A windowed Fourier method for approximation of non-periodic functions on equispaced nodes

Rodrigo B. Platte

Abstract A windowed Fourier method is proposed for approximation of non-periodic functions on equispaced nodes. Spectral convergence is obtained in most of the domain, except near the boundaries, where polynomial least-squares is used to correct the approximation. Because the method can be implemented using partition of unit and domain decomposition, it is suitable for adaptive and parallel implementations and large scale computations. Computations can be carried out using fast Fourier transforms. Comparisons with Fourier extension, rational interpolation and least-squares methods are presented.

1 Introduction

The recovery of a function from a finite set of its values is a common problem in scientific computing and is one of the main underlying problems in the numerical solution of partial differential equations. This manuscript focuses on the special case of approximating functions from values sampled at evenly distributed points.

It is known that polynomial interpolants of smooth functions at equally spaced points do not necessarily converge, even if the function is analytic. Instead one may see wild oscillations near the endpoints, an effect known as the Runge phenomenon. Associated with this phenomenon is the exponential growth of the condition number of the interpolation process. Several other methods have been proposed for recovering smooth functions from uniform data, such as polynomial least-squares, rational interpolation, and radial basis functions; to name but a few. It is now known that these methods cannot converge at geometric (exponential) rates and remain stable for large data sets [11]. In practice, however, some methods perform remarkably well.

School of Mathematical and Statistical Sciences, Arizona State University, e-mail: rbp@asu.edu.
This work was supported in part by AFOSR FA9550-12-1-0393.

In this work we present a hybrid method based on *windowed Fourier* (WF) approximations combined with polynomial least-squares corrections near boundaries. The algorithm is an adaptation of the method presented in [10], in which a hybrid grid was used in the approximations – uniform nodes in the interior of the domain combined with Chebyshev points near the endpoints. In contrast, the algorithm proposed here relies strictly on equispaced grids. Besides describing a criterium for choosing parameters in the algorithm (window size, boundary layer correction, and polynomial degree), a generalized version of Hermite's error formula is used to compare the accuracy of the proposed algorithm with other known methods for the approximation of analytic functions.

The WF scheme can be closely related with Fourier continuation (or extension/embedding) methods [4, 6], which have been extensively explored recently. It is important to point out that Fourier extension requires a periodic continuation of the target function outside the domain of interest. The extension is not unique and different strategies have been presented in the literature to generate them. In [6], for example, an SVD based least squares approximation is used, while in [3] polynomials are used to periodically extend the function. Although an FFT based implementation is available for the SVD approach [8], it is restricted to extended domains that are at least twice as large (in 1D) as the domain of interest, a limitation that has implications on the oversampling rate for stable approximations. A detailed study of the tradeoffs between amount of oversampling and numerical stability has been recently presented in [1, 2]. The WF method, on the other hand, does not require least-squares approximations on the interior of the domain, with Fourier coefficients being computed by interpolation.

Along these lines, several other methods have been proposed to approximate functions from equispaced nodes with spectral-like accuracy. Examples can be found, for instance, in [5, 11]. Here we focus on describing the WF method and providing numerical experiments to demonstrate its performance.

2 Background and algorithm

For simplicity, we describe the scheme for approximations on a bounded interval. Computations in higher dimensions are carried out using tensor products. The WF method for equispaced points is motivated by [10], where a similar strategy was proposed as an alternative to traditional spectral methods. In that paper, polynomial approximations near the edges of the domain were computed on Chebyshev nodes, as the main focus was the solution of partial differential equations. In the present work, we replace polynomial interpolation with least-squares, and relax the restriction on the node distribution near the boundary.

A windowed Fourier approximation is illustrated in Figure 1. To approximate a non-periodic function u, using Fourier expansions, a smooth window function w is used. The window and its derivatives are close to zero at boundary points and the product uw can be accurately approximated by a truncated trigonometric series. The

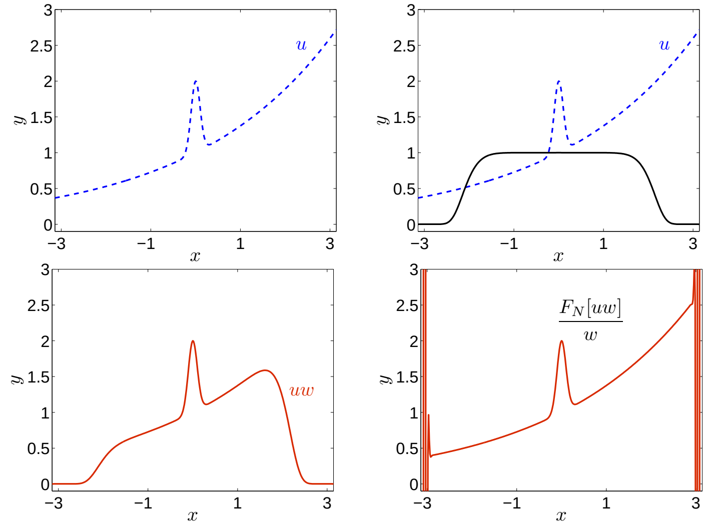

Fig. 1 From left to right: the function $u(x) = \exp(x/\pi) + \exp(-50x^2)$ (dashed), the window function w (solid), the product uw, and the Fourier approximation of uw devided by w.

function u can be recovered from this approximation by dividing it by w. Since w is close to zero near the boundaries, errors are amplified in that region. To correct the approximation near the ends of the interval, local polynomial approximations are used.

Although not pursued here, it is also possible to take advantage of the windowing process to decompose the domain if variable resolution is required. There are many possible window functions. As in [10], the method proposed here uses super-Gaussian window functions, $w(x) = \exp\left(-\alpha(x/\pi)^{2\lambda}\right)$, $x \in [-\pi, \pi]$, where λ is a positive integer and $\alpha \approx 52\ln 2$ is used in double precision. This choice of α ensures that $w(\pm\pi) \approx 2^{-52}$, which is the machine epsilon.

The difficulty in finding a suitable window is that a nonzero function that has all derivatives vanishing at a point cannot be analytic on any neighborhood of the interval in consideration. There must be a compromise between enforcing periodicity and the stiffness of the product uw. This aspect is investigated in detail in [10]. Figure 2 shows super–Gaussian window functions for $\lambda = 2,\ 5$, and 20. Notice that $\lambda = 20$ gives large support for the Fourier approximation but also stiff gradients near the boundaries. Although super-Gaussians are not compactly supported in infinite precision, we found that their Fourier sums have better convergence properties than C^∞ compactly supported window functions, such as $\exp(-1/(1-(x/\pi)^2)^\lambda)$. The lack of exact periodicity in derivatives in the case of super-Gaussians only affects convergence rates once the error falls below machine precision and their Fourier sums converge geometrically (exponentially) for all practical purposes.

Throughout this work we consider the truncated Fourier series,

$$F_N[u](x) = \sum_{n=-N}^{N}{}' \hat{u}_n \exp(inx), \qquad -\pi \le x \le \pi, \tag{1}$$

where the coefficients $\hat{u}_n$ are computed so that $u(x_j) = F_N[u](x_j)$ at $2N$ equally spaced nodes. The prime indicates that the terms $n = \pm N$ are multiplied by $\frac{1}{2}$. It is well known that the series converges exponentially fast, as $N \to \infty$, to smooth periodic functions.

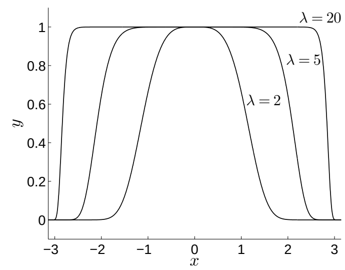

Fig. 2 Super-Gaussian window functions: $\exp(-32(x/\pi)^{2\lambda})$.

The accuracy and performance of the WF method depends on how well the product uw is approximated. Analysis presented in [10] shows that the number of modes in (1) required to approximate w, to a fixed accuracy, is linearly proportional to λ. In particular, approximations of w accurate to almost machine precision can be obtained with $N > 12\lambda$. This linear dependence can be explained using standard error estimates for Fourier interpolation and we refer to [10] for details.

Once the product uw is approximated, the recovery of u can be obtained by a point-wise division by w. Notice that the magnitude of the error in this process is inversely proportional to the values of $w(x)$, i.e.

$$|u(x) - F_N[uw](x)/w(x)| = |u(x)w(x) - F_N[uw](x)|\,/w(x).$$

Therefore, if corrections are made to the approximation in the regions where $w(x) < 0.05$, the error at the cutoff points would be about twenty times larger than at the center of the interval. Choosing the cutoff points, x_a and x_b, to satisfy $w(x_a) = w(x_b) = 0.05$, gives

$$x_a = -\pi\left((\ln 20)/\alpha\right)^{1/2\lambda} \text{ and } x_b = \pi\left((\ln 20)\alpha\right)^{1/2\lambda}. \tag{2}$$

Asymptotically, this means that the number of nodes in the correction regions remains nearly constant as $N \to \infty$.

For the 1–D case, the **algorithm** can be summarized as follows.

- Given $2N$ equispaced nodes on $[-\pi, \pi]$, choose λ to be $N/12$.
- Set the cutoff points according to (2).

- Approximate the product wu using (1). The approximation in $[x_a, x_b]$ is then given by $(F_N[uw](x))/w(x)$.
- Correct approximations in $[-\pi, x_a]$ and $[x_b, \pi]$ using polynomial least-squares. Here we choose the polynomial degree to be half the number of nodes in the correction regions.

The rate of convergence of scheme is limited by the least-squares polynomial correction near the ends of the domain. That region, however, shrinks as N is increased. Spectral accuracy is attained in most of the domain and fast convergence is expected for functions free of singularities or steep gradients near the boundary.

3 Accuracy for analytic functions

In this section we use a generalization of Hermite's error formula to compare the accuracy of different methods. To this end, we consider a general linear approximation of an analytic function f from its data values as

$$\mathcal{L}_{f,N}(x) := \sum_{j=1}^{N} f(x_j) L_j(x), \tag{3}$$

where the L_j are bounded functions. In the case of interpolation, L_j would be a cardinal functions.

Theorem 1. *Suppose f is analytic in a closed simply connected region R and C is a simple, closed, rectifiable curve that lies in R and encloses the interpolation points x_j, $j = 1, \ldots, N$. The error at x in the approximation (3) is*

$$f(x) - \mathcal{L}_{f,N}(x) = \frac{1}{2\pi i} \int_C \frac{f(z)}{z-x} r_N(z,x) dz, \tag{4}$$

$$r_N(z,x) = 1 - (z-x) \sum_{j=1}^{N} \frac{L_j(x)}{z - x_j}. \tag{5}$$

From (5) we can see that $r_N(z,x)$ can be interpreted as the relative error at x in approximating the function $1/(z-x)$ with (3). Additional details, including the proof, can be found in [9].

Using (4) we can bound the error. Under the assumptions of Theorem 1, assume that x and all interpolation points are in $[-1,1]$, then

$$\left\| f - \mathcal{L}_{f,N} \right\|_{[-1,1]} \le M_f \max_{x \in [-1,1]} |r_N(z,x)|, \text{ where } M_f = \frac{\text{arclength}(C) \max_{z \in C} |f(z)|}{2\pi \min_{x \in [-1,1],\, z \in C} |z-x|}.$$

Notice that M_f depends only on the function being approximated and the accuracy of the approximating scheme, including node distribution, is captured in r_N.

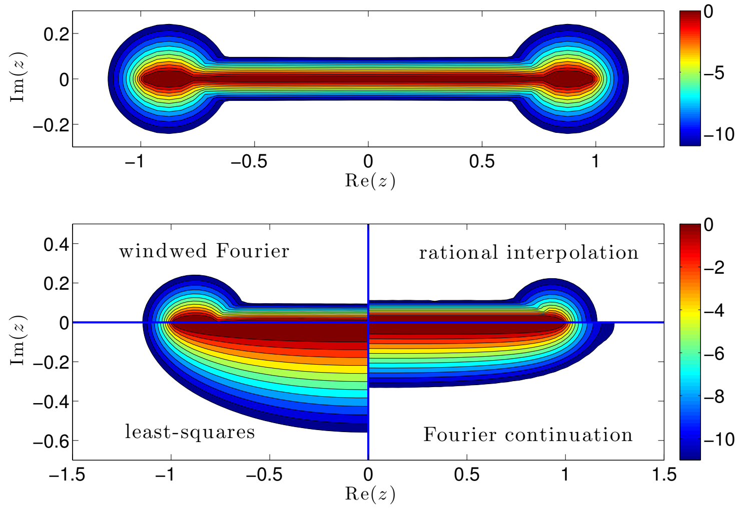

Fig. 3 **Top:** Level curves of $R_N(z) = \max_{x\in[-1,1]} |r_N(z,x)|$ in a log scale for WF with $N = 200$. Contours shown correspond to $10^{-11}, 10^{-10}, \ldots, 1$. **Bottom:** Same as the top plot, but with three other methods for comparison.

Figure 3 shows the contour levels of $R_N(z) := \max_{x\in[-1,1]} |r_N(z,x)|$. Results are shown for $N = 200$ and are qualitatively similar for other values of N. The top panel in Fig. 3 shows R_{200} for the WF method. It shows that approximations are very accurate in the interior of the domain. As could be expected, R_N decays more slowly near the end points due to the local polynomial correction in that region.

For reference, the bottom panel of Fig. 3 shows the corresponding values of R_N for three other methods: polynomial least-squares, rational interpolation, and Fourier continuation. These methods are known to work well on equispaced nodes. For the least-squares approximation, the degree of the polynomial was chosen to be approximately $4\sqrt{N}$ to ensure stable results (see e.g. [11, 12]). The rational interpolation approximation is computed using the method presented Floater and Hormann in [7] with degree 15. The Fourier continuation scheme was computed using the SVD approach and oversampled least-squares [8]. Extended domains were twice as large as each subdomain (see [8, 1] for details). Notice that the WF method compares favorably to all three methods as R_N takes smaller values in a larger part of the complex plane. This figure also shows that R_N is similar for the WF and rational interpolation methods.

Finally, we use four functions to test the performance of the WF method. Figure 4 shows the error as a function of N, the number of equispaced nodes, for each function. For reference, the error in polynomial interpolation on N Chebyshev nodes is also included. As predicted by the contours in Fig. 3, WF and rational interpolation present similar accuracy. Due to better resolution in the interior of the domain, the WF approximation of f_2 converges faster than polynomial interpolation on Chebyshev nodes. While Fourier continuation outperforms WF approximations for the oscillatory function f_1, WF is significantly more accurate for f_2 and f_3. Notice that f_4 has a singularity close to a boundary point, at $x = 1.05$, and as a consequence all methods converge at sub-geometric rates (and are much less accurate than interpolation at Chebyshev points). This result is also in good agreement with Fig 3.

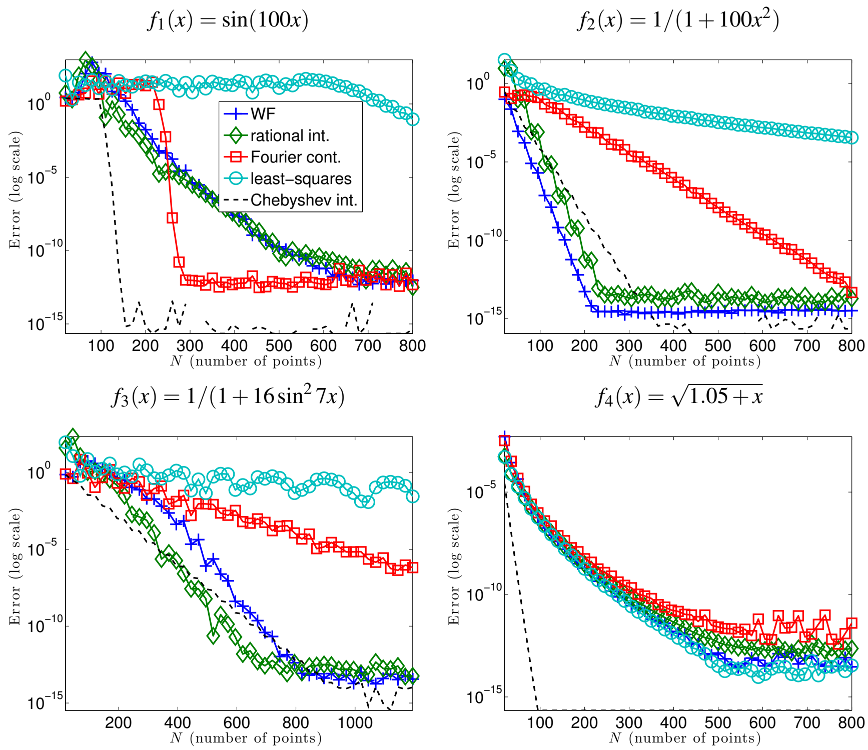

Fig. 4 Error decay in the approximation of four functions on the interval $[-1,1]$.

4 Concluding Remarks

It is important to point out that the WF method is not exponentially convergent for all analytic functions. The results presented in Figure 4 do not contradict the theorem in [11], which asserts that stable methods cannot converge at geometric rates for all functions that analytic in regions enclosing the interval of approximation. This is evident in the approximation of $f_4(x) = \sqrt{1.05+x}$, where the error plot shows sub-geometric decay for all methods (except for Chebyshev interpolation).

The approximation order of the method presented here is dominated by the polynomial least-squares approximation near the endpoints. Because the number of nodes on these regions is nearly constant (the size of boundary layer shrinks as the overall number of points is increased), the polynomial degree remains nearly constant as $N \to \infty$. For the parameters used to generate the plots in Figure 4, the degree for the polynomial corrections is 21 when $N \geq 150$. The algebraic convergence near the endpoints of the interval is reflected in the lobes of Fig. 3. Although the formal convergence of the method is sub-geometric, in many practical cases, the error decays exponentially fast for practical values of N . In the case of $f_2(x) = 1/(1+100x^2)$, for instance, the approximation error is dictated by the singularity near $x = 0$ and the correction regions are very accurately approximated by a high order polynomial.

In contrast to the other three numerical schemes used to obtain Fig. 4, the polynomial correction step in the algorithm being proposed here leads to an approximation that is not continuous. The jump size in the resulting approximation is of the order of the approximation error and can be a source of instability when solving PDEs. This issue has been partially addressed in [10], when the polynomial correction is carried out using Chebyshev interpolation, but has not yet been addressed when equispaced nodes are used.

References

1. B. Adcock, D. Huybrechs, and J. Martín-Vaquero. On the numerical stability of Fourier extensions. *Found. Comput. Math.*, pages 1–53, 2012.
2. B. Adcock and J. Ruan. Parameter selection and numerical approximation properties of Fourier extensions from fixed data. *J. Comput. Phys.*, 273:453–471, 2014.
3. N. Albin and O. P. Bruno. A spectral FC solver for the compressible Navier–Stokes equations in general domains i: Explicit time-stepping. *J. Comput. Phys.*, 230(16):6248–6270, 2011.
4. 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.
5. 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. *Commun. Comput. Phys.*, 5(2-4):484–497, 2009.
6. 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.
7. M. S. Floater and K. Hormann. Barycentric rational interpolation with no poles and high rates of approximation. *Numer. Math.*, 107:315–331, 2007.
8. M. Lyon. A fast algorithm for fourier continuation. *SIAM J. Sci. Comp.*, 33(6):3241–3260, 2011.
9. R. B. Platte. How fast do radial basis function interpolants of analytic functions converge? *IMA J. Numer. Anal.*, 31(4):1578–1597, 2011.
10. R. B. Platte and A. Gelb. A hybrid Fourier-Chebyshev method for partial differential equations. *J. Sci. Comput.*, 39(2):244–264, 2009.
11. R. B. Platte, L. N. Trefethen, and A. B. J. Kuijlaars. Impossibility of fast stable approximation of analytic functions from equispaced samples. *SIAM Rev.*, 53:308–318, 2011.
12. E. A. Rakhmanov. Bounds for polynomials with a unit discrete norm. *Ann. of Math. (2)*, 165(1):55–88, 2007.