On the adaptive spectral approximation of functions using redundant sets and frames

Vincent Coppé* Daan Huybrechs†

KU Leuven
Department of Computer Science
Celestijnenlaan 200A
3001 Leuven, Belgium

October 14, 2020

Abstract

The approximation of smooth functions with a spectral basis typically leads to rapidly decaying coefficients where the rate of decay depends on the smoothness of the function and vice-versa. The optimal number of degrees of freedom in the approximation can be determined with relative ease by truncating the coefficients once a threshold is reached. Recent approximation schemes based on redundant sets and frames extend the applicability of spectral approximations to functions defined on irregular geometries and to certain non-smooth functions. However, due to their inherent redundancy, the expansion coefficients in frame approximations do not necessarily decay even for very smooth functions. In this paper, we highlight this lack of equivalence between smoothness and coefficient decay and we explore approaches to determine an optimal number of degrees of freedom for such redundant approximations.

1 Introduction

The approximation of a function $f \in H$ in a Hilbert space H using a spectral basis $\Phi \triangleq \{\phi_k\}_{k \in \mathbb{N}_0}$ for H , such as a family of orthogonal polynomials in $L^2([-1, 1])$, leads to an infinite expansion that can be truncated after finitely many terms N :

$$f(x) \approx f_N(x) = \sum_{k=1}^N c_k \phi_k(x). \quad (1)$$

For spectral bases, the coefficients c_k typically decay rapidly if f is smooth. The optimal number of degrees of freedom N can then be determined with relative ease by truncating the coefficients once the coefficient size reaches a given threshold.

A popular spectral basis for function approximation is the sequence of Chebyshev polynomials $\{T_n\}_{n=0}^\infty$ [31,]. They are extensively used in software packages such as Chebfun [22, 39,] and ApproxFun [34,]. The optimal coefficient vector length in ChebFun is decided by the method called `standardChop`¹. This routine is more involved than a simple truncation based on coefficient size, but it is similar in spirit [4,]. A list of some existing truncation techniques for Chebyshev interpolants is presented in Chapter 3 of [7].

*Email: vincent.coppe@cs.kuleuven.be. Website: <https://people.cs.kuleuven.be/~vincent.coppe>.

†Email: daan.huybrechs@cs.kuleuven.be. Website: <https://people.cs.kuleuven.be/~daan.huybrechs>.

¹This is the name of the routine in Chebfun version 5.3.

For any orthonormal basis $\Phi = \{\phi_k\}_{k \in \mathbb{N}_0}$, the coefficients in (1) are given by $c_k = \langle f, \phi_k \rangle$. Some decay of coefficients in orthonormal bases is therefore guaranteed by Bessel's inequality [38, eq. (2.1.17)]

$$\sum_{k=1}^N |\langle f, \phi_k \rangle|^2 \leq \|f\|_{\mathbb{H}}^2, \quad \forall N \in \mathbb{N}_0. \quad (2)$$

This inequality implies that the series $\sum_{k=1}^{\infty} |\langle f, \phi_k \rangle|^2$ is convergent and bounded by $\|f\|_{\mathbb{H}}^2$. Therefore, the expansion coefficients associated with an orthonormal basis satisfy at least

$$\lim_{k \rightarrow \infty} c_k = \lim_{k \rightarrow \infty} \langle f, \phi_k \rangle = 0.$$

Hence, a truncation strategy can in principle be used for expansions in any orthonormal basis.

However, orthonormal bases are not very flexible, and they might not be available for many problems [12, Ch. 4]. For example, it may be difficult or even impossible to construct a basis offering spectral approximation accuracy on a geometrically complicated bounded domain Ω . On the other hand, one can easily create a bounding box Ξ such that $\Omega \subset \Xi$ and create a basis Φ for $L^2(\Xi)$ — e.g., a tensor product of Fourier series. The basis Φ for $L^2(\Xi)$ is not a basis for $L^2(\Omega)$. Instead, the restriction of an orthonormal basis to a subdomain yields a *frame* [1, 12,]. We will refer to such frames obtained from Fourier series on a bounding box as *Fourier extension* frames. Other types of frames that are easily created include the augmentation of an orthonormal basis by a finite number of additional functions which are bounded, and the concatenation of two or more orthonormal bases. This enables the spectral approximation of wide classes of functions including certain non-smooth functions.

A frame is more general and more flexible than a basis. Precise definitions are given in §2. Importantly, like a basis, a frame $\{\phi_k\}_{k \in \mathbb{N}_0}$ for a Hilbert space $\mathbb{H}$ is complete, such that any function $f \in \mathbb{H}$ can be written as an expansion $f = \sum_{k \in \mathbb{N}_0} c_k \phi_k$. However, unlike a basis, a frame may be redundant, such that the expansion coefficients $\{c_k\}_{k \in \mathbb{N}_0}$ are not necessarily unique.

Many types of frames that have been studied in signal processing [27, 28,] or in the field of wavelets [20,] are associated with a so-called *canonical dual frame* $\{\tilde{\phi}_k\}_{k \in \mathbb{N}_0}$, which has the property that

$$f = \sum_{k \in \mathbb{N}_0} \langle f, \tilde{\phi}_k \rangle \phi_k, \quad f \in \mathbb{H}. \quad (3)$$

This expansion satisfies an analogue of (2). Thus, in many cases of interest, it can be truncated after finitely many terms, and result in an approximation to f . However, this is not always the preferred approach. Sometimes, the dual frame is not known or can not easily be computed. Even if it is known, the expansion (3) may not converge rapidly after truncation with increasing N [1,].

Function approximation using the truncation of an infinite frame, i.e., using $\Phi_N \triangleq \{\phi_k\}_{k=1}^N$, was analysed by [1, 2]. The approximation coefficients are computed using a regularised singular value decomposition of an associated linear system, which is ill-conditioned if Φ_N is redundant. Loosely speaking, the ill-conditioning reflects the fact that the coefficients are not unique. Yet, it was shown that highly accuracy and numerically stable computation of an approximation

$$f \approx f_N = \sum_{k=1}^N c_{N,k} \phi_k$$

can be achieved. Fast algorithms are known for univariate and multivariate Fourier extensions [29, 32, 33,], and for extension approximations based on B-splines [16,] and wavelets [17,]. However, as reflected in our notation, the coefficients c_N depend on N : all coefficients may change in nearly arbitrary ways if N is increased, and in practice they usually do. In fact, since the algorithms suggested in [29, 32, 33, 16, 17] involve routines from randomised linear algebra, the coefficients actually exhibit a fair degree of randomness. Also, perhaps surprisingly, in spite of near-geometric convergence of Fourier extension approximation f_N to f for analytic f and increasing N [26, 3, 40, 23,], the coefficients do not necessarily decay.

The computation of spectral approximations using redundant sets with coefficient decay has not yet received a lot of attention. [30] proposes the solution of a weighted least squares problem for univariate Fourier extension, and adapts the fast method of [29] for an efficient implementation. Two algorithms were proposed by [25], again for Fourier extension, that achieve coefficient decay. The first algorithm uses weighted least squares, similar to [30]. It is stated that in case a good Fourier extension exists, i.e., an extension with rapidly decaying Fourier coefficients, then it will be found numerically [25, Lemma 2.3]. The second method is based on Hermite interpolation (interpolation of function values and derivatives at the endpoints). A similar approach using Hermite interpolation was explored earlier by [36]. The algorithms using weights impose a decay rate a priori, and therefore assume prior knowledge of the smoothness of the function. The algorithms based on Hermite interpolation yield good results for univariate approximations, but are difficult to extend to higher-dimensional domains with general shape. Furthermore, their stability depends on the stability of the Hermite interpolation problem.

Compared to these references, the aim in this paper is the truncation at an optimal number of degrees of freedom. We only assume an algorithm to compute approximation coefficients $\mathbf{c}_N$, we do not assume or enforce that these coefficients decay. (In the process we do suggest one way to do so, without assuming a priori knowledge of f .) We aim for approximations using the truncation of an infinite frame or, more generally, an infinite dictionary² $\Phi \triangleq \{\phi_k\}_{k \in \mathbb{N}_0}$. In view of the above discussion, truncation based on coefficient decay is not a suitable strategy. Our criterion is to adaptively determine a value of N such that

$$\|f - f_N\|_{\mathcal{H}} \leq \delta \|f\|_{\mathcal{H}}, \quad (4)$$

for a given desired relative accuracy δ . The smallest value of N for which the condition holds is the optimal value. In addition, for reasons outlined further on but mainly for improved stability, we aim for the coefficients to satisfy a bound

$$\|\mathbf{c}_N\| \leq \mu \|f\|_{\mathcal{H}}, \quad (5)$$

for a given parameter $\mu > 0$. For expansions in orthonormal bases (4) typically implies (5), but for approximation in the presence of redundancy that is not necessarily the case.

In our implementation, (5) is enforced *implicitly* rather than explicitly since it is not trivial to choose an appropriate value for μ without intricate knowledge of (frame properties of) the dictionary Φ , and we would like to avoid such expert knowledge. Furthermore, the continuous norms in (5) are replaced by discrete approximations, and a considerable part of the paper is devoted to a justification of this approximation. Fortunately, compared to the many subtleties of the analysis, the algorithms we arrive at are relatively straightforward to use and implement.

We focus on optimal truncations of an infinite frame after N terms. This is in contrast to existing literature on adaptive computations using sparse approximations or best N -term approximations in which frames are also sometimes used (see, e.g., [13, 14, 37, 19]). The best N -term approximation u_N to u is based on an optimal subset of N coefficients out of an infinite set. Such adaptive approximations to the solutions of an operator equation $Lu = g$ were explored in [13, 14] using Riesz bases of wavelet type and extended in [37] to wavelet frames and in [19] to Gelfand frames. Efficient algorithms in this context are possible owing to the compression properties of wavelets. In this paper, we fix the subset $1, 2, \dots, N$ of the first N coefficients. This simpler approach may not yield the optimal results of best N -term approximations, but is closer in spirit to the truncation strategies of spectral methods. The flexibility of frames and dictionaries enables the extension of spectral methods to much wider classes of functions.

An alternative approach may be to reorthogonalize a truncated frame, in order to obtain a stable basis for the same space. An example for polynomial approximation problems is given in [9], based on applying Arnoldi iterations to an ill-conditioned Vandermonde matrix. Through a change of variables, the method of [9] is also applicable to univariate Fourier extension problems

²A *dictionary* is a general term in approximation theory to describe a collection of functions that has no discernible structure or properties like a basis or frame, most often used when that collection is overcomplete [11,].

that are used as an example several times in this paper. This comes at the expense of having to compute and store the orthogonalized basis itself along with the coefficients. Yet, by reducing the problem to an orthogonal projection, coefficient decay is introduced and conventional spectral truncation algorithms may be applied.

In §2, we briefly recall frames and frame approximations, and we highlight the particular case of Fourier extension (or *Fourier continuation*). This approximation scheme was first proposed in [6, 10]. In §3, we illustrate the problems associated with truncation based on coefficient decay. We discuss an inefficient way to calculate a frame approximation with decaying coefficients such that truncation can be used, and explore an alternative truncation strategy based on the residual of a least squares problem. This strategy forms the basis of the algorithms in §4 that determine the optimal approximation of length N adaptively. The first algorithm guarantees optimal N but is costly, the second is much more efficient but only approximates the optimal N .

2 Preliminaries

We recall the definition of frames for a Hilbert space H , introduce examples and quote the relevant theory of function approximation with truncated frames [1, 2,]. This theory motivates the stopping criterion (4)–(5) for the adaptive frame approximation, as well as the use of a regularised solver for the ill-conditioned linear system that is used to construct the approximation.

We hasten to add that the algorithms of the paper apply even when $\Phi \triangleq \{\phi_k\}_{k \in \mathbb{N}_0}$ is not a frame. However, in that case, one can not guarantee the existence of stable approximations for all functions in H . In particular, we can not always guarantee (5).

2.1 Frames

Frames are a generalisation of orthogonal and Riesz bases. A frame for a separable Hilbert space H is defined as a sequence $\Phi = \{\phi_k\}_{k \in \mathbb{N}_0}$ that satisfies the so-called *frame condition*

$$A\|f\|_H^2 \leq \sum_{k \in \mathbb{N}_0} |\langle f, \phi_k \rangle|^2 \leq B\|f\|_H^2, \quad \forall f \in H, \quad (6)$$

with constants $0 < A, B < \infty$. For example, the augmented Fourier basis

$$\Phi = \{\phi_k\}_{k \in \mathbb{Z}} \cup \{T_n\}_{n=0}^K, \quad \phi_k(t) = \frac{1}{\sqrt{2}} e^{ik\pi t},$$

where $\{T_n\}_{n=0}^K$ is a finite number of Chebyshev polynomials, is a frame on $H = L^2(-1, 1)$ with lower frame bound $A = 1$ and a finite upper frame bound B that depends on K .

A finite upper frame bound B implies boundedness of the associated *Gram operator*. A positive lower frame bound $A > 0$ implies that, like a basis, the frame is complete in H . The frame condition ensures that the samples $\langle f, \phi_k \rangle$ carry sufficient information to reconstruct the underlying function f (the associated *frame operator* is bounded and boundedly invertible). See [12] for details and the preliminary section of [1] for a concise overview.

Unlike a basis, the frame condition (6) allows the frame elements ϕ_k to be linearly dependent, or nearly linearly dependent. This can be seen with a simple example. Consider the concatenation of two different orthonormal bases $\Phi = \{\phi_k\}_{k \in \mathbb{N}_0}$ and $\Psi = \{\psi_k\}_{k \in \mathbb{N}_0}$ for the same space H . This concatenated set is clearly complete, because both Φ and Ψ are complete. Applying the Riesz identity twice, i.e., $\|f\|_H^2 = \sum_{k \in \mathbb{N}_0} |\langle f, \phi_k \rangle|^2 = \sum_{k \in \mathbb{N}_0} |\langle f, \psi_k \rangle|^2$, we see that the set satisfies the frame condition (6) with $A = B = 2$. To see that there are an infinite number of representations for every $f \in H$, it is enough to note that the zero function can be represented in infinitely many ways. For all $g \in H$, we can find the basis coefficients $\{a_k(g)\}_{k \in \mathbb{N}_0}$ and $\{b_k(g)\}_{k \in \mathbb{N}_0}$ for the bases Φ and Ψ respectively. Then, for each $g \in H$, the concatenation of coefficients $\{a_k(g)\}_{k \in \mathbb{N}_0} \cup \{-b_k(g)\}_{k \in \mathbb{N}_0}$ consists of exact coefficients in the frame $\Phi \cup \Psi$ of the zero function.

Figure 1: Approximation of $f(x) = x$ on $[-1, 1]$ by a classical Fourier series on $[-1, 1]$ (left panel) and by a Fourier extension frame on $[-2, 2]$ (right panel). The frame approximation does not suffer from the Gibbs phenomenon.

2.1.1 Example: Fourier extensions

A Fourier extension frame is a type of frame suitable for spectral function approximation on domains with arbitrarily complex geometries. In order to approximate a function on a bounded domain Ω , this domain is embedded into a bounding box Ξ , $\Omega \subset \Xi$. A tensor product Fourier series on Ξ is an orthogonal basis for $L^2(\Xi)$, which is used to approximate functions in $L^2(\Omega)$. Implicitly, this corresponds to extending a function from Ω to a periodic function on Ξ . One can think of the redundancy in this frame as corresponding to the many different ways in which a smooth function can be extended to another smooth function on a larger domain.

A 1-D Fourier extension example is shown in Fig. 1. The function $f(x) = x$ is smooth but non-periodic on $[-1, 1]$ and famously gives rise to the Gibbs phenomenon when approximated by a Fourier series on $[-1, 1]$ [24,]. In contrast, when using Fourier series on $\Xi = [-2, 2]$, periodicity on $\Omega = [-1, 1]$ is not required and extensions exist that exhibit no Gibbs phenomenon and offer rapid numerical convergence. One such extension is shown in the right panel of Fig. 1.

2.2 Computing frame and dictionary approximations

We summarise the methodology of [1] and [2]. Suppose we want to approximate $f \in \mathcal{H}$ using the truncation $\Phi_N = \{\phi_n\}_{n=1}^N$ of an infinite frame for $\mathcal{H}$. Assume that the data on $f \in \mathcal{H}$ is given by M functionals $\{l_{m,M}(f)\}_{m=1}^M$ applied to f . Two examples of families of functionals include inner products: $l_{m,M}(f) = \langle f, \phi_m \rangle$, and point evaluations: $l_{m,M}(f) = f(x_{m,M})$ for a set of M points $\{x_{m,M}\}_{m=1}^M$. It is advocated in [2] to consider $M > N$, i.e., to compute an approximation based on oversampling.

The approximation problem leads to the following linear (least squares if $M > N$) system:

$$A\mathbf{c} = \mathbf{b}, \quad \mathbf{b} = \{l_{m,M}(f)\}_{m=1}^M, \quad A = \{l_{m,M}(\phi_n)\}_{m=1, n=1}^{M,N} \in \mathbb{C}^{M \times N}. \quad (7)$$

This system is typically ill-conditioned, owing to the redundancy in the set Φ_N . The theme of the references is that regularisation may still lead to accurate and stable approximations. The particular regularisation proposed consists of truncating the singular values of A below a threshold ϵ :

$$A = U\Sigma V^T \approx U\Sigma_\epsilon V^T =: A_\epsilon,$$

where Σ_ϵ is Σ with all singular values smaller than ϵ replaced by 0. If we define the *synthesis operator*

$$\mathcal{T}_N \mathbf{a} = \sum_{n=1}^N a_n \phi_n,$$

then we can write the regularised solution of (7) as $\mathbf{c}^\epsilon = A_\epsilon^\dagger \mathbf{b}$ and denote the corresponding function approximation succinctly by

$$\mathcal{P}_{M,N}^\epsilon f = \mathcal{T}_N \mathbf{c}^\epsilon = \sum_{n=1}^N c_n^\epsilon \phi_n.$$

We comment on the computational cost of this scheme further on in §2.6.

2.3 Error bound

The SVD-regularized method yields the error bound [2, Theorem 1.3]:

$$\|f - \mathcal{P}_{M,N}^\epsilon f\|_{\mathcal{H}} \leq \inf_{\mathbf{z} \in \mathbb{C}^N} \left\{ \|f - \mathcal{T}_N \mathbf{z}\|_{\mathcal{H}} + \kappa_{M,N}^\epsilon \|f - \mathcal{T}_N \mathbf{z}\|_M + \epsilon \lambda_{M,N}^\epsilon \|\mathbf{z}\| \right\}. \quad (8)$$

Here, $\kappa_{M,N}^\epsilon$ and $\lambda_{M,N}^\epsilon$ are constants that depend on the sampling functionals $l_{m,M}$ and on the frame elements ϕ_n , and which we discuss later. The M -norm $\|\cdot\|_M^2 \triangleq \sum_{m=1}^m |l_{m,M}(\cdot)|^2$ is a data-dependent norm. From its definition one can conclude that

$$\|f - \mathcal{T}_N \mathbf{z}\|_M = \|f - \sum_{n=1}^N z_n \phi_n\|_M = \|A\mathbf{z} - \mathbf{b}\|, \quad (9)$$

i.e., the M -norm appearing in the right hand side of (8), is nothing but the discrete ℓ_2 norm of the residual vector of the linear system. Its connection to the norm of the function space $\mathcal{H}$ itself, i.e. the connection between the first two terms in the right hand side of (8), is also further explored later on.

Most importantly, the error bound (8) shows the following. If an accurate approximation to f exists in the span of Φ_N , and if its expansion coefficients $\mathbf{z}$ are bounded in norm, then an approximation will be found numerically that is at least as good. This is the case regardless of the condition number of A . It is often desirable to have small coefficients in post-processing routines. However, the theme of this paper is that, while the coefficients returned typically do have bounded norm $\|\mathbf{z}\|$, they may not necessarily decay even when approximating smooth functions.

Remark 2.1. The error bound (8) holds for all sets Φ_N of bounded functions in $\mathcal{H}$, even if the infinite system Φ is not a frame. The frame condition on Φ ensures that *all* functions $f \in \mathcal{H}$ have at least one representation with bounded coefficients, hence all functions can be approximated numerically (with sufficiently large N) in spite of ill-conditioning. Equivalently, a vector $\mathbf{z}$ exists that makes the right hand side of (8) small. If Φ is not a frame, then the approximation method may still be successful for many functions, but it may also fail for some others.

We note that the coefficients $\mathbf{c}^\epsilon$ of the approximation $\mathcal{P}_{M,N}^\epsilon f = \sum_{k=1}^N c_k^\epsilon \phi_k$ can also be bounded [2, Theorem 3.8]:

$$\|\mathbf{c}^\epsilon\| \leq \inf_{\mathbf{z} \in \mathbb{C}^N} \left\{ \frac{1}{\epsilon} \|f - \mathcal{T}_N \mathbf{z}\|_M + \|\mathbf{z}\| \right\}. \quad (10)$$

This bound implies that, unless f can be well approximated in Φ_N such that the first term above is small, the coefficients found by solving (7) may be as large as ϵ^{-1} . As soon as N is sufficiently large to resolve f , the first term in the bound may decrease and the coefficient norm settles down. The smallest possible coefficient norm is achieved by the dual frame expansion coefficients (3), satisfying

$$\|\mathbf{c}\| \leq \frac{1}{\sqrt{A}} \|f\|_{\mathcal{H}}, \quad (11)$$

where A is the lower frame bound. After the onset of convergence, in view of (10) the regularized coefficient norm can be expected to approach the upper bound in (11).

2.4 Discrete approximation and the stable sampling rate

The analysis of the $\kappa_{M,N}^\epsilon$ and $\lambda_{M,N}^\epsilon$ constants that appear in (8) is highly involved and frame-specific. However, if the sampling functionals are sufficiently ‘rich’ for increasing M , a concept that is defined below, it is guaranteed that both constants are bounded for M sufficiently large relative to N [2, Theorem 1.3]. Furthermore, it is shown that there exists a *stable sampling rate* $M = \Theta^\epsilon(N, \theta)$, with $1 < \theta < \infty$, such that for $M \geq \Theta^\epsilon(N, \theta)$ both constants satisfy the bound

$$\kappa_{M,N}^\epsilon, \lambda_{M,N}^\epsilon \leq \frac{\theta}{A'}, \quad 1 < \theta < \infty. \quad (12)$$

The constant A' appearing in the denominator is determined by the family of sampling functionals. It arises from the ‘Marcinkiewicz-Zygmund’ condition [2, eq. (1.5)] with $0 < A' \leq B' < \infty$,

$$A' \|f\|_H^2 \leq \liminf_{M \rightarrow \infty} \|f\|_M^2 \leq \limsup_{M \rightarrow \infty} \|f\|_M^2 \leq B' \|f\|_H^2, \quad \forall f \in G. \quad (13)$$

The space $G \subset H$ in (13) is a subspace of H on which the sampling functionals are well-defined. In the main examples of this paper $H = L^2(\Omega)$ and $G = C(\Omega)$, since the latter space allows point evaluations (and hence discrete sampling) but the former does not. On the other hand, H is a Hilbert space but G is not. Loosely speaking, in the case of discrete sampling, condition (13) simply means that the points have to be sufficiently well spread as M increases. In addition, ideally the samples of a function f are normalized in such a way that the discrete norm of the samples approximately agrees with the continuous norm of f in H . This is difficult to achieve when $M = N$, but far easier in the limit $M \rightarrow \infty$ for fixed N .

2.5 Practical choices

In practice, we will assume that the stable sampling rate is known. For all examples in this paper, we simply choose a linear oversampling rate

$$M = \gamma N, \quad \gamma > 1, \quad (14)$$

for a value of γ that is deemed sufficiently large based on experiments for a given frame (typically $\gamma = 2$). Because of regularisation, linear oversampling — even with equispaced data — is possible for spectral approximation despite the *impossibility theorem* which states that stable and exponentially converging algorithms using equispaced data require quadratic oversampling (see [2, 35, 8]). Note that regularization precludes true exponential convergence even for analytic functions: the right hand side of (8) is never smaller than $O(\epsilon)$.

Remark 2.2. If the stable sampling rate is not known, and linear oversampling is deemed insufficient, it may be computed numerically [2, §3.6]. However, numerical computation of the stable sampling rate may be expensive, and it requires inner products of the basis functions even if the approximation problem itself is formulated using discrete function evaluations only.

All our examples are based on the function space $H = L^2(\Omega)$ for some domain Ω , and as sampling functionals we choose weighted point evaluations of the form

$$l_{m,M}(f) = w_{m,M} f(x_{m,M}), \quad f \in G = L^\infty(\Omega).$$

The points belong to the domain Ω at hand, and the weights $w_{m,M} > 0$ are chosen such that

$$\|f\|_M^2 = \sum_{m=1}^M w_{m,M}^2 |f(x_{m,M})|^2 \rightarrow \int_{\Omega} |f(x)|^2 dx = \|f\|_{L^2(\Omega)}^2, \quad M \gg 1. \quad (15)$$

This is realised by forming a Riemann sum for $\|f\|_{L^2(\Omega)}$ in the limit $M \rightarrow \infty$. We stress that the agreement between $\|f\|_M$ and $\|f\|_H$ need not be highly accurate. As it will merely be used to detect convergence later on, it does not play a role in the approximation accuracy. Still, we have conveniently achieved that $A' = B' = 1$, referring to the constants in condition (13). When using equispaced points, a suitable choice is

$$w_{m,M} = \sqrt{\frac{|\Omega|}{M}}.$$

Importantly, these choices do not restrict the domain shape Ω . Even for a very irregular geometry, an equispaced grid on a bounding box can be restricted to Ω and the weights defined above yield a Riemann sum in the limit $M \rightarrow \infty$.

Finally, in the adaptive scheme we will not assume that the user knows exactly what the frame bounds A and B are, nor what the constants A' and B' are in (13).

2.6 Computational cost

Because a truncated SVD is generally used to solve (7) with regularization, the computational complexity to find the least squares solutions is typically $\mathcal{O}(MN^2)$. In the regime $M \sim N$ that we adopt in all examples of this paper, this means the cost is cubic in N . However, as mentioned in the introduction, more efficient algorithms exist. The AZ algorithm is the most general form of these efficient algorithms and it was introduced in [18]. The AZ algorithm is compatible with the techniques developed below, for several types of frames. Its computational complexity is much more favourable. For example, it reduces the computational complexity to $\mathcal{O}(N \log^2 N)$ for univariate Fourier extensions. The complexity for higher-dimensional problems remains higher than (log-)linear, but better than cubic – we refer to [18] for details.

The examples in this paper can be reproduced with Jupyter notebooks using our publicly available Julia implementation of frame approximations. We provide two versions of the notebook that creates the figures in this paper, one with and one without the use of the AZ algorithm. They are available here: [15]. The results are entirely similar. In the remainder of this article we focus on convergence characteristics, rather than on the way an approximation is computed for a given value of N .

3 Truncation strategies for frame approximations

An essential element of any adaptive approach for finding the optimal function set size N is the criterion that gives an indication on whether or not the approximation f_N of size N is sufficiently accurate. We say that a solution is accurate if, given $\delta > 0$,

$$\|f - f_N\|_{\mathcal{H}} \leq \delta \|f\|_{\mathcal{H}}. \quad (16)$$

As mentioned in the introduction, a stopping criterion based on coefficient decay is possible in general when approximating with a spectral basis, but that is no longer the case when approximating with truncated frames. We illustrate this restriction in §3.1 and consider an alternative criterion in §3.3–§3.4.

3.1 Truncation based on coefficient size

The decreasing coefficient size in the approximation of a smooth function using an orthogonal basis is shown in the top-left panel of Fig. 2. The panel shows the size of the expansion coefficients for e^x using Chebyshev polynomials on $[-1, 1]$. For spectral approximation schemes, the coefficients decay at a rate that increases with increasing smoothness of the function: since e^x is entire, its expansion coefficients decay super exponentially [5,].

This property is not seen when approximating with frames. In the right panel of Fig. 2, e^x is approximated on $[-1, 1]$ with Chebyshev polynomials that were scaled to $[-2, 2]$. The approximation is computed by solving (7) with regularised least squares, using $\epsilon = 10^{-14}$. This could be called a Chebyshev extension problem: $f(x) = e^x$ is implicitly extended from $[-1, 1]$ to a polynomial of degree $N - 1$ on $[-2, 2]$. Here, some decay of the coefficients is observed for the two approximations shown, corresponding to $N = 20$ and $N = 40$. However, the approximation with larger N actually has larger coefficients and slower decay. Thus, the decay rate of the coefficients does not offer clear information on the optimal value of N .

Furthermore, the size of the coefficients does not even correspond closely to the approximation error. The approximation errors for the two experiments are shown in the bottom row of the figure. In both cases, machine precision accuracy is roughly reached, and the optimal values of N are quite close for the basis approximation (left) and the frame approximation (right). However, the size of the error in the Chebyshev extension approximation problem is much smaller than the size of its expansion coefficients. All coefficients shown in the top-right panel, including the late ones, are orders of magnitude larger than machine precision.

One may be led to believe that the lack of coefficient decay is an artefact of the formulation of the problem as a least squares approximation in (7). Indeed, the current experiment shows

Figure 2: Top: The coefficient sizes of the approximation of $f(x) = e^x$ on $[-1, 1]$ with Chebyshev polynomials on $[-1, 1]$ (left), and with Chebyshev polynomials on $[-2, 2]$ (right) with 20 and 40 degrees of freedom in red squares and blue dots respectively. Bottom: The uniform approximation error for increasing approximation size.

absence of coefficient decay of only one particular solution to the linear system. In view of the redundancy of the frame, there might be other representations of the same function with rapidly decaying coefficients. No doubt this is the case. Thus, the fact that rapidly decaying coefficients are not recovered by the regularised least squares solver means that such coefficient vectors must yield a larger right hand side in the bound (8). However, the question remains how to compute those approximations. The ill-conditioning of the system matrix (in the current example the matrix is singular to working precision once $N > 10$) does not prevent computation of vectors that correspond to a small residual (backward error), but it does prevent computation of any specific solution vector to high accuracy (forward error). In §3.2, we will describe one way to compute approximations with rapidly decaying coefficients, albeit a costly one.

Here, we illustrate the inherent difficulty with a modification of the linear system, in which we penalise the late coefficients. Thus, rather than solving $A\mathbf{c}^\epsilon = \mathbf{b}$ for the coefficient vector $\mathbf{c}^\epsilon$, we solve a weighted least squares problem

$$\begin{aligned} AD\mathbf{y}^\epsilon &= \mathbf{b} \\ \mathbf{c}^\epsilon &= D\mathbf{y}^\epsilon \end{aligned} \tag{17}$$

using a truncated SVD, with threshold ϵ , in the first step. Matrix D is a diagonal matrix with diagonal entries chosen to be, say, $d_{nn} = n^{-\alpha}$, with $\alpha > 0$. The original least squares problem yields, among all possible vectors with small residual, the vector with smallest norm $\|\mathbf{c}^\epsilon\|$. The weighted least squares problem yields a vector with small norm $\|\mathbf{y}^\epsilon\|$, hence $\mathbf{c}^\epsilon = D\mathbf{y}^\epsilon$ may be expected to exhibit $c_k \sim k^{-\alpha}$ decay. This approach is similar to the method of [30] and to the first algorithm in [25]. One difference is that, in [25], the ill-conditioned system is solved using iterative refinement. A disadvantage of this weighted scheme is that one has to decide a priori on the decay rate α , which assumes knowledge of the smoothness properties of f .

The weighted least squares formulation is illustrated in Fig. 3 (blue line) for a Chebyshev extension approximation. The diagonal matrix has entries that decay algebraically with k , the degree of the corresponding Chebyshev polynomial. The smoothed Chebyshev extension approximation does have coefficients that decay quicker than in the original non-weighted formulation of the problem. However, the two issues identified above remain: approximations with larger N have larger coefficients, giving no indication of the optimal value of N , and the size of the coefficients remains significantly larger than the actual approximation error on $[-1, 1]$.

Figure 3: Approximation of $f(x) = e^x$ in the interval $[-1, 1]$ with Chebyshev polynomials on $[-2, 2]$. Left: the approximation $\mathcal{P}_{M,N}^\epsilon f(t)$ with extension, $\epsilon = 10^{-14}$. Middle: the pointwise error $|f(t) - \mathcal{P}_{M,N}^\epsilon(t)|$, $N = 61$. Right: the sizes of $\mathbf{c}_{k,N}$, $N = 61$ (full line, filled marks) and $N = 101$ (dashed line). Red (circle): Approximation, without weight, see (7). Blue (square): Approximation with weight D , $D(k, k) = (10^{-4} + |k| + |k|^2 + |k|^3)^{-1}$, see (17). Green (diamond): Incrementally weighted least squares, see Algorithm 1.

3.2 Optimal coefficient decay

Optimal coefficient decay without a priori knowledge of the smoothness of f can also be achieved with a straightforward, albeit computationally expensive, modification to the weighted least squares problem. It is achieved by choosing the weight of coefficient c_{i+1} proportional to the best approximation error using i degrees of freedom.

Assume that we have solved the frame approximation problem with i degrees of freedom, i.e., we have solved $A_i \mathbf{c}_i^\epsilon = \mathbf{b}_i$ with a residual norm $r_i \triangleq \|A_i \mathbf{c}_i^\epsilon - \mathbf{b}_i\|$. This means that the approximation error $f - f_i$ is on the order of r_i , at least pointwise. If Φ is normalised, it is reasonable to assume that we can approximate the tail $f - f_i$ with coefficients of size r_i or smaller. Thus, we can use the residual corresponding to i degrees of freedom as a weight for the later degrees of freedom in the weighted least squares problem. Doing so for $i = 1, \dots, N$ leads to Algorithm 1.

Algorithm 1 Incrementally weighted least squares

Input: $A_1, \dots, A_N = A$, $\mathbf{b}_1, \dots, \mathbf{b}_N = \mathbf{b}$

Output: $\mathbf{c}^\epsilon$ such that $A\mathbf{c}^\epsilon \approx \mathbf{b}$

```

Solve  $A_1 \mathbf{c}_1^\epsilon = \mathbf{b}_1$ 
 $e_1 \leftarrow \|A_1 \mathbf{c}_1^\epsilon - \mathbf{b}_1\|$ 
 $D_2 = \text{diag}(\|\mathbf{b}\|, e_1)$ 
for  $i = 2, \dots, N$  do
    Solve  $A_i D_i \mathbf{y}_i^\epsilon = \mathbf{b}_i$ 
     $e_i \leftarrow \|A_i D_i \mathbf{y}_i^\epsilon - \mathbf{b}_i\|$ 
     $D_{i+1} = \text{diag}(\|\mathbf{b}\|, e_1, e_2, \dots, e_{i-1}, e_i)$ 
end for
 $\mathbf{c}^\epsilon \leftarrow D_N \mathbf{y}_N^\epsilon$ 

```

We present no rigorous analysis of this algorithm, but we illustrate its potential with an experiment in Fig. 3 (green line). In this experiment, we obtained coefficients that decay down to the approximation error, at the maximal rate allowed by the smoothness of the function and the accuracy of the regularised solver. The method is generally applicable, but comes at the cost of solving approximation problems with increasing number of degrees of freedom. This could be optimised by using the same weight for a range of degrees of freedom, thereby reducing the number of approximation problems that have to be solved. We do not pursue such optimisations further for the time being. Instead, we focus on the optimal truncation problem in the remainder of the paper.

3.3 Truncation based on the residual

We have observed that the coefficients of frame approximations alone do not offer sufficient information about optimal truncation or about the accuracy of the approximation. Fortunately, there is a simple alternative: the size of the residual of the least squares problem. The residual is of little use when computing an interpolant using a basis. In that case, it is always zero [21, Chapter XI]. If all goes well, interpolation yields a well-conditioned square linear system that can be solved to high accuracy, regardless of the accuracy of the corresponding function approximation (i.e., independent of the error between the interpolation points). In the setting of this paper, i.e., with oversampling, the residual does correspond more closely to the approximation error. Recall (9) and note that $f - \mathcal{T}_N \mathbf{z} = f - \sum_{k=1}^N z_k \phi_k$ is the difference between f and its approximation. The discrete residual vector is the sampled version of the continuous residual function, and it remains to verify the connection between $\|\mathbf{A}\mathbf{c} - \mathbf{b}\| = \|f - \mathcal{T}_N \mathbf{z}\|_M$ and $\|f - \mathcal{T}_N \mathbf{z}\|_H$.

A good approximation in the continuous norm immediately implies small error in the M -norm as well, if the functionals $l_{m,M}$ are bounded. If we denote the sampling operator by

$$\mathcal{L}_M(f) : G \rightarrow \mathbb{C}^M : f \rightarrow \{l_{m,M}(f)\}_{m=1}^M, \quad (18)$$

then indeed we have

$$\|f - \mathcal{T}_N \mathbf{z}\|_M = \|\mathcal{L}_M(f - \mathcal{T}_N \mathbf{z})\| \leq L_M \|f - \mathcal{T}_N \mathbf{z}\|_H,$$

where $L_M = \|\mathcal{L}_M\|$. Note that if (13) holds, we expect $L_M \leq \sqrt{B'}$.

However, the question at hand is the converse. Does a small discrete residual imply a small continuous residual? That is, can we find a bound of the form:

$$\|f - \mathcal{T}_N \mathbf{z}\|_H \leq C_M \|f - \mathcal{T}_N \mathbf{z}\|_M,$$

with some constant $C_M < \infty$? The answer to this question is, unfortunately, *no*, since for most spaces G and H one can find a function $f \in G$ satisfying $\|f\|_H > 0$ but such that $\mathcal{L}_M f = 0$. Yet, this is a problem for *all* adaptive numerical methods based on a discretisation. Indeed, even for Chebyshev approximation, the function $f(x) = e^x + T_{100}(x)$ would produce exactly the same convergence plot as in the top-left panel of Fig. 2. Truncation based on coefficient size at $N < 101$ would incur $\mathcal{O}(1)$ error of the approximation, invisible to the discretisation.³

Fortunately, since we oversample, we can at least make a statement about C_M in the asymptotic limit $M \rightarrow \infty$ as a direct consequence of (13).

Lemma 3.1. *Let $f \in G \subset H$ and $\mathbf{z} \in \mathbb{C}^N$, let $\Phi_N \triangleq \{\phi_k\}_{k=1}^N$ be the truncation of a frame Φ for H with all its elements satisfying $\phi_k \in G$, and let a sampling operator $\mathcal{L}_M : G \rightarrow \mathbb{C}^M$ be given by (18). If the associated norm $\|f\|_M \triangleq \|\mathcal{L}_M f\|$ satisfies (13), then*

$$\|f - \mathcal{T}_N \mathbf{z}\|_H \leq \frac{1}{\sqrt{A'}} \sqrt{\liminf_{M \rightarrow \infty} \|f - \mathcal{T}_N \mathbf{z}\|_M}.$$

Proof. This is a consequence of (13) for the function $g \triangleq f - \mathcal{T}_N \mathbf{z}$, noting that $g \in G$. $\square$

We have already made one crucial assumption about M , namely that $M = \Theta^\epsilon(N, \theta)$ follows a stable sampling rate as a function of N . Loosely speaking, this condition guarantees that $\|f\|_M$ carries sufficient information to recover $f \in H$ if the latter can be well approximated in $H_N \triangleq \text{span } \Phi_N$. Hence, we are within the asymptotic regime in which the discrete residual is a good substitute for the continuous residual. However, as long as f is unresolved (N is too small), we can not rely on the residual alone. We will have to augment any check of the residual with other checks to verify the onset of convergence.

³The constructor in Chebfun samples the function at a few random points in order to avoid such scenarios with a certain probability. We will adopt a similar procedure shortly.

Figure 4: Evolution of the H-norm approximation error (dots), residual (squares) and coefficient size (triangles) of the Fourier extension approximation to $f(x) = e^{\cos(8\pi x)}$ for $\epsilon = 10^{-12}, 10^{-9}, 10^{-6}, 10^{-3}$ (indicated with dashed black horizontal line). The asymptotic limit $\|f\|/\sqrt{A}$ for the coefficient size is indicated with the solid black horizontal line.

3.4 A stopping criterion

After solving the approximation problem $\mathbf{A}\mathbf{c}^\epsilon = \mathbf{b}$, the computable tools at our disposal are the coefficient norm $\|\mathbf{c}^\epsilon\|$, the right hand side norm $\|\mathbf{b}\|$ and the residual norm $\|\mathbf{A}\mathbf{c}^\epsilon - \mathbf{b}\|$. Recall that $\|\mathbf{b}\| = \|\mathcal{L}_M f\| = \|f\|_M$ and $\|\mathbf{A}\mathbf{c}^\epsilon - \mathbf{b}\| = \|f - \mathcal{T}_N \mathbf{c}_N^\epsilon\|_M$.

We aim for a combination of (4) and (5), e.g.,

$$\|f - f_N\|_H \leq \delta \|f\|_H \quad \text{and} \quad \|\mathbf{c}_N\| \leq \mu \|f\|_H.$$

3.4.1 Coefficient norm versus residual norm

We set out to illustrate that the coefficient norm and the residual norm convey similar information, yet the residual norm is preferable to use in practice.

It is clear from the previous subsections that it seems a good idea to consider a residual-based criterion

$$\|\mathbf{A}\mathbf{c}^\epsilon - \mathbf{b}\| \leq \delta \|\mathbf{b}\|. \quad (19)$$

Indeed, the residual is equivalent to the pointwise approximation error $\|f - \mathcal{T}_N \mathbf{c}\|_M$ and, once M is sufficiently large, it tends to (a multiple of) $\|f - \mathcal{T}_N \mathbf{c}\|_H$ by Lemma 3.1. Analogously, $\|\mathbf{b}\| = \|f\|_M$ tends to $\|f\|_H$. These expressions return in (4). We can thus approximately replace the true stopping criterion (4) by the discrete and computable condition (19), at least for large M . For simplicity, we use the same constant δ in both inequalities.

As is evident from the coefficient bound (10), the coefficient norm can grow large in the regime before the onset of convergence. From Corollary 3.6 in [2], we also obtain under the stable sampling rate that (recall also (11))

$$\limsup_{N \rightarrow \infty} \|\mathbf{c}^\epsilon\| \leq \frac{1}{\sqrt{A}} \|f\|, \quad (20)$$

where A is the lower frame bound of the frame. The stability requirement (5) can be meaningfully replaced by a corresponding condition involving computable quantities⁴

$$\|\mathbf{c}^\epsilon\| \leq \mu \|\mathbf{b}\|. \quad (21)$$

Here, we again assume that M is sufficiently large. Unfortunately, it is not straightforward to set a suitable value for the parameter μ , in particular because an optimal choice depends on analytical

⁴Recall from Remark 2.1 that coefficients can not be bounded for all functions f if Φ is not a frame, in particular if there is no lower frame bound $A > 0$. In that case, (5) may not be achievable for all functions, but it can still be achieved for some.

knowledge of the frame and the sampling which might not be available in all applications. In particular, one wants to know the quantitative relation between $\|\mathbf{b}\|$ and $\|f\|_{\mathbf{H}}$, as well as the value of A in (11) and (20). If μ is set too small, the stopping criterion may be too strict which impedes the termination of the algorithms suggested below. Luckily, a quantitative choice for μ can be avoided as discussed further on.

Fig. 4 illustrates the evolution of the coefficient norm (triangles) for increasing N along with the H-norm error (dots) and the residual (squares) for $\epsilon = 10^{-12}$ (top left), $\epsilon = 10^{-9}$ (top right), $\epsilon = 10^{-6}$ (bottom left), and $\epsilon = 10^{-3}$ (bottom right). The residual tends to the truncation parameter ϵ (dashed horizontal line) and the coefficient norm to $\|f\|_{\mathbf{H}}/\sqrt{A}$ (full horizontal line). This is in full agreement with the stated theory.

The figure illustrates several important features. Firstly, the residual alone is no good estimate for the H-norm error for small N (and small M). The error of the approximation is invisible to the discretisation. Secondly, the coefficient norm does indeed attain the worst case $\mathcal{O}(\epsilon^{-1})$, suggesting that the second criterion (21) might be needed. However, the figure also shows that the coefficient norm reaches $\|f\|_{\mathbf{H}}/\sqrt{A}$ approximately for the same value of N where the residual reaches ϵ . Loosely speaking, if $\delta \approx \epsilon$, criteria (19) and (21) convey similar information with a well chosen μ . This is to be expected, as the coefficient bound given by (10) indicates that the coefficient size is related to the residual norm.

However, a suitable choice of μ depends on the frame dependent constant A , which may not be known in practice. In contrast, it is far easier to ensure that $A' \approx B' \approx 1$ by weighting the sample points as a Riemann sum, independently of what Φ_N is, to ensure close correspondence between the H and M norms. Also, the coefficient norm tends to decrease very slowly towards its limit, creating a risk of obtaining an optimal approximation size that is way too large if μ is chosen just a bit too small. We therefore prefer the residual-based criterion.

Another important piece of information conveyed by Fig. 4 is that choosing δ and ϵ far away leads to other problems. Firstly, if δ is too large with respect to ϵ , and using the residual-based criterion (19) only, we risk that the coefficient norm is high: $\|\mathbf{c}^\epsilon\| = \mathcal{O}(\delta/\epsilon)$. E.g., in Fig. 4 we see that the choice of $\epsilon = 10^{-12}$, $\delta = 10^{-3}$ and the residual-based criterion leads to the optimal value optimal $N = 150$, but with a coefficient norm of $\mathcal{O}(10^9)$. Secondly, if ϵ is too large with respect to δ , we can never solve the underlying linear system at the accuracy required by δ . We conclude that ϵ and δ should be chosen comparable in size, and that ϵ should be small enough to allow solutions that reach the desired accuracy δ as N increases. We will choose $\epsilon 10^{-2}$ times than δ in the examples below.

3.4.2 Residual-based stopping criterion for all N

We adopt the residual-based criterion (19) and discard the coefficient-based criterion (21). However, since (19) is valid only in the regime of convergence, we need to augment the conditions to avoid returning an unacceptable solution for small N that happens to pass the test by chance. Following the example of Chebfun, we do so by evaluating in a few additional random points.

Given constants δ , δ' and ϵ , we accept a solution $\mathbf{c}_N^\epsilon$ if all of the following conditions are met:

1. $\|\mathbf{A}\mathbf{c}^\epsilon - \mathbf{b}\| \leq \delta\|\mathbf{b}\|$, and
2. $|f(t_i) - f_N(t_i)| \leq \delta'\|\mathbf{b}\|$, $i = 1, \dots, Q$, where $\{t_i\}_{i=1}^Q$ are random points in Ω .

The conditions are verified in order.

As argued in §3.3, we can discard a check on the coefficient norm if δ and ϵ are close. It is the second condition that, with some probability of success, catches the case where the residual is small in the M -norm, yet the approximation is large in between the sampling points. This is achieved by verifying for the desired accuracy in a few random points. For all examples in this paper, this condition is active only at very small values of N and M , but it may also detect problems with degenerate cases such as $f(x) = e^x + \phi_{100}(x)$ where $\phi_{100} \notin \text{span}\{\Phi_{99}\}$.

Some care is needed for the choice of δ' . On the one hand, since the second restriction might be more stringent than the first if δ' is too small. On the other hand, if δ' is too large we might approximate $f(x) = e^x + \delta\phi_{100}(x)$ inaccurately. In the examples below, it is safe to choose $\delta = \delta'$.

Figure 5: Illustration of the optimal value of N and the coefficient norm for varying δ and ϵ according to the stopping criterion of §3.4.2. Shown are the results for $f_1(x) = e^{\cos(8\pi)}$ (first two panels) and $f_2(x) = 10^6 f_1(x)$ (last two panels). The functions are approximated on $[0, 1/2]$ using Fourier series on $[0, 1]$. In both cases, the left panel depicts N and the right panel shows $\|\mathbf{x}\|$, both in logarithmic scales. Even if no sufficient solution is found, the search for optimal N stops at $N = 4096$. The exact algorithm used is Algorithm 3 with $\delta' = \delta$, $Q = 3$, $M \approx 2N$.

Figure 6: Approximation of a function with noise, $f(x) = e^x + \sigma \cos(2000\pi x)$, on $[0, 1/2]$ using Fourier series on $[0, 1]$. Shown are the residual in logarithmic scale, $\log_{10}(\|A_N \mathbf{c}_N^\epsilon - \mathbf{b}_N\|)$ (left panel), and the optimal value of N , $\log_{10}(N)$ (right panel), according to the criteria in §3.4.2 and for varying δ and σ . The exact algorithm used is Algorithm 3 with $\delta' = \delta$, $Q = 3$, $\epsilon = \delta/100$ and $M \approx 2N$. The coefficient norm is large whenever δ is smaller than σ . Best results are obtained for $\delta \approx \sigma$, shown along the anti-diagonal.

These criteria are illustrated in Fig. 5. We compare the optimal N and the coefficient norm $\|\mathbf{c}_N^\epsilon\|$ for $\epsilon, \delta \in [10^{-10}, 10^{-3}]$, with $\delta' = \delta$ and $Q = 3$. If ϵ is too large compared to δ , no optimal value of N can be found. In the figure, these solutions have the maximal value $N = 4096$, and they can be found in the top left triangles in the first and third panel. If δ is large and ϵ is small, the coefficient norm grows large. This is visible in the bottom right corners of panels 2 and 4.

Fig. 5 also illustrates the scale-invariance of the relative stopping criterion (19). The optimal value of N found for the functions $f_1(x) = e^{\cos(8\pi)}$ and $f_2(x) = 10^6 f_1(x)$ are comparable while the coefficient norm is clearly 10^6 times larger for the latter function.

Finally, in Fig. 6 we show results for a function with noise. We choose

$$f(x) = e^x + \sigma \cos(2000\pi x).$$

In this case we need to choose δ appropriately, ideally comparable in size to the noise level. If δ is too small, then we try to approximate the noise, which results in large values of N . We do not include a description in this paper of methods for the detection of noise levels, instead relying on the user to choose δ appropriately.

4 Algorithms

We aim to find the smallest value of N such that the approximation to a given function f with N degrees of freedom satisfies all conditions in §3.4.2:

$$N(f; \delta, \delta', \epsilon, Q) = \arg \min_{\substack{\|A_N \mathbf{c}_N^\epsilon - \mathbf{b}\| \leq \delta \|\mathbf{b}\| \\ |f(t_i) - f_N(t_i)| \leq \delta' \|\mathbf{b}\|, i=1, \dots, Q}} N. \quad (22)$$

Recall that we assume a stable sampling rate $M = \Theta(N, \theta)$ and ϵ sufficiently small, but close to δ . In all examples below, equispaced point samples were used and $M = 2N$ is a sufficiently large oversampling rate.

4.1 Brute force incremental strategy

The simplest algorithm attempts each value of N in order, until the optimal value $N(f; \delta, \delta', \epsilon, Q)$ is found. This approach yields the optimal value of N in all cases. Obviously, it is very computationally expensive, as it takes the solution of $\mathcal{O}(N_{\text{opt}})$ approximation problems to find the answer. Still, we have implemented this approach as a benchmark to compare with later on. The approach is specified in Algorithm 2 where, as in Algorithm 1, we use $\{A_i\}_{i=1}^\infty$ and $\{\mathbf{b}_i\}_{i=1}^\infty$ to denote the linear systems that need to be solved for a frame approximation using i coefficients. The number of samples M is implicit.

Algorithm 2 The incremental strategy

Input: Given $\delta, \delta', \epsilon, Q, \{A_i\}_{i=1}^\infty, \{\mathbf{b}_i = \mathcal{L}_i f\}_{i=1}^\infty$

Output: Optimal N

```

 $N = 1$ 
Solve  $A_N \mathbf{c}_N^\epsilon = \mathbf{b}_N$ 
while  $\|A_N \mathbf{c}_N^\epsilon - \mathbf{b}_N\| > \delta \|\mathbf{b}\|$  or  $|f(t_i) - f_N(t_i)| > \delta' \|\mathbf{b}_N\|, i = 1, \dots, Q$  do
     $N \leftarrow N + 1$ 
    Solve  $A_N \mathbf{c}_N^\epsilon = \mathbf{b}_N$ 
end while
return  $N$ 

```

Note that Algorithm 2 is exactly like Algorithm 1, but without the adaptive choice of weights. If the weighted approximation problem can be solved efficiently, then Algorithm 1 is not much more expensive than Algorithm 2.

4.2 Bisection strategy

The incremental strategy can be optimised with a relatively small modification that leads to substantially improved computational cost. We consider a method that consists of two stages. It is specified in Algorithm 3.

In the first stage, the method finds a size $\hat{N}$ that meets the stopping criterion, but which may not be optimal. This is achieved by doubling the number of degrees of freedom in every iteration until $\hat{N}$ is found. Hence, it takes $\mathcal{O}(\log N_{\text{opt}})$ steps where $N_{\text{opt}} = N(f; \delta, \delta', \epsilon, Q)$ is the optimal value. In practice, we also stop at a maximal value of N_{max} to ensure that the algorithm terminates.

In the second stage, the optimal step size is known to be between $N_h = \hat{N}$ and $N_l = \hat{N}/2$. The precise value is found using bisection. This takes an additional $\mathcal{O}(\log N_{\text{opt}})$ steps. Hence, in total the scheme converges to an answer based on the solution of $\mathcal{O}(\log N_{\text{opt}})$ approximation problems, rather than N_{opt} approximation problems in Algorithm 2.

The first doubling stage is also employed in the software packages Chebfun [22, 39,] and ApproxFun [34,]. The second stage is different from the aforementioned packages, in which the coefficients can be truncated based on their decay and based on plateau detection to determine a noise level [4,]. Plateau detection is impractical in our case, because we can only observe a plateau in the residuals for increasing number of degrees of freedom, and that is expensive to compute. In comparison, we have to solve a small number of additional approximation problems for values of N between N_h and $\frac{N_h}{2}$.

In the description of the algorithm and the estimation of its cost, we have assumed for simplicity that N can take all values between 1 and N_{max} . In our implementation, we have merely assumed that there is a sequence of sets Φ_j with strictly increasing cardinality $N(j)$. This is considerably

Algorithm 3 Doubling and bisection

Input: Given $\delta, \delta', \epsilon, Q, \{A_i\}_{i=1}^\infty, \{\mathbf{b}_i = \mathcal{L}_i f\}_{i=1}^\infty$ **Output:** Optimal N

```
 $N = 1$ 
Solve  $A_N \mathbf{c}_N^\epsilon = \mathbf{b}_N$ 
while  $\|A_N \mathbf{c}_N^\epsilon - \mathbf{b}_N\| > \delta \|\mathbf{b}\|$  or  $|f(t_i) - f_N(t_i)| > \delta' \|\mathbf{b}_N\|, i = 1, \dots, Q$  do ▷ Doubling
     $N = 2N$ 
    Solve  $A_N \mathbf{c}_N^\epsilon = \mathbf{b}_N$ 
end while
 $\hat{N} \leftarrow N, N_h \leftarrow N, N_l \leftarrow N/2$ 
while  $|N_h - N_l| > 0$  do ▷ Bisection
     $N = \lceil \frac{N_l + N_h}{2} \rceil$ 
    Solve  $A_N \mathbf{c}_N^\epsilon = \mathbf{b}_N$ 
    if  $\|A_N \mathbf{c}_N^\epsilon - \mathbf{b}_N\| > \delta \|\mathbf{b}_{\hat{N}}\|$  or  $|f(t_i) - f_N(t_i)| > \delta' \|\mathbf{b}_{\hat{N}}\|, i = 1, \dots, Q$  then
         $N_l \leftarrow N$ 
    else
         $N_h \leftarrow N$ 
    end if
end while
return  $N$ 
```

more general, but highly similar to the algorithm as stated here. The final computational cost does depend on the relation between $N(j)$ and j .

For example, in the case of a multivariate approximation with tensor product form, we may have that $N = \prod_{i=1}^d N_i$, using N_i degrees of freedom in dimension i . If each N_i is doubled in phase 1 of the algorithm, then N increases by a factor of 2^d . Alternatively, each N_i can be multiplied by $2^{1/d}$, and subsequently rounded to the nearest integer. This results in roughly an increase by a factor of 2 for the total number of degrees of freedom. Yet, even then, not all integers N can be reached – at least the prime numbers are always excluded. Care has to be taken in the implementation of the bisection algorithm to ensure that it terminates, and that only valid values of N are used.

Another example is the approximation using a function dictionary of the form

$$\Phi_N \triangleq \{T_k(x)\}_{k=0}^{N_1-1} \cup \{(1-x)^\alpha T_k(x)\}_{k=0}^{N_2-1}.$$

This dictionary is well suited to approximate functions with an algebraic singularity of order α at $x = 1$. Here, $N = N_1 + N_2$ is the sum of the lengths of the constituent sets (rather than their product as in the tensor product example). Depending on the application, it may be desirable to vary N_1 while keeping N_2 fixed, or one may vary N_1 and N_2 simultaneously. Our own implementation is configurable and allows both options, and in fact even allows the combination of tensor product dictionaries and weighted dictionaries, see Fig. 10.

4.3 On the monotonicity of convergence

The bisection strategy is quite clearly more performant than the incremental strategy. However, there is an important qualitative difference: it is not a priori guaranteed that the bisection approach actually yields the optimal value of N . Assuming that random sampling is sufficient to detect all degenerate cases, the incremental approach always yields the minimal N , since it is based on the enumeration of all possibilities. The bisection approach may miss the optimal value, if the decrease of the error is not monotonic.

Say the optimal value in an approximation problem lies between $N_l = 16$ and $N_h = 32$, and say, $N_{\text{opt}} = 18$. If the bisection algorithm notices that convergence criteria are not satisfied for $N = (N_l + N_h)/2 = 24$, then it will continue to look for an optimal value in the interval $[25, 32]$. Whatever it returns will be larger than 18.

We explore the extent to which regularised frame-approximations such as those described in §2.2 and §2.4 yield monotonic convergence rates. In the absence of regularisation, monotonicity is often guaranteed. Indeed, if the bases Φ_N form a nested sequence, i.e., $\text{span } \Phi_{N-1} \subset \text{span } \Phi_N$, then the best approximation error decays monotonically. This is the case if Φ_N arises from the truncation of an infinite set.

Lemma 4.1. *Let $\{\Phi_N\}_{N=1}^\infty$ form a nested sequence that is ultimately dense in a Hilbert space H . Let $f \in H$, and let the orthogonal projection of f on the truncated space $H_N \triangleq \text{span } \Phi_N$ be $\mathcal{P}_N f$, then*

$$\|f - \mathcal{P}_N f\|_H \leq \|f - \mathcal{P}_{N-1} f\|_H. \quad (23)$$

Proof. The orthogonal projection in H onto H_N satisfies by construction

$$\|f - \mathcal{P}_N f\|_H \leq \|f - \mathcal{T}_N \mathbf{z}\|_H, \quad \forall \mathbf{z} \in \mathbb{C}^N.$$

Since $H_{N-1} \subseteq H_N$, this also holds for all $\mathcal{T}_{N-1} \mathbf{z}$, with $\mathbf{z} \in \mathbb{C}^{N-1}$, and more specifically for

$$\mathcal{P}_{N-1} f = \arg \min_{\mathbf{z} \in \mathbb{C}^{N-1}} \|f - \mathcal{T}_{N-1} \mathbf{z}\|_H.$$

□

Monotonicity may be violated if the sequence is not nested, or if the discrete solution does not accurately reflect the best approximation. The latter may be a consequence of a coincidental bad choice of sampling points, but is affected more substantially by the regularisation.

Indeed, the lemma above does not hold for the regularised approximation spaces $H_N^\epsilon \subset H_N$ that result from the regularisation. Let

$$A_{M,N} = \{\ell_{k,M}(\phi_l)\}_{k,l=1}^{M,N}$$

be the discretisation matrix of a dictionary Φ_N using a sampling operator

$$f \rightarrow \{\ell_{k,M}(f)\}_{k=1}^M,$$

with singular value decomposition $A_{M,N} = U \Sigma V'$, where $U \in \mathbb{C}^{M \times M}$, and $V \in \mathbb{C}^{N \times N}$ are unitary, and $\Sigma \in \mathbb{R}^{M \times N}$ is a diagonal matrix with positive entries σ_k . Using a truncated SVD to regularise the system (by discarding all singular values smaller than a given threshold ϵ), we define H_N^ϵ as the span of the right singular vectors v_k corresponding to singular values σ_k larger than ϵ :

$$H_N^\epsilon \triangleq \text{span}\{\mathcal{T}_N v_k | \sigma_k > \epsilon\}. \quad (24)$$

Since H_{N-1}^ϵ is not necessarily included in H_N^ϵ , even when $H_{N-1} \subset H_N$, a decrease in approximation error is not guaranteed. However, using the analysis in [1] and [2], we can make the following statements. We consider first the computation of best approximations, i.e., based on orthogonal projection, in combination with regularization.

Theorem 4.2. *Let $f \in H$, and the orthogonal projection of f on H_N^ϵ be $\mathcal{P}_N^\epsilon f$, i.e., we choose the sampling functionals $l_{m,M}(f) = \langle f, \phi_m \rangle$ and set $M = N$. Then*

$$\|f - \mathcal{P}_N^\epsilon f\|_H \leq \|f - \mathcal{P}_{N-1}^\epsilon f\|_H + \sqrt{\epsilon} \|\mathbf{c}_{N-1}^\epsilon\|_{\ell^2(\mathbb{C}^{N-1})}, \quad (25)$$

where $\mathbf{c}_{N-1}^\epsilon$ is the coefficient vector of $\mathcal{P}_{N-1}^\epsilon f$.

Proof. The error bound (8) holds for any collection of sampling functionals $\{l_{m,M}\}_{m=1}^M$. An orthogonal projection corresponds more specifically to sampling via inner products with the frame elements, i.e., $l_{m,M}(f) = \langle f, \phi_m \rangle$, as well as $M = N$. A specific bound for orthogonal projection with SVD-based regularization is given in [1, Theorem 13]:

$$\|f - \mathcal{P}_N^\epsilon f\|_H \leq \|f - \mathcal{T}_N \mathbf{z}_N\|_H + \sqrt{\epsilon} \|\mathbf{z}_N\|_{\ell^2(\mathbb{C}^N)}, \quad \forall \mathbf{z}_N \in \mathbb{C}^N. \quad (26)$$

Figure 7: Illustration of Theorems 4.2 and 4.3 using a Fourier extension frame approximation for $f(x) = 1$ (left panels) and $f(x) = \exp(\cos(5x))$ (right panels), with truncation parameter $\epsilon = 10^{-5}$. Top row: approximation error $\|f - \mathcal{P}_N^\epsilon f\|_H$ (red squares) and the upper bound (25) (blue dots). Bottom row: Approximation error $\|f - \mathcal{P}_{M',N}^\epsilon f\|_H$ (squares) and the upper bound (27), here simplified to be $\|f - \mathcal{P}_{M',N-1}^\epsilon f\|_H + \epsilon \|c\|_{\ell^2(\mathbb{C}^{N-1})}$ (dots). Convergence is largely monotonic, but not completely, and no convergence occurs below the regularization threshold.

Since this holds for any $\mathbf{z}_N \in \mathbb{C}^N$, we can choose any $\mathbf{z}_N$ for which the last element is zero and obtain

$$\|f - \mathcal{P}_N^\epsilon f\|_H \leq \|f - \mathcal{T}_{N-1} \mathbf{z}_{N-1}\|_H + \sqrt{\epsilon} \|\mathbf{z}_{N-1}\|_{\ell^2(\mathbb{C}^{N-1})}, \quad \forall \mathbf{z}_{N-1} \in \mathbb{C}^{N-1},$$

and more specifically

$$\|f - \mathcal{P}_N^\epsilon f\|_H \leq \|f - \mathcal{P}_{N-1}^\epsilon f\|_H + \sqrt{\epsilon} \|\mathbf{c}\|_{\ell^2(\mathbb{C}^{N-1})},$$

where $\mathbf{c}$ is the coefficient vector of $\mathcal{P}_{N-1}^\epsilon f$. $\square$

Therefore, bisection can not guarantee the most optimal solution, but it will give one with an error that is close to the optimal error. Note that, compared to (8), the result of Theorem 4.2 only demonstrates $\mathcal{O}(\sqrt{\epsilon})$ accuracy. This is due not to the projection and the corresponding choice of inner product samples, but to the choice $M = N$. A similar result can be stated for the bound (8), which applies to more general samples, and which does show accuracy on the order of $\mathcal{O}(\epsilon)$ based on oversampling at the stable sampling rate.

Theorem 4.3. Assume that M and M' satisfy the stable sampling rate for N and $N - 1$ respectively, i.e., $M \geq \Theta^\epsilon(N, \theta)$ and $M' \geq \Theta^\epsilon(N - 1, \theta)$. Denote by $\mathcal{P}_{M,N}^\epsilon f$ the ϵ -regularised solution of the discrete least squares problem, and let A' be the lower bound constant in (13), then

$$\|f - \mathcal{P}_{M,N}^\epsilon f\|_H \leq \|f - \mathcal{P}_{M',N-1}^\epsilon f\|_H + \frac{\theta}{A'} (\|f - \mathcal{P}_{M',N-1}^\epsilon f\|_M + \epsilon \|\mathbf{c}_{N-1}^\epsilon\|_{\ell^2(\mathbb{C}^{N-1})}), \quad (27)$$

where $\mathbf{c}_{N-1}^\epsilon$ is the coefficient vector of $\mathcal{P}_{M',N-1}^\epsilon f$.

Proof. The stable sampling rate is satisfied, so by (12)

$$\|f - \mathcal{P}_{M,N}^\epsilon f\|_H \leq \|f - \mathcal{T}_N \mathbf{z}_N\|_H + \frac{\theta}{A'} (\|f - \mathcal{T}_N \mathbf{z}_N\|_M + \epsilon \|\mathbf{z}_N\|), \quad \forall \mathbf{z}_N \in \mathbb{C}^N.$$

Since this holds for any $\mathbf{z}_N \in \mathbb{C}^N$, we can choose a vector $\mathbf{z}_{N-1} \in \mathbb{C}^{N-1}$ and append a zero element at the end. This leads to

$$\|f - \mathcal{P}_{M,N}^\epsilon f\|_H \leq \|f - \mathcal{T}_{N-1} \mathbf{z}_{N-1}\|_H + \frac{\theta}{A'} (\|f - \mathcal{T}_{N-1} \mathbf{z}_{N-1}\|_M + \epsilon \|\mathbf{z}_{N-1}\|), \quad \forall \mathbf{z}_{N-1} \in \mathbb{C}^{N-1}.$$

We obtain the result by choosing $\mathbf{z}_{N-1}$ equal to $\mathbf{c}_{N-1}^\epsilon$, the coefficient vector of $\mathcal{P}_{M',N-1}^\epsilon f$. $\square$

Figure 8: Adaptive approximation of $f(x) = \cos(px)$ on $[-1, 1]$ using Fourier series on $[-2, 2]$ with $\delta = \epsilon = 10^{-10}$ and $Q = 3$. Left: The difference $N_{\text{bis}} - N_{\text{inc}}$ between the optimal value for N obtained by the bisection algorithm compared to the brute-force incremental algorithm, as a function of p . The difference is very small. Right: the ratio of the computation time of the adaptive bisection approach over the time it takes to approximate f once using the optimal number of degrees of freedom. To avoid the influence of the random point evaluation, the median is taken over 7 experiments for every data point.

The theorem shows that the sequence is almost monotonic. Lack of monotonicity can be caused by the two rightmost terms in (27). The first of those, $\|f - P_{M', N-1}^\epsilon f\|_M$, is expected to be on the same order as $\|f - \mathcal{P}_{M', N-1}^\epsilon f\|_H$ in the regime of convergence, i.e., for sufficiently large N . The second term may cause minor jumps in the approximation error on the order of ϵ .

The behaviour of the H-norm error and the upper bounds of Theorems 4.2 and 4.3 are illustrated in Fig. 7. Theorem 4.2 is illustrated in the top row for the constant function $f(x) = 1$ (left panel) and a generic smooth function (right panel), using $\epsilon = 10^{-5}$. Theorem 4.3 is illustrated in the bottom row for the same two functions. In all cases, the function is approximated using a Fourier extension frame.

Consider the constant function first (left panels). The regularisation yields a sequence of function spaces $H_1^\epsilon, H_2^\epsilon, \dots$. Only the first of these actually contains constant functions, the subsequent spaces do not. Hence, the error for $f(x) = 1$ is machine precision at $N = 1$, but is significantly larger afterwards. The upper bounds of the theorems are on the order of $\sqrt{\epsilon}$ and ϵ respectively. The true approximation error stays below this bound, but as expected no convergence can be achieved beyond the accuracy determined by the regularisation parameter.

The results for the smooth function show nearly monotonic behaviour, though not exactly, both for the upper bounds and the true residuals. The bounds allow for jumps on the order of $\sqrt{\epsilon}$ (top) and ϵ (bottom). For Theorem 4.3, we have simplified the bound to $\|f - \mathcal{P}_{M', N-1}^\epsilon f\|_H + \epsilon \|\mathbf{c}\|_{\ell^2(\mathbb{C}^{N-1})}$. Jumps are still present, but less visible and in the range allowed by the bound.

4.4 Numerical experiments

4.4.1 A univariate Fourier extension example

We start with a simple univariate example. In Fig. 8, we approximate $f(x) = \cos(px)$ on $[-1, 1]$ using Fourier series on $[2, 2]$, i.e., using Fourier extension. Larger values of p render the function more oscillatory, and hence should lead to larger optimal values of N . We choose $p \in [0, 500]$, $\delta = 10^{-10}$, $\epsilon = 10^{-12}$, $Q = 3$.

We compare the optimal N given by the incremental approach and the bisection approach. The bisection approach results in a slight overestimation of the true optimal N , obtained by the greedy approach. As seen in the figure, the difference is in this example never larger than 3 degrees of freedom. For $p = 500$, the optimal value of N is 642.

We also compare the computational cost of the bisection approach with that of a single function approximation using the optimal number of degrees of freedom. Let t_{biss} be the time it takes to find N_{opt} , and let $t_{N_{\text{opt}}}$ be the time to compute an approximation with N_{opt} degrees of freedom. The figure shows the ratio

$$\frac{t_{\text{biss}}}{t_{N_{\text{opt}}}}.$$

Figure 9: Adaptive approximation of $f(x, y, z) = \exp(y \cos(x) \sin(z - 1))$ on the ball with radius 0.3 and centre $(0.5, 0.5, 0.5)$ using an extension frame with B-splines of order 5 on the bounding box $[0, 1]^3$, $\delta = 10^{-6}$, $\epsilon = 10^{-8}$ and $Q = 3$. Top: evolution of residual (left) and coefficient norm (right) during the adaptive process. The doubling step occurs once, the following iterations ensure that the optimal value was the one after doubling. Bottom left and right are respectively the optimal approximant ($\mathbf{N} = (14, 14, 14)$) and the $\log_{10}$ of its uniform error

We expect this ratio to be bounded by $\mathcal{O}(\log t_{N_{\text{opt}}})$ since the bisection approach computes only logarithmically many additional function approximations. In this example, the computational cost is cubic in N for each approximation problem. The approximations with a large N therefore take a much larger fraction of the total clock time than those with small N . The implication is that the logarithmic growth in the number of iterations is not actually visible in the figure.

4.4.2 Three-dimensional non-spectral approximation on a non-rectangular domain

The emphasis in the presentation of the paper lies with spectral approximation. However, the algorithm for adaptive approximation is not necessarily restricted to bases or frames with spectral approximation properties. It may also be applied to higher-dimensional problems.

In order to illustrate these properties, we approximate the smooth function

$$f(x, y, z) = \exp(y \cos(x) \sin(z - 1)) \quad (28)$$

in three dimensions with B-splines of order 5 in Fig. 9. The approximation domain is the ball with centre $(0.5, 0.5, 0.5)$ and radius 0.3, so this is again an example of an extension frame. See [16] for more details on fast algorithms for spline extension based on the AZ algorithm.

In the figure, we show the evolution of the residual error and coefficient norm during the bisection algorithm. Only one doubling step is present in which the size of the coefficient array grows from $(8, 7, 7)$ to $(14, 14, 14)$. Thereafter, the bisection iterations conclude that $(14, 14, 14)$ is the optimal choice.

4.4.3 Two-dimensional spectral approximation of a singular function on a non-rectangular domain

We conclude with a more involved example that combines several difficulties. We choose a function that has an algebraic point singularity in 2-D,

$$f(x, y) = \cos(p\pi(x + y)) + \sqrt{x^2 + y^2} \sin(1 + p\pi(x + y)). \quad (29)$$

Figure 10: We approximate the singular function (29) on the disk with centre $[0, 0]$ and radius 0.9 using the weighted Fourier Extension frame (30). Upper left: The optimal number of degrees of freedom found by the adaptive approach using $\delta = 10^{-6}$ and $\epsilon = 10^{-8}$. Upper right: the ratio of the timings of the adaptive approach w.r.t. the timings of an approximation using the same number degrees of freedom as the optimal solution.. Lower left and right: the approximation and the $\log_{10}$ of the uniform error of the approximant found with $p = 3, \delta = 10^{-6}, \epsilon = 10^{-8}$.

As before, the parameter p controls the oscillatory nature of the function. Furthermore, we set out to approximate this function on a non-rectangular domain, a disk with radius 0.9. To that end, first we choose a Fourier extension frame by restricting a Fourier series $\Psi \triangleq \{\psi_{k,l}\}_{k,l \in \mathbb{Z}}$ on $[-1, 1]^2$ to the disk. Next, we use that frame for the smooth and singular parts of the function:

$$\Phi \triangleq \Psi \cup \sqrt{x^2 + y^2} \Psi = \{\psi_{k,l}(x, y)\}_{k,l \in \mathbb{Z}} \cup \{\sqrt{x^2 + y^2} \psi_{k,l}(x, y)\}_{k,l \in \mathbb{Z}}. \quad (30)$$

We define a truncation Φ_N with increasing length N in terms of the truncated frames Ψ_{N_1} and $\sqrt{x^2 + y^2} \Psi_{N_2}$, with $N = N_1 + N_2$. When increasing N , we alternate between increasing N_1 and N_2 , i.e., $N_1 = 1 + \lceil \frac{N-1}{2} \rceil$ and $N_2 = 1 + \lfloor \frac{N-1}{2} \rfloor$. In turn, Ψ_{N_i} has size $N_i^{(1)} \times N_i^{(2)}$ and we similarly alternate between increasing $N_i^{(1)}$ and $N_i^{(2)}$ when increasing N_i . The first set Ψ_1 has size 1×1 .

In Fig. 10, we have repeated the experiment of Fig. 8. As before, we increase p to obtain adaptive approximations with higher N and compare timings and optimal N . The example is completed with a plot of the optimal approximant (bottom left) and the $\log_{10}$ of the uniform error (bottom right) for $p = 3, \epsilon = 10^{-8}, \delta = 10^{-6}, Q = 3$.

References

- [1] B. Adcock and D. Huybrechs. Frames and numerical approximation. *SIAM Rev.*, 61(3):443–473, 2019.
- [2] B. Adcock and D. Huybrechs. Frames and numerical approximation II: generalized sampling. *arXiv e-prints*, page arXiv:1802.01950, Jul 2020.
- [3] B. Adcock, D. Huybrechs, and J. Martin-Vaquero. On the numerical stability of Fourier extensions. *Found. Comput. Math.*, 14(4):635–687, 2014.
- [4] J. L. Aurentz and L. N. Trefethen. Chopping a Chebyshev series. *ACM Trans. Math. Software*, 43(4):1–21, jan 2017.

- [5] J. P. Boyd. *Chebyshev and Fourier spectral methods*. Courier Dover Publications, Mineola, NY, 2001.
- [6] J. P. Boyd. A comparison of numerical algorithms for Fourier extension of the first, second, and third kinds. *J. Comput. Phys.*, 178(1):118–160, 2002.
- [7] J. P. Boyd. *Solving Transcendental Equations: The Chebyshev Polynomial Proxy and Other Numerical Rootfinders, Perturbation Series, and Oracles*, volume 139. SIAM, Philadelphia, 2014.
- [8] J. P. Boyd and J. R. Ong. Exponentially-convergent strategies for defeating the Runge phenomenon for the approximation of non-periodic functions, part I: Single-interval schemes. *Communications in Computational Physics*, 5(2-4):484–497, 2009.
- [9] P. D. Brubeck, Y. Nakatsukasa, and L. N. Trefethen. Vandermonde with arnoldi. Technical Report arXiv:1911.09988, University of Oxford, Oxford, 2019.
- [10] 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):1094–1125, 2007.
- [11] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20(1):33–61, 1999.
- [12] O. Christensen. *An Introduction to Frames and Riesz Bases*. Birkhauser, Switzerland, 2 edition, 2016.
- [13] A. Cohen, W. Dahmen, and R. DeVore. Adaptive wavelet methods for elliptic operator equations: Convergence rates. *Math. Comput.*, 70(233):27–76, may 2000.
- [14] A. Cohen, W. Dahmen, and R. DeVore. Adaptive wavelet methods II—beyond the elliptic case. *Found. Comput. Math.*, 2(3):203–202, aug 2002.
- [15] V. Coppé. AdaptiveFrameExperiments.jl. *GitHub repository*, 2020.
- [16] V. Coppé and D. Huybrechs. Efficient function approximation on general bounded domains using splines on a cartesian grid. Technical Report arXiv:1911.07894, KU Leuven, Leuven, 2019.
- [17] V. Coppé and D. Huybrechs. Efficient function approximation on general bounded domains using wavelets on a cartesian grid. Technical Report arXiv:2004.03537, KU Leuven, Leuven, 2020.
- [18] V. Coppé, D. Huybrechs, R. Matthysen, and M. Webb. The AZ algorithm for least squares systems with a known incomplete generalized inverse. *SIAM J. Matrix Anal. Appl.*, page to appear, 2020.
- [19] S. Dahlke, M. Fornasier, and T. Raasch. Adaptive frame methods for elliptic operator equations. *Adv. Comput. Math.*, 27(1):27–63, may 2007.
- [20] I. Daubechies. *Ten Lectures of Wavelets*, volume 61. SIAM, Philadelphia, 1992.
- [21] P. J. Davis. *Interpolation and approximation*. Courier Corporation, New York, 1975.
- [22] T. A. Driscoll, N. Hale, and L. N. Trefethen. *Chebfun guide*. Pafnuty Publications, Oxford, 2014.
- [23] J. S. Geronimo and K. Liechty. The fourier extension method and discrete orthogonal polynomials on an arc of the circle. *Adv. Math.*, 365:107064, 2020.

- [24] J. W. Gibbs. Fourier’s series. *Nature*, 59(1522):200, 1898.
- [25] N. Gruberger and D. Levin. Two algorithms for periodic extension on uniform grids. *Numer. Algorithms*, pages 1–20, 2020.
- [26] D. Huybrechs. On the Fourier extension of nonperiodic functions. *SIAM J. Numer. Anal.*, 47(6):4326–4355, 2010.
- [27] J. Kovačević and A. Chebira. An introduction to frames. *Foundations and Trends® in Signal Processing*, 2(1):1–94, oct 2007.
- [28] L. Lovisolo and E. A. da Silva. Frames in signal processing. *Academic Press Library in Signal Processing*, 1:561–590, jan 2014.
- [29] M. Lyon. A fast algorithm for Fourier continuation. *SIAM J. Sci. Comput.*, 33(6):3241–3260, 2011.
- [30] M. Lyon. Sobolev smoothing of SVD-based Fourier continuations. *Appl. Math. Lett.*, 25(12):2227–2231, 2012.
- [31] J. C. Mason and D. C. Handscomb. *Chebyshev Polynomials*. Chapman and Hall/CRC, Boca Raton, 2003.
- [32] R. Matthysen and D. Huybrechs. Fast algorithms for the computation of Fourier extensions of arbitrary length. *SIAM J. Sci. Comput.*, 38(2):A899–A922, 2016.
- [33] R. Matthysen and D. Huybrechs. Function approximation on arbitrary domains using Fourier extension frames. *SIAM J. Numer. Anal.*, 56(3):1360–1385, 2018.
- [34] S. Olver et al. ApproxFun v0.10.0 Julia package. URL <https://github.com/JuliaApproximation/ApproxFun.jl>, 2019.
- [35] R. B. Platte, L. N. Trefethen, and A. B. Kuijlaars. Impossibility of fast stable approximation of analytic functions from equispaced samples. *SIAM Review*, 53(2):308–318, 2011.
- [36] D. Potts and N. Van Buggenhout. Fourier extension and sampling on the sphere. In *2017 International Conference on Sampling Theory and Applications (SampTA)*, pages 82–86, Estonia, 2017. IEEE, Web.
- [37] R. Stevenson. Adaptive solution of operator equations using wavelet frames. *SIAM J. Numer. Anal.*, 41(3):1074–1100, jan 2003.
- [38] G. Szegő. *Orthogonal polynomials*, volume 23. American Mathematical Soc., New York, 1939.
- [39] L. N. Trefethen. *Approximation theory and approximation practice*, volume 128. SIAM, Philadelphia, 2013.
- [40] M. Webb, V. Coppé, and D. Huybrechs. Pointwise and uniform convergence of Fourier extensions. *Constr. Approx.*, pages 1–37, 2019.