L^2$  projection on polynomial spaces. *SIAM J. Sci. Comput.* 35 (3), A1440–A1460.
- Migliorati, G., Nobile, F., von Schwerin, E., Tempone, R., 2014. Analysis of the discrete  $L^2$  projection on polynomial spaces with random evaluations. *Found. Comput. Math.* 14, 419–456.
- Monga, V., Li, Y., Eldar, Y.C., 2021. Algorithm unrolling: interpretable, efficient deep learning for signal and image processing. *IEEE Signal Process. Mag.* 38 (2), 18–44.
- Montanelli, H., Du, Q., 2019. New error bounds for deep ReLU networks using sparse grids. *SIAM J. Math. Data Sci.* 1 (1), 78–92.
- Montanelli, H., Yang, H., Du, Q., 2021. Deep ReLU networks overcome the curse of dimensionality for bandlimited functions. *J. Comput. Math.* 39 (6), 801–815.
- Neyra-Nesterenko, M., Adcock, B., 2023. NESTANets: stable, accurate and efficient neural networks for analysis-sparse inverse problems. *Sampl. Theory Signal Process. Data Anal.* 21 (4).
- Nobile, F., Tempone, R., Webster, C.G., 2008a. An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* 46 (5), 2411–2442.
- Nobile, F., Tempone, R., Webster, C.G., 2008b. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* 46 (5), 2309–2345.
- Novak, E., 1988. *Deterministic and Stochastic Error Bounds in Numerical Analysis*. Number 1. Springer Berlin, Heidelberg.
- Novak, E., Woźniakowski, H., 2008. *Tractability of Multivariate Problems, Volume I: Linear Information*, vol. 6. European Math. Soc. Publ. House, Zürich.
- Novak, E., Woźniakowski, H., 2010. *Tractability of Multivariate Problems, Volume II: Standard Information for Functionals*, vol. 12. European Math. Soc., Zürich.
- Ohn, I., Kim, Y., 2019. Smooth function approximation by deep neural networks with general activation functions. *Entropy* 21 (7), 627.
- Opschoor, J.A.A., Petersen, P.C., Schwab, C., 2020. Deep ReLU networks and high-order finite element methods. *Anal. Appl. (Singap.)* 18 (5), 715–770.

- Opschoor, J.A.A., Schwab, C., 2023. Deep ReLU networks and high-order finite element methods II: Chebyshev emulation. arXiv:2310.07261.
- Opschoor, J.A.A., Schwab, Ch., Zech, J., 2022. Exponential ReLU DNN expression of holomorphic maps in high dimension. *Constr. Approx.* 55, 537–582.
- Peng, J., Hampton, J., Doostan, A., 2014. A weighted  $\ell_1$ -minimization approach for sparse polynomial chaos expansions. *J. Comput. Phys.* 267, 92–111.
- Petersen, P., Voigtlaender, F., 2018. Optimal approximation of piecewise smooth functions using deep ReLU neural networks. *Neural Netw.* 108, 296–330.
- Pinkus, A., 1999. Approximation theory of the MLP model in neural networks. *Acta Numer.* 8, 143–195.
- Poggio, T., Mhaskar, H., Rosasco, L., Miranda, B., Liao, Q., 2017. Why and when can deep-but not shallow-networks avoid the curse of dimensionality: a review. *Int. J. Autom. Comput.* 14, 503–519.
- Rauhut, H., Schwab, C., 2017. Compressive sensing Petrov-Galerkin approximation of high-dimensional parametric operator equations. *Math. Comput.* 86, 661–700.
- Rauhut, H., Ward, R., 2012. Sparse Legendre expansions via  $\ell_1$ -minimization. *J. Approx. Theory* 164 (5), 517–533.
- Rauhut, H., Ward, R., 2016. Interpolation via weighted  $\ell^1$  minimization. *Appl. Comput. Harmon. Anal.* 40 (2), 321–351.
- Roulet, V., d’Aspremont, A., 2020. Sharpness, restart, and acceleration. *SIAM J. Optim.* 30 (1), 262–289.
- Scarabosio, L., 2022. Deep neural network surrogates for nonsmooth quantities of interest in shape uncertainty quantification. *SIAM/ASA J. Uncertain. Quantificat.* 10 (3), 975–1011.
- Schmidt-Hieber, J., 2020. Nonparametric regression using deep neural networks with ReLU activation function. *Ann. Stat.* 48 (4), 1875–1897.
- Schwab, C., Stein, A., Zech, J., 2023. Deep operator network approximation rates for Lipschitz operators. arXiv:2307.09835.
- Schwab, C., Zech, J., 2019. Deep learning in high dimension: neural network expression rates for generalized polynomial chaos expansions in UQ. *Anal. Appl. (Singap.)* 17 (1), 19–55.
- Schwab, C., Zech, J., 2023. Deep learning in high dimension: neural network expression rates for analytic functions in  $L^2(\mathbb{R}^d, \gamma_d)$ . *SIAM/ASA J. Uncertain. Quantificat.* 11 (1), 199–234.
- Settles, B., 2012. *Active Learning. Synthesis Lectures on Artificial Intelligence and Machine Learning.* Springer, Cham, Switzerland.
- Shaham, U., Cloninger, A., Coifman, R.R., 2018. Provable approximation properties for deep neural networks. *Appl. Comput. Harmon. Anal.* 44, 537–557.
- Smith, R.C., 2013. *Uncertainty Quantification: Theory, Implementation, and Applications.* Comput. Sci. Eng. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Sullivan, T.J., 2015. *Introduction to Uncertainty Quantification.* Texts Appl. Math., vol. 63. Springer, Cham, Switzerland.
- Suzuki, T., 2019. Adaptivity of deep ReLU network for learning in Besov and mixed smooth Besov spaces: optimal rate and curse of dimensionality. In: *ICLR*.
- Tang, S., Li, B., Haijun, Y., 2019. ChebNet: efficient and stable constructions of deep neural networks with rectified power units via Chebyshev approximation. arXiv:1911.05467.
- Telgarsky, M., 2017. Neural networks and rational functions. In: *ICML*.
- Temlyakov, V.N., 2018. The Marcinkiewicz-type discretization theorems. *Constr. Approx.* 48 (2), 337–369.
- Todor, R.A., Schwab, C., 2007. Convergence rates for sparse chaos approximations of elliptic problems with stochastic coefficients. *IMA J. Numer. Anal.* 27 (2), 232–261.
- Tran, H., Webster, C.G., Zhang, G., 2017. Analysis of quasi-optimal polynomial approximations for parameterized PDEs with deterministic and stochastic coefficients. *Numer. Math.* 137 (2), 451–493.
- Traub, J.F., Woźniakowski, H., Wasilkowski, G.W., 1988. *Information-Based Complexity.* Elsevier Science and Technology Books.

- Trefethen, L.N., 2013. *Approximation Theory and Approximation Practice*. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Vidyasagar, M., 2019. *An Introduction to Compressed Sensing*. Comput. Sci. Eng. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Xiu, D., Hesthaven, J.S., 2005. High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.* 27 (3), 1118–1139.
- Yan, L., Guo, L., Xiu, D., 2012. Stochastic collocation algorithms using  $\ell_1$ -minimization. *Int. J. Uncertain. Quantificat.* 2 (3), 279–293.
- Yang, X., Karniadakis, G.E., 2013. Reweighted  $\ell_1$  minimization method for stochastic elliptic differential equations. *J. Comput. Phys.* 248, 87–108.
- Yarotsky, D., 2017. Error bounds for approximations with deep ReLU networks. *Neural Netw.* 94, 103–114.
- Yarotsky, D., 2018. Optimal approximation of continuous functions by very deep ReLU networks. In: Bubeck, Sébastien, Perchet, Vianney, Rigollet, Philippe (Eds.), *Proceedings of the 31st Conference on Learning Theory*. In: *Proceedings of Machine Learning Research*, vol. 75. PMLR, pp. 639–649.
- Zhang, G., Webster, C.G., Gunzburger, M., Burkardt, J., 2016. Hyperspherical sparse approximation techniques for high-dimensional discontinuity detection. *SIAM Rev.* 58 (3), 517–551.