UNIVERSITY OF WISCONSIN-MADISON
CENTER FOR THE MATHEMATICAL SCIENCES

Computing the Sobolev regularity of refinable functions by the Arnoldi Method

Amos Ron, Zuowei Shen and Kim-Chuan Toh

Computer Science Department
University of Wisconsin-Madison
1210 West Dayton Street
Madison, Wisconsin 53706, USA
amos@cs.wisc.edu

Department of Mathematics
National University of Singapore
10 Kent Ridge Crescent
Singapore 119260
matzuows@math.nus.edu.sg

Department of Mathematics
National University of Singapore
10 Kent Ridge Crescent
Singapore 119260
mattohkc@math.nus.edu.sg

ABSTRACT

The recent paper [RS1] provides a complete characterization of the L_2-smoothness of a refinable function in terms of the spectrum of an associated operator. Based on this theory, we devise in this paper a numerically stable algorithm for calculating that smoothness parameter, employing the deflated Arnoldi method to this end. The algorithm is coded in `Matlab`, and details of the numerical implementation are discussed, together with some of the numerical experiments. The algorithm is designed to handle large masks, as well as masks of refinable function with unstable shifts. This latter case is particularly important, in view of the recent developments in the area of wavelet frames.

AMS (MOS) Subject Classifications: Primary 42C15, Secondary 39B99, 46E35

Key Words: refinable functions, wavelets, smoothness, regularity, transition operators, transfer operators, Arnoldi's method

This work was supported by the National Science Foundation under Grants DMS-9626319 and DMS-9872890, by the U.S. Army Research Office under Contracts DAAH04-95-1-0089 and DAAG55-98-1-0443, by the National Institute of Health and by the Strategic Wavelet Program Grant from the National University of Singapore.

Computing the Sobolev regularity of refinable functions by the Arnoldi Method

Amos Ron, Zuowei Shen and Kim-Chuan Toh

1. Introduction

We are interested in the computation of the smoothness parameter of refinable functions. *Refinable functions* (known also as 'scaling functions') are solutions of special functional equations that are known as *refinement equations*. The refinement equation expresses a dilate of the solution as the convolution product of that solution with a discrete kernel, the latter being known as the *mask* (cf. (2.2) for the precise definition). The smoothness of refinable functions is important in two subareas of analysis. In the area of *subdivision algorithms*, it determines the smoothness of the limit curve/surface of the subdivision process; in the area of *wavelets*, the smoothness of the refinable function is passed on to all wavelet systems that are derived from it (via the *multiresolution analysis* vehicle). In most practical cases, the refinable function is not known explicitly, and the available information consists, primarily, of the mask. Therefore, the determination of the smoothness of the solution from properties of the mask is one of the key problems in the above-mentioned areas.

Our efforts in this paper are focused on the study of the above problem via the *transfer/transition operator* approach. The analysis of the regularity of refinable functions in terms the transfer operator was developed by several authors (cf., e.g., [D], [DD], [E], and [V] for the univariate case, [RiS1], [CGV], [J], [LMW], [RS1] and [R2] for the multivariate case). In the L_2-case, the regularity estimates are in terms of a specific eigenpair of an associated operator, hence seem to be computationally feasible. However, while the smoothness parameter of some examples was successfully computed by some authors (see e.g. [HJ] and [RS1]), there has not been (to best of our knowledge) a reliable (i.e., robust) numerical algorithm that works without significant restrictions on the mask.

Our method is based on the characterizations of the L_2-smoothness parameter given in [RS1], a detailed account of which is given in §2. For the discussion here, it suffices to note that the characterization is given in terms of the restriction of a certain linear operator (the *transfer operator*) to a finite-dimensional invariant subspace H of it (the elements of H are trigonometric polynomials). In order to compute the smoothness using this approach, one has to overcome four different obstacles, two of which are of theoretical nature and the other two of numerical nature. First, one needs a characterization of the space H, a characterization that applies to a wide range of refinement equations; specifically, one should avoid restrictions on the refinement equations that either cannot be verified numerically, or exclude examples of interest. Second, the characterization of the invariant space H must be computationally verifiable; we found that in most practical cases it is not feasible to compute a basis for H, hence one must have an alternative method for checking whether a given function belongs to that invariant space. That alternative method employs a superspace H_0 of H which is also an invariant subspace of the transfer operator, and which has an easily computable basis. The algorithm then finds in H_0 eigenvectors of the transfer operator, and uses a subtle criterion to determine whether the eigenvector found also lies in H. The success of this approach relies on the ability to recover accurately many eigenvectors, and not only few dominant ones. Thus, our third obstacle is the necessity of choosing and implementing carefully the eigen-solver. Fourth, a direct implementation of the theory converts 'small' problems (measured, say, in terms of the support of the mask) to a huge numerical mess, unless properly approached. For example, the matrix involved in computing one of the bivariate interpolatory refinable functions constructed in [RiS1] has an order of about 4×10^3, leading, thereby, to a numerically prohibitive eigenproblem.

We present our algorithm and its implementation in four stages. In the first (§2), we survey the results of [RS1] on the regularity of refinable functions, results that serve as the main stimulus for the present endeavor. As is seen there, the characterization of [RS1] can be implemented in many different ways, and we carefully devise in the second stage (§3) what we consider to be the 'winning algorithm' (designed to be fast for the average problem and robust for other cases). The algorithm requires a supplementary stable method for computing eigenvalues of linear operators. In the third stage of the presentation (§4), we describe a variation of the Arnoldi method [A] that is used to that end, and provide a rough sketch of our `Matlab` code. We document in §5 a sample of the numerical experiments. Finally, proofs of some results used in §3 are given in §6.

One must keep in mind that it is rather hard to devise a good universal numerical algorithm for this problem, since the numerical challenge in computing the smoothness has many, conflicting, faces. For example, in the construction of compactly supported bivariate interpolatory subdivision schemes, as well as in the related construction of certain orthogonal and biorthogonal refinable functions (see e.g.[DGL], [DDD], [CD], [CS], [RiS1-2], [JRS], [HL] [HJ], [KS], [BW]), one expects to have a relatively large mask, hence one has to cope with the sheer size of the problem. In contrast, in the theory of *wavelet frames*, and in the subsequent constructions of tight wavelet frames and bi-frames, (cf. [RS2-5] and in particular [GR]), good wavelet systems (e.g., tight frames), are derived from a multiresolution analysis based on a refinable function with unstable shifts. While that refinable function may be very attractive for applications (having many alternative properties such as high smoothness, good approximation order and small support), the problem of finding its smoothness without the stability assumption is a theoretical challenge (which was overcome for the first time in [RS1]), and is also a computational challenge.

2. The Sobolev regularity of refinable functions

Since the main objective of this paper is to convert (some of) the results in [RS1] from theory to practice, we naturally review first the pertinent results of that paper. The presentation here is confined to the setup of the present paper. We only consider here *scalar* refinable functions (PSI case) in one or two variables whose refinement masks are finitely supported. (The characterizations of [RS1] apply to the vector (FSI) case, to any number of dimensions, and do not assume the mask to be finitely supported.) A complete list of the assumptions made in this paper is provided in the sequel.

Let s be a $d \times d$ integer matrix that satisfies

$$s^* s = \lambda I, \tag{2.1}$$

for some $\lambda > 1$. We refer to such matrix s as a **dilation matrix** or, more precisely, as an *isotropic dilation matrix.* Let ϕ be a compactly supported L_2-function in d variables (or, more generally, a compactly supported distribution). We say that ϕ is **refinable with respect to the dilation matrix** s, if there exists a finitely supported sequence a such that

$$\phi(x) = |\det s| \sum_{j\in\mathbb{Z}^d} a(j)\phi(sx-j), \quad x \in \mathbb{R}^d. \tag{2.2}$$

The equivalent formulation of this condition on the Fourier domain is

$$\widehat{\phi}(s^*\cdot) = \widehat{a}\widehat{\phi}, \tag{2.3}$$

with $\widehat{a}$ the symbol of the sequence a, i.e.,

$$\widehat{a}(\omega) = \sum_{j\in\mathbb{Z}^d} a(j) \exp(-ij\omega).$$

The sequence a (as well as its symbol $\widehat{a}$) is called the **refinement mask** of ϕ. The L_2**-regularity parameter** $\alpha(\phi)$ of ϕ is defined by

$$\alpha(\phi) := \sup\{\alpha \in \mathbb{R} : \ \phi \in W_2^\alpha(\mathbb{R}^d)\}.$$

Here,

$$W_2^\alpha(\mathbb{R}^d)$$

is the usual Sobolev space.

As it turns out, the regularity of ϕ is determined by properties of a related function known as **the autocorrelation function** $\phi^{\#}$ of ϕ, and which is defined as follows:

$$\phi^{\#} : t \mapsto \int_{\mathbb{R}^d} \phi(x)\phi(x-t)\, dx.$$

It is easy to see that the Fourier transform of $\phi^{\#}$ is $|\widehat{\phi}|^2$. Hence, $\phi^{\#}$ is refinable with mask

$$\widehat{b} := |\widehat{a}|^2.$$

The 2π-periodization of the Fourier transform of $\phi^{\#}$, i.e., the $L_1(\mathbb{T}^d)$-function

$$\Phi := \sum_{j\in\mathbb{Z}^d} |\widehat{\phi}(\cdot + 2\pi j)|^2, \tag{2.4}$$

plays a pivotal role in our discussion. Since $\phi^{\#}$ is compactly supported (by the fact that ϕ is), the Poisson summation formula implies that Φ is a trigonometric polynomial whose spectrum (i.e., frequencies) all belong to the set

$$(\operatorname{supp} \phi^{\#}) \cap \mathbb{Z}^d = (\operatorname{supp} \phi - \operatorname{supp} \phi) \cap \mathbb{Z}^d. \tag{2.5}$$

Next, we define the transfer operator. Let

$$\Gamma \tag{2.6}$$

be any representer set of the quotient group $2\pi(s^{*-1}\mathbb{Z}^d/\mathbb{Z}^d)$. The **transfer** or **transition operator** T is defined as:

$$T : L_2(\mathbb{T}^d) \mapsto L_2(\mathbb{T}^d) : f \mapsto \sum_{\gamma\in\Gamma} (\widehat{b}f)(s^{*-1} \cdot + \gamma). \tag{2.7}$$

For example, if the spatial dimension is 1, and the dilation is dyadic (i.e., $s = 2I$), Γ can be chosen as $\{0, \pi\}$, and T becomes

$$(Tf)(\omega) = (\widehat{b}f)(\frac{\omega}{2}) + (\widehat{b}f)(\frac{\omega}{2} + \pi).$$

As was already alluded to in the introduction, the L_2-smoothness of ϕ is characterized by the spectral radius of the restriction of T to a certain invariant space H of it, with H finite dimensional and consisting of trigonometric polynomials. In general, the space H does not have a simple structure. As a first step, we would like to construct a finite dimensional superspace of H (made also of trigonometric polynomials) which on one hand will be T-invariant, while, on the other hand, will have a simple structure.

To this end, let

$$\mathcal{Z}_\phi \;:=\; \{j \in \mathbb{Z}^d \;:\; \|j\|_2 \le r\},$$

where r is any (fixed) number larger than or equal to

$$\frac{1}{\sqrt{\lambda}-1} \max\{\|j\|_2 \;:\; b_j \neq 0\}$$

with λ defined in (2.1) and with (b_j) being the mask coefficient of the autocorrelation function. Then, since

$$\|s^{*-1}x\|_2 = \frac{1}{\sqrt{\lambda}}\|x\|_2,$$

the space

$$H_\phi \tag{2.8}$$

of all trigonometric polynomials whose band lies in that set (i.e., the space spanned by the exponentials $\exp(ij\cdot)$, $j \in \mathcal{Z}_\phi$) is a T-invariant subspace, and that all eigenvectors of T that are trigonometric polynomials must lie in H_ϕ. Moreover, given *any* trigonometric polynomial f, we have that $T^k f \in H_\phi$ for all sufficiently large k (see [LLS1]). This last property implies that H_ϕ must contain each eigenvector f of T, provided that f is a trigonometric polynomial, and that its associated eigenvalue is non-zero. We use these basic facts in the sequel without further notice.

Theorem 2.2 of [RS1] states that the regularity parameter $\alpha(\phi)$ of ϕ is,

$$\alpha(\phi) = -\frac{\log_\lambda \rho}{2},$$

where λ is given by (2.1), and $\rho = |\mu|$ with μ an eigenvalue of the transfer operator (and with the associated eigenvector being a trigonometric polynomial). Hence, the key to the numerical computation of the regularity parameter $\alpha(\phi)$ is to compute the eigenpair (μ, f_μ) of T. We will describe in this paper a reliable and numerically stable algorithm that computes this eigenpair of T, thereby finds $\alpha(\phi)$. The algorithm is based on the characterization of ρ as the spectral radius of the restriction of T to H, with H a certain T-invariant subspace (that is defined below) of H_ϕ. One should note that H, as any subspace of H_ϕ, consists of trigonometric polynomials, each of which can be finitely represented in terms of its Fourier coefficients. However, in order to compute ρ directly from the above description, we also need a robust method for constructing a basis for H; since the methods we could find for constructing a basis for H are highly unstable, we will study the action of T on the larger space H_ϕ, and we will actually find ρ by other means. But, first, we recall the description of the space H from [RS1].

The space H is defined as $H := H_\phi \cap I_\phi$, with I_ϕ an ideal of trigonometric polynomials defined below. To this end, we set

$$\Pi$$

for the space of all d-variate (algebraic) polynomials, and Π_ϕ for the following subspace of it:

$$\Pi_\phi := \{p \in \Pi : \sum_{j \in \mathbb{Z}^d} p(j)\phi^{\#}(\cdot - j) \in \Pi\}.$$

Definition 2.9: the ideal I_ϕ. Let ϕ be a compactly supported L_2-function with $\widehat{\phi}(0) \neq 0$. Let $\phi^{\#}$ be the autocorrelation function of ϕ and let Φ be the 2π-periodization of the Fourier transform of $\phi^{\#}$ as given in (2.4). The ideal I_ϕ is the collection of all trigonometric polynomials (in $L_2(\mathbb{T}^d)$) f that satisfy:
(i) $f/\Phi \in L_\infty(\mathbb{T}^d)$.
(ii) f is annihilated by Π_ϕ in the sense that $p(-iD)f(0) = 0$, for all $p \in \Pi_\phi$.

With the definition of I_ϕ, the results of [RS1] that are used in the present paper for computing the regularity parameter $\alpha(\phi)$ are summarized as follows:

Result 2.10. *Let ϕ be a compactly supported refinable function corresponding to the isotropic dilation matrix s, with $\widehat{\phi}(0) \neq 0$ and let T be its associated transfer operator. Further, let the space H_ϕ and the ideal I_ϕ be given as in (2.8) and Definition 2.9, respectively. Then:*
(i) I_ϕ is T-invariant.
(ii) The regularity parameter $\alpha(\phi)$ is

$$-(\log_\lambda \rho)/2,$$

where ρ is the maximal modulus of the eigenvalues of the restriction of T to I_ϕ.
(iii) For ρ in (ii), there is an eigenpair (μ, f) of T such that $\rho = |\mu|$ and $f \in H_\phi \cap I_\phi$.

Indeed, the T-invariance of I_ϕ is proved in Theorem 2.4 of [RS1]. That theorem also shows that the regularity parameter $\alpha(\phi)$ is determined by any dominant eigenpair (μ, f) of T restricted on $H_\phi \cap I_\phi$, in the sense that $\alpha(\phi) = -(\log_\lambda |\mu|)/2$. This gives (ii). Recalling that all the T-eigenvectors in I_ϕ are either in H_ϕ or in $\ker T$, we get (iii). □

3. An algorithm for computing the regularity parameter

Result 2.10 suggests that in order to compute the regularity parameter of the refinable function, we 'merely' need to find the spectral radius of the restriction of T to H, where $H = H_\phi \cap I_\phi$. However, the result cannot be implemented directly, due to the fact that there is no 'good' method for constructing a basis for H.

Before we advance the discussion any further, we seek the following 'terminological relief': from now on, given any linear space S, and any linear bounded operator T from S into a superspace of it, the notion of *the spectral radius of* T is meant as *the spectral radius of the restriction of* T *to the largest* T*-invariant subspace of* S.

Result 2.10 suggests the following 'direct algorithm':

A simple method for computing $\alpha(\phi)$. Given the transfer operator T associated with the compactly supported refinable ϕ, do the following:

(i) Choose a T-invariant superspace H_0 of $H_\phi \cap I_\phi$ (one which is convenient for computations).

(ii) Find all eigenvalues ν of $T_{|_{H_0}}$.

(iii) For each eigenvalue ν, find the corresponding eigenspace V_ν, then check whether $V_\nu \cap I_\phi \neq 0$.

(iv) The desired ρ is $\max\{|\nu| : \ V_\nu \cap I_\phi \neq 0\}$.

Various improvements of this direct algorithm are possible. The most obvious one is to avoid finding all the eigenvalues ((ii) above), and instead finding them one by one in decreasing modulus of the eigenvalue; stop when the first eigenvector in I_ϕ is found. That approach suits the Arnoldi method of computing eigenvalues and eigenvectors. However, even with that improvement, the above 'direct method' suffers from the following drawbacks:

(a) If the critical eigenpair (μ, f_μ) is preceded by many other eigenpairs (whose eigenvalues have greater magnitudes), the approximation provided by the Arnoldi method for the critical eigenvector f_μ may be crude, and it may be hard to determine numerically whether $f_\mu \in I_\phi$.

(b) The necessity to compute a bulk of eigenpairs makes the process relatively slow.

(c) Even if the eigenvector is computed with high accuracy, it may still be hard to determine whether it belongs to I_ϕ. This problem (which exists in other approaches too, but to a lesser extent) is particularly troubling in the case of a multiple eigenvalue, since then we must check whether I_ϕ has a non-zero intersection with the eigenspace, a task which is almost always a numerical challenge (unless the eigenspace lies entirely in I_ϕ).

The above discussion reveals the following three different aspects that a successful algorithm has to deal with:

Aspect I: the eigenproblem aspect. We need to recover an eigenpair of a linear operator. The eigenpair that we look for may be dominated by many other pairs; nonetheless, we need a fast and accurate recovery of the eigenpair. It would be best if all/many/most of the eigenpairs that dominate the critical one can somehow be avoided. A fast implementation also requires a savvy conversion of the problem to matrix computations.

Aspect II: Φ and Π_ϕ. One of the key steps in any algorithm that computes the regularity parameter is to determine whether a given trigonometric polynomial f is in I_ϕ. For this, one needs to (1) find the polynomial Φ, and (2) find the space Π_ϕ (see the definition of the ideal I_ϕ). The first task is relatively modest: once we adopt a mild assumption (the *E-condition*, see below), it becomes truly simple to compute Φ accurately. As to the second task, viz. computing a basis for Π_ϕ, it is hampered by the fact that Π_ϕ, in general, does not have a simple structure (e.g., may not have a monomial basis), which makes it 'unpleasant' even under some additional conditions (e.g., stability). To overcome this difficulty, we use subtle theoretical facts that allow us to get away with only partial computation of Π_ϕ. Moreover, under 'favorable conditions' (which are far less demanding than stability), the approach yields a substantial shortcut in the search of the critical eigenvalue.

Aspect III: testing a given eigenvector. In order to check whether a given eigenvector f is in I_ϕ, one needs to check whether both (i) and (ii) in the definition of I_ϕ are satisfied. As we will see, the algorithm used here frees us from checking the second condition in the definition of I_ϕ. Furthermore, when the trigonometric polynomial Φ is positive *everywhere* (a condition which is known as 'the stability of the shifts of ϕ'), the first condition in the definition of the ideal I_ϕ is automatically satisfied. Hence, under this stability assumption, the process of checking whether the eigenvector in hand is in I_ϕ is fast and very robust. Without the stability assumption, we have to check whether f/Φ is bounded or not. This problem is on par with the classical NA problem: determining whether a small number is 0 or not. As said, this problem is particularly acute for multiple eigenvalues.

The first and third aspects above are problems that belong to the area Numerical Algebra, and we will discuss them in the next section, as a part of the discussion on the implementation and the code. To have an optimal treatment of the second aspect, we need some additional discussion concerning the regularity of refinable functions (beyond the general discussion of the previous section).

The discussion is divided into two parts: the first is about the computation of Φ and the second deals with Π_ϕ.

Computing the trigonometric polynomial Φ. We start with a finitely supported mask a. For a given mask, we want to know whether there exists a compactly supported solution to the corresponding refinement equation. If there is a solution, whether the solution is unique and whether the solution is in L_2. The following result provides satisfactory answers.

Result 3.1. *Let a be a finitely supported mask, and let T be the associated transfer operator.*

(i) If $\sum_{\alpha\in\mathbb{Z}^d} a(\alpha) = 1$ (i.e. $\widehat{a}(0) = 1$), there exists a compactly supported distribution ϕ that solves the refinement equation. It is the unique solution that satisfies $\widehat{\phi}(0) = 1$.

(ii) If the restriction of T to H_ϕ has spectral radius 1, and if all the eigenvalues (of that restriction) that lie on the unit circle are non-defective, then the corresponding solution of the refinement equation must lie in L_2.

(iii) If the solution ϕ of the corresponding refinement equation is in L_2, then $(1, \Phi)$ is an eigenpair of T.

The first statement is proved by showing that the infinite expansion $\prod_{j=1}^{\infty} \widehat{a}(s^{*-j}\omega)$ converges, uniformly on compact sets, to a tempered distribution. The last assertion is a straightforward exercise. The proof of the second assertion can be found in [LLS2] as well as in [R2].

Corollary 3.2. *Let a be a finite mask satisfying $\widehat{a}(0) = 1$. Assume that the corresponding transfer operator T satisfies (ii) of Result 3.1. Then T must have an eigenpair $(1, f)$, with f a non-negative trigonometric polynomial.*

Proof: By (ii) of Result 3.1, the corresponding refinement equation has an L_2-solution ϕ. Assertion (iii) in Result 3.1 then implies that $(1, \Phi)$ is an eigenpair of T. The polynomial Φ is obviously non-negative. $\square$

The condition that appears in part (ii) of Result 3.1 is not necessary for the solution ϕ to be in L_2 (cf. [RS1]), but refinable functions whose transfer operator violate this condition are quite 'pathological'. In our algorithm, we assume a bit more, namely that the eigenvalue 1 is *simple*:

Definition: the weak E-condition. Let a be a given finite mask with $\widehat{a}(0) = 1$ and let ϕ be the corresponding compactly supported solution. Let T be the transfer operator associated with ϕ. We say that a (or ϕ, or T) satisfies the weak E-condition, if the restriction of T to H_ϕ has spectral radius 1, all the eigenvalues on the unit circle are non-defective, and 1 (which is then necessarily an eigenvalue) is a simple eigenvalue.

Remark. The previous discussion implies that, under the weak E-condition, the refinement equation has a unique compactly supported solution, ϕ, that lies in L_2 and satisfies $\widehat{\phi}(0) = 1$. Further, Φ (i.e., the 2π-periodization of the Fourier transform of the autocorrelation of the solution) is the unique eigenvector (up to a constant) of the eigenvalue 1 of the transfer operator.

Remark. If we add to the weak E-condition the additional assumption that $T_{|_{H_\phi}}$ has a *unique* dominant eigenvalue, we obtain a condition known as the E-condition (which is useful in the analysis of various problems: for example, [LLS2] proves that the E-condition characterizes the L_2-convergence of the cascade algorithm; see also §3.1 of [R2]). This explains our usage of 'weak E condition'. Finally, we point out that it is not difficult to show that if T satisfies the weak E condition on H_ϕ, then it satisfies weak E condition on any T-invariant superspace H_0 of H_ϕ that consists of trigonometric polynomial (see [LLS2]).

In the first step of the algorithm, we select a convenient T-invariant superspace H_0 of H_ϕ. Then, the algorithm checks whether T satisfies the weak E-condition. If the weak E-condition is satisfied, it computes the eigenvector associated with the eigenvalue 1. The (normalized) symbol of that eigenvector is the function Φ.

Doing without Π_ϕ. We now elaborate on the second condition in the definition of I_ϕ. Let $\mathcal{Z}$ be some finite, fixed, subset of $\mathbb{Z}^d$, and let $H_\mathcal{Z}$ be the span of the exponentials $\omega \mapsto \exp(ij\cdot\omega)$, $j \in \mathcal{Z}$, endowed with the $L_2(\mathbb{T}^d)$-inner product $\langle\cdot,\cdot\rangle$. Given any (algebraic) polynomial p and any $f \in H_\mathcal{Z}$, one observes that $p(-iD)f(0) = \sum_{j\in\mathcal{Z}} p(j)f_j$, with $(f_j)_j$ the Fourier coefficients of f. Hence, the linear functional (in $H^*_\mathcal{Z}$) $f \mapsto p(-iD)f(0)$ is represented by the trigonometric polynomial $t_p(\omega) := \sum_{j\in\mathcal{Z}} p(j)\exp(ij\cdot\omega)$, i.e.,

$$p(-iD)f(0) = \langle t_p, f\rangle = \sum_{j\in H_\mathcal{Z}} p(j)f_j, \quad \forall f \in H_\mathcal{Z}. \tag{3.3}$$

We now connect the above abstract discussion to our concrete problem. In this discussion, we use, for a given subspace $Q \subset \Pi$ of algebraic polynomials, the notation

$$P_Q$$

for the orthogonal projector from $H_\mathcal{Z}$ onto $\{t_p : p \in Q\}$.

When $\mathcal{Z}$ above is $\mathcal{Z}_\phi$ (see (2.5)), the space $H_\mathcal{Z}$ becomes H_ϕ (of (2.8)). Furthermore, by choosing Q above to be Π_ϕ, the second condition in the definition of I_ϕ simply says that the critical eigenvector lies in the orthogonal complement (in H_ϕ) of $\{t_p :\ p \in \Pi_\phi\}$. Thus, if we set

$$P_\phi := P_{\Pi_\phi}$$

for the orthogonal projection of H_ϕ onto $\{t_p :\ p \in \Pi_\phi\}$, condition (ii) in the definition of I_ϕ will be automatically satisfied if we iterate (in the search for the critical eigenvector) with the operator $(1 - P_\phi)T$, instead of iterating with the transfer operator itself. This allows us to restate Result 2.10 in the following equivalent (yet more practical) way:

Restatement of Result 2.10. *In the notations and assumptions of Result 2.10, let I'_ϕ be the ideal of all trigonometric polynomials of the form $t\Phi$, $t \in L_\infty(\mathbb{T}^d)$ (i.e., those that satisfy the first condition in the definition of I_ϕ). Then the spectral radius ρ in Result 2.10 is the same as the spectral radius of the restriction of $(1 - P_\phi)T$ to $H_\phi \cap I'_\phi$.*

The discussion still leaves us with the need of finding a basis for Π_ϕ (in order to compute the projector P_ϕ). As we alluded to before, this can be partially circumvented: suppose that Q is some subspace of Π_ϕ, and suppose that we replace the operator $(1 - P_\phi)T$ by the operator $(1 - P_Q)T$. The latter one will fail to suppress some of the eigenvalues that the former one does; however, that apparent fault is harmless if we know that all these 'unsuppressed' eigenvalues are smaller than the critical one. But, do we have such a space Q, which, in addition, has a simple basis?

In order to answer the above question, we define

$$m_\phi := \max\{m \in \mathbb{N} : \Pi_m \subset \Pi_\phi\}, \tag{3.4}$$

where Π_m is the space of d-variate algebraic polynomials with degree $\le m$. We will show that we can replace the space Π_ϕ by the space Π_{m_ϕ}, and, moreover, we can sometimes do with Π_m, for $m < m_\phi$. In addition, we show a way to compute m_ϕ from the given data, viz., the mask a and the trigonometric polynomial Φ. We begin with that latter issue.

Proposition 3.5. *Let ϕ be a refinable compactly supported L_2-function with mask a. Let Γ be the set defined as in (2.6). Then $\Pi_m \subset \Pi_\phi$ if and only if $\widehat{b}\Phi$ $(= |\widehat{a}|^2\Phi)$ has a zero of order $m+1$ at each of the points in $\Gamma\backslash 0$.*

The spaces Π_m, $m \le m_\phi$, are certainly subspaces of Π_ϕ and have a simple structure. The next result studies the suitability of the choice $Q := \Pi_m$. For notational convenience we set, for any non-negative integer m,

$$P_m := P_{\Pi_m}.$$

Proposition 3.6. *Let ϕ be a refinable functions with corresponding mask a and transfer operator T. Let ρ_m be the spectral radius of the restriction of $(1 - P_m)T$ to $H_\phi \cap I'_\phi$. Then:*
(a) $\rho_{m_\phi} = \rho$.
(b) For an odd $m \le m_\phi$, we still have $\rho_m = \rho$, unless $\rho_m \le \lambda^{-m-1}$.

We prove the above propositions in the last section, hence are ready to present here our algorithm.

Algorithm: Step I. Compute the T-invariant space H_ϕ. Then, check whether T satisfies (on H_ϕ) the weak E-condition. If 1 is not an eigenvalue of T, return the message `‘‘There is no` L_2`-solution to the refinement equation’’`, and quit. If, otherwise, the weak E-condition is still violated, give another appropriate rejection message (that indicates that the solution may still not be in L_2) and quit. If the weak E-condition is satisfied, compute the eigenvector associated with the eigenvalue 1. Check (for consistency only) that the eigenvector is non-negative (or non-positive). The (normalized) symbol of that eigenvector is the function Φ.

Algorithm: Step II. Set $m_{\phi,\gamma} + 1$ to be the order of the zero that $|\widehat{a}|^2\Phi$ has at γ, and set

$$m_\phi := \min\{m_{\phi,\gamma} : \gamma \in \Gamma\backslash 0\}.$$

□

Algorithm: Step III. Find the eigenpairs (in H_ϕ) of $(1 - P_{\Pi_{m_\phi}})T$, one by one, ordered according to the eigenvalue modulus. Stop when finding the first eigenpair (μ, f_μ) for which f_μ/Φ is bounded. The L_2-regularity of $\alpha(\phi)$ is then $-\frac{\log_\lambda(|\mu|)}{2}$. □

Remark. We note that no differentiation is really conducted in Step II. Instead, one uses the fact that

$$p(-iD)f(\gamma) = \langle p\, \exp(i\gamma\cdot), f\rangle$$

(compare with 3.3). Further, since the maximal order of zeros of $\widehat{b}\Phi$ is even, m_ϕ is odd. □

Remark. We note that if Φ does not vanish at $\Gamma\backslash 0$, then the space $(1-P_{\Pi_{m_\phi}})H_\phi$ is T-invariant. In contrast, if Φ vanishes at a point of $\Gamma\backslash 0$, $(1 - P_{\Pi_{m_\phi}})H_\phi$ may not be T-invariant any more. Nonetheless, Proposition 3.6 always holds. Its proof relies on the fact that the subspace $(1 - P_{\Pi_{m_\phi}})(H_\phi \cap I_\phi)$ is *always* T-invariant. □

The algorithm checks for possible shortcuts; stability. In many cases of interest, the shifts of the refinable function are **stable**. A convenient way to define the stability here (which is entirely equivalent to the more standard definitions) is that $\Phi > 0$ (everywhere). Since our algorithm computes Φ in any event, it checks whether Φ is everywhere positive. In that event, it performs two shortcuts. The major one is that the first condition in the definition of I_ϕ becomes *superfluous*, and hence the iterations with $(1 - P_{\Pi_{m_\phi}})T$ search for a dominant eigenvalue. This not only accelerates the algorithm, but also results in a dramatic improvement of its numerical stability. Indeed, in this case we do not need to determine whether a large value of f/Φ should be interpreted as finite or infinite. Note that, since Φ is the dominant eigenvector, we are able to compute Φ with great accuracy. Hence, it is possible to have stable numerical algorithm to check whether $\Phi > 0$.

In the case of stability, another, less important, shortcut occurs: in the computation of m_ϕ, we look in general for the order of the zeros of $\widehat{b}\Phi$ on $\Gamma\backslash 0$. If Φ vanishes nowhere, these zeros coincide with those of $\widehat{b}$, and we do not need to compute $\widehat{b}\Phi$ (i.e., to convolve their Fourier coefficients.) For that shortcut, we only need Φ to be non-zero on $\Gamma\backslash 0$ (and indeed we implement that shortcut under that mere latter condition).

4. Numerical implementation details

In the actual numerical implementation, we treat the transfer operator as acting on *sequences*, i.e., we use the operator $\mathcal{T}$ defined by

$$\mathcal{T}c := (T\widehat{c})^\vee.$$

The sequence c is always defined on $\mathbb{Z}^d$ and has finite support. We use the pairing

$$\langle \theta, c\rangle := \sum_{j\in\mathbb{Z}^d} \theta(j)\overline{c(j)}, \tag{4.1}$$

in which c is finitely supported and θ is any sequence defined on $\mathbb{Z}^d$, or more generally, a function in $C(\mathbb{R}^d)$.

Next, we provide some details about the steps in the algorithm given in the previous section.

For the first step, we find the set $\mathcal{Z}_\phi$ as in Section 2. We then compute, via the deflated Arnoldi method, a basis for the dominant eigenspace of the transfer operator $\mathcal{T} : \ell_2(\mathcal{Z}_\phi) \longrightarrow \ell_2(\mathcal{Z}_\phi)$. Then, we check whether the transfer operator $\mathcal{T}$ satisfies the weak E-condition. If the weak E-condition is satisfied, we compute the eigenvector corresponding to the eigenvalue 1: its Fourier series is the function Φ; else, the weak E-condition is violated, and we quit.

For the second step, we first check whether Φ vanishes on $\Gamma\backslash\{0\}$. If it does not, we find the largest integer m such that

$$\langle \exp(i\gamma\cdot)\, p, a\rangle = 0, \qquad \forall\, p \in \Pi_m,\ \gamma \in \Gamma\backslash\{0\}, \tag{4.2}$$

where a is the refinement mask of ϕ. We then set $m_\phi := 2m + 1$. If Φ vanishes on $\Gamma\backslash\{0\}$, then we find the largest integer m such that

$$\langle \exp(i\gamma\cdot)\, p, c\rangle = 0 \qquad \forall\, p \in \Pi_m,\ \gamma \in \Gamma\backslash\{0\}, \tag{4.3}$$

where $c = b * h$, b is the mask of the autocorrelation function $\phi^\#$, and h is the Fourier coefficients of Φ, i.e., $\widehat{h} = \Phi$. We then set $m_\phi := m$.

For these, it is sufficient to check that (4.2) or (4.3) holds for a basis of Π_m. However, it is important to choose a well-conditioned basis. The usual monomial basis of Π_m is very ill-conditioned, therefore is inappropriate for our purpose. We choose here instead a suitable orthonormal basis. That orthonormal basis is described in the sequel.

For the third step, if Φ vanishes nowhere, we compute the dominant eigenvalue μ of $(I - P_{\Pi m_\phi})T$ via the deflated Arnoldi method as detailed below. Then, set $\alpha(\phi) = -\frac{\log_\lambda |\mu|}{2}$.

If Φ vanishes anywhere (in $[-\pi, \pi]^d$ as this function is 2π-periodic), then we proceed as follows:

(i) We compute the next group of the distinct dominant eigenvalues of $(I - P_{\Pi m_\phi})T$ via the deflated Arnoldi method. Then we order the eigenvalues according to decreasing magnitudes of their values as

$$|\mu_1| \geq |\mu_2| \geq \cdots.$$

(ii) We compute a basis for the eigenspace associated with each of the eigenvalues computed in (i) via the deflated Arnoldi method. Denote them as $\{f_1, \dots, f_L\}$.

(iii) If there exists scalars $t_1, \dots, t_L$ not all zero such that

$$\sum_{i=1}^{L} t_i \widehat{f}_i / \Phi$$

is bounded, then set $\alpha(\phi) = -\frac{\log_\lambda |\mu_k|}{2}$; stop. Otherwise, go back to step (i).

We discuss now the following numerical methods used to implement the algorithm.

The action of $\mathcal{T}$ on a vector.
Let c be an arbitrary sequence in $\ell_2(\mathcal{Z}_\phi)$. The action of the transfer operator $\mathcal{T}$ on c is as follows. First, generate a new sequence $b * c$ by convolution, then reparameterize the sequence $(b * c)_{j\in\mathbb{Z}^d}$ to a sequence defined on $s^{-1}\mathbb{Z}^d$. Finally, the image $\mathcal{T}c$ is the restriction to $\mathbb{Z}^d$ of the sequence $(b * c)_{j\in s^{-1}\mathbb{Z}^d}$. The resulting sequence $\mathcal{T}c$ is still supported in $\mathcal{Z}_\phi$. Once $\mathcal{T}c \in \ell_2(\mathcal{Z}_\phi)$ is obtained, it is relatively easy to compute orthogonal projections of it onto various subspaces, provided that we also have an orthonormal basis for these subspaces.

Construction of an orthonormal basis for Π_n.
The standard construction of an orthonormal basis (ON) for Π_m is done by applying the Gram-Schmidt process to the monomial basis $\{(j^\beta)_{j\in\mathcal{Z}_\phi} : |\beta| \le m\}$. However, this standard construction is numerically unstable. A more stable process (known as the modified Gram-Schmidt) can be devised by modifying the Gram-Schmidt process, which we describe now in the bivariate case. Set $N := \#\mathcal{Z}_\phi$.

Modified Gram-Schmidt:
  Let $v^{(0,0)} = \frac{1}{\sqrt{N}}(1)_{j\in\mathcal{Z}_\phi}$.
  for $k = 1, 2, \ldots, m$
    for $\beta_1 = 0, 1, \ldots, k$
      if $\beta_1 = 0$
        $w = (j(2)\, v^{(0,k-1)}(j))_{j\in\mathcal{Z}_\phi}$.
      else
        $w = (j(1)\, v^{(\beta_1-1,k-\beta_1)}(j))_{j\in\mathcal{Z}_\phi}$.
      Orthogonalized w against all previously generated ON vectors v to get w'.
      Set $v^{(\beta_1,k-\beta_1)} = w'/\|w'\|_2$.
  Let

$$B_m := \{v^{(\beta)} : |\beta| \le m\}. \tag{4.4}$$

Now, we describe here how to apply the Deflated Arnoldi Method [S] to our case. The method may not be as robust as other more sophisticated methods for the same purpose, such as the implicitly restarted Arnoldi [LSVY], [LS], the Jacobi-Davidson Method [SV], and the truncated RQ iteration [SY]. Nonetheless, as our examples in the next section show, even with this relatively simple method, our proposed algorithm works well. Of course, for a more robust implementation, one should replace the Deflated Arnoldi Method by one of the more robust dominant eigenspaces solvers just mentioned.

The Deflated Arnoldi Method.
We first note that the operator $(I - P_\phi)\mathcal{T}$ can be viewed as an operator on $\mathbb{R}^N$ with $N = |\mathcal{Z}_\phi|$; we just need to order the points in $\mathcal{Z}_\phi$, and identify $\ell_2(\mathcal{Z}_\phi)$ with $\mathbb{R}^N$, the latter equipped with the standard inner product on $\mathbb{R}^N$. Let A be an arbitrary linear endomorphism of $\mathbb{R}^N$. The Deflated Arnoldi Method is described in the following steps:

(1) Choose an initial vector $v_1 \in \mathbb{R}^N$ with $\|v_1\|_2 = 1$. Set $k = 1$. Select the number m of Arnoldi iterations to be performed in each pass.
(2) Arnoldi iteration:

  for $j = k, k+1, \ldots, m$
    compute $w = Av_j$
    for $i = 1, 2, \ldots, j$

$$
\begin{aligned}
&\quad h_{ij} = \langle w, v_i \rangle \\
&\quad w = w - h_{ij} v_i \\
&h_{j+1,j} = \|w\|_2 \\
&v_{j+1} = w / h_{j+1,j}
\end{aligned}
$$

Let V_m be the matrix whose kth column is the vector v_k and $H_m = (h_{ij})$ be the $m \times m$ upper Hessenberg matrix constructed above. The vectors v_j generated by the Arnoldi iteration satisfy the following relation:

$$AV_m = V_m H_m + h_{m+1,m}\, v_{m+1} e_m^T.$$

Suppose (μ, y) is an eigenpair of H_m. Then $(\mu, V_m y)$ is an approximate eigenpair of A.

(3) Compute a basis, say $\{y_1, y_2, \cdots, y_t\}$, for the eigenspace associated with the dominant eigenvalue μ of H_m. Let $\rho_k = \|AV_m y_k - \mu y_k\|_2$, for $k = 1, \ldots, t$. Suppose we order the basis such that the residual norms ρ_k satisfy $\rho_1 \leq \rho_2 \leq \cdots \leq \rho_t$. If $\{y_1, y_2, \ldots, y_r\}$ (where $r \leq t$) are the vectors such that the corresponding residual norms are small enough, then $\{u_1 = V_m y_1, u_2 = V_m y_2, \cdots, u_r = V_m y_r\}$ are approximate eigenvectors of A associated with the dominant eigenvalue μ.

(4) *Deflation*: Suppose $\{y_1, y_2, \ldots, y_r\}$ are eigenvectors of H_m corresponding to converged eigenvectors $\{u_1, u_2, \ldots, u_r\}$ of A associated with the dominant eigenvalue μ. This step is to deflate these converged eigenvectors from the Arnoldi iteration so that additional eigenvectors of A associated with the eigenvalue μ can be found, whenever they exist.

(i) Compute the QR factorization of the matrix $(y_1, \ldots, y_r)$ using Householder matrices:

$$(y_1 \quad \ldots \quad y_r) = Q \begin{pmatrix} R_r \\ 0 \end{pmatrix}$$

where Q is an $m \times m$ orthogonal matrix and R_r is an $r \times r$ upper triangular matrix.

(ii) Update the factorization

$$H_m \leftarrow Q^T H_m Q$$

$$V_m \leftarrow V_m Q.$$

It can be shown that the matrices V_m and H_m satisfies the relation

$$AV_m = \; V_m H_m + h_{m+1,m}\, v_{m+1} e_m^T + h_{m+1,m}\, v_{m+1} w^T,$$

for some vector w such that $\|w\|_2$ is close to the machine epsilon if the condition number of R_r is modest. Furthermore, the columns of V_m together with v_{m+1} form an orthonormal set, and the first r vectors of V_m lie in the eigenspace of A associated with μ. That is, the first r vectors of V_m are Schur vectors for the eigenspace of A associated with μ. The $r \times r$ principal minor of H_m is the upper triangular matrix R_r.

(iii) Exit Step 4. Discard the vectors $v_{r+1}, \cdots, v_m$ in V_m. Set $k = r + 1$ and $v_{r+1} = v_{m+1}$, repeat Step 2 through Step 4; stop if a basis for the eigenspace of A associated with μ has been found. Note that this process is equivalent to applying a new deflated Arnoldi iteration with initial vector v_{r+1} to the operator $(I - P_r)A$, where P_r is the orthogonal projector onto the subspace spanned by the Schur vectors $\{v_1, \cdots, v_r\}$ of A.

Remark. For simplicity, our discussion above focused on finding the dominant eigenspace of A, but this restriction is not necessary. In practice, one can find the eigenspaces associated with several dominant eigenvalues simultaneously.

Checking the boundedness of f/Φ

The major and the most difficult substep of Step III is to check the boundedness of f/Φ, with f a given trigonometric polynomial. When $\Phi > 0$ (i.e., when the shifts of ϕ are stable,) f/Φ is always bounded, and this substep is omitted. Thus, under the stability assumption, our algorithm (and code) is very robust, for both univariate and bivariate cases. As a proof of evidence, our code successfully computed the regularity of the $(8,8,8)$ bivariate interpolatory mask of [RiS1], whose autocorrelation mask is support on the square $[-31,31]\times[-31,31]$. The matrix representation of the associated transfer operation has an order of about 4000, and a brute force calculation of the regularity using the transfer operator would require one to find hundreds of eigenvalues of a huge matrix and decide later which of the eigenvalues is the critical one. In contrast, since our algorithm does not use the matrix representation of the transfer operator explicitly, the size of the memory we need is only a small fraction of that required by a direction calculation. Also, by suppressing *a priori* hundreds of eigenvectors corresponding to polynomial reproduction, we need only to calculate the dominant eigenvalue of the operator $(I-P_\phi)\mathcal{T}$ instead of a multitude of eigenvalues of $\mathcal{T}$.

For the univariate case, since the function Φ only has finitely many isolated zeros and the multiplicity of each zero is relatively easy to find, the boundedness of f/Φ can be completely settled. Consequently, the algorithm and the derived code provide in this case the exact regularity parameter.

For the bivariate case, it is much more difficult to compute numerically the multiplicity space of the zeros of Φ. The current version of the code can only handle the case when Φ has finitely many zeros (which we find as an acceptable assumption: refinable functions with unstable shifts may be very useful in the construction of *framelets* with 'customized' properties (cf. [R1], [DRS]). It is very unlikely that any of these constructs will violate the 'finitely many zeros' condition). Already for this case, the reliability of our code depends on (i) the number of the zeros and their distribution, (ii) the 'degree' of the multiplicity space of each zero. However, for all of the interesting examples we tested, we did obtain reliable smoothness parameters. Even for an 'extremely bad' refinable function (i.e., whose Φ vanishes at many points and to high degrees) the code is able to provide 'good' lower bounds on the regularity, much better than the lower bound obtained by ignoring the dependence relation effect.

Given a trigonometric polynomial f, in order to check whether f/Φ is bounded in $[-\pi,\pi]^d$, one needs only to check whether it is bounded in local neighborhoods of the zeros of Φ.

Let ξ be a zero of Φ in $[-\pi,\pi]^2$ of exact order m. (The number m can be computed numerically.) Thus, all the derivatives of Φ up to order $m-1$ vanish at ξ, but some derivatives of order m do not. The Taylor expansion of Φ at ξ has then the form

$$\Phi(\xi+\eta) \;=\; \sum_{|\beta|=m} \frac{D^\beta\Phi(\xi)}{\beta!}\,\eta^\beta \;+\; \mathcal{O}(\|\eta\|^{m+1}).$$

Now, if f/Φ were to be bounded in a local neighborhood of ξ, then it would be necessary for f to satisfy the condition:

$$D^\beta f(\xi) \;=\; 0 \quad \forall\ |\beta| \le m-1. \tag{4.5}$$

Hence, we can reject those eigenvalues whose eigenspace contains no eigenvectors that satisfy (4.5) (for all ξ).

Next we discuss how the condition (4.5) can be checked numerically in our algorithm.

Suppose $\{g_1,\ldots,g_L\}$ is a basis for the eigenspace associated with an eigenvalue μ. Consider the eigenvector $g=\sum_{i=1}^L c_i g_i$, where not all the coefficients are zero. The condition that $D^\beta\widehat{g}(\xi)=0$ for all $|\beta|\le m-1$ is equivalent to

$$\langle g, S_\xi\rangle = 0,$$

where

$$S_\xi \;=\; \{p\,\exp(i\xi\cdot)\;:\;p\in\Pi_{m-1}\}\,.$$

Since g is supported on $\mathcal{Z}_\phi$, the functions in S_ξ may be regarded as sequences defined on that domain, too. Thus, we may interpret the above condition as saying that g lies in the null space of B_ξ^*, for a suitable matrix B_ξ (whose columns span S_ξ). Thus, if G is a corresponding matrix representation for the basis $\{g_1,\cdots,g_L\}$, we need to find the null space of $B_\xi^* G$.

If Φ has more than one zero, say $\xi^{(1)}, \ldots, \xi^{(K)}$, then in order to find g such that $\widehat{g}/\Phi$ is bounded on $[-\pi, \pi]^d$, we seek a non-trivial null space for

$$\mathcal{G} := \left[B_{\xi^{(1)}} \ \cdots \ B_{\xi^{(K)}}\right]^* G.$$

In our implementation, we find c in the null space of $\mathcal{G}$ by computing the SVD (singular value decomposition) of $\mathcal{G}$. If the minimal singular value $\sigma_{\min}(\mathcal{G})$ of $\mathcal{G}$ is sufficiently small, then we conclude that the null space of $\mathcal{G}$ is non-trivial and take c to be a minimal singular vector of $\mathcal{G}$.

Suppose that f satisfy (4.5), and that $d = 2$. If, in addition, the following two polynomials

$$t \mapsto \sum_{|\beta|=q} \frac{D^\beta \Phi(\xi)}{\beta!} t^{\beta_2}, \qquad t \mapsto \sum_{|\beta|=q} \frac{D^\beta \Phi(\xi)}{\beta!} t^{\beta_1}, \tag{4.6}$$

are strictly positive on the interval $[-1, 1]$, then f/Φ is bounded. To see this, we analyze the ratio

$$\frac{f(\xi+\eta)}{\Phi(\xi+\eta)} = \frac{\sum_{|\beta|=q} \frac{D^\beta f(\xi)}{\beta!} \eta^\beta + \mathcal{O}(\|\eta\|^{q+1})}{\sum_{|\beta|=q} \frac{D^\beta \Phi(\xi)}{\beta!} \eta^\beta + \mathcal{O}(\|\eta\|^{q+1})}, \tag{4.7}$$

for sufficiently small non-zero vector η. Suppose $|\eta_2| \le |\eta_1|$. Then $\eta_2 = t\,\eta_1$ for $t \in [-1, 1]$, and substituting this into (4.7) would lead to

$$\frac{f(\xi+\eta)}{\Phi(\xi+\eta)} = \frac{\sum_{|\beta|=q} \frac{D^\beta f(\xi)}{\beta!} t^{\beta_2} + \mathcal{O}(\eta_1)}{\sum_{|\beta|=q} \frac{D^\beta \Phi(\xi)}{\beta!} t^{\beta_2} + \mathcal{O}(\eta_1)}.$$

Hence, whenever the polynomials in (4.6) are strictly positive on $[-1, 1]$, f/Φ is bounded in a neighborhood of ξ (the above argument applies to the case $|\eta_2| \le |\eta_1|$, and the complementary case is obtained by symmetry.) Finally, we remark whether the polynomials in (4.6) are strictly positive can be checked numerically.

It must be emphasized that multiplicity of the zero of Φ at a given point ξ, while necessarily of finite-dimension (since the zero is isolated), is not always of a total degree form. The present version of our code, however, computes only the total degree subspace of that multiplicity space, hence provides in such cases *lower bounds* on the smoothness parameter.

5. Examples

We record some of our numerical experiments that we conducted as a test for the code.

The first class of examples are taken from the bivariate interpolatory refinable functions that were constructed by [RiS1] ('interpolatory' means that $\phi(j) = \delta_j$, $j \in \mathbb{Z}^d$ and is a stronger property than stability). These examples demonstrate that the code can handle very large masks of stable refinable functions.

Example 5.1. The mask a_r of an interpolatory refinable function ϕ_r in [RiS1] is obtained by convoluting the mask m_r of a centered three directional box spline with mask q_r of a carefully chosen distribution. The symbol of m_r (for an even r) is

$$\widehat{m}_r(\omega) = \Big(\cos\big(\frac{\omega_1}{2}\big)\cos\big(\frac{\omega_2}{2}\big)\cos\big(\frac{\omega_1+\omega_2}{2}\big)\Big)^r.$$

The mask m_r is of a box spline that lies in $C^{2r-2}(\mathbb{R}^2)$. The smoothness of ϕ_r also increases with r, but not at the same rate as its box spline factor. The distribution factor q_r, while having a negative effect on the smoothness, is necessary in order to achieve the interpolatory property of ϕ_r. For $r = 2$, the corresponding q_2 is:

$$\widehat{q}_2(\omega) = \Big(5 - \cos(\omega_1) - \cos(\omega_2) - \cos(\omega_1+\omega_2)\Big)/2,$$

The L_2-regularity of ϕ_2 is 2.440765. We computed the smoothness α_r of the other interpolatory refinable functions ϕ_r, $r = 3, 4, \ldots 8$. They are as follows:

r	3	4	5	6	7	8
α_r	3.175132	3.793134	4.344014	4.862018	5.362768	5.852746

As a second test class, we tested four directional box splines. It is well-known (cf. [BHR]) that the shifts of the four directional box spline are not stable. At the same time, their smoothness is explicitly known.

Example 5.2. The symbols of the masks of the four direction box splines considered here are:

$$\widehat{m}_r(\omega) = \Big(\cos\big(\frac{\omega_1}{2}\big)\cos\big(\frac{\omega_2}{2}\big)\cos\big(\frac{\omega_1+\omega_2}{2}\big)\cos\big(\frac{\omega_1-\omega_2}{2}\big)\Big)^r.$$

Our code computed, for $r = 1, 2, 3, 4$ the corresponding smoothness of

$$2.5,\ 5.5,\ 8.5,\ 11.5.$$

These are, indeed, the exact smoothness parameters of these splines.

□

The third set of examples is taken from [JS]. The pertinent refinable functions are univariate, interpolatory, and correspond to dilation $s = 3, 4$. The shifts of these functions form an orthonormal system.

Example 5.3. The mask a_n of interpolatory refinable function ϕ_n whose shifts form an orthonormal basis is obtained by convoluting a B-spline of order n with the mask q_n of some distribution. The smoothness of the examples in [JS] with dilation $s = 3$, and with B-spline factor of order 2 and 3 are 0.963825 and 1.098068 respectively. The smoothness of the examples in [JS] with dilation $s = 4$ and a B-spline factor of order $2, 3, 4$ are 0.890339, 1.21178, and 1.303449. □

Example 5.4. The next example is a univariate refinable function whose shifts are unstable, with mask given by

$$\widehat{m}(\omega) = \cos^j\big(\frac{\omega}{2}\big)\ \big(2\cos(\omega) - 1\big)^k.$$

For $(j, k) = (4, 3)$ and $(j, k) = (4, 2)$, the smoothness of the refinable functions computed is 3.5. This agrees with the fact that both functions are cubic splines. We note that for this examples the lower bound estimates (that ignore the first condition in the definition of I_ϕ) fail to yield the correct smoothness. □

The last example shows the difficulties in getting the exact regularity of refinable functions, in the case the corresponding dominant eigenvector Φ of T has many zeros. However, a good lower bound of the regularity is still possible to obtain.

Example 5.5. The mask is

$$\widehat{m}_r(\omega) = \cos\big(\frac{\omega_1}{2}\big)\cos\big(\frac{\omega_2}{2}\big)\cos\big(\frac{\omega_1+\omega_2}{2}\big)\cos\big(\frac{\omega_1-\omega_2}{2}\big)\ \Big(\frac{1+e^{i(6\omega_1+5\omega_2)}}{2}\Big)\ \Big(\frac{1+e^{i(-3\omega_1+5\omega_2)}}{2}\Big).$$

The operator $(I - P_\phi)T$ has the following dominant eigenvalues

$$\mu = 2^{-6} \quad \text{with the dimension of eigenspace} = 6;$$
$$\mu = 2^{-7} \text{ with the dimension of eigenspace} = 12;$$
$$\mu = 2^{-8} \text{ where dimension of eigenspace} = 52.$$

Thus, a straightforward lower bound on the smoothness is 3. The function Φ has about 79 zeros in $[-\pi, \pi) \times [-\pi, \pi)$. In our computations, we were able to compute accurately the following zeros

$$(-\pi, -\pi) \quad \pm(2\pi/3, 0) \quad \pm(0, 0.8\pi) \quad \pm(0, 0.4\pi) \quad \pm(2\pi/3, 0.8\pi).$$

Each is verified to have total order 4. Based on these zeros, we were able to reject the eigenvalues 2^{-6} and 2^{-7} as 'false' eigenvalues. Thus a lower bound on the regularity is 4. The refinable function is this case is a box spline whose exact L_2-smoothness is $\alpha = 4.5$.

□

6. Proofs of Propositions 3.5 and 3.6

Proof of Proposition 3.5: Approximation Theory basics (cf., e.g., [BDR] and [BR]) imply that $\Pi_m \subset \Pi_\phi$ if and only if $|\widehat{\phi}|^2$ has a zero of order $m+1$ at each $j \in 2\pi\mathbb{Z}^d\backslash 0$. Set $\mathcal{L} := 2\pi(\mathbb{Z}^d\backslash(s^*\mathbb{Z}^d))$ (to get a feeling for that set: in one dimension, dyadic dilations, this is the set of 2π-odd integers). Given a non-zero 2π-integer j, we write it as $j = s^{*k}j'$, $j' \in \mathcal{L}$, and use k times the refinement equation to conclude that

$$\widehat{\phi}(\omega + j) = \widehat{\phi}(s^{*-k}\omega + j') \prod_{n=1}^{k} \widehat{a}(s^{*-n}(\omega + j)).$$

This means that $|\widehat{\phi}|^2$ has a zero of order $m+1$ as each point of $2\pi\mathbb{Z}^d\backslash 0$ if and only if it has such a zero at each point of (the smaller set) $\mathcal{L}$.

We proceed by stating the following lemma, whose proof is postponed until after the proposition is proved:

Lemma. *Let ϕ be a compactly supported L_2-function. Let $\gamma \in \mathbb{R}^d$. Then $|\widehat{\phi}|^2$ vanishes to order m at each $j \in \gamma + 2\pi\mathbb{Z}^d$ if and only if its 2π-periodization Φ has such zero at γ.*

In order to complete the proof of the proposition, note that $\mathcal{L}$ is the disjoint union of the cosets $s^*(\gamma + 2\pi\mathbb{Z}^d)$, $\gamma \in \Gamma\backslash 0$. For $j \in 2\pi\mathbb{Z}^d$, the 2π-periodicity of $\widehat{b}$ implies that $|\widehat{\phi}(s^*(\gamma + j))|^2 = \widehat{b}(\gamma)|\widehat{\phi}(\gamma + j)|^2$. The 2π-periodization of the right-hand-side is $\widehat{b}(\gamma)\Phi(\gamma)$, thus the lemma applies to show that $|\widehat{\phi}|^2$ has a zero of order $m+1$ at each of $s^*(\gamma + 2\pi\mathbb{Z}^d)$ if and only if $\widehat{b}\Phi$ has such a zero at γ. Varying that conclusion over all $\gamma \in \Gamma\backslash 0$, we obtain the desired result.

It remains now to prove the lemma. One implication here is trivial: since $|\widehat{\phi}|^2$ is non-negative, its 2π-periodization can have a zero of a certain order at γ only if each of the summands has a corresponding zero.

Assume conversely that $|\widehat{\phi}|^2$ has a zero of order m at each $\gamma + j$, $j \in 2\pi\mathbb{Z}^d$, and note that (since $\widehat{\phi}$ is smooth) m must be even. Let Ω be a small neighborhood of γ. Since ϕ is compactly supported, we have that $\widehat{\phi} \in W_2^\rho(\mathbb{R}^d)$ for any ρ. Now, since $\widehat{\phi}$ has a zero of order $m/2$ at $\gamma + j$, we have (with D^β, $\beta \in \mathbb{Z}^d$, the usual partial differentiation)

$$|\widehat{\phi}(\omega + \gamma + j)| \le c|\omega|^{m/2} \max_{|\beta|=m/2} \|D^\beta\widehat{\phi}\|_{L_\infty(\Omega+j)}, \quad \text{for } \omega \in \Omega. \tag{6.1}$$

Choosing $\rho > m/2 + d/2$, the Sobolev embedding theorem implies that $W_2^\rho(\Omega + j)$ is continuously embedded in the Sobolev space $W_\infty^{m/2}(\Omega + j)$. Thus,

$$\max_{0\le|\beta|\le m/2} \|D^\beta\widehat{\phi}\|_{L_\infty(\Omega+j)} \le c_1\|\widehat{\phi}\|_{W_2^\rho(\Omega+j)},$$

with c_1 *independent of* j (since all the $\Omega + j$ sets are translates of each other). Substituting this into (6.1) we obtain that

$$|\widehat{\phi}(\omega + \gamma + j)| \le c_2|\omega|^{m/2}\|\widehat{\phi}\|_{W_2^\rho(\Omega+j)}, \quad \omega \in \Omega, j \in 2\pi\mathbb{Z}^d.$$

Squaring the last inequality and summing over $j \in 2\pi\mathbb{Z}^d$ (and assuming, for simplicity and without loss, that ρ is an integer) we obtain that

$$\Phi(\omega + \gamma) \le c_3|\omega|^m\|\widehat{\phi}\|^2_{W_2^\rho(\mathbb{R}^d)}.$$

□

Proof of Proposition 3.6: Statement (a) follows from (b): choosing in (b) m to be (the odd number) m_ϕ, we get (a) unless $\rho_{m_\phi} \le \lambda^{-m_\phi-1}$. However, in the event that this latter inequality holds, we get that $\rho \le \rho_m \le \lambda^{-m_\phi-1}$, implying thereby that $\alpha(\phi) \ge \frac{m_\phi+1}{2}$. This implies (cf. [R1]) that the shifts of ϕ span all polynomials of degree $\frac{m_\phi+1}{2}$, hence that the shifts of $\phi^\#$ span all polynomials of degree $m_\phi+1$, in contradiction to the very definition of m_ϕ.

In order to prove (b), let $f \in H_\phi \cap I'_\phi$ be an eigenvector of the operator $(1-P_m)T$, with an associated eigenvalue μ. Assume also that $|\mu| > \lambda^{-m-1}$.

We first prove that f is actually an eigenvector of T. For that, we first observe that $\widehat{b}f$ has a zero of order $m+1$ at each of the points of Γ: for $\gamma \in \Gamma\backslash 0$, this follows from the fact that $\widehat{b}f = \widehat{b}\Phi t$, for a bounded t (since $f \in I'_\phi$), together with Proposition 3.5. For $\gamma = 0$, this follows from that fact that, by assumption, $(1-P_m)Tf = \mu f$, hence that f lies in the range of $(1-P_m)$ (and every function in that range vanishes to order $m+1$ at the origin). Thus, indeed, $\widehat{b}f$ vanishes to order $m+1$ on Γ. We conclude from the definition of T that Tf vanishes to order $m+1$ at the origin, hence that $\mu f = (1-P_m)Tf = Tf$.

We will now establish (b) by a chain of (in)equalities. First, by our assumptions on μ,

$$m+1 < \log_\lambda(|\mu|).$$

Second, once we know that (μ, f) is an eigenpair of T, we can write

$$\log_\lambda(|\mu|) = \lim_{k\to\infty} \frac{\log_\lambda \|T^k(f)\|_{L_1(\mathbb{T}^d)}}{k}.$$

Since $|Tf| \le T|f|$ (regardless of the nature of f), we get that

$$\lim_{k\to\infty} \frac{\log_\lambda \|T^k(f)\|_{L_1(\mathbb{T}^d)}}{k} \le \limsup_{k\to\infty} \frac{\log_\lambda \|T^k|f|\|_{L_1(\mathbb{T}^d)}}{k}. \tag{6.2}$$

Let

$$u : \omega \mapsto \Big(\sum_{j=1}^d \sin^2(\omega_j/2)\Big)^{(m+1)/2}.$$

Then, our assumptions here imply that the function

$$g := |f|/u$$

is bounded, and that, moreover, g/Φ is also bounded. Invoking (b) of Corollary 2.10 of [RS1] (with g there being our g here, and with ℓ there being $(m+1)/2$ here; the corollary requires that the right hand side of (6.2) is greater than $m+1$, something that we have already proved), we get that

$$\limsup_{k\to\infty} \frac{\log_\lambda \|T^k(|f|)\|_{L_1(\mathbb{T}^d)}}{k} \le -2\alpha(\phi).$$

Finally,

$$-2\alpha(\phi) = \log_\lambda \rho.$$

We thus conclude that $\log_\lambda |\mu| \le \log_\lambda \rho$, hence that $\rho_m \le \rho$. The converse inequality is trivial. $\square$

References

[A] W.E. Arnoldi, The principle of minimized iterations in the solution of the matrix eigenvalue problem, Quart. Appl. Math., **9** (1951), 17–29.

[BDR] C. de Boor, R. DeVore and A. Ron, Approximation from shift-invariant subspaces of $L_2(\mathbb{R}^d)$ Trans. Amer. Math. Soc. **341** (1994), 787–806. Ftp site: `ftp://ftp.cs.wisc.edu/Approx` file: `l2shift.ps`

[BR] C. de Boor and A. Ron, The exponentials in the span of the integer translates of a compactly supported function: approximation orders and quasi-interpolation, J. London Math. Soc. **45** (1992), 519-535. Ftp site: `ftp://ftp.cs.wisc.edu/Approx` file: `quasi.ps`

[BW] E. Belogay and Yang Wang, Arbitrarily smooth orthogonal nonseparable wavelets in $\mathbb{R}^2$, SIAM J. Math. Anal. **30** (1999), 678-697.

[CD] A. Cohen and I. Daubechies, Nonseparable bidimensional wavelet bases, Rev. Mat. Iberoamericana **9** (1993), 51–137.

[CGV] A. Cohen, K. Gröchenig and L. Villemoes, Regularity of multivariate refinable functions, Constr. Approx. **15** (1999), 241–255.

[CS] A. Cohen and J.-M. Schlenker, Compactly supported bidimensional wavelet bases with hexagonal symmetry, Constr. Approx. **9** (1993), 209–236.

[D] I. Daubechies, *Ten lectures on wavelets*, CBMS Conference Series in Applied Mathematics, **61**, SIAM, Philadelphia, 1992.

[DD] G. Deslaurier, and S. Dubuc, Symmetric iterative interpolation processes, Constr. Approx. **5** (1989), 49-68.

[DDD] G. Deslaurier, J. Dubois and S. Dubuc, Multidimensional iterative interpolation, Canad. J. Math, **43** (1991), 297–312.

[DRS] I. Daubechies, A. Ron and Z. Shen, Construction and analysis of framelets in one and several dimensions, (1999), manuscript.

[DGL] N. Dyn, J. A. Gregory and D. Levin, A butterfly subdivision scheme for surface interpolation with tension control, ACM Trans. on Graphics, **9** (1990), pp 160-169.

[E] T. Eirola, Sobolev characterization of solutions of dilation equations. SIAM J. Math. Anal. **23** (1992), 1015–1030.

[GR] K. Gröchenig and A. Ron, Tight compactly supported wavelet frames of arbitrary high smoothness, *Proc. Amer. Math. Soc.*, **126** (1998), 1101–1107. Ftp site: `ftp://ftp.cs.wisc.edu/Approx` file: `cg.ps`

[HL] Wenjie He and Ming-Jun Lai, Construction of bivariate compactly supported biorthogonal box spline wavelets with arbitrarily high regularities, Appl. Comput. Harm. Anal. **6** (1999), 53-74.

[HJ] Bin Han and Rong-Qing Jia, Optimal interpolatory subdivision schemes in multidimensional spaces, SIAM J. Numer. Anal., **36** (1999), 105-124.

[J] Rong-Qing Jia, Characterization of smoothness of multivariate refinable functions in Sobolev spaces, Trans. Amer. Math. Soc. **351** (1999), 4089-4112.

[JRS] Hui Ji, S. D. Riemenschneider and Z. Shen, Multivariate Compactly supported fundamental refinable functions, duals and biorthogonal wavelets, Studies in Appl. Math., **102** (1999), 173-204.

[JS] Hui Ji and Z. Shen, Compactly supported (bi)orthogonal wavelets generated by interpolatory refinable functions, Adv. Comp. Math. **11** (1999), 81-104.

[KS] J. Kovacević and W. Sweldens, Wavelet families increasing order in arbitrary dimensions, (1997) preprint.

[LMW] K. S. Lau, M. F. Ma and J. Wang, On some sharp regularity estimations of L^2- scaling functions, SIAM J. Math. Anal. **27** (1996), 835-864.

[LLS1] W. Lawton, S. L. Lee and Z. Shen, Stability and orthonormality of multivariate refinable functions, SIAM J. Math. Anal. **28** (1997), 999–1014.

[LLS2] W. Lawton, S. L. Lee and Z. Shen, Convergence of multidimensional cascade algorithm, Numer. Math., **78** (1998), 427-438.

[LS] R.B. Lehoucq and D.C. Sorensen, Deflation techniques for an implicitly restarted Arnoldi iteration, SIAM J. Matrix Anal. Appl. **17** (1996), 789-821.

[LSVY] R.B. Lehoucq, D.C. Sorensen, P. Vu, and C. Yang, ARPACK: An implementation of the implicitly restarted Arnoldi iteration that computes some of the eigenvalues and eigenvectors of a large sparse matrix, 1995. Available, from `ftp.caam.rice.edu` from the directory `pub/software/ARPACK`.

[R1] A. Ron, Smooth refinable functions provide good approximation orders, SIAM J. Math. Anal. **28** (1997), 731–748. Ftp: `ftp://ftp.cs.wisc.edu/Approx` File: `smoothwav.ps`

[R2] A. Ron, Wavelets and their associated operators, in *Approximation Theory IX* Vol. II, C. K. Chui, L. L. Schumaker eds, Vanderbilt University Press (1998), 283–317. Ftp: `ftp://ftp.cs.wisc.edu/Approx` File: `texas9.ps`

[RiS1] S. D. Riemenschneider and Z. Shen, Multidimensional interpolatory subdivision schemes, SIAM J. Numer. Anal., **34** (1997), 2357-2381.

[RiS2] S. D. Riemenschneider and Z. Shen, Construction of compactly supported biorthogonal wavelets in $L_2(\mathbb{R}^s)$, preprint (1997).

[RS1] A. Ron and Z. Shen, The Sobolev regularity of refinable functions, preprint (1997). `ftp://ftp.cs.wisc.edu/Approx` file `reg.ps`.

[RS2] A. Ron and Z. Shen, Affine systems in $L_2(\mathbb{R}^d)$: the analysis of the analysis operator, J. Functional Anal., **148** (1997), 408-447. Ftp site: `ftp://ftp.cs.wisc.edu/Approx` file `affine.ps`.

[RS3] A. Ron and Z. Shen, Affine systems in $L_2(\mathbb{R}^d)$ II: dual system, J. Fourier Anal. App., **3** (1997), 617-637. Ftp site: `ftp://ftp.cs.wisc.edu/Approx` file `dframe.ps`.

[RS4] A. Ron and Z. Shen, Compactly supported tight affine spline frames in $L_2(\mathbb{R}^d)$, Math. Comp., **67** (1998), 191–207. Ftp site: `ftp://ftp.cs.wisc.edu/Approx` file `tight.ps`.

[RS5] A. Ron and Z. Shen, Construction of compactly supported affine frames in $L_2(\mathbf{R}^\mathbf{d})$, in *Advances in wavelets*, Ka-Sing Lau ed, Springer-Verlag, 1998, 27-49.

[S] Y. Saad, Numerical Methods for Large Eigenvalue Problems, Manchester U. Press, Manchester, UK, 1992.

[SV] G.L. Sleijpen and H.A. Van der Vorst, A Jacobi-Davidson iteration method for linear eigenvalue problems, SIAM J. Matrix Anal. Appl. **17** (1996), 401-425.

[SY] D.C. Sorensen and C. Yang, A truncated RQ iteration for large scale eigenvalue calculations, SIAM J. Matrix Anal. Appl. **19** (1998), 1045-1073.

[V] L.F. Villemoes, Wavelet analysis of refinable equations, SIAM J. Math. Anal., **25** (1994), 1433-1460.