



APPROXIMATING SMOOTH, MULTIVARIATE FUNCTIONS ON IRREGULAR DOMAINS

BEN ADCOCK¹ and DAAN HUYBRECHS²

¹ Department of Mathematics, Simon Fraser University, 8888 University Drive, Burnaby, BC V5A 1S6, Canada;

email: ben_adcock@sfu.ca

² Department of Computer Science, KU Leuven, Celestijnenlaan 200A, BE-3001 Leuven, Belgium;

email: daan.huybrechs@cs.kuleuven.be

Received 26 November 2018; accepted 3 March 2020

Abstract

In this paper, we introduce a method known as *polynomial frame approximation* for approximating smooth, multivariate functions defined on irregular domains in d dimensions, where d can be arbitrary. This method is simple, and relies only on orthogonal polynomials on a bounding tensor-product domain. In particular, the domain of the function need not be known in advance. When restricted to a subdomain, an orthonormal basis is no longer a basis, but a frame. Numerical computations with frames present potential difficulties, due to the near-linear dependence of the truncated approximation system. Nevertheless, well-conditioned approximations can be obtained via regularization, for instance, truncated singular value decompositions. We comprehensively analyze such approximations in this paper, providing error estimates for functions with both classical and mixed Sobolev regularity, with the latter being particularly suitable for higher-dimensional problems. We also analyze the sample complexity of the approximation for sample points chosen randomly according to a probability measure, providing estimates in terms of the corresponding *Nikolskii inequality* for the domain. In particular, we show that the sample complexity for points drawn from the uniform measure is quadratic (up to a log factor) in the dimension of the polynomial space, independently of d , for a large class of nontrivial domains. This extends a well-known result for polynomial approximation in hypercubes.

2010 Mathematics Subject Classification: 41A10, 41A25 (primary); 41A63, 41A17 (secondary)

1. Introduction

Many problems in computational science call for the approximation of smooth, multivariate functions. This problem is often challenging, due to the *curse of dimensionality*. Yet significant strides have been made over the last several decades toward its mitigation, typically by assuming some anisotropic behavior of the function being approximated. Approaches such as sparse grids [16] have enjoyed substantial success in the numerical solution of high-dimensional partial differential equations (PDEs), and more recently techniques based on computing multivariate polynomial approximations—often referred to as generalized polynomial chaos expansions [50]—have begun to be widely used for problems in uncertainty quantification (UQ) (see [4, 18, 19, 22, 23, 27, 39, 52] and references therein).

The majority of algorithms for high-dimensional approximation assume the underlying function f is defined over a tensor-product domain. The key benefit of doing so is simplicity. Indeed, the orthogonal polynomials on a tensor-product domain with respect to a tensor-product measure are precisely tensor products of the corresponding one-dimensional orthogonal polynomials. Yet there are many practical instances where the domain of interest is not of tensor-product type. One example is surrogate model construction in UQ. In practice, it is often the case that the random variables are correlated [49], which leads to an irregular domain. Alternatively or in addition, the given forward model may not be well defined over the whole of the assumed tensor-product domain, or may produce values in certain regions that are known to be unphysical (for example, negative pressures). This in effect leads to failed evaluations, resulting once more in an irregular domain [45]. Similarly, in model order reduction, techniques such as active subspaces [26] lead to approximation problems over irregular domains. For example, when a function defined on a high-dimensional hypercube is projected to a function of a reduced set of parameters, the resulting domain (the projection of the hypercube) is generally polyhedral, a so-called zonotope [46]. Finally, many applications in UQ also involve forward models which are piecewise smooth (see [44, 45] and references therein). Unless such discontinuities happen to be aligned along coordinate axes, this results in an approximation problem involving two or more smooth functions defined over irregular domains.

With this issue in mind, the purpose of this paper is to present a systematic study of a simple but effective technique for approximating high-dimensional functions defined on irregular domains. It is based on using tensor-product orthogonal polynomials on a bounding box, and is referred to as *polynomial frame approximation*. The approach corresponds to approximation in a frame, rather than a basis, since there are potentially many ways the unknown function on the irregular domain can be represented in a basis on the bounding box. Our main

results demonstrate that this procedure achieves (to a significant degree) the four primary criteria for a numerical approximation scheme: namely, *simplicity*, *accuracy*, *stability* and *efficiency*. We elaborate on the meaning of these terms in the next section; however, we note in passing that simplicity means that the same procedure can be applied to a broad class of irregular domains. In particular, no costly parametrization of the domain or its boundary (a potentially infeasible task in high dimensions) is required to construct the approximation. Instead, we will make the less restrictive assumption that samples can be randomly drawn from a certain measure on Ω , related to the orthogonality measure on the bounding box. Typically, this is taken as the uniform measure on Ω . Note that our focus in this paper is neither on the best choice of measure nor on the question of how to sample from a given measure. These are challenging, and potentially highly domain-dependent, issues, whereas in this work we strive for generality. We return briefly to this question in Section 9.

The main contribution of this paper is the rigorous analysis of polynomial frame approximations. Central to this is the notion of frames of Hilbert spaces, as opposed to more conventional orthogonal bases. We stress at this point that our technique does not attempt to orthogonalize a basis. Instead, it relies on the particular properties of frames to achieve accurate and stable approximations. A key facet of frame approximations (not just of polynomial type) is that they lead to highly ill-conditioned linear systems of equations. However, by using regularization, we are able to obtain a mapping from the sample points to the polynomial space that is both well conditioned and accurate. We also determine approximation rates and sample complexity estimates that scale well with the underlying dimension, thus (on the proviso that samples can be drawn efficiently from the desired measure—see above) mitigating the curse of dimensionality to a significant extent.

Before we proceed further, it is worth noting that polynomial frame approximation, and variations thereof, are in essence already used in many of the aforementioned applications. Indeed, any approach to surrogate model construction in UQ which computes a generalized polynomial chaos expansion from function evaluations which are limited (due to the particular problem at hand) to a nontensorial subdomain is equivalent to polynomial frame approximation. See Section 2.4 for further details. However, a thorough analysis of the accuracy, stability and efficiency of such approximations—in particular, exploiting the connections to frame theory as we do in this paper—is, to the best of our knowledge, lacking. Besides providing the first clear theoretical explanation for why these algorithms work in the practical setting of irregular domains, we also expect the results of this paper to shed light on ways in which to improve them. For example, the problem of designing better sampling sets for irregular domains is a topic of significant practical interest.

2. Overview of the paper

We commence with a short overview of the paper.

2.1. Polynomial frame approximations. This paper concerns the approximation of a smooth multivariate function $f : \Omega \rightarrow \mathbb{C}$ defined over a nontensor product domain $\Omega \subset \mathbb{R}^d$. The approximation is based on four key steps:

- (i) Choose a tensor-product domain D such that $\Omega \subseteq D$.
- (ii) Choose a tensor-product probability measure ν on D , a tensor-product orthonormal basis $\{\psi_n\}_n$ of $L^2(D, \nu)$ and a finite index set Λ with $|\Lambda| = N$.
- (iii) Take M samples of f of the form $f(\mathbf{y}_1), \dots, f(\mathbf{y}_M)$, where $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_M\} \subset \Omega$.
- (iv) Compute an approximation to f of the form $f_{\mathcal{Y}, \Lambda} = \sum_{n \in \Lambda} c_n \phi_n$, where $\phi_n = \psi_n|_{\Omega}$.

This immediately raises a number of questions, which are now discussed:

1. *How to compute the approximation.* There are several options for doing this, including interpolation if $M = |\Lambda| = N$, sparse regularization (that is, compressed sensing) if $M < N$ and least-squares fitting if $M > N$. We consider the latter. Interpolation requires good choices of nodes $\mathbf{y}_1, \dots, \mathbf{y}_N$ so as to maintain small Lebesgue constants, and it is unclear how to design such nodes for general irregular domains. Compressed sensing is an interesting option, however beyond the scope of this paper (see Section 9 for some further discussion). Least-squares fitting, on the other hand, is a popular tool for high-dimensional approximation on tensor-product domains [18, 25, 34, 37–39, 53], and has the twin benefits of being simple to both implement and analyze. Note that the least-squares approximation $f_{\mathcal{Y}, \Lambda}$ is given by

$$f_{\mathcal{Y}, \Lambda} = \operatorname{argmin}_{p \in P_{\Lambda}} \frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |f(\mathbf{y}) - p(\mathbf{y})|^2, \quad (2.1)$$

where $P_{\Lambda} = \operatorname{span}\{\phi_n : n \in \Lambda\}$ is the finite-dimensional approximation space. Equivalently, the coefficients $\mathbf{c} = (c_n)_{n \in \Lambda}$ of $f_{\mathcal{Y}, \Lambda}$ are the solution of the algebraic least-squares problem

$$\mathbf{c} = \operatorname{argmin}_{\mathbf{x} \in \mathbb{C}^N} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2, \quad (2.2)$$

where $\mathbf{A} = \left\{ \frac{1}{\sqrt{M}} \phi_n(\mathbf{y}) \right\}_{\mathbf{y} \in \mathcal{Y}, n \in \Lambda} \in \mathbb{C}^{M \times N}$ and $\mathbf{b} = \left\{ \frac{1}{\sqrt{M}} f(\mathbf{y}) \right\}_{\mathbf{y} \in \mathcal{Y}} \in \mathbb{C}^M$.

2. *How to choose the orthonormal basis $\{\psi_n\}_n$ and index set Λ .* Smooth functions are typically well approximated by polynomials, so we shall generally take $\{\psi_n\}_n$ to be an orthonormal tensor-product polynomial basis. Our main numerical examples consider tensor-product Legendre polynomials. We also highlight the possibility of nonpolynomial approximations, for example, using a cosine basis when Ω is compactly contained in $D = (-1, 1)^d$. Given the basis $\{\psi_n\}_n$, we consider several standard choices for Λ , including total degree and hyperbolic cross index sets, or more generally, so-called *lower* sets. These sets have been used quite extensively for multivariate polynomial approximations in tensor-product domains (see [3, 18–20, 24, 35, 37] and references therein).

3. *How to choose the sample points Υ .* Our primary concern in this regard lies with the *sampling efficiency* (or *sample complexity*) of the approximation: namely, how large M must be in relation to $N = |\Lambda|$ to ensure a good approximation. The problem of designing optimal sampling points for high-dimensional polynomial approximation remains open even in tensor-product domains (although we note in passing some recent quasioptimal constructions [25]). We shall therefore not attempt to solve it for irregular domains. Instead, we consider straightforward random samplings. Specifically, we draw $\mathbf{y}_1, \dots, \mathbf{y}_M$ independently according to a suitable probability measure on Ω (for example, the uniform measure whenever Ω is compact). We throughout assume that it is computationally feasible to draw samples from this measure. Although simple, this approach permits concrete sample complexity estimates for a large class of domains Ω which are quadratic (up to a log factor) in $N = |\Lambda|$ for any dimension d . Up to a domain-dependent constant which we determine, this sample complexity is the same as the corresponding result for compact tensor-product domains when the sample points are drawn from the uniform measure [18] (this scaling is essentially sharp: as discussed in [9], based on a result of [43], in one dimension if the sample points are deterministic and equispaced, then the least-squares approximation is ill-conditioned unless the number of sample points scales quadratically in the polynomial degree N).

2.2. Conditioning and stability. The approach outlined above is certainly simple, and it is tempting to think that it can achieve high accuracy. After all, the method computes a polynomial approximation in a domain, albeit an irregular one. Unfortunately, there is an issue. The matrix \mathbf{A} of system (2.2) is extremely ill-conditioned, even when $M \gg N$ (we estimate this ill-conditioning later in the paper for relevant examples). This is due to the fact that the set $\{\phi_n\}_n$ is not a basis for the space of square-integrable functions over Ω , but rather a frame. See Section 3.3 for the definition of a frame. Frames are typically redundant, meaning

that any function f has infinitely many expansions of the form $f = \sum_n c_n \phi_n$ with coefficients $\{c_n\}_n$ in ℓ^2 . When translated to the finite setting, this redundancy means that the truncated Gram matrix

$$\mathbf{G}_\Lambda = \{(\phi_m, \phi_n)_{L^2(\Omega, \mu)}\}_{m, n \in \Lambda}, \quad (2.3)$$

where μ is the measure defined in (3.2), is typically extremely poorly conditioned for large N [7]. Note that $\mathbb{E}(\mathbf{A}^* \mathbf{A}) = \mathbf{G}_\Lambda$ if the sample points y_i are drawn independently according to μ . Hence the least-squares matrix \mathbf{A} is expected to inherit similar ill-conditioning.

In the face of such ill-conditioning, one would usually expect it to be impossible to achieve high accuracy in floating point arithmetic. Fortunately, this expectation turns out to be incorrect. The frame property endows the problem with sufficient structure so that accurate, well-conditioned approximations can be computed via a simple regularization procedure. In this paper, we show that regularized least-squares solutions, computed via hard thresholding of the singular values of \mathbf{A} , yield well-conditioned approximations which converge rapidly down to the thresholding parameter ϵ . This parameter is typically set according to some desired target accuracy.

REMARK 2.1. We stress that the frame property is crucial in endowing the approximation with these properties, hence why we refer to this approach as *polynomial frame approximations*. Choosing $\{\phi_n\}_n$ to be the monomial basis also leads to an exceedingly ill-conditioned problem, but one where high accuracy may not be possible. The underlying reason for this is that the frame property guarantees existence of expansions $f = \sum_n c_n \phi_n$ for which the coefficients $\{c_n\}_n$ decay (accuracy) and have bounded ℓ^2 -norm (stability). See [6, 7] for further discussion.

2.3. Main results. We now summarize our main results.

Accuracy and conditioning. Our first result concerns the accuracy and condition number of the regularized least-squares approximation. As mentioned above, this approximation is constructed using a truncated Singular Value Decomposition (SVD) of the least-squares matrix \mathbf{A} with a threshold parameter $\epsilon > 0$. We write $f_{\mathcal{Y}, \Lambda, \epsilon}$ for this approximation and \mathbf{c}^ϵ for its coefficients in the system $\{\phi_n\}_{n \in \Lambda}$.

THEOREM 2.2 (Accuracy and conditioning). *There exists a constant $C_{\mathcal{Y}, \Lambda, \epsilon} > 0$ such that*

$$\|f - f_{\mathcal{Y}, \Lambda, \epsilon}\|_{L^2(\Omega, \mu)} \leq (1 + C_{\mathcal{Y}, \Lambda, \epsilon}) E_{\Lambda, \epsilon}(f), \quad (2.4)$$

where

$$E_{\Lambda, \epsilon}(f) = \inf\{\|f - p\|_{L^\infty(\Omega)} + \epsilon\|p\|_{L^2(D, \nu)} : p \in P_\Lambda\},$$

and μ is the measure given by (3.2). Moreover, the coefficients \mathbf{c}^ϵ of $f_{\mathcal{Y}, \Lambda, \epsilon}$ satisfy

$$\|\mathbf{c}^\epsilon\|_2 = \|f_{\mathcal{Y}, \Lambda, \epsilon}\|_{L^2(D, \nu)} \leq \frac{E_{\Lambda, \epsilon}(f)}{\epsilon}, \quad (2.5)$$

and the absolute (ℓ^2, L^2) -condition number of the reconstruction operator $\mathcal{L}_{\mathcal{Y}, \Lambda, \epsilon} : \mathbb{C}^M \rightarrow P_\Lambda$, $\mathbf{b} \mapsto f_{\mathcal{Y}, \Lambda, \epsilon}$, where \mathbf{b} is as in (2.2), is at most $C_{\mathcal{Y}, \Lambda, \epsilon}$.

See Section 4. Several remarks are in order. First, bound (2.4) separates the accuracy of the regularized least-squares approximation into an approximation error term $E_{\Lambda, \epsilon}(f)$ depending only on ϵ and the space P_Λ and independent of the samples \mathcal{Y} , and a constant $C_{\mathcal{Y}, \Lambda, \epsilon}$ depending on ϵ , \mathcal{Y} and P_Λ . In other words, $E_{\Lambda, \epsilon}(f)$ determines the rate of approximation, whereas $C_{\mathcal{Y}, \Lambda, \epsilon}$ (more specifically, the requirement that $C_{\mathcal{Y}, \Lambda, \epsilon} \lesssim 1$) determines the sample complexity.

Second, note that $E_{\Lambda, \epsilon}(f)$ depends on how well f can be approximated in Ω by polynomials $p \in P_\Lambda$ that do not grow too large on D . The latter requirement—which stems from the regularization carried out—is an expression of stability, since a polynomial growing large on D would necessarily have large coefficients. Our main estimates for $E_{\Lambda, \epsilon}(f)$, given below, are derived by constructing polynomials which approximate f at specified rates in Ω (depending on the smoothness of f), and which remain bounded on D .

Third, note that (2.5) ensures the stored values—namely, the coefficients \mathbf{c}^ϵ —cannot be too large in magnitude. This would otherwise result in ill-conditioning of the evaluation map $\mathbf{c}^\epsilon \mapsto f_{\mathcal{Y}, \Lambda, \epsilon}(\mathbf{x})$. While $\|\mathbf{c}^\epsilon\|_2$ may be of magnitude roughly $1/\epsilon$ initially, once the approximation error $E_{\Lambda, \epsilon}(f)$ reaches close to the target accuracy ϵ , we have $\|\mathbf{c}^\epsilon\|_2 \lesssim 1$.

Rate of approximation. In Section 5, we analyze $E_{\Lambda, \epsilon}(f)$ for the main example considered in this paper, Legendre polynomials on $D = (-1, 1)^d$. We consider two standard choices of index sets Λ : the total degree index set $\Lambda = \Lambda_n^{\text{TD}}$ defined in (3.5) and the hyperbolic cross index set $\Lambda = \Lambda_n^{\text{HC}}$ defined in (3.6). The former is suitable for low-dimensional problems, but quickly becomes too large as d increases. The cardinality of the latter on the other hand scales much more mildly with d .

Our main results are split into two cases:

(i) *f smooth in Ω only.* In the first case, f is smooth in Ω but may be nonsmooth, or even undefined in $D \setminus \Omega$. If Ω is a Lipschitz domain and $f \in H^m(\Omega, \mu)$, where $H^m(\Omega, \mu)$ is the classical Sobolev space of order m (see (5.1)), then we

show that

$$E_{\Lambda, \epsilon}(f) \leq \begin{cases} c_{m,d,\Omega}(n^{d-m} + \epsilon) \|f\|_{H^m(\Omega, \mu)} & \Lambda = \Lambda_n^{\text{TD}}, \\ c_{m,d,\Omega}(n^{\frac{d-m}{d}} + \epsilon) \|f\|_{H^m(\Omega, \mu)} & \Lambda = \Lambda_n^{\text{HC}}, \end{cases} \tag{2.6}$$

where $c_{m,d,\Omega} > 0$ is a constant depending on m, d and Ω but independent of f . See Theorem 5.1 (we note in passing that the factor $d - m$ can be improved slightly to $\theta(d) - m$, where $\theta(d)$ is a particular constant satisfying $\theta(d) \leq d$). This result asserts convergence at an algebraic rate depending on the smoothness of f in Ω only. However, it also exhibits the familiar curse of dimensionality. In the case of the total degree index set Λ^{TD} , the cardinality $N = |\Lambda^{\text{TD}}| \asymp n^d$ as $n \rightarrow \infty$, and therefore

$$n^{d-m} \asymp N^{\frac{d-m}{d}}, \quad n \rightarrow \infty,$$

whereas for the hyperbolic cross (wherein $N = |\Lambda_n^{\text{HC}}| \asymp n(\log(n))^{d-1}$), one has

$$n^{\frac{d-m}{d}} \asymp N^{\frac{d-m}{d}} (\log(N))^{\frac{(m-d)(d-1)}{d}}, \quad n \rightarrow \infty.$$

(ii) *f smooth in D.* In high-dimensional approximation, a standard way to overcome the d -dependence in results such as (2.6) is to assume certain anisotropic smoothness. As we discuss in Section 9, it is currently unknown how to do this within the setting of case (i). However, when f has appropriate regularity over the whole of D —or equivalently, f is the restriction to Ω of some appropriately regular function defined on D —then we have the following result. If $f \in H_{\text{mix}}^m(D, \nu)$, where $H_{\text{mix}}^m(D, \nu)$ is the Sobolev space of dominating mixed smoothness on D (see (5.2)), then

$$E_{\Lambda, \epsilon}(f) \leq \begin{cases} c_{m,d} \|f\|_{H_{\text{mix}}^m(D, \nu)} n^{1-m} + \epsilon \|f\|_{L^2(D, \nu)} & \Lambda = \Lambda_n^{\text{TD}}, \\ c_{m,d} \|f\|_{H_{\text{mix}}^m(D, \nu)} n^{1-m} (\log(n))^{\frac{d-1}{2}} + \epsilon \|f\|_{L^2(D, \nu)} & \Lambda = \Lambda_n^{\text{HC}}, \end{cases} \tag{2.7}$$

where $c_{m,d} > 0$ is a constant depending on m and d but independent of Ω and f . See Theorem 5.3. Observe that

$$n^{1-m} \asymp N^{\frac{1-m}{d}}, \quad N = |\Lambda_n^{\text{TD}}|,$$

whereas

$$n^{1-m} (\log(n))^{\frac{d-1}{2}} \asymp N^{1-m} (\log(N))^{(d-1)(m-1/2)}, \quad N = |\Lambda_n^{\text{HC}}|.$$

Hence, up to the log factor, the hyperbolic cross index set Λ_n^{HC} achieves an algebraic rate of convergence that is independent of the dimension d , and therefore suitable for higher-dimensional problems. Our numerical results in Section 8 show computations using this index set for dimensions up to $d = 15$.

Let us make several remarks. First, we note that case (ii) requires absolutely no conditions on the domain Ω , besides being measurable. In particular, the domain can be extremely rough, as long as f is smooth over the whole extended domain D . In Section 8, we show some numerical results of this type. Second, (2.7) behaves like n^{1-m} , not n^{-m} as might be expected. The additional power of n stems from the presence of the $L^\infty(\Omega)$ -norm in $E_{\Lambda,\epsilon}(f)$. This factor can be improved whenever Ω is compactly contained in D , in which case one obtains a factor of the form $n^{1/2-m}$ (see Theorems 5.1 and 5.3). Third, when the sample points are drawn randomly and independently (as they are in this paper), it is possible to prove an estimate in expectation for the squared L^2 -error of a slightly modified least-squares estimator (see Section 7) involving the L^2 -norm approximation error

$$\tilde{E}_{\Lambda,\epsilon}(f) = \inf\{\|f - p\|_{L^2(\Omega,\mu)} + \epsilon\|p\|_{L^2(D,\nu)} : p \in P_\Lambda\}.$$

See Theorem 7.1. Analogous to (2.6) and (2.7), this quantity admits the following estimates. First, if Ω is Lipschitz and $f \in H^m(\Omega, \mu)$, then

$$\tilde{E}_{\Lambda,\epsilon}(f) \leq \begin{cases} c_{m,d,\Omega}(n^{-m} + \epsilon)\|f\|_{H^m(\Omega,\mu)} & \Lambda = \Lambda_n^{\text{TD}} \\ c_{m,d,\Omega}(n^{-\frac{m}{d}} + \epsilon)\|f\|_{H^m(\Omega,\mu)} & \Lambda = \Lambda_n^{\text{HC}}. \end{cases} \quad (2.8)$$

Conversely, if $f \in H_{\text{mix}}^m(D, \nu)$, then

$$\tilde{E}_{\Lambda,\epsilon}(f) \leq c_{m,d}\|f\|_{H_{\text{mix}}^m(D,\nu)}n^{-m} + \epsilon\|f\|_{L^2(D,\nu)}, \quad \Lambda = \Lambda_n^{\text{TD}} \text{ or } \Lambda = \Lambda_n^{\text{HC}}. \quad (2.9)$$

See Theorems 7.2 and 7.3. As with $E_{\Lambda,\epsilon}(f)$ above, this latter result for the hyperbolic cross index set shows how polynomial frame approximation can mitigate the curse of dimensionality.

Sample complexity. Our final result concerns efficiency, that is, sample complexity, of the approximation. In view of Theorem 2.2, this corresponds to determining how large M must be in order for the condition $C_{\mathcal{R},\Lambda,\epsilon} \lesssim 1$ to hold. Our main contribution is for the following class of domains Ω .

DEFINITION 2.3 (λ -rectangle property). A compact domain Ω has the λ -rectangle property for some $0 < \lambda < 1$ if it can be written as a (possibly overlapping and uncountable) union

$$\Omega = \bigcup_{R \in \mathcal{R}} R,$$

of hyperrectangles R satisfying

$$\inf_{R \in \mathcal{R}} \text{Vol}(R) = \lambda \text{Vol}(\Omega).$$

Note that many domains of practical interest have this property. However, there are notable exceptions, including simplices and balls. See Section 6.3 for further discussion.

As we show in Section 6.2, when the samples y_m are chosen randomly and independently according to the uniform measure on Ω , the sample complexity of the approximation can in general be related to the constant of the $(L^2(\Omega, \mu), L^\infty(\Omega))$ -Nikolskii inequality for the space P_Λ . We use the λ -rectangle property to get concrete estimates for this constant, culminating in the following result.

THEOREM 2.4 (Sample complexity). *Suppose that $\Omega \subseteq (-1, 1)^d$ has the λ -rectangle property and let P_Λ be constructed from the tensor Legendre polynomial basis on $(-1, 1)^d$, where $\Lambda \subset \mathbb{N}_0^d$ is any lower set (see Definition 3.1) of cardinality $|\Lambda| = N$. Let $0 < \delta, \gamma < 1$ and y_1, \dots, y_M be independent and randomly drawn according to the uniform probability measure on Ω . Then*

$$C_{\gamma, \Lambda, \epsilon} \leq \frac{1}{\sqrt{1 - \delta}}, \quad \forall \epsilon > 0,$$

with probability at least $1 - \gamma$, provided

$$M \geq N^2 \lambda^{-1} ((1 - \delta) \log(1 - \delta) + \delta)^{-1} \log(N/\gamma).$$

See Corollary 6.6. This result establishes log-quadratic scaling of the number of samples with the dimension of the polynomial space, extending a well-known result for tensor-product domains to a large class of irregular domains. Note that this result holds for all lower sets, and in particular, the total degree and hyperbolic cross index sets discussed above.

2.4. Related work. The idea of approximating a function on an irregular domain by using an orthogonal basis on a bounding tensor-product domain is well established within the context of embedded or fictitious domain methods in numerical PDEs [41] (see also [12]). The so-called *Fourier extensions* or *Fourier continuations* were studied in detail in [13, 15]. Applications to surface parametrization and numerical PDEs in complex geometries were considered in [15] and [11, 14, 33], respectively. Our work can be considered an extension of [8] from the univariate to the multivariate setting, although we use algebraic as opposed to trigonometric polynomials since these are more common in applications such as UQ. Our work also extends recent research on computing polynomial approximations of functions defined on high-dimensional tensor-product domains. This approach has received substantial interest recently, due to its applications in, notably, UQ. See [4, 18, 19, 22, 23, 27, 39, 52] and

references therein. A consequence of this paper is that an irregular domain (either known or unknown) often presents no barrier to polynomial approximation of high-dimensional functions. As noted, polynomial approximations are frequently used in practical UQ studies even when the domain is nontensorial (see [44, 45] and references therein). Our work therefore provides a theoretical basis for these approaches. Finally, we note that polynomial frame approximation is just one example of so-called *numerical frame approximation*. For a broader perspective on the uses of frames in numerical analysis and approximation, see [6, 7].

3. Polynomial frame approximations

3.1. Notation. We first require some further notation. Throughout this paper, $D \subseteq \mathbb{R}^d$ will be a domain with a probability measure ν . Typically, D will be of tensor-product type, that is,

$$D = [a_1, b_1] \otimes \cdots \otimes [a_d, b_d] \subseteq \mathbb{R}^d, \quad (3.1)$$

where $-\infty \leq a_k < b_k \leq \infty$ and $\nu = \nu^{(1)} \otimes \cdots \otimes \nu^{(d)}$ will be a tensor product of one-dimensional probability measures. We write $L^2(D, \nu)$ for the space of square-integrable functions on D .

The d -dimensional variable is denoted by $\mathbf{y} = (y_1, \dots, y_d) \in \mathbb{R}^d$. Given D , we let $\Omega \subseteq D$ be a domain and define the probability measure μ by

$$d\mu(\mathbf{y}) = \frac{\mathbb{I}_\Omega(\mathbf{y})}{\nu_\Omega} d\nu(\mathbf{y}), \quad \nu_\Omega = \int_\Omega d\nu, \quad (3.2)$$

where \mathbb{I}_Ω is the indicator function of Ω . We write $L^2(\Omega, \mu)$ for the space of square-integrable functions on Ω with respect to μ .

Throughout, $\mathbf{n} = (n_1, \dots, n_d) \in \mathbb{N}_0^d$ denotes a multi-index. Let $I \subseteq \mathbb{N}_0^d$ be a countable set of multi-indices and $\{\psi_{\mathbf{n}}\}_{\mathbf{n} \in I}$ be an orthonormal basis of $L^2(D, \nu)$. If D is of the form (3.1), then this basis will usually be of tensor-product type, that is,

$$\psi_{\mathbf{n}}(\mathbf{y}) = \prod_{k=1}^d \psi_{n_k}^{(k)}(y_k),$$

where $\{\psi_{n_k}^{(k)}\}$ is an orthonormal basis of $L^2((a_k, b_k), \nu^{(k)})$. Given $\{\psi_{\mathbf{n}}\}_{\mathbf{n} \in I}$, we let

$$\phi_{\mathbf{n}} = \psi_{\mathbf{n}}|_{\Omega}, \quad \mathbf{n} \in I, \quad (3.3)$$

be the corresponding functions defined on Ω .

3.2. Multi-index sets. Our interest lies in computing finite approximations in system (3.3). To this end, let $\Lambda \subset I$ be a finite multi-index set and define

$$P_\Lambda = \text{span}\{\phi_n : \mathbf{n} \in \Lambda\} \subset L^2(\Omega, \mu),$$

as the finite-dimensional space within which we seek an approximation to f . We consider the following three standard choices of multi-index sets: the *tensor-product* set

$$\Lambda = \Lambda_n^{\text{TP}} = \{\mathbf{n} \in \mathbb{N}_0^d : |\mathbf{n}|_\infty \leq n\}, \quad (3.4)$$

where $n \in \mathbb{N}_0$ and $|\mathbf{n}|_\infty = \max_{k=1, \dots, d} |n_k|$; the *total degree* set

$$\Lambda = \Lambda_n^{\text{TD}} = \{\mathbf{n} \in \mathbb{N}_0^d : |\mathbf{n}|_1 \leq n\}, \quad (3.5)$$

where $|\mathbf{n}|_1 = |n_1| + \dots + |n_d|$; and the (isotropic) *hyperbolic cross* set

$$\Lambda = \Lambda_n^{\text{HC}} = \{\mathbf{n} \in \mathbb{N}_0^d : |\mathbf{n}|_{\text{hc}} \leq n + 1\}, \quad |\mathbf{n}|_{\text{hc}} = \prod_{k=1}^d (n_k + 1). \quad (3.6)$$

Note that the cardinality $N = |\Lambda_n^{\text{TP}}| = (n + 1)^d$ usually grows too quickly with n in high dimensions to be practical. The total degree set, with cardinality

$$N = |\Lambda_n^{\text{TD}}| = \binom{n + d}{d},$$

mitigates this issue to some extent, but still typically grows too rapidly for moderate- to high-dimensional problems. Hyperbolic cross index sets are a practical alternative in this case. An exact formula for the cardinality of the hyperbolic cross Λ_n^{HC} in terms of n and d is not known, but there are a variety of upper bounds, including

$$|\Lambda_n^{\text{HC}}| \leq \lfloor (n + 1)(1 + \log(n + 1))^{d-1} \rfloor.$$

See, for example, [34, Proposition A.1].

The above three multi-index sets are all examples of so-called *lower* sets (also known as *downward closed* or *monotone* sets—see, for example, [23, 28]).

DEFINITION 3.1. A multi-index set $\Lambda \subseteq \mathbb{N}_0^d$ is lower if whenever

$$\mathbf{n} = (n_1, \dots, n_d) \in \Lambda \quad \text{and} \quad \mathbf{n}' = (n'_1, \dots, n'_d)$$

satisfies $n'_k \leq n_k$ for all k , then $\mathbf{n}' \in \Lambda$.

In our main results regarding efficiency, we establish sample complexity estimates which are valid for arbitrary lower sets. While we shall not do it in this paper, such generality allows for the possibility of considering other multi-index sets, for example, anisotropic hyperbolic cross index sets, which may be defined by *a priori* or *a posteriori* estimates, or computed adaptively.

3.3. Polynomial frames. We first recall the definition of a frame (see, for example, [21]).

DEFINITION 3.2. A countable set $\{\phi_n\}_{n \in I}$ of a Hilbert space H is a frame if there exist constants $0 < A \leq B < \infty$ such that

$$A\|f\|^2 \leq \sum_{n \in I} |\langle f, \phi_n \rangle|^2 \leq B\|f\|^2, \quad \forall f \in H, \tag{3.7}$$

where $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ are the inner product and norm, respectively, on H .

Let $\{\phi_n\}_{n \in I}$ be the system constructed in (3.3). It is straightforward to see that this is a frame for $H = L^2(\Omega, \mu)$. Indeed, let $f \in L^2(\Omega, \mu)$ and \tilde{f} be its extension by zero to D . Then by Parseval’s relation for the orthonormal basis $\{\psi_n\}_{n \in I}$,

$$\sum_{n \in I} |\langle f, \phi_n \rangle_{L^2(\Omega, \mu)}|^2 = \sum_{n \in I} \frac{|\langle \tilde{f}, \psi_n \rangle_{L^2(D, \nu)}|^2}{(v_\Omega)^2} = \frac{\|\tilde{f}\|_{L^2(D, \nu)}^2}{(v_\Omega)^2} = \frac{\|f\|_{L^2(\Omega, \mu)}^2}{v_\Omega},$$

where v_Ω is given by (3.2). Hence (3.7) holds with $A = B = 1/v_\Omega$, making this system a frame. Frames such as this for which $A = B$ are known as *tight* frames.

A general property of frames is their redundancy: any $f \in H$ can have infinitely many expansions $f = \sum_{n \in I} c_n \phi_n$ with coefficients $\{c_n\}_{n \in I} \in \ell^2(I)$. It is easy to see how redundancy occurs in the polynomial frame. Indeed, let \tilde{f} be any extension of f to $L^2(D, \nu)$ and define

$$c_n = \langle \tilde{f}, \psi_n \rangle_{L^2(D, \nu)},$$

as the coefficients of \tilde{f} in the orthonormal basis $\{\psi_n\}_{n \in I}$. Then $\{c_n\}_{n \in I} \in \ell^2(I)$ and

$$\sum_{n \in I} c_n \phi_n = \sum_{n \in I} c_n \psi_n \Big|_\Omega = \tilde{f}|_\Omega = f.$$

Since there are infinitely many extensions of f to $L^2(D, \nu)$, each with distinct coefficients $\{c_n\}_{n \in I}$, it follows that there are infinitely many representations of f in the frame $\{\phi_n\}_{n \in I}$.

3.4. Least-squares polynomial frame approximations. Let $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_M\} \subset \Omega$ be a set of M distinct points (for the moment, we choose not to specify their distribution) and Λ be a finite set of multi-indices of size $|\Lambda| = N$, where $N \leq M$. Consider the approximation to f in the space P_Λ by discrete least-squares fitting:

$$f_{\mathcal{Y},\Lambda} = \operatorname{argmin}_{p \in P_\Lambda} \frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |f(\mathbf{y}) - p(\mathbf{y})|^2.$$

If $f_{\mathcal{Y},\Lambda}$ is expressed as

$$f_{\mathcal{Y},\Lambda} = \sum_{\mathbf{n} \in \Lambda} c_{\mathbf{n}} \phi_{\mathbf{n}},$$

then this is equivalent to the algebraic least-squares problem

$$\mathbf{c} = (c_{\mathbf{n}})_{\mathbf{n} \in \Lambda} = \operatorname{argmin}_{\mathbf{x} \in \mathbb{C}^N} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2, \quad (3.8)$$

where

$$\mathbf{A} = \mathbf{A}_{\mathcal{Y},\Lambda} = \left(\frac{1}{\sqrt{M}} \phi_{\mathbf{n}}(\mathbf{y}) \right)_{\mathbf{y} \in \mathcal{Y}, \mathbf{n} \in \Lambda} \in \mathbb{C}^{M \times N}, \quad \mathbf{b} = \mathbf{b}_{\mathcal{Y}} = \left(\frac{1}{\sqrt{M}} f(\mathbf{y}) \right)_{\mathbf{y} \in \mathcal{Y}} \in \mathbb{C}^M.$$

Note that \mathbf{A} may fail to be full rank—for example, if the points \mathcal{Y} are chosen poorly or the functions $\phi_{\mathbf{n}}$, $\mathbf{n} \in \Lambda$, are linearly dependent—in which case (3.8) does not have a unique solution. However, even if it is full rank, as mentioned in Section 2.2 and shown explicitly in Section 4.1, \mathbf{A} is typically severely ill-conditioned for large N . Hence it is necessary to regularize (3.8). We shall do this via truncated singular value decompositions (that is, spectral filtering). (Related strategies such as Tikhonov regularization could be used instead, with some changes to the ensuing presentation.)

To this end, suppose that \mathbf{A} has singular values $\{\sigma_{\mathbf{n}}\}_{\mathbf{n} \in \Lambda}$ and singular value decomposition $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, where $\mathbf{U} \in \mathbb{C}^{M \times M}$, $\mathbf{\Sigma} \in \mathbb{R}^{M \times N}$ and $\mathbf{V} \in \mathbb{C}^{N \times N}$. Define

$$\mathbf{A}_\epsilon = \mathbf{A}_{\mathcal{Y},\Lambda,\epsilon} = \mathbf{U}\mathbf{\Sigma}_\epsilon\mathbf{V}^*,$$

where the diagonal matrix $\mathbf{\Sigma}_\epsilon$ has n th entry $\sigma_{\mathbf{n}}$ if $\sigma_{\mathbf{n}} > \epsilon$ and zero otherwise. Then the truncated SVD least-squares approximation is defined as

$$f_{\mathcal{Y},\Lambda,\epsilon} = \sum_{\mathbf{n} \in \Lambda} (\mathbf{c}_\epsilon)_{\mathbf{n}} \phi_{\mathbf{n}}, \quad (3.9)$$

where its coefficients \mathbf{c}_ϵ are given by

$$\mathbf{c}_\epsilon = (\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{b}_{\mathcal{Y}} = \mathbf{V}(\mathbf{\Sigma}_\epsilon)^\dagger \mathbf{U}^* \mathbf{b}_{\mathcal{Y}}.$$

Here \dagger denotes the pseudoinverse. We consider this approximation from now on. Note that the regularization parameter ϵ is usually set in relation to some desired target accuracy (see Section 8).

3.5. Main example. We end this section by introducing our main example. This is the case where Ω is bounded and, without loss of generality, contained in $D = (-1, 1)^d$, and where $\{\psi_n\}_{n \in \mathbb{N}_0^d}$ is the tensor Legendre polynomial basis on D corresponding to the uniform probability measure $d\nu(\mathbf{y}) = 2^{-d} d\mathbf{y}$. When normalized with respect to the uniform probability measure on D , this basis is defined by

$$\psi_n(\mathbf{y}) = \prod_{k=1}^d \sqrt{2n_k + 1} P_{n_k}(y_k),$$

where P_n is the n th classical Legendre polynomial (see [Appendix A](#)). For the truncated index set, we let $\Lambda = \Lambda_n^{\text{TD}}$ or $\Lambda = \Lambda_n^{\text{HC}}$ be either the total degree (3.5) or hyperbolic cross (3.6) index set with index n . We also assume that the sampling points $\mathbf{y}_1, \dots, \mathbf{y}_M$ are drawn independently according to the measure μ , which in this case is the uniform probability measure on Ω :

$$d\mu(\mathbf{y}) = \frac{1}{\text{Vol}(\Omega)} d\mathbf{y}. \quad (3.10)$$

While this approach leads to concrete, d -independent sample complexity estimates for many domains, we do not claim that it is an optimal sampling procedure. See [Sections 8 and 9](#) for further discussion.

REMARK 3.3. As mentioned, we assume that it is computationally feasible to draw samples from μ . For the numerical examples shown later, this is achieved by rejection sampling. Depending on the domain, however, and especially in high dimensions, this may be a substantial challenge.

4. Accuracy and conditioning

We now investigate the accuracy and conditioning of approximation (3.9). In [Section 4.1](#), we show that least-squares matrix \mathbf{A} is ill-conditioned for large N , thus explaining why regularization is needed. Next, in [Section 4.2](#), we introduce the key constant $C_{\mathcal{L}, \Lambda, \epsilon}$, and in [Section 4.3](#), we give the main result of this section. Note that the approach in [Sections 4.2 and 4.3](#) follows that of [\[6\]](#) (which applies to general frames).

4.1. Ill-conditioning of the matrix \mathbf{A} . Unless the frame happens to be a Riesz basis (which is not the case in our setting), frame approximations always lead to ill-conditioned least-squares matrices for sufficiently large truncation space Λ [[7](#), Lemma 5]. In the case of the polynomial frame, this is related to the Remez inequality for the polynomial space P_Λ over Ω and D . To see this, observe that

the minimal and maximal singular values of \mathbf{A} are

$$\sigma_{\min}(\mathbf{A}) = \inf_{\substack{p \in P_\Lambda \\ p \neq 0}} \left\{ \frac{\sqrt{\frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |p(\mathbf{y})|^2}}{\|p\|_{L^2(D, \nu)}} \right\}, \quad \sigma_{\max}(\mathbf{A}) = \sup_{\substack{p \in P_\Lambda \\ p \neq 0}} \left\{ \frac{\sqrt{\frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |p(\mathbf{y})|^2}}{\|p\|_{L^2(D, \nu)}} \right\}.$$

For simplicity, assume that the constant function is contained in P_Λ . This will hold in all examples considered later. Letting $p(\mathbf{y}) = 1$, we get $\sigma_{\max}(\mathbf{A}) \geq 1$. Conversely, note that $\frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |p(\mathbf{y})|^2 \leq \|p\|_{L^\infty(\Omega)}^2$ and let $N(P_\Lambda, D, \nu) > 0$ be the optimal constant such that

$$\|p\|_{L^\infty(D)} \leq N(P_\Lambda, D, \nu) \|p\|_{L^2(D, \nu)}, \quad \forall p \in P_\Lambda.$$

We refer to this as an $(L^2(D, \nu), L^\infty(D))$ -Nikolskii inequality for the space P_Λ . Inequalities such as these will be discussed further in Section 6, since they are pivotal in estimating the sample complexity of the approximation. This gives

$$\frac{1}{\sigma_{\min}(\mathbf{A})} \geq (N(P_\Lambda, D, \nu))^{-1} \sup \left\{ \frac{\|p\|_{L^\infty(D)}}{\|p\|_{L^\infty(\Omega)}} : p \in P_\Lambda, p \neq 0 \right\},$$

and therefore

$$\text{cond}(\mathbf{A}) \geq \frac{R(P_\Lambda, \Omega, D)}{N(P_\Lambda, D, \nu)}, \tag{4.1}$$

where $R(P_\Lambda, \Omega, D)$ is the constant in the Remez inequality for the domains Ω and D :

$$\|p\|_{L^\infty(D)} \leq R(P_\Lambda, \Omega, D) \|p\|_{L^\infty(\Omega)}, \quad p \in P_\Lambda.$$

Note that bound (4.1) is completely deterministic, and independent of the samples \mathcal{Y} .

Typically, the right-hand side of (4.1) will grow rapidly with N . To see why, note first that the Nikolskii constant is usually at most algebraic in $N = |\Lambda|$. In particular, if ν is the uniform measure on D and Λ is a lower set, then $N(P_\Lambda, D, \nu) \leq N^2$ [35, Theorem 6] (see also the proof of Theorem 6.5). Similar bounds are found in [35] for other ultraspherical and Jacobi measures. Conversely, the constant $R(P_\Lambda, \Omega, D)$ is typically exponentially large in N . Its precise behavior depends on the domain Ω and the index set Λ , and for the sake of brevity, we will not consider this issue in depth. However, we note in passing that in the one-dimensional case, for example, if $\Lambda = \{0, \dots, N - 1\}$ and $D = (-1, 1)$, then

$$R(P_\Lambda, \Omega, (-1, 1)) \leq T_{N-1}(4/|\Omega| - 1),$$

where T_{N-1} is the $(N - 1)$ th Chebyshev polynomial and $|\Omega|$ denotes the Lebesgue measure of Ω . Moreover, equality holds if $\Omega = [-1, -1 + |\Omega|]$ in which case

one has the exponential growth

$$R(P_A, \Omega, (-1, 1)) \geq \frac{1}{2} \left(\frac{4}{|\Omega|} - 1 \right)^{N-1}.$$

We refer to [29] for further information, including results in higher dimensions, as well as to [47] for results on multivariate Remez inequalities for hyperbolic cross index sets.

4.2. Key constants. For convenience, we now define the following operator

$$\mathcal{T}_A : \mathbb{C}^N \rightarrow P_A, \quad \mathbf{c} = \{c_n\}_{n \in \Lambda} \mapsto \sum_{n \in \Lambda} c_n \phi_n.$$

This is commonly referred to as the *synthesis* operator in frame theory. We now let

$$C_{\mathcal{Y}, \Lambda, \epsilon} = \max\{C'_{\mathcal{Y}, \Lambda, \epsilon}, C''_{\mathcal{Y}, \Lambda, \epsilon}\}, \tag{4.2}$$

where

$$\begin{aligned} C'_{\mathcal{Y}, \Lambda, \epsilon} &= \max_{\substack{\mathbf{b} \in \mathbb{C}^M \\ \|\mathbf{b}\|_2=1}} \|\mathcal{T}_A(\mathbf{A}_{\mathcal{Y}, \Lambda, \epsilon})^\dagger \mathbf{b}\|_{L^2(\Omega, \mu)}, \\ C''_{\mathcal{Y}, \Lambda, \epsilon} &= \epsilon^{-1} \max_{\substack{\mathbf{d} \in \mathbb{C}^N \\ \|\mathbf{d}\|_2=1}} \|\mathcal{T}_A \mathbf{d} - \mathcal{T}_A(\mathbf{A}_{\mathcal{Y}, \Lambda, \epsilon})^\dagger \mathbf{A}_{\mathcal{Y}, \Lambda} \mathbf{d}\|_{L^2(\Omega, \mu)}. \end{aligned} \tag{4.3}$$

It is useful to interpret these constants. First, define the *reconstruction* operator

$$\mathcal{L}_{\mathcal{Y}, \Lambda, \epsilon} : \mathbb{C}^M \rightarrow P_A; \quad \mathbf{b} \mapsto \mathcal{T}_A(\mathbf{A}_{\mathcal{Y}, \Lambda, \epsilon})^\dagger \mathbf{b}. \tag{4.4}$$

This operator takes a vector of samples $\mathbf{b} \in \mathbb{C}^M$ to its truncated SVD approximation in P_A . In particular, if

$$\mathcal{S}_{\mathcal{Y}} : L^\infty(\Omega) \rightarrow \mathbb{C}^M; \quad f \mapsto \left\{ \frac{1}{\sqrt{M}} f(\mathbf{y}) \right\}_{\mathbf{y} \in \mathcal{Y}} \tag{4.5}$$

is the operator taking a function f to its samples, then

$$f_{\mathcal{Y}, \Lambda, \epsilon} = \mathcal{L}_{\mathcal{Y}, \Lambda, \epsilon} \mathcal{S}_{\mathcal{Y}} f. \tag{4.6}$$

The constant $C'_{\mathcal{Y}, \Lambda, \epsilon}$ is precisely the operator norm—or equivalently, since it is a linear operator, the absolute condition number—of $\mathcal{L}_{\mathcal{Y}, \Lambda, \epsilon}$ with respect to the ℓ^2 - and $L^2(\Omega, \mu)$ -norms:

$$C'_{\mathcal{Y}, \Lambda, \epsilon} = \max_{\substack{\mathbf{b} \in \mathbb{C}^M \\ \|\mathbf{b}\|_2=1}} \|\mathcal{L}_{\mathcal{Y}, \Lambda, \epsilon} \mathbf{b}\|_{L^2(\Omega, \mu)}.$$

In other words, boundedness of $C_{\mathcal{Y}, \Lambda, \epsilon}$ implies robustness of the approximation to perturbations in the data (for example, noise). On the other hand, $C''_{\mathcal{Y}, \Lambda, \epsilon}$ also has

the equivalent definition

$$C''_{\mathcal{Y},\Lambda,\epsilon} = \epsilon^{-1} \sup\{\|p - p_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(\Omega,\mu)} : p \in P_\Lambda, \|p\|_{L^2(D,v)} = 1\}.$$

In particular, $C''_{\mathcal{Y},\Lambda,0} = 0$ since the unregularized mapping $f \mapsto f_{\mathcal{Y},\Lambda,0}$ is a projection onto P_Λ . When $\epsilon > 0$, this constant measures how close the map $f \mapsto f_{\mathcal{Y},\Lambda,\epsilon}$ is to being a projection onto P_Λ .

4.3. Main result on accuracy and conditioning

THEOREM 4.1. *Let $f \in L^\infty(\Omega, \mu)$ and suppose that $f_{\mathcal{Y},\Lambda,\epsilon}$ is the truncated SVD least-squares approximation. Then*

$$\begin{aligned} \|f - f_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(\Omega,\mu)} &\leq (1 + C'_{\mathcal{Y},\Lambda,\epsilon})\|f - p\|_{L^\infty(\Omega)} + \epsilon C''_{\mathcal{Y},\Lambda,\epsilon}\|p\|_{L^2(D,v)} \\ &\leq (1 + C_{\mathcal{Y},\Lambda,\epsilon})E_{\Lambda,\epsilon}(f), \end{aligned}$$

where $C'_{\mathcal{Y},\Lambda,\epsilon}$, $C''_{\mathcal{Y},\Lambda,\epsilon}$ and $C_{\mathcal{Y},\Lambda,\epsilon}$ are as in (4.3) and (4.2), respectively, and

$$E_{\Lambda,\epsilon}(f) = \inf\{\|f - p\|_{L^\infty(\Omega)} + \epsilon\|p\|_{L^2(D,v)} : p \in P_\Lambda\}. \quad (4.7)$$

Moreover, the coefficients \mathbf{c}^ϵ of $f_{\mathcal{Y},\Lambda,\epsilon}$ satisfy

$$\|\mathbf{c}^\epsilon\|_2 = \|f_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)} \leq \frac{E_{\Lambda,\epsilon}(f)}{\epsilon}.$$

Proof. Let $p = \mathcal{T}_\Lambda \mathbf{c}$ for some $\mathbf{c} \in \mathbb{C}^N$. Then, recalling the definitions of the constants $C'_{\mathcal{Y},\Lambda,\epsilon}$ and $C''_{\mathcal{Y},\Lambda,\epsilon}$, we have

$$\begin{aligned} \|f - f_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(\Omega,\mu)} &\leq \|f - p\|_{L^2(\Omega,\mu)} + \|p_{\mathcal{Y},\Lambda,\epsilon} - f_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(\Omega,\mu)} + \|p - p_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(\Omega,\mu)} \\ &\leq \|f - p\|_{L^2(\Omega,\mu)} + C'_{\mathcal{Y},\Lambda,\epsilon}\|\mathcal{S}_{\mathcal{Y}}(f - p)\|_2 + \epsilon C''_{\mathcal{Y},\Lambda,\epsilon}\|p\|_{L^2(D,v)} \\ &\leq (1 + C'_{\mathcal{Y},\Lambda,\epsilon})\|f - p\|_{L^\infty(\Omega)} + \epsilon C''_{\mathcal{Y},\Lambda,\epsilon}\|p\|_{L^2(D,v)}, \end{aligned}$$

which gives the first result. Note that in the third step, we use (4.5) to deduce that $\|\mathcal{S}_{\mathcal{Y}}(f - p)\|_2 \leq \|f - p\|_{L^\infty(\Omega)}$ and the fact that μ is a probability measure, which implies that $\|f - p\|_{L^2(\Omega,\mu)} \leq \|f - p\|_{L^\infty(\Omega)}$. For the second result, we first use Parseval's identity to give $\|\mathbf{c}^\epsilon\|_2 = \|f_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)}$ and then write

$$\|f_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)} \leq \|f_{\mathcal{Y},\Lambda,\epsilon} - p_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)} + \|p_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)}. \quad (4.8)$$

Consider the first term. By (4.6), we have

$$\begin{aligned} \|f_{\mathcal{Y},\Lambda,\epsilon} - p_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)} &= \|\mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathcal{S}_{\mathcal{Y}}(f - p)\|_{L^2(D,v)} \\ &= \|(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathcal{S}_{\mathcal{Y}}(f - p)\|_2 \\ &\leq \frac{1}{\epsilon}\|\mathcal{S}_{\mathcal{Y}}(f - p)\|_2 \leq \frac{1}{\epsilon}\|f - p\|_{L^\infty(\Omega)}. \end{aligned} \quad (4.9)$$

Here in the second step we use Parseval’s identity, in the third step we use standard properties of the SVD and in the fourth step we use (4.5). Now consider the second term of (4.8). Observe that $\mathbf{A}_{\mathcal{Y},\Lambda} = \mathcal{S}_{\mathcal{Y}}\mathcal{T}_{\Lambda}$. Hence, if $p = \mathcal{T}_{\Lambda}\mathbf{c}$, then using standard properties of the SVD once more, we get

$$\begin{aligned} \|p_{\mathcal{Y},\Lambda,\epsilon}\|_{L^2(D,v)} &= \|\mathcal{T}_{\Lambda}(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^{\dagger}\mathcal{S}_{\mathcal{Y}}\mathcal{T}_{\Lambda}\mathbf{c}\|_{L^2(D,v)} \\ &= \|(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^{\dagger}\mathbf{A}_{\mathcal{Y},\Lambda}\mathbf{c}\|_2 \\ &\leq \|\mathbf{c}\|_2 = \|p\|_{L^2(D,v)}. \end{aligned}$$

Combining this with (4.9) and substituting both into (4.8) now gives the second result. □

A few remarks are in order. First, to guarantee accuracy and good (absolute) conditioning of the approximation, we need to ensure that $C_{\mathcal{Y},\Lambda,\epsilon} \lesssim 1$. This constant depends on the polynomial space, the data and the threshold ϵ , but is independent of the function f . In Section 6, we derive bounds for this constant. Second, once $C_{\mathcal{Y},\Lambda,\epsilon}$ is bounded, the approximation error is determined via the term $E_{\Lambda,\epsilon}(f)$, which depends on f and the polynomial space but is independent of the data. We estimate this term for functions in certain Sobolev spaces in Section 5.

Third, we note that the coefficients of the ensuing approximation are bounded by the approximation error divided by ϵ . Thus, although the coefficients may initially be $\mathcal{O}(1/\epsilon)$, they are $\mathcal{O}(1)$ in the limit as the dimension N of the approximation space P_{Λ} tends to infinity. Note that bounded coefficients are particularly important for practical computations, since these are the values that will be stored. Indeed, if the coefficients could grow arbitrarily large in relation to the function f , then the pointwise evaluation operator $\mathbf{c}^{\epsilon} \mapsto f_{\mathcal{Y},\Lambda,\epsilon}(\mathbf{x})$ would be ill-conditioned.

Fourth and finally, we note that $C_{\mathcal{Y},\Lambda,\epsilon} \leq 1/(\sqrt{v_{\Omega}}\epsilon)$ for any \mathcal{Y}, Λ and $\epsilon > 0$ [6, Proposition 4.6]. In other words, the ill-conditioning of the reconstruction operator scales at worst like $1/\epsilon$.

5. Approximation error for Legendre polynomial frames

We now consider the approximation error $E_{\Lambda,\epsilon}(f)$, defined by equation (4.7). In doing so, we treat the following two scenarios separately:

- (i) f defined and smooth over D ;
- (ii) f undefined or nonsmooth over D .

We first require several notions of smoothness. Let

$$H^m(\Omega, \mu) = \{f \in L^2(\Omega, \mu) : \mathcal{D}^j f \in L^2(\Omega, \mu) : |j|_1 \leq m\} \tag{5.1}$$

be the classical Sobolev spaces of index $m \geq 0$ on Ω , with norm

$$\|f\|_{H^m(\Omega, \mu)} = \sqrt{\sum_{|j|_1 \leq m} \|\mathcal{D}^j f\|_{L^2(\Omega, \mu)}^2}.$$

Here $\mathcal{D}^j = \frac{\partial^{|j|_1}}{\partial y_1^{j_1} \dots \partial y_d^{j_d}}$ is the partial derivative operator of order j . These spaces are suitable for approximations using the tensor product or total degree spaces in low dimensions. For moderate to high dimensions, we instead consider Sobolev spaces of dominating mixed smoothness:

$$H_{\text{mix}}^m(D, \nu) = \{f \in L^2(D, \nu) : \mathcal{D}^j f \in L^2(D, \nu) : |j|_\infty \leq m\}, \tag{5.2}$$

with norm

$$\|f\|_{H_{\text{mix}}^m(D, \nu)} = \sqrt{\sum_{|j|_\infty \leq m} \|\mathcal{D}^j f\|_{L^2(D, \nu)}^2}.$$

5.1. Results for the classical Sobolev spaces H^m . We first consider the tensor product and total degree index sets.

THEOREM 5.1. *Let P_Λ be constructed from the tensor Legendre polynomial basis on $L^2(D, \nu)$, where $D = (-1, 1)^d$, ν is the uniform measure on D , and $\Lambda = \Lambda_n$ is either the tensor product (3.4) or total degree (3.5) index set of degree n . If $\Omega \subseteq D$ and $f \in H^m(D, \nu)$ for some $m > d/2$, then*

$$E_{\Lambda, \epsilon}(f) \leq c_{m,d} \|f\|_{H^m(D, \nu)} n^{\theta(d)-m} + \epsilon \|f\|_{L^2(D, \nu)},$$

where

$$\theta(d) = \begin{cases} \frac{d(2d+1)}{2d+2} & \text{odd } d \\ \frac{d(2d+3)}{2d+4} & \text{even } d. \end{cases} \tag{5.3}$$

Conversely, if $\Omega \subseteq D$ is Lipschitz and $f \in H^m(\Omega, \mu)$, where μ is the uniform measure on Ω and $m > d/2$, then

$$E_{\Lambda, \epsilon}(f) \leq c_{m,d,\Omega} (n^{\theta(d)-m} + \epsilon) \|f\|_{H^m(\Omega, \mu)}.$$

Proof. Let $\Lambda = \Lambda_n^{\text{TP}}$. In the first case, since f is defined over the whole of D , we may let $p = f_\Lambda$ be its orthogonal projection onto $\text{span}\{\psi_n : \mathbf{n} \in \Lambda\} \subset L^2(D, \nu)$. Then

$$E_{\Lambda, \epsilon}(f) \leq \|f - f_\Lambda\|_{L^\infty(D)} + \epsilon \|f_\Lambda\|_{L^2(D, \nu)} \leq \|f - f_\Lambda\|_{L^\infty(D)} + \epsilon \|f\|_{L^2(D, \nu)}. \quad (5.4)$$

It remains to estimate the first term. For this, we first use the Gagliardo–Nirenberg inequality (see, for example, [30]) to give

$$\|f - f_\Lambda\|_{L^\infty(D)} \leq c_{k,d} \|f - f_\Lambda\|_{H^k(D, \nu)}^{\frac{d}{2k}} \|f - f_\Lambda\|_{L^2(D, \nu)}^{1 - \frac{d}{2k}}, \quad d < 2k \leq 2m.$$

We now use the estimate

$$\|f - f_\Lambda\|_{H^l(D, \nu)} \leq c_{l,m,d} n^{\sigma(l)-m} \|f\|_{H^m(D, \nu)},$$

where $\sigma(l) = 0$ for $l = 0$ and $\sigma(l) = 2l - 1/2$ for $l > 0$ (see, for example, [17, (5.8.11)]). Hence

$$\|f - f_\Lambda\|_{L^\infty(D)} \leq c_{k,m,d} n^{\frac{d(2k-1/2-m)}{2k} - m(1-\frac{d}{2k})} \|f\|_{H^m(D, \nu)} = c_{k,m,d} n^{d(1-\frac{1}{4k})-m} \|f\|_{H^m(D, \nu)}.$$

Setting $k = \frac{d+1}{2}$ (odd d) or $k = \frac{d+2}{2}$ (even d) and substituting into (5.4) yields the first result for $\Lambda = \Lambda_n^{\text{TP}}$. For the total degree index set $\Lambda = \Lambda_n^{\text{TD}}$, we first recall that $\Lambda_{n/d}^{\text{TP}} \subseteq \Lambda_n^{\text{TP}}$. We therefore let $p = f_{\Lambda_{n/d}^{\text{TP}}}$ so that

$$E_{\Lambda_n^{\text{TD}}, \epsilon}(f) \leq \|f - f_{\Lambda_{n/d}^{\text{TP}}}\|_{L^\infty(D)} + \epsilon \|f\|_{L^2(D, \nu)}.$$

The result for this index set now follows from the previous bound for Λ_n^{TD} .

Now consider the case where Ω is Lipschitz and $f \in H^m(\Omega, \mu)$. We follow the argument of [7, Proposition 5.8]. We first note that there is an extension g of f to $H^m(D, \nu)$ satisfying

$$\|g\|_{H^m(D, \nu)} \leq c_{m,d,\Omega} \|f\|_{H^m(\Omega, \mu)}.$$

Now let $p = g_\Lambda$ be the orthogonal projection of g onto $\text{span}\{\psi_n : \mathbf{n} \in \Lambda\}$. Then

$$\|p\|_{L^2(D, \nu)} \leq \|g\|_{L^2(D, \nu)} \leq c_{m,d,\Omega} \|f\|_{H^m(\Omega, \mu)},$$

and, by the previously derived result,

$$\begin{aligned} \|f - p\|_{L^\infty(\Omega)} &\leq \|g - g_\Lambda\|_{L^\infty(D)} \leq c_{m,d,\Omega} n^{\theta(d)-m} \|g\|_{H^m(D, \nu)} \\ &\leq c_{m,d,\Omega} n^{\theta(d)-m} \|f\|_{H^m(\Omega, \mu)}. \end{aligned}$$

This gives the second result. \square

Unsurprisingly, in scenario (i), one obtains a slightly better error bound, where the constant in the ϵ term involves the smaller L^2 -norm as opposed to the H^m -norm. For completeness, we now also consider the hyperbolic cross index set.

THEOREM 5.2. *Let P_Λ be constructed from the tensor Legendre polynomial basis on $L^2(D, \nu)$, where $D = (-1, 1)^d$, ν is the uniform measure on D , and $\Lambda = \Lambda_n$ is the hyperbolic cross index set (3.6) of degree n . If $\Omega \subseteq D$ and $f \in H^m(D, \nu)$ for some $m > d/2$, then*

$$E_{\Lambda, \epsilon}(f) \leq c_{m,d} \|f\|_{H^{md}(D, \nu)} n^{\frac{\theta(d)-m}{d}} + \epsilon \|f\|_{L^2(D, \nu)},$$

where $\theta(d)$ is as in (5.3). Conversely, if $\Omega \subseteq D$ is Lipschitz and $f \in H^m(\Omega, \mu)$, where μ is the uniform measure on Ω and $m > d/2$, then

$$E_{\Lambda, \epsilon}(f) \leq c_{m,d,\Omega} \left(n^{\frac{\theta(d)-m}{d}} + \epsilon \right) \|f\|_{H^m(\Omega, \mu)}.$$

Proof. Let $n^* = \lfloor (n+1)^{1/d} - 1 \rfloor$ and observe that $\Lambda_{n^*}^{\text{TP}} \subseteq \Lambda_n^{\text{HC}}$. We now use the arguments from the proof of the previous theorem. \square

As is to be expected, Theorems 5.1 and 5.2, which assume only classical Sobolev regularity, all exhibit the curse of dimensionality. This can be seen by noting that

$$n^{\theta(d)-m} \asymp N^{\frac{d-m}{d}}, \quad n \rightarrow \infty,$$

for fixed d whenever Λ is the total degree or tensor product index set, since in both cases $N = |\Lambda| \asymp n^d$. Conversely, for the hyperbolic cross index set, one has

$$n^{\frac{d-m}{d}} \asymp N^{\frac{d-m}{d}} (\log(N))^{\frac{(m-d)(d-1)}{d}},$$

since in this case $N \asymp n(\log(n))^{d-1}$.

5.2. Results for the mixed Sobolev spaces H_{mix}^m . Seeking to mitigate the curse of dimensionality when using the hyperbolic cross index set, we now consider the mixed Sobolev spaces $H_{\text{mix}}^m(D, \nu)$.

THEOREM 5.3. *Let P_Λ be constructed from the tensor Legendre polynomial basis on $L^2(D, \nu)$, where $D = (-1, 1)^d$ and ν is the uniform measure on D . If $f \in H_{\text{mix}}^m(D, \nu)$ for some $m \geq 1$, then*

$$E_{\Lambda, \epsilon}(f) \leq \begin{cases} c_{m,d} \|f\|_{H_{\text{mix}}^m(D, \nu)} n^{1-m} + \epsilon \|f\|_{L^2(D, \nu)} & \Lambda = \Lambda_n^{\text{TP}} \text{ or } \Lambda = \Lambda_n^{\text{TD}} \\ c_{m,d} \|f\|_{H_{\text{mix}}^m(D, \nu)} n^{1-m} (\log n)^{\frac{d-1}{2}} + \epsilon \|f\|_{L^2(D, \nu)} & \Lambda = \Lambda_n^{\text{HC}}. \end{cases}$$

Furthermore, if Ω is compactly contained in D , then

$$E_{\Lambda, \epsilon}(f) \leq \begin{cases} c_{m,d,\Omega} \|f\|_{H_{\text{mix}}^m(D,v)} n^{1/2-m} + \epsilon \|f\|_{L^2(D,v)} & \Lambda = \Lambda_n^{\text{TP}} \text{ or } \Lambda = \Lambda_n^{\text{TD}} \\ c_{m,d,\Omega} \|f\|_{H_{\text{mix}}^m(D,v)} n^{1/2-m} (\log n)^{\frac{d-1}{2}} + \epsilon \|f\|_{L^2(D,v)} & \Lambda = \Lambda_n^{\text{HC}}. \end{cases}$$

Proof. Since $f \in L^2(D, v)$, we may let $p = f_\Lambda$ be its orthogonal projection onto $\text{span}\{\psi_n : n \in \Lambda\}$. Then, using (A.1), (A.2) and (A.4), we obtain

$$\begin{aligned} \|f - f_\Lambda\|_{L^\infty(D)} &\leq \sum_{n \notin \Lambda} \prod_{k=1}^d \sqrt{2n_k + 1} |\langle f, \psi_n \rangle_{L^2(D,v)}| \\ &\leq \left(\sum_{n \notin \Lambda} \chi_{n,m}^{\text{mix}} |\langle f, \psi_n \rangle_{L^2(D,v)}|^2 \right)^{1/2} \left(\sum_{n \notin \Lambda} \frac{\prod_{k=1}^d (2n_k + 1)}{\chi_{n,m}^{\text{mix}}} \right)^{1/2}, \\ &\leq \|f\|_{\tilde{H}_{\text{mix}}^m(D,v)} \left(\sum_{n \notin \Lambda} \frac{\prod_{k=1}^d (2n_k + 1)}{\chi_{n,m}^{\text{mix}}} \right)^{1/2}, \end{aligned}$$

where $\tilde{H}_{\text{mix}}^m(D, v)$ and $\chi_{n,m}^{\text{mix}}$ are as in (A.3) and (A.5), respectively. Observe that

$$\chi_{n,m}^{\text{mix}} \geq c_{m,d} \left(\prod_{k=1}^d (n_k + 1) \right)^{2m}$$

for some constant $c_{m,d}$, and therefore

$$\|f - f_\Lambda\|_{L^\infty(D)} \leq c_{m,d} \|f\|_{H_{\text{mix}}^m(D,v)} \left(\sum_{n \notin \Lambda} \prod_{k=1}^d (n_k + 1)^{1-2m} \right)^{1/2}, \tag{5.5}$$

where we also note that $\|f\|_{\tilde{H}_{\text{mix}}^m(D,v)} \leq c_{m,d} \|f\|_{H_{\text{mix}}^m(D,v)}$. We now specify the index set. First, suppose that $\Lambda = \Lambda_n^{\text{TP}}$. Let $[d]$ denote the set of ordered tuples with entries in $\{1, \dots, d\}$. Then

$$\begin{aligned} \sum_{n \notin \Lambda} \prod_{k=1}^d (n_k + 1)^{1-2m} &= \sum_{\substack{\sigma \in [d] \\ k \notin \sigma}} \sum_{\substack{n_k=0 \\ k \in \sigma}}^n \prod_{k=1}^d (n_k + 1)^{1-2m} \\ &= \sum_{\sigma \in [d]} \left(1 + \sum_{l=1}^n l^{1-2m} \right)^{d-|\sigma|} \left(\sum_{l \geq n} l^{1-2m} \right)^{|\sigma|} \leq c_{m,d} n^{2-2m}. \end{aligned}$$

Substituting into (5.5) now gives the result for $\Lambda = \Lambda_n^{\text{TP}}$. Moreover, the result for the total degree index set now follows as well, after noting that $\Lambda_n^{\text{TD}} \supseteq \Lambda_n^{\text{TP}}$.

Finally, for the hyperbolic cross index $\Lambda = \Lambda_n^{\text{HC}}$ set we use, for example, [1, Lemma 2.30] to get

$$\sum_{n \notin \Lambda} \prod_{k=1}^d (n_k + 1)^{1-2m} \leq c_{m,d} n^{2-2m} (\log(n))^{d-1},$$

as required.

It remains to consider the case where Ω is compactly contained in D . We first recall that univariate Legendre polynomials are uniformly bounded in compact subintervals of $(-1, 1)$:

$$|\psi_n(y)| \leq c_r, \quad -1 + r \leq y \leq 1 - r, \quad \forall n \in \mathbb{N}_0, \quad 0 < r < 1,$$

for some $c_r > 0$. Hence $\|\phi_n\|_{L^\infty(\Omega)} \leq c_\Omega, \forall n \in \mathbb{N}_0^d$. Letting $p = f_\Lambda$ and arguing as before, we get

$$\begin{aligned} \|f - f_\Lambda\|_{L^\infty(\Omega)} &\leq c_{m,d,\Omega} \|f\|_{H_{\text{mix}}^m(D,v)} \left(\sum_{n \notin \Lambda} \frac{1}{\chi_{n,m}^{\text{mix}}} \right)^{1/2} \\ &\leq c_{m,d,\Omega} \|f\|_{H_{\text{mix}}^m(D,v)} \left(\sum_{n \notin \Lambda} \prod_{k=1}^d (n_k + 1)^{-2m} \right)^{1/2}. \end{aligned}$$

We now proceed in the same way, replacing the exponent $1 - 2m$ by $-2m$ throughout. \square

6. Sample complexity

In this section, we consider the efficiency of the approximation. In view of Theorem 4.1, this requires estimating the constant $C_{\mathcal{Y},\Lambda,\epsilon}$ defined in (4.2). Our main results are twofold. First, in Section 6.2, we show that when the sample points are drawn independently according to a suitable measure on Ω , then the sample complexity can always be related to the constant of a certain Nikolskii inequality for the polynomial space P_Λ . Second, in Section 6.3, we show that for domains satisfying a suitable property, this constant is at most log-quadratic in the dimension N of the polynomial space P_Λ .

6.1. The constant $C_{\mathcal{Y},\Lambda}$. It is difficult to analyze $C_{\mathcal{Y},\Lambda,\epsilon}$ directly, since it is defined in terms of the singular values and singular vectors of the matrix \mathbf{A} . In order to provide concrete bounds, we now consider

$$C_{\mathcal{Y},\Lambda} = \sup \left\{ \|p\|_{L^2(\Omega,\mu)} : p \in P_\Lambda, \frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |p(\mathbf{y})|^2 = 1 \right\}.$$

Note that $C_{\mathcal{Y},\Lambda}$ depends only on the samples \mathcal{Y} and the space P_Λ . Unlike $C_{\mathcal{Y},\Lambda,\epsilon}$, it is independent of functions ϕ_n , $\mathbf{n} \in \Lambda$, used to span this space and consequently the domain D as well. We also have the following.

LEMMA 6.1. *Let $C_{\mathcal{Y},\Lambda,\epsilon}$ be as in (4.2). Then $C_{\mathcal{Y},\Lambda,\epsilon} \leq C_{\mathcal{Y},\Lambda}$ and moreover $C_{\mathcal{Y},\Lambda} = C_{\mathcal{Y},\Lambda,\epsilon}$ whenever the minimum singular value of $\mathbf{A} = \mathbf{A}_{\mathcal{Y},\Lambda}$ satisfies $\sigma_{\min}(\mathbf{A}) > \epsilon$.*

Proof. Recall that $C_{\mathcal{Y},\Lambda,\epsilon}$ is the maximum of $C'_{\mathcal{Y},\Lambda,\epsilon}$ and $C''_{\mathcal{Y},\Lambda,\epsilon}$. Consider $C'_{\mathcal{Y},\Lambda,\epsilon}$. Let $\mathbf{b} \in \mathbb{C}^M$, $\|\mathbf{b}\|_2 = 1$, and note that $\mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{b} \in P_\Lambda$. Hence, by the definition of $C_{\mathcal{Y},\Lambda}$ and (4.5), we have

$$\|\mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{b}\|_{L^2(\Omega,\mu)} \leq C_{\mathcal{Y},\Lambda} \|\mathcal{S}_\mathcal{Y} \mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{b}\|_2.$$

Note that $\mathbf{A} = \mathbf{A}_{\mathcal{Y},\Lambda} = \mathcal{S}_\mathcal{Y} \mathcal{T}_\Lambda$. By standard properties of the SVD,

$$\|\mathcal{S}_\mathcal{Y} \mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{b}\|_2 \leq \|\mathbf{b}\|_2 = 1,$$

and therefore $\|\mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{b}\|_{L^2(\Omega,\mu)} \leq C_{\mathcal{Y},\Lambda}$. Since \mathbf{b} was arbitrary, we get $C'_{\mathcal{Y},\Lambda,\epsilon} \leq C_{\mathcal{Y},\Lambda}$.

On the other hand, suppose that $\sigma_{\min}(\mathbf{A}) > \epsilon$. Let $p = \mathcal{T}_\Lambda \mathbf{c} \in P_\Lambda$ with $\|\mathcal{S}_\mathcal{Y} p\|_2 = 1$. Let $\mathbf{b} = \mathcal{S}_\mathcal{Y} p$ and write $p = \mathcal{T}_\Lambda \mathbf{c}$. Then

$$\begin{aligned} C'_{\mathcal{Y},\Lambda,\epsilon} &\geq \|\mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathcal{S}_\mathcal{Y} \mathcal{T}_\Lambda \mathbf{c}\|_{L^2(\Omega,\mu)} = \|\mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{A}_{\mathcal{Y},\Lambda} \mathbf{c}\|_{L^2(\Omega,\mu)} \\ &= \|\mathcal{T}_\Lambda \mathbf{c}\|_{L^2(\Omega,\mu)} = \|p\|_{L^2(\Omega,\mu)}, \end{aligned}$$

where in the third step we use the fact that $\mathbf{A}_{\mathcal{Y},\Lambda}$ is full rank. Hence, since p was arbitrary, we get $C'_{\mathcal{Y},\Lambda,\epsilon} \geq C_{\mathcal{Y},\Lambda}$, and therefore $C'_{\mathcal{Y},\Lambda,\epsilon} = C_{\mathcal{Y},\Lambda}$ in this case.

Finally, consider $C''_{\mathcal{Y},\Lambda,\epsilon}$. Let $\mathbf{d} \in \mathbb{C}^N$, $\|\mathbf{d}\|_2 = 1$. Then

$$\begin{aligned} \|\mathcal{T}_\Lambda \mathbf{d} - \mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathcal{T}_\Lambda \mathbf{d}\|_{L^2(\Omega,\mu)} &\leq C_{\mathcal{Y},\Lambda} \|\mathcal{S}_\mathcal{Y}(\mathcal{T}_\Lambda \mathbf{d} - \mathcal{T}_\Lambda(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{A}_{\mathcal{Y},\Lambda} \mathbf{d})\|_2 \\ &= C_{\mathcal{Y},\Lambda} \|(\mathbf{A}_{\mathcal{Y},\Lambda} - \mathbf{A}_{\mathcal{Y},\Lambda}(\mathbf{A}_{\mathcal{Y},\Lambda,\epsilon})^\dagger \mathbf{A}_{\mathcal{Y},\Lambda}) \mathbf{d}\|_2 \\ &= C_{\mathcal{Y},\Lambda} \|\mathbf{U}(\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_\epsilon) \mathbf{V}^* \mathbf{d}\|_2 \leq C_{\mathcal{Y},\Lambda} \epsilon \|\mathbf{d}\|_2. \end{aligned}$$

Since \mathbf{d} was arbitrary, we get $C''_{\mathcal{Y},\Lambda,\epsilon} \leq C_{\mathcal{Y},\Lambda}$ as required. On the other hand, if $\sigma_{\min}(\mathbf{A}) > \epsilon$, then $\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_\epsilon = \mathbf{0}$. Hence $C''_{\mathcal{Y},\Lambda,\epsilon} = 0$. \square

6.2. Random sampling for compact domains and Nikolskii inequalities.

We now show that $C_{\mathcal{Y},\Lambda}$ can be bounded using the constant of a suitable Nikolskii inequality for the space $P_\Lambda \subset L^2(\Omega, \mu)$. To this end, let $N(P_\Lambda, \Omega, \mu)$ be the smallest positive number in the $(L^2(\Omega, \mu), L^\infty(\Omega))$ -Nikolskii inequality

$$\|p\|_{L^\infty(\Omega)} \leq N(P_\Lambda, \Omega, \mu) \|p\|_{L^2(\Omega,\mu)}, \quad \forall p \in P_\Lambda. \quad (6.1)$$

Then we have the following result.

THEOREM 6.2. *Let $0 < \delta, \gamma < 1$ and suppose that $\mathbf{y}_1, \dots, \mathbf{y}_M$ are independent and randomly drawn according to the probability measure μ defined by (3.2). If*

$$M \geq (\mathbf{N}(P_\Lambda, \Omega, \mu))^2 ((1 - \delta) \log(1 - \delta) + \delta)^{-1} \log(N/\gamma),$$

where $N = |\Lambda|$ and $\mathbf{N}(P_\Lambda, \Omega, \mu)$ is the constant of the Nikolskii inequality (6.1), then with probability at least $1 - \gamma$ the quantity $C_{\mathcal{Y}, \Lambda}$ satisfies

$$C_{\mathcal{Y}, \Lambda} \leq \frac{1}{\sqrt{1 - \delta}}.$$

Proof. Our proof is based on essentially the same arguments as those used in previous works (see, for example, [22]). First, let $\{\Phi_n\}_{n \in \Lambda}$ be an orthonormal basis for P_Λ in $L^2(\Omega, \mu)$. Let $p \in P_\Lambda$ be arbitrary and write $p = \sum_{n \in \Lambda} c_n \Phi_n$, so that

$$\|p\|_{L^2(\Omega, \mu)}^2 = \int_{\Omega} |p(\mathbf{y})|^2 d\mu(\mathbf{y}) = \|\mathbf{c}\|_2^2,$$

where $\mathbf{c} = (c_n)_{n \in \Lambda}$, and $\frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} |p(\mathbf{y})|^2 = \mathbf{c}^* \mathbf{B} \mathbf{c}$, where $\mathbf{B} \in \mathbb{C}^{N \times N}$ is the self-adjoint matrix with

$$(\mathbf{B})_{m, n} = \frac{1}{M} \sum_{\mathbf{y} \in \mathcal{Y}} \overline{\Phi_m(\mathbf{y})} \Phi_n(\mathbf{y}), \quad m, n \in \Lambda.$$

It follows that $C_{\mathcal{Y}, \Lambda} = 1/\sqrt{\lambda_{\min}(\mathbf{B})}$, where $\lambda_{\min}(\mathbf{B})$ is the minimal eigenvalue of \mathbf{B} . We estimate this quantity by writing it in the usual way as the sum of random matrices:

$$\mathbf{B} = \sum_{m=1}^M \mathbf{X}_m, \quad \mathbf{X}_m = \left\{ \frac{1}{M} \overline{\Phi_m(\mathbf{y}_m)} \Phi_n(\mathbf{y}_m) \right\}_{m, n \in \Lambda}.$$

By construction, these matrices are independent, nonnegative definite and satisfy $\mathbb{E}(\mathbf{X}_m) = \frac{1}{M} \mathbf{I}$, where \mathbf{I} is the identity matrix. Moreover, for any $\mathbf{c} \in \mathbb{C}^N$ we have

$$\begin{aligned} \mathbf{c}^* \mathbf{X}_m \mathbf{c} &= \frac{1}{M} \left| \sum_{n \in \Lambda} c_n \Phi_n(\mathbf{y}_m) \right|^2 \leq \frac{(\mathbf{N}(P_\Lambda, \Omega, \mu))^2}{M} \left\| \sum_{n \in \Lambda} c_n \Phi_n \right\|_{L^2(\Omega, \mu)}^2 \\ &= \frac{(\mathbf{N}(P_\Lambda, \Omega, \mu))^2}{M} \|\mathbf{c}\|_2^2. \end{aligned}$$

The matrix Chernoff bound (see, for example, [48, Theorem 1.1]) now gives

$$\mathbb{P}(\lambda_{\min}(\mathbf{X}) \leq (1 - \delta)) \leq N \exp\left(-\frac{(1 - \delta) \log(1 - \delta) + \delta}{M^{-1}(\mathbf{N}(P_\Lambda, \Omega, \mu))^2}\right).$$

Setting the right-hand side equal to γ and rearranging yields the result. \square

This leads to the following result on accuracy of the truncated SVD least-squares approximation.

COROLLARY 6.3. *Let $0 < \delta, \gamma < 1$ and suppose that $\mathbf{y}_1, \dots, \mathbf{y}_M$ are independent and randomly drawn according to the probability measure μ defined by (3.2). Let*

$$M \geq (N(P_\Lambda, \Omega, \mu))^2((1 - \delta) \log(1 - \delta) + \delta)^{-1} \log(N/\gamma),$$

where $N = |\Lambda|$ and $N(P_\Lambda, \Omega, \mu)$ is the constant of the Nikolskii inequality (6.1). Then with probability at least $1 - \gamma$ the truncated SVD least-squares approximation $f_{\mathcal{R}, \Lambda, \epsilon}$ of $f \in L^\infty(\Omega)$ satisfies

$$\|f - f_{\mathcal{R}, \Lambda, \epsilon}\|_{L^2(\Omega, \mu)} \leq \left(1 + \frac{1}{\sqrt{1 - \delta}}\right) E_{\Lambda, \epsilon}(f),$$

where $E_{\Lambda, \epsilon}(f)$ is as in (4.7).

6.3. The λ -rectangle property and log-quadratic sample complexity. We now consider $N(P_\Lambda, \Omega, \mu)$. Estimating this constant for general irregular domains in arbitrarily many dimensions is an open problem. We shall not attempt to resolve it in full generality here (see Section 9 for some further discussion). Instead, we show that this constant is at most log-quadratic for a large class of irregular domains whenever μ is the uniform measure.

The types of domain we now consider are those satisfying the so-called following property.

DEFINITION 6.4 (λ -rectangle property). A compact domain Ω has the λ -rectangle property for some $0 < \lambda < 1$ if it can be written as a (possibly overlapping and uncountable) union

$$\Omega = \bigcup_{R \in \mathcal{R}} R$$

of hyperrectangles R satisfying

$$\inf_{R \in \mathcal{R}} \text{Vol}(R) = \lambda \text{Vol}(\Omega).$$

There are many domains of interest that have this property. We now list several examples:

- *L-shaped domains.* These are unions of two rectangles, so they clearly have this property.

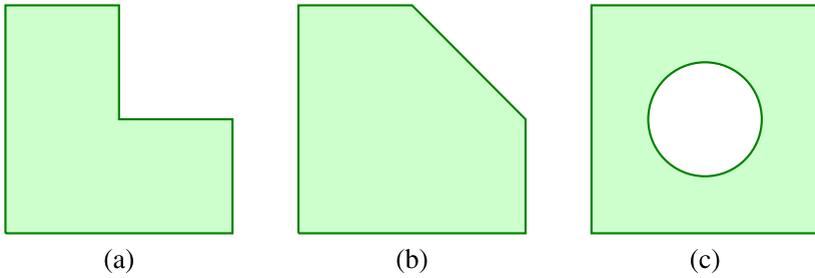


Figure 1. Examples of domains that have the λ -rectangle property: (a) is an L -shaped domain, (b) is a domain with a linear constraint, and (c) is a domain with an exclusion.

- *Domains with linear constraints.* The domain

$$\Omega = \{-1 \leq y_1, y_2 \leq 1, y_1 + y_2 \leq 1\},$$

along with its various higher-dimensional generalizations, can be expressed as

$$\Omega = \bigcup_{x \in [0,1]} R_x, \quad R_x = [-1, x] \otimes [-1, 1 - x].$$

Hence it has the λ -rectangle property with $\lambda = 4/7$. Note that such domains can occur in problems such as surrogate forward model construction in parameter studies; for instance, whenever two parameters y_1 and y_2 , rather than being independent, satisfy a (possibly *a priori* unknown) linear relation.

- *Domains with exclusions.* The domain

$$\Omega = \{-1 \leq y_1, y_2 \leq 1, y_1^2 + y_2^2 \geq 1/2\},$$

along with various generalizations, also satisfies the λ -rectangle property. Note that such domains correspond to practical scenarios where, due to certain physical constraints, $f(\mathbf{y})$ can only be evaluated for \mathbf{y} not too close to zero.

See Figure 1 for illustrations. On the other hand, there are a number of notable domains that do not have this property. These include the unit Euclidean ball $\{\mathbf{y} \in \mathbb{R}^d : \|\mathbf{y}\|_2 \leq 1\}$ and the simplex $\{\mathbf{y} \in \mathbb{R}^d : 0 \leq y_1, \dots, y_{d-1} \leq 1, 0 \leq y_d \leq 1 - (y_1 + \dots + y_{d-1})\}$. See Section 9 for additional details.

THEOREM 6.5. *Suppose that $\Omega \subseteq (-1, 1)^d$ has the λ -rectangle property and let P_A be constructed from the tensor Legendre polynomial basis on $(-1, 1)^d$,*

where Λ is any lower set (see Definition 3.1) of cardinality $|\Lambda| = N$. Let μ be the uniform probability measure on Ω and $N(P_\Lambda, \Omega, \mu)$ be the constant in the Nikolskii inequality (6.1). Then

$$(N(P_\Lambda, \Omega, \mu))^2 \leq \frac{N^2}{\lambda}.$$

Proof. We first claim that $P_\Lambda = \text{span}\{\phi_n : n \in \Lambda\}$ coincides with the space

$$\mathbb{P}_\Lambda = \text{span}\{\mathbf{y} \in \Omega \mapsto \mathbf{y}^n : \mathbf{n} = (n_1, \dots, n_d) \in \Lambda\}.$$

Here we use the notation $\mathbf{y}^n = y_1^{n_1} \cdots y_d^{n_d}$. Since $\phi_n(\mathbf{y})$ is a tensor Legendre polynomial, we have

$$\phi_n(\mathbf{y}) = \prod_{k=1}^d \psi_{n_k}^{(k)}(y_k) = \prod_{k=1}^d \left(\sum_{m_k=0}^{n_k} a_{m_k, n_k} y_k^{m_k} \right) = \sum_{m_1=0}^{n_1} \cdots \sum_{m_d=0}^{n_d} a_{m, n} \mathbf{y}^m,$$

where a_{m_k, n_k} are the coefficients of $\psi_{n_k}^{(k)}$ in the monomial basis and $a_{m, n} = \prod_{k=1}^d a_{m_k, n_k}$. Since $m_k \leq n_k$ for all k , it follows from the lower set assumption that $\mathbf{m} \in \Lambda$ and therefore $\phi_n \in \mathbb{P}_\Lambda$. Hence $P_\Lambda \subseteq \mathbb{P}_\Lambda$. In a similar manner, one also finds that $\mathbf{y} \mapsto \mathbf{y}^n$ is in P_Λ , and therefore $\mathbb{P}_\Lambda \subseteq P_\Lambda$, as required.

Now let $p \in \mathbb{P}_\Lambda$ and $\mathbf{y} \in \Omega$ with $\mathbf{y} \in R$ for some $R \in \mathcal{R}$. Define the uniform measure on R as

$$d\tilde{\mu}(\mathbf{y}) = \frac{1}{\text{Vol}(R)} d\mathbf{y},$$

and note that $|p(\mathbf{y})| \leq N(\mathbb{P}_\Lambda, R, \tilde{\mu}) \|p\|_{L^2(R, \tilde{\mu})}$, where $N(\mathbb{P}_\Lambda, R, \tilde{\mu})$ is the Nikolskii constant for the space \mathbb{P}_Λ with respect to $L^2(R, \tilde{\mu})$. It is known that $(N(\mathbb{P}_\Lambda, R, \tilde{\mu}))^2 \leq N^2$ [35, Theorem 6]. Also

$$\|p\|_{L^2(R, \tilde{\mu})}^2 = \frac{1}{\text{Vol}(R)} \int_R |p(\mathbf{y})|^2 d\mathbf{y} \leq \frac{\text{Vol}(\Omega)}{\text{Vol}(R)} \int_\Omega |p(\mathbf{y})|^2 d\mu(\mathbf{y}) \leq \frac{1}{\lambda} \|p\|_{L^2(\Omega, \mu)}^2.$$

Hence $|p(\mathbf{y})|^2 \leq \frac{N^2}{\lambda} \|p\|_{L^2(\Omega, \mu)}^2$. Since $\mathbf{y} \in \Omega$ and $p \in \mathbb{P}_\Lambda$ were arbitrary, we now get the result. □

Combining this with Corollary 6.3 now gives the following.

COROLLARY 6.6. *Suppose that $\Omega \subseteq (-1, 1)^d$ has the λ -rectangle property and let P_Λ be constructed from the tensor Legendre polynomial basis on $(-1, 1)^d$, where Λ is any lower set of cardinality $|\Lambda| = N$. Let $0 < \delta, \gamma < 1$ and suppose*

that $\mathbf{y}_1, \dots, \mathbf{y}_M$ are independent and randomly drawn according to the uniform probability measure on Ω . If

$$M \geq N^2 \lambda^{-1} ((1 - \delta) \log(1 - \delta) + \delta)^{-1} \log(N/\gamma),$$

where $N = |\Lambda|$, then with probability at least $1 - \gamma$ the truncated SVD least-squares approximation $f_{\mathcal{R}, \Lambda, \epsilon}$ of $f \in L^\infty(\Omega, \mu)$ satisfies

$$\|f - f_{\mathcal{R}, \Lambda, \epsilon}\|_{L^2(\Omega, \mu)} \leq \left(1 + \frac{1}{\sqrt{1 - \delta}}\right) E_{\Lambda, \epsilon}(f),$$

where $E_{\Lambda, \epsilon}(f)$ is as in (4.7).

7. Truncated estimators and L^2 -error bounds

The error bounds proved in Corollary 6.3 and elsewhere have the limitation of relating (in probability) the L^2 -norm of the error to an approximation error $E_{\Lambda, \epsilon}(f)$ measured in the L^∞ -norm. In this penultimate section, we show that it is possible to bound the L^2 -norm of a related estimator in expectation in terms of the L^2 -norm approximation error

$$\tilde{E}_{\Lambda, \epsilon}(f) = \inf\{\|f - p\|_{L^2(\Omega, \mu)} + \epsilon \|p\|_{L^2(D, \nu)} : p \in P_\Lambda\}. \tag{7.1}$$

We follow the approach of [22]. First, suppose that $f \in L^\infty(\Omega, \mu)$ and let $L \geq 0$ be such that $\|f\|_{L^\infty(\Omega)} \leq L$. Now define the truncation operator

$$T_L(g)(\mathbf{y}) = \text{sign}(g(\mathbf{y})) \min\{|g(\mathbf{y})|, L\},$$

where $\text{sign}(z)$ denotes the complex sign of $z \in \mathbb{C}$. If $f_{\mathcal{R}, \Lambda, \epsilon}$ is the truncated SVD least-squares approximation, we now consider the new approximation

$$f_{\mathcal{R}, \Lambda, \epsilon, L} = T_L(f_{\mathcal{R}, \Lambda, \epsilon}). \tag{7.2}$$

Our main result is now the following.

THEOREM 7.1. *Let $0 < \delta, \gamma < 1$ and $f \in L^\infty(\Omega)$ with $\|f\|_{L^\infty(\Omega)} \leq L$ for some $L \geq 0$. Let $\mathbf{y}_1, \dots, \mathbf{y}_M$ be independent and randomly drawn according to μ and $f_{\mathcal{R}, \Lambda, \epsilon, L}$ be as in (7.2). If $\tilde{E}_{\Lambda, \epsilon}(f)$ is as in (7.1) and*

$$M \geq (\mathbf{N}(P_\Lambda, \Omega, \mu))^2 ((1 - \delta) \log(1 - \delta) + \delta)^{-1} \log(N/\gamma), \tag{7.3}$$

where $N = |\Lambda|$ and $\mathbf{N}(P_\Lambda, \Omega, \mu)$ is the constant of the Nikolskii inequality (6.1), then

$$\mathbb{E}(\|f - f_{\mathcal{R}, \Lambda, \epsilon, L}\|_{L^2(\Omega, \mu)}^2) \leq 3 \frac{2 - \delta}{1 - \delta} (\tilde{E}_{\Lambda, \epsilon}(f))^2 + 4L^2\gamma.$$

Proof. The proof is based on [22, Theorem 2]. Let E be the event $C_{\gamma,\Lambda,\epsilon} \leq \frac{1}{\sqrt{1-\delta}}$, where $C_{\gamma,\Lambda,\epsilon}$ is as in (4.2). Lemma 6.1, Theorem 6.2 and the measurement condition (7.3) give that $\mathbb{P}(E^c) \leq \gamma$. Now let $d\mu$ be the uniform measure on Ω and $d\mu_M = d\mu \otimes \cdots \otimes d\mu$ be the probability measure of the draw $\mathbf{y}_1, \dots, \mathbf{y}_M$. Then

$$\begin{aligned} & \mathbb{E}(\|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2) d\mu_M \\ &= \int_E \|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2 d\mu_M + \int_{E^c} \|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2 d\mu_M \\ &\leq \int_E \|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2 d\mu_M + 4L^2\gamma. \end{aligned} \quad (7.4)$$

It remains to bound the first term. Assume the event E occurs and let $p \in P_\Lambda$ be such that $\|f - p\|_{L^2(\Omega,\mu)} + \epsilon\|p\|_{L^2(D,v)} = \tilde{E}_{\Lambda,\epsilon}(f)$ (it is straightforward to show that such a minimizer exists, since P_Λ is finite-dimensional). Then, arguing as in the proof of Theorem 4.1 and using the fact that $C_{\gamma,\Lambda,\epsilon} \leq \frac{1}{\sqrt{1-\delta}}$, we have

$$\begin{aligned} & \|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2 \\ &\leq (\|f - p\|_{L^2(\Omega,\mu)} + C'_{\gamma,\Lambda,\epsilon}\|\mathcal{S}_\gamma(f - p)\|_2 + \epsilon C''_{\gamma,\Lambda,\epsilon}\|p\|_{L^2(D,v)})^2 \\ &\leq 3\|f - p\|_{L^2(\Omega,\mu)}^2 + \frac{3}{1-\delta}\|\mathcal{S}_\gamma(f - p)\|_2^2 + \frac{3\epsilon^2}{1-\delta}\|p\|_{L^2(D,v)}^2. \end{aligned}$$

Hence

$$\begin{aligned} \int_E \|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2 d\mu_M &\leq 3\|f - p\|_{L^2(\Omega,\mu)}^2 + \frac{3}{1-\delta}\mathbb{E}(\|\mathcal{S}_\gamma(f - p)\|_2^2) \\ &\quad + \frac{3\epsilon^2}{1-\delta}\|p\|_{L^2(D,v)}^2. \end{aligned}$$

Observe that $\mathbb{E}(\|\mathcal{S}_\gamma(f - p)\|_2^2) = \mathbb{E}|f(\mathbf{y}) - p(\mathbf{y})|^2 = \|f - p\|_{L^2(\Omega,\mu)}^2$. Therefore we obtain

$$\begin{aligned} \int_E \|f - f_{\gamma,\Lambda,\epsilon,L}\|_{L^2(\Omega,\mu)}^2 d\mu_M &\leq 3\frac{2-\delta}{1-\delta}(\|f - p\|_{L^2(\Omega,\mu)}^2 + \epsilon^2\|p\|_{L^2(D,v)}^2) \\ &\leq 3\frac{2-\delta}{1-\delta}(\tilde{E}_{\Lambda,\epsilon}(f))^2. \end{aligned}$$

Substituting this into (7.4) now gives the result. \square

Much like in Section 5, we can establish bounds for $\tilde{E}_{\Lambda,\epsilon}(f)$ under different regularity conditions.

THEOREM 7.2. *Let P_Λ be constructed from the tensor Legendre polynomial basis on $L^2(D, \nu)$, where $D = (-1, 1)^d$, ν is the uniform measure on D , and $\Lambda = \Lambda_n$ is either the tensor product (3.4) or total degree (3.5) index set of degree n . If $\Omega \subseteq D$ and $f \in H^m(D, \nu)$ for some $m \geq 1$, then*

$$\tilde{E}_{\Lambda, \epsilon}(f) \leq c_{m,d} \|f\|_{H^m(D, \nu)} n^{-m} + \epsilon \|f\|_{L^2(D, \nu)}.$$

Conversely, if $\Omega \subseteq (-1, 1)^d$ is Lipschitz and $f \in H^m(\Omega, \mu)$, where μ is the uniform measure on Ω and $m \geq 1$, then

$$\tilde{E}_{\Lambda, \epsilon}(f) \leq c_{m,d, \Omega} (n^{-m} + \epsilon) \|f\|_{H^m(\Omega, \mu)}.$$

Finally, if $\Lambda = \Lambda_n^{\text{HC}}$ is the hyperbolic cross index set (3.6), then the same results hold with n^{-m} replaced by $n^{-m/d}$.

Proof. As in the proof of Theorem 5.1, if $f \in H^m(D, \nu)$, we let f_Λ be the orthogonal projection of f onto $\text{span}\{\psi_n : \mathbf{n} \in \Lambda\} \subset L^2(D, \nu)$. Then by Parseval’s identity and (A.4),

$$\begin{aligned} \|f - f_\Lambda\|_{L^2(D, \nu)}^2 &= \sum_{\mathbf{n} \notin \Lambda} |\langle f, \psi_n \rangle_{L^2(D, \nu)}|^2 \leq \frac{1}{\min_{\mathbf{n} \notin \Lambda} \{\chi_{n,m}\}} \sum_{\mathbf{n} \notin \Lambda} \chi_{n,m} |\langle f, \psi_n \rangle_{L^2(D, \nu)}|^2 \\ &\leq \frac{1}{\min_{\mathbf{n} \notin \Lambda} \{\chi_{n,m}\}} \|f\|_{\tilde{H}^m(D, \nu)}^2. \end{aligned}$$

It remains to bound $\min_{\mathbf{n} \notin \Lambda} \{\chi_{n,m}\}$ for the three index sets. Using (A.5), we first observe that

$$\chi_{n,m} = \sum_{|j|_1 \leq m} \prod_{k=1}^d (n_k(n_k + 1))^{j_k} \geq |\mathbf{n}|_\infty^{2m} > n^{2m}, \quad \mathbf{n} \notin \Lambda_n^{\text{TP}}.$$

Similarly, for the total degree index set

$$\chi_{n,m} \geq c_{m,d} |\mathbf{n}|_1^{2m} > c_{m,d} n^{2m}, \quad \mathbf{n} \notin \Lambda_n^{\text{TD}},$$

and for the hyperbolic cross

$$\chi_{n,m} \geq c_{m,d} |\mathbf{n}|_{\text{hc}}^{2m/d} > c_{m,d} n^{2m/d}, \quad \mathbf{n} \notin \Lambda_n^{\text{HC}}.$$

This gives the first result. For the second result, we argue as in the proof of Theorem 5.1 to construct an extension $g \in H^m(D, \nu)$ of f , and then use the previously derived bounds. \square

THEOREM 7.3. *Let P_Λ be constructed from the tensor Legendre polynomial basis on $L^2(D, \nu)$, where $D = (-1, 1)^d$ and ν is the uniform measure on D . If $f \in H_{\text{mix}}^m(D, \nu)$ for some $m \geq 1$, then*

$$\tilde{E}_{\Lambda, \epsilon}(f) \leq c_{m,d} \|f\|_{H_{\text{mix}}^m(D, \nu)} n^{-m} + \epsilon \|f\|_{L^2(D, \nu)}$$

when $\Lambda = \Lambda_n^{\text{TP}}$, $\Lambda = \Lambda_n^{\text{TD}}$ or $\Lambda = \Lambda_n^{\text{HC}}$.

Proof. Consider the setup of the previous proof. We have

$$\|f - f_\Lambda\|_{L^2(D, \nu)}^2 \leq \frac{1}{\min_{\mathbf{n} \notin \Lambda} \{\chi_{\mathbf{n}, m}^{\text{mix}}\}} \|f\|_{\tilde{H}_{\text{mix}}^m(D, \nu)}^2,$$

where $\chi_{\mathbf{n}, m}^{\text{mix}}$ is as in (A.5). We now observe that

$$\chi_{\mathbf{n}, m}^{\text{mix}} = \sum_{|j|_\infty \leq m} \prod_{k=1}^d (n_k(n_k + 1))^{j_k} \geq c_{m,d} |\mathbf{n}|_{\text{hc}}^{2m} > c_{m,d} n^{2m}, \quad \mathbf{n} \notin \Lambda,$$

where Λ is any of the three index sets considered. The result now follows immediately. □

8. Numerical results and discussion

We conclude this paper with several numerical experiments illustrating the theoretical results. Unless otherwise stated, we use Legendre polynomials on $D = (-1, 1)^d$, hyperbolic cross index sets, samples drawn independently from the uniform measure on Ω and a threshold parameter $\epsilon = 10^{-8}$.

8.1. Function regularity. We first consider the approximation of several bivariate functions. The left panel of Figure 2 shows the approximation of a smooth function on the domain $\Omega = \{\mathbf{y} : f(\mathbf{y}) \geq 0\}$. The function is singular on $D \setminus \Omega$. Yet, as predicted by the results of Section 5, this does not hamper its approximation on Ω . The right panel shows the approximation of a function defined on the Mandelbrot set. This domain is not Lipschitz, but since the function has a smooth extension to the whole of D , an accurate approximation is obtained. This also agrees with the results of Section 5. Note that in neither case does the domain need to be known in advance in order to compute the approximation. It is defined implicitly by the data.

8.2. Sample complexity. In Figure 3, we examine the sample complexity of polynomial frame approximations for a two-dimensional circular domain. This

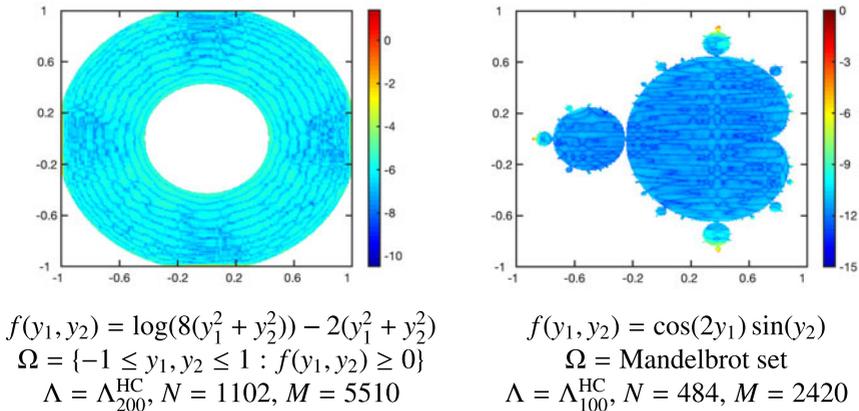


Figure 2. Pointwise error for polynomial frame approximations over two bivariate domains.

requires computing the constant $C_{\mathcal{Y}, \Lambda, \epsilon}$, which is discussed in the remark below. Figure 3(a) suggests that quadratic oversampling is sufficient in this case, even though the domain is not of λ -rectangle type. Moreover, linear or log-linear oversampling results in exponential increase of $C_{\mathcal{Y}, \Lambda, \epsilon}$, up to roughly $1/\epsilon$ (recall that $C_{\mathcal{Y}, \Lambda, \epsilon} \lesssim 1/\epsilon$; see Section 4.3). On the other hand, Figures 3(b),(c) suggest that log-linear oversampling is sufficient whenever domain Ω does not touch the bounding cube D . Furthermore, the constant $C_{\mathcal{Y}, \Lambda, \epsilon}$ gets smaller (for the same level of oversampling) as the radius r of the domain decreases, or in other words, as the distance between the boundary of Ω and the boundary of D grows. These interesting observations, which are at odds with the log-quadratic rates predicted in Section 6, have been thoroughly documented in the one-dimensional case in the related setting when trigonometric polynomials are used instead of algebraic polynomials [8, 10].

While we currently have no proof, it is possible to give an intuitive explanation for this phenomenon. The sample complexity relates to the maximal growth of a polynomial (in an L^2 -sense) on Ω when it is bounded at M points in Ω . A polynomial that grows large in this sense must also be large on $D \setminus \Omega$, and therefore have large coefficients when represented in the Legendre basis. Yet, when regularizing via the truncated SVD (which prohibits large coefficients), such polynomials are excluded from the resulting approximation space. This also explains why the constant $C_{\mathcal{Y}, \Lambda, \epsilon}$ decreases as r decreases: for $r = 1$ the boundaries of Ω and D intersect, but as r decreases these boundaries are increasingly separated. Formalizing this intuition into a proof is an open problem.

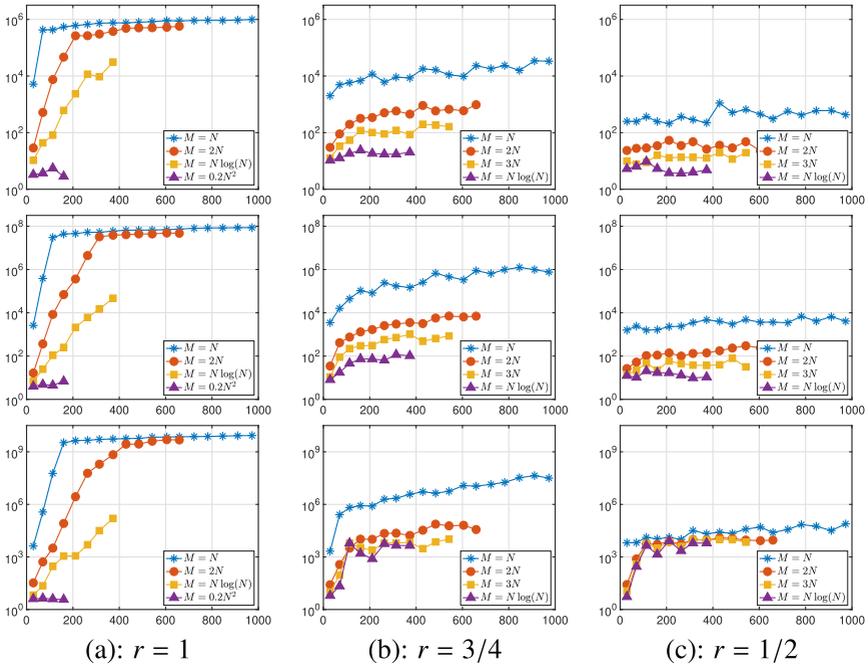


Figure 3. The constant $C_{\gamma, \Lambda, \epsilon}$ against N for $\epsilon = 10^{-6}$ (top row), $\epsilon = 10^{-8}$ (middle row) and $\epsilon = 10^{-10}$ (bottom row). The domain Ω is a circle of radius r in $d = 2$ dimensions. Computations were averaged over 20 trials with the median value taken. The computation of $C_{\gamma, \Lambda, \epsilon}$ was done as in Remark 8.1, using a precomputed grid of $K = 10\,000$ Monte Carlo points in Ω .

REMARK 8.1. As shown in [6], the constants $C'_{\gamma, \Lambda, \epsilon}$ and $C''_{\gamma, \Lambda, \epsilon}$ can be expressed as

$$C'_{\gamma, \Lambda, \epsilon} = \sqrt{\lambda_{\max}((\mathbf{B}')^* \mathbf{G} \mathbf{B}')}, \quad C''_{\gamma, \Lambda, \epsilon} = \epsilon^{-1} \sqrt{\lambda_{\max}((\mathbf{B}'')^* \mathbf{G} \mathbf{B}'')}, \quad (8.1)$$

where $\mathbf{G} = \mathbf{G}_\Lambda$ is the Gram matrix of the truncated frame (2.3), $\mathbf{B}' = (\mathbf{A}_{\gamma, \Lambda, \epsilon})^\dagger = \mathbf{V}(\boldsymbol{\Sigma}_\epsilon)^\dagger \mathbf{U}^*$ and $\mathbf{B}'' = \mathbf{V} \mathbf{I}_\epsilon^\perp \mathbf{V}^*$. Here $\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^*$ is the SVD of \mathbf{A} , and $\mathbf{I}_\epsilon^\perp$ is the diagonal matrix with n th entry 1 if $\sigma_n \leq \epsilon$ and zero otherwise. Computing the Gram matrix \mathbf{G} over an irregular domain is difficult, but it can be done approximately via Monte Carlo integration. Specifically, $\mathbf{G} \approx \mathbf{H}^* \mathbf{H}$, where

$$\mathbf{H} = \mathbf{H}_{K, \Lambda} = \left(\frac{1}{\sqrt{K}} \phi_n(\mathbf{z}_k) \right)_{k=1, \dots, K, n \in \Lambda} \in \mathbb{C}^{K \times N},$$

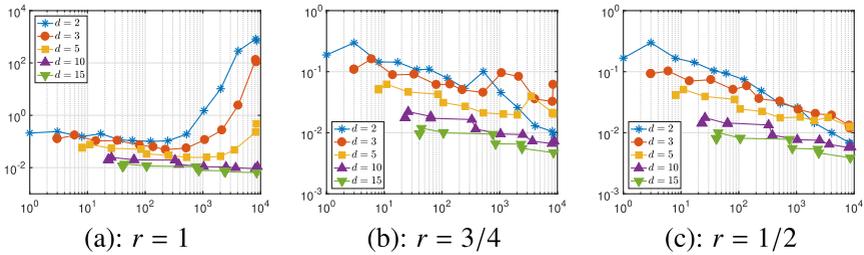


Figure 4. The median error over 20 trials versus M for approximating the function $f(\mathbf{y}) = 1/\sum_{i=1}^d \sqrt{|y_i|}$ on the annular domain $\Omega = \{\mathbf{y} : r/4 \leq \|\mathbf{y}\|_2 \leq r\}$. For each M , the value of N is chosen as the largest such that $N \log(N) \leq M$.

and $\mathbf{z}_1, \dots, \mathbf{z}_K$ are drawn independently from μ . Replacing \mathbf{G} by $\mathbf{H}^* \mathbf{H}$ in (8.1) and using standard properties of singular values leads to the simpler approximate expressions

$$C'_{\mathcal{Y}, \Lambda, \epsilon} \approx \|\mathbf{H}\mathbf{V}(\boldsymbol{\Sigma}_\epsilon)^\dagger\|_2, \quad C''_{\mathcal{Y}, \Lambda, \epsilon} \approx \epsilon^{-1} \|\mathbf{H}\mathbf{V}\mathbf{I}_\epsilon^\perp\|_2,$$

where $\|\cdot\|_2$ denotes the matrix 2-norm.

8.3. Higher dimensions. In Figure 4, we consider the approximation error in various different dimensions. This figure shows the approximation error versus M for an annular region of several different radii. In view of the previous discussion, log-linear oversampling was used throughout. It is noticeable that when $r = 1$, meaning that Ω touches the boundary of D , the approximation is ill-conditioned, and the error duly increases for large enough M . As is to be expected, this increase is most severe in lower dimensions (since the cardinality of the polynomial space is largest in this setting). Conversely, as soon as Ω is compactly contained in D , the approximation error decreases as M increases. Note that the function being approximated is smooth in Ω but singular at $\mathbf{y} = \mathbf{0} \in D \setminus \Omega$. As predicted by the results of Section 5, the approximation error decreases rapidly despite this singularity.

8.4. Choice of ϵ . In this section, we discuss the influence of the regularization parameter on the approximation. First, we note that the approximation is fairly robust to the choice of the parameter ϵ . In the noiseless setting, ϵ can be considered a *target accuracy* for the method: namely, for sufficiently large M and N , the approximation error will be of the order of ϵ (provided, of course, ϵ is larger than machine epsilon, since floating point error will always limit the

best achievable accuracy in practice). Indeed, under the mild conditions that the subspaces $\Lambda = \Lambda_N$ satisfy $\Lambda_1 \subset \Lambda_2 \subset \dots$ and $\bigcup_N \Lambda_N = \mathbb{N}_0^d$ (which certainly holds for all choices considered in this paper), one has

$$\limsup_{N \rightarrow \infty} \tilde{E}_{\Lambda_N, \epsilon}(f) \leq \epsilon \|f\|_{L^2(\Omega, \mu)}.$$

This follows by choosing $p = g_\Lambda$ in (7.1), where $g \in L^2(D, \nu)$ is the extension of f by zero to D and g_Λ is its orthogonal projection (a similar conclusion holds for $E_{\Lambda, \epsilon}(f)$ under slightly stronger regularity assumptions, since this quantity involves an L^∞ -norm as opposed to the L^2 -norm).

This robustness is in stark contrast to the setting of ill-posed problems, where a careful choice of regularization parameter is usually crucial (see, for example, [32, 40]). In such problems, one is typically interested in a specific solution \mathbf{c} of the linear system, and the regularization parameter needs to be carefully chosen to strike a balance between the residual of the linear system and some desired property of \mathbf{c} (for example, smoothness). On the other hand, our concern lies not with the vector \mathbf{c} , but rather with how well the function $f_{\mathcal{Y}, \Lambda, \epsilon}$ approximates f , without preference for one set of coefficients over another, and this implies that success is measured largely by the size of the residual only. Furthermore, success is guaranteed for any f by increasing N due to the completeness of the polynomial frame.

The situation is slightly different if the function samples $f(\mathbf{y})$, $\mathbf{y} \in \mathcal{Y}$, are corrupted by noise. In the setting of ill-posed problems, an optimal choice of the regularization parameter often involves the corner of the L-curve [31]. Yet, the method of this paper remains robust in the noisy setting: the presence of noise merely implies that the limiting accuracy is determined by the maximum of ϵ and the noise level.

The above discussion assumes sufficient oversampling so that the constant $C_{\mathcal{Y}, \Lambda, \epsilon}$ in Theorem 4.1 satisfies $C_{\mathcal{Y}, \Lambda, \epsilon} \lesssim 1$. The parameter ϵ also affects this constant. Generally, $C_{\mathcal{Y}, \Lambda, \epsilon}$ increases as ϵ decreases, reflecting the fact that as ϵ decreases more singular values are retained and the regularized approximation space becomes larger. Hence, smaller ϵ generally means worse conditioning and accuracy of the approximation for fixed M . Or equivalently, a larger M is required to maintain the same level of conditioning and accuracy. Note that this is not reflected in the sample complexity analysis conducted in Section 6, wherein the dependence on ϵ was ignored (recall Lemma 6.1). Nevertheless, this has a practical impact. If one requires only low accuracy (or if accuracy is limited by noise in data), then it is disadvantageous to take ϵ any smaller than needed.

These assertions are confirmed in Figure 3. For all choices of M and N , a larger ϵ implies a smaller constant $C_{\mathcal{Y}, \Lambda, \epsilon}$. These phenomena have also been investigated for the closely related *Fourier extension* approximation in the one-dimensional

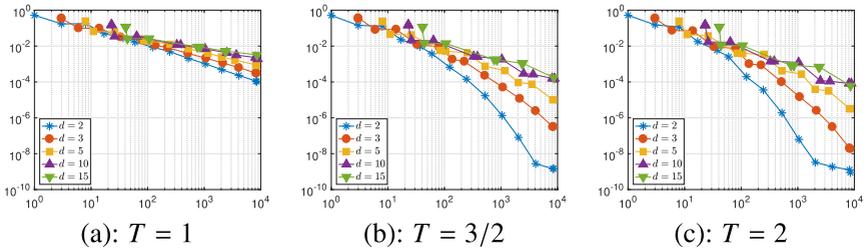


Figure 5. The median error over 20 trials versus M for approximating the function $f(\mathbf{y}) = \exp(-\sum_{i=1}^d y_i/d)$ on the corner domain $\Omega = \{\mathbf{y} \in (-1, 1)^d : y_1 + \dots + y_d \leq 1\}$ using the tensor cosine basis on $[-T, T]^d$. For each M , the value of N is chosen as the largest such that $N \log(N) \leq M$.

setting; see [10]. See also Section 9 for some further comments and formulation of open problems.

8.5. Other bases. Finally, in Figures 5 and 6, we use different orthogonal bases on the extended domain. First, in Figure 5, we consider the tensor cosine basis defined on $D = (-T, T)^d$, where $T \geq 1$ is a parameter. The basis elements in this case are tensor products on the univariate functions

$$\phi_n(y) = \cos(n\pi(y + T)/(2T)), \quad n = 0, 1, 2, \dots$$

When $T = 1$, the domain Ω is not compactly contained in D and the approximation error decreases slowly, at the rate of N^{-1} . This stems from the fact that cosine expansions, much like Fourier expansions, only converge rapidly for smooth functions that satisfy additional boundary conditions [1]. When there is no gap between Ω and the boundary of D , there are no smooth extensions of f satisfying these boundary conditions. Conversely, once $T > 1$ and Ω is compactly contained in D , such extensions exist, and we witness correspondingly faster convergence. Error estimates similar to those proved in Section 5 can also be established for these approximations. See [2] for further details.

In Figure 6, we consider Chebyshev polynomials on $D = (-1, 1)^d$ and random sampling according to the tensor Chebyshev density restricted to Ω with log-linear oversampling. Note that the $d = 2$ approximation exhibits instability. We conjecture that this is related to distribution of the sample points. Points drawn on a cube according to the Chebyshev density cluster quadratically near the boundary of the cube, a property which generally permits a lower sample complexity. However, points drawn according to the same density when restricted to a subdomain Ω do not necessarily cluster in this way over the boundary of Ω .

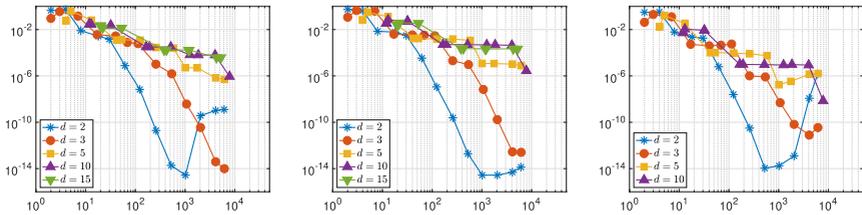


Figure 6. The median error over 20 trials versus M for approximating the function $f(\mathbf{y}) = \cos(\sum_{i=1}^d y_i/d)$ on the domains $\Omega = \{\mathbf{y} \in (-1, 1)^d : y_1 + \dots + y_d \leq 1\}$ (left), $\Omega = \{\mathbf{y} \in (-1, 1)^d : \|\mathbf{y}\|_2 \geq 1/2\}$ (middle) and $\Omega = \{\mathbf{y} \in (-1, 1)^d : \|\mathbf{y}\|_2 \leq 1\}$ (right) using Chebyshev polynomials on $D = (-1, 1)^d$ and samples drawn randomly according to the Chebyshev measure restricted to Ω . For each M , the value of N is chosen as the largest such that $\frac{1}{2}N \log(N) \leq M$.

Unless Ω is compactly contained in D , it appears the severity of the instability is related to the amount of boundary Ω and D share.

9. Conclusions and challenges

In this work, we have introduced and analyzed a method, known as polynomial frame approximation, for approximating multivariate functions on irregular domains. Among the various results proved, we have shown that for functions of mixed Sobolev regularity, the regularized least-squares polynomial frame approximation is well conditioned and converges algebraically fast down to a given threshold parameter ϵ . Moreover, for a large class of domains, the sample complexity is provably quadratic in the dimension of the approximation space, up to a log factor.

This paper marks only a first foray into the broader topic of multivariate polynomial approximation on irregular domains. Consequently, there are a number of interesting challenges for future research. We conclude by highlighting three directions for further work:

1. *Sample complexity estimates.* When sampling from the uniform measure, we have shown log-quadratic sample complexity for λ -rectangle domains, with the factor λ^{-1} appearing in the sample complexity bound. It is unknown whether or not this factor is sharp. Moreover, as mentioned, many domains do not have this property. We conjecture that the same sample complexity holds for a much more general class of domains which includes spheres and simplices (two notable domains which do not have the λ -rectangle property) and which is potentially also

invariant under rotations (rotations generally destroy the λ -rectangle property). This remains an open problem. Moreover, as discussed in Section 8, log-linear sample complexity appears to be sufficient whenever Ω is compactly contained in D . While there is intuition behind this observation, we currently have no proof.

2. *Choice of ϵ .* As discussed in Section 8.4, the method is robust to the choice of ϵ . Yet this parameter does affect the sample complexity. Understanding the intricate relationship between the sample complexity, the domains Ω and D , the subspace P_Λ and the parameter ϵ is very much an open problem. As also noted in Section 8.4, in this paper, we consider a fixed ϵ chosen according to some desired target accuracy (in our experiments, we have simply taken $\epsilon = 10^{-8}$). The possibility of *adaptive* strategies, choosing ϵ depending on Λ and f , is a topic for future work.

3. *Optimal sampling.* Recent work has identified densities for random sampling which achieve near-optimal log-linear sample complexities for least-squares approximations [25]. While these densities can be defined over irregular domains, it becomes challenging to sample efficiently from them in the case where the domain is not of tensor-product type. One solution to this problem is to employ discrete measures, supported over a fine grid that suitably fills Ω . This strategy, which uses ideas of [25], has been recently developed in [5, 36]. Yet this procedure requires the domain Ω to be known in advance, and requires a fine grid to first be generated. This may not be possible in all applications, especially in higher dimensions. For instance, the case $\Omega = \{\mathbf{y} : f(\mathbf{y}) \geq 0\}$, which arises in practical surrogate model construction problems (see Section 1), presents clear difficulties. Developing efficient sampling procedures for such problems remains a topic for future investigation.

4. *Compressed-sensing-based polynomial approximations.* Polynomial-based compressed sensing approaches have recently proved effective for high-dimensional approximation in regular domains (see [3, 4, 20, 39, 42, 51] and references therein). A problem for future work is to extend these approaches to irregular domains. Note that since polynomial frames are redundant, the usual compressed sensing theory for orthogonal bases does not apply.

Acknowledgements

A preliminary version of this work was presented at the Oberwolfach conference on ‘Multiscale and High-Dimensional Problems’. The authors would like to thank the organizers and participants for the useful discussions and

feedback received during the conference. They would also like to thank Claudio Canuto, Juan M. Cardenas, Albert Cohen, Paul Constantine, Ingrid Daubechies, Bert Debuschere, Alireza Doostan, Wolfgang Hackbusch, Sebastian Moraga, Vladimir Temlyakov and Tino Ullrich. The first author is supported by NSERC grant 611675, as well as an Alfred P. Sloan Research Fellowship. The second author is supported by FWO-Flanders projects G.0641.11 and G.A004.14, as well as by KU Leuven project C14/15/055.

Conflict of Interest: None.

Appendix A. Background on Legendre polynomials

This section contains some ancillary results on Legendre polynomials used earlier in the paper. Let $\{\psi_n\}_{n=0}^\infty$ be the orthonormal Legendre polynomial basis with respect to the uniform measure on $(-1, 1)$. This is defined by

$$\psi_n(y) = \sqrt{2n+1} P_n(y), \quad (\text{A.1})$$

where P_n is the classical Legendre polynomial with normalization $P_n(1) = 1$.

A.1. One-dimensional Legendre–Sobolev spaces. Recall that $\psi_n(y)$ are the eigenfunctions of the Sturm–Liouville operator \mathcal{L} , defined by

$$\mathcal{L}f(y) = ((1 - y^2)f'(y))'.$$

Specifically, $\mathcal{L}\psi_n(y) = n(n+1)\psi_n(y)$. The operator \mathcal{L} is compact, self-adjoint and nonnegative definite. Note that

$$\langle \mathcal{L}f, g \rangle_{L^2(D, \nu)} = \langle f', g' \rangle_{L^2(D, \rho)} = \langle f, \mathcal{L}g \rangle_{L^2(D, \nu)},$$

where $D = (-1, 1)$, ν is the uniform measure on $(-1, 1)$ and $d\rho(y) = \frac{1-y^2}{2} dy$. The operator \mathcal{L} has a well-defined square root, which we write as $\mathcal{S} = \mathcal{L}^{1/2}$. Note that

$$\|\mathcal{S}f\|_{L^2(D, \nu)}^2 = \|f'\|_{L^2(D, \rho)}^2 = \langle \mathcal{L}f, f \rangle_{L^2(D, \nu)}.$$

With this in hand, for $j \in \mathbb{N}$, let $\mathcal{S}^j = \mathcal{S} \circ \mathcal{S} \circ \dots \circ \mathcal{S}$ be the j -fold composition of \mathcal{S} and define the Legendre–Sobolev space

$$\tilde{H}^m(D, \nu) = \{f \in L^2(D, \nu) : \mathcal{S}^j f \in L^2(D, \nu), j = 0, \dots, m\},$$

with inner product and norm

$$\langle f, g \rangle_{\tilde{H}^m(D, \nu)} = \sum_{j=0}^m \langle \mathcal{S}^j f, \mathcal{S}^j g \rangle_{L^2(D, \nu)}, \quad \|f\|_{\tilde{H}^m(D, \nu)} = \sqrt{\sum_{j=0}^m \|\mathcal{S}^j f\|_{L^2(D, \nu)}^2}.$$

The set $\{\psi_n\}_{n \in \mathbb{N}_0}$ is an orthogonal basis for $\tilde{H}^m(D, \nu)$, and one has the expression

$$\|f\|_{\tilde{H}^m(D, \nu)} = \sqrt{\sum_{n=0}^{\infty} \chi_{n,m} |\langle f, \psi_n \rangle_{L^2(D, \nu)}|^2}, \quad \chi_{n,m} = \sum_{j=0}^m (n(n+1))^j.$$

Here we use the convention $0^0 = 1$.

A.2. Multidimensional Legendre–Sobolev spaces. Let $D = (-1, 1)^d$ be the unit cube and define the tensor Legendre polynomial basis $\{\psi_n\}_{n \in \mathbb{N}_0^d}$ as

$$\psi_n(\mathbf{y}) = \prod_{k=1}^d \psi_{n_k}(y_k), \quad \mathbf{n} = (n_1, \dots, n_d) \in \mathbb{N}_0^d, \quad \mathbf{y} = (y_1, \dots, y_d) \in D. \quad (\text{A.2})$$

For $k = 1, \dots, d$, let \mathcal{L}_k be the compact, self-adjoint nonnegative definite operator

$$\mathcal{L}_k f(\mathbf{y}) = \frac{\partial}{\partial y_k} \left((1 - y_k^2) \frac{\partial f}{\partial y_k} \right),$$

with corresponding square root $\mathcal{S}_k = \mathcal{L}_k^{1/2}$ and powers $\mathcal{S}_k^j = \mathcal{S}_k \circ \dots \circ \mathcal{S}_k$. Now let $\mathbf{j} = (j_1, \dots, j_d) \in \mathbb{N}_0^d$ be a multi-index. We define the operator

$$\mathcal{S}^{\mathbf{j}} = \mathcal{S}_1^{j_1} \circ \dots \circ \mathcal{S}_d^{j_d}.$$

With this in hand, we now define the d -dimensional Legendre–Sobolev spaces

$$\tilde{H}^m(D, \nu) = \{f \in L^2(D, \nu) : \mathcal{S}^{\mathbf{j}} f \in L^2(D, \nu), |\mathbf{j}|_1 \leq m\},$$

with inner product and norm

$$\langle f, g \rangle_{\tilde{H}^m(D, \nu)} = \sum_{|\mathbf{j}|_1 \leq m} \langle \mathcal{S}^{\mathbf{j}} f, \mathcal{S}^{\mathbf{j}} g \rangle_{L^2(D, \nu)}, \quad \|f\|_{\tilde{H}^m(D, \nu)} = \sqrt{\sum_{|\mathbf{j}|_1 \leq m} \|\mathcal{S}^{\mathbf{j}} f\|_{L^2(D, \nu)}^2}.$$

We also define the mixed d -dimensional Legendre–Sobolev spaces as

$$\tilde{H}_{\text{mix}}^m(D, \nu) = \{f \in L^2(D, \nu) : \mathcal{S}^{\mathbf{j}} f \in L^2(D, \nu), |\mathbf{j}|_{\infty} \leq m\}, \quad (\text{A.3})$$

with inner product and norm

$$\langle f, g \rangle_{\tilde{H}^m_{\text{mix}}(D, \nu)} = \sum_{|j|_{\infty} \leq m} \langle \mathcal{S}^j f, \mathcal{S}^j g \rangle_{L^2(D, \nu)}, \quad \|f\|_{\tilde{H}^m_{\text{mix}}(D, \nu)} = \sqrt{\sum_{|j|_{\infty} \leq m} \|\mathcal{S}^j f\|_{L^2(D, \nu)}^2}.$$

Both these norms can be characterized in terms of Legendre polynomial coefficients. Specifically,

$$\begin{aligned} \|f\|_{\tilde{H}^m(D, \nu)} &= \sqrt{\sum_{n \in \mathbb{N}_0^d} \chi_{n, m} |\langle f, \psi_n \rangle_{L^2(D, \nu)}|^2}, \\ \|f\|_{\tilde{H}^m_{\text{mix}}(D, \nu)} &= \sqrt{\sum_{n \in \mathbb{N}_0^d} \chi_{n, m}^{\text{mix}} |\langle f, \psi_n \rangle_{L^2(D, \nu)}|^2}, \end{aligned} \tag{A.4}$$

where

$$\chi_{n, m} = \sum_{|j|_1 \leq m} \prod_{k=1}^d (n_k (n_k + 1))^{j_k}, \quad \chi_{n, m}^{\text{mix}} = \sum_{|j|_{\infty} \leq m} \prod_{k=1}^d (n_k (n_k + 1))^{j_k}. \tag{A.5}$$

Finally, we note that one has the continuous embeddings $H^m(D, \nu) \hookrightarrow \tilde{H}^m(D, \nu)$ and $H^m_{\text{mix}}(D, \nu) \hookrightarrow \tilde{H}^m_{\text{mix}}(D, \nu)$.

References

- [1] B. Adcock, ‘Modified Fourier expansions: theory, construction and applications’, PhD Thesis, University of Cambridge, 2010.
- [2] B. Adcock, ‘Multivariate modified Fourier series and application to boundary value problems’, *Numer. Math.* **115**(4) (2010), 511–552.
- [3] B. Adcock, ‘Infinite-dimensional compressed sensing and function interpolation’, *Found. Comput. Math.* **18**(3) (2018), 661–701.
- [4] B. Adcock, S. Brugiapaglia and C. G. Webster, ‘Compressed sensing approaches for polynomial approximation of high-dimensional functions’, in *Compressed Sensing and its Applications* (Birkhäuser, Cham, 2017).
- [5] B. Adcock and J. M. Cardenas, ‘Near-optimal sampling strategies for multivariate function approximation on general domains’, *SIAM J. Math. Data Sci.* (2020), (to appear).
- [6] B. Adcock and D. Huybrechs, ‘Frames and numerical approximation II: generalized sampling’, Preprint, 2018, [arXiv:1802.01950](https://arxiv.org/abs/1802.01950).
- [7] B. Adcock and D. Huybrechs, ‘Frames and numerical approximation’, *SIAM Rev.* **61**(3) (2019), 443–473.
- [8] B. Adcock, D. Huybrechs and J. Martín-Vaquero, ‘On the numerical stability of Fourier extensions’, *Found. Comput. Math.* **14**(4) (2014), 635–687.
- [9] B. Adcock, R. Platte and A. Shadrin, ‘Optimal sampling rates for approximating analytic functions from pointwise samples’, *IMA J. Numer. Anal.* **39**(3) (2019), 1360–1390.

- [10] B. Adcock and J. Ruan, 'Parameter selection and numerical approximation properties of Fourier extensions from fixed data', *J. Comput. Phys.* **273** (2014), 453–471.
- [11] N. Albin and O. P. Bruno, 'A spectral FC solver for the compressible Navier–Stokes equations in general domains I: explicit time-stepping', *J. Comput. Phys.* **230**(16) (2011), 6248–6270.
- [12] J. Boyd, 'Fourier embedded domain methods: extending a function defined on an irregular region to a rectangle so that the extension is spatially periodic and C^∞ ', *Appl. Math. Comput.* **161**(2) (2005), 591–597.
- [13] J. P. Boyd, 'A comparison of numerical algorithms for Fourier Extension of the first, second, and third kinds', *J. Comput. Phys.* **178** (2002), 118–160.
- [14] O. Bruno and M. Lyon, 'High-order unconditionally stable FC-AD solvers for general smooth domains I. Basic elements', *J. Comput. Phys.* **229**(6) (2010), 2009–2033.
- [15] O. P. Bruno, Y. Han and M. M. Pohlman, 'Accurate, high-order representation of complex three-dimensional surfaces via Fourier continuation analysis', *J. Comput. Phys.* **227**(2) (2007), 1094–1125.
- [16] H.-J. Bungartz and M. Griebel, 'Sparse grids', *Acta Numer.* **13** (2004), 147–269.
- [17] C. Canuto, M. Y. Hussaini, A. Quarteroni and T. A. Zang, *Spectral Methods: Fundamentals in Single Domains* (Springer, Berlin Heidelberg, 2006).
- [18] 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**(3) (2015), 815–837.
- [19] A. Chkifa, A. Cohen and C. Schwab, 'Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs', *J. Math. Pures Appl.* **103** (2015), 400–428.
- [20] 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), 1415–1450.
- [21] O. Christensen, *An Introduction to Frames and Riesz Bases* (Birkhäuser, Basel, 2003).
- [22] A. Cohen, M. A. Davenport and D. Leviatan, 'On the stability and accuracy of least squares approximations', *Found. Comput. Math.* **13** (2013), 819–834.
- [23] A. Cohen and R. A. DeVore, 'Approximation of high-dimensional parametric PDEs', *Acta Numer.* **24** (2015), 1–159.
- [24] A. Cohen, R. A. DeVore and C. Schwab, 'Convergence rates of best N -term Galerkin approximations for a class of elliptic sPDEs', *Found. Comput. Math.* **10** (2010), 615–646.
- [25] A. Cohen and G. Migliorati, 'Optimal weighted least-squares methods', *SMAI J. Comput. Math.* **3** (2017), 181–203.
- [26] P. G. Constantine, *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies* (SIAM, Philadelphia, PA, 2015).
- [27] A. Doostan and H. Owhadi, 'A non-adapted sparse approximation of PDEs with stochastic inputs', *J. Comput. Phys.* **230**(8) (2011), 3015–3034.
- [28] N. Dyn and M. Floater, 'Multivariate polynomial interpolation on lower sets', *J. Approx. Theory* **177** (2013), 34–42.
- [29] M. I. Ganzburg, 'Polynomial inequalities on measurable sets and their applications', *Constr. Approx.* **17** (2001), 275–306.
- [30] W. Hackbusch, ' L^∞ estimation of tensor truncations', *Numer. Math.* **125** (2013), 419–440.
- [31] P. C. Hansen, 'Analysis of discrete ill-posed problems by means of the l-curve', *SIAM Rev.* **34**(4) (1992), 561–580.
- [32] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion* (SIAM, Philadelphia, PA, 1998).

- [33] M. Lyon and O. Bruno, ‘High-order unconditionally stable FC-AD solvers for general smooth domains II. Elliptic, parabolic and hyperbolic PDEs; theoretical considerations’, *J. Comput. Phys.* **229**(9) (2010), 3358–3381.
- [34] G. Migliorati, ‘Polynomial approximation by means of the random discrete L^2 projection and application to inverse problems for PDEs with stochastic data’, PhD Thesis, Politecnico di Milano, 2013.
- [35] G. Migliorati, ‘Multivariate Markov-type and Nikolskii-type inequalities for polynomials associated with downward closed multi-index sets’, *J. Approx. Theory* **189** (2015), 137–159.
- [36] G. Migliorati, ‘Multivariate approximation of functions on irregular domains by weighted least-squares methods’, Preprint, 2019, [arXiv:1907.12304](https://arxiv.org/abs/1907.12304).
- [37] G. Migliorati, F. Nobile, E. von Schwerin and R. Tempone, ‘Analysis of the discrete L^2 projection on polynomial spaces with random evaluations’, *Found. Comput. Math.* **14** (2014), 419–456.
- [38] A. Narayan, J. D. Jakeman and T. Zhou, ‘A Christoffel function weighted least squares algorithm for collocation approximations’, *Math. Comp.* **86** (2017), 1913–1947.
- [39] A. Narayan and T. Zhou, ‘Stochastic collocation on unstructured multivariate meshes’, *Commun. Comput. Phys.* **18**(1) (2015), 1–36.
- [40] A. Neumaier, ‘Solving ill-conditioned and singular linear systems: a tutorial on regularization’, *SIAM Rev.* **40**(3) (1998), 636–666.
- [41] R. Pasquetti and M. Elghaoui, ‘A spectral embedding method applied to the advection–diffusion equation’, *J. Comput. Phys.* **125** (1996), 464–476.
- [42] J. Peng, J. Hampton and A. Doostan, ‘A weighted ℓ_1 -minimization approach for sparse polynomial chaos expansions’, *J. Comput. Phys.* **267** (2014), 92–111.
- [43] R. Platte, L. N. Trefethen and A. Kuijlaars, ‘Impossibility of fast stable approximation of analytic functions from equispaced samples’, *SIAM Rev.* **53**(2) (2011), 308–318.
- [44] K. Sargsyan, C. Safta, B. J. Debusschere and H. Najm, ‘Uncertainty quantification given discontinuous model response and a limited number of model runs’, *SIAM J. Sci. Comput.* **34**(1) (2012), B44–B64.
- [45] K. Sargsyan, C. Safta, H. Najm, B. J. Debusschere, D. Ricciuto and P. Thornton, ‘Dimensionality reduction for complex models via Bayesian compressive sensing’, *Int. J. Uncertain. Quantif.* **4**(1) (2014), 63–93.
- [46] K. Stinson, D. F. Gleich and P. G. Constantine, ‘A randomized algorithm for enumerating zonotope vertices’, Preprint, 2016, [arXiv:1602.06620](https://arxiv.org/abs/1602.06620).
- [47] V. Temlyakov and S. Tikhonov, ‘Remez-type inequalities for the hyperbolic cross polynomials’, *Constr. Approx.* **46**(3) (2017), 593–615.
- [48] J. A. Tropp, ‘User friendly tail bounds for sums of random matrices’, *Found. Comput. Math.* **12** (2012), 389–434.
- [49] J. A. S. Witteveen and G. Iaccarino, ‘Simplex stochastic collocation with random sampling and extrapolation for nonhypercube probability spaces’, *SIAM J. Sci. Comput.* **34**(2) (2012), A814–A838.
- [50] D. Xiu and G. E. Karniadakis, ‘The Wiener–Askey polynomial chaos for stochastic differential equations’, *SIAM J. Sci. Comput.* **24**(2) (2002), 619–644.
- [51] L. Yan, L. Guo and D. Xiu, ‘Stochastic collocation algorithms using ℓ_1 -minimization’, *Int. J. Uncertain. Quantif.* **2**(3) (2012), 279–293.
- [52] X. Yang and G. E. Karniadakis, ‘Reweighted ℓ_1 minimization method for stochastic elliptic differential equations’, *J. Comput. Phys.* **248** (2013), 87–108.
- [53] T. Zhou, A. Narayan and D. Xiu, ‘Weighted discrete least-squares polynomial approximation using randomized quadratures’, *J. Comput. Phys.* **1** (2015), 787–800.