

—
Ben Adcock, Simone Brugiapaglia,
Nick Dexter, Sebastian Moraga

**On Efficient Algorithms for
Computing Near-Best
Polynomial Approximations to
High-Dimensional,
Hilbert-Valued Functions from
Limited Samples**

EM
S ■
PRESS

Memoirs of the European Mathematical Society

Edited by

Anton Alekseev (Université de Genève)

Hélène Esnault (Freie Universität Berlin)

Gerard van der Geer (Universiteit van Amsterdam)

Ari Laptev (Imperial College London)

Laure Saint-Raymond (Institut des Hautes Études Scientifiques)

Susanna Terracini (Università degli Studi di Torino)

The *Memoirs of the European Mathematical Society* publish outstanding research contributions in individual volumes, in all areas of mathematics and with a particular focus on works that are longer and more comprehensive than usual research articles.

The collection's editorial board consists of the editors-in-chief of the *Journal of the European Mathematical Society* and the *EMS Surveys in Mathematical Sciences*, along with editors of book series of the publishing house of the EMS as well as other distinguished mathematicians.

All submitted works go through a highly selective peer-review process.

Previously published in this series:

J.-M. Delort, N. Masmoudi, *Long-Time Dispersive Estimates for Perturbations of a Kink Solution of One-Dimensional Cubic Wave Equations*

G. Cotti, *Cyclic Stratum of Frobenius Manifolds, Borel–Laplace (α, β) -Multitransforms, and Integral Representations of Solutions of Quantum Differential Equations*

A. Kostenko, N. Nicolussi, *Laplacians on Infinite Graphs*

A. Carey, F. Gesztesy, G. Levitina, R. Nichols, F. Sukochev, D. Zanin, *The Limiting Absorption Principle for Massless Dirac Operators, Properties of Spectral Shift Functions, and an Application to the Witten Index of Non-Fredholm Operators*

J. Kigami, *Conductive Homogeneity of Compact Metric Spaces and Construction of p -Energy*

A. Buium, L. E. Miller, *Purely Arithmetic PDEs Over a p -Adic Field: δ -Characters and δ -Modular Forms*

M. Duerinckx, A. Gloria, *On Einstein's Effective Viscosity Formula*

R. Willett, G. Yu, *The Universal Coefficient Theorem for C^* -Algebras with Finite Complexity*

B. Janssens, K.-H. Neeb, *Positive Energy Representations of Gauge Groups I. Localization*

S. Dipierro, G. Giacomini, E. Valdinoci, *The Lévy Flight Foraging Hypothesis in Bounded Regions*

A. Naor, *Extension, Separation and Isomorphic Reverse Isoperimetry*

N. Lerner, *Integrating the Wigner Distribution on Subsets of the Phase Space, a Survey*



Ben Adcock, Simone Brugiapaglia,
Nick Dexter, Sebastian Moraga

**On Efficient Algorithms for
Computing Near-Best
Polynomial Approximations to
High-Dimensional,
Hilbert-Valued Functions from
Limited Samples**



Authors

Ben Adcock
Department of Mathematics
Simon Fraser University
8888 University Drive
Burnaby BC, V5A 1S6 Canada
Email: ben_adcock@sfu.ca

Simone Brugiapaglia
Department of Mathematics and Statistics
Concordia University
J.W. McConnell Building
1400 De Maisonneuve Blvd. W
Montréal, QC, H3G 1M8 Canada
Email: simone.brugiapaglia@concordia.ca

Nick Dexter
Department of Mathematics
Simon Fraser University
8888 University Drive
Burnaby BC, V5A 1S6 Canada
Email: nicholas_dexter@sfu.ca

Sebastian Moraga
Department of Mathematics
Simon Fraser University
8888 University Drive
Burnaby BC, V5A 1S6 Canada
Email: sebastian_moraga_scheuermann@sfu.ca

Each volume of the *Memoirs of the European Mathematical Society* is available individually or as part of an annual subscription. It may be ordered from your bookseller, subscription agency, or directly from the publisher via subscriptions@ems.press.

ISSN 2747-9080, eISSN 2747-9099

ISBN 978-3-98547-070-9, eISBN 978-3-98547-570-4, DOI 10.4171/MEMS/13

© © The content of this volume is licensed under the CC BY 4.0 license, with the exception of the logo and branding of the European Mathematical Society and EMS Press, and where otherwise noted.

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>.

Published by EMS Press, an imprint of the

European Mathematical Society – EMS – Publishing House GmbH
Institut für Mathematik
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin, Germany

<https://ems.press>

© 2024 European Mathematical Society

Typesetting: AfricanType, Cairo, Egypt

Printed in Germany

© Printed on acid free paper

Abstract

Sparse polynomial approximation has become an indispensable technique for approximating smooth, high- or infinite-dimensional functions from limited samples. This is a key task in computational science and engineering – e.g., surrogate modeling uncertainty quantification, wherein the underlying function is the solution map of a parametric or stochastic Differential Equation (DE). Yet, sparse polynomial approximation lacks a complete theory. On the one hand, there is a well-developed theory of *best s -term polynomial approximation*, which asserts exponential or algebraic rates of convergence for holomorphic functions. On the other hand, there are increasingly mature methods such as (weighted) ℓ^1 -minimization for computing such approximations. While the sample complexity of these methods has been analyzed through compressed sensing theory, whether they achieve the rates of the best s -term approximation is not fully understood. Furthermore, these methods are not algorithms per se, since they involve exact minimizers of nonlinear optimization problems. This work closes these gaps. Specifically, we consider the following question: *are there robust, efficient algorithms for computing sparse polynomial approximations to finite- or infinite-dimensional, holomorphic and Hilbert-valued functions from limited samples that achieve the same rates as the best s -term approximation?* We answer this affirmatively by introducing algorithms with exponential or algebraic convergence rates that are also robust to *sampling, algorithmic and physical discretization* errors. We tackle both scalar- and Hilbert-valued functions, this being particularly relevant in parametric or stochastic DEs. Our results involve several significant developments of existing techniques, including a novel restarted primal-dual iteration for solving weighted ℓ^1 -minimization problems in Hilbert spaces. Our theory is supplemented by numerical experiments demonstrating the practical efficacy of these algorithms.

Mathematics Subject Classification (2020). Primary 65D15; Secondary 41A10

Keywords. High-dimensional approximation, polynomial approximation, best s -term approximation, compressed sensing, parametric PDEs

Funding. BA acknowledges the support of NSERC through grant RGPIN-2021-611675. SB acknowledges the support of NSERC through grant RGPIN-2020-06766, the Faculty of Arts and Science of Concordia University and the CRM Applied Math Lab. ND acknowledges support from a PIMS postdoctoral fellowship.

Contents

1	Introduction	1
1.1	High-dimensional function approximation from limited samples	1
1.2	Smoothness and best s -term polynomial approximation	2
1.3	Computing sparse polynomial approximations	3
1.4	Problem and main contributions	3
1.5	Discussion and further contributions	6
1.6	Related work	8
1.7	Outline	9
2	Preliminaries	11
2.1	Notation	11
2.2	Domains and function spaces	12
2.3	Holomorphy	13
2.4	Orthogonal polynomials, polynomial expansions and best s -term polynomial approximation	14
2.5	Sequence spaces and best s -term approximation of sequences	15
2.6	Rates of best s -term polynomial approximation	16
2.7	Lower and anchored sets	17
3	Problem statement and main results	19
3.1	Samples	19
3.2	Problem statement	20
3.3	Main results	21
4	Construction of the algorithms	31
4.1	Recovery via Hilbert-valued, weighted ℓ^1 -minimization	31
4.2	Reformulation as a matrix recovery problem and the mappings in Theorems 3.4, 3.7 and 3.10	33
4.3	The primal-dual iteration	35
4.4	The primal-dual iteration for the weighted SR-LASSO problem	37
4.5	The algorithms in Theorems 3.5, 3.8 and 3.11	41
4.6	An efficient restarting procedure for the primal-dual iteration and the algorithms used in Theorems 3.6, 3.9 and 3.12	42
5	Numerical experiments	47
5.1	Experimental setup	47

5.2	Numerical results 1: The optimization error	51
5.3	Numerical results 2: Approximation error and run time	54
6	Overview of the proofs	59
7	Hilbert-valued compressed sensing	61
7.1	Weighted sparsity and weighted best approximation	61
7.2	The weighted robust null space property	62
7.3	The weighted rNSP and weighted restricted isometry property	63
8	Error bounds for polynomial approximation via the Hilbert-valued, weighted SR-LASSO	67
8.1	The weighted RIP for the polynomial approximation problem	67
8.2	Bounds for polynomial approximations obtained as inexact minimizers	68
9	Error bounds and the restarting scheme for the primal-dual iteration	75
9.1	Error bounds for the primal-dual iteration	75
9.2	The restarting scheme	76
10	Final arguments	81
10.1	Algebraic rates of convergence, finite dimensions	81
10.2	Algebraic rates of convergence, infinite dimensions	85
10.3	Exponential rates of convergence, finite dimensions	86
11	Conclusions	89
A	Best polynomial approximation rates for holomorphic functions	91
A.1	The finite-dimensional case	91
A.2	The infinite-dimensional case	92
	References	95

Chapter 1

Introduction

A fundamental task in computational science and engineering involves accurately approximating a smooth target function from limited data. Such a task arises notably in the study of parametric models of physical processes. Here the variables represent the parameters in the system, e.g., material properties, forcing terms, or boundary information, and the parametric model is often represented as a (system of) Differential Equations (DEs) or Partial Differential Equations (PDEs) depending on these parameters. Important objectives involve understanding how the choice of such parameters affect the output(s) of the system and, in the stochastic setting, understanding how uncertainty in the parameter values propagates to its output – the latter being one of the key tasks in computational Uncertainty Quantification (UQ) [62, 91, 128, 131].

1.1 High-dimensional function approximation from limited samples

Abstractly, this task can be recast as that of approximating an unknown target function

$$f : \mathcal{U} \rightarrow \mathcal{V}, \quad y \mapsto f(y),$$

from *sample values* (or *snapshots*)

$$f(\mathbf{y}_1), \dots, f(\mathbf{y}_m). \tag{1.1}$$

Here, the input space \mathcal{U} is typically a subset of \mathbb{R}^d (in the finite-dimensional case) or $\mathbb{R}^{\mathbb{N}}$ (in the infinite-dimensional case). The output space \mathcal{V} could either be a scalar field, a finite-dimensional vector space or an infinite-dimensional Banach or Hilbert space.

This problem is challenging in a number of ways. First, the dimension d is high, since modern parametric models typically involve many parameters. It may also be infinite, e.g., in the case of a random field represented via its Karhunen–Loève expansion. Therefore, care must be taken to design methods that scale well with dimension. In addition, the amount of samples m is often highly limited. For example, in the parametric DE setting, each evaluation of f involves an expensive computational simulation. The data (1.1) is also always corrupted by errors, due to noise in physical experiments or numerical error in solving a DE. And finally, since the output $f(y)$ is often the solution of DE parametrized by the vector y , it may consequently take values in an infinite-dimensional Banach or Hilbert space. While it is commonplace to circumvent this issue in practice by considering scalar-valued *quantities of interest* (i.e., functions of the form $g(y) = \mathcal{Q}(f(y))$ for some known map $\mathcal{Q} : \mathcal{V} \rightarrow \mathbb{C}$),

approximating the full function f is both of theoretical interest and practical importance [52].

Remark 1.1. As a further consideration, we note that in many scenarios one may have substantial flexibility to choose the sample points $y_1, \dots, y_m \in \mathcal{U}$ in (1.1). However, in other scenarios they may be fixed, e.g., when dealing with legacy data. In this work, we consider Monte Carlo sampling – which may be considered either as a designed sampling strategy or a fixed one, depending on the setting. Here, the samples are drawn randomly and independently of an underlying probability measure on \mathcal{U} . This is very common in practice, in particular in UQ settings.

1.2 Smoothness and best s -term polynomial approximation

A key characteristic of parametric model problems is that the target function f is often smooth. There is now a large body of literature that has established that solution maps of a wide range of different parametric DEs are *holomorphic* (i.e., *analytic*) functions of their parameters. We mention in passing problems such as elliptic PDEs with affine and (certain) nonaffine parametric dependence, parabolic PDEs, PDEs over parametrized domains and shape uncertainty, parametric Initial Value Problems (IVPs), parametric hyperbolic problems and parametric control problems. Classical results in this area can be found in, e.g., [146]. For surveys of more recent results, we refer to [41] and [8] and references therein.

In tandem with the effort to establish holomorphic regularity of parametric DEs, there has also been a focus on applying polynomial methods, and in particular, *best s -term polynomial approximation* to construct finite approximations to such functions. In best s -term approximation, the function f is approximated by an s -term expansion corresponding to its largest s coefficients (measured in the \mathcal{V} -norm) with respect to a polynomial basis. Common choices include multivariate Taylor polynomials, tensor-product Legendre and Chebyshev polynomials on bounded hypercubes or tensor-product Hermite and Laguerre polynomials on \mathbb{R}^d or $[0, \infty)^d$, respectively. Over the last fifteen years, there have been significant developments in the approximation theory of such techniques (see Section 1.6). Signature results have established *exponential* and *algebraic* convergence rates for the best s -term approximation. The former assert that the error decays at least exponentially fast in $s^{1/d}$ in finite dimensions for any holomorphic function. The latter assert that the error decays algebraically fast; specifically, like $s^{1/2-1/p}$ for some $0 < p < 1$. These algebraic rates also hold in infinite dimensions, thus establishing best s -term approximation as a (theoretical) means to approximate holomorphic functions of infinitely many variables. We review several such results in Section 2.6.

1.3 Computing sparse polynomial approximations

Unfortunately, the best s -term approximation cannot usually be computed from the samples (1.1). Indeed, constructing it in theory involves computing and then searching over infinitely many coefficients. Both tasks are generally impossible. Therefore, there has also been a focus on methods to compute accurate polynomial approximations from sample values.

One line of work focuses on *least-squares* methods, wherein a polynomial approximation (or sequence of approximations) is computed in a fixed polynomial subspace (or sequence of nested subspaces). See Section 1.6 for relevant references. Such methods are essentially optimal if a (sequence of) polynomial subspace that gives a quasi-best s -term approximation is known.

However, this information is generally unavailable in practice (although it may be for certain simple parametric DEs). It essentially equates to knowing the region of holomorphy of the underlying function, which is itself similar to knowing the order of importance of the parametric variables, and their relative strengths. To counter this, there are *adaptive* least-squares methods [33, 35, 41, 45, 63, 101, 102]. Here one strives to construct such subspaces adaptively using the given data (1.1), typically via a greedy procedure. However, these methods currently lack theoretical guarantees [35, 41].

To overcome this limitation, there has also been a substantial focus on methods inspired by *compressed sensing* [13, 61, 145]. See Section 1.6 once more for relevant references. These methods seek a polynomial approximation in a larger subspace, whose coefficients are defined as a minimizer of an ℓ^1 - or weighted ℓ^1 -minimization problem. A key component of this endeavour has been to determine the *sample complexity* of such schemes, i.e., quantifying how many (typically Monte Carlo) samples m are sufficient to obtain an approximation with a certain guaranteed error bound, involving a (weighted) best approximation error plus a truncation error. Yet, precise rates of approximation (i.e., algebraic or exponential in m) have typically not been derived for these schemes in previous works. Another key limitation of past work is that such methods are not algorithms per se. Indeed, they consider exact minimizers of nonlinear optimization problems, which cannot be computed exactly in finitely many arithmetic operations.

1.4 Problem and main contributions

Least-squares and compressed sensing techniques are commonly applied to compute polynomial approximations to parametric and stochastic DEs. However, as explained above, there is a key gap between theory and practice. The theory of the best s -term approximation asserts the existence of polynomial approximations that attain specific

algebraic or exponential rates of convergence for arbitrary holomorphic functions. Yet, it is currently unknown whether similar rates in terms of the number of samples m can be obtained via an algorithm that computes a polynomial approximation from the samples (1.1) in finitely many arithmetic operations. The purpose of this work is to close this gap.

We now describe the problem considered in this work. Let $\mathcal{U} = [-1, 1]^d$, where $d \in \mathbb{N}$ or $d = \infty$, and \mathcal{V} be an arbitrary separable Hilbert space. Let ϱ be either the uniform or Chebyshev (arcsine) measure and consider the associated tensor-product Legendre or Chebyshev polynomials. Now let $f : \mathcal{U} \rightarrow \mathcal{V}$ be the unknown target function that we seek to approximate, draw m sample points $\mathbf{y}_1, \dots, \mathbf{y}_m$ i.i.d. from ϱ and let

$$d_i = f(\mathbf{y}_i) + n_i, \quad i = 1, \dots, m, \quad (1.2)$$

be m noisy samples of f . Then, informally stated, the problem we study in this work is the following: *devise algorithms that take (1.2) as input and compute the coefficients of a polynomial approximation \hat{f} to f with guarantees on both the computational complexity and the error $f - \hat{f}$.* Note that the formal problem statement involves several technicalities (in particular, the definition of an algorithm), so we defer it to Section 3.2.

Our main contributions are on the existence of such algorithms (see Tables 4.2 and 4.3 and Algorithms 2 and 5). Let $\|\cdot\|_{L^2_\varrho(\mathcal{U};\mathcal{V})}$ be the Lebesgue–Bochner norm. Then, in all cases, we establish an error bound of the form

$$\|f - \hat{f}\|_{L^2_\varrho(\mathcal{U};\mathcal{V})} \lesssim E_{\text{app}} + E_{\text{samp}} + E_{\text{disc}} + E_{\text{alg}}, \quad (1.3)$$

with probability at least $1 - \epsilon$ with respect to the (Monte Carlo) draw of the sample points \mathbf{y}_i . This bound provides a complete accounting for the main sources of error in the problem.

- E_{app} is a *polynomial approximation error* term. Depending on the specific setup, it decays algebraically (Theorems 3.4–3.9) or exponentially (Theorems 3.10–3.12) with respect to m (up to several log terms). For instance, in the infinite-dimensional setting (Theorems 3.7–3.9), this term is given by

$$E_{\text{app}} = C \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}, \quad L = \log(m) \cdot (\log^3(m) + \log(\epsilon^{-1})), \quad (1.4)$$

where $c_0 \geq 1$ is a universal constant, C is a constant depending on (the region of holomorphy of) f only, $p \in (0, 1)$ is a parameter determined by the region of holomorphy of f and $0 < \epsilon < 1$ is the failure probability of (1.3). It is completely equivalent to the corresponding algebraic decay rate (Theorem 2.5) for the best s -term approximation error, except with s replaced by $m/(c_0 L)$.

- E_{samp} is the *sampling error*. It is equal to

$$E_{\text{samp}} = \sqrt{\frac{1}{m} \sum_{i=1}^m \|n_i\|_{\mathcal{V}}^2},$$

i.e., the norm of the error in the samples (1.2). The presence of this term means that the algorithms are robust to errors in the samples.

- E_{disc} is the *physical discretization error*. This term accounts for the fact that an algorithm cannot work with (i.e., take as input, or perform computations in) \mathcal{V} when it is an infinite-dimensional Hilbert space. The algorithms (see Tables 4.2 and 4.3 and Algorithms 2 and 5) therefore work in a finite-dimensional discretization space $\mathcal{V}_h \subseteq \mathcal{V}$. This is a standard step in parametric DEs, where discretization is often performed via techniques such as the Finite Element Method (FEM). In this case, \mathcal{V}_h is a finite element space. The term E_{disc} quantifies the effect of this error. It is given by

$$E_{\text{disc}} = \|f - \mathcal{P}_h(f)\|_{L^\infty(\mathcal{U}; \mathcal{V})},$$

where $\mathcal{P}_h : \mathcal{V} \rightarrow \mathcal{V}_h$ is the orthogonal projection onto \mathcal{V}_h . In other words, the effect of working in \mathcal{V}_h instead of \mathcal{V} is determined by the error of the (pointwise) best approximation $\mathcal{P}_h(f)$ to f from \mathcal{V}_h . If \mathcal{V} has finite dimension, we may set

$$\mathcal{V}_h = \mathcal{V},$$

which implies that $E_{\text{disc}} = 0$ in this case.

- E_{alg} is the *algorithmic error*. It depends on the number of iterations t performed by the algorithm that computes the coefficients of the polynomial approximation \hat{f} . We construct one type of algorithm (see Table 4.2 and Algorithm 2) where this term is $\mathcal{O}(1/t)$ as $t \rightarrow \infty$. This decay is relatively slow, especially in the regime where E_{app} is exponentially small in m . However, we also present an *efficient* algorithm (Table 4.3 and Algorithm 5) for which this term decays exponentially fast in t (specifically, $\mathcal{O}(e^{-t})$ as $t \rightarrow \infty$), subject to an additional theoretical constraint. This constraint is seemingly an artefact of the proof. Our numerical experiments suggest it is unnecessary in practice.

We also determine the computational cost of the algorithms in all cases. Here, we draw two main conclusions.

- In the infinite-dimensional case (Theorems 3.8–3.9), the computational cost is *subexponential* in m . Specifically, after t iterations of the algorithm, it is

$$\mathcal{O}\left(t \cdot m^{1+(\alpha+1)\log(4m)/\log(2)}\right), \quad m \rightarrow \infty,$$

where $\alpha = 1$ (Legendre) or $\alpha = \log(3)/\log(4) \approx 0.79$ (Chebyshev).

- In the finite-dimensional, exponential setting (Theorems 3.11–3.12), the computational cost is *algebraic* in m for fixed d : namely,

$$\mathcal{O}(t \cdot m^{\alpha+2} (\log(m))^{(d-1)(\alpha+1)}), \quad m \rightarrow \infty,$$

for the same values of α .

Note that these computational cost estimates also depend polynomially on the dimension of the discretization space \mathcal{V}_h .

1.5 Discussion and further contributions

This work bridges a gap between the best s -term polynomial approximation theory and algorithms for computing such approximations from sample values. In particular, it asserts algebraic and exponential rates with respect to the number of samples m that are highly similar to those of the best approximation. In other words, polynomial approximations of holomorphic functions can be achieved in a *sample efficient* manner. Furthermore, they can be computed in supexponential or algebraic computational cost.

Our main results assume holomorphy of the underlying function in order to attain these rates. However, they require no *a priori* knowledge of the region of holomorphy. As discussed, if such information is available, then least-squares methods can be used more straightforwardly to compute an approximation. The holomorphy assumption is made in order to have concrete algebraic and exponential rates. However, our algorithms exist independently of this smoothness assumption. It would be possible to also provide rates for other classes of functions, e.g., those possessing finite orders of (mixed) smoothness. We use holomorphy as our assumption due to its strong connections with the theory of parametric DEs.

Our algorithms and analysis are based on compressed sensing theory and involve computing approximate minimizers of certain weighted ℓ^1 -minimization problems. Here we make several additional contributions.

- (i) We provide precise error rates for polynomial approximation via compressed sensing. As noted, most prior work on compressed sensing involves quantifying the sample complexity to obtain a certain (weighted) best approximation error. We impose the holomorphy assumption to obtain specific algebraic and exponential rates.
- (ii) Prior works consider polynomial approximations formed by exact minimizers of nonlinear optimization problems. We introduce novel, efficient algorithms to compute approximate minimizers in finite computational time (see also below).

- (iii) While these algorithms are motivated by the desire to have full theoretical guarantees, they are also completely practical. We present a series of numerical experiments demonstrating their practical efficacy. In fact, our experiments show that these algorithms work even better than our theoretical results suggest.
- (iv) Most prior works on compressed sensing (with the exception of [52]) focus on scalar-valued functions, e.g., quantities of interest of parametric DEs. We develop algorithms that work in the Hilbert-valued setting, and, crucially, provide error bounds that take into account physical discretization error (see above).

More precisely, we first formulate the approximation problem as the recovery of a finite, Hilbert-valued vector (i.e., an element of \mathcal{V}^N) via a so-called weighted, Square-Root LASSO (SR-LASSO) optimization problem. The use of the SR-LASSO, as opposed to the classical LASSO or various constrained formulations, is crucial to this work. It is *noise-blind*. Hence, it allows us to devise algorithms that do not require any a priori (and generally unavailable) estimates on the measurement error n_i in (1.2) or the truncation error with respect to the finite polynomial space in which the approximation is constructed.

To develop algorithms, we employ two key ideas. First, we use a powerful, general-purpose first-order optimization method for solving nonsmooth, convex optimization problems. Second, we use the technique of *restarts* to drastically accelerate its convergence. For the former, we employ the *primal-dual iteration* (also known as the Chambolle–Pock algorithm) [30, 31]. We present error bounds for this method for solving the Hilbert-valued, weighted SR-LASSO, which decay like $\mathcal{O}(1/t)$, where t is the iteration number. Next, we use a novel restarting procedure, recently introduced in [47, 48], to obtain faster, exponential decay of the form $\mathcal{O}(e^{-t})$.

To the best of our knowledge, this is the first time either the primal-dual iteration or a restarting scheme has been applied to the problem of sparse polynomial approximation. Many existing works use blackbox solvers such as SPGL1 [143, 144]. See [52] for a forward-backwards splitting technique in combination with Bregman iterations and fixed-point continuation and [142] for an approach based on Douglas–Rachford splitting. Besides its amenability to theoretical analysis, the primal-dual scheme is also particularly attractive because of its insensitivity to parameter choices and the possibility of performing acceleration via restarts.

As noted, polynomial-based methods have become popular tools for the practical approximation high-dimensional, holomorphic functions arising in problems in computational science and engineering. However, they are by no means the only method. Other popular techniques include Gaussian processes (also known as kriging) [128, 131], radial basis methods [84, 128], reduced-order methods [78, 118] and,

recently, methods based on deep neural networks and deep learning [5, 6, 11, 49, 59, 60, 76, 92, 114, 115, 125]. Our goal in this work is to develop algorithms for constructing polynomial approximations that achieve the same rates as the theoretical benchmark provided by the best s -term polynomial approximation. An important consideration that we do not address in this work is *tractability* and the *information complexity* [111, 113] of these classes of functions and, in particular, whether polynomial-based methods constitute *optimal algorithms*. This question has been studied in the infinite-dimensional case in recent work [12]. Here, it is shown that the rate $m^{1/2-1/p}$ is a lower bound for the (*adaptive*) m -width for such classes, i.e., no combination of m (*adaptive*) linear samples and a (potentially nonlinear) reconstruction map can achieve an approximation error decaying faster than this rate. Notice that this rate is the same, up to constants and logarithmic factors, as (1.4). Unfortunately, this does not imply our algorithms are near optimal for this problem – and, moreover, that *standard information*, i.e., pointwise samples, constitutes near-optimal information – because our theoretical results in the infinite-dimensional case are *nonuniform*. See Remark 3.13 for further discussion on this point, and Chapter 11 for further comments on tractability.

1.6 Related work

The systematic study of best s -term polynomial approximation of high- or infinite-dimensional holomorphic functions began around 2010 with the works of [25, 42, 43, 75, 139]. For reviews, see [41] and [8, Chapter 3]. Note that many of these works assume the function is a solution of a parametric PDE, and therefore first demonstrate that such a function is holomorphic. However, other works avoid this step and use specific properties of the DE to obtain refined estimates. See, e.g., [19, 20] for results of this type. Other recent works such as [8, 27] also study the problem without assuming the function is a solution of a parametric PDE.

The study of least-squares method for constructing such approximations from sample points began in the early 2010s [34, 40, 100, 105]. There has since been significant research on this topic. Many subsequent works have pursued extensions, such as enhanced sampling strategies [65, 104, 107, 126, 135, 155, 156], near-optimal sampling strategies [9, 44, 71], optimal sampling strategies [21, 54, 56, 85, 93, 137], methods for general domains [14, 55, 103], optimal and adaptive methods [46, 101, 102] and multilevel strategies [70]. See [45, 67, 69] and [8, Chapter 5] for reviews.

Compressed sensing was introduced in the context of image and signal processing by modeling image and signals as sparse vectors [13, 29, 57, 61]. Its use in polynomial approximation started early in the last decade with the works of [26, 58, 99, 120, 149]. This has also led to substantial research. See [51, 52, 58, 99, 119, 151] and references therein for applications to parametric PDEs. Various extensions include refined sampling strategies [17, 53, 68, 72, 83, 94, 134], iterative methods and basis selection

techniques [16, 74, 142, 152, 152–154], nonconvex optimization methods [64, 140, 148, 150], sublinear-time algorithms [38, 39], gradient-enhanced minimization techniques [15, 66, 82, 117, 130, 133], methods for dealing with corrupted samples [3, 7, 80, 127] and multilevel and multifidelity strategies [28, 109]. For additional information and reviews, see [73, 86, 97, 98, 108] and [8, Chapter 7].

Our work combines and extends several key elements of this literature. First, weighted ℓ^1 -minimization, which was developed in [1–4, 37, 116, 121, 151] and [8, Chapter 6–7]. Second, the notions of *lower* and *anchored* sets (see Section 2.7). These have been extensively studied in the best s -term polynomial approximation literature. Compressed sensing techniques aiming to exploit such structures were first considered in [2, 3, 37] and [8, Chapter 7]. Third, the extension of classical compressed sensing theory from vectors in \mathbb{R}^N (or \mathbb{C}^N) to Hilbert-valued vectors in \mathcal{V}^N . This was first developed in [52]. In order to prove our main results, we also extend this work to the weighted setting.

See [30, 31, 31] for more on the primal-dual iteration and [122–124] for the general notion of restarts in continuous optimization. Note that there are also various non-optimization based techniques in the compressed sensing literature (see, e.g., [61]), including iterative thresholding and greedy methods. The latter are closely related to the adaptive least-squares methods discussed earlier [8, Section 6.2.5]. However, such techniques currently do not possess theoretical guarantees in the weighted setting.

There have been several previous attempts to connect compressed sensing theory for analyzing the sample complexity of polynomial approximations via (weighted) ℓ^1 -minimization and best s -term polynomial approximation theory. In [119], the authors consider approximating scalar quantities of interest of solutions to affine parametric operator equations in Banach spaces. Assuming a certain weighted summability criterion, they first show holomorphy of the parametric solution map and then use a weighted ℓ^1 -minimization procedure in combination with Chebyshev polynomials to derive algebraic rates of convergence, similar to (1.4). Our work is more general, since its starting point is a holomorphic function, not a solution of a parametric operator equation. We also consider Hilbert-valued functions, i.e., the whole solution map, not a scalar quantity of interest of it. Moreover, the work of [119] is based on exact minimizers of certain constrained, weighted ℓ^1 -minimization problems, whereas we construct full algorithms. Recently, at the same time as writing this work, some similar results were presented in the book [8] written by two of the authors. However, these only consider the scalar-valued case and do not address algorithms, which is the main focus of this work.

1.7 Outline

The remainder of this work proceeds as follows. We commence in Chapter 2 by introducing preliminaries, including key notation and best s -term polynomial

approximation theory. Next, in Chapter 3 we first formally define the problem and then state our main results on the existence of algorithms. In Chapter 4 we derive these algorithms. Then in Chapter 5 we present numerical experiments demonstrating their practical performance. Chapters 6–10 are devoted to the proofs of the main results. See Chapter 6 for a detailed overview of these chapters. Finally, in Chapter 11 we present our conclusions.

Chapter 2

Preliminaries

In this chapter, we introduce key preliminary material needed later. After some initial notation, we define the domains (the symmetric hypercubes), probability measures (the uniform and Chebyshev measures, respectively) and the Lebesgue–Bochner spaces. We next formalize our main smoothness assumption: namely, holomorphy in suitable (unions of) Bernstein polyellipses. We then introduce orthogonal polynomial expansions and best s -term polynomials approximations, before discussing sequence spaces and best s -term approximations of sequences. Finally, we conclude by reviewing algebraic and exponential rates of convergence for best s -term polynomial approximations, before a short discussion on lower and anchored sets.

2.1 Notation

We first introduce some notation. For $d \in \mathbb{N}$, we write $[d] = \{1, \dots, d\}$. We also extend this to allow for $d = \infty$, in which case $[d] = \mathbb{N}$ is the set of positive integers. For $d \in \mathbb{N} \cup \{\infty\}$, we write \mathbf{e}_j , $j \in [d]$, for the standard basis vectors, i.e., $\mathbf{e}_j = (\delta_{jk})_{k \in [d]}$. Also for $d \in \mathbb{N} \cup \{\infty\}$, we write \mathbb{R}^d or \mathbb{C}^d for the vector space of real or complex vectors of length d . Note that when $d = \infty$, \mathbb{R}^d and \mathbb{C}^d are the vector spaces $\mathbb{R}^{\mathbb{N}}$ and $\mathbb{C}^{\mathbb{N}}$ of real- or complex-valued sequences indexed over \mathbb{N} .

For $1 \leq p \leq \infty$, we write $\|\cdot\|_p$ for the usual vector ℓ^p -norm and for the induced matrix ℓ^p -norm. When $0 < p < 1$, we use the same notation to denote the ℓ^p -quasinorm. For $1 \leq p, q < \infty$ we define the matrix $\ell^{p,q}$ -norm of an $m \times n$ matrix $\mathbf{G} = (G_{ij})_{i,j=1}^{m,n}$ as $\|\mathbf{G}\|_{p,q}^q := \sum_{j=1}^n (\sum_{i=1}^m |G_{ij}|^p)^{q/p}$, and similarly when $p = \infty$ or $q = \infty$.

Throughout this work, we consider sets of multi-indices. Let $d \in \mathbb{N}$. Then we define the multi-index set \mathcal{F} as the set of nonnegative multi-indices, i.e.,

$$\mathcal{F} := \mathbb{N}_0^d = \{\mathbf{v} = (v_k)_{k=1}^d : v_k \in \mathbb{N}_0\}, \quad d < \infty. \quad (2.1)$$

When $d = \infty$, we consider multi-indices in $\mathbb{N}_0^{\mathbb{N}}$ with at most finitely many nonzero terms, i.e., we define

$$\mathcal{F} := \{\mathbf{v} = (v_k)_{k=1}^{\infty} \in \mathbb{N}_0^{\mathbb{N}} : |\{k : v_k \neq 0\}| < \infty\}, \quad d = \infty. \quad (2.2)$$

In either finite or infinite dimensions, we write $\mathbf{0}$ and $\mathbf{1}$ for the multi-indices consisting of all zeros and all ones, respectively. Finally, the inequality $\boldsymbol{\mu} \leq \mathbf{v}$ is understood componentwise for any multi-indices $\boldsymbol{\mu}$ and \mathbf{v} .

2.2 Domains and function spaces

Let $\varrho = \varrho^{(1)}$ be a probability measure on $[-1, 1]$. In this work, we focus on two main examples, the uniform and Chebyshev (arcsine) measures. These are defined by

$$d\varrho(y) = 2^{-1} dy, \quad \text{and} \quad d\varrho(y) = \frac{1}{\pi \sqrt{1-y^2}} dy, \quad y \in \mathcal{U}, \quad (2.3)$$

respectively. See Chapter 11 for a short discussion on other domains and measures. In finite dimensions, we let $\mathcal{U} = [-1, 1]^d$ be the symmetric d -dimensional hypercube and write $\mathbf{y} = (y_1, \dots, y_d) \in \mathcal{U}$ for the variable in this domain. We define a probability measure on \mathcal{U} as the product measure

$$\varrho = \varrho^{(d)} := \varrho^{(1)} \times \dots \times \varrho^{(1)}.$$

In particular, the d -dimensional uniform and Chebyshev measures are given by

$$d\varrho(\mathbf{y}) = 2^{-d} d\mathbf{y} \quad \text{and} \quad d\varrho(\mathbf{y}) = \prod_{k=1}^d \frac{1}{\pi \sqrt{1-y_k^2}} d\mathbf{y}, \quad \forall \mathbf{y} \in \mathcal{U},$$

respectively. In infinite dimensions, we consider the domain $\mathcal{U} = [-1, 1]^{\mathbb{N}}$ and write $\mathbf{y} = (y_1, y_2, \dots) \in \mathcal{U}$ for the variable in this domain. The Kolmogorov extension theorem (see, e.g., [136, Section 2.4]) guarantees the existence of a tensor-product probability measure on \mathcal{U} , which we denote as

$$\varrho = \varrho^{(\infty)} = \prod_{k \in \mathbb{N}} \varrho^{(1)}.$$

In either finite or infinite dimensions, for $1 \leq p \leq \infty$ we write $L_{\varrho}^p(\mathcal{U})$ for the corresponding weighted Lebesgue spaces of complex scalar-valued functions over \mathcal{U} and $\|\cdot\|_{L_{\varrho}^p(\mathcal{U})}$ for their norms.

Throughout, we let \mathcal{V} be a separable Hilbert space over \mathbb{C} (it presents few difficulties to consider a complex field instead of the real field). We write $\langle \cdot, \cdot \rangle_{\mathcal{V}}$ and $\|\cdot\|_{\mathcal{V}}$ for its inner product and norm. We define the weighted (Lebesgue-)Bochner space $L_{\varrho}^p(\mathcal{U}; \mathcal{V})$ as the space consisting of (equivalence classes of) strongly ϱ -measurable functions $f : \mathcal{U} \rightarrow \mathcal{V}$ for which $\|f\|_{L_{\varrho}^p(\mathcal{U}; \mathcal{V})} < \infty$, where

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

Note that $L_{\varrho}^p(\mathcal{U})$ is a special case of $L_{\varrho}^p(\mathcal{U}; \mathcal{V})$ corresponding to $\mathcal{V} = (\mathbb{C}, |\cdot|)$.

When \mathcal{V} is infinite dimensional, we usually cannot work directly with it. Hence, we consider a finite-dimensional discretization

$$\mathcal{V}_h \subseteq \mathcal{V}. \quad (2.4)$$

Here $h > 0$ denotes a discretization parameter, e.g., the mesh size in the case of a finite element discretization (as is common in parametric DEs). In the context of finite elements, assuming (2.4) corresponds to considering so-called *conforming* discretizations. We let $\{\varphi_k\}_{k=1}^K$ be a (not necessarily orthonormal) basis of \mathcal{V}_h , where $K = K(h) = \dim(\mathcal{V}_h)$. We write $\mathcal{P}_h : \mathcal{V} \rightarrow \mathcal{V}_h$ for the orthogonal projection onto \mathcal{V}_h and, for $f \in L^2_{\mathcal{O}}(\mathcal{U}; \mathcal{V})$, we let $\mathcal{P}_h f \in L^2_{\mathcal{O}}(\mathcal{U}; \mathcal{V}_h)$ be the function defined almost everywhere as

$$(\mathcal{P}_h f)(y) = \mathcal{P}_h(f(y)), \quad y \in \mathcal{U}. \quad (2.5)$$

2.3 Holomorphy

Here we recall the definitions of holomorphy and holomorphic extension for Hilbert-valued functions. We note that equivalent definitions are possible (see, e.g., [77, Chapter 2]) and that the definition employed in this work is based on the notion of the Gateaux partial derivative. For other details on differentiability of Hilbert-valued functions we refer to [22, Chapter 17], and the references therein. Note the following definitions apply in both the finite- ($d \in \mathbb{N}$) and infinite- ($d = \infty$) dimensional settings, where we recall that $[d] = \mathbb{N}$ and $\mathbb{C}^d = \mathbb{C}^{\mathbb{N}}$ when $d = \infty$.

Definition 2.1 (Holomorphy; finite- or infinite-dimensional case). Let $d \in \mathbb{N} \cup \{\infty\}$, $\mathcal{O} \subseteq \mathbb{C}^d$ be an open set and \mathcal{V} be a separable Hilbert 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 $z \in \mathcal{O}$ and any $j \in [d]$, the following limit exists in \mathcal{V} :

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

Let $f : \mathcal{U} \rightarrow \mathcal{V}$ and $\mathcal{U} \subset \mathcal{O} \subseteq \mathbb{C}^d$ be an open set. If there is a function $\tilde{f} : \mathcal{O} \rightarrow \mathcal{V}$ that is holomorphic in \mathcal{O} and for which $\tilde{f}|_{\mathcal{U}} = f$, then we say that f has a *holomorphic extension to \mathcal{O}* , or simply, that f is *holomorphic in \mathcal{O}* . In this case, we also define $\|f\|_{L^\infty(\mathcal{O}; \mathcal{V})} := \|\tilde{f}\|_{L^\infty(\mathcal{O}; \mathcal{V})}$ 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} .

We are interested in approximating Hilbert-valued functions $f : \mathcal{U} \rightarrow \mathcal{V}$ that are holomorphic in suitable complex regions containing \mathcal{U} – specifically, regions defined by *Bernstein (poly)ellipses*. When $d = 1$ the Bernstein ellipse of parameter $\rho > 1$ is given by

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

This is an ellipse with ± 1 as its foci and major and minor semi-axis lengths $\frac{1}{2}(\rho \pm \rho^{-1})$. For $d \in \mathbb{N} \cup \{\infty\}$, given $\boldsymbol{\rho} = (\rho_j)_{j=1}^d \in \mathbb{R}^d$ with $\boldsymbol{\rho} > \mathbf{1}$, we define

the Bernstein polyellipse as the Cartesian product

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

We denote the class of Hilbert-valued functions that are holomorphic in $\mathcal{E}(\boldsymbol{\rho})$ with norm at most one as

$$\mathcal{B}(\boldsymbol{\rho}) = \{f : \mathcal{U} \rightarrow \mathcal{V}, f \text{ holomorphic in } \mathcal{E}(\boldsymbol{\rho}), \|f\|_{L^\infty(\mathcal{E}(\boldsymbol{\rho}); \mathcal{V})} \leq 1\}. \quad (2.6)$$

In infinite dimensions, we also consider a class of functions that are holomorphic in a certain union of Bernstein polyellipses. Let $0 < p < 1$, $\varepsilon > 0$ and $\mathbf{b} = (b_j)_{j \in \mathbb{N}} \in \ell^p(\mathbb{N})$. We define

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

In analogy with $\mathcal{B}(\boldsymbol{\rho})$, we write

$$\mathcal{B}(\mathbf{b}, \varepsilon) = \{f : \mathcal{U} \rightarrow \mathcal{V}, f \text{ holomorphic in } \mathcal{R}(\mathbf{b}, \varepsilon), \|f\|_{L^\infty(\mathcal{R}(\mathbf{b}, \varepsilon); \mathcal{V})} \leq 1\} \quad (2.7)$$

for the corresponding space of functions that are holomorphic in $\mathcal{R}(\mathbf{b}, \varepsilon)$ with norm at most one.

2.4 Orthogonal polynomials, polynomial expansions and best s -term polynomial approximation

Under mild assumptions on $\varrho^{(1)}$ (see, e.g., [106, Section 2.1] or [132, Section 2.2]), there exists a unique orthonormal polynomial basis $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0}$ of $L^2_{\varrho}([-1, 1])$, where $\Psi_{\mathbf{v}} = \Psi_{\mathbf{v}}^{(1)}$ is a polynomial of degree \mathbf{v} . For the measures (2.3), these are the Legendre and Chebyshev polynomials, respectively. Given the corresponding tensor-product measure ϱ on $\mathcal{U} = [-1, 1]^d$, we construct an orthonormal basis

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

of $L^2_{\varrho}(\mathcal{U})$ via tensorization

$$\Psi_{\mathbf{v}}(\mathbf{y}) = \prod_{k \in [d]} \Psi_{v_k}(y_k), \quad \mathbf{y} \in \mathcal{U}, \mathbf{v} \in \mathcal{F}.$$

Note that $\Psi_{\mathbf{0}}^{(1)} = 1$ since $\varrho^{(1)}$ is a probability measure. Therefore, since $\mathbf{v} \in \mathcal{F}$ has only finitely many nonzero entries, in infinite dimensions this equivalent to

$$\Psi_{\mathbf{v}}(\mathbf{y}) = \prod_{k: v_k \neq 0} \Psi_{v_k}(y_k),$$

which is a product of finitely many terms.

Let $f \in L^2_\varrho(\mathcal{U}; \mathcal{V})$. Then it has the convergent expansion (in $L^2_\varrho(\mathcal{U}; \mathcal{V})$) given by

$$f = \sum_{\mathbf{v} \in \mathcal{F}} c_{\mathbf{v}} \Psi_{\mathbf{v}}, \quad c_{\mathbf{v}} := \int_{\mathcal{U}} f(\mathbf{y}) \Psi_{\mathbf{v}}(\mathbf{y}) \, d\varrho(\mathbf{y}) \in \mathcal{V}, \quad (2.8)$$

where the *coefficients* $c_{\mathbf{v}}$ are elements of \mathcal{V} . Now let $S \subset \mathcal{F}$ be a finite index set and

$$\mathcal{P}_{S; \mathcal{V}} = \left\{ \sum_{\mathbf{v} \in S} c_{\mathbf{v}} \Psi_{\mathbf{v}} : c_{\mathbf{v}} \in \mathcal{V} \right\} \subset L^2_\varrho(\mathcal{U}; \mathcal{V}). \quad (2.9)$$

Then the $L^2(\mathcal{U}; \mathcal{V})$ -norm *best s -term polynomial approximation* f_s of f is defined as

$$f_s \in \operatorname{argmin} \{ \|f - g\|_{L^2_\varrho(\mathcal{U}; \mathcal{V})} : g \in \mathcal{P}_{S; \mathcal{V}}, S \subset \mathcal{F}, |S| = s \}. \quad (2.10)$$

Note that f_s has the explicit expression

$$f_s = \sum_{\mathbf{v} \in S^*} c_{\mathbf{v}} \Psi_{\mathbf{v}},$$

where $S^* \subset \mathcal{F}$, $|S^*| = s$, is a set of consisting of the multi-indices of the largest s values of the coefficient norms $(\|c_{\mathbf{v}}\|_{\mathcal{V}})_{\mathbf{v} \in \mathbb{N}_0^d}$. By Parseval's identity, the error of this approximation satisfies

$$\|f - f_s\|_{L^2_\varrho(\mathcal{U}; \mathcal{V})} = \sqrt{\sum_{\mathbf{v} \notin S^*} \|c_{\mathbf{v}}\|_{\mathcal{V}}^2}. \quad (2.11)$$

2.5 Sequence spaces and best s -term approximation of sequences

The expression (2.11) motivates studying s -term approximation of sequences of polynomial coefficients. To do this, we now introduce some further notation.

Let $\Lambda \subseteq \mathcal{F}$ denote a (possibly infinite) multi-index set. We write $\mathbf{v} = (v_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$ for a sequence with \mathcal{V} -valued entries $v_{\mathbf{v}} \in \mathcal{V}$. For $1 \leq p \leq \infty$, we define the space $\ell^p(\Lambda; \mathcal{V})$ as the set of those sequences $\mathbf{v} = (v_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$ for which $\|\mathbf{v}\|_{p; \mathcal{V}} < \infty$, where

$$\|\mathbf{v}\|_{p; \mathcal{V}} := \begin{cases} (\sum_{\mathbf{v} \in \Lambda} \|v_{\mathbf{v}}\|_{\mathcal{V}}^p)^{1/p} & 1 \leq p < \infty, \\ \sup_{\mathbf{v} \in \Lambda} \|v_{\mathbf{v}}\|_{\mathcal{V}} & p = \infty. \end{cases}$$

Note that $\ell^2(\Lambda; \mathcal{V})$ is a Hilbert space with inner product

$$\langle \mathbf{u}, \mathbf{v} \rangle_{2; \mathcal{V}} = \sum_{\mathbf{v} \in \Lambda} \langle u_{\mathbf{v}}, v_{\mathbf{v}} \rangle_{\mathcal{V}}.$$

On occasion, we will consider complex, scalar-valued sequences. In this case, $\mathcal{V} = (\mathbb{C}, |\cdot|)$ in the various definitions above. For ease of notation, we simply write $\ell^p(\Lambda)$, $\|\cdot\|_p$, $\langle \cdot, \cdot \rangle_2$ and so forth in this case.

Definition 2.2 (Sparsity). Let $\Lambda \subseteq \mathcal{F}$ and $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$ be a \mathcal{V} -valued sequence. The *support* of \mathbf{c} is the set

$$\text{supp}(\mathbf{c}) = \{\mathbf{v} \in \Lambda : \|c_{\mathbf{v}}\|_{\mathcal{V}} \neq 0\}.$$

A sequence is *s-sparse* for some $s \in \mathbb{N}_0$ satisfying $s \leq |\Lambda|$ if it has at most s nonzero entries, i.e.,

$$|\text{supp}(\mathbf{c})| \leq s.$$

Definition 2.3 (Best s -term approximation error). 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 ℓ^p -norm best s -term approximation error of \mathbf{c} is

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

Let $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \mathcal{F}}$ be the coefficients of some function $f \in L^2_{\rho}(\mathcal{U}; \mathcal{V})$, as defined in (2.8). Then, when $p = 2$, we have the following:

$$\sigma_s(\mathbf{c})_{2; \mathcal{V}} = \|f - f_s\|_{L^2_{\rho}(\mathcal{U}; \mathcal{V})},$$

where f_s is its best s -term polynomial approximation (2.10). Therefore, we can study the error of f_s by studying the quantity $\sigma_s(\mathbf{c})_{2; \mathcal{V}}$. For notational purposes, we denote this quantity in terms of the coefficients \mathbf{c} . However, on some occasions, this term is expressed as $\sigma_s(f)_{2; \mathcal{V}}$ instead.

2.6 Rates of best s -term polynomial approximation

As noted, best s -term polynomial approximation of holomorphic functions is a well-studied subject, especially in the context of solutions of parametric DEs. See, e.g., [23–25, 27, 36, 42, 43, 75, 115, 139, 141] and, in particular, [41] and [8, Chapter 3]. In this section, we recap two standard types of error decay rates for this approximation, those of *algebraic* and *exponential* type, respectively. Note that these results are for Chebyshev and Legendre polynomial approximations – the main focus of the work. The latter type of decay rate holds in finite dimensions, while the former holds in both finite and infinite dimensions. In this work, these error decay rates serve as the optimal benchmark against which to compare the approximations computed from sample values.

The following two results are standard, and have appeared in various different guises in the aforementioned works.

Theorem 2.4 (Algebraic rates of convergence; finite-dimensional case). *Let $0 < p \leq 1$ and $f \in \mathcal{B}(\boldsymbol{\rho})$ for some $\boldsymbol{\rho} > \mathbf{1}$. Let $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \mathbb{N}_0^d}$ be as in (2.8). Then, for every $s \geq 1$ there are sets $S_1, S_2 \subset \mathcal{F}$, $|S_1|, |S_2| \leq s$, such that*

$$\|f - f_{S_1}\|_{L^2_{\rho}(\mathcal{U}; \mathcal{V})} \leq C \cdot s^{1/2-1/p}, \quad \|f - f_{S_2}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \leq C \cdot s^{1-1/p}, \quad (2.12)$$

where $f_{S_i} = \sum_{\mathbf{v} \in S_i} c_{\mathbf{v}} \Psi_{\mathbf{v}}$ for $i = 1, 2$ and $C = C(d, p, \rho) > 0$ depends on d , p and ρ only.

Theorem 2.5 (Algebraic rates of convergence; infinite-dimensional case). *Let $0 < p < 1$, $\varepsilon > 0$, $\mathbf{b} = (b_j)_{j \in \mathbb{N}} \in \ell^p(\mathbb{N})$ and $f \in \mathcal{B}(\mathbf{b}, \varepsilon)$, where $\mathcal{B}(\mathbf{b}, \varepsilon)$ is as in (2.7). Then, for every $s \geq 1$ there are sets $S_1, S_2 \subset \mathcal{F}$, $|S_1|, |S_2| \leq s$, such that*

$$\|f - f_{S_1}\|_{L^2_{\theta}(\mathbf{u}; \mathbf{v})} \leq C \cdot s^{1/2-1/p}, \quad \|f - f_{S_2}\|_{L^{\infty}(\mathbf{u}; \mathbf{v})} \leq C \cdot s^{-1/p}, \quad (2.13)$$

where $f_{S_i} = \sum_{\mathbf{v} \in S_i} c_{\mathbf{v}} \Psi_{\mathbf{v}}$ for $i = 1, 2$ and $C = C(\mathbf{b}, \varepsilon, p) > 0$ depends on \mathbf{b} , ε and p only.

Observe that the curse of dimensionality is not avoided in the constant $C(d, p, \rho)$ in (2.12), but it is avoided in the rate. Conversely, (2.13) holds in infinite dimensions.

We next state a result on exponential convergence in finite dimensions. Such rates have been established in various different works (see, e.g., [23, 24, 41, 115, 141]). The following result is a minor modification of [8, Theorem 3.25], in which we allow arbitrary $s \geq 1$ at the expense of a constant C in the error bound.

Theorem 2.6 (Exponential rates of convergence; finite-dimensional case). *Let $f \in \mathcal{B}(\rho)$ for some $\rho > 1$ and $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \mathbb{N}_0^d}$ be as in (2.8). Then, for every $s \geq 1$ there is a set $S \subset \mathcal{F}$, $|S| \leq s$, such that*

$$\|f - f_S\|_{L^2_{\theta}(\mathbf{u}; \mathbf{v})} \leq \|f - f_S\|_{L^{\infty}(\mathbf{u}; \mathbf{v})} \leq C \cdot \exp(-\gamma s^{1/d}), \quad (2.14)$$

for all

$$0 < \gamma < (d+1)^{-1} \left(d! \prod_{j=1}^d \ln(\rho_j) \right)^{1/d}, \quad (2.15)$$

where $f_S = \sum_{\mathbf{v} \in S} c_{\mathbf{v}} \Psi_{\mathbf{v}}$ and $C = C(d, \gamma, p, \rho) > 0$ is a constant depending on d , γ , p and ρ only.

In Appendix A we show how these three theorems can be obtained as immediate consequences of several more general results.

Remark 2.7. It is possible to improve the rate (2.14) by removing the $(d+1)^{-1}$ factor in (2.15) [141]. The difficulty in doing this is that such rates are not necessarily attained in lower sets (this is, however, true if ρ is sufficiently large – see [8, Lemma 7.20]). As we discuss next, lower sets are a crucial ingredient in our analysis. Fortunately, the rates described in Theorem 2.6 can always be attained in lower sets.

2.7 Lower and anchored sets

Our objective in this work is to construct a polynomial approximation that attains error bounds that are similar to those of the best s -term approximation f_s , for any

holomorphic function f . Hence, ideally, we would have access to the multi-index set S corresponding to the largest s coefficients of f (measured in the \mathcal{V} -norm). As discussed, this is not possible in general, since the only information we have about f is its values at a finite number of sample points. Another problem is that such coefficients could occur at arbitrarily large multi-indices, thus necessitating a search over infinitely many multi-indices. Fortunately, it is well known that near-best s -term polynomial approximations can be constructed using sets of multi-indices with additional structure. These are *lower* sets (used in the finite-dimensional case) and *anchored* sets (used in the infinite-dimensional case). Classical references for lower and anchored sets include [50, 90, 96, 138]. More recently, these structures have been used extensively in the construction of interpolation, least-squares and compressed sensing schemes for polynomial approximation with desirable sample complexity bounds (see, e.g., [8] and references therein).

Definition 2.8. A set $\Lambda \subseteq \mathcal{F}$ is *lower* if the following holds for every $\mathbf{v}, \boldsymbol{\mu} \in \mathcal{F}$:

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

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

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

Lower sets are typically used in finite-dimensional settings, with anchored sets being employed in infinite dimensions. They are key concepts that we exploit in this work. To underscore their usefulness, we remark in passing that the rates articulated in not just Theorem 2.6, but also Theorems 2.4 and 2.5, can all be attained using s -term approximations in lower or anchored sets, subject to some modifications. See Appendix A.

Chapter 3

Problem statement and main results

In this chapter, we first formally define the problem we aim to solve before stating our main results. This work concerns algorithms for computing approximation of Hilbert-valued functions from finitely many sample values. We define this concept formally in a moment. For now, though, we consider that an algorithm must take a finite input and produce a finite output. Hence, in order to discuss algorithms, we first need to define what these finite inputs and outputs are in our setting.

3.1 Samples

Let $f \in L^2_{\varrho}(\mathcal{U}; \mathcal{V})$ be the function we seek to approximate. Throughout this work, we consider m *sample points* $\mathbf{y}_1, \dots, \mathbf{y}_m \in \mathcal{U}$ drawn randomly and independently according to the probability measure ϱ . Corresponding to each sample point, we consider the noisy *sample values*

$$d_i = f(\mathbf{y}_i) + n_i \in \mathcal{V}_h, \quad i = 1, \dots, m,$$

where $\mathbf{n} = (n_i)_{i=1}^m \in \mathcal{V}^m$ is an error term, referred to as the *sampling error*. Observe that the samples values d_i are assumed to be elements of the finite-dimensional space \mathcal{V}_h . This is a natural assumption to make. Indeed, in the context of parametric DEs, the value $f(\mathbf{y})$ (the solution of the DE with parameter value \mathbf{y}) is typically computed via a (finite element) discretization of the DE, thus yielding an element of \mathcal{V}_h , which is the corresponding discrete (finite element) space.

As a result of the assumption $d_i \in \mathcal{V}_h$, the error term n_i encompasses the error involved in approximating $f(\mathbf{y}_i) \in \mathcal{V}$ by an element of \mathcal{V}_h , e.g., the (finite element) discretization error in the context of a parametric DE. Note that we do not specify precisely how such an approximation is performed, nor how large an error this results in. In other words, we consider the computation that evaluates f at \mathbf{y}_i as a black box. A particular case of interest is when the d_i are the orthogonal projections of the exact sample values $f(\mathbf{y}_i)$, i.e.,

$$d_i = \mathcal{P}_h(f(\mathbf{y}_i)), \quad i = 1, \dots, m.$$

However, we do not assume this in what follows, since in practice the numerical procedure that yields the d_i may not involve computing the projection \mathcal{P}_h . Our goal is to develop algorithms for which the error scales linearly in $\|\mathbf{n}\|_{2;\mathcal{V}}$, the norm of the noise, thus accounting for any black box mechanism for computing the samples.

Recall that we consider a basis $\{\varphi_k\}_{k=1}^K$ for \mathcal{V}_h . We assume that the computation that evaluates $f(\mathbf{y}_i)$ produces the coefficients of the sample values d_i in this basis (i.e., the finite element coefficients in the aforementioned example). Therefore, we now write the sample values as

$$d_i = f(\mathbf{y}_i) + n_i = \sum_{k=1}^K d_{ik} \varphi_k, \quad i = 1, \dots, m, \quad (3.1)$$

and consider the values $d_{ik} \in \mathbb{C}$ as the *data* we obtain by sampling f .

3.2 Problem statement

We now formally define the input and output of the algorithm. The *input* of the algorithm is the collection of sample points $(\mathbf{y}_i)_{i=1}^m$ and the array of mK values $(d_{i,k})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ defined by (3.1). We next define the output. To this end, we first fix a multi-index set $\Lambda \subset \mathcal{F}$ of size $|\Lambda| = N$ for some $N \geq 1$. This set defines a polynomial space $\mathcal{P}_{\Lambda; \mathcal{V}_h}$, as in (2.9), within which we shall construct the resulting polynomial approximation. Hence, we consider an approximation of the form $\hat{f} \in \mathcal{P}_{\Lambda; \mathcal{V}_h}$ given by

$$\hat{f} : \mathbf{y} \mapsto \sum_{j=1}^N \left(\sum_{k=1}^K \hat{c}_{jk} \varphi_k \right) \Psi_{\mathbf{v}_j}(\mathbf{y}), \quad (3.2)$$

where $\hat{c}_{j,k} \in \mathbb{C}$ for $j \in [N]$, $k \in [K]$ and $\mathbf{v}_1, \dots, \mathbf{v}_N$ is some indexing of the multi-indices in Λ . In this way, we define formally the *output* of the algorithm as the array of coefficients $(\hat{c}_{jk})_{j,k=1}^{N,K} \in \mathbb{C}^{N \times K}$.

Finally, in order to define an algorithm we need one additional ingredient. Let

$$\mathbf{G} = (\langle \varphi_j, \varphi_k \rangle_{\mathcal{V}})_{j,k=1}^K \in \mathbb{C}^{K \times K} \quad (3.3)$$

denote the Gram matrix of the basis $\{\varphi_k\}_{k=1}^K \subset \mathcal{V}_h$. Note that \mathbf{G} is self adjoint and positive definite. However, \mathbf{G} is only equal to the identity when $\{\varphi_k\}_{k=1}^K$ is orthonormal. In what follows, we assume that it is possible to perform matrix-vector multiplications with \mathbf{G} . In other words, we have access to the function

$$\mathcal{T}_{\mathbf{G}} : \mathbb{C}^K \rightarrow \mathbb{C}^K, \quad \mathbf{x} \mapsto \mathbf{G}\mathbf{x}.$$

We also write $F(\mathbf{G})$ for the maximum number of arithmetic operations and comparisons required to evaluate $\mathcal{T}_{\mathbf{G}}(\mathbf{x})$ for arbitrary \mathbf{x} . Note that $F(\mathbf{G}) \leq K^2$ in general. However, this may be smaller when \mathbf{G} is structured. For instance, in the case of a finite element discretization, this computation can often be performed in $\mathcal{O}(K)$ operations.

Definition 3.1 (Algorithm for polynomial approximation of Hilbert-valued functions). Let $\Lambda \subset \mathcal{F}$ of size $|\Lambda| = N$ be given, along with an indexing $\mathbf{v}_1, \dots, \mathbf{v}_N$ of the multi-indices in Λ . An *algorithm for polynomial approximation of Hilbert-valued functions from sample values* is a mapping

$$\mathcal{A} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K}, \quad ((\mathbf{y}_i)_{i=1}^m, (d_{i,k})_{i,k=1}^{m,K}) \mapsto (\hat{c}_{jk})_{j,k=1}^{N,K},$$

for which the evaluation of $\mathcal{A}((\mathbf{y}_i), (d_{i,k}))$ involves only finitely many arithmetic operations (including square roots), comparisons and evaluations of the matrix-vector multiplication function $\mathcal{T}_{\mathcal{G}}$. If $(d_{i,k})$ is as in (3.1) for some $f \in L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})$, then the resulting approximation \hat{f} of f is given by (3.2), where $(\hat{c}_{jk}) = \mathcal{A}((\mathbf{y}_i), (d_{i,k}))$. The *computational cost* of an algorithm \mathcal{A} is the maximum number of arithmetic operations and comparisons (including those used in the evaluation of $\mathcal{T}_{\mathcal{G}}$) used to compute the output from any input.

Remark 3.2. As formulated above, it is up to the user to choose a suitable multi-index set Λ . Fortunately, as we see in our main results below, this multi-index set is given simply and explicitly in terms of m and another parameter ϵ (a failure probability). In particular, no ‘oracle’ knowledge of the function being approximated is required. Thus, one can also make the stronger assertion in what follows in which the algorithm takes the same input, but outputs both the desired index set Λ and the polynomial coefficients. For ease of presentation, we shall not do this.

Remark 3.3. When $d = \infty$ each sample point \mathbf{y}_i is an infinite sequence of real numbers. It is implicit in Definition 3.1 that the algorithm only accesses finitely many entries of this sequence. This does not cause any problems. As noted, the polynomial approximation is obtained in the index set Λ , which is a finite subset of \mathcal{F} . Hence, the multi-indices in Λ are nonzero only in their first n entries, for some finite n . Therefore, it is only necessary to access the first n entries of each sequence \mathbf{y}_i . More concretely, in our main results below, the polynomial approximation in infinite dimensions is obtained in a multi-index set $\Lambda = \Lambda_n^{\text{HCl}}$ in which only the first n terms can be nonzero, where n is an integer given explicitly in terms of m and ϵ .

3.3 Main results

We now present the main results of this work. We reiterate at this stage that these results are formulated for Chebyshev and Legendre polynomials. See Chapter 11 for some further discussion on other polynomial systems.

As noted above, these results employ specific choices of the index set Λ in order to obtain the desired approximation rates. Specifically, in finite dimensions, we

consider the *hyperbolic cross* index set

$$\Lambda = \Lambda_{n,d}^{\text{HC}} = \left\{ \mathbf{v} = (v_k)_{k=1}^d \in \mathbb{N}_0^d : \prod_{k=1}^d (v_k + 1) \leq n \right\} \subset \mathbb{N}_0^d. \quad (3.4)$$

We term n the *order* of the hyperbolic cross. Note that it is common to consider (3.4) as the hyperbolic cross of order $n - 1$. We use n here as it is slightly more convenient for this work. When defined this way, $\Lambda_{n,d}^{\text{HC}}$ is in fact the union of all lower sets (see Definition 2.8) in d dimensions of size at most n (see, e.g., [8, Proposition 2.5]). Thus, this set is a natural choice for polynomial approximation.

In infinite dimensions, we define the following index set

$$\Lambda = \Lambda_n^{\text{HCl}} = \left\{ \mathbf{v} = (v_k)_{k=1}^{\infty} \in \mathcal{F} : \prod_{j=1}^n (v_k + 1) \leq n, v_k = 0, k > n \right\} \subset \mathcal{F}.$$

Similarly, the union of all anchored sets (Definition 2.8) of size at most n in infinite dimensions is a subset of Λ_n^{HCl} (see, e.g., [8, Proposition 2.18]). Note that Λ_n^{HCl} is isomorphic to $\Lambda_{n,n}^{\text{HC}}$ under the restriction map

$$\mathbf{v} = (v_k)_{k=1}^{\infty} \in \mathcal{F} \mapsto (v_k)_{k=1}^n \in \mathbb{N}_0^d.$$

For convenience, we now also define

$$N = \Theta(n, d) = \begin{cases} |\Lambda_{n,d}^{\text{HC}}| & d < \infty, \\ |\Lambda_n^{\text{HCl}}| = |\Lambda_{n,n}^{\text{HC}}| & d = \infty, \end{cases} \quad (3.5)$$

as the cardinality of the index set employed. In general, the exact behaviour of $\Theta(n, d)$ is unknown. However, it admits a variety of different bounds. These are summarized as follows for $d < \infty$:

$$N = |\Lambda_{n,d}^{\text{HC}}| \leq \min \left\{ 2n^3 4^d, en^{2+\log(d)/\log(2)}, \frac{n(\log(n) + d \log(2))^{d-1}}{(d-1)!} \right\}. \quad (3.6)$$

The bounds are based on [32, 89]. See also [8, Lemmas B.3–B.5].

Finally, we also define

$$\alpha = \begin{cases} 1 & \text{Legendre,} \\ \log(3)/\log(4) & \text{Chebyshev,} \end{cases} \quad (3.7)$$

and, given $m \geq 3$ and $\epsilon \in (0, 1)$, $L = L(m, d, \epsilon)$ as

$$L = \begin{cases} \log(m) \cdot (\log(m) \cdot \min\{\log(m) + d, \log(ed) \cdot \log(m)\}) + \log(\epsilon^{-1}) & d < \infty, \\ \log(m) \cdot (\log^3(m) + \log(\epsilon^{-1})) & d = \infty. \end{cases} \quad (3.8)$$

3.3.1 Algebraic rates of convergence, finite dimensions

Theorem 3.4 (Existence of a mapping; algebraic case, finite dimensions). *Let $d \in \mathbb{N}$, $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0^d} \subset L^2_{\varrho}(\mathcal{U})$ be either the orthonormal Chebyshev or Legendre basis and $\{\varphi_k\}_{k=1}^K$ be a basis for \mathcal{V}_h . Then for every $m \geq 3$, $0 < \epsilon < 1$ and $K \geq 1$, there is a mapping*

$$\mathcal{M} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

where $N = \Theta(n, d)$ is as in (3.5) with $n = \lceil m/L \rceil$ and $L = L(m, d, \epsilon)$ as in (3.8), with the following property. Let $f \in \mathcal{B}(\boldsymbol{\rho})$ for arbitrary $\boldsymbol{\rho} > \mathbf{1}$, draw $\mathbf{y}_1, \dots, \mathbf{y}_m$ randomly and independently according to ϱ and let $(d_{ik})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ be as in (3.1) for arbitrary noise terms $\mathbf{n} = (n_i)_{i=1}^m \in \mathcal{V}$. Let $(\hat{c}_{jk}) = \mathcal{M}((\mathbf{y}_i), (d_{ik}))$ and define the approximation \hat{f} as in (3.2) based on the index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$. Then the following holds with probability at least $1 - \epsilon$. The error satisfies

$$\|f - \hat{f}\|_{L^2_{\varrho}(\mathcal{U}; \mathcal{V})} \leq c_1 \cdot \zeta, \quad \|f - \hat{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot \zeta, \quad (3.9)$$

for any $0 < p \leq 1$, where

$$\zeta := C \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p} + \frac{\|\mathbf{n}\|_{2; \mathcal{V}}}{\sqrt{m}} + \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathcal{U}; \mathcal{V})}, \quad (3.10)$$

$c_0, c_1, c_2 \geq 1$ are universal constants and $C = C(d, p, \boldsymbol{\rho})$ depends on d , p and $\boldsymbol{\rho}$ only.

We now make several remarks about this result. The same remarks apply (with obvious modifications) to all subsequent results as well. First, notice how the index set Λ in which the approximation is constructed is given completely explicitly in terms of m , d and ϵ . Thus, as claimed in Remark 3.2, no ‘oracle’ information about the function being approximated is required. Indeed, notice that the mapping described in this theorem is *universal* in the sense that it applies equally to *any* function $f \in \mathcal{B}(\boldsymbol{\rho})$ and *any* $\boldsymbol{\rho} > \mathbf{1}$.

A key aspect of this theorem is the factor ζ , defined in (3.10), which determines the error bounds (3.9). As claimed in Section 1.4, this incorporates three main key errors arising in the approximation process.

- (i) *The approximation error.* This is the algebraically decaying term in ζ : namely, $E_{\text{app}} = C \cdot (m/(c_0 L))^{1/2-1/p}$. It is completely equivalent to the best s -term approximation error bound in Theorem 2.4, except with s replaced by $m/(c_0 L)$.
- (ii) *The sampling error.* This is the term $E_{\text{samp}} = \|\mathbf{n}\|_{2; \mathcal{V}}/\sqrt{m}$, where $\mathbf{n} = (n_i)_{i=1}^m$ is as in (3.1). In other words, the effect of any errors in computing the sample values $f(\mathbf{y}_i)$ enters linearly in the overall error bound.

(iii) *The physical discretization error.* This is the term

$$E_{\text{disc}} = \|f - \mathcal{P}_h(f)\|_{L^\infty(\mathfrak{U};\mathfrak{V})}.$$

It describes the effect of working in the finite-dimensional subspace \mathcal{V}_h , instead of the full space \mathcal{V} . Critically, it depends on the orthogonal projection (best approximation) $\mathcal{P}_h(f)$ of f from \mathcal{V}_h .

Notice that (i) also describes the *sample complexity* of the scheme. Indeed, Theorem 3.4 asserts that there is a polynomial approximation that can be obtained from m samples that attains the best s -term rate $s^{1/2-1/p}$, where $s = m/(c_0L)$ scales like m up to the polylogarithmic factor L .

Theorem 3.4 asserts the existence of a mapping that takes samples values as its input and produces the coefficients of a polynomial approximation attaining a desired error bound as its output. The mapping, as we see later, arises as a minimizer of a certain weighted ℓ^1 -minimization problem. Thus, it is not an algorithm in the sense of Definition 3.1. In the next two theorems we assert the existence of algorithms that attain the same error, plus additional algorithmic error terms.

Theorem 3.5 (Existence of an algorithm; algebraic case, finite dimensions). *Consider the setup of Theorem 3.4. Then, for every $t \geq 1$, there exists an algorithm*

$$\mathcal{A}_t : \mathfrak{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

in the sense of Definition 3.1 such that the same property holds, except with (3.9) replaced by

$$\begin{aligned} \|f - \hat{f}\|_{L^2_{\mathfrak{D}}(\mathfrak{U};\mathfrak{V})} &\leq c_1 \cdot \left(\zeta + \frac{1}{t}\right), \\ \|f - \hat{f}\|_{L^\infty(\mathfrak{U};\mathfrak{V})} &\leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot \left(\zeta + \frac{1}{t}\right), \end{aligned}$$

where $c_1, c_2 \geq 1$ are as in (3.9) and ζ is as in (3.10). The computational cost of the algorithm is bounded by

$$c_3 \cdot [m \cdot \Theta(n, d) \cdot d + t \cdot (m \cdot \Theta(n, d) \cdot K + (\Theta(n, d) + m) \cdot (F(\mathbf{G}) + K)) \cdot (\Theta(n, d))^\alpha], \quad (3.11)$$

where $n = \lceil m/L \rceil$ is as in Theorem 3.4, $\Theta(n, d)$ is as in (3.5), α is as in (3.7) and $c_3 > 0$ is a universal constant.

The key element of this theorem is that the same error bound as in Theorem 3.4 is attained, up to an additional term. In particular, we have the three sources of errors (i)–(iii), plus the following:

(iv) *The algorithmic error.* This is the error $E_{\text{alg}} = 1/t$ committed by the algorithm \mathcal{A}_t in approximately computing the output of the mapping \mathcal{M} in Theorem 3.4. It is given in terms of the parameter t , which also enters linearly into the computational cost estimate (3.11).

Unfortunately, the $1/t$ decay rate of the algorithmic error is slow. Hence, it may be computationally expensive to compute an approximation to within a desired error bound. Fortunately, as we now explain, it is possible to improve it to e^{-t} subject to an additional technical assumption.

Theorem 3.6 (Existence of an efficient algorithm; algebraic case, finite dimensions). *Consider the setup of Theorem 3.4. Then for every $t \geq 1$ and $\zeta' > 0$ there exists an algorithm*

$$\mathcal{A}_{t,\zeta'} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

in the sense of Definition 3.1 such that the same property holds whenever $\zeta' \geq \zeta$, except with (3.9) replaced by

$$\begin{aligned} \|f - \hat{f}\|_{L^2_\varrho(\mathcal{U}; \mathcal{V})} &\leq c_1 \cdot (\zeta + \zeta' + e^{-t}), \\ \|f - \hat{f}\|_{L^\infty(\mathcal{U}; \mathcal{V})} &\leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot (\zeta + \zeta' + e^{-t}), \end{aligned} \quad (3.12)$$

where $c_1, c_2 \geq 1$ are as in (3.9) and ζ is as in (3.10). The computational cost of the algorithm is bounded by

$$c_3 \cdot [m \cdot \Theta(n, d) \cdot d + t \cdot (m \cdot \Theta(n, d) \cdot K + (\Theta(n, d) + m) \cdot (F(\mathbf{G}) + K)) \cdot (\Theta(n, d))^\alpha],$$

where $n = \lceil m/L \rceil$ is as in Theorem 3.4, $\Theta(n, d)$ is as in (3.5), α is as in (3.7) and $c_3 > 0$ is a universal constant.

We refer to this as an “efficient” algorithm, since the parameter t enters linearly in the computational cost but the algorithmic error scales like e^{-t} . The main limitation of this result is that the algorithm parameter ζ' needs to be an upper bound for the true error bound ζ in order for (3.12) to hold. This is a technical assumption for the proof, and does not appear necessary in practice. We demonstrate this phenomenon through numerical experiment in Chapter 5.

3.3.2 Algebraic rates of convergence, infinite dimensions

We now consider algebraic rates of convergence in the infinite-dimensional setting. The next three results should be compared against the corresponding best s -term approximation result, Theorem 2.5.

Theorem 3.7 (Existence of a mapping; algebraic case, infinite dimensions). *Let $d = \infty$, $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0^d} \subset L^2_\varrho(\mathcal{U})$ be either the orthonormal Chebyshev or Legendre basis and $\{\varphi_k\}_{k=1}^K$ be a basis for \mathcal{V}_h . Then for every $m \geq 3$, $0 < \epsilon < 1$ and $K \geq 1$, there is a mapping*

$$\mathcal{M} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

where $N = \Theta(n, d)$ is as in (3.5) with $n = \lceil m/L \rceil$, where $L = L(m, d, \epsilon)$ is as in (3.8), with the following property. Let $\epsilon > 0$, $0 < p < 1$ and $\mathbf{b} \in \ell^p(\mathbb{N})$, $\mathbf{b} > \mathbf{0}$, be monotonically nonincreasing. Let $f \in \mathcal{B}(\mathbf{b}, \epsilon)$, draw $\mathbf{y}_1, \dots, \mathbf{y}_m$ randomly and independently according to ϱ and let $(d_{ik})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ be as in (3.1) for arbitrary noise terms $\mathbf{n} = (n_i)_{i=1}^m \in \mathcal{V}$. Let $(\hat{c}_{jk}) = \mathcal{M}((\mathbf{y}_i), (d_{ik}))$ and define the approximation \hat{f} as in (3.2) based on the index set $\Lambda = \Lambda_n^{\text{HCl}}$. Then the following holds with probability at least $1 - \epsilon$. The error satisfies

$$\|f - \hat{f}\|_{L^2_{\varrho}(\mathfrak{U}; \mathfrak{V})} \leq c_1 \cdot \zeta, \quad \|f - \hat{f}\|_{L^\infty(\mathfrak{U}; \mathfrak{V})} \leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot \zeta, \quad (3.13)$$

where

$$\zeta := C \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p} + \frac{\|\mathbf{n}\|_{2; \mathfrak{V}}}{\sqrt{m}} + \|f - \mathcal{P}_h(f)\|_{L^\infty(\mathfrak{U}; \mathfrak{V})}, \quad (3.14)$$

$c_0, c_1, c_2 \geq 1$ are universal constants and $C = C(\mathbf{b}, \epsilon, p)$ depends on \mathbf{b} , ϵ and p only.

Theorem 3.8 (Existence of an algorithm; algebraic case, infinite dimensions). *Consider the setup of Theorem 3.7. Then, for every $t \geq 1$, there exists an algorithm*

$$\mathcal{A}_t : \mathfrak{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

in the sense of Definition 3.1 such that the same property holds, except with (3.13) replaced by

$$\begin{aligned} \|f - \hat{f}\|_{L^2_{\varrho}(\mathfrak{U}; \mathfrak{V})} &\leq c_1 \cdot \left(\zeta + \frac{1}{t}\right), \\ \|f - \hat{f}\|_{L^\infty(\mathfrak{U}; \mathfrak{V})} &\leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot \left(\zeta + \frac{1}{t}\right), \end{aligned}$$

where $c_1, c_2 \geq 1$ are as in (3.13) and ζ is as in (3.14). The computational cost of the algorithm is bounded by

$$c_3 \cdot [m \cdot \Theta(n, \infty) \cdot n + t \cdot (m \cdot \Theta(n, \infty) \cdot K + (\Theta(n, \infty) + m) \cdot (F(\mathbf{G}) + K)) \cdot (\Theta(n, \infty))^\alpha],$$

where $n = \lceil m/L \rceil$ is as in Theorem 3.7, $\Theta(n, \infty)$ is as in (3.5), α is as in (3.7) and $c_3 > 0$ is a universal constant.

In finite dimensions, the computational cost estimate (3.11) is somewhat difficult to interpret, since its behaviour depends on the relative sizes of m and d . Fortunately, in infinite dimensions we can give a more informative assessment. Suppose, for simplicity, that K is fixed (for example, $K = 1$ in the case of a scalar-valued function approximation problem). Then the computational cost is bounded by

$$c \cdot m \cdot \Theta(n, \infty) \cdot n + c_K \cdot t \cdot m \cdot \Theta(n, \infty)^{\alpha+1},$$

where $c > 0$ is a universal constant $c_K > 0$ is a constant depending on K only. Recall from (3.5) that $\Theta(n, \infty) = |\Lambda_n^{\text{HCl}}| = |\Lambda_{n,n}^{\text{HC}}|$. Now, when $d = n$ and n is sufficiently large, the minimum in (3.6) is attained by the second term $e n^{2+\log(n)/\log(2)}$. Substituting this into the above expression and recalling that $n = \lceil m/L \rceil$, where $L = L(m, \infty, \epsilon)$, we deduce that the computational cost is bounded by

$$c_K \cdot t \cdot m \cdot g(m)^{(\alpha+1)\log(4g(m))/\log(2)}, \quad g(m) := \left\lceil \frac{m}{\log(m) \cdot (\log^3(m) + \log(\epsilon^{-1}))} \right\rceil.$$

Since $m \geq 3$ by assumption, we have $\log(m) \geq 1$ and therefore $g(m) \leq m$. Hence, this admits the slightly looser upper bound

$$c_K \cdot t \cdot m^{1+(\alpha+1)\log(4m)/\log(2)}.$$

We deduce that the computational cost (for fixed K and t) is *subexponential* in m . Further, if we choose $t = m^{1/p-1/2}$ in accordance with the algebraically decaying term in (3.14), then we conclude the following: *it is possible to approximate a holomorphic function of infinitely many variables with error decaying algebraically fast in m via an algorithm whose computational cost is subexponential in m* . Whether this can be reduced to an algebraic cost is an open problem.

Theorem 3.9 (Existence of an efficient algorithm; algebraic case, infinite dimensions). *Consider the setup of Theorem 3.7. Then, for every $t \geq 1$ and $\zeta' > 0$ there exists an algorithm*

$$\mathcal{A}_{t,\zeta'} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

in the sense of Definition 3.1 such that the same property holds whenever $\zeta' \geq \zeta$, except with (3.13) replaced by

$$\begin{aligned} \|f - \hat{f}\|_{L^2_0(\mathcal{U}; \mathcal{V})} &\leq c_1 \cdot (\zeta + \zeta' + e^{-t}), \\ \|f - \hat{f}\|_{L^\infty(\mathcal{U}; \mathcal{V})} &\leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot (\zeta + \zeta' + e^{-t}), \end{aligned}$$

where $c_1, c_2 \geq 1$ are as in (3.13) and $\zeta \leq \zeta'$ is as in (3.14). The computational cost of the algorithm is bounded by

$$c_3 \cdot [m \cdot \Theta(n, \infty) \cdot n + t \cdot (m \cdot \Theta(n, \infty) \cdot K + (\Theta(n, \infty) + m) \cdot (F(\mathbf{G}) + K)) \cdot (\Theta(n, \infty))^\alpha],$$

where $n = \lceil m/L \rceil$ is as in Theorem 3.7, $\Theta(n, \infty)$ is as in (3.5), α is as in (3.7) and $c_3 > 0$ is a universal constant.

3.3.3 Exponential rates of convergence, finite dimensions

Finally, we consider exponential rates of convergence in finite dimensions. The following results should be compared against Theorem 2.6.

Theorem 3.10 (Existence of a mapping; exponential case, finite dimensions). *Let $d \in \mathbb{N}$, $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0^d} \subset L^2_{\varrho}(\mathcal{U})$ be either the orthonormal Chebyshev or Legendre basis and $\{\varphi_k\}_{k=1}^K$ be a basis for \mathcal{V}_h . Then for every $m \geq 3$, $0 < \epsilon < 1$ and $K \geq 1$, there is a mapping*

$$\mathcal{M} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

where $N = \Theta(n, d)$ is as in (3.5) with

$$n = \begin{cases} \lceil \sqrt{m/L} \rceil & \text{Legendre,} \\ \lceil m/(2^d L) \rceil & \text{Chebyshev,} \end{cases} \quad (3.15)$$

and L as in (3.8), with the following property. Draw $\mathbf{y}_1, \dots, \mathbf{y}_m$ randomly and independently according to ϱ . Then, with probability at least $1 - \epsilon$, the following holds. Let $f \in \mathcal{B}(\boldsymbol{\rho})$ for arbitrary $\boldsymbol{\rho} > \mathbf{1}$, $(d_{ik})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ be as in (3.1) for arbitrary noise terms $\mathbf{n} = (n_i)_{i=1}^m \in \mathcal{V}$, $(\hat{c}_{jk})_{j,k=1}^{N,K} = \mathcal{M}((\mathbf{y}_i)_{i=1}^m, (d_{ik})_{i,k=1}^{m,K})$ and define the approximation \hat{f} as in (3.2) based on the index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$. Then the error satisfies

$$\|f - \hat{f}\|_{L^2_{\varrho}(\mathcal{U}; \mathcal{V})} \leq c_1 \cdot \zeta, \quad \|f - \hat{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot \zeta, \quad (3.16)$$

for any

$$0 < \gamma < (d + 1)^{-1} \left(d! \prod_{j=1}^d \log(\rho_j) \right)^{1/d},$$

where

$$\zeta := C \cdot \begin{cases} \exp\left(-\frac{\gamma}{2} \left(\frac{m}{c_0 L}\right)^{\frac{1}{d}}\right) & \text{Chebyshev} \\ \exp\left(-\gamma \left(\frac{m}{c_0 L}\right)^{\frac{1}{2d}}\right) & \text{Legendre} \end{cases} + \frac{\|\mathbf{n}\|_{2; \mathcal{V}}}{\sqrt{m}} + \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathcal{U}; \mathcal{V})}, \quad (3.17)$$

$c_0, c_1, c_2 \geq 1$ are universal constants and $C = C(d, \gamma, \boldsymbol{\rho})$ depends on d , γ and $\boldsymbol{\rho}$ only.

Theorem 3.11 (Existence of an algorithm; exponential case, finite dimensions). *Consider the setup of Theorem 3.10. Then, for every $t \geq 1$, there exists an algorithm*

$$\mathcal{A}_t : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

in the sense of Definition 3.1 such that the same property holds, except with (3.16) replaced by

$$\|f - \hat{f}\|_{L^2_{\varrho}(\mathcal{U}; \mathcal{V})} \leq c_1 \cdot \left(\zeta + \frac{1}{t} \right),$$

$$\|f - \hat{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \leq c_2 \cdot \sqrt{\frac{m}{L}} \cdot \left(\zeta + \frac{1}{t} \right),$$

where $c_1, c_2 \geq 1$ are as in (3.16) and ζ is as in (3.17). The computational cost of the algorithm is bounded by

$$c_3 \cdot [m \cdot \Theta(n, d) \cdot n + t \cdot (m \cdot \Theta(n, d) \cdot K + (\Theta(n, d) + m) \cdot (F(\mathbf{G}) + K)) \cdot (\Theta(n, d))^\alpha],$$

where n is as in (3.15), $\Theta(n, d)$ is as in (3.5), α is as in (3.7) and $c_3 > 0$ is a universal constant.

Theorem 3.12 (Existence of an efficient algorithm; exponential case, finite dimensions). *Consider the setup of Theorem 3.10. Suppose that there is a known upper bound $\zeta' \geq \zeta$, where ζ is as in (3.17). Then, for every $t \geq 1$ and $\zeta' > 0$ there exists an algorithm*

$$\mathcal{A}_{t, \zeta'} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K},$$

in the sense of Definition 3.1 for which the same property holds whenever $\zeta' \geq \zeta$, except with (3.16) replaced by

$$\begin{aligned} \|f - \hat{f}\|_{L^2_{\sigma}(\mathcal{U}; \mathcal{V})} &\leq c_1 \cdot (\zeta + \zeta' + e^{-t}), \\ \|f - \hat{f}\|_{L^\infty(\mathcal{U}; \mathcal{V})} &\leq c_2 \cdot \sqrt{\frac{m}{L}} (\zeta + \zeta' + e^{-t}), \end{aligned}$$

where $c_1, c_2 \geq 1$ are as in (3.16) and ζ is as in (3.17). The computational cost of the algorithm is bounded by

$$c_3 \cdot [m \cdot \Theta(n, d) \cdot n + t \cdot (m \cdot \Theta(n, d) \cdot K + (\Theta(n, d) + m) \cdot (F(\mathbf{G}) + K)) \cdot (\Theta(n, d))^\alpha],$$

where n is as in (3.15), $\Theta(n, d)$ is as in (3.5), α is as in (3.7) and $c_3 > 0$ is a universal constant.

As before, suppose that K is fixed and, since we consider exponential rates, that d is also fixed. Then, using the third estimate in (3.6), we deduce that the computational cost of this algorithm is bounded by

$$c_{K,d} \cdot (m \cdot n^2 \cdot (\log(n))^{d-1} + t \cdot m \cdot (n \cdot (\log(n))^{d-1})^{\alpha+1}).$$

Using the crude bound $n \leq m$, we obtain the bound

$$c_{K,d} \cdot (t \cdot m^{\alpha+2} (\log(m))^{(d-1)(\alpha+1)}).$$

Thus, for fixed t , the computational cost is polynomial in m as $m \rightarrow \infty$. In particular, with the efficient algorithm of Theorem 3.12 (subject to the caveat that an upper bound for the error is known) we deduce the following: *in fixed dimension d , it is possible to approximate a holomorphic function with error decaying exponentially fast in m via an algorithm whose computational cost is polynomial in m .* Whether the polynomial growth rate described above is sharp is an open problem.

Remark 3.13. There is a subtle difference between the algebraic and exponential results. The former are *nonuniform* in the sense that a single draw of the sample points $\mathbf{y}_1, \dots, \mathbf{y}_m$ is sufficient for recovery of a fixed function f with high probability up to the specified error bound. The latter are *uniform*, since a single draw of the sample points $\mathbf{y}_1, \dots, \mathbf{y}_m$ is sufficient for recovery of any function with high probability up to the specified error bound. The reason for this difference stems from bounding a *discrete* error term (8.10), which is a random variable depending on f and the sample points. In the algebraic case, in order to obtain the desired algebraic exponent $1/2 - 1/p$ we bound this term with high probability for each fixed f . See Step 4 of the proof of Theorem 8.2. This renders the ensuing result nonuniform. Conversely, in the exponential case (where the appearance of small algebraic factors is not a concern, since they can be absorbed into the exponentially decaying term) we bound this term with probability one for *any* f . See Step 4 of the proof of Theorem 8.4. Note that one could also derive uniform guarantees in the algebraic case by considering a fixed value of p and letting \mathcal{M} and \mathcal{A} depend on p , or by considering a restricted range $0 < p \leq p^* < 1$. Both strategies involve a larger value of n , with its size depending on p or p^* . See [8, Section 7.6.2] for further discussion.

Chapter 4

Construction of the algorithms

In this chapter, we describe the construction of the algorithms asserted in our main results. These are based on techniques from compressed sensing [8, 13, 61] on the premise that the polynomial coefficients of a holomorphic function are approximately sparse. There are several main differences between standard compressed sensing and what we develop below. First, following [2, 7, 8, 37, 119, 121], we work in a weighted setting in order to promote sparsity in lower or anchored sets (recall Section 2.7). Second, following [52], we work with Hilbert-valued vectors, whose entries take values in the Hilbert space \mathcal{V} . Finally, so as to avoid unrealistic assumptions on the functions being approximated, we use consider *noise-blind* decoders, as in [3]. See also Remark 4.1.

4.1 Recovery via Hilbert-valued, weighted ℓ^1 -minimization

We first require some additional notation. Given $N \in \mathbb{N}$ we let \mathcal{V}^N be the vector space of Hilbert-valued vectors of length N , i.e., $\mathbf{v} = (v_i)_{i=1}^N$ where $v_i \in \mathcal{V}$, $i = 1, \dots, N$. Next, given $\Lambda \subseteq \mathcal{F}$ and a vector of positive weights $\mathbf{w} = (w_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$, where $\mathbf{w} > \mathbf{0}$, we define the weighted $\ell_w^p(\Lambda; \mathcal{V})$ space, $0 < p \leq 2$, as the set of \mathcal{V} -valued sequences $\mathbf{v} = (v_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$ for which

$$\|\mathbf{v}\|_{p, \mathbf{w}; \mathcal{V}} := \left(\sum_{\mathbf{v} \in \Lambda} w_{\mathbf{v}}^{2-p} \|v_{\mathbf{v}}\|_{\mathcal{V}}^p \right)^{1/p} < \infty.$$

Notice that $\ell_w^2(\Lambda; \mathcal{V})$ coincides with the unweighted space $\ell^2(\Lambda; \mathcal{V})$.

Now, let $\Lambda \subset \mathcal{F}$ be a finite multi-index set of size $|\Lambda| = N$ and consider the ordering $\Lambda = \{\mathbf{v}_1, \dots, \mathbf{v}_N\}$. Note that we will, in practice, choose either $\Lambda = \Lambda_{n,d}^{\text{HC}}$ when $d < \infty$ or $\Lambda = \Lambda_n^{\text{HC1}}$ when $d = \infty$, where the order n is as described in the corresponding theorem (Theorems 3.4–3.12). With this in mind, given $f \in L_{\mathcal{G}}^2(\mathcal{U}; \mathcal{V})$, define

$$f_{\Lambda} = \sum_{\mathbf{v} \in \Lambda} c_{\mathbf{v}} \Psi_{\mathbf{v}} \tag{4.1}$$

as the truncated expansion of f based on the index set Λ and

$$\mathbf{c}_{\Lambda} = (c_{\mathbf{v}_j})_{j=1}^N \in \mathcal{V}^N \tag{4.2}$$

as the finite vector of coefficients of f with indices in Λ . As explained in Section 3.2, our objective is, in effect, to approximate these coefficients.

We do this as follows. Given $\mathbf{y}_1, \dots, \mathbf{y}_m \in \mathcal{U}$, we define the normalized measurement matrix

$$\mathbf{A} = \left(\frac{\Psi_{\mathbf{v}_j}(\mathbf{y}_i)}{\sqrt{m}} \right)_{i,j=1}^{m,N} \in \mathbb{C}^{m \times N} \quad (4.3)$$

and the normalized measurement and error vectors

$$\mathbf{b} = \frac{1}{\sqrt{m}}(f(\mathbf{y}_i) + n_i)_{i=1}^m \in \mathcal{V}_h^m, \quad \mathbf{e} = \frac{1}{\sqrt{m}}(n_i)_{i=1}^m \in \mathcal{V}^m. \quad (4.4)$$

Notice that any $m \times N$ matrix $\mathbf{A} = (a_{ij})_{i,j=1}^{m,N}$ extends to a bounded linear operator $\mathcal{V}^N \rightarrow \mathcal{V}^m$ (or $\mathcal{V}_h^N \rightarrow \mathcal{V}_h^m$) in the obvious way, i.e.,

$$\mathbf{x} = (x_i)_{i=1}^N \in \mathcal{V}^N \mapsto \mathbf{A}\mathbf{x} = \left(\sum_{j=1}^N a_{ij}x_j \right)_{i=1}^m \in \mathcal{V}^m.$$

For ease of notation, we make no distinction between the matrix $\mathbf{A} \in \mathbb{C}^{m \times N}$ and the linear operator $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ (or $\mathbf{A} \in \mathcal{B}(\mathcal{V}_h^N, \mathcal{V}_h^m)$) in what follows. Using this, we obtain

$$\mathbf{A}\mathbf{c}_\Lambda = \frac{1}{\sqrt{m}}(f_\Lambda(\mathbf{y}_i))_{i=1}^m = \frac{1}{\sqrt{m}}(f(\mathbf{y}_i))_{i=1}^m - \frac{1}{\sqrt{m}}(f(\mathbf{y}_i) - f_\Lambda(\mathbf{y}_i))_{i=1}^m,$$

and therefore

$$\mathbf{A}\mathbf{c}_\Lambda + \mathbf{e} + \mathbf{e}' = \mathbf{b}, \quad (4.5)$$

where

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

We have now formulated the recovery of \mathbf{c}_Λ as the solution of a noisy linear system (4.5), where the noise term $\mathbf{e} + \mathbf{e}'$ encompasses both the noise $\mathbf{e} = (n_i)_{i=1}^m / \sqrt{m}$ in the sample values and the error \mathbf{e}' due to the truncation (4.1) of the infinite expansion (2.8) via the index set Λ .

Due to the discussion in Sections 2.5–2.7, we expect the coefficients \mathbf{c}_Λ to not only be approximately sparse, but also well approximated by a subset of s coefficients whose indices define a lower or anchored set. In classical compressed sensing, one exploits sparse structure via minimizing an ℓ^1 -norm. To exploit sparse and lower structure, we follow ideas of [2, 7, 8, 37] and use a weighted ℓ^1 -norm penalty. Specifically, we now compute an approximate solution via the Hilbert-valued, *weighted Square-Root LASSO (SR-LASSO)* optimization problem

$$\min_{\mathbf{z} \in \mathcal{V}_h^N} \mathcal{G}(\mathbf{z}), \quad \mathcal{G}(\mathbf{z}) := \lambda \|\mathbf{z}\|_{1, \mathbf{w}; \mathcal{V}} + \|\mathbf{A}\mathbf{z} - \mathbf{b}\|_{2; \mathcal{V}}. \quad (4.6)$$

Here $\lambda > 0$ is a tuning hyperparameter.

Remark 4.1. As an alternative to solve this Hilbert-valued compressed sensing problem, we could use a formulation based on a constrained basis pursuit or unconstrained LASSO problem. However, we consider the SR-LASSO problem (4.6) instead. While other approaches are arguably more common, based on [3] the SR-LASSO has the desirable property that the optimal values of its hyperparameter λ is independent of the noise term (in this case $e + e'$). This is not the case for other formulations, whose hyperparameters need to be chosen in terms of the (unknown) magnitude of the noise in order to ensure good theoretical and practical performance (see, e.g., [13, Chapter 6]). This is particularly problematic in the setting of function approximation, where such terms are function dependent (for instance, the term e' depends on the expansion tail $f - f_\Lambda$) and therefore generally unknown. See [3] and [8, Section 6.6] for further discussion.

Notice that (4.6) is solved over \mathcal{V}_h^N not \mathcal{V}^N , since the latter would not be numerically solvable in general. As we see below, it can be reformulated an optimization problem over $\mathbb{C}^{N \times K}$, where $K = \dim(\mathcal{V}_h)$. However, since the true coefficients of f are elements of \mathcal{V} and not \mathcal{V}_h , this discretization inevitably results in an additional error, which must also be accounted for in the analysis. This leads precisely to the physical discretization error (term (iii) in Section 3.3.1).

Finally, we now also specify the weights. Following [2, 7, 37] (see also [8, Remark 2.14]), a good choice of weights (for promoting lower or anchored structure) is given by the so-called *intrinsic weights*

$$\mathbf{w} = \mathbf{u} = (u_{\mathbf{v}})_{\mathbf{v} \in \Lambda}, \quad u_{\mathbf{v}} = \|\Psi_{\mathbf{v}}\|_{L^\infty(\mathcal{U})}, \quad \mathbf{v} \in \Lambda. \quad (4.7)$$

In particular, for Chebyshev and Legendre polynomials these are given explicitly by

$$u_{\mathbf{v}} = \|\Psi_{\mathbf{v}}\|_{L^\infty(\mathcal{U})} = \begin{cases} \prod_{j=1}^d \sqrt{2v_j + 1}, & \text{Legendre,} \\ 2^{\|\mathbf{v}\|_0/2}, & \text{Chebyshev,} \end{cases}$$

where $\|\mathbf{v}\|_0 := |\text{supp}(\mathbf{v})|$. Typically, we index these weights over the multi-indices $\mathbf{v} \in \Lambda$. However, we will, for convenience, often write w_i instead of $w_{\mathbf{v}_i}$ in what follows, where, as above, $\{\mathbf{v}_1, \dots, \mathbf{v}_N\}$ is an ordering of Λ .

4.2 Reformulation as a matrix recovery problem and the mappings in Theorems 3.4, 3.7 and 3.10

We now describe the mappings whose existence is asserted in Theorems 3.4, 3.7 and 3.10. These maps all arise via exact solutions of weighted SR-LASSO optimization problems. However, since (4.6) yields a vector in \mathcal{V}_h^N and the mappings should yield outputs in $\mathbb{C}^{N \times K}$, we first need to reformulate (4.6) using the basis $\{\varphi_i\}_{i=1}^K$ for \mathcal{V}_h .

Notice first that any vector of coefficients $\mathbf{c} = (c_{\mathbf{v}_i})_{i=1}^N \in \mathcal{V}_h^N$ is equivalent to a matrix of coefficients

$$\mathbf{C} = (c_{ik})_{i,k=1}^{N,K} \in \mathbb{C}^{N \times K},$$

via the relation

$$c_{\mathbf{v}_i} = \sum_{k=1}^K c_{ik} \varphi_k, \quad i \in [N].$$

Next, observe that if $g = \sum_{k=1}^K d_k \varphi_k \in \mathcal{V}_h$ then

$$\|g\|_{\mathcal{V}} = \|\mathbf{d}\|_{\mathbf{G}} = \sqrt{\mathbf{d}^* \mathbf{G} \mathbf{d}},$$

where $\mathbf{d} = (d_k)_{k=1}^K \in \mathbb{C}^K$ and $\mathbf{G} \in \mathbb{C}^{K \times K}$ is the Gram matrix for $\{\varphi_k\}_{k=1}^K$, given by (3.3). Since \mathbf{G} is positive definite, it has a unique positive definite square root matrix $\mathbf{G}^{1/2}$. Hence, we may write

$$\|g\|_{\mathcal{V}} = \|\mathbf{G}^{1/2} \mathbf{d}\|_2.$$

We now use some additional notation. Given $1 \leq p \leq \infty$ and $1 \leq q \leq 2$, we define the weighted $\ell_w^{p,q}$ -norm of a matrix $\mathbf{C} = (c_{ik})_{i,k=1}^{N,K} \in \mathbb{C}^{N \times K}$ as

$$\|\mathbf{C}\|_{p,q,\mathbf{w}} = \left(\sum_{i=1}^N w_i^{2-p} \left(\sum_{k=1}^K |c_{ik}|^q \right)^{p/q} \right)^{1/p}.$$

Note that this is precisely the weighted ℓ_w^p -norm of the vector of $(\|c_i\|_q)_{i=1}^N$, where $c_i = (c_{ik})_{k=1}^K \in \mathbb{C}^K$ is the i th row of \mathbf{C} . Further, if $p = q = 2$, then this is just the unweighted $\ell^{2,2}$ -norm of a matrix (which is simply its Frobenius norm). In this case, we typically write $\|\cdot\|_{2,2}$.

Now let $\mathbf{z} \in \mathcal{V}_h^N$ be arbitrary, $\mathbf{Z} \in \mathbb{C}^{N \times K}$ be the corresponding matrix and $\mathbf{z}_i \in \mathbb{C}^K$ be the i th row of \mathbf{Z} . Then

$$\|\mathbf{z}\|_{1,\mathbf{w};\mathcal{V}} = \sum_{i=1}^N w_i \|z_{\mathbf{v}_i}\|_{\mathcal{V}} = \sum_{i=1}^N w_i \|\mathbf{G}^{1/2} \mathbf{z}_i\|_2 = \|\mathbf{Z} \mathbf{G}^{1/2}\|_{2,1,\mathbf{w}}.$$

Similarly, let $\mathbf{A} = (a_{ij})_{i,j=1}^{m,N} \in \mathbb{C}^{m \times N}$ and $\mathbf{b} = (b_i)_{i=1}^m \in \mathcal{V}_h^m$ be as in (4.3) and (4.4), respectively, and let $\mathbf{B} \in \mathbb{C}^{m \times K}$ be the matrix corresponding to \mathbf{b} . Then

$$\|\mathbf{A} \mathbf{z} - \mathbf{b}\|_{2;\mathcal{V}}^2 = \sum_{i=1}^m \left\| \sum_{j=1}^N a_{ij} z_{\mathbf{v}_j} - b_i \right\|_{\mathcal{V}}^2 = \|(\mathbf{A} \mathbf{Z} - \mathbf{B}) \mathbf{G}^{1/2}\|_{2,2}^2.$$

Therefore, we now consider the minimization problem

$$\min_{\mathbf{Z} \in \mathbb{C}^{N \times K}} \left\{ \lambda \|\mathbf{Z}\|_{2,1,\mathbf{w}} + \|(\mathbf{A} \mathbf{Z} - \mathbf{B}) \mathbf{G}^{1/2}\|_{2,2} \right\}. \quad (4.8)$$

-
- Let m , ϵ and n be as given in the particular theorem and set $\Lambda = \Lambda_{n,d}^{\text{HC}}$ (Theorems 3.4 and 3.10) or $\Lambda = \Lambda_n^{\text{HCl}}$ (Theorem 3.7).
 - Set $\lambda = (4\sqrt{m/L})^{-1}$, where $L = L(m, d, \epsilon)$ is as in (3.8).
 - Let $\mathbf{D} = (d_{ik})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ and $\mathbf{Y} = (y_i)_{i=1}^m$ be an input, as in (3.1), and set $\mathbf{B} = \frac{1}{\sqrt{m}}\mathbf{D}$.
 - Let \mathbf{G} , \mathbf{A} and \mathbf{w} be as in (3.3), (4.3) and (4.7), respectively.
 - Define the output $\hat{\mathbf{C}} = \mathcal{M}(\mathbf{Y}, \mathbf{D})$ as the minimizer of (4.8) with smallest $\ell^{2,2}$ -norm.
-

Table 4.1. The mappings $\mathcal{M} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K}$ used in Theorems 3.4, 3.7 and 3.10.

This is equivalent to (4.6) in the following sense. A vector $\hat{\mathbf{c}} = (\hat{c}_{v_i})_{i=1}^N \in \mathcal{V}_h^N$ is a minimizer of (4.6) if and only if the matrix $\hat{\mathbf{C}} = (\hat{c}_{ik})_{i,k=1}^{N,K} \in \mathbb{C}^{N \times K}$ with entries defined by the relation

$$\hat{c}_{v_i} = \sum_{k=1}^K \hat{c}_{ik} \varphi_k, \quad i \in [N],$$

is a minimizer of (4.8).

With this in hand, we are now ready to define the mappings used in Theorems 3.4, 3.7 and 3.10. These are described in Table 4.1. Note that these are indeed well-defined mappings, since the minimizer of (4.8) with smallest $\ell^{2,2}$ -norm is unique (this follows from the facts that (4.8) is a convex problem, therefore its set of minimizers is a convex set, and the function $\mathbf{Z} \mapsto \|\mathbf{Z}\|_{2,2}^2$ is strongly convex). This particular choice is arbitrary, and is made solely so as to have a well-defined mapping. It is of no consequence whatsoever in our analysis, since the various error bounds we prove later hold for any minimizer of (4.8).

4.3 The primal-dual iteration

To derive the algorithms described in the other main theorems, we need methods for approximately solving the optimization problems (4.6) and (4.8). We use the *primal-dual iteration* [30] (also known as the Chambolle–Pock algorithm) to this end. We first briefly describe the primal-dual iteration in the general case (see [30, 31, 31], as well as [13, Section 7.5] for more detailed treatments), before specializing to the weighted SR-LASSO problem in the next section.

Let $(\mathcal{X}, \langle \cdot, \cdot \rangle_{\mathcal{X}})$ and $(\mathcal{Y}, \langle \cdot, \cdot \rangle_{\mathcal{Y}})$ be (complex) Hilbert spaces, $g : \mathcal{X} \rightarrow \mathbb{R} \cup \{\infty\}$, $h : \mathcal{Y} \rightarrow \mathbb{R} \cup \{\infty\}$ be proper, lower semicontinuous and convex functions and

$A \in \mathcal{B}(\mathcal{X}, \mathcal{Y})$ be a bounded linear operator satisfying

$$\text{dom}(h) \cap A(\text{dom}(g)) \neq \emptyset.$$

The primal-dual iteration is a general method for solving the convex optimization problem

$$\min_{x \in \mathcal{X}} \{g(x) + h(A(x))\}. \quad (4.9)$$

Under this setting the (Fenchel–Rockafeller) dual problem is

$$\min_{\xi \in \mathcal{Y}} \{g^*(A^*\xi) + h^*(-\xi)\},$$

where g^* and h^* are the convex conjugate functions of g and h , respectively. Recall that, for a function $f : \mathcal{X} \rightarrow \mathbb{R} \cup \{\infty\}$, its convex conjugate is defined by

$$f^*(z) = \sup_{x \in \mathcal{X}} (\text{Re}\langle x, z \rangle_{\mathcal{Y}} - f(x)), \quad z \in \mathcal{X}. \quad (4.10)$$

The Lagrangian of (4.9) is defined by

$$\mathcal{L}(x, \xi) = g(x) + \text{Re}\langle A(x), \xi \rangle_{\mathcal{Y}} - h^*(\xi), \quad x \in \text{dom}(g), \xi \in \text{dom}(h^*), \quad (4.11)$$

and $\mathcal{L}(x, \xi) = \infty$ if $x \notin \text{dom}(g)$ or $\mathcal{L}(x, \xi) = -\infty$ if $\xi \notin \text{dom}(h^*)$. This in turn leads to the saddle-point formulation of the problem

$$\min_{x \in \mathcal{X}} \max_{\xi \in \mathcal{Y}} \mathcal{L}(x, \xi).$$

The primal-dual iteration seeks a solution $(\hat{x}, \hat{\xi})$ of the saddle-point problem by solving the following fixed-point equation

$$\begin{aligned} \hat{x} &= \text{prox}_{\tau g}(\hat{x} - \tau A^*(\hat{\xi})), \\ \hat{\xi} &= \text{prox}_{\sigma h^*}(\hat{\xi} + \sigma A(\hat{x})), \end{aligned}$$

where $\tau, \sigma > 0$ are stepsize parameters and prox is the proximal operator, which is defined by

$$\text{prox}_f(z) = \arg \min_{x \in \mathcal{X}} \left\{ f(x) + \frac{1}{2} \|x - z\|_{\mathcal{X}}^2 \right\}, \quad z \in \text{dom}(f).$$

To be precise, given initial values $(x^{(0)}, \xi^{(0)}) \in \mathcal{X} \times \mathcal{Y}$ the primal-dual iteration defines a sequence $\{(x^{(n)}, \xi^{(n)})\}_{n=1}^{\infty} \subset \mathcal{X} \times \mathcal{Y}$ as follows:

$$\begin{aligned} x^{(n+1)} &= \text{prox}_{\tau g}(x^{(n)} - \tau A^*(\xi^{(n)})), \\ \xi^{(n+1)} &= \text{prox}_{\sigma h^*}(\xi^{(n)} + \sigma A(x^{(n+1)} - x^{(n)})). \end{aligned} \quad (4.12)$$

4.4 The primal-dual iteration for the weighted SR-LASSO problem

We now apply this scheme to (4.6) and (4.8). We first describe an algorithm to approximately solve the Hilbert-valued problem (4.6), before using the equivalence between elements of \mathcal{V}_h^N and $\mathbb{C}^{N \times K}$ to obtain an algorithm for approximately solving (4.8).

Consider (4.6). We define $\mathcal{X} = (\mathcal{V}_h^N, \langle \cdot, \cdot \rangle_{2;\mathcal{V}})$, $\mathcal{Y} = (\mathcal{V}_h^m, \langle \cdot, \cdot \rangle_{2;\mathcal{V}})$ and $g : \mathcal{X} \rightarrow \mathbb{R} \cup \{\infty\}$, $h : \mathcal{Y} \rightarrow \mathbb{R} \cup \{\infty\}$ as the proper, lower semicontinuous and convex functions

$$g(\mathbf{x}) = \lambda \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}}, \quad h(\mathbf{y}) = \|\mathbf{y} - \mathbf{b}\|_{2;\mathcal{V}}, \quad \mathbf{x} \in \mathcal{V}_h^N, \quad \mathbf{y} \in \mathcal{V}_h^m.$$

We first find the proximal maps of g and h^* . Using (4.10), we see that

$$\begin{aligned} h^*(\boldsymbol{\xi}) &= \sup_{\mathbf{v} \in \mathcal{V}_h^m} (\operatorname{Re}\langle \mathbf{v}, \boldsymbol{\xi} \rangle_{\mathcal{V}} - \|\mathbf{v} - \mathbf{b}\|_{2;\mathcal{V}}) \\ &= \operatorname{Re}\langle \mathbf{b}, \boldsymbol{\xi} \rangle_{\mathcal{V}} + \sup_{\mathbf{v} \in \mathcal{V}_h^m} (\operatorname{Re}\langle \mathbf{v}, \boldsymbol{\xi} \rangle_{\mathcal{V}} - \|\mathbf{v}\|_{2;\mathcal{V}}), \end{aligned}$$

for all $\boldsymbol{\xi} \in \mathcal{V}_h^m$. From [22, Examples 13.3 and 13.4] it follows that

$$(\|\cdot\|_{\mathcal{V}})^* = \delta_B, \quad B := \{\boldsymbol{\xi} \in \mathcal{V}_h^m : \|\boldsymbol{\xi}\|_{2;\mathcal{V}} \leq 1\},$$

where δ_B is the indicator function of the set B , taking value $\delta_B(\boldsymbol{\xi}) = 0$ when $\boldsymbol{\xi} \in B$ and $+\infty$ otherwise. Hence,

$$h^*(\boldsymbol{\xi}) = \operatorname{Re}\langle \mathbf{b}, \boldsymbol{\xi} \rangle_{\mathcal{V}} + \delta_B(\boldsymbol{\xi}). \quad (4.13)$$

Using this, we obtain

$$\begin{aligned} \operatorname{prox}_{\sigma h^*}(\boldsymbol{\xi}) &= \arg \min_{\mathbf{z} \in \mathcal{V}_h^m} \left\{ \sigma \delta_B(\mathbf{z}) + \sigma \operatorname{Re}\langle \mathbf{b}, \mathbf{z} \rangle_{\mathcal{V}} + \frac{1}{2} \|\mathbf{z} - \boldsymbol{\xi}\|_{2;\mathcal{V}}^2 \right\} \\ &= \arg \min_{\mathbf{z} : \|\mathbf{z}\|_{2;\mathcal{V}} \leq 1} \left\{ \frac{1}{2} \|\mathbf{z} - (\boldsymbol{\xi} - \sigma \mathbf{b})\|_{2;\mathcal{V}}^2 \right\} \\ &= \operatorname{proj}_B(\boldsymbol{\xi} - \sigma \mathbf{b}), \end{aligned}$$

where proj_B is the projection onto B , which is given explicitly by

$$\operatorname{proj}_B(\boldsymbol{\xi}) = \min \left\{ 1, \frac{1}{\|\boldsymbol{\xi}\|_{2;\mathcal{V}}} \right\} \boldsymbol{\xi}.$$

On the other hand, applying the definition of the proximal operator to the function τg with parameter $\tau > 0$, we deduce that

$$(\operatorname{prox}_{\tau g}(\mathbf{x}))_i = \operatorname{prox}_{\tau w_i \lambda \|\cdot\|_{\mathcal{V}}}(x_i), \quad i = 1, \dots, N, \quad \text{where } \mathbf{x} = (x_i)_{i=1}^N \in \mathcal{V}_h^N.$$

Algorithm 1: primal-dual-wSRLASSO – the primal-dual iteration for the weighted SR-LASSO problem (4.6)

inputs : measurement matrix $A \in \mathbb{C}^{m \times N}$, measurements $\mathbf{b} \in \mathcal{V}_h^N$, positive weights $\mathbf{w} = (w_i)_{i=1}^N$, parameter $\lambda > 0$, stepsizes $\tau, \sigma > 0$, maximum number of iterations $T \geq 1$, initial values $\mathbf{c}^{(0)} \in \mathcal{V}_h^N$, $\boldsymbol{\xi}^{(0)} \in \mathcal{V}_h^m$

output : $\bar{\mathbf{c}} = \text{primal-dual-wSRLASSO}(A, \mathbf{b}, \mathbf{w}, \lambda, \tau, \sigma, T, \mathbf{c}^{(0)}, \boldsymbol{\xi}^{(0)})$, an approximate minimizer of (4.6)

initialize: $\bar{\mathbf{c}}^{(0)} = \mathbf{0} \in \mathcal{V}_h^N$

```

1 for  $n = 0, 1, \dots, T - 1$  do
2    $\mathbf{p} = (p_i)_{i=1}^N = \mathbf{c}^{(n)} - \tau A^* \boldsymbol{\xi}^{(n)}$ 
3    $\mathbf{c}^{(n+1)} = \left( \max\{\|p_i\|_{\mathcal{V}} - \tau \lambda w_i, 0\} \frac{p_i}{\|p_i\|_{\mathcal{V}}} \right)_{i=1}^N$ 
4    $\mathbf{q} = \boldsymbol{\xi}^{(n)} + \sigma A(2\mathbf{c}^{(n+1)} - \mathbf{c}^{(n)}) - \sigma \mathbf{b}$ 
5    $\boldsymbol{\xi}^{(n+1)} = \min\left\{1, \frac{1}{\|\mathbf{q}\|_{2:\mathcal{V}}}\right\} \mathbf{q}$ 
6    $\bar{\mathbf{c}}^{(n+1)} = \frac{n}{n+1} \bar{\mathbf{c}}^{(n)} + \frac{1}{n+1} \mathbf{c}^{(n+1)}$ 
7 end
8  $\bar{\mathbf{c}} = \bar{\mathbf{c}}^{(T)}$ 

```

Moreover, a simple adaptation of [22, Example 14.5] with the $\|\cdot\|_{\mathcal{V}}$ -norm gives

$$\text{prox}_{\tau\|\cdot\|_{\mathcal{V}}}(x) = \max\{\|x\|_{\mathcal{V}} - \tau, 0\} \frac{x}{\|x\|_{\mathcal{V}}}, \quad \forall x \in \mathcal{V}_h \setminus \{0\}.$$

Hence,

$$\text{prox}_{\tau g}(\mathbf{x}) = \left(\max\{\|x_i\|_{\mathcal{V}} - \tau \lambda w_i, 0\} \frac{x_i}{\|x_i\|_{\mathcal{V}}} \right)_{i=1}^N, \quad \mathbf{x} = (x_i)_{i=1}^N \in \mathcal{V}_h^N \setminus \{\mathbf{0}\}.$$

With this in hand, we are now ready to define the primal-dual iteration for (4.6). As we see later, the analysis of convergence for the primal-dual iteration is given in terms of the *ergodic* sequence

$$\bar{\mathbf{c}}^{(n)} = \frac{1}{n} \sum_{i=1}^n \mathbf{c}^{(i)}, \quad n = 1, 2, \dots,$$

where $\mathbf{c}^{(i)} \in \mathcal{V}_h^N$ is the primal variable obtained at the i th step of the iteration. Hence, we now include the computation of these sequences in the primal-dual iteration for the weighted SR-LASSO problem (4.6), and take this as the output. The resulting procedure is described in Algorithm 1.

Algorithm 2: primal-dual-wSRLASSO-C – the primal-dual iteration for the weighted SR-LASSO problem (4.8)

inputs : measurement matrix $\mathbf{A} \in \mathbb{C}^{m \times N}$, measurements $\mathbf{B} \in \mathbb{C}^{m \times K}$, positive weights $\mathbf{w} = (w_i)_{i=1}^N$, Gram matrix $\mathbf{G} \in \mathbb{C}^{K \times K}$, parameter $\lambda > 0$, stepsizes $\tau, \sigma > 0$, maximum number of iterations $T \geq 1$, initial values $\mathbf{C}^{(0)} \in \mathbb{C}^{N \times K}$, $\mathbf{\Xi}^{(0)} \in \mathbb{C}^{m \times K}$

output : $\bar{\mathbf{C}} =$
 primal-dual-wSRLASSO-C($\mathbf{A}, \mathbf{b}, \mathbf{w}, \mathbf{G}, \lambda, \tau, \sigma, T, \mathbf{C}^{(0)}, \mathbf{\Xi}^{(0)}$),
 an approximate minimizer of (4.8)

initialize: $\bar{\mathbf{C}}^{(0)} = \mathbf{0} \in \mathbb{C}^{N \times K}$

```

1 for  $n = 0, 1, \dots, T - 1$  do
2    $\mathbf{P} = (p_{ik})_{j,k=1}^{N,K} = \mathbf{C}^{(n)} - \tau \mathbf{A}^* \mathbf{\Xi}^{(n)}$ 
3   for  $i = 1, \dots, N$  do
4      $\mathbf{p}_i = (p_{ik})_{k=1}^K$ 
5      $(c_{ik}^{(n+1)})_{k=1}^K = \max\{\|\mathbf{G}^{1/2} \mathbf{p}_i\|_2 - \tau \lambda w_i, 0\} \frac{\mathbf{p}_i}{\|\mathbf{G}^{1/2} \mathbf{p}_i\|_2}$ 
6   end
7    $\mathbf{C}^{(n+1)} = (c_{ik}^{(n+1)})_{i,k=1}^{N,K}$ 
8    $\mathbf{Q} = \mathbf{\Xi}^{(n)} + \sigma \mathbf{A} (2\mathbf{C}^{(n+1)} - \mathbf{C}^{(n)}) - \sigma \mathbf{B}$ 
9    $\mathbf{\Xi}^{(n+1)} = \min\left\{1, \frac{1}{\|\mathbf{Q}\mathbf{G}^{1/2}\|_{2,2}}\right\} \mathbf{Q}$ 
10   $\bar{\mathbf{C}}^{(n+1)} = \frac{n}{n+1} \bar{\mathbf{C}}^{(n)} + \frac{1}{n+1} \mathbf{C}^{(n+1)}$ 
11 end
12  $\bar{\mathbf{C}} = \bar{\mathbf{C}}^{(T)}$ 

```

Having done this, we next adapt Algorithm 1 in the way mentioned previously to obtain an algorithm for (4.8). This is given in Algorithm 2.

Remark 4.2. Note that even though the square-root matrix $\mathbf{G}^{1/2}$ is used in Algorithm 2, this matrix does not need to be computed. Indeed,

$$\|\mathbf{G}^{1/2} \mathbf{d}\|_2 = \sqrt{\mathbf{d}^* \mathbf{G} \mathbf{d}}, \quad \mathbf{d} \in \mathbb{C}^K,$$

and for a matrix $\mathbf{C} \in \mathbb{C}^{N \times K}$, we have

$$\|\mathbf{C}\mathbf{G}^{1/2}\|_{2,2} = \sqrt{\sum_{i=1}^N \|\mathbf{G}^{1/2} \mathbf{c}_i\|_2^2} = \sqrt{\sum_{i=1}^N \mathbf{c}_i^* \mathbf{G} \mathbf{c}_i},$$

where $\mathbf{c}_i \in \mathbb{C}^K$ is the i th row of \mathbf{C} . In particular, computing $\|\mathbf{G}^{1/2} \mathbf{d}\|$ involves

at most $c(F(\mathbf{G}) + K)$ arithmetic operations, and computing $\|\mathbf{C}\mathbf{G}^{1/2}\|_{2,2}$ involves $cm(F(\mathbf{G}) + K)$ arithmetic operations, for some universal constant $c > 0$.

To conclude this section, we now state and prove a lemma on the computational cost of Algorithm 2. This will be used later when proving the main theorems.

Lemma 4.3 (Computational cost of Algorithm 2). *The computational cost of Algorithm 2 is bounded by*

$$c \cdot (m \cdot N \cdot K + (m + N) \cdot (F(\mathbf{G}) + K)) \cdot T,$$

where $c > 0$ is a universal constant.

Proof. We proceed line-by-line. Line 2 involves a matrix-matrix multiplication and matrix subtraction, for a total of at most

$$c \cdot m \cdot N \cdot K \quad (\text{line 2})$$

arithmetic operations for some universal constant c . Now consider lines 3–5. By the previous remark, we may calculate $\|\mathbf{G}^{1/2}\mathbf{p}_i\|_2 = \sqrt{\mathbf{p}_i^*\mathbf{G}\mathbf{p}_i}$ using one multiplication with the matrix \mathbf{G} , one inner product of vectors of length K and one square root (recall from Definition 3.1 that we count square roots as arithmetic operations). This involves at most $c \cdot (F(\mathbf{G}) + K)$ arithmetic operations. Hence, the cost of line 5 is at most

$$c \cdot (F(\mathbf{G}) + K) \quad (\text{line 5}),$$

for a possibly different universal constant c . Therefore, the total cost of lines 3–5 is

$$c \cdot (F(\mathbf{G}) + K) \cdot N \quad (\text{lines 3–5}).$$

Line 7 involves no arithmetic operations and line 8 involves at most

$$c \cdot m \cdot N \cdot K \quad (\text{line 8})$$

operations. Consider line 9. Because of the previous remark, the computation of the term $\|\mathbf{Q}\mathbf{G}^{1/2}\|_{2,2}$ can be performed in at most $c \cdot m \cdot (F(\mathbf{G}) + K)$ operations (since \mathbf{Q} is of size $m \times K$). Hence, line 9 involves at most

$$c \cdot m \cdot (F(\mathbf{G}) + K) \quad (\text{line 9})$$

operations. Finally, line 10 involves at most

$$c \cdot N \cdot K \quad (\text{line 10})$$

operations. After simplifying, we deduce that lines 2–10 involve at most

$$c \cdot (m \cdot N \cdot K + (K + F(\mathbf{G})) \cdot (N + m)) \quad (\text{lines 2–10})$$

operations. The result follows by multiplying this by the number of iterations T . ■

Algorithm 3: construct-A – constructing the measurement matrix (4.3)

inputs : sample points $y_1, \dots, y_m \in \mathcal{U}^d$, finite index set
 $\Lambda = \{\mathbf{v}_1, \dots, \mathbf{v}_N\} \subset \mathcal{F}$

output : $A = \text{construct-A}((y_i)_{i=1}^m, \Lambda) \in \mathbb{C}^{m \times N}$, the measurement matrix (4.3)

initialize: $\bar{C}^{(0)} = \mathbf{0} \in \mathbb{C}^{N \times K}$

- 1 $k = \max\{j : (\mathbf{v}_i)_j \neq 0, i = 1, \dots, N, j = 1, \dots, d\}$
- 2 $n = \max\{(\mathbf{v}_i)_j : i = 1, \dots, N, j = 1, \dots, n\}$
- 3 **for** $i = 1, \dots, m$ **do**
- 4 Set $\mathbf{z} = (z_j)_{j=1}^k = ((y_i)_j)_{j=1}^k$
- 5 $b_{ij} = \Psi_j(z_i), i = 1, \dots, k, j = 0, \dots, n,$
- 6 **for** $j = 1, \dots, N$ **do**
- 7 $a_{ij} = \prod_{l=1}^n b_{l, (\mathbf{v}_j)_l}$
- 8 **end**
- 9 **end**
- 10 $A = \frac{1}{\sqrt{m}} (a_{ij})_{i,j=1}^{m,N}$

4.5 The algorithms in Theorems 3.5, 3.8 and 3.11

We are now almost ready to specify the algorithms used in Theorems 3.5, 3.8 and 3.11. Notice that Algorithms 1 and 2 require the measurement matrix A as an input. Hence, we first describe the computation of this matrix for Chebyshev and Legendre polynomials. This is summarized in Algorithm 3. Notice that line 5 of this algorithm involves evaluating the first k one-dimensional Chebyshev or Legendre polynomials. This can be done efficiently via the three-term recurrence relation, as explained in the proof of the following result.

Lemma 4.4 (Computational cost of Algorithm 3). *The computational cost of Algorithm 3 is bounded by*

$$c \cdot m \cdot (n + N) \cdot k,$$

where $c > 0$ is a universal constant and k and n are as in lines 1 and 2 of the algorithm.

Proof. Consider line 5 of the algorithm. Evaluation of the first $k + 1$ Chebyshev or Legendre polynomials can be done via the three-term recurrence relation. In the Chebyshev case, this is given by

$$\Psi_0(z) = 1, \quad \Psi_1(z) = \sqrt{2}z, \quad \Psi_{j+1}(z) = 2z\Psi_j(z) - c_j\Psi_{j-1}(z), \quad j = 1, \dots, k,$$

-
- Let m, ϵ, n and t be as given in the particular theorem and set:
 - $\Lambda = \Lambda_{n,d}^{\text{HC}}$ (Theorems 3.4 and 3.10) or $\Lambda = \Lambda_n^{\text{HCl}}$ (Theorem 3.7),
 - $\lambda = (4\sqrt{m/L})^{-1}$, where $L = L(m, d, \epsilon)$ is as in (3.8),
 - $\tau = \sigma = (\Theta(n, d))^{-\alpha}$, where $\Theta(n, d)$ and α are as in (3.5) and (3.7), respectively,
 - $T = \lceil 2(\Theta(n, d))^{\alpha} t \rceil$.
 - Let $\mathbf{D} = (d_{ik})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ and $\mathbf{Y} = (y_i)_{i=1}^m$ be an input, as in (3.1), and set $\mathbf{B} = \frac{1}{\sqrt{m}} \mathbf{D}$.
 - Compute $\mathbf{A} = \text{construct-A}(\mathbf{Y}, \Lambda)$.
 - Let \mathbf{G} and \mathbf{w} be as in (3.3) and (4.7), respectively.
 - Define the output $\bar{\mathbf{C}} = \mathcal{A}(\mathbf{D})$, where

$$\mathcal{A}(\mathbf{D}) = \text{primal-dual-wSRLASSO-C}(\mathbf{A}, \mathbf{B}, \mathbf{w}, \mathbf{G}, \lambda, \tau, \sigma, T, \mathbf{0}, \mathbf{0})$$

Table 4.2. The algorithms $\mathcal{A} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K}$ used in Theorems 3.5, 3.8 and 3.11.

where $c_j = 1$ if $j \geq 1$ and $1/\sqrt{2}$ otherwise, and in the Legendre case, it is given by

$$\begin{aligned} \Psi_0(z) &= 1, & \Psi_1(z) &= \sqrt{3}z, \\ \Psi_{j+1}(z) &= \frac{\sqrt{j+3/2}}{j+1} \left(\frac{2j+1}{\sqrt{j+1/2}} z \Psi_j(z) - \frac{j}{\sqrt{j-1/2}} \Psi_{j-1}(z) \right), & j &= 2, \dots, k, \end{aligned}$$

(recall that these polynomials are normalized with respect to their respective probability measures). Hence, the computational cost for line 5 is bounded by $c \cdot n \cdot k$. The computational cost for lines 6–8 is precisely $N \cdot (k - 1)$. Hence, the computational cost for forming each row of \mathbf{A} is bounded by $c \cdot (n \cdot k + N \cdot k)$. The result now follows. \blacksquare

With this in hand, we are now ready to specify the algorithms used in Theorems 3.5, 3.8 and 3.11. These are given in Table 4.2.

4.6 An efficient restarting procedure for the primal-dual iteration and the algorithms used in Theorems 3.6, 3.9 and 3.12

While the primal-dual iteration converges under very general conditions, it typically does so very slowly, with the error in the objective function decreasing like $\mathcal{O}(1/n)$,

Algorithm 4: primal-dual-rst-wSRLASSO – the restarted primal-dual iteration for the weighted SR-LASSO problem (4.6)

inputs : measurement matrix $\mathbf{A} \in \mathbb{C}^{m \times N}$, measurements $\mathbf{b} \in \mathcal{V}_h^N$, positive weights $\mathbf{w} = (w_i)_{i=1}^N$, parameter $\lambda > 0$, stepsizes $\tau, \sigma > 0$, number of primal-dual iterations $T \geq 1$, number of restarts $R \geq 1$, tolerance $\zeta' > 0$, scale parameter $0 < r < 1$, constant $s > 0$, initial values $\mathbf{c}^{(0)} = \mathbf{0} \in \mathcal{V}_h^N$, $\boldsymbol{\xi}^{(0)} = \mathbf{0} \in \mathcal{V}_h^m$.

output : $\tilde{\mathbf{c}} = \text{primal-dual-rst-wSRLASSO}(\mathbf{A}, \mathbf{b}, \mathbf{w}, \lambda, \tau, \sigma, T, R, \zeta', r, s)$, an approximate minimizer of (4.6)

initialize: $\tilde{\mathbf{c}}^{(0)} = \mathbf{0} \in \mathcal{V}_h^N$, $\varepsilon_0 = \|\mathbf{b}\|_{2;\mathcal{V}}$

1 **for** $l = 0, \dots, R - 1$ **do**

2 $\varepsilon_{l+1} = r(\varepsilon_l + \zeta')$

3 $a_l = s\varepsilon_{l+1}$

4 $\tilde{\mathbf{c}}^{(l+1)} = a_l \cdot \text{primal-dual-wSRLASSO}(\mathbf{A}, \mathbf{b}/a_l, \mathbf{w}, \lambda, \tau, \sigma, T, \tilde{\mathbf{c}}^{(l)}/a_l, \mathbf{0})$

5 **end**

6 $\tilde{\mathbf{c}} = \tilde{\mathbf{c}}^{(R)}$

where n is the iteration number. To obtain exponential convergence (down to some controlled tolerance) we employ a restarting procedure. This is based on recent work of [47, 48].

Restarting is a general concept in optimization, where the output of an algorithm after a fixed number of steps is then fed into the algorithm as input, after suitably scaling the parameters of the algorithm [122–124]. In the case of the primal-dual iteration for the weighted SR-LASSO problem, this procedure involves three hyperparameters: a *tolerance* $\zeta' > 0$ and *scale* parameters $0 < r < 1$ and $s > 0$. After applying one step of the primal-dual iteration (Algorithms 1 or 2) yielding an output $\mathbf{c}^{(1)}$, it then scales this vector and the right-hand side vector \mathbf{b} by an exponentially decaying factor a_l (defined in terms of ζ' , r and s), before feeding in these values into the primal-dual iteration as input.

We explain the motivations behind the specific form of the restart procedure for the primal-dual iteration later in Section 9.2. For now, we simply state the resulting procedures in the case of the weighted SR-LASSO problems (4.6) and (4.8). These are given in Algorithms 4 and 5, respectively. With these in hand, we can also give the algorithms used in Theorems 3.6, 3.9 and 3.12. See Table 4.3.

Note that these algorithms involve a number c^* , which is a universal constant. It is possible to provide a precise numerical value of this constant by carefully tracking the constants in several of the proof steps. Since doing so is not especially illuminative,

Algorithm 5: primal-dual-rst-wSRLASSO-C – the restarted primal-dual iteration for the weighted SR-LASSO problem (4.8)

inputs : measurement matrix $\mathbf{A} \in \mathbb{C}^{m \times N}$, measurements $\mathbf{B} \in \mathbb{C}^{N \times K}$, positive weights $\mathbf{w} = (w_i)_{i=1}^N$, Gram matrix $\mathbf{G} \in \mathbb{C}^{K \times K}$, parameter $\lambda > 0$, stepsizes $\tau, \sigma > 0$, number of primal-dual iterations $T \geq 1$, number of restarts $R \geq 1$, tolerance $\zeta' > 0$, scale parameter $0 < r < 1$, constant $s > 0$, initial values $\mathbf{C}^{(0)} = \mathbf{0} \in \mathbb{C}^{N \times K}$, $\mathbf{\Xi}^{(0)} = \mathbf{0} \in \mathbb{C}^{m \times K}$

output : $\tilde{\mathbf{C}} =$
primal-dual-rst-wSRLASSO-C($\mathbf{A}, \mathbf{b}, \mathbf{w}, \mathbf{G}, \lambda, \tau, \sigma, T, R, \zeta', r, s$),
an approximate minimizer of (4.8)

initialize: $\tilde{\mathbf{C}}^{(0)} = \mathbf{0} \in \mathbb{C}^{N \times K}$, $\varepsilon_0 = \|\mathbf{B}\mathbf{G}^{1/2}\|_{2;2}$

- 1 **for** $l = 0, \dots, R - 1$ **do**
- 2 $\varepsilon_{l+1} = r(\varepsilon_l + \zeta)$
- 3 $a_l = s\varepsilon_{l+1}$
- 4 $\tilde{\mathbf{C}}^{(l+1)} =$
 $a_l \cdot \text{primal-dual-wSRLASSO-C}(\mathbf{A}, \mathbf{B}/a_l, \mathbf{w}, \mathbf{G}, \lambda, \tau, \sigma, T, \tilde{\mathbf{C}}^{(l)}/a_l, \mathbf{0})$
- 5 **end**
- 6 $\tilde{\mathbf{C}} = \tilde{\mathbf{C}}^{(R)}$

we forgo this additional effort. Instead, we now give a little more detail on this constant.

Remark 4.5. From (10.10) we see that

$$c^* = 3296\sqrt{c_0},$$

where c_0 is the universal constant that arises in (3.10). As shown in the proof of Theorem 8.2, the constant c_0 needs to be chosen sufficiently large so that the measurement matrix \mathbf{A} satisfies the so-called *weighted RIP*. In particular, it is related to the universal constant $c > 0$ defined in Lemma 8.1. See, in particular, (8.2). A numerical value for this constant can indeed be found using results shown in [37]. With this in hand, one can then keep track of the constant c_0 in the proof of Theorem 8.2 to find its numerical value. This discussion also highlights why tracking the value of c^* is not particularly illuminative. Indeed, it is well known that universal constants appearing in RIP estimates in compressed sensing are generally very pessimistic [8, 13, 61].

-
- Let m, ϵ, n, t and ζ' be as given in the particular theorem and set:
 - $\Lambda = \Lambda_{n,d}^{\text{HC}}$ (Theorems 3.6 and 3.12) or $\Lambda = \Lambda_n^{\text{HCl}}$ (Theorem 3.9),
 - $\lambda = (4\sqrt{m/L})^{-1}$, where $L = L(m, d, \epsilon)$ is as in (3.8),
 - $\tau = \sigma = (\Theta(n, d))^{-\alpha}$, where $\Theta(n, d)$ and α are as in (3.5) and (3.7), respectively,
 - $T = \lceil (\Theta(n, d))^{\alpha} c^* \rceil$, where c^* is a universal constant,
 - $R = t$
 - $r = e^{-1}$
 - $s = \frac{(\Theta(n,d))^{\alpha} T}{2}$
 - Let $\mathbf{D} = (d_{ik})_{i,k=1}^{m,K} \in \mathbb{C}^{m \times K}$ and $\mathbf{Y} = (y_i)_{i=1}^m$ be an input, as in (3.1), and set $\mathbf{B} = \frac{1}{\sqrt{m}} \mathbf{D}$.
 - Compute $\mathbf{A} = \text{construct-A}(\mathbf{Y}, \Lambda)$.
 - Let \mathbf{G}, \mathbf{A} and \mathbf{w} be as in (4.3), (3.3) and (4.7), respectively.
 - Define the output $\tilde{\mathbf{C}} = \mathcal{A}(\mathbf{D})$, where

$$\mathcal{A}(\mathbf{D}) = \text{primal-dual-rst-wSRLASSO-C}(\mathbf{A}, \mathbf{B}, \mathbf{w}, \mathbf{G}, \lambda, \tau, \sigma, T, R, \zeta, r, c)$$

Table 4.3. The algorithms $\mathcal{A} : \mathcal{U}^m \times \mathbb{C}^{m \times K} \rightarrow \mathbb{C}^{N \times K}$ used in Theorems 3.6, 3.9 and 3.12.

Chapter 5

Numerical experiments

In this chapter, we present a number of numerical experiments to examine the practical performance of the previously developed algorithms.

5.1 Experimental setup

We first describe the experimental setup.

5.1.1 Hyperparameter values

The algorithms used in the main theorems (see Tables 4.2 and 4.3) are designed to ensure the desired error bounds. In our numerical experiments, we deviate from these values in a number of minor ways. However, our hyperparameter choices are still closely based on theory. We now discuss the precise hyperparameter choices used in the experiments. These choices are also summarized in Table 5.1.

First, we take the parameter λ to be $\lambda = (\sqrt{25m})^{-1}$. This differs somewhat from the value $\lambda = (4\sqrt{m/L})^{-1}$ used in the theoretical algorithms. The rationale behind doing this is that L is, in practice, a polylogarithmic factor that arises from the compressed sensing theory. It is well known that logarithmic factors appearing in compressed sensing theory are generally quite pessimistic [8, 13, 61]. Therefore, we avoid using L . The choice $\lambda = (5\sqrt{m})^{-1}$ was obtained in [8, Appendix A] after manual tuning.

As shown later, the primal-dual iteration converges subject to the condition $\|\mathbf{A}\|_2^2 \leq (\tau\sigma)^{-1}$. See Lemma 9.2. Since the error bound (9.2) scales linearly in τ^{-1} and σ^{-1} , a standard choice for these parameters is

$$\tau = \sigma = 1/\|\mathbf{A}\|_2. \tag{5.1}$$

In Tables 4.2 and 4.3 we choose

$$\tau = \sigma = (\Theta(n, d))^{-\alpha},$$

since the latter is an upper bound for $\|\mathbf{A}\|_2$, i.e., $\|\mathbf{A}\|_2 \leq (\Theta(n, d))^\alpha$. See (10.9). This bound is arguably quite crude. The reason for using it in our main theorems is to avoid having to compute $\|\mathbf{A}\|_2$, since this generally cannot be done in finitely many arithmetic operations. However, in our numerical experiments we simply use (5.1) instead, as it is simpler and $\|\mathbf{A}\|_2$ can be efficiently approximated in practice.

Parameter	Value	Notes
λ	$(\sqrt{25m})^{-1}$	Based on [8, Appendix A]
σ	$\ \mathbf{A}\ _2^{-1}$	Based on Lemma 9.2
τ	$\ \mathbf{A}\ _2^{-1}$	Based on Lemma 9.2
r	e^{-1}	Based on Theorem 9.4
T	$\left\lceil \frac{2\ \mathbf{A}\ _2}{r} \right\rceil$	Based on Theorem 9.4, assuming $C = 1$
s	$\frac{T}{2\ \mathbf{A}\ _2}$	Based on Theorem 9.4

Table 5.1. Hyperparameter values used in the numerical experiments. The first three parameters are used in both the unrestarted and restarted primal-dual iterations. The final three parameters are used in the restarted scheme only.

For the restarting scheme, we also have the scale parameter $0 < r < 1$, the constant $s > 0$ and the number of inner iterations T . These parameters are inferred from Theorem 9.4. This result shows that the error in the restarted primal dual iteration after l restarts is bounded by

$$r^l \|\mathbf{b}\|_{2;\mathcal{V}} + \frac{r}{1-r} \zeta', \quad (5.2)$$

provided

$$T = \left\lceil \frac{2C}{r\sqrt{\sigma\tau}} \right\rceil, \quad a_l = \frac{1}{2} \sigma \varepsilon_{l+1} T, \quad l = 0, 2, \dots$$

Here, as discussed in Theorem 9.4, $C > 0$ is a numerical constant that arises from the compressed sensing theory. This and the choice (5.1) leads immediately to the following value for s :

$$s = \frac{T}{2\|\mathbf{A}\|_2}.$$

Unfortunately, the constant C is difficult to determine exactly (it is closely related to the constant c^* discussed in Remark 4.5). In our experiments, we simply pick the value $C = 1$. This immediately yields

$$T = \left\lceil \frac{2\|\mathbf{A}\|_2}{r} \right\rceil.$$

Finally, to determine a value of r we consider the error bound (5.2). This argument is based on [48]. After l restarts, the total number of iterations $t = Tl$. Substituting the value of T , we see that

$$r^l = \exp(\log(r)t/T) = \exp\left(\log(r) \left\lceil \frac{2\|\mathbf{A}\|_2}{r} \right\rceil^{-1} t\right). \quad (5.3)$$

Ignoring the ceiling function, it therefore makes sense to choose $0 < r < 1$ to minimize the function $r \mapsto r \log(r)$. This attains its minimum value of $-e^{-1}$ at $r = e^{-1}$, which is the value we now use.

5.1.2 Test functions

We consider four test functions. The first two are scalar-valued functions, given by

$$f_1(\mathbf{y}) = \exp\left(-\frac{1}{2d} \sum_{k=1}^d y_k\right), \quad \forall \mathbf{y} \in \mathcal{U}, \quad \text{with } d = 2, \quad (5.4)$$

and

$$f_2(\mathbf{y}) = \exp\left(-\frac{2}{d} \sum_{k=1}^d (y_k - w_k)^2\right), \quad \forall \mathbf{y} \in \mathcal{U}, \quad (5.5)$$

$$\text{with } w_k = \frac{(-1)^k}{k+1}, \quad \forall k \in [d] \text{ and } d = 16.$$

These are standard test functions (see, e.g., [8, Appendix A.1]). The first function varies very little with respect to \mathbf{y} . Hence, it is expected to be very well approximated by a sparse polynomial approximation. The second has more variation in \mathbf{y} , therefore we expect a larger approximation error.

We also consider two Hilbert-valued functions. These both arise as solutions of the parametric elliptic diffusion equation

$$\begin{aligned} -\nabla \cdot (a(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{x}, \mathbf{y})) &= g(\mathbf{x}), \quad \forall \mathbf{x} \in D, \mathbf{y} \in \mathcal{U}, \\ u(\mathbf{x}, \mathbf{y}) &= 0, \quad \forall \mathbf{x} \in \partial D, \mathbf{y} \in \mathcal{U}, \end{aligned}$$

which is a standard problem in the parametric PDE literature. We take the physical domain D as

$$D = (0, 1)^2.$$

For simplicity, we also choose

$$g(\mathbf{x}) = 10$$

to be constant. In this case, the solution map

$$\mathcal{U} \rightarrow \mathcal{V}, \quad \mathbf{y} \mapsto u(\cdot, \mathbf{y}), \quad \mathcal{V} = H_0^1(D),$$

is a Hilbert-valued function with codomain being the Sobolev space $H_0^1(D)$. We consider two different setups, leading to smooth and less smooth Hilbert-valued functions, which we denote as f_3 and f_4 , respectively. The first is a simple two-dimensional problem with lognormal diffusion coefficient:

$$f_3: \quad d = 2, \quad a(\mathbf{x}, \mathbf{y}) = 5 + \exp(x_1 y_1 + x_2 y_2). \quad (5.6)$$

For the second, we consider the diffusion coefficient from [5, equation (24)], modified from an earlier example from [110, equation (5.2)], with 30-dimensional parametric dependence and one-dimensional (layered) spatial dependence given by

$$f_4 : d = 30, a(x, y) = \exp\left(1 + y_1 \left(\frac{\sqrt{\pi}\beta}{2}\right)^{1/2} + \sum_{i=2}^d \zeta_i \vartheta_i(x) y_i\right),$$

$$\zeta_i := (\sqrt{\pi}\beta)^{1/2} \exp\left(\frac{-(\lfloor \frac{i}{2} \rfloor \pi \beta)^2}{8}\right), \quad \vartheta_i(x) := \begin{cases} \sin(\lfloor \frac{i}{2} \rfloor \pi x_1 / \beta_p) & i \text{ even,} \\ \cos(\lfloor \frac{i}{2} \rfloor \pi x_1 / \beta_p) & i \text{ odd,} \end{cases}$$

$$\beta_c = 1/8, \quad \beta_p = \max\{1, 2\beta_c\}, \quad \beta = \beta_c / \beta_p. \quad (5.7)$$

5.1.3 Error metrics and finite element discretization

In our experiments, we consider the relative $L^2_{\bar{\theta}}(\mathcal{U})$ -norm error

$$\frac{\|f - \hat{f}\|_{L^2_{\bar{\theta}}(\mathcal{U})}}{\|f\|_{L^2_{\bar{\theta}}(\mathcal{U})}}, \quad (5.8)$$

for the scalar-valued functions f_1 and f_2 and the relative $L^2_{\bar{\theta}}(\mathcal{U}; H_0^1(D))$ -norm error

$$\frac{\|f - \hat{f}\|_{L^2_{\bar{\theta}}(\mathcal{U}; H_0^1(D))}}{\|f\|_{L^2_{\bar{\theta}}(\mathcal{U}; H_0^1(D))}}, \quad (5.9)$$

for the Hilbert-valued functions f_3 and f_4 . To (approximately) compute this error we use a high-order isotropic Smolyak sparse grid quadrature rule based on Clenshaw–Curtis points. This rule is generated using the TASMANIAN software package [129]. We set the level of the quadrature rule in each experiment as large as possible within the constraints of computational time and memory.

We now describe the discretization \mathcal{V}_h for the Hilbert-valued functions f_3 and f_4 . This is obtained via the finite element method as implemented by Dolfin [95], and accessed through the python FEniCS project [18]. We generate a regular triangulation \mathcal{T}_h of \bar{D} composed of triangles T of equal diameter $h_T = h$. We consider a conforming discretization, which results in a finite-dimensional subspace $\mathcal{V}_h \subset \mathcal{V} = H_0^1(D)$, where \mathcal{V}_h is the space spanned by the usual Lagrange finite elements $\{\varphi_i\}_{i=1}^K$ of order $k = 1$. We rely on the Dolfin `UnitSquareMesh` method to generate a mesh with 33 nodes per side, corresponding to a finite element triangulation with $K = 1089$ nodes, 2048 elements and meshsize $h = \sqrt{2}/32$. See [5, 52] for further implementation details.

Explicit forms of the Hilbert-valued functions f_3 and f_4 are not available. Therefore, computing the relative error requires first computing a reference solution. This is




Figure 5.1. Approximation error versus iteration number for the function f_1 from (5.4). This figure shows the relative L^2 errors of the polynomial approximations obtained from (left) the iterates $e^{(n)}$ and (right) the ergodic sequence $\bar{e}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and $m = 250$ sample points drawn randomly and independently from the uniform measure. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 2$ and $n = 184$, which gives a basis of cardinality $N = |\Lambda| = 997$. We compare the primal dual iteration “PD” and the restarted primal dual iteration “PDR” for various values of the tolerance ζ' . We also plot the theoretical error curve (5.10), where t is the iteration number. The quadrature rule used to compute the relative error is a sparse grid rule of level 11 consisting of $M = 7169$ points.

usually done by using a finite element discretization with meshsize an order of magnitude smaller than that used to compute the various approximations. However, our main focus in these experiments is on the polynomial approximation and algorithmic errors E_{app} and E_{alg} . Since our theoretical results assert that the approximations are robust to physical discretization error, we do not perform this additional (and costly) computational step. Instead, we compute reference solutions using the same finite element discretization as that used to construct the various approximations. In other words, there is no physical discretization error present in these experiments.

5.2 Numerical results 1: The optimization error

Our first experiments, Figures 5.1–5.4, compare the behaviour of the unrestarted primal-dual iteration to the restarted primal-dual iteration with several different values of the tolerance parameter ζ' . In all cases, we observe a consistent improvement from the restarted scheme. This is particularly noticeable for the functions f_1 and f_3 , since the underlying approximation error ζ is smaller in these cases. Recall that these functions are well approximated by polynomials. As predicted by our theoretical results,

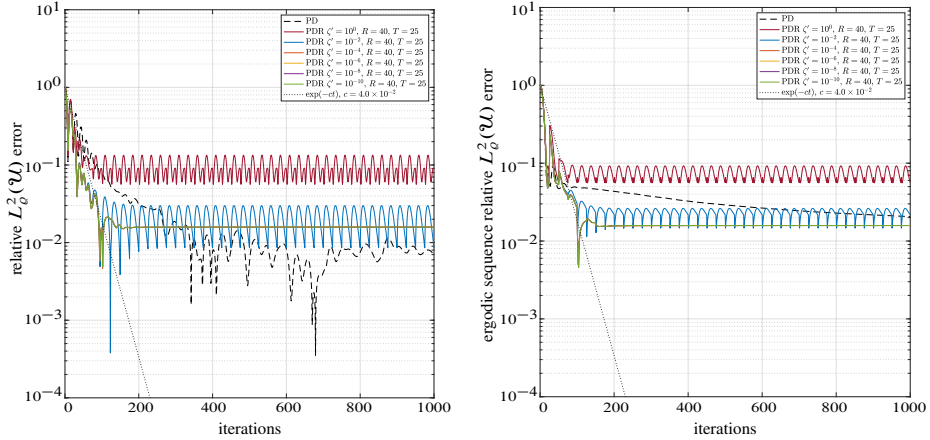


Figure 5.2. Approximation error versus iteration number for the function f_2 from (5.5). This figure shows the relative L^2 errors of the polynomial approximations obtained from (left) the iterates $c^{(n)}$ and (right) the ergodic sequence $\bar{c}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and $m = 2000$ sample points drawn randomly and independently from the uniform measure. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 16$ and $n = 16$, which gives a basis of cardinality $N = |\Lambda| = 8277$. We compare the primal dual iteration “PD” and the restarted primal dual iteration “PDR” for various values of the tolerance ζ' . We also plot the theoretical error curve (5.10), where t is the iteration number. The quadrature rule used to compute the relative error is a sparse grid rule of level 5 consisting of $M = 51137$ points.

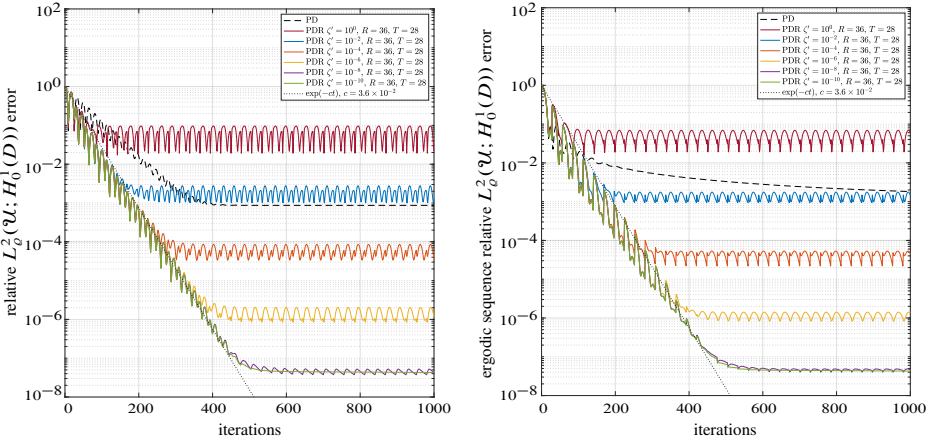


Figure 5.3. Approximation error versus iteration number for the function f_3 from (5.6). This figure shows the relative L^2 errors of the polynomial approximations obtained from (left) the iterates $c^{(n)}$ and (right) the ergodic sequence $\bar{c}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and $m = 250$ sample points drawn randomly and independently from the uniform measure. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 2$ and $n = 184$, which gives a basis of cardinality $N = |\Lambda| = 997$. We compare the primal dual iteration “PD” and the restarted primal dual iteration “PDR” for various values of the tolerance ζ' . We also plot the theoretical error curve (5.10), where t is the iteration number. The quadrature rule used to compute the relative error is a sparse grid rule of level 9 consisting of $M = 1537$ points.




Figure 5.4. Approximation error versus iteration number for the function f_4 from (5.7). This figure shows the relative L^2 errors of the polynomial approximations obtained from (left) the iterates $c^{(n)}$ and (right) the ergodic sequence $\bar{c}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and $m = 1000$ sample points drawn randomly and independently from the uniform measure. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 30$ and $n = 10$, which gives a basis of cardinality $N = |\Lambda| = 7841$. We compare the primal dual iteration “PD” and the restarted primal dual iteration “PDR” for various values of the tolerance ζ' . We also plot the theoretical error curve (5.10), where t is the iteration number. The quadrature rule used to compute the relative error is a sparse grid rule of level 3 consisting of $M = 1861$ points.

the error for the restarted scheme decays exponentially fast with respect to the number of iterations to this limiting accuracy. For example, in the case of f_1 the restarted scheme (with sufficiently small ζ') achieves a relative error of less than 10^{-6} using only 500 iterations. However, the unrestarted scheme only achieves an error of around 10^{-3} after 1000 iterations.

An important takeaway from these experiments is the insensitivity of the algorithm to the parameter ζ' . Our theoretical results only show exponential convergence (with respect to iteration number) when $\zeta' \geq \zeta$, where ζ is a certain upper bound for the error. This appears unnecessary in practice. For instance, in Figures 5.2 and 5.4 we expect the underlying error ζ to be roughly 10^{-2} in magnitude, since this is the limiting error achieved by the unrestarted scheme. Yet setting $\zeta' = 10^{-10}$ has no noticeable effect on the performance of the restarted scheme. Moreover, for $\zeta' \in \{10^{-4}, 10^{-6}, 10^{-8}, 10^{-10}\}$ the results are nearly identical in both Figures 5.2 and 5.4, and hence the plot lines are overlaid in the case of the restarted scheme.

Another noticeable feature of these experiments is the close agreement between the theorized rate of exponential decay of the restarted scheme, which is given by the right-hand side of (5.3) and what is observed in practice. Since the value $r = e^{-1}$ is




Figure 5.5. (left) Approximation error and (right) average run time versus number of samples m for the function f_1 from (5.4). This figure shows the relative L^2 errors of the polynomial approximations obtained from the ergodic sequence $\bar{c}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and various sets of m sample points drawn randomly and independently from the uniform measure for each trial. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 2$ and $n = 184$, which gives a basis of cardinality $N = |\Lambda| = 997$. We use the restarted primal dual iteration “PDR” with $\zeta' = 10^{-8}$, and display the average error over 50 trials measured in the sample mean in blue and the corrected sample standard deviation after a log transformation in shaded blue, see [8, Appendix A.1.3] for more details. The quadrature rule used to compute the relative error is a sparse grid rule of level 11 consisting of $M = 7169$ points.

used in these experiments, in Figures 5.1–5.4 we also plot the function

$$\exp(-ct), \quad c := \lceil 2e\|A\|_2 \rceil^{-1} \quad (5.10)$$

versus the iteration number t . This theoretical curve exactly predicts the observed rate of exponential decay of the restarted schemes.

Finally, in all four figures we also show the error of the (restarted) primal-dual iterates, as well as the ergodic sequences. Despite the theoretical results only holding for the latter, we see similar error decay for the iterates. In fact, the iterates give slightly better performance in the case of the unrestarted scheme. As expected, the ergodic sequence PDR reduces the variation in the error for the restarted scheme. Moreover, plotting the ergodic sequence we can see more clearly the benefit of using restarts over not restarting.

5.3 Numerical results 2: Approximation error and run time

In the second set of experiments, our aim is to study the approximation error versus the number of samples m . Having compared different solvers in the previous




Figure 5.6. (left) Approximation error and (right) average run time versus number of samples m for the function f_2 from (5.5). This figure shows the relative L^2 errors of the polynomial approximations obtained from the ergodic sequence $\tilde{\mathbf{c}}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and various sets of m sample points drawn randomly and independently from the uniform measure for each trial. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 16$ and $n = 16$, which gives a basis of cardinality $N = |\Lambda| = 8277$. We use the restarted primal dual iteration “PDR” with $\zeta' = 10^{-4}$, and display the average error over 50 trials measured in the sample mean in blue and the corrected sample standard deviation after a log transformation in shaded blue, see [8, Appendix A.1.3] for more details. The quadrature rule used to compute the relative error is a sparse grid rule of level 5 consisting of $M = 51137$ points.

experiments, we now limit our attention to the restarted primal-dual iteration. The only modification we make is to introduce a stopping criterion for the number of restarts. Specifically, given a tolerance ζ' , we halt the iteration if the difference between two consecutive iterates is less than $5 \cdot \zeta'$. Specifically, if

$$\|\tilde{\mathbf{c}}^{(l)} - \tilde{\mathbf{c}}^{(l-1)}\|_2 \leq 5 \cdot \zeta',$$

in the scalar-valued case or

$$\|\tilde{\mathbf{c}}^{(l)} - \tilde{\mathbf{c}}^{(l-1)}\|_{2;\mathbf{v}} \leq 5 \cdot \zeta',$$

in the Hilbert-valued case, where $\tilde{\mathbf{c}}^{(l)}$ is the output of the restarted primal-dual iteration after l restarts, then we halt and take $\tilde{\mathbf{c}}^{(l)}$ as the polynomial coefficients of the resulting approximation.

In the following experiments, we perform multiple trials for each value of m . For each trial, we generate a set of sample Monte Carlo points $\mathbf{y}_1, \dots, \mathbf{y}_m$, then compute the relative error (5.8) or (5.9) of the approximation using a sparse grid quadrature as before. Having done this, we then compute the sample mean and (corrected) sample




Figure 5.7. Approximation error versus number of samples m for the function f_3 from (5.6). This figure shows the relative L^2 errors of the polynomial approximations obtained from the ergodic sequence $\bar{c}^{(m)}$. These approximations are constructed using the Legendre polynomial basis and various sets of m sample points drawn randomly and independently from the uniform measure for each trial. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 2$ and $n = 184$, which gives a basis of cardinality $N = |\Lambda| = 997$. We compare the restarted primal dual iteration “PDR” with $\zeta' = 10^{-8}$ with the average performance over 50 trials measured in the sample mean in blue and the corrected sample standard deviation after a log transformation in shaded blue, see [8, Appendix A.1.3] for more details. The quadrature rule used to compute the relative error is a sparse grid rule of level 11 consisting of $M = 7169$ points.

standard deviation after a log transformation. See [8, Appendix A.1.3] for further discussion and rationale behind this computation.

The results for the four functions f_1, f_2, f_3, f_4 are shown in Figures 5.5–5.8. Figure 5.5 shows the average approximation error and run times for f_1 . As discussed, this function is expected to be well approximated by polynomials. In accordance, the error decreases rapidly, achieving roughly 10^{-7} relative L^2 error when $m \approx 200$. This is in broad agreement with the exponential decay rate of the error shown in our main theorems. In Figure 5.6 we consider the more challenging, higher-dimensional function f_2 , plotting the average approximation error and run time. Here, as expected, the error decreases significantly more slowly. Both figures exhibit a linear scaling of the run time with the number of samples m . This is consistent with our analysis, since




Figure 5.8. Approximation error versus number of samples m for the function f_4 from (5.7). This figure shows the relative L^2 errors of the polynomial approximations obtained from the ergodic sequence $\bar{c}^{(n)}$. These approximations are constructed using the Legendre polynomial basis and various sets of m sample points drawn randomly and independently from the uniform measure for each trial. The index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where $d = 30$ and $n = 10$, which gives a basis of cardinality $N = |\Lambda| = 7841$. We compare the restarted primal dual iteration “PDR” with $\zeta' = 10^{-4}$ with the average performance over 50 trials measured in the sample mean in blue and the corrected sample standard deviation after a log transformation in shaded blue, see [8, Appendix A.1.3] for more details. The quadrature rule used to compute the relative error is a sparse grid rule of level 3 consisting of $M = 1861$ points.

each algorithm iteration involves dense matrix-vector multiplications with an $m \times N$ matrix. Also, comparing Figures 5.5 and 5.6 when $m = 250$, we notice the run time is roughly 16 times larger for the latter. This is also in agreement with our analysis. Indeed, $N \approx 1000$ in Figure 5.5 while $N \approx 8000$ in Figure 5.6. However, the number of inner iterations $T = \lceil 2\|A\|_2/r \rceil$ is roughly twice as large in Figure 5.6, where $\|A\|_2 \approx 13$ when $m = 250$, as it is in Figure 5.5, where $\|A\|_2 \approx 7$. The combination of these two factors accounts for the roughly 16-fold increase in run time.

Figure 5.7 displays the performance of the restarted scheme on the Hilbert-valued function f_3 . Here we also observe rapid decrease in the error with respect to increasing number of samples m , with relative L^2 error approximately 10^{-6} when $m \approx 200$. Finally, Figure 5.8 shows the results for the less smooth high-dimensional Hilbert-valued function f_4 . For this function, we expect slower decrease in the error with

respect to m , which is reflected in this set of results. Nonetheless, despite its high dimensionality ($d = 30$) we still achieve two digits of relative accuracy using only $m \approx 1000$ samples.

Chapter 6

Overview of the proofs

The rest of this work is devoted to proving the main results. Since these involve a number of technical steps, we now give a brief overview of how these proofs proceed.

We commence in Chapter 7 by developing compressed sensing theory for Hilbert-valued vectors. We introduced the so-called *weighted robust Null Space Property (rNSP) over \mathcal{V}* , and then show in Lemma 7.4 that it implies certain error bounds for inexact minimizers of the Hilbert-valued, weighted SR-LASSO problem. Next, we introduced the *weighted Restricted Isometry Property (RIP)* and then in Lemma 7.6 we show that this property over \mathbb{C} implies the weighted rNSP over \mathcal{V} .

In Chapter 8 we focus on the polynomial approximation problem. We first give a sufficient condition in terms of m for the measurement matrix (4.3) to satisfy the weighted RIP with high probability (Lemma 8.1). Next, we state and prove three general results (Theorems 8.2–8.4) that give error bounds for polynomial approximations obtained as inexact minimizers of the Hilbert-valued, weighted SR-LASSO problem. These results are split into the three cases considered in our main results, i.e., the algebraic and finite-dimensional case, the algebraic and infinite-dimensional case, and the exponential case. The error bounds in these results split into terms corresponding to the polynomial approximation error, the physical discretization error, the sampling error, and the error in the objective function at the inexact minimizer.

With this in mind, in Chapter 9, we first present error bounds for inexact minimizers obtained by finitely many iterations of the primal-dual iteration. See Lemma 9.2. Having done this, we then have the ingredients needed to derive the restarting scheme. We derive this scheme and present an error bound for it in Theorem 9.4.

We conclude with in Chapter 10 with the final arguments. We use the three key theorems (Theorems 8.2–8.4) and then proceed to estimate each of the aforementioned error terms. For the polynomial approximation error we appeal to several results that are given in Appendix A. For the error in the objective function we use the results shown in Chapter 9. After straightforwardly bounding the other two error terms, we finally obtain the main results.

Chapter 7

Hilbert-valued compressed sensing

In this chapter, we develop Hilbert-valued compressed sensing theory. Here, rather than the classical setting of a vector in \mathbb{C}^N , one seeks to recover a Hilbert-valued vector in \mathcal{V}^N . This was considered in [52] in the for the classical sparsity model with ℓ^1 -minimization. We now consider the more general weighted sparsity model and weighted ℓ^1 -minimization. This model was first developed in [121] in the scalar-valued case. See also [2, 37] and [8, Chapter 6]. Note that in this chapter, we shall write \mathcal{V} rather than \mathcal{V}_h , as is done in (4.6). Since \mathcal{V} is arbitrary, all the results shown below will also apply in the case of \mathcal{V}_h .

7.1 Weighted sparsity and weighted best approximation

Let $\Lambda \subseteq \mathcal{F}$ and $\mathbf{w} = (w_\nu)_{\nu \in \Lambda} > \mathbf{0}$ be positive weights. Given a set $S \subseteq \Lambda$, we define its weighted cardinality as

$$|S|_{\mathbf{w}} := \sum_{i \in S} w_i^2.$$

The following two definitions extend Definitions 2.2 and 2.3 to the weighted setting.

Definition 7.1 (Weighted sparsity). Let $\Lambda \subseteq \mathcal{F}$. A \mathcal{V} -valued sequence $\mathbf{c} = (c_\nu)_{\nu \in \Lambda}$ is *weighted (k, \mathbf{w}) -sparse* for some $k \geq 0$ and weights $\mathbf{w} = (w_\nu)_{\nu \in \Lambda} > \mathbf{0}$ if

$$|\text{supp}(\mathbf{c})|_{\mathbf{w}} \leq k,$$

where $\text{supp}(\mathbf{z}) = \{\nu : \|z_\nu\|_{\mathcal{V}} \neq 0\}$ is the *support* of \mathbf{z} . The set of such vectors is denoted by $\Sigma_{k, \mathbf{w}}$.

Definition 7.2 (Weighted best (k, \mathbf{w}) -term approximation error). Let $\Lambda \subseteq \mathcal{F}$, $0 < p \leq 2$, $\mathbf{w} > \mathbf{0}$, $\mathbf{c} \in \ell_{\mathbf{w}}^p(\Lambda; \mathcal{V})$ and $k \geq 0$. The $\ell_{\mathbf{w}}^p$ -norm *weighted best (k, \mathbf{w}) -term approximation error* of \mathbf{c} is

$$\sigma_k(\mathbf{c})_{p, \mathbf{w}; \mathcal{V}} = \min\{\|\mathbf{c} - \mathbf{z}\|_{p, \mathbf{w}; \mathcal{V}} : \mathbf{z} \in \Sigma_{k, \mathbf{w}}\}.$$

Notice that this is equivalent to

$$\sigma_k(\mathbf{c})_{p, \mathbf{w}; \mathcal{V}} = \inf\{\|\mathbf{c} - \mathbf{c}_S\|_{p, \mathbf{w}; \mathcal{V}} : S \subseteq \Lambda, |S|_{\mathbf{w}} \leq k\}. \quad (7.1)$$

Here and elsewhere, for a sequence $\mathbf{c} = (c_\nu)_{\nu \in \Lambda}$ and a set $S \subseteq \Lambda$, we define \mathbf{c}_S as the sequence with ν th entry equal to c_ν if $\nu \in S$ and zero otherwise.

7.2 The weighted robust null space property

For the rest of this chapter, we consider the index set $\Lambda = \{1, \dots, N\}$ for some $N \in \mathbb{N}$. Our analysis of the weighted SR-LASSO problem is presented in terms of the so-called weighted robust null space property. Let $\mathbf{w} > \mathbf{0}$ and $k > 0$. A bounded linear operator $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ has the *weighted robust Null Space Property (rNSP)* over \mathcal{V} of order (k, \mathbf{w}) with constants $0 < \rho < 1$ and $\gamma > 0$ if

$$\|\mathbf{x}_S\|_{2;\mathcal{V}} \leq \frac{\rho \|\mathbf{x}_{S^c}\|_{1,\mathbf{w};\mathcal{V}}}{\sqrt{k}} + \gamma \|\mathbf{A}\mathbf{x}\|_{2;\mathcal{V}}, \quad \forall \mathbf{x} \in \mathcal{V}^N,$$

for any $S \subseteq [N]$ with $|S|_{\mathbf{w}} \leq k$.

Importantly, the weighted rNSP implies distance bounds in the $\ell_{\mathbf{w}}^1$ - and ℓ^2 -norms. The following lemma is standard in the scalar case (see, e.g., [8, Lemma 6.24]). We omit the proof of its extension to the Hilbert-valued case, since it follows almost exactly the same arguments.

Lemma 7.3 (Weighted rNSP implies $\ell_{\mathbf{w}}^1$ and ℓ^2 distance bounds). *Suppose that $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ has the weighted rNSP over \mathcal{V} of order (k, \mathbf{w}) with constants $0 < \rho < 1$ and $\gamma > 0$. Let $\mathbf{x}, \mathbf{z} \in \mathcal{V}^N$. Then*

$$\begin{aligned} \|\mathbf{z} - \mathbf{x}\|_{1,\mathbf{w};\mathcal{V}} &\leq C_1(2\sigma_k(\mathbf{x})_{1,\mathbf{w};\mathcal{V}} + \|\mathbf{z}\|_{1,\mathbf{w};\mathcal{V}} - \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}}) \\ &\quad + C_2\sqrt{k}\|\mathbf{A}(\mathbf{z} - \mathbf{x})\|_{2;\mathcal{V}}, \\ \|\mathbf{z} - \mathbf{x}\|_{2;\mathcal{V}} &\leq \frac{C'_1}{\sqrt{k}}(2\sigma_k(\mathbf{x})_{1,\mathbf{w};\mathcal{V}} + \|\mathbf{z}\|_{1,\mathbf{w};\mathcal{V}} - \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}}) \\ &\quad + C'_2\|\mathbf{A}(\mathbf{z} - \mathbf{x})\|_{2;\mathcal{V}}, \end{aligned}$$

where the constants are given by

$$C_1 = \frac{(1 + \rho)}{(1 - \rho)}, \quad C_2 = \frac{2\gamma}{(1 - \rho)}, \quad C'_1 = \left(\frac{(1 + \rho)^2}{1 - \rho} \right) \quad \text{and} \quad C'_2 = \left(\frac{(3 + \rho)\gamma}{1 - \rho} \right).$$

Lemma 7.3 can be used to show distance bounds for exact minimizers of the Hilbert-valued weighted SR-LASSO problem

$$\min_{\mathbf{z} \in \mathcal{V}^N} \mathcal{G}(\mathbf{z}), \quad \mathcal{G}(\mathbf{z}) := \lambda \|\mathbf{z}\|_{1,\mathbf{w};\mathcal{V}} + \|\mathbf{A}\mathbf{z} - \mathbf{b}\|_{2;\mathcal{V}}. \quad (7.2)$$

Fortunately, it also implies bounds for approximate minimizers, such as those obtained by a finite number of steps of the primal-dual iteration.

Lemma 7.4 (Weighted rNSP implies error bounds for inexact minimizers). *Suppose that $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ has the weighted rNSP over \mathcal{V} of order (k, \mathbf{w}) with constants $0 < \rho < 1$ and $\gamma > 0$. Let $\mathbf{x} \in \mathcal{V}^N$, $\mathbf{b} \in \mathcal{V}^m$ and $\mathbf{e} = \mathbf{A}\mathbf{x} - \mathbf{b} \in \mathcal{V}^m$, and consider*

the problem (7.2) with parameter

$$0 < \lambda \leq \frac{(1 + \rho)^2}{(3 + \rho)\gamma} k^{-1/2}. \quad (7.3)$$

Then, for any $\tilde{\mathbf{x}} \in \mathcal{V}^N$,

$$\begin{aligned} \|\tilde{\mathbf{x}} - \mathbf{x}\|_{1, \mathbf{w}; \mathcal{V}} &\leq C_1 \left(2\sigma_k(\mathbf{x})_{1, \mathbf{w}; \mathcal{V}} + \frac{\mathcal{G}(\tilde{\mathbf{x}}) - \mathcal{G}(\mathbf{x})}{\lambda} \right) + \left(\frac{C_1}{\lambda} + C_2 \sqrt{k} \right) \|\mathbf{e}\|_{2; \mathcal{V}}, \\ \|\tilde{\mathbf{x}} - \mathbf{x}\|_{2; \mathcal{V}} &\leq \frac{C'_1}{\sqrt{k}} \left(2\sigma_k(\mathbf{x})_{1, \mathbf{w}; \mathcal{V}} + \frac{\mathcal{G}(\tilde{\mathbf{x}}) - \mathcal{G}(\mathbf{x})}{\lambda} \right) + \left(\frac{C'_1}{\sqrt{k}\lambda} + C'_2 \right) \|\mathbf{e}\|_{2; \mathcal{V}}, \end{aligned}$$

where C_1, C_2, C'_1 and C'_2 are as in Lemma 7.3.

Proof. First notice that $C'_1/C'_2 \leq C_1/C_2$ since $0 < \rho < 1$, where C_1, C_2, C'_1 and C'_2 are as in Lemma 7.3. Hence the condition on λ implies that

$$\lambda \leq \min\{C_1/C_2, C'_1/C'_2\} k^{-1/2},$$

Using this lemma and this bound, we deduce that

$$\begin{aligned} \|\tilde{\mathbf{x}} - \mathbf{x}\|_{1, \mathbf{w}; \mathcal{V}} &\leq 2C_1\sigma_k(\mathbf{x})_{1, \mathbf{w}; \mathcal{V}} + \frac{C_1}{\lambda} (\lambda \|\tilde{\mathbf{x}}\|_{1, \mathbf{w}; \mathcal{V}} + \|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_{2; \mathcal{V}} - \lambda \|\mathbf{x}\|_{1, \mathbf{w}; \mathcal{V}}) \\ &\quad + C_2 \sqrt{k} \|\mathbf{e}\|_{2; \mathcal{V}}. \end{aligned}$$

The definition of \mathcal{G} in (7.2) gives

$$\|\tilde{\mathbf{x}} - \mathbf{x}\|_{1, \mathbf{w}; \mathcal{V}} \leq 2C_1\sigma_k(\mathbf{x})_{1, \mathbf{w}; \mathcal{V}} + \frac{C_1}{\lambda} (\mathcal{G}(\tilde{\mathbf{x}}) - \mathcal{G}(\mathbf{x}) + \|\mathbf{e}\|_{2; \mathcal{V}}) + C_2 \sqrt{k} \|\mathbf{e}\|_{2; \mathcal{V}},$$

which is the first result. The second follows in an analogous manner. \blacksquare

7.3 The weighted rNSP and weighted restricted isometry property

In the next chapter, we give explicit conditions in terms of m under which the measurement matrices (4.3) satisfy the weighted rNSP over \mathcal{V} . It is well known that showing the (weighted) rNSP directly can be difficult. In the classical, scalar setting, this is overcome by showing that the (weighted) rNSP is implied by the so-called (weighted) restricted isometry property. Hence, in this section, we first introduced this property and describe its relation to the (weighted) rNSP.

Let $\mathbf{w} > \mathbf{0}$ and $k > 0$. A bounded linear operator $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ has the *weighted Restricted Isometry Property (RIP)* over \mathcal{V} of order (k, \mathbf{w}) if there exists a constant $0 < \delta < 1$ such that

$$(1 - \delta) \|\mathbf{z}\|_{2; \mathcal{V}}^2 \leq \|\mathbf{A}\mathbf{z}\|_{2; \mathcal{V}}^2 \leq (1 + \delta) \|\mathbf{z}\|_{2; \mathcal{V}}^2, \quad \forall \mathbf{z} \in \Sigma_{k, \mathbf{w}} \subseteq \mathcal{V}^N.$$

The smallest constant such that this property holds is called the (k, \mathbf{w}) th *weighted Restricted Isometry Constant (wRIC)* of \mathbf{A} , and is denoted as $\delta_{k, \mathbf{w}}$.

It is first convenient to show an equivalence between the scalar weighted RIP over \mathbb{C} and the Hilbert-valued weighted RIP over \mathcal{V} .

Lemma 7.5 (Weighted RIP over \mathbb{C} is equivalent to the weighted RIP over \mathcal{V}). *Let $\mathbf{w} > \mathbf{0}$, $k > 0$ and $\mathbf{A} = (a_{ij})_{i,j=1}^{m,N} \in \mathbb{C}^{m \times N}$ be a matrix. Then \mathbf{A} satisfies the weighted RIP over \mathbb{C} of order (k, \mathbf{w}) with constant $0 < \delta < 1$ if and only if the corresponding bounded linear operator $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ defined by*

$$\mathbf{x} = (x_i)_{i=1}^N \in \mathcal{V}^N \mapsto \mathbf{A}\mathbf{x} := \left(\sum_{i=1}^N a_{ij} x_j \right)_{i=1}^m \in \mathcal{V}^m,$$

satisfies the weighted RIP over \mathcal{V} of order (k, \mathbf{w}) with the same constant δ .

Proof. We follow similar arguments to [52, Remark 3.5]. First, we rewrite the equivalence as follows:

$$(1 - \delta) \|\mathbf{x}\|_{2; \mathcal{V}}^2 \leq \|\mathbf{A}\mathbf{x}\|_{2; \mathcal{V}}^2 \leq (1 + \delta) \|\mathbf{x}\|_{2; \mathcal{V}}^2, \quad \forall \mathbf{x} \in \mathcal{V}^N, |\text{supp}(\mathbf{x})|_{\mathbf{w}} \leq k, \quad (7.4)$$

if and only if

$$(1 - \delta) \|\mathbf{x}\|_2^2 \leq \|\mathbf{A}\mathbf{x}\|_2^2 \leq (1 + \delta) \|\mathbf{x}\|_2^2, \quad \forall \mathbf{x} \in \mathbb{C}^N, |\text{supp}(\mathbf{x})|_{\mathbf{w}} \leq k. \quad (7.5)$$

Suppose that (7.5) holds. Let $\mathbf{x} = (x_j)_{j=1}^N \in \mathcal{V}^N$ be (k, \mathbf{w}) -sparse and $\{\phi_i\}_i$ be an orthonormal basis of \mathcal{V} . Then, for each $i \in [N]$, $x_i \in \mathcal{V}$ can be uniquely represented as

$$x_i = \sum_j \alpha_{ij} \phi_j, \quad \alpha_{ij} \in \mathbb{C}.$$

Let $\mathbf{x}_j = (\alpha_{ij})_{i=1}^N \in \mathbb{C}^N$. Then $\text{supp}(\mathbf{x}_j) \subseteq \text{supp}(\mathbf{x})$ and therefore \mathbf{x}_j is (k, \mathbf{w}) -sparse. Hence (7.5) gives

$$(1 - \delta) \|\mathbf{x}_j\|_2^2 \leq \|\mathbf{A}\mathbf{x}_j\|_2^2 \leq (1 + \delta) \|\mathbf{x}_j\|_2^2. \quad (7.6)$$

Now observe that

$$\sum_j \|\mathbf{x}_j\|_2^2 = \sum_{i=1}^N \sum_j |\alpha_{ij}|^2 = \sum_{i=1}^N \|x_i\|_{\mathcal{V}}^2 = \|\mathbf{x}\|_{2; \mathcal{V}}^2$$

and

$$\sum_j \|\mathbf{A}\mathbf{x}_j\|_2^2 = \sum_j \sum_{i=1}^m \left| \sum_{k=1}^N a_{ik} \alpha_{kj} \right|^2 = \sum_{i=1}^m \left\| \sum_{k=1}^N a_{ik} x_k \right\|_{\mathcal{V}}^2 = \|\mathbf{A}\mathbf{x}\|_{2; \mathcal{V}}^2.$$

Summing (7.6) over j , we deduce that (7.4) holds.

Conversely, suppose that (7.4) holds and let $\mathbf{z} = (z_i)_{i=1}^N \in \mathbb{C}^N$ with $|\text{supp}(\mathbf{z})|_{\mathbf{w}} \leq k$. Define $\mathbf{x} = (z_i \phi_i) \in \mathcal{V}^N$ and notice that $\|\mathbf{x}\|_{2;\mathcal{V}} = \|\mathbf{z}\|_2$ and $\|\mathbf{A}\mathbf{x}\|_{2;\mathcal{V}} = \|\mathbf{A}\mathbf{z}\|_2$. Since $\text{supp}(\mathbf{x}) = \text{supp}(\mathbf{z})$ and $|\text{supp}(\mathbf{z})|_{\mathbf{w}} \leq k$, we now apply (7.4) to deduce that $(1 - \delta)\|\mathbf{z}\|_2^2 \leq \|\mathbf{A}\mathbf{z}\|_2^2 \leq (1 + \delta)\|\mathbf{z}\|_2^2$. We conclude that (7.5) holds. ■

The following result shows that the weighted RIP is a sufficient condition for the weighted rNSP. This result is well known in the scalar-valued case (see, e.g., [8, Theorem 6.26]). Since its extension to the Hilbert-valued case is straightforward, we omit the proof.

Lemma 7.6 (Weighted RIP implies the weighted rNSP). *Let $\mathbf{w} > \mathbf{0}$, $k > 0$ and suppose that $\mathbf{A} \in \mathbb{C}^{m \times N}$ has the weighted RIP over \mathcal{V} of order $(2k, \mathbf{w})$ with constant $\delta_{2k, \mathbf{w}} < (2\sqrt{2} - 1)/7$. Then \mathbf{A} has the weighted rNSP of order (k, \mathbf{w}) over \mathcal{V} with constants $\rho = 2\sqrt{2}\delta_{2k, \mathbf{w}}/(1 - \delta_{2k, \mathbf{w}})$ and $\gamma = \sqrt{1 + \delta_{2k, \mathbf{w}}}/(1 - \delta_{2k, \mathbf{w}})$.*

Chapter 8

Error bounds for polynomial approximation via the Hilbert-valued, weighted SR-LASSO

Having developed the necessary tools for Hilbert-valued compressed sensing, we now specialize to the case introduced in Section 4.1 of polynomial approximation via the Hilbert-valued, weighted SR-LASSO problem (4.6). Our main results in this chapter, Theorems 8.2–8.4, yield error bounds for (inexact) minimizers of this problem in terms of the best polynomial approximation error, the Hilbert space discretization error and the noise.

8.1 The weighted RIP for the polynomial approximation problem

In order to obtain these results, we first need to assert conditions on m under which the relevant measurement matrix satisfies the weighted RIP. As in Section 4.1, we let $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathcal{F}} \subset L^2_{\varrho}(\mathcal{U})$ be either the tensor Chebyshev or Legendre polynomial basis,

$$\Lambda = \begin{cases} \Lambda_{n,d}^{\text{HC}} & d < \infty, \\ \Lambda_n^{\text{HCl}} & d = \infty, \end{cases} \quad (8.1)$$

be the hyperbolic cross index set and draw $\mathbf{y}_1, \dots, \mathbf{y}_m$ independently and identically from the measure ϱ . Then we define the measurement matrix \mathbf{A} exactly as in (4.3).

Lemma 8.1 (Weighted RIP for orthogonal polynomials). *Let $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0^d}$ be the orthonormal tensor Legendre or Chebyshev polynomial basis of $L^2_{\varrho}(\mathcal{U})$, Λ be as in (8.1) for some $n \geq 1$ and $\mathbf{y}_1, \dots, \mathbf{y}_m$ be drawn independently and identically from the measure ϱ . Let $0 < \epsilon < 1$, $k > 0$, \mathbf{u} be the intrinsic weights (4.7), $L' = L'(k, n, d, \epsilon)$ be given by*

$$L' = \begin{cases} \log(2k) \cdot (\log(2k) \cdot \min\{\log(n) + d, \log(ed) \cdot \log(2n)\} + \log(\epsilon^{-1})) & d < \infty, \\ \log(2k) \cdot (\log(2k) \cdot \log^2(2n) + \log(\epsilon^{-1})) & d = \infty, \end{cases}$$

and suppose that

$$m \geq c \cdot k \cdot L'(k, n, d, \epsilon), \quad (8.2)$$

where $c > 0$ is a universal constant. Then, with probability at least $1 - \epsilon$, the matrix \mathbf{A} defined in (4.3) satisfies the weighted RIP of order (k, \mathbf{u}) with constant $\delta_{k, \mathbf{u}} \leq 1/4$.

Proof. The proof uses ideas that are now standard. The matrix \mathbf{A} is a specific type of measurement matrix associated to the bounded orthonormal system $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \Lambda}$ (see,

e.g., [8, Section 6.4.3] or [61, Chapter 12]). Such a matrix satisfies the weighted RIP of order $k > 0$ with constant $\delta_{k,\mathbf{u}} \leq \delta$ whenever

$$m \geq c \cdot k \cdot \delta^{-2} \cdot \log\left(\frac{2k}{\delta^2}\right) \cdot \left[\frac{1}{\delta^4} \log\left(\frac{2k}{\delta^2}\right) \cdot \log(2N) + \frac{1}{\delta} \log(\epsilon^{-1}) \right], \quad (8.3)$$

where $c > 0$ is a universal constant. See, e.g., [8, Theorem 6.27 and equation (6.36)] (this result is based on [37]). To obtain the result, we set $\delta = 1/4$. Hence, (8.3) is implied by

$$m \geq c \cdot k \cdot \log(2k) \cdot [\log(2k) \cdot \log(2N) + \log(\epsilon^{-1})],$$

for a potentially different universal constant c . Next, we use (3.6) (and recall that $|\Lambda_n^{\text{HCl}}| = |\Lambda_{n,n}^{\text{HC}}|$) to estimate

$$\log(2N) \leq c \begin{cases} \min\{d + \log(n), \log(2d) \cdot \log(2n)\} & d < \infty, \\ \log^2(2n) & d = \infty, \end{cases}$$

for a potentially different universal constant. The result now follows after substituting this into the previous expression. \blacksquare

Note that the choice of $1/4$ in this lemma is arbitrary. Any value less than

$$(2\sqrt{2} - 1)/7 \approx 0.261$$

(see Lemma 7.6) will suffice.

8.2 Bounds for polynomial approximations obtained as inexact minimizers

We now present the main results of this chapter. These three results provide error bounds for polynomial approximations that are obtained as (inexact) minimizers to the weighted SR-LASSO problem (4.6). Each theorem corresponds to one of the three scenarios in our main results in Section 3.3. Hence, we label them accordingly as algebraic and finite-dimensional, algebraic and infinite-dimensional, and exponential. In order to state these results, we now define some additional notation. Given $f \in L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})$ and $\Lambda \subseteq \mathcal{F}$, where \mathcal{F} is as in (2.1)–(2.2), we let

$$E_{\Lambda,2}(f) = \|f - f_{\Lambda}\|_{L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})}, \quad E_{\Lambda,\infty}(f) = \|f - f_{\Lambda}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})},$$

where f_{Λ} is as in (4.1), and, given a subspace $\mathcal{V}_h \subseteq L^2_{\mathcal{Q}}(\mathcal{U}; \mathcal{V})$, we let

$$E_{h,\infty}(f) = \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathcal{U}; \mathcal{V})},$$

where $\mathcal{P}_h(f)$ is as in (2.5).

Theorem 8.2 (Error bounds for inexact minimizers, algebraic and finite-dimensional case). *Let $d \in \mathbb{N}$, $m \geq 3$, $0 < \epsilon < 1$, $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0^d} \subset L_{\varrho}^2(\mathcal{U})$ be either the orthonormal Chebyshev or Legendre basis, $\mathcal{V}_h \subseteq L_{\varrho}^2(\mathcal{U})$ be a subspace of $L_{\varrho}^2(\mathcal{U})$ and $\Lambda = \Lambda_{n,d}^{\text{HC}}$ be the hyperbolic cross index set with $n = \lceil m/L \rceil$ where $L = L(m, d, \epsilon)$ is as in (3.8). Let $f \in L_{\varrho}^2(\mathcal{U}; \mathcal{V})$, draw y_1, \dots, y_m randomly and independently according to ϱ and suppose that \mathbf{A} , \mathbf{b} and \mathbf{e} are as in (4.3) and (4.4). Consider the Hilbert-valued, weighted SR-LASSO problem (4.6) with weights $\mathbf{w} = \mathbf{u}$ as in (4.7) and $\lambda = (4\sqrt{m/L})^{-1}$. Then there exists universal constants $c_0, c_1, c_2 \geq 1$ such that the following holds with probability at least $1 - \epsilon$. Any $\tilde{\mathbf{c}} = (\tilde{c}_{\mathbf{v}})_{\mathbf{v} \in \Lambda} \in \mathbb{C}^N$ satisfies*

$$\|f - \tilde{f}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} \leq c_1 \cdot \xi, \quad \|f - \tilde{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \leq c_2 \cdot \sqrt{k} \cdot \xi, \quad \tilde{f} := \sum_{\mathbf{v} \in \Lambda} \tilde{c}_{\mathbf{v}} \Psi_{\mathbf{v}},$$

where

$$\xi = \frac{\sigma_k(\mathbf{c}_{\Lambda})_{1, \mathbf{u}; \mathcal{V}}}{\sqrt{k}} + \frac{E_{\Lambda, \infty}(f)}{\sqrt{k}} + E_{\Lambda, 2}(f) + E_{h, \infty}(f) + \mathcal{G}(\tilde{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_{\Lambda})) + \frac{\|\mathbf{n}\|_{2; \mathcal{V}}}{\sqrt{m}},$$

\mathbf{c}_{Λ} is as in (4.2), $\mathcal{P}_h(\mathbf{c}_{\Lambda}) = (\mathcal{P}_h(c_{\mathbf{v}}))_{\mathbf{v} \in \Lambda}$, $k = m/(c_0 L)$ for $L = L(m, d, \epsilon)$ as in (3.8), and \mathbf{n} is as in (4.4).

Proof. We divide the proof into several steps.

Step 1: Splitting the error into separate terms. Consider the $L_{\varrho}^2(\mathcal{U}; \mathcal{V})$ -norm error first. By the triangle inequality and the fact that \mathcal{P}_h is a projection, we have

$$\begin{aligned} & \|f - \tilde{f}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} \\ & \leq \|f - \mathcal{P}_h(f)\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} + \|\mathcal{P}_h(f) - \mathcal{P}_h(f_{\Lambda})\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} + \|\mathcal{P}_h(f_{\Lambda}) - \tilde{f}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} \\ & \leq \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} + \|f - f_{\Lambda}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} + \|\mathcal{P}_h(f_{\Lambda}) - \tilde{f}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} \\ & = E_{h, \infty}(f) + E_{\Lambda, 2}(f) + \|\mathcal{P}_h(f_{\Lambda}) - \tilde{f}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})}. \end{aligned}$$

Then, by orthonormality, we have

$$\|f - \tilde{f}\|_{L_{\varrho}^2(\mathcal{U}; \mathcal{V})} \leq E_{h, \infty}(f) + E_{\Lambda, 2}(f) + \|\mathcal{P}_h(\mathbf{c}_{\Lambda}) - \tilde{\mathbf{c}}\|_{2; \mathcal{V}}.$$

Similarly, for the $L^{\infty}(\mathcal{U}; \mathcal{V})$ -norm error, we have

$$\begin{aligned} & \|f - \tilde{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \\ & \leq \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} + \|\mathcal{P}_h(f) - \mathcal{P}_h(f_{\Lambda})\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} + \|\mathcal{P}_h(f_{\Lambda}) - \tilde{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \\ & \leq \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} + \|f - f_{\Lambda}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} + \|\mathcal{P}_h(f_{\Lambda}) - \tilde{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \\ & = E_{h, \infty}(f) + E_{\Lambda, \infty}(f) + \|\mathcal{P}_h(f_{\Lambda}) - \tilde{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})}. \end{aligned}$$

Using the definition (4.7) of the weights \mathbf{u} , we deduce that

$$\|f - \tilde{f}\|_{L^\infty(\mathcal{U}; \mathcal{V})} \leq E_{h, \infty}(f) + E_{\Lambda, \infty}(f) + \|\mathcal{P}_h(\mathbf{c}_\Lambda) - \tilde{\mathbf{c}}\|_{1, \mathbf{u}; \mathcal{V}}.$$

Therefore, the rest of the proof is devoted to showing the following bounds:

$$\|\mathcal{P}_h(\mathbf{c}_\Lambda) - \tilde{\mathbf{c}}\|_{2; \mathcal{V}} \leq c_1 \cdot \xi, \quad \|\mathcal{P}_h(\mathbf{c}_\Lambda) - \tilde{\mathbf{c}}\|_{1, \mathbf{u}; \mathcal{V}} \leq c_2 \cdot \sqrt{k} \cdot \xi. \quad (8.4)$$

We do this in the next two steps by first asserting that \mathbf{A} has the weighted rNSP (Step 2) and then by applying the error bounds of Lemma 7.4 (Steps 3 and 4).

Step 2: Asserting the weighted rNSP. We now show that \mathbf{A} has the weighted rNSP over \mathcal{V}_h of order (k, \mathbf{u}) with probability at least $1 - \epsilon/2$. This is based on Lemma 8.1. First observe that

$$L = L(m, d, \epsilon) \geq \log^2(3) \cdot \min\{\log(3) + 1, \log(3) \cdot \log(e)\} \geq 1,$$

since $m \geq 3$. This implies that $m \geq m/L \geq m/(c_0 L) = k$ since $c_0 \geq 1$ as well. Since $n = \lceil m/L \rceil \leq m/L + 1 \leq 2m$, we get

$$\begin{aligned} & \log(4k) \cdot (\log(4k) \cdot \min\{\log(n) + d, \log(ed) \cdot \log(2n)\} + \log(2/\epsilon)) \\ & \leq \log(4m) \cdot (\log(4m) \cdot \min\{\log(2m) + d, \log(ed) \cdot \log(4m)\} + \log(2/\epsilon)) \\ & \leq c_0 L(m, d, \epsilon)/2 \end{aligned}$$

for a suitably large choice of c_0 . Hence

$$m = c_0 k L(m, d, \epsilon) \geq 2c_0 k L'(2k, d, \epsilon/2),$$

where L' is defined as in Lemma 8.1, and therefore (again assuming a suitably large choice of c_0) (8.2) holds with k replaced by $2k$. We deduce that \mathbf{A} satisfies the weighted RIP over \mathbb{C} of order $(2k, \mathbf{u})$ with constant $\delta_{2k, \mathbf{u}} \leq 1/4$, with probability at least $1 - \epsilon/2$. Then, we deduce from Lemmas 7.5 and 7.6 that \mathbf{A} has (with the same probability) the weighted rNSP over \mathcal{V}_h of order (k, \mathbf{u}) with constants $\rho = 2\sqrt{2}/3$ and $\gamma = 2\sqrt{5}/3$.

Step 3: Bounding $\mathcal{P}_h(\mathbf{c}_\Lambda) - \tilde{\mathbf{c}}$ using the weighted rNSP. We use Lemma 7.4. First, consider the value of λ . Since $c_0 \geq 1$ we have $m/L \geq m/(c_0 L) = k$. Hence, recalling the values for ρ and γ obtained in the previous step, we have

$$\frac{1}{4\sqrt{c_0}} \frac{1}{\sqrt{k}} = \frac{1}{4\sqrt{m/L}} = \lambda \leq \frac{1}{4\sqrt{k}} < \frac{(1 + \rho)^2}{(3 + \rho)\gamma} \frac{1}{\sqrt{k}}. \quad (8.5)$$

Therefore, (7.3) holds. We now apply this lemma with $\mathcal{V} = \mathcal{V}_h$, $\mathbf{x} = \mathcal{P}_h(\mathbf{c}_\Lambda)$, $\tilde{\mathbf{x}} = \tilde{\mathbf{c}}$ and $\mathbf{e} = \mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}$. Notice first that the best (k, \mathbf{u}) -approximation error (7.1)

satisfies

$$\begin{aligned} \sigma_k(\mathcal{P}_h(\mathbf{c}_\Lambda))_{1,\mathbf{u};\mathfrak{V}} &= \inf \left\{ \sum_{\mathbf{v} \in \Lambda \setminus S} u_{\mathbf{v}} \|\mathcal{P}_h(\mathbf{c}_{\mathbf{v}})\|_{\mathfrak{V}} : S \subseteq \Lambda, |S|_{\mathbf{u}} \leq k \right\} \\ &\leq \sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{u};\mathfrak{V}}, \end{aligned} \quad (8.6)$$

since \mathcal{P}_h is a projection. Hence, applying Lemma 7.4 and using the lower bound in (8.5), we get

$$\begin{aligned} \|\tilde{\mathbf{c}} - \mathcal{P}_h(\mathbf{c}_\Lambda)\|_{2;\mathfrak{V}} &\leq c_1 \left[\frac{\sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{w};\mathfrak{V}}}{\sqrt{k}} + \mathcal{G}(\tilde{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) + \|\mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{2;\mathfrak{V}} \right], \\ \|\tilde{\mathbf{c}} - \mathcal{P}_h(\mathbf{c}_\Lambda)\|_{1,\mathbf{u};\mathfrak{V}} &\leq c_2 \left[\sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{w};\mathfrak{V}} + \sqrt{k}(\mathcal{G}(\tilde{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda))) \right. \\ &\quad \left. + \sqrt{k} \|\mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{2;\mathfrak{V}} \right], \end{aligned} \quad (8.7)$$

with probability at least $1 - \epsilon/2$. Therefore, to show (8.4) and therefore complete the proof, it suffices to show that the following holds with probability at least $1 - \epsilon/2$:

$$\|\mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{2;\mathfrak{V}} \leq \sqrt{2} \left(\frac{E_{\Lambda,\infty}(f)}{\sqrt{k}} + E_{\Lambda,2}(f) \right) + E_{h,\infty}(f) + \frac{\|\mathbf{n}\|_{2;\mathfrak{V}}}{\sqrt{m}}. \quad (8.8)$$

The overall result then follows by the union bound.

Step 4: Showing that (8.8) holds. Observe that

$$\begin{aligned} &\sqrt{m} \|\mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_i \\ &\leq \|\mathcal{P}_h(f_\Lambda)(\mathbf{y}_i) - f(\mathbf{y}_i)\|_{\mathfrak{V}} + \|\mathbf{n}_i\|_{\mathfrak{V}} \\ &\leq \|\mathcal{P}_h(f_\Lambda)(\mathbf{y}_i) - \mathcal{P}_h(f)(\mathbf{y}_i)\|_{\mathfrak{V}} + \|f(\mathbf{y}_i) - \mathcal{P}_h(f)(\mathbf{y}_i)\|_{\mathfrak{V}} + \|\mathbf{n}_i\|_{\mathfrak{V}} \\ &\leq \|f(\mathbf{y}_i) - f_\Lambda(\mathbf{y}_i)\|_{\mathfrak{V}} + E_{h,\infty}(f) + \|\mathbf{n}_i\|_{\mathfrak{V}}. \end{aligned}$$

Therefore,

$$\|\mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{\mathfrak{V};2} \leq E_{\Lambda,\text{disc}}(f) + E_{h,\infty}(f) + \frac{\|\mathbf{n}\|_{2;\mathfrak{V}}}{\sqrt{m}}, \quad (8.9)$$

where

$$E_{\Lambda,\text{disc}}(f) = \sqrt{\frac{1}{m} \sum_{i=1}^m \|f(\mathbf{y}_i) - f_\Lambda(\mathbf{y}_i)\|_{\mathfrak{V}}^2}. \quad (8.10)$$

For this final step, we follow near-identical arguments to those found in [8, Lemma 7.11]. This shows that

$$E_{\Lambda,\text{disc}}(f) \leq \sqrt{2} \left(\frac{E_{\Lambda,\infty}(f)}{\sqrt{k}} + E_{\Lambda,2}(f) \right),$$

with probability at least $1 - \epsilon/2$, provided $m \geq 2k \log(2/\epsilon)$. However, this follows due to the assumptions on m and the arguments given in Step 2. Thus we obtain (8.8) and the proof is complete. \blacksquare

Theorem 8.3 (Error bounds for inexact minimizers, algebraic and infinite-dimensional case). *Let $d = \infty$, $m \geq 3$, $0 < \epsilon < 1$, $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathcal{F}} \subset L^2_{\varrho}(\mathcal{U})$ be either the orthonormal Chebyshev or Legendre basis, $\mathcal{V}_h \subseteq L^2_{\varrho}(\mathcal{U})$ be a subspace of $L^2_{\varrho}(\mathcal{U})$ and $\Lambda = \Lambda_n^{\text{HC}}$ be the hyperbolic cross index set with $n = \lceil m/L \rceil$ where $L = L(m, d, \epsilon)$ is as in (3.8). Let $f \in L^2_{\varrho}(\mathcal{U}; \mathcal{V})$, draw $\mathbf{y}_1, \dots, \mathbf{y}_m$ randomly and independently according to ϱ and suppose that \mathbf{A} , \mathbf{b} and \mathbf{e} are as in (4.3) and (4.4). Consider the Hilbert-valued, weighted SR-LASSO problem (4.6) with weights $\mathbf{w} = \mathbf{u}$ as in (4.7) and $\lambda = (4\sqrt{m/L})^{-1}$. Then there exists universal constants $c_0, c_1, c_2 \geq 1$ such that the following holds with probability at least $1 - \epsilon$. Any $\tilde{\mathbf{c}} = (\tilde{c}_{\mathbf{v}})_{\mathbf{v} \in \Lambda} \in \mathbb{C}^N$ satisfies*

$$\|f - \tilde{f}\|_{L^2_{\varrho}(\mathcal{U}; \mathcal{V})} \leq c_1 \cdot \xi, \quad \|f - \tilde{f}\|_{L^{\infty}(\mathcal{U}; \mathcal{V})} \leq c_2 \cdot \sqrt{k} \cdot \xi, \quad \tilde{f} := \sum_{\mathbf{v} \in \Lambda} \tilde{c}_{\mathbf{v}} \Psi_{\mathbf{v}},$$

where

$$\xi = \frac{\sigma_k(\mathbf{c}_{\Lambda})_{1, \mathbf{u}; \mathcal{V}}}{\sqrt{k}} + \frac{E_{\Lambda, \infty}(f)}{\sqrt{k}} + E_{\Lambda, 2}(f) + E_{h, \infty}(f) + \mathcal{G}(\tilde{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_{\Lambda})) + \frac{\|\mathbf{n}\|_{2; \mathcal{V}}}{\sqrt{m}},$$

\mathbf{c}_{Λ} is as in (4.2), $\mathcal{P}_h(\mathbf{c}_{\Lambda}) = (\mathcal{P}_h(\mathbf{c}_{\mathbf{v}}))_{\mathbf{v} \in \Lambda}$, $k = m/(c_0 L)$ for $L = L(m, d, \epsilon)$ as in (3.8), and \mathbf{n} is as in (4.4).

Proof. The proof has the same structure as that of the previous theorem. Steps 1, 3 and 4 are identical. The only differences occur in Step 2. We now describe these changes. Once more we observe that $L = L(m, \infty, \epsilon) \geq 1$ since $m \geq 3$. Hence, $m \geq m/L \geq m/(c_0 L) = k$ since $c_0 \geq 1$. We also have $n = \lceil m/L \rceil \leq 2m$. Therefore,

$$\begin{aligned} \log(4k) \cdot (\log(4k) \cdot \log^2(2n) + \log(2/\epsilon)) &\leq \log(4m) \cdot (\log^3(4m) + \log(2/\epsilon)) \\ &\leq c_0 L(m, \infty, \epsilon)/2 \end{aligned}$$

for a suitably large choice of c_0 . We deduce that

$$m = c_0 k L(m, \infty, \epsilon) \geq 2c_0 k L'(2k, \infty, \epsilon/2),$$

where L' is as in Lemma 8.1. An application of this lemma now shows that \mathbf{A} has the weighted RIP of order $(2k, \mathbf{u})$ with constant $\delta_{2k, \mathbf{u}} \leq 1/4$, as required. \blacksquare

Theorem 8.4 (Error bounds for inexact minimizers, exponential case). *Let $d \in \mathbb{N}$, $m \geq 3$, $0 < \epsilon < 1$, $\{\Psi_{\mathbf{v}}\}_{\mathbf{v} \in \mathbb{N}_0^d} \subset L^2_{\varrho}(\mathcal{U})$ be either the orthonormal Chebyshev or Legendre basis, $\mathcal{V}_h \subseteq L^2_{\varrho}(\mathcal{U})$ be a subspace of $L^2_{\varrho}(\mathcal{U})$ and $\Lambda = \Lambda_{n, d}^{\text{HC}}$ be the hyperbolic cross index set with n as in (3.15). Draw $\mathbf{y}_1, \dots, \mathbf{y}_m$ randomly and independently according to ϱ . Then, with probability at least $1 - \epsilon$, the following holds. Let $f \in L^2_{\varrho}(\mathcal{U}; \mathcal{V})$ and suppose that \mathbf{A} , \mathbf{b} and \mathbf{e} are as in (4.3) and (4.4). Consider the Hilbert-valued, weighted SR-LASSO problem (4.6) with weights $\mathbf{w} = \mathbf{u}$ as in (4.7) and $\lambda = (4\sqrt{m/L})^{-1}$. Then there exists universal constants $c_0, c_1, c_2 \geq 1$ such that*

any $\tilde{\mathbf{c}} = (\tilde{c}_{\mathbf{v}})_{\mathbf{v} \in \Lambda} \in \mathbb{C}^N$ satisfies

$$\|f - \tilde{f}\|_{L^2_{\theta}(\mathcal{U}; \mathcal{V})} \leq c_1 \cdot \xi, \quad \|f - \tilde{f}\|_{L^\infty(\mathcal{U}; \mathcal{V})} \leq c_2 \cdot \sqrt{k} \cdot \xi, \quad \tilde{f} := \sum_{\mathbf{v} \in \Lambda} \tilde{c}_{\mathbf{v}} \Psi_{\mathbf{v}},$$

where

$$\xi = \frac{\sigma_k(\mathbf{c}_\Lambda)_{1, \mathbf{u}; \mathcal{V}}}{\sqrt{k}} + E_{\Lambda, \infty}(f) + E_{h, \infty}(f) + \mathcal{G}(\tilde{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) + \frac{\|\mathbf{n}\|_{2; \mathcal{V}}}{\sqrt{m}},$$

\mathbf{c}_Λ is as in (4.2), $\mathcal{P}_h(\mathbf{c}_\Lambda) = (\mathcal{P}_h(c_{\mathbf{v}}))_{\mathbf{v} \in \Lambda}$, $k = m/(c_0 L)$ for $L = L(m, d, \epsilon)$ as in (3.8), and \mathbf{n} is as in (4.4).

Proof. The proof has the same structure as that of Theorem 8.2. Step 1 is identical, and reduces the proof to showing that (8.4) holds. We now describe the modifications needed in Steps 2–4.

Step 2: Asserting the weighted rNSP. We now show that \mathbf{A} has the weighted rNSP over \mathcal{V}_h of order (k, \mathbf{u}) with probability at least $1 - \epsilon$. This step is essentially the same, except for the choice of n and the probability $1 - \epsilon$ instead of $1 - \epsilon/2$.

Step 3: Bounding $\mathcal{P}_h(\mathbf{c}_\Lambda) - \tilde{\mathbf{c}}$ using the weighted rNSP. Since λ and k are the same as in Theorem 8.2, the bound (8.5) also holds in this case. We then follow the same arguments, leading to (8.7) holding with probability at least $1 - \epsilon$. Finally, rather than (8.8), we ask for the slightly modified bound

$$\|\mathbf{A} \mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{2; \mathcal{V}} \leq E_{\Lambda, \infty}(f) + E_{h, \infty}(f) + \frac{\|\mathbf{n}\|_{2; \mathcal{V}}}{\sqrt{m}}, \quad (8.11)$$

to hold with probability one.

Step 4: Showing (8.11) holds. By the same argument, we see that (8.9) holds. Instead of the probabilistic bound for $E_{\Lambda, \text{disc}}(f)$, we now simply bound it as

$$E_{\Lambda, \text{disc}}(f) \leq \|f - f_\Lambda\|_{L^\infty(\mathcal{U}; \mathcal{V})} = E_{\Lambda, \infty}(f).$$

This immediately implies (8.11).

Finally, we observe that we can simplify the previous estimates in this case using the bound $E_{\Lambda, 2}(f) \leq E_{\Lambda, \infty}(f)$. \blacksquare

Chapter 9

Error bounds and the restarting scheme for the primal-dual iteration

Theorems 8.2–8.4 reduce the problem of proving the main results (Theorems 3.4–3.12) to two tasks. The first involves bounding the error in the objective function, i.e., the term

$$\mathcal{G}(\tilde{c}) - \mathcal{G}(\mathcal{P}_h(c_\Lambda)),$$

where \tilde{c} is either an exact minimizer or an approximate minimizer obtained via the primal dual iteration. The second involves the various approximation error terms depending on f and its polynomial coefficients.

In this chapter, we address the first task. We first provide an error bound for the (unrestarted) primal-dual iteration when applied to Hilbert-valued weighted SR-LASSO problem (7.2), and then use this to derive the specific restart scheme.

9.1 Error bounds for the primal-dual iteration

We now return to the general setting of the primal-dual iteration, where it is applied to the problem (4.9) and takes the form (4.12). The following result from [31, Theorem 5.1] establishes an important error bound for the Lagrangian difference.

Theorem 9.1. *Let $\tau, \sigma > 0$, initial points $(x^{(0)}, \xi^{(0)}) \in \mathcal{X} \times \mathcal{Y}$ and a bounded linear operator $A \in \mathcal{B}(\mathcal{X}, \mathcal{Y})$, be such that $\|A\|_{\mathcal{B}(\mathcal{X}, \mathcal{Y})}^2 \leq (\tau\sigma)^{-1}$. Consider the sequence $\{(x^{(n)}, \xi^{(n)})\}_{n=1}^\infty$ generated by the primal-dual iteration (4.12). Then, for any $(x, \xi) \in \mathcal{X} \times \mathcal{Y}$,*

$$\mathcal{L}(\bar{x}^{(n)}, \xi) - \mathcal{L}(x, \bar{\xi}^{(n)}) \leq \frac{\tau^{-1} \|x - x^{(0)}\|_{2; \mathcal{Y}}^2 + \sigma^{-1} \|\xi - \xi^{(0)}\|_{2; \mathcal{Y}}^2}{n}, \quad (9.1)$$

where

$$\bar{x}^{(n)} = \frac{1}{n} \sum_{k=1}^n x^{(k)} \quad \text{and} \quad \bar{\xi}^{(n)} = \frac{1}{n} \sum_{k=1}^n \xi^{(k)},$$

are the ergodic sequences and \mathcal{L} is the Lagrangian (4.11).

The following lemma shows a decay rate of $1/n$ on the objective function in the case of the primal-dual iteration when applied to the problem (7.2). It is an extension of [13, Lemma 8.6] to the weighted and Hilbert-valued setting.

Lemma 9.2. *Let $A \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ and $\tau, \sigma > 0$ be such that $\|A\|_{\mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)}^2 \leq (\tau\sigma)^{-1}$. Consider the sequence $\{(\mathbf{x}^{(n)}, \boldsymbol{\xi}^{(n)})\}_{n=1}^{\infty}$ generated by the primal-dual iteration in (4.12) applied to (7.2) with $\mathbf{x}^{(0)} \in \mathcal{V}^N$ and $\boldsymbol{\xi}^{(0)} = \mathbf{0} \in \mathcal{V}^m$. Then, for any $\mathbf{x} \in \mathcal{V}^N$,*

$$\mathcal{G}(\bar{\mathbf{x}}^{(n)}) - \mathcal{G}(\mathbf{x}) \leq \frac{\tau^{-1} \|\mathbf{x} - \mathbf{x}_0\|_{2;\mathcal{V}}^2 + \sigma^{-1}}{n}, \quad \bar{\mathbf{x}}^{(n)} = \frac{1}{n} \sum_{k=1}^n \mathbf{x}^{(k)}. \quad (9.2)$$

Proof. Using (4.11) and (4.13), the left-hand side of (9.1) is given by

$$\begin{aligned} \mathcal{T}_n(\mathbf{x}, \boldsymbol{\xi}) &:= (\lambda \|\bar{\mathbf{x}}^{(n)}\|_{1,\mathbf{w};\mathcal{V}} + \operatorname{Re}\langle A\bar{\mathbf{x}}^{(n)} - \mathbf{b}, \boldsymbol{\xi} \rangle_{2;\mathcal{V}} + \delta_B(\boldsymbol{\xi})) \\ &\quad - (\lambda \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}} + \operatorname{Re}\langle A\mathbf{x} - \mathbf{b}, \bar{\boldsymbol{\xi}}^{(n)} \rangle_{2;\mathcal{V}} + \delta_B(\bar{\boldsymbol{\xi}}^{(n)})), \end{aligned}$$

where B is the unit ball in \mathcal{V}^m . Observe that the term $\boldsymbol{\xi}^{(n)}$ produced by this iteration satisfies $\|\boldsymbol{\xi}^{(n)}\|_{2;\mathcal{V}} \leq 1$. This follows from the observation shown in Section 4.4 that the proximal mapping

$$\operatorname{prox}_{\sigma h^*}(\boldsymbol{\xi}) = \operatorname{proj}_B(\boldsymbol{\xi} - \sigma \mathbf{b})$$

involves the projection onto the unit ball B . Hence, the ergodic sequence $\bar{\boldsymbol{\xi}}^{(n)}$ satisfies $\|\bar{\boldsymbol{\xi}}^{(n)}\|_{2;\mathcal{V}} \leq 1$ as well. Suppose now that $A\mathbf{x}^{(n)} - \mathbf{b} \neq \mathbf{0}$ and set

$$\boldsymbol{\xi} = \frac{A\mathbf{x}^{(n)} - \mathbf{b}}{\|A\mathbf{x}^{(n)} - \mathbf{b}\|_{2;\mathcal{V}}}.$$

Then $\delta_B(\boldsymbol{\xi}) = \delta_B(\bar{\boldsymbol{\xi}}^{(n)}) = 1$ and therefore

$$\begin{aligned} \mathcal{T}_n(\mathbf{x}, \boldsymbol{\xi}) &= (\lambda \|\bar{\mathbf{x}}^{(n)}\|_{1,\mathbf{w};\mathcal{V}} + \|A\bar{\mathbf{x}}^{(n)} - \mathbf{b}\|_{2;\mathcal{V}}) - (\lambda \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}} + \operatorname{Re}\langle A\mathbf{x} - \mathbf{b}, \bar{\boldsymbol{\xi}}^{(n)} \rangle_{2;\mathcal{V}}) \\ &\geq (\lambda \|\bar{\mathbf{x}}^{(n)}\|_{1,\mathbf{w};\mathcal{V}} + \|A\bar{\mathbf{x}}^{(n)} - \mathbf{b}\|_{2;\mathcal{V}}) - (\lambda \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}} + \|A\mathbf{x} - \mathbf{b}\|_{2;\mathcal{V}}). \end{aligned}$$

Clearly, the same bound also holds in the case $A\mathbf{x}^{(n)} - \mathbf{b} = \mathbf{0}$ where $\boldsymbol{\xi}$ is an arbitrary unit vector. Hence Theorem 9.1 and the fact that

$$\|\boldsymbol{\xi} - \boldsymbol{\xi}_0\|_{2;\mathcal{V}} = \|\boldsymbol{\xi}\|_{2;\mathcal{V}} = 1$$

gives the result. ■

9.2 The restarting scheme

For convenience, we now introduce new and slightly modify some existing notation. First, we redefine the objective function \mathcal{G} of the Hilbert-valued weighted SR-LASSO problem (7.2) to make the dependence on the term \mathbf{b} explicit: namely, we set

$$\mathcal{G}(\mathbf{x}, \mathbf{b}) = \lambda \|\mathbf{x}\|_{1,\mathbf{w};\mathcal{V}} + \|A\mathbf{x} - \mathbf{b}\|_{2;\mathcal{V}}, \quad \mathbf{x} \in \mathcal{V}^N, \quad \mathbf{b} \in \mathcal{V}^m.$$

We then let

$$\mathcal{E}(z, \mathbf{x}, \mathbf{b}) = \mathcal{G}(z, \mathbf{b}) - \mathcal{G}(\mathbf{x}, \mathbf{b}), \quad \mathbf{x}, z \in \mathcal{V}^N, \mathbf{b} \in \mathcal{V}^m.$$

Now consider the ergodic sequence $\bar{\mathbf{x}}^{(n)}$ produced by n iterations of the primal-dual iteration (4.12) applied to (7.2) with parameters $\tau, \sigma > 0$, $\mathbf{x}_0 \in \mathcal{V}^N$ and $\boldsymbol{\xi}_0 = \mathbf{0} \in \mathcal{V}^m$. For reasons that will become clear in a moment, we now make the dependence on the vector \mathbf{b} in (7.2), the number of iterations $\bar{\mathbf{x}}^{(n)}$ and the initial vector \mathbf{x}_0 explicit, by defining

$$\mathcal{P}(\mathbf{x}_0, \mathbf{b}, n) = \bar{\mathbf{x}}^{(n)}.$$

With this in hand, we conclude this discussion by noting the following two scaling properties:

$$\mathcal{G}(a\mathbf{x}, \mathbf{b}) = a\mathcal{G}(\mathbf{x}, \mathbf{b}/a), \quad \mathcal{E}(a\mathbf{z}, \mathbf{x}, \mathbf{b}) = a\mathcal{E}(z, \mathbf{x}/a, \mathbf{b}/a). \quad (9.3)$$

These hold for any $a > 0$ and for any $\mathbf{x}, z \in \mathcal{V}^N$ and $\mathbf{b} \in \mathcal{V}^m$.

Lemma 9.3. *Suppose that $A \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ has the weighted rNSP over \mathcal{V} of order (k, \mathbf{w}) with constants $0 < \rho < 1$ and $\gamma > 0$. Consider the Hilbert-valued weighted SR-LASSO problem (7.2) with parameter $\lambda = c/\sqrt{k}$, where*

$$0 < c \leq \frac{(1 + \rho)^2}{(3 + \rho)\gamma}.$$

Let \mathcal{E} and \mathcal{P} be as defined above, τ, σ satisfy $\|A\|_{\mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)}^2 \leq (\tau\sigma)^{-1}$ and $\mathbf{x}, \mathbf{x}_0 \in \mathcal{V}^N$, $\mathbf{b} \in \mathcal{V}^m$, $a > 0$. Then

$$\mathcal{E}(a\mathcal{P}(\mathbf{x}_0/a, \mathbf{b}/a, n), \mathbf{x}, \mathbf{b}) \leq \frac{C^2}{a\tau n} (\mathcal{E}(\mathbf{x}_0, \mathbf{x}, \mathbf{b}) + \xi)^2 + \frac{a}{\sigma n},$$

where

$$C = 2 \max\{C'_1/c, C'_2\}, \quad (9.4)$$

C'_1, C'_2 are as in Lemma 7.3 and

$$\xi = \xi(\mathbf{x}, \mathbf{b}) = \frac{\sigma_k(\mathbf{x})_{1, \mathbf{w}; \mathcal{V}}}{\sqrt{k}} + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_{2; \mathcal{V}}. \quad (9.5)$$

Proof. The scaling property (9.3) and Lemma 9.2 give

$$\begin{aligned} \mathcal{E}(a\mathcal{P}(\mathbf{x}_0/a, \mathbf{b}/a, n), \mathbf{x}, \mathbf{b}) &= a\mathcal{E}(\mathcal{P}(\mathbf{x}_0/a, \mathbf{b}/a, n), \mathbf{x}/a, \mathbf{b}/a) \\ &\leq a \left(\frac{\tau^{-1} \|\mathbf{x}/a - \mathbf{x}_0/a\|_{2; \mathcal{V}}^2 + \sigma^{-1}}{n} \right) \\ &= \frac{\|\mathbf{x} - \mathbf{x}_0\|_{2; \mathcal{V}}^2}{a\tau n} + \frac{a}{\sigma n}. \end{aligned}$$

Now consider the term $\|\mathbf{x} - \mathbf{x}_0\|_{2;\mathcal{V}}$. Since \mathbf{A} has the weighted rNSP and λ satisfies (7.3), we may use Lemma 7.4 to get

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}_0\|_{2;\mathcal{V}} &\leq \frac{C'_1}{\sqrt{k}} \left(2\sigma_k(\mathbf{x})_{1,\mathbf{w};\mathcal{V}} + \frac{\mathcal{G}(\mathbf{x}_0, \mathbf{b}) - \mathcal{G}(\mathbf{x}, \mathbf{b})}{\lambda} \right) + \left(\frac{C'_1}{\sqrt{k}\lambda} + C'_2 \right) \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_{2;\mathcal{V}} \\ &= \frac{C'_1}{\sqrt{k}\lambda} \mathcal{E}(\mathbf{x}_0, \mathbf{x}, \mathbf{b}) + 2C'_1 \frac{\sigma_k(\mathbf{x})_{1,\mathbf{w};\mathcal{V}}}{\sqrt{k}} + \left(\frac{C'_1}{\sqrt{k}\lambda} + C'_2 \right) \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_{2;\mathcal{V}} \\ &\leq 2 \max\{C'_1/c, C'_2\} (\mathcal{E}(\mathbf{x}_0, \mathbf{x}, \mathbf{b}) + \xi). \end{aligned}$$

Substituting this into the previous expression now gives the result. \blacksquare

This lemma gives the rationale behind the restarted scheme. It says the error in the objective function of the scaled output $a\mathcal{P}(\mathbf{x}_0/a, \mathbf{b}/a, n)$ of the primal-dual iteration with initial value \mathbf{x}_0 can be bounded in terms of the error in the objective function at the initial value, plus terms depending on the scaling parameter a , the number of iterations n and the compressed sensing error term ξ . By choosing these parameters suitably and iterating this procedure, we obtain the restarting scheme. We summarize this in the following theorem.

Theorem 9.4 (Restarting scheme). *Suppose that $\mathbf{A} \in \mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)$ has the weighted rNSP over \mathcal{V} of order (k, \mathbf{w}) with constants $0 < \rho < 1$ and $\gamma > 0$. Consider the Hilbert-valued weighted SR-LASSO problem (7.2) with parameter $\lambda = c/\sqrt{k}$, where $0 < c \leq \frac{(1+\rho)^2}{(3+\rho)\gamma}$. Let $\mathbf{x} \in \mathcal{V}^N$, $\mathbf{b} \in \mathcal{V}^m$, $\zeta' \geq \xi$, where ξ is as in (9.5), $0 < r < 1$ and define the sequence*

$$\varepsilon_0 = \|\mathbf{b}\|_{2;\mathcal{V}}, \quad \varepsilon_{k+1} = r(\varepsilon_k + \zeta'), \quad k = 0, 1, 2, \dots$$

Let \mathcal{E} and \mathcal{P} be as defined above, τ, σ satisfy $\|\mathbf{A}\|_{\mathcal{B}(\mathcal{V}^N, \mathcal{V}^m)}^2 \leq (\tau\sigma)^{-1}$ and set

$$n = \left\lceil \frac{2C}{r\sqrt{\sigma\tau}} \right\rceil, \quad a_k = \frac{1}{2}\sigma\varepsilon_{k+1}n, \quad k = 0, 1, 2, \dots,$$

where C is as in (9.4). Then the iteration $\tilde{\mathbf{x}}^{(0)}, \tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \dots$, defined by

$$\tilde{\mathbf{x}}^{(0)} = \mathbf{0}, \quad \tilde{\mathbf{x}}^{(k+1)} = a_k \mathcal{P}(\tilde{\mathbf{x}}^{(k)}/a_k, \mathbf{b}/a_k, n), \quad k = 0, 1, 2, \dots,$$

satisfies

$$\mathcal{E}(\mathbf{x}_k^*, \mathbf{x}, \mathbf{b}) \leq \varepsilon_k \leq r^k \|\mathbf{b}\|_{2;\mathcal{V}} + \frac{r}{1-r} \zeta', \quad k = 0, 1, 2, \dots$$

Proof. We use induction on k . Suppose first that $k = 0$. Then, by definition,

$$\mathcal{E}(\tilde{\mathbf{x}}^{(k)}, \mathbf{x}, \mathbf{b}) = \mathcal{E}(\mathbf{0}, \mathbf{x}, \mathbf{b}) \leq \mathcal{G}(\mathbf{0}, \mathbf{b}) = \|\mathbf{b}\|_{2;\mathcal{V}} = \varepsilon_0.$$

Now suppose that the result holds for k . The previous lemma gives

$$\begin{aligned}\mathcal{E}(\tilde{\mathbf{x}}^{(k+1)}, \mathbf{x}, \mathbf{b}) &= \mathcal{E}(a_k \mathcal{P}(\tilde{\mathbf{x}}^{(k)}/a_k, \mathbf{b}/a_k, n), \mathbf{x}, \mathbf{b}) \\ &\leq \frac{C^2}{a_k \tau n} (\mathcal{E}(\tilde{\mathbf{x}}^{(k)}, \mathbf{x}, \mathbf{b}) + \zeta)^2 + \frac{a_k}{\sigma n} \\ &\leq \frac{C^2}{a_k \tau n} (\varepsilon_k + \zeta)^2 + \frac{a_k}{\sigma n}.\end{aligned}$$

We now substitute the values of n and a_k to obtain

$$\mathcal{E}(\tilde{\mathbf{x}}^{(k+1)}, \mathbf{x}, \mathbf{b}) = \frac{2C^2(\varepsilon_k + \zeta)}{r\sigma\tau n^2} + \frac{1}{2}r(\varepsilon_k + \zeta) \leq \frac{1}{2}r(\varepsilon_k + \zeta) + \frac{1}{2}r(\varepsilon_k + \zeta) = \varepsilon_{k+1}.$$

This completes the proof. ■

This theorem states that the restarted primal-dual iteration $\tilde{\mathbf{x}}^{(0)}, \tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \dots$ yields an objective function error $\mathcal{E}(\tilde{\mathbf{x}}^{(k)}, \mathbf{x}, \mathbf{b})$ that converges exponentially fast in the number of restarts k . Further, each (inner) primal-dual iteration involves a number of steps n that depends on the parameters C, r, σ and τ . In other words, n is a constant independent of k . Hence, the restarted scheme converges exponentially fast in the total number of primal-dual iterations as well.

As discussed in Section 5.1.1, it is typical to use this theorem to optimize the choice of r . Recall that this leads to the explicit choice $r = e^{-1}$. We use this value in our algorithms – see Table 4.3.

Chapter 10

Final arguments

We are now ready to prove the main results, Theorems 3.4–3.12. In several of these proofs, we require the following definition. Let $s \in \mathbb{N}$ and set

$$k(s) = \max\{|S|_{\mathbf{u}} : S \subset \mathbb{N}_0^d, |S| \leq s, S \text{ lower}\}, \quad (10.1)$$

where \mathbf{u} are the intrinsic weights (4.7) (recall the definition of a lower set from Definition 2.8). It can be shown that

$$k(s) = s^2, \quad (\text{Legendre}), \quad k(s) \leq \min\{2^d s, s^{\log(3)/\log(2)}\}, \quad (\text{Chebyshev}). \quad (10.2)$$

See, e.g., [8, equation (7.42), Propositions 5.13 and 5.17]. We will use this property several times in what follows.

10.1 Algebraic rates of convergence, finite dimensions

Proof of Theorem 3.4. The mapping was described in Table 4.1. As shown therein, we can write the corresponding approximation as $\hat{f} = \sum_{\mathbf{v} \in \Lambda} \hat{c}_{\mathbf{v}} \Psi_{\mathbf{v}}$, where $\hat{c} = (\hat{c}_{\mathbf{v}})_{\mathbf{v} \in \Lambda}$ is a minimizer of (4.6). Next, due to the various assumptions made, we may apply Theorem 8.2. Setting $\tilde{f} = \hat{f}$ and $\tilde{c} = \hat{c}$, we deduce that

$$\|f - \hat{f}\|_{L^2_{\rho}(\mathbf{u}; \mathbf{v})} \leq c_1 \cdot \xi, \quad \|f - \hat{f}\|_{L^{\infty}(\mathbf{u}; \mathbf{v})} \leq c_2 \cdot \sqrt{k} \cdot \xi, \quad (10.3)$$

where (after writing out the term $E_{h, \infty}(f)$ explicitly)

$$\begin{aligned} \xi &= \frac{\sigma_k(\mathbf{c}_{\Lambda})_{1, \mathbf{u}; \mathbf{v}}}{\sqrt{k}} + \frac{E_{\Lambda, \infty}(f)}{\sqrt{k}} + E_{\Lambda, 2}(f) + \|f - \mathcal{P}_h(f)\|_{L^{\infty}(\mathbf{u}; \mathbf{v})} \\ &\quad + \mathcal{G}(\hat{c}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_{\Lambda})) + \frac{\|\mathbf{n}\|_{2; \mathbf{v}}}{\sqrt{m}}, \end{aligned} \quad (10.4)$$

and $k = m/(c_0 L)$ with $c_0 \geq 1$ a universal constant. We now bound each term separately.

Step 1. The terms $\sigma_k(\mathbf{c}_{\Lambda})_{1, \mathbf{u}; \mathbf{v}}/\sqrt{k}$, $E_{\Lambda, \infty}(f)/\sqrt{k}$ and $E_{\Lambda, 2}(f)$. The term

$$\sigma_k(\mathbf{c}_{\Lambda})_{1, \mathbf{u}; \mathbf{v}}/\sqrt{k}$$

is estimated via (ii) of Theorem A.1 with $q = 1$. This gives

$$\frac{\sigma_k(\mathbf{c}_{\Lambda})_{1, \mathbf{u}; \mathbf{v}}}{\sqrt{k}} \leq C(d, p, \rho) \cdot k^{1/2-1/p} = C(d, p, \rho) \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}. \quad (10.5)$$

We estimate the term $E_{\Lambda,2}(f)$ by first recalling that $\Lambda = \Lambda_{n,d}^{\text{HC}}$ is the union of all lower sets (see Definition 2.8) of size at most $n = \lceil m/L \rceil$ (see Section 3.2). Hence, using (i) of Theorem A.1 with $s = n$ and $q = 2$, we get

$$\begin{aligned} E_{\Lambda,2}(f) &= \|\mathbf{c} - \mathbf{c}_{\Lambda}\|_{2;\mathcal{V}} \leq \|\mathbf{c} - \mathbf{c}_S\|_{2;\mathcal{V}} \leq C(d, p, \rho) \cdot n^{1/2-1/p} \\ &\leq C(d, p, \rho) \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}. \end{aligned} \quad (10.6)$$

Here, in the last step we recall that $n = \lceil m/L \rceil$ and $c_0 \geq 1$.

It remains to consider $E_{\Lambda,\infty}(f)/\sqrt{k}$. Due to the choice of weights, we have $E_{\Lambda,\infty}(f) \leq \|\mathbf{c} - \mathbf{c}_{\Lambda}\|_{1,\mathbf{u};\mathcal{V}}$. We now apply (i) of Theorem A.1 once more, with $s = n$ and $q = 1$, to get

$$E_{\Lambda,\infty}(f) \leq \|\mathbf{c} - \mathbf{c}_S\|_{1,\mathbf{u};\mathcal{V}} \leq C(d, p, \rho) \cdot n^{1-1/p}.$$

Since $n = \lceil m/L \rceil \geq m/(c_0 L) = k$, we obtain

$$\frac{E_{\Lambda,\infty}(f)}{\sqrt{k}} \leq C(d, p, \rho) \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}. \quad (10.7)$$

Step 2. The term $\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_{\Lambda}))$. Since $\hat{\mathbf{c}}$ is a minimizer of (4.6) and $\mathcal{P}_h(\mathbf{c}_{\Lambda}) \in \mathcal{V}_h^N$ is feasible for (4.6), this term satisfies

$$\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_{\Lambda})) \leq 0. \quad (10.8)$$

Step 3. Conclusion. We now substitute the bounds (10.5)–(10.8) into (10.4). Since $k \leq m/L$, we deduce that $\xi \leq \zeta$, where ζ is given by (3.10). This completes the proof. \blacksquare

Proof of Theorem 3.5. The argument is similar to that of the previous theorem. Recall from Section 4.5 that, in this case the approximation $\hat{f} = \sum_{\mathbf{v} \in \Lambda} \tilde{c}_{\mathbf{v}} \Psi_{\mathbf{v}}$, where $\hat{\mathbf{c}} = \tilde{\mathbf{c}}^{(T)}$ is the ergodic sequence obtained after T steps of the primal-dual iteration applied to (4.6). Hence, the only difference is the estimation of $\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_{\Lambda}))$ in Step 2.

We now do this using Lemma 9.2. In order to apply this lemma we first need to estimate $\|\mathbf{A}\|_{\mathcal{B}(\mathcal{V}_h^N, \mathcal{V}_h^m)}$. Let $\mathbf{x} = (x_{\mathbf{v}})_{\mathbf{v} \in \Lambda} \in \mathcal{V}_h^N$ and define $p(\mathbf{y}) = \sum_{\mathbf{v} \in \Lambda} x_{\mathbf{v}} \Psi_{\mathbf{v}}$. Then

$$\|\mathbf{A}\mathbf{x}\|_{2;\mathcal{V}} = \sqrt{\frac{1}{m} \sum_{i=1}^m \|p(\mathbf{y}_i)\|_{\mathcal{V}}^2} \leq \sup_{\mathbf{y} \in \mathcal{U}} \|p(\mathbf{y})\|_{\mathcal{V}} \leq \sum_{\mathbf{v} \in \Lambda} \|x_{\mathbf{v}}\|_{\mathcal{V}} u_{\mathbf{v}} \leq \|\mathbf{x}\|_{2;\mathcal{V}} \sqrt{|\Lambda|_{\mathbf{u}}}.$$

Now the set Λ is lower and of cardinality $|\Lambda| = \Theta(n, d)$. Hence, by (10.2) with $s = N$, we have $|\Lambda|_{\mathbf{u}} \leq (\Theta(n, d))^{2\alpha}$, where α is as in (3.7). Since \mathbf{x} was arbitrary, we get

$$\|\mathbf{A}\|_{2;\mathcal{V}} \leq (\Theta(n, d))^{\alpha}. \quad (10.9)$$

Since the primal-dual iteration in Section 4.5 is used with $\tau = \sigma = (\Theta(n, d))^{-\alpha}$, we have that

$$\|\mathbf{A}\|_{2;\mathbf{v}}^2 \leq (\tau\sigma)^{-1}.$$

Hence, we may apply Lemma 9.2. Since the iteration is also initialized with the zero vector and run for a total of $T = \lceil 2(\Theta(n, d))^{\alpha} t \rceil$ iterations (see Section 4.5 once more), this gives

$$\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) \leq (\Theta(n, d))^\alpha \frac{\|\mathcal{P}_h(\mathbf{c}_\Lambda)\|_{2;\mathbf{v}}^2 + 1}{T}.$$

Observe that

$$\|\mathcal{P}_h(\mathbf{c}_\Lambda)\|_{2;\mathbf{v}} \leq \|\mathbf{c}_\Lambda\|_{2;\mathbf{v}} \leq \|\mathbf{c}\|_{c;\mathbf{v}} = \|f\|_{L^2_\rho(\mathbf{u};\mathbf{v})} \leq 1.$$

Here, in the last step, we use the fact that $f \in \mathcal{B}(\rho)$, and therefore

$$\|f\|_{L^2_\rho(\mathbf{u};\mathbf{v})} \leq \|f\|_{L^\infty(\mathbf{u};\mathbf{v})} \leq 1.$$

Using this and the value of T , we deduce that

$$\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) \leq \frac{1}{t}.$$

Substituting this into (10.4) and combining with the other estimates (10.5)–(10.7) derived in Step 2 of the proof of Theorem 3.4 now gives the desired error bound.

It remains to estimate the computational cost. We do this via Lemmas 4.3 and 4.4. First observe that the value k in Lemma 4.4 is equal to $k = d$ in this case, since the index set $\Lambda = \Lambda_{n,d}^{\text{HC}}$ is a d -dimensional hyperbolic cross index set. Similarly, the value n in Lemma 4.4 is bounded by the order n of this hyperbolic cross. As Λ is a lower set, we also have $n \leq N$. Hence, the computational cost for forming the matrix \mathbf{A} is bounded by $c \cdot m \cdot N \cdot d$. We now use Lemma 4.3 to bound the computational cost of the algorithm. Finally, we recall that $N = \Theta(n, d)$ and $T = \lceil 2(\Theta(n, d))^{\alpha} t \rceil$ in this case. ■

Proof of Theorem 3.6. As in the previous proof, we only need to estimate the term $\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda))$. Recall from Table 4.3 that in this case $\hat{\mathbf{c}} = \tilde{\mathbf{c}}^{(R)}$ is the output of the restarted primal-dual iteration with R restarts. Our goal is to use Theorem 9.4 applied to the problem (4.6) with weights $\mathbf{w} = \mathbf{u}$ as in (4.7), $\lambda = (4\sqrt{m/L})^{-1}$ and $\mathbf{x} = \mathcal{P}_h(\mathbf{c}_\Lambda)$.

We first show that the conditions of this theorem hold. Recall from Step 2 of the proof of Theorem 8.2 that the matrix \mathbf{A} has the weighted rNSP of order (k, \mathbf{u}) over \mathcal{V}_h with constants $\rho = 2\sqrt{2}/3$ and $\gamma = 2\sqrt{5}/3$. In particular,

$$\frac{(1 + \rho)^2}{(3 + \rho)\gamma} \geq 0.64.$$

We now use (8.5) to see that

$$\lambda = \frac{1}{4\sqrt{c_0}} \frac{1}{\sqrt{k}} \leq \frac{(1+\rho)^2}{(3+\rho)\gamma} \frac{1}{\sqrt{k}},$$

for a sufficiently large choice of c_0 .

Next, with this choice of \mathbf{x} , we see that

$$\xi(\mathbf{x}, \mathbf{b}) = \frac{\sigma_k(\mathcal{P}_h(\mathbf{c}_\Lambda))_{1,\mathbf{u};\mathcal{V}}}{\sqrt{k}} + \|\mathbf{A} \mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{2;\mathcal{V}}.$$

Using (8.6) and (8.8), we get

$$\xi(\mathbf{x}, \mathbf{b}) \leq \frac{\sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{w};\mathcal{V}}}{\sqrt{k}} + \sqrt{2} \left(\frac{E_{\Lambda,\infty}(f)}{\sqrt{k}} + E_{\Lambda,2}(f) \right) + E_{h,\infty}(f) + \frac{\|\mathbf{n}\|_{2;\mathcal{V}}}{\sqrt{m}},$$

with probability at least $1 - \epsilon$. Using (10.5)–(10.7), we deduce that

$$\xi(\mathbf{x}, \mathbf{b}) \leq \zeta,$$

with probability at least $1 - \epsilon$, where ζ is as in (3.10). Hence, $\xi(\mathbf{x}, \mathbf{b}) \leq \zeta'$.

Next, recall from Table 4.3 that $\tau = \sigma = (\Theta(n, d))^{-\alpha}$ in this case. Due to (10.9), we see that $\|\mathbf{A}\|_{2;\mathcal{V}} \leq (\tau\sigma)^{-1}$ as well.

Now consider the constant C defined in (9.4). The values for ρ and γ give that $C'_1 \leq C'_2 \leq 103$. Since $\lambda = c/\sqrt{k}$ with $c = 1/(4\sqrt{c_0})$, we see that

$$4C \leq 812/c = 3296\sqrt{c_0} := c^*. \quad (10.10)$$

Therefore, recalling that $r = 1/2$ and $\tau = \sigma = (\Theta(n, d))^{-\alpha}$, we see that

$$\left\lceil \frac{2C}{r\sqrt{\sigma\tau}} \right\rceil = \lceil (\Theta(n, d))^\alpha c^* \rceil = T,$$

where T is as specified in Table 4.3, and

$$\frac{1}{2} r \sigma (\varepsilon_k + \zeta') T = \frac{(\Theta(n, d))^\alpha T}{4} \varepsilon_{k+1} = s \varepsilon_{k+1} = a_k,$$

where s and a_k are as specified in Table 4.3 and Algorithm 4, respectively.

With this in hand, we are now finally in a position to apply Theorem 9.4. We deduce that

$$\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) = \mathcal{E}(\tilde{\mathbf{c}}^{(R)}, \mathcal{P}_h(\mathbf{c}_\Lambda), \mathbf{b}) \leq \varepsilon_k = e^{-R} \|\mathbf{b}\|_{2;\mathcal{V}} + \zeta'.$$

To complete the proof of the error bound (3.12), we simply note that $\|\mathbf{b}\|_{2;\mathcal{V}} \leq \|f\|_{L^\infty(\mathcal{U};\mathcal{V})} \leq 1$, since $f \in \mathcal{B}(\rho)$.

It remains to estimate the computational cost. As before, the computational cost for forming the matrix \mathbf{A} is bounded by $c \cdot m \cdot N \cdot d$. Next, by construction, we

observe that the algorithm consists of $R = t$ primal-dual iterations, each involving $T = \lceil (\Theta(n, d))^{\alpha} c^* \rceil$ steps. Therefore, by Lemma 4.3 the computational cost for the algorithm is

$$c \cdot (m \cdot N \cdot K + (m + N) \cdot (F(\mathbf{G}) + K)) \cdot \lceil (\Theta(n, d))^{\alpha} c^* \rceil \cdot t.$$

Since $N = \Theta(n, d)$ and c^* is a universal constant, the result follows. \blacksquare

10.2 Algebraic rates of convergence, infinite dimensions

Proof of Theorem 3.7. The proof is similar to that of Theorem 3.4, except that it uses Theorem 8.3 in place of Theorem 8.2. In particular, we see that (10.3) also holds in this case with ξ as in (10.4) and $k = m/(c_0 L)$.

Step 2 is identical. The only differences occur in Step 1. We now describe the changes needed in this case. First consider the term $\sigma_k(\mathbf{c}_\Lambda)_{1, \mathbf{u}; \mathbf{v}} / \sqrt{k}$. To bound this, we use (ii) of Theorem A.3 with $q = 1 > p$. This gives

$$\frac{\sigma_k(\mathbf{c}_\Lambda)_{1, \mathbf{u}; \mathbf{v}}}{\sqrt{k}} \leq C(\mathbf{b}, \varepsilon, p) \cdot k^{1/2-1/p} = C(\mathbf{b}, \varepsilon, p) \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}.$$

To estimate $E_{\Lambda, 2}(f)$, recall that $\Lambda = \Lambda_n^{\text{HCl}}$ contains all anchored sets (see Definition 2.8) of size at most $n = \lceil m/L \rceil$ (see Section 3.2). Hence, using (iii) of Theorem A.3 with $s = n$ and $q = 2 > p$, we get

$$\begin{aligned} E_{\Lambda, 2}(f) &= \|\mathbf{c} - \mathbf{c}_\Lambda\|_{2; \mathbf{v}} \leq \|\mathbf{c} - \mathbf{c}_S\|_{2; \mathbf{v}} \leq C(\mathbf{b}, \varepsilon, p) \cdot n^{1/2-1/p} \\ &\leq C(\mathbf{b}, \varepsilon, p) \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}. \end{aligned}$$

Finally, for $E_{\Lambda, \infty}(f)$, we use (iii) of Theorem A.3 once more (with $q = 1 > p$) to get

$$\begin{aligned} \frac{E_{\Lambda, \infty}(f)}{\sqrt{k}} &\leq \frac{\|\mathbf{c} - \mathbf{c}_S\|_{1, \mathbf{u}; \mathbf{v}}}{\sqrt{k}} \leq C(\mathbf{b}, \varepsilon, p) \cdot k^{1/2-1/p} \\ &= C(\mathbf{b}, \varepsilon, p) \cdot \left(\frac{m}{c_0 L}\right)^{1/2-1/p}. \end{aligned}$$

Having done this, we also observe that $\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) \leq 0$ in this case, since $\hat{\mathbf{c}}$ is once more an exact minimizer. Using this and the previously derived bounds, we conclude that $\xi \leq \zeta$, where ζ is as in (3.14). This gives the result. \blacksquare

Proof of Theorem 3.8. The argument is similar to that of Theorem 3.5. Here $\hat{\mathbf{c}} = \bar{\mathbf{c}}^{(T)}$ is the ergodic sequence obtained after T steps of the primal-dual iteration applied to (4.6) as well.

We recall that the set Λ is lower and of cardinality $|\Lambda| = \Theta(n, d)$ with $d = \infty$. Hence, by (10.2) with $s = N$, we have $|\Lambda|_{\mathbf{u}} \leq (\Theta(n, d))^{2\alpha}$, where α is as in (3.7). Using this, we get

$$\|\mathbf{A}\|_{2;\mathbf{v}} \leq (\Theta(n, d))^\alpha,$$

as before. Since the primal-dual iteration in Table 4.3 is used with

$$\tau = \sigma = (\Theta(n, d))^{-\alpha},$$

we have that $\|\mathbf{A}\|_{2;\mathbf{v}}^2 \leq (\tau\sigma)^{-1}$. Hence, following the same steps we deduce that

$$\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) \leq \frac{1}{t}.$$

Substituting this into (10.4) and combining with the other estimates (10.5)–(10.7) derived in Step 2 of the proof of Theorem 3.4 now gives the desired error bound.

The computational cost estimate is similar to the that in the proof of Theorem 3.5. In this case, observe that the value k in Lemma 4.4 is equal to n . Hence, the computational cost of forming \mathbf{A} is bounded by $c \cdot m \cdot N \cdot n$ in this case. The computational cost for the algorithm is given by Lemma 4.3. To complete the estimate, we substitute the values $N = \Theta(n, d)$ and $T = \lceil 2(\Theta(n, d))^{\alpha t} \rceil$, as before. ■

Proof of Theorem 3.9. The proof is similar to that of Theorem 3.6 and involves estimating the term $\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda))$. Using the same steps, we deduce that

$$\xi(\mathbf{x}, \mathbf{b}) \leq \zeta,$$

with probability at least $1 - \epsilon/2$, where ζ is as in (3.14). Hence, $\xi(\mathbf{x}, \mathbf{b}) \leq \zeta'$.

Next, recall from Table 4.3 that $\tau = \sigma = (\Theta(n, d))^{-\alpha}$ with $d = \infty$ in this case. Due to (10.9), we see that $\|\mathbf{A}\|_{2;\mathbf{v}} \leq (\tau\sigma)^{-1}$ holds. We now apply Theorem 9.4 to obtain

$$\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) = \mathcal{E}(\hat{\mathbf{c}}^{(R)}, \mathcal{P}_h(\mathbf{c}_\Lambda), \mathbf{b}) \leq \epsilon_R = e^{-R} \|\mathbf{b}\|_{2;\mathbf{v}} + \zeta'.$$

To complete the proof of the error bound (3.12), we simply note that $\|\mathbf{b}\|_{2;\mathbf{v}} \leq \|f\|_{L^\infty(\mathbf{u};\mathbf{v})} \leq 1$, since $f \in \mathcal{B}(\mathbf{b}, \epsilon)$.

The computational cost estimate is as in the previous proof. ■

10.3 Exponential rates of convergence, finite dimensions

Proof of Theorem 3.10. The proof has the same structure to that of Theorem 3.4, the only differences being the use of Theorem 8.4 instead of Theorem 8.2 and the estimation of the various terms in Step 1. Suppose first that $m \geq c_0 2^{d+2} L$ and define the

following:

$$s = \begin{cases} \lceil \sqrt{m/(4c_0L)} \rceil & \text{Legendre,} \\ \lceil m/(4c_02^dL) \rceil & \text{Chebyshev.} \end{cases}$$

Observe that

$$s \leq \begin{cases} \sqrt{m/(c_0L)} & \text{Legendre,} \\ m/(c_02^dL) & \text{Chebyshev,} \end{cases}$$

and therefore the quantity $k(s)$ defined in (10.1) satisfies

$$k(s) \leq \frac{m}{c_0L} = k.$$

Now consider the term $\sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{u};\mathbf{v}}/\sqrt{k}$. Using this and (iii) of Theorem A.1 with $p = 1$ we have

$$\frac{\sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{u};\mathbf{v}}}{\sqrt{k}} \leq \frac{\sigma_{k(s)}(\mathbf{c})_{1,\mathbf{u};\mathbf{v}}}{\sqrt{k}} \leq \frac{C(d, \gamma, \boldsymbol{\rho}) \cdot \exp(-\gamma s^{1/d})}{\sqrt{k}}.$$

Note that this is possible since any lower set S of size at most s satisfies $|S|_{\mathbf{u}} \leq k(s)$ by definition.

Now consider $E_{\Lambda,\infty}(f)$. Recall that $\Lambda = \Lambda_{n,d}^{\text{HC}}$, where n is as in (3.15). Clearly $n \geq s$, since $c_0 \geq 1$. Hence Λ contains all lower sets of size at most s . We deduce that

$$E_{\Lambda,\infty}(f) \leq \|\mathbf{c} - \mathbf{c}_S\|_{1,\mathbf{u};\mathbf{v}},$$

for any lower set of size s . We now use (iii) of Theorem A.1 with $p = 1$ once more, to get

$$E_{\Lambda,\infty}(f) \leq C(d, \gamma, \boldsymbol{\rho}) \cdot \exp(-\gamma s^{1/d}).$$

We now combine this with the previous bound to deduce that the quantity ξ in Theorem 8.4 satisfies

$$\xi \leq C(d, \gamma, \boldsymbol{\rho}) \cdot \exp(-\gamma s^{1/d}) + E_{h,\infty}(f) + \frac{\|\mathbf{n}\|_{2;\mathbf{v}}}{\sqrt{m}},$$

(here, we also recall that the term $\mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)) \leq 0$, as in the proof of Theorem 3.4). Using the value of s and recalling that $m \geq c_02^{d+2}L$, we deduce that

$$\begin{aligned} \xi \leq C(d, \gamma, \boldsymbol{\rho}) \cdot & \begin{cases} \exp\left(-\frac{\gamma}{2}\left(\frac{m}{4c_0L}\right)^{\frac{1}{d}}\right) & \text{Chebyshev} \\ \exp\left(-\gamma\left(\frac{m}{4c_0L}\right)^{\frac{1}{2d}}\right) & \text{Legendre} \end{cases} + \frac{\|\mathbf{n}\|_{2;\mathbf{v}}}{\sqrt{m}} \\ & + \|f - \mathcal{P}_h(f)\|_{L^\infty(\mathbf{u};\mathbf{v})}, \end{aligned}$$

for $m \geq c_02^{d+2}L$. However, this bound also clearly holds for all $m \geq 1$, up to a change in the constant $C(d, \gamma, \boldsymbol{\rho})$. After relabeling the universal constant $4c_0$ as c_0 , we deduce that $\xi \leq \zeta$, where ζ is as in (3.17). This concludes the proof. \blacksquare

Proof of Theorem 3.11. The argument is the same as the proof of Theorem 3.5. The difference relies on the fact that now ζ has the following bound

$$\xi \leq C(d, \gamma, \rho) \cdot \begin{cases} \exp\left(-\frac{\gamma}{2}\left(\frac{m}{4c_0L}\right)^{\frac{1}{d}}\right) & \text{Chebyshev} \\ \exp\left(-\gamma\left(\frac{m}{4c_0L}\right)^{\frac{1}{2d}}\right) & \text{Legendre} \end{cases} + \frac{\|\mathbf{n}\|_{2;\mathcal{V}}}{\sqrt{m}} \\ + \|f - \mathcal{P}_h(f)\|_{L^\infty(\mathcal{U};\mathcal{V})} + \mathcal{G}(\hat{\mathbf{c}}) - \mathcal{G}(\mathcal{P}_h(\mathbf{c}_\Lambda)).$$

To estimate the final term, we argue exactly as in the proof of Theorem 3.5. The computational cost estimate is likewise identical. ■

Proof of Theorem 3.12. The proof is similar to that of Theorem 3.6, except we use Theorem 8.4 instead. Recall from Step 2 of the proof of Theorem 8.4 that the matrix A has the weighted rNSP of order (k, \mathbf{u}) over \mathcal{V}_h with constants $\rho = 2\sqrt{2}/3$ and $\gamma = 2\sqrt{5}/3$ with probability $1 - \epsilon$. In particular,

$$\frac{(1 + \rho)^2}{(3 + \rho)\gamma} \geq 0.64.$$

We now use (8.5) to see that

$$\lambda = \frac{1}{4\sqrt{c_0}} \frac{1}{\sqrt{k}} \leq \frac{(1 + \rho)^2}{(3 + \rho)\gamma} \frac{1}{\sqrt{k}},$$

for a sufficiently large choice of c_0 , as before.

Next, with the choice $\mathbf{x} = \mathcal{P}_h(\mathbf{c}_\Lambda)$ as before, we see that

$$\xi(\mathbf{x}, \mathbf{b}) = \frac{\sigma_k(\mathcal{P}_h(\mathbf{c}_\Lambda))_{1,\mathbf{u};\mathcal{V}}}{\sqrt{k}} + \|\mathbf{A}\mathcal{P}_h(\mathbf{c}_\Lambda) - \mathbf{b}\|_{2;\mathcal{V}}.$$

Using (8.11), we get

$$\xi(\mathbf{x}, \mathbf{b}) \leq \frac{\sigma_k(\mathbf{c}_\Lambda)_{1,\mathbf{w};\mathcal{V}}}{\sqrt{k}} + E_{\Lambda,\infty}(f) + E_{h,\infty}(f) + \frac{\|\mathbf{n}\|_{2;\mathcal{V}}}{\sqrt{m}},$$

with probability $1 - \epsilon$. It now follows from the proof of Theorem 3.10 that

$$\xi(\mathbf{x}, \mathbf{b}) \leq \zeta,$$

with probability at least $1 - \epsilon$, where ζ is as in (3.17). Hence, $\xi(\mathbf{x}, \mathbf{b}) \leq \zeta'$.

The rest of the proof follows the same steps as the proof of Theorem 3.6. ■

Chapter 11

Conclusions

Sparse polynomial approximation is a useful tool in parametric model problems, including surrogate model construction in UQ. The theory of best s -term approximation supports the use of polynomial-based methods, and techniques such as least squares and compressed sensing are known to have desirable sample complexity bounds for obtaining polynomial approximations. In this work, we have closed a key gap between these two areas of research, by showing the existence of algorithms that achieve the algebraic and exponential rates of the best s -term approximation with respect to the number of samples m . Thus, sparse polynomial approximation can be practically realized in a provably sample-efficient manner. As our numerical experiments confirm, our algorithms are practical, and actually perform better than the theory suggests.

There are a number of avenues for further research. First, this work has focused on Chebyshev and Legendre polynomials on the hypercube $[-1, 1]^d$. It is plausible that it can be extended to general ultraspherical or Jacobi polynomials. A more significant challenge involves Hermite or Laguerre polynomials on \mathbb{R}^d or $[0, \infty)^d$, respectively. This is an interesting problem for future research.

It is notable that the algorithms developed in this work do not generally compute m -term polynomial approximations. Indeed, (inexact) minimizers of the SR-LASSO problem will generally be nonsparse vectors of length $N = \Theta(n, d)$. It is interesting to investigate whether one can develop algorithms that achieve the same error bounds while computing m -term polynomial approximations. In classical compressed sensing, one can typically compute sparse solutions by using a greedy or iterative procedure (see, e.g., [61]). Unfortunately, it is not clear how to extend these procedures to the weighted case with theoretical guarantees. Nonetheless, certain weighted greedy methods appear to work well in practice for sparse polynomial approximation [4].

Another motivation for considering different algorithms is to see if the computational cost estimates can be reduced. While this is often not the main computational bottleneck in parametric model problems (generally, computing the samples is the most computationally intensive step), it is still an important issue. We have shown that the computational cost is at worst subexponential in m in infinite dimensions, and algebraic in m (for fixed d) in finite dimensions. Whether these are optimal is an interesting open problem. Here, ideas from sublinear-time algorithms [38, 39] may be particularly useful.

In the case of the exponential rates, it is notable that the best s -term approximation error is exponentially small in $\gamma \cdot s^{1/d}$ (see Theorem 2.6), whereas the exponents in

Section 3.3.3 are $(\gamma/2) \cdot (m/(c_0L))^{1/d}$ (Chebyshev) and $\gamma \cdot (m/(c_0L))^{1/(2d)}$ (Legendre). The reason for this can be traced to the sample complexity estimate for computing a sparse (and lower) polynomial approximations via compressed sensing with Monte Carlo sampling, i.e., $m \approx c_0 \cdot 2^d \cdot s \cdot L$ (Chebyshev) or $m \approx c_0 \cdot s^2 \cdot L$ (Legendre). To see why this is the case, combine Lemma 8.1 with (10.2). In the setting of least squares, in which the desired polynomial subspace is known, it is possible to change the sampling measure to obtain sample complexity bounds that are log-linear in s and therefore near optimal. See, e.g., [9, 44, 71]. More recently, several works [21, 54, 56, 85, 93, 137] have also introduced sampling schemes that achieve linear sample complexity in s – i.e., optimal up to a constant. Unfortunately, it is unknown whether linear or log-linear sample complexity possible in the compressed sensing setting, where the target subspace is unknown. See [10] for further discussion on this issue.

Finally, as previously noted in Section 1.5, this work focuses on polynomial approximation, and not on fundamental issues pertaining to tractability and the information complexity of the classes of multivariate holomorphic functions considered. For some related work in this direction, see [81, 112, 147] and references therein. A question of particular interest is whether pointwise samples (i.e., *standard information*), and more specifically, i.i.d. pointwise samples (i.e., *random information*) constitutes optimal or near-optimal information for these classes of functions. These questions have recently been considered in a broader context in [79, 88]. See also [87] for the case of functions in Sobolev spaces. As we observed in Section 1.5, in a recent work [12] we derived lower bounds for the (adaptive) m -widths for classes of $(\mathbf{b}, \varepsilon)$ -holomorphic functions in infinite dimensions. Showing that the algorithms (or small modifications thereof) developed in this work also attain (nearly) matching upper bounds – and, consequently, that i.i.d. pointwise samples constitute (near) optimal information – is an interesting problem for future work.

Appendix A

Best polynomial approximation rates for holomorphic functions

In this appendix, we recap a series of standard best approximation error bounds for polynomial approximation of holomorphic functions. These are used in Chapter 10 to estimate the various error terms appearing in Theorems 8.2–8.4.

A.1 The finite-dimensional case

We first consider the finite-dimensional case, where $\mathcal{U} = [-1, 1]^d$ for $d < \infty$ and $f : \mathcal{U} \rightarrow \mathcal{V}$ is a Hilbert-valued function (in fact, the following results also apply in the more general setting of Banach-valued functions; however, we shall not consider this explicitly). We now summarize the various approximation error bounds in the following theorem. This result combines various well-known results in the literature. It is essentially the same as [8, Theorem 3.25]. However, we have made a number of minor edits to fit the notation and setup of this work (see Remark A.2 below).

Theorem A.1 (Best s -term decay rates; finite dimensions). *Let $d \in \mathbb{N}$, $f \in \mathcal{B}(\boldsymbol{\rho})$ for some $\boldsymbol{\rho} > \mathbf{1}$, where $\mathcal{B}(\boldsymbol{\rho})$ is as in (2.6), and $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \mathbb{N}_0^d}$ be its Chebyshev or Legendre coefficients. Then the following best s -term decay rates hold:*

- (i) *for any $0 < p \leq q \leq 2$ and $s \in \mathbb{N}$, there exists a lower set $S \subset \mathbb{N}_0^d$ of size $|S| \leq s$ such that*

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

where $\sigma_s(\mathbf{c})_{q;\mathcal{V}}$ is as in Definition 2.3 (with $\Lambda = \mathbb{N}_0^d$), \mathbf{u} is as in (4.7) and $C = C(d, p, \boldsymbol{\rho}) > 0$ depends on d , p and $\boldsymbol{\rho}$ only;

- (ii) *for any $0 < p \leq q \leq 2$ and $k > 0$,*

$$\sigma_k(\mathbf{c})_{q,\mathbf{u};\mathcal{V}} \leq C \cdot k^{1/q-1/p},$$

where $\sigma_k(\mathbf{c})_{q,\mathbf{u};\mathcal{V}}$ is as in Definition 7.2, \mathbf{u} is as in (4.7) and

$$C = C(d, p, \boldsymbol{\rho}) > 0$$

depends on d , p and $\boldsymbol{\rho}$ only;

- (iii) *for any $0 < p \leq 2$,*

$$0 < \gamma < (d + 1)^{-1} \left(d! \prod_{j=1}^d \log(\rho_j) \right)^{1/d},$$

and $s \in \mathbb{N}$, there exists a lower set $S \subset \mathbb{N}_0^d$ of size $|S| \leq s$ such that

$$\sigma_s(\mathbf{c})_{p;\mathcal{V}} \leq \|\mathbf{c} - \mathbf{c}_S\|_{p;\mathcal{V}} \leq \|\mathbf{c} - \mathbf{c}_S\|_{p,\mathbf{u};\mathcal{V}} \leq C \cdot \exp(-\gamma s^{1/d}),$$

where $\sigma_s(\mathbf{c})_{p;\mathcal{V}}$ is as in Definition 2.3 (with $\Lambda = \mathbb{N}_0^d$), \mathbf{u} is as in (4.7) and $C = C(d, \gamma, p, \rho) > 0$ depends on d, γ, p and ρ only.

Remark A.2. There are several differences between Theorem A.1 and [8, Theorem 3.25]. A minor difference is that we do not specify the various constants C appearing in the result. Another difference is in the presentation of (iii). Here we allow arbitrary $s \geq 1$ (instead of $s \geq \bar{s}$) at the expense of a larger (and unspecified) constant C . The main difference, however, is the additional term $\|\mathbf{c} - \mathbf{c}_S\|_{q,\mathbf{u};\mathcal{V}}$ appearing in (i). This can be shown as follows. First, one defines the sequence $\bar{\mathbf{c}} = (u_{\mathbf{v}}^{2/q-1} c_{\mathbf{v}})_{\mathbf{v} \in \mathbb{N}_0^d}$ so that $\|\mathbf{c} - \mathbf{c}_S\|_{q,\mathbf{u};\mathcal{V}} = \|\bar{\mathbf{c}} - \bar{\mathbf{c}}_S\|_{q;\mathcal{V}}$ and then uses Stechkin's inequality in lower sets (see, e.g., [8, Lemma 3.9]) to show that $\|\bar{\mathbf{c}} - \bar{\mathbf{c}}_S\|_{q;\mathcal{V}} \leq s^{1/q-1/p} \|\bar{\mathbf{c}}\|_{p,M;\mathcal{V}}$, where $\|\cdot\|_{p,M;\mathcal{V}}$ is the norm on the majorant ℓ^p space $\ell_M^p(\mathbb{N}_0^d; \mathcal{V})$ (see, e.g., [8, Definition 3.8]). Finally, it can be shown that $\|\bar{\mathbf{c}}\|_{p,M;\mathcal{V}} \leq C(d, p, \rho)$ using standard arguments. See, e.g., [8, Lemma 7.19] (this lemma only considers the scalar-valued case; however the extension to the Hilbert-valued case is straightforward).

Note that Theorem A.1 immediately implies Theorems 2.4 and 2.6. For the former, we note that

$$\|f - f_{S_1}\|_{L^2_{\bar{\theta}}(\mathcal{U};\mathcal{V})} = \|\mathbf{c} - \mathbf{c}_{S_1}\|_{2;\mathcal{V}} \quad \text{and} \quad \|f - f_{S_2}\|_{L^\infty(\mathcal{U};\mathcal{V})} \leq \|\mathbf{c} - \mathbf{c}_{S_2}\|_{1,\mathbf{u};\mathcal{V}}.$$

We then apply (i) with $q = 2$ or $q = 1$. For the latter, we use (iii) with $p = 1$.

A.2 The infinite-dimensional case

We now consider the infinite-dimensional case, where $d = \infty$ and $\mathcal{U} = [-1, 1]^{\mathbb{N}}$.

Theorem A.3 (Best s -term decay rates; infinite-dimensional case). *Let $d = \infty$, $0 < p < 1$, $\varepsilon > 0$, $\mathbf{b} \in \ell^p(\mathbb{N})$ with $\mathbf{b} > \mathbf{0}$ and $f \in \mathcal{B}(\mathbf{b}, \varepsilon)$, where $\mathcal{B}(\mathbf{b}, \varepsilon)$ is as in (2.7). Let $\mathbf{c} = (c_{\mathbf{v}})_{\mathbf{v} \in \mathcal{F}}$ be the Chebyshev or Legendre coefficients of f . Then the following best s -term decay rates hold:*

- (i) For any $p \leq q < \infty$ and $s \in \mathbb{N}$, there exists a lower set $S \subset \mathcal{F}$ of size $|S| \leq s$ such that

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

where $\sigma_s(\mathbf{c})_{q;\mathcal{V}}$ is as in Definition 2.3 (with $\Lambda = \mathcal{F}$), \mathbf{u} is as in (4.7) and $C = C(\mathbf{b}, \varepsilon, p) > 0$ depends on \mathbf{b}, ε and p only.

(ii) For any $p \leq q \leq 2$ and $k > 0$,

$$\sigma_k(\mathbf{c})_{q,\mathbf{u};\mathcal{V}} \leq C \cdot k^{1/q-1/p},$$

where $\sigma_k(\mathbf{c})_{q,\mathbf{u};\mathcal{V}}$ is as in Definition 7.2, \mathbf{u} is as in (4.7) and $C = (\mathbf{b}, \varepsilon, p) > 0$ depends on \mathbf{b} , ε and p only.

(iii) Suppose that \mathbf{b} is monotonically nonincreasing. Then, for any $p \leq q < \infty$ and $s \in \mathbb{N}$, there exists an anchored set $S \subset \mathcal{F}$ of size $|S| \leq s$ such that

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

where $\sigma_s(\mathbf{c})_{q;\mathcal{V}}$ is as in Definition 2.3 (with $\Lambda = \mathcal{F}$), \mathbf{u} is as in (4.7) and $C = (\mathbf{b}, \varepsilon, p) > 0$ depends on \mathbf{b} , ε and p only.

Remark A.4. This theorem is based on [8, Theorems 3.29 and 3.33]. Besides the term $\|\mathbf{c} - \mathbf{c}_S\|_{q,\mathbf{u};\mathcal{V}}$, parts (i) and (iii) can be found in [8, Theorem 3.29] and [8, Theorem 3.33], respectively. As in the finite-dimensional case (see Remark A.2), the main difference is the assertion of the bound on $\|\mathbf{c} - \mathbf{c}_S\|_{q,\mathbf{u};\mathcal{V}}$. This can be established through similar arguments, using either the majorant ℓ^p space $\ell_M^p(\mathcal{F}; \mathcal{V})$ or the anchored ℓ^p space $\ell_A^p(\mathcal{F}; \mathcal{V})$ (see, e.g., [8, Definition 3.31]) and then Stechkin's inequality in lower or anchored sets (see, e.g., [8, Lemma 3.32]). See also [8, Lemma 7.23] (this lemma only considers the scalar-valued case; however the extension to the Hilbert-valued case is straightforward).

Note that neither [8, Theorem 3.29] nor [8, Theorem 3.33] asserts part (ii) of Theorem A.3. This can be shown via the weighted Stechkin's inequality (see, e.g., [8, Lemma 3.12]), which gives the bound $\sigma_k(\mathbf{c})_{q,\mathbf{u};\mathcal{V}} \leq \|\mathbf{c}\|_{p,\mathbf{u};\mathcal{V}} \cdot k^{1/q-1/p}$, and then by showing that $\|\mathbf{c}\|_{p,\mathbf{u};\mathcal{V}} \leq C(\mathbf{b}, \varepsilon, p)$. This latter fact can be obtained by the straightforward extension of [8, Lemma 7.23] to the Hilbert-valued setting.

Note that Theorem A.3 implies Theorem 2.5. This follows from (i) by setting either $q = 2$ or $q = 1$.

References

- [1] B. Adcock, [Infinite-dimensional \$\ell^1\$ minimization and function approximation from pointwise data](#). *Constr. Approx.* **45** (2017), no. 3, 345–390
- [2] B. Adcock, [Infinite-dimensional compressed sensing and function interpolation](#). *Found. Comput. Math.* **18** (2018), no. 3, 661–701
- [3] B. Adcock, A. Bao, and S. Brugiapaglia, [Correcting for unknown errors in sparse high-dimensional function approximation](#). *Numer. Math.* **142** (2019), no. 3, 667–711
- [4] B. Adcock and S. Brugiapaglia, [Sparse approximation of multivariate functions from small datasets via weighted orthogonal matching pursuit](#). In *Spectral and high order methods for partial differential equations—ICOSAHOM 2018*, pp. 611–621, Lect. Notes Comput. Sci. Eng. 134, Springer, Cham, 2020
- [5] B. Adcock, S. Brugiapaglia, N. Dexter, and S. Moraga, [Deep neural networks are effective at learning high-dimensional Hilbert-valued functions from limited data](#). In *Proceedings of the second annual conference on mathematical and scientific machine learning*, pp. 1–36, Proc. Mach. Learn. Res. (PMLR) 145, 2021
- [6] B. Adcock, S. Brugiapaglia, N. Dexter, and S. Moraga, [Near-optimal learning of Banach-valued, high-dimensional functions via deep neural networks](#). [v1] 2022, [v2] 2023, arXiv:2211.12633v2
- [7] B. Adcock, S. Brugiapaglia, and C. G. Webster, [Compressed sensing approaches for polynomial approximation of high-dimensional functions](#). In *Compressed sensing and its applications*, pp. 93–124, Appl. Numer. Harmon. Anal., Birkhäuser/Springer, Cham, 2017
- [8] B. Adcock, S. Brugiapaglia, and C. G. Webster, [Sparse polynomial approximation of high-dimensional functions](#). *Comput. Sci. Eng.* 25, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2022
- [9] B. Adcock and J. M. Cardenas, [Near-optimal sampling strategies for multivariate function approximation on general domains](#). *SIAM J. Math. Data Sci.* **2** (2020), no. 3, 607–630
- [10] B. Adcock, J. M. Cardenas, N. Dexter, and S. Moraga, [Towards optimal sampling for learning sparse approximations in high dimensions](#). In *High-dimensional optimization and probability—with a view towards data science*, pp. 9–77, Springer Optim. Appl. 191, Springer, Cham, 2022
- [11] B. Adcock and N. Dexter, [The gap between theory and practice in function approximation with deep neural networks](#). *SIAM J. Math. Data Sci.* **3** (2021), no. 2, 624–655
- [12] B. Adcock, N. Dexter, and S. Moraga, [Optimal approximation of infinite-dimensional holomorphic functions](#). *Calcolo* **61** (2024), no. 1, article no. 12
- [13] B. Adcock and A. C. Hansen, [Compressive imaging: Structure, sampling, learning](#). Cambridge University Press, Cambridge, 2021
- [14] B. Adcock and D. Huybrechs, [Approximating smooth, multivariate functions on irregular domains](#). *Forum Math. Sigma* **8** (2020), article no. e26

- [15] B. Adcock and Y. Sui, [Compressive Hermite interpolation: Sparse, high-dimensional approximation from gradient-augmented measurements](#). *Constr. Approx.* **50** (2019), no. 1, 167–207
- [16] N. Alemazkour and H. Meidani, [Divide and conquer: An incremental sparsity promoting compressive sampling approach for polynomial chaos expansions](#). *Comput. Methods Appl. Mech. Engrg.* **318** (2017), 937–956
- [17] N. Alemazkour and H. Meidani, [A near-optimal sampling strategy for sparse recovery of polynomial chaos expansions](#). *J. Comput. Phys.* **371** (2018), 137–151
- [18] S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells, [The FEniCS Project Version 1.5](#). *Archive of Numerical Software* **3** (2015), no. 100, 9–23
- [19] M. Bachmayr, A. Cohen, R. DeVore, and G. Migliorati, [Sparse polynomial approximation of parametric elliptic PDEs. Part II: Lognormal coefficients](#). *ESAIM Math. Model. Numer. Anal.* **51** (2017), no. 1, 341–363
- [20] M. Bachmayr, A. Cohen, and G. Migliorati, [Sparse polynomial approximation of parametric elliptic PDEs. Part I: Affine coefficients](#). *ESAIM Math. Model. Numer. Anal.* **51** (2017), no. 1, 321–339
- [21] F. Bartel, M. Schäfer, and T. Ullrich, [Constructive subsampling of finite frames with applications in optimal function recovery](#). *Appl. Comput. Harmon. Anal.* **65** (2023), 209–248
- [22] H. H. Bauschke and P. L. Combettes, [Convex analysis and monotone operator theory in Hilbert spaces](#). 2nd edn., CMS Books Math./Ouvrages Math. SMC, Springer, Cham, 2017
- [23] J. Beck, F. Nobile, L. Tamellini, and R. Tempone, [Convergence of quasi-optimal stochastic Galerkin methods for a class of PDEs with random coefficients](#). *Comput. Math. Appl.* **67** (2014), no. 4, 732–751
- [24] J. Beck, R. Tempone, F. Nobile, and L. Tamellini, [On the optimal polynomial approximation of stochastic PDEs by Galerkin and collocation methods](#). *Math. Models Methods Appl. Sci.* **22** (2012), no. 9, article no. 1250023
- [25] M. Bieri, R. Andreev, and C. Schwab, [Sparse tensor discretization of elliptic SPDEs](#). *SIAM J. Sci. Comput.* **31** (2009/10), no. 6, 4281–4304
- [26] G. Blatman and B. Sudret, [Adaptive sparse polynomial chaos expansion based on least angle regression](#). *J. Comput. Phys.* **230** (2011), no. 6, 2345–2367
- [27] A. Bonito, R. DeVore, D. Guignard, P. Jantsch, and G. Petrova, [Polynomial approximation of anisotropic analytic functions of several variables](#). *Constr. Approx.* **53** (2021), no. 2, 319–348
- [28] J.-L. Bouchot, H. Rauhut, and C. Schwab, [Multi-level compressed sensing Petrov–Galerkin discretization of high-dimensional parametric PDEs](#). [v1] 2017, [v2] 2017, arXiv:1701.01671v2
- [29] E. J. Candès, J. Romberg, and T. Tao, [Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information](#). *IEEE Trans. Inform. Theory* **52** (2006), no. 2, 489–509

- [30] A. Chambolle and T. Pock, [A first-order primal-dual algorithm for convex problems with applications to imaging](#). *J. Math. Imaging Vision* **40** (2011), no. 1, 120–145
- [31] A. Chambolle and T. Pock, [An introduction to continuous optimization for imaging](#). *Acta Numer.* **25** (2016), 161–319
- [32] A. Chernov and D. Dũng, [New explicit-in-dimension estimates for the cardinality of high-dimensional hyperbolic crosses and approximation of functions having mixed smoothness](#). *J. Complexity* **32** (2016), no. 1, 92–121
- [33] A. Chkifa, A. Cohen, R. DeVore, and C. Schwab, [Sparse adaptive Taylor approximation algorithms for parametric and stochastic elliptic PDEs](#). *ESAIM Math. Model. Numer. Anal.* **47** (2013), no. 1, 253–280
- [34] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, and R. Tempone, [Discrete least squares polynomial approximation with random evaluations—application to parametric and stochastic elliptic PDEs](#). *ESAIM Math. Model. Numer. Anal.* **49** (2015), no. 3, 815–837
- [35] A. Chkifa, A. Cohen, and C. Schwab, [High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs](#). *Found. Comput. Math.* **14** (2014), no. 4, 601–633
- [36] A. Chkifa, A. Cohen, and C. Schwab, [Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs](#). *J. Math. Pures Appl. (9)* **103** (2015), no. 2, 400–428
- [37] A. Chkifa, N. Dexter, H. Tran, and C. G. Webster, [Polynomial approximation via compressed sensing of high-dimensional functions on lower sets](#). *Math. Comp.* **87** (2018), no. 311, 1415–1450
- [38] B. Choi, M. Iwen, and T. Volkmer, [Sparse harmonic transforms II: Best \$s\$ -term approximation guarantees for bounded orthonormal product bases in sublinear-time](#). *Numer. Math.* **148** (2021), no. 2, 293–362
- [39] B. Choi, M. A. Iwen, and F. Krahmer, [Sparse harmonic transforms: A new class of sublinear-time algorithms for learning functions of many variables](#). *Found. Comput. Math.* **21** (2021), no. 2, 275–329
- [40] A. Cohen, M. A. Davenport, and D. Leviatan, [On the stability and accuracy of least squares approximations](#). *Found. Comput. Math.* **13** (2013), no. 5, 819–834
- [41] A. Cohen and R. DeVore, [Approximation of high-dimensional parametric PDEs](#). *Acta Numer.* **24** (2015), 1–159
- [42] A. Cohen, R. DeVore, and C. Schwab, [Convergence rates of best \$N\$ -term Galerkin approximations for a class of elliptic sPDEs](#). *Found. Comput. Math.* **10** (2010), no. 6, 615–646
- [43] A. Cohen, R. DeVore, and C. Schwab, [Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's](#). *Anal. Appl. (Singap.)* **9** (2011), no. 1, 11–47
- [44] A. Cohen and G. Migliorati, [Optimal weighted least-squares methods](#). *SMAI J. Comput. Math.* **3** (2017), 181–203
- [45] A. Cohen and G. Migliorati, [Multivariate approximation in downward closed polynomial spaces](#). In *Contemporary computational mathematics—a celebration of the 80th birthday of Ian Sloan. Vol. 1, 2*, pp. 233–282, Springer, Cham, 2018

- [46] A. Cohen, G. Migliorati, and F. Nobile, [Discrete least-squares approximations over optimized downward closed polynomial spaces in arbitrary dimension](#). *Constr. Approx.* **45** (2017), no. 3, 497–519
- [47] M. J. Colbrook, [WARPd: a linearly convergent first-order primal-dual algorithm for inverse problems with approximate sharpness conditions](#). *SIAM J. Imaging Sci.* **15** (2022), no. 3, 1539–1575
- [48] M. J. Colbrook, V. Antun, and A. C. Hansen, [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** (2022), no. 12, article no. e2107151119
- [49] J. Daws and C. Webster, [Analysis of deep neural networks with Quasi-optimal polynomial approximation rates](#). 2019, arXiv:1912.02302
- [50] C. de Boor and A. Ron, [Computational aspects of polynomial interpolation in several variables](#). *Math. Comp.* **58** (1992), no. 198, 705–727
- [51] N. Dexter, [Sparse reconstruction techniques for solutions of high-dimensional parametric PDEs](#). Ph.D. thesis, University of Tennessee, USA, 2018
- [52] N. Dexter, H. Tran, and C. Webster, [A mixed \$\ell_1\$ regularization approach for sparse simultaneous approximation of parameterized PDEs](#). *ESAIM Math. Model. Numer. Anal.* **53** (2019), no. 6, 2025–2045
- [53] P. Diaz, A. Doostan, and J. Hampton, [Sparse polynomial chaos expansions via compressed sensing and D-optimal design](#). *Comput. Methods Appl. Mech. Engrg.* **336** (2018), 640–666
- [54] M. Dolbeault and A. Cohen, [Optimal pointwise sampling for \$L^2\$ approximation](#). *J. Complexity* **68** (2022), article no. 101602
- [55] M. Dolbeault and A. Cohen, [Optimal sampling and Christoffel functions on general domains](#). *Constr. Approx.* **56** (2022), no. 1, 121–163
- [56] M. Dolbeault, D. Krieg, and M. Ullrich, [A sharp upper bound for sampling numbers in \$L_2\$](#) . *Appl. Comput. Harmon. Anal.* **63** (2023), 113–134
- [57] D. L. Donoho, [Compressed sensing](#). *IEEE Trans. Inform. Theory* **52** (2006), no. 4, 1289–1306
- [58] A. Doostan and H. Owhadi, [A non-adapted sparse approximation of PDEs with stochastic inputs](#). *J. Comput. Phys.* **230** (2011), no. 8, 3015–3034
- [59] D. Dũng and V. K. Nguyen, [Deep ReLU neural networks in high-dimensional approximation](#). *Neural Networks* **142** (2021), 619–635
- [60] D. Dũng, V. K. Nguyen, and M. X. Thao, [Computation complexity of deep ReLU neural networks in high-dimensional approximation](#). *J. Crohn’s Colitis* **37** (2021), no. 3, 291–320
- [61] S. Foucart and H. Rauhut, [A mathematical introduction to compressive sensing](#). Appl. Numer. Harmon. Anal., Birkhäuser/Springer, New York, 2013
- [62] R. Ghanem, D. Higdon, and H. Owhadi, [Handbook of uncertainty quantification](#). Springer, Switzerland, 2017

- [63] C. J. Gittelson, [An adaptive stochastic Galerkin method for random elliptic operators](#). *Math. Comp.* **82** (2013), no. 283, 1515–1541
- [64] L. Guo, Y. Liu, and L. Yan, [Sparse recovery via \$\ell_q\$ -minimization for polynomial chaos expansions](#). *Numer. Math. Theory Methods Appl.* **10** (2017), no. 4, 775–797
- [65] L. Guo, A. Narayan, L. Yan, and T. Zhou, [Weighted approximate Fekete points: sampling for least-squares polynomial approximation](#). *SIAM J. Sci. Comput.* **40** (2018), no. 1, A366–A387
- [66] L. Guo, A. Narayan, and T. Zhou, [A gradient enhanced \$\ell_1\$ -minimization for sparse approximation of polynomial chaos expansions](#). *J. Comput. Phys.* **367** (2018), 49–64
- [67] L. Guo, A. Narayan, and T. Zhou, [Constructing least-squares polynomial approximations](#). *SIAM Rev.* **62** (2020), no. 2, 483–508
- [68] L. Guo, A. Narayan, T. Zhou, and Y. Chen, [Stochastic collocation methods via \$\ell_1\$ minimization using randomized quadratures](#). *SIAM J. Sci. Comput.* **39** (2017), no. 1, A333–A359
- [69] M. Hadigol and A. Doostan, [Least squares polynomial chaos expansion: A review of sampling strategies](#). *Comput. Methods Appl. Mech. Engrg.* **332** (2018), 382–407
- [70] A.-L. Haji-Ali, F. Nobile, R. Tempone, and S. Wolfers, [Multilevel weighted least squares polynomial approximation](#). *ESAIM Math. Model. Numer. Anal.* **54** (2020), no. 2, 649–677
- [71] J. Hampton and A. Doostan, [Coherence motivated sampling and convergence analysis of least squares polynomial Chaos regression](#). *Comput. Methods Appl. Mech. Engrg.* **290** (2015), 73–97
- [72] J. Hampton and A. Doostan, [Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies](#). *J. Comput. Phys.* **280** (2015), 363–386
- [73] J. Hampton and A. Doostan, [Compressive sampling methods for sparse polynomial chaos expansions](#). In *Handbook of uncertainty quantification. Vol. 1, 2, 3*, pp. 827–855, Springer, Cham, 2017
- [74] J. Hampton and A. Doostan, [Basis adaptive sample efficient polynomial chaos \(BASE-PC\)](#). *J. Comput. Phys.* **371** (2018), 20–49
- [75] M. Hansen and C. Schwab, [Analytic regularity and nonlinear approximation of a class of parametric semilinear elliptic PDEs](#). *Math. Nachr.* **286** (2013), no. 8-9, 832–860
- [76] L. Herrmann, J. A. A. Opschoor, and C. Schwab, [Constructive deep ReLU neural network approximation](#). *J. Sci. Comput.* **90** (2022), no. 2, article no. 75
- [77] M. Hervé, [Analyticity in infinite-dimensional spaces](#). De Gruyter Stud. Math. 10, Walter de Gruyter, Berlin, 1989
- [78] J. S. Hesthaven, G. Rozza, and B. Stamm, [Certified reduced basis methods for parametrized partial differential equations](#). SpringerBriefs Math., Springer, Cham; BCAM Basque Center for Applied Mathematics, Bilbao, 2016
- [79] A. Hinrichs, D. Krieg, E. Novak, J. Prochno, and M. Ullrich, [On the power of random information](#). In *Multivariate algorithms and information-based complexity*, pp. 43–64, Radon Ser. Comput. Appl. Math. 27, De Gruyter, Berlin, 2020

- [80] L. S. T. Ho, H. Schaeffer, G. Tran, and R. Ward, [Recovery guarantees for polynomial coefficients from weakly dependent data with outliers](#). *J. Approx. Theory* **259** (2020), article no. 105472
- [81] F. L. Huang and S. Zhang, [Approximation of infinitely differentiable multivariate functions is not strongly tractable](#). *J. Complexity* **23** (2007), no. 1, 73–81
- [82] J. D. Jakeman, M. S. Eldred, and K. Sargsyan, [Enhancing \$\ell_1\$ -minimization estimates of polynomial chaos expansions using basis selection](#). *J. Comput. Phys.* **289** (2015), 18–34
- [83] J. D. Jakeman, A. Narayan, and T. Zhou, [A generalized sampling and preconditioning scheme for sparse approximation of polynomial chaos expansions](#). *SIAM J. Sci. Comput.* **39** (2017), no. 3, A1114–A1144
- [84] J.-H. Jung, S. Gottlieb, S. O. Kim, C. L. Bresten, and D. Higgs, [Recovery of high order accuracy in radial basis function approximations of discontinuous problems](#). *J. Sci. Comput.* **45** (2010), no. 1-3, 359–381
- [85] L. Kämmerer, T. Ullrich, and T. Volkmer, [Worst-case recovery guarantees for least squares approximation using random samples](#). *Constr. Approx.* **54** (2021), no. 2, 295–352
- [86] I. A. Kougioumtzoglou, I. Petromichelakis, and A. F. Psaros, [Sparse representations and compressive sampling approaches in engineering mechanics: A review of theoretical concepts and diverse applications](#). *Probabilistic Eng. Mech.* **61** (2020), article no. 103082
- [87] D. Krieg, E. Novak, and M. Sonnleitner, [Recovery of Sobolev functions restricted to iid sampling](#). *Math. Comp.* **91** (2022), no. 338, 2715–2738
- [88] D. Krieg, P. Siedlecki, M. Ullrich, and H. Woźniakowski, [Exponential tractability of \$L_2\$ -approximation with function values](#). *Adv. Comput. Math.* **49** (2023), no. 2, article no. 18
- [89] T. Kühn, W. Sickel, and T. Ullrich, [Approximation of mixed order Sobolev functions on the \$d\$ -torus: asymptotics, preasymptotics, and \$d\$ -dependence](#). *Constr. Approx.* **42** (2015), no. 3, 353–398
- [90] J. Kuntzman, *Méthodes Numériques – Interpolation, Dérivées*. Dunod, Paris, France, 1959
- [91] O. P. Le Maître and O. M. Knio, *Spectral methods for uncertainty quantification*. Sci. Comput., Springer, New York, 2010
- [92] B. Li, S. Tang, and H. Yu, [Better approximations of high dimensional smooth functions by deep neural networks with rectified power units](#). *Commun. Comput. Phys.* **27** (2020), no. 2, 379–411
- [93] I. Limonova and V. Temlyakov, [On sampling discretization in \$L_2\$](#) . *J. Math. Anal. Appl.* **515** (2022), no. 2, article no. 126457
- [94] Y. Liu and L. Guo, [Stochastic collocation via \$l_1\$ -minimisation on low discrepancy point sets with application to uncertainty quantification](#). *East Asian J. Appl. Math.* **6** (2016), no. 2, 171–191
- [95] A. Logg and G. N. Wells, [DOLFIN: Automated finite element computing](#). *ACM Trans. Math. Software* **37** (2010), no. 2, article no. 20

- [96] G. G. Lorentz and R. A. Lorentz, [Solvability problems of bivariate interpolation. I. *Constr. Approx.* **2** \(1986\), no. 2, 153–169](#)
- [97] N. Lüthen, S. Marelli, and B. Sudret, [Sparse polynomial chaos expansions: Literature survey and benchmark. *SIAM/ASA J. Uncertain. Quantif.* **9** \(2021\), no. 2, 593–649](#)
- [98] N. Lüthen, S. Marelli, and B. Sudret, [Automatic selection of basis-adaptive sparse polynomial chaos expansions for engineering applications. *Int. J. Uncertain. Quantif.* **12** \(2022\), no. 3, 49–74](#)
- [99] L. Mathelin and K. A. Gallivan, [A compressed sensing approach for partial differential equations with random input data. *Commun. Comput. Phys.* **12** \(2012\), no. 4, 919–954](#)
- [100] G. Migliorati, [Polynomial approximation by means of the random discrete \$L^2\$ projection and application to inverse problems for PDEs with stochastic data.](#) Ph.D. thesis, Politecnico di Milano, Italy, 2013
- [101] G. Migliorati, Adaptive polynomial approximation by means of random discrete least squares. In *Numerical mathematics and advanced applications—ENUMATH 2013*, pp. 547–554, Lect. Notes Comput. Sci. Eng. 103, Springer, Cham, 2015
- [102] G. Migliorati, [Adaptive approximation by optimal weighted least-squares methods. *SIAM J. Numer. Anal.* **57** \(2019\), no. 5, 2217–2245](#)
- [103] G. Migliorati, [Multivariate approximation of functions on irregular domains by weighted least-squares methods. *IMA J. Numer. Anal.* **41** \(2021\), no. 2, 1293–1317](#)
- [104] G. Migliorati and F. Nobile, [Analysis of discrete least squares on multivariate polynomial spaces with evaluations at low-discrepancy point sets. *J. Complexity* **31** \(2015\), no. 4, 517–542](#)
- [105] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone, [Analysis of discrete \$L^2\$ projection on polynomial spaces with random evaluations. *Found. Comput. Math.* **14** \(2014\), no. 3, 419–456](#)
- [106] A. Narayan, [Computation of induced orthogonal polynomial distributions. *Electron. Trans. Numer. Anal.* **50** \(2018\), 71–97](#)
- [107] A. Narayan, J. D. Jakeman, and T. Zhou, [A Christoffel function weighted least squares algorithm for collocation approximations. *Math. Comp.* **86** \(2017\), no. 306, 1913–1947](#)
- [108] A. Narayan and T. Zhou, [Stochastic collocation on unstructured multivariate meshes. *Commun. Comput. Phys.* **18** \(2015\), no. 1, 1–36](#)
- [109] L. W.-T. Ng and M. Eldred, [Multifidelity uncertainty quantification using noninvasive polynomial chaos and stochastic collocation.](#) In *53rd AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics and materials conference*, 45, AIAA, 2012
- [110] F. Nobile, R. Tempone, and C. G. Webster, [A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* **46** \(2008\), no. 5, 2309–2345](#)
- [111] E. Novak and H. Woźniakowski, [Tractability of multivariate problems. Vol. 1: Linear information.](#) EMS Tracts Math. 6, European Mathematical Society (EMS), Zürich, 2008
- [112] E. Novak and H. Woźniakowski, [Approximation of infinitely differentiable multivariate functions is intractable. *J. Complexity* **25** \(2009\), no. 4, 398–404](#)

- [113] E. Novak and H. Woźniakowski, *Tractability of multivariate problems. Volume II: Standard information for functionals*. EMS Tracts Math. 12, European Mathematical Society (EMS), Zürich, 2010
- [114] J. A. A. Opschoor, C. Schwab, and J. Zech, [Deep learning in high dimension: ReLU neural network expression for Bayesian PDE inversion](#). In *Optimization and control for partial differential equations—uncertainty quantification, open and closed-loop control, and shape optimization*, pp. 419–462, Radon Ser. Comput. Appl. Math. 29, De Gruyter, Berlin, 2022
- [115] J. A. A. Opschoor, C. Schwab, and J. Zech, [Exponential ReLU DNN expression of holomorphic maps in high dimension](#). *Constr. Approx.* **55** (2022), no. 1, 537–582
- [116] J. Peng, J. Hampton, and A. Doostan, [A weighted \$\ell_1\$ -minimization approach for sparse polynomial chaos expansions](#). *J. Comput. Phys.* **267** (2014), 92–111
- [117] J. Peng, J. Hampton, and A. Doostan, [On polynomial chaos expansion via gradient-enhanced \$\ell_1\$ -minimization](#). *J. Comput. Phys.* **310** (2016), 440–458
- [118] A. Quarteroni, A. Manzoni, and F. Negri, *Reduced basis methods for partial differential equations*. Unitext 92, Springer, Cham, 2016
- [119] H. Rauhut and C. Schwab, [Compressive sensing Petrov–Galerkin approximation of high-dimensional parametric operator equations](#). *Math. Comp.* **86** (2017), no. 304, 661–700
- [120] H. Rauhut and R. Ward, [Sparse Legendre expansions via \$\ell_1\$ -minimization](#). *J. Approx. Theory* **164** (2012), no. 5, 517–533
- [121] H. Rauhut and R. Ward, [Interpolation via weighted \$\ell_1\$ minimization](#). *Appl. Comput. Harmon. Anal.* **40** (2016), no. 2, 321–351
- [122] J. Renegar and B. Grimmer, [A simple nearly optimal restart scheme for speeding up first-order methods](#). *Found. Comput. Math.* **22** (2022), no. 1, 211–256
- [123] V. Roulet, N. Boumal, and A. d’Aspremont, [Computational complexity versus statistical performance on sparse recovery problems](#). *Inf. Inference* **9** (2020), no. 1, 1–32
- [124] V. Roulet and A. d’Aspremont, [Sharpness, restart, and acceleration](#). *SIAM J. Optim.* **30** (2020), no. 1, 262–289
- [125] C. Schwab and J. Zech, [Deep learning in high dimension: Neural network expression rates for analytic functions in \$L^2\(\mathbb{R}^d, \gamma_d\)\$](#) . *SIAM/ASA J. Uncertain. Quantif.* **11** (2023), no. 1, 199–234
- [126] P. Seshadri, A. Narayan, and S. Mahadevan, [Effectively subsampled quadratures for least squares polynomial approximations](#). *SIAM/ASA J. Uncertain. Quantif.* **5** (2017), no. 1, 1003–1023
- [127] Y. Shin and D. Xiu, [Correcting data corruption errors for multivariate function approximation](#). *SIAM J. Sci. Comput.* **38** (2016), no. 4, A2492–A2511
- [128] R. C. Smith, *Uncertainty quantification*. Comput. Sci. Eng. 12, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2014
- [129] M. Stoyanov, *User manual: Tasmanian sparse grids*. Technical report ornl/tm-2015/596, Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge, TN, 2015

- [130] Y. Sui, *Weighted ℓ^1 minimization techniques for compressed sensing and their applications*. Ph.D. thesis, Simon Fraser University, Canada, 2020
- [131] T. J. Sullivan, *Introduction to uncertainty quantification*. Texts Appl. Math. 63, Springer, Cham, 2015
- [132] G. Szegő, *Orthogonal Polynomials*. Amer. Math. Soc. Colloq. Publ. 23, American Mathematical Society, New York, 1939
- [133] G. Tang, *Methods for High Dimensional Uncertainty Quantification: Regularization, Sensitivity Analysis, and Derivative Enhancement*. Ph.D. thesis, Stanford University, 2013
- [134] G. Tang and G. Iaccarino, [Subsampled Gauss quadrature nodes for estimating polynomial chaos expansions](#). *SIAM/ASA J. Uncertain. Quantif.* **2** (2014), no. 1, 423–443
- [135] T. Tang and T. Zhou, [On discrete least-squares projection in unbounded domain with random evaluations and its application to parametric uncertainty quantification](#). *SIAM J. Sci. Comput.* **36** (2014), no. 5, A2272–A2295
- [136] T. Tao, *An introduction to measure theory*. Grad. Stud. Math. 126, American Mathematical Society, Providence, RI, 2011
- [137] V. Temlyakov, [On optimal recovery in \$L_2\$](#) . *J. Complexity* **65** (2021), article no. 101545
- [138] V. N. Temlyakov, Approximation of periodic functions of several variables with a bounded mixed derivative. *Trudy Mat. Inst. Steklov.* **156** (1980), article no. 233–260
- [139] R. A. Todor and C. Schwab, [Convergence rates for sparse chaos approximations of elliptic problems with stochastic coefficients](#). *IMA J. Numer. Anal.* **27** (2007), no. 2, 232–261
- [140] H. Tran and C. Webster, [A class of null space conditions for sparse recovery via nonconvex, non-separable minimizations](#). *Results Appl. Math.* **3** (2019), article no. 100011
- [141] H. Tran, C. G. Webster, and G. Zhang, [Analysis of quasi-optimal polynomial approximations for parameterized PDEs with deterministic and stochastic coefficients](#). *Numer. Math.* **137** (2017), no. 2, 451–493
- [142] P. Tsilifis, X. Huan, C. Safta, K. Sargsyan, G. Lacaze, J. C. Oefelein, H. N. Najm, and R. G. Ghanem, [Compressive sensing adaptation for polynomial chaos expansions](#). *J. Comput. Phys.* **380** (2019), 29–47
- [143] E. van den Berg and M. P. Friedlander, [Probing the Pareto frontier for basis pursuit solutions](#). *SIAM J. Sci. Comput.* **31** (2008/09), no. 2, 890–912
- [144] E. van den Berg and M. P. Friedlander, A solver for large-scale sparse reconstruction. 2019, <https://friedlander.io/spg11>, visited on 5 March 2024
- [145] M. Vidyasagar, *An introduction to compressed sensing*. Comput. Sci. Eng. 22, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2020
- [146] W. Walter, *Ordinary differential equations*. Grad. Texts in Math. 182, Springer, New York, 1998
- [147] G. Xu, [On weak tractability of the Smolyak algorithm for approximation problems](#). *J. Approx. Theory* **192** (2015), 347–361
- [148] Y. Xu, A. Narayan, H. Tran, and C. G. Webster, [Analysis of the ratio of \$\ell_1\$ and \$\ell_2\$ norms in compressed sensing](#). *Appl. Comput. Harmon. Anal.* **55** (2021), 486–511

- [149] L. Yan, L. Guo, and D. Xiu, [Stochastic collocation algorithms using \$\ell_1\$ -minimization](#). *Int. J. Uncertain. Quantif.* **2** (2012), no. 3, 279–293
- [150] L. Yan, Y. Shin, and D. Xiu, [Sparse approximation using \$\ell_1 - \ell_2\$ minimization and its application to stochastic collocation](#). *SIAM J. Sci. Comput.* **39** (2017), no. 1, A214–A239
- [151] X. Yang and G. E. Karniadakis, [Reweighted \$\ell_1\$ minimization method for stochastic elliptic differential equations](#). *J. Comput. Phys.* **248** (2013), 87–108
- [152] X. Yang, H. Lei, N. A. Baker, and G. Lin, [Enhancing sparsity of Hermite polynomial expansions by iterative rotations](#). *J. Comput. Phys.* **307** (2016), 94–109
- [153] X. Yang, W. Li, and A. Tartakovsky, [Sliced-inverse-regression-aided rotated compressive sensing method for uncertainty quantification](#). *SIAM/ASA J. Uncertain. Quantif.* **6** (2018), no. 4, 1532–1554
- [154] X. Yang, X. Wan, L. Lin, and H. Lei, [A general framework for enhancing sparsity of generalized polynomial chaos expansions](#). *Int. J. Uncertain. Quantif.* **9** (2019), no. 3, 221–243
- [155] T. Zhou, A. Narayan, and D. Xiu, [Weighted discrete least-squares polynomial approximation using randomized quadratures](#). *J. Comput. Phys.* **298** (2015), 787–800
- [156] T. Zhou, A. Narayan, and Z. Xu, [Multivariate discrete least-squares approximations with a new type of collocation grid](#). *SIAM J. Sci. Comput.* **36** (2014), no. 5, A2401–A2422

Ben Adcock, Simone Brugiapaglia,
Nick Dexter, Sebastian Moraga

On Efficient Algorithms for Computing Near-Best Polynomial Approximations to High-Dimensional, Hilbert-Valued Functions from Limited Samples

Sparse polynomial approximation is an important tool for approximating high-dimensional functions from limited samples – a task commonly arising in computational science and engineering. Yet, it lacks a complete theory. There is a well-developed theory of *best s -term polynomial approximation*, which asserts exponential or algebraic rates of convergence for holomorphic functions. There are also increasingly mature methods such as (weighted) ℓ^1 -minimization for practically computing such approximations. However, whether these methods achieve the rates of the best s -term approximation is not fully understood. Moreover, these methods are not algorithms per se, since they involve exact minimizers of nonlinear optimization problems. This paper closes these gaps by affirmatively answering the following question: *are there robust, efficient algorithms for computing sparse polynomial approximations to finite- or infinite-dimensional, holomorphic and Hilbert-valued functions from limited samples that achieve the same rates as the best s -term approximation?* We do so by introducing algorithms with exponential or algebraic convergence rates that are also robust to *sampling, algorithmic and physical discretization* errors. Our results involve several developments of existing techniques, including a new restarted primal-dual iteration for solving weighted ℓ^1 -minimization problems in Hilbert spaces. Our theory is supplemented by numerical experiments demonstrating the efficacy of these algorithms.

<https://ems.press>

ISSN 2747-9080

ISBN 978-3-98547-070-9

