

Parameter selection and numerical approximation properties of Fourier extensions from fixed data

Ben Adcock
Department of Mathematics
Purdue University
USA

Joseph Ruan
Department of Mathematics
Purdue University
USA

August 16, 2018

Abstract

Fourier extensions have been shown to be an effective means for the approximation of smooth, nonperiodic functions on bounded intervals given their values on an equispaced, or in general, scattered grid. Related to this method are two parameters. These are the extension parameter T (the ratio of the size of the extended domain to the physical domain) and the oversampling ratio η (the number of sampling nodes per Fourier mode). The purpose of this paper is to investigate how the choice of these parameters affects the accuracy and stability of the approximation. Our main contribution is to document the following interesting phenomenon: namely, if the desired condition number of the algorithm is fixed in advance, then the particular choice of such parameters makes little difference to the algorithm's accuracy. As a result, one is free to choose T without concern that it is suboptimal. In particular, one may use the value $T = 2$ – which corresponds to the case where the extended domain is precisely twice the size of the physical domain – for which there is known to be a fast algorithm for computing the approximation. In addition, we also determine the resolution power (points-per-wavelength) of the approximation to be equal to $T\eta$, and address the trade-off between resolution power and stability.

1 Introduction

In many problems, one is faced with the task of recovering a smooth function $f : [-1, 1] \rightarrow \mathbb{C}$ to high accuracy from its pointwise samples on an equispaced, or in general, scattered grid. This problem is challenging, unless the grid points have a specific distribution, since it is difficult to simultaneously ensure both rapid convergence and numerical stability. In particular, for equispaced data a result of Trefethen, Platte & Kuijlaars states that no stable method can converge faster than root-exponentially in the number of data points [13], and that any method with more rapid convergence must necessarily be unstable.

Nevertheless, it has been widely reported that so-called *Fourier extensions* (also known as *Fourier continuations*) lead to effective methods in practice for reconstructions from equispaced or scattered data [4, 5, 6, 7, 9, 10]. This was confirmed recently in [4] wherein it was shown that Fourier extensions (henceforth abbreviated to FEs) circumvent the stability barrier of [13] in a certain sense. Specifically, they converge down to a finite, but user-controlled, maximal accuracy.

In the method of FEs a function $f : [-1, 1] \rightarrow \mathbb{C}$ is approximated by a Fourier series of degree N defined on a larger domain $[-T, T]$, where $T > 1$ is a user-controlled parameter. Suppose now that the number of equispaced data points is equal to $2M + 1$ for some $M \geq N$. As shown in [4, 10], if $M = \eta N$ for some fixed *oversampling* ratio $\eta \geq 1$, then one may compute an accurate and

stable FE approximation of f from this data by a simple least-squares fit (see also §2 for details). Moreover, when the extension parameter T is equal to 2 – that is, the extended domain $[-T, T]$ is precisely twice the size of the physical domain $[-1, 1]$ – an algorithm developed by M. Lyon allows for the computation of the FE approximation in only $\mathcal{O}(M(\log M)^2)$ operations [11]. Note that this fast algorithm relies on the particular symmetries of FEs when $T = 2$.

From a practical standpoint, it is vitally important to understand how to choose the parameters T and η . The purpose of this paper is to address this issue. In particular, we seek to determine how choices of these parameters affect both the stability and the accuracy of the FE approximation. Note that, due to the aforementioned fast algorithm, there is a seeming advantage to using the value $T = 2$. However, is choice optimal vis-a-vis the other properties of the algorithm, namely, convergence and stability? Or does another choice (albeit lacking a fast algorithm) give better numerical performance in these respects? We shall provide answers to these questions.

It is known that different choices of T affect the intrinsic approximation properties of the FE approximation system, i.e. the space of trigonometric polynomials of degree N on the extended domain $[-T, T]$ [3, 4, 5]. For example, when T is close to 1, the FE approximation system possesses better *resolution* power for oscillatory functions [3]. However, it is perhaps not surprising that such choices also require larger oversampling parameters η to maintain the algorithm’s stability. When M is fixed, a larger η means a smaller parameter N , and therefore the best approximation error in the above subspace, which is determined by the size of N , is correspondingly larger.

From this argument, it is apparent that a balance must be struck between T and η so as to preserve accuracy and stability. The main result we obtain in this paper through numerical experiment is that these two effects precisely counteract each other. Specifically, if the desired condition number of the algorithm is fixed in advance, then, provided T is not too large, the precise choice of T makes no substantial difference to the accuracy of the Fourier extension algorithm. Smaller T is exactly offset by the requirement of a larger value of η to preserve stability.

From this result, we are able to draw two main conclusions. First, any attempt to optimize T will only bring limited, and most likely highly function-dependent, success. Second, since the choice of T makes little difference, one may safely use $T = 2$, and the resulting fast algorithm, without worrying that this choice may be suboptimal. We remark in passing that some previous insight into the effect of the parameters was given by Bruno et al. in [7]. However, this was largely carried out for specific functions.

On the face of it, the conclusion we draw may appear surprising, or at the very least, a peculiar phenomenon isolated to the particular choice of equispaced data. After further numerical experiments, we conclude that this phenomenon is actually quite widespread. Specifically, we show exactly the same results for both scattered nonequispaced data, as well as Fourier data. Hence we conclude that unless the data is chosen specifically to favour a particular choice of T (see §5 for an example of such data), the value of T makes little difference to the algorithm. We note here that, much as in the case of equispaced data, a fast Fourier extension algorithm for scattered data has also been developed in the case $T = 2$ [12].

In some applications, including the numerical solution of PDEs, an important question about an approximation algorithm is that of resolution power. Specifically, how many measurements (e.g. equispaced function samples) are required to recover an oscillation of frequency ω . This topic was first investigated rigorously by Gottlieb & Orszag [8], who popularized the concept of *points-per-wavelength*. Through our experiments we establish that this quantity for FE approximations is given by the product of T and η . We give theoretical arguments as to why this should be the case, and in the case $T = 2$ (and therefore, by the above discussion, all values of T) provide numerical results assessing the tradeoff between resolution power and numerical stability.

As the reader will have noticed, our aim in this paper is to investigate FEs through numerical

experiment. The main conclusion we draw is based solely on the result of these experiments. Although we do present some mathematical insight as to why it should hold, this is a ways short of a proof. We leave this as a topic for future work. As we explain in §6, this will likely require an intricate analysis of the singular values and singular vectors of a certain matrix related to Slepian’s prolate matrix [14, 17], which is beyond the scope of this paper. Nevertheless, we feel the conclusion we draw, albeit without a proof, is of substantial independent interest for anyone seeking to use FEs in practice.

2 Fourier extensions

Our concern in this paper is the approximation of functions defined on compact intervals, which without loss of generality we take to be $[-1, 1]$. The method of Fourier extensions (FEs) is based on approximating such functions using a Fourier series defined on an extended interval $[-T, T]$, where $T > 1$ is the so-called extension parameter. In other words, given $N \in \mathbb{N}$ we compute an approximation to f from the subspace

$$\mathcal{G}_N^{(T)} := \text{span} \{ \phi_n : |n| \leq N \}, \quad \phi_n(x) = e^{i \frac{n\pi}{T} x}.$$

2.1 Approximation properties of the subspace $\mathcal{G}_N^{(T)}$

We now present several results concerning the intrinsic approximation properties of the subspace $\mathcal{G}_N^{(T)}$. We use the notation $H^k(-1, 1)$ for the standard Sobolev space on an interval $(-1, 1)$, where $k \geq 0$. We denote the corresponding norm by $\|\cdot\|_{H^k(-1,1)}$.

Theorem 2.1 ([3]). *Let $T_0 > 1$ and suppose that $f \in H^k(-1, 1)$ for some $k \geq 0$. Then, for each $N \in \mathbb{N}$ and $T \geq T_0$, there exists a $\phi \in \mathcal{G}_N^{(T)}$ such that*

$$\|f - \phi\|_{L^2(-1,1)} \leq C(k, T_0) \left(\frac{N\pi}{T} \right)^{-k} \|f\|_{H^k(-1,1)}, \quad \|\phi\|_{L^2(-T,T)} \leq C(k, T_0) \|f\|_{H^k(-1,1)},$$

for some constant $C(k, T_0)$ depending on k and T_0 only.

This theorem asserts *algebraic* convergence of the best approximations in $\mathcal{G}_N^{(T)}$ when f has k derivatives, and *superalgebraic* convergence whenever f is smooth. Note that it also implies the existence of a function $\phi \in \mathcal{G}_N^{(T)}$ which gives such convergence rates, and which cannot grow too large on the extended domain $[-T, T]$. This will be of significance in §2.3.

Our next result confirms *geometric* convergence of best approximations in $\mathcal{G}_N^{(T)}$ in the case that f is analytic. To state this result, we first recall the definition of a Bernstein ellipse:

$$\mathcal{B}(\rho) = \left\{ \frac{1}{2} \left(\rho^{-1} e^{i\theta} + \rho e^{-i\theta} \right) : \theta \in [-\pi, \pi] \right\} \subseteq \mathbb{C}, \quad \rho > 1.$$

As discussed in [4, 9], Fourier extensions can be viewed as polynomial approximations in the mapped variable $z = m(x)$, where

$$m(x) = 2 \frac{\cos \frac{\pi}{T} x - \cos \frac{\pi}{T}}{1 - \cos \frac{\pi}{T}} - 1. \tag{2.1}$$

Note that m maps $[0, 1]$ to $[-1, 1]$ bijectively. Since the convergence of polynomial approximations of analytic functions is determined by Bernstein ellipses, it makes sense to introduce the new regions

$$\mathcal{D}(\rho) = m^{-1}(\mathcal{B}(\rho)), \quad \rho > 1,$$

We now have the the following theorem:

Theorem 2.2 ([9]). *Let $T > 1$ be given and suppose that f is analytic in $\mathcal{D}(\rho')$ for some $\rho' > 1$ and continuous on its boundary. Then, for each $N \in \mathbb{N}$, there exists a $\phi \in \mathcal{G}_N^{(T)}$ such that*

$$\|f - \phi\|_{L^\infty(-1,1)} \leq \frac{c_f(T)}{1 - \rho} \rho^{-N}, \quad (2.2)$$

where $c_f(T) > 0$ is proportional to $\max_{z \in D(\rho)} |f(z)|$,

$$\rho = \min \{ \rho', E(T) \},$$

and $E(T) = \cot^2 \left(\frac{\pi}{4T} \right)$. Moreover, ϕ satisfies

$$\|\phi\|_{L^\infty(-T,T)} \leq c_f(T) (E(T)/\rho)^N. \quad (2.3)$$

This theorem establishes geometric convergence of best approximations in $\mathcal{G}_N^{(T)}$. However, (2.3) suggests that in order to obtain such a convergence rate, one may have to allow for exponential growth of the corresponding ϕ in the extended domain $[-T, T]$ whenever $\rho < E(T)$. We shall return to this observation in §2.3.

We remark also that Theorem 2.2 asserts that the maximal rate of geometric convergence is limited to $E(T)$, even if f is entire. This is due to the mapping m^{-1} which introduces a square-root type singularity and thereby limits the overall rate of convergence. See [4, 9] for a discussion.

2.2 Fourier extensions from equispaced data

The concern of the majority of this paper is the approximation of a function f from its values

$$f(m/M), \quad m = -M, \dots, M,$$

on an equispaced grid of $2M + 1$ points. For convenience, let us define the operator

$$S_M : L^\infty(-1, 1) \rightarrow \mathbb{C}^{2M+1}, \quad f \mapsto \frac{1}{\sqrt{M}} (f(m/M))_{m=-M}^M.$$

We refer to S_M as the *sampling* operator. Given the vector $S_M(f)$ of samples of f , we construct its FE approximation in the standard way via a least-squares data fit [4, 5, 7, 9, 10]. Let $N \leq M$ be given. Then we define the FE approximation as follows:

$$F_{N,M}^{(T)}(f) := \operatorname{argmin}_{\phi \in \mathcal{G}_N^{(T)}} \sum_{|m| \leq M} |f(m/M) - \phi(m/M)|^2, \quad (2.4)$$

or more succinctly,

$$F_{N,M}^{(T)}(f) := \operatorname{argmin}_{\phi \in \mathcal{G}_N^{(T)}} |S_M(f - \phi)|,$$

where $|\cdot|$ denotes the usual Euclidean norm on \mathbb{C}^{2M+1} . Note that $F_{N,M}$ is an operator with domain $L^\infty(-1, 1)$ and range $\mathcal{G}_N^{(T)}$. Moreover, if we denote

$$F_{N,M}^{(T)}(f) = \sum_{|n| \leq N} a_n e^{i \frac{n\pi}{T} x},$$

then the vector $\mathbf{a} = (a_n)_{n=-N}^N$ of FE coefficients is the solution of the least squares problem

$$\mathbf{a} = \operatorname{argmin}_{\mathbf{c} \in \mathbb{C}^{2N+1}} \left| A^{(T)} \mathbf{c} - S_M(f) \right|, \quad (2.5)$$

where $A^{(T)} \in \mathbb{C}^{(2M+1) \times (2N+1)}$ has entries

$$(A^{(T)})_{m,n} = \frac{1}{\sqrt{M}} e^{i \frac{nm\pi}{T}}, \quad |m| \leq M, \quad |n| \leq N.$$

Note that the normalization $1/\sqrt{M}$ in both S_M and $A^{(T)}$ means that the entries of the normal matrix $(A^{(T)})^* A^{(T)}$ are Riemann sum approximations to the Gram matrix of the functions $\phi_n(x)$. This ensures that the singular values of $A^{(T)}$ lie between 0 and 1 for large M , which, since we typically consider truncated SVDs with a fixed truncation parameter (see later), ensures that there is no linear drift in the error for large M .

For convenience, let us now introduce some additional notation. Let

$$L_{N,M}^{(T)} : \mathbb{C}^{2M+1} \rightarrow \mathbb{C}^{2N+1},$$

be defined by

$$L_{N,M}^{(T)}(\mathbf{b}) = \operatorname{argmin}_{\mathbf{c} \in \mathbb{C}^{2N+1}} \left| A^{(T)} \mathbf{c} - \mathbf{b} \right|, \quad (2.6)$$

and let

$$R_N^{(T)} : \mathbb{C}^{2N+1} \rightarrow \mathcal{G}_N^{(T)}, \quad \mathbf{a} = (a_n)_{n=-N}^N \mapsto \sum_{|n| \leq N} a_n e^{i \frac{n\pi}{T} x}.$$

Note that $F_{N,M}^{(T)} = R_N^{(T)} \circ L_{N,M}^{(T)} \circ S_M$.

As discussed in [4], the algebraic least-squares problem to be solved in (2.6) is highly ill-conditioned. When applied to (2.6), different numerical algorithms may consequently give somewhat different results. For this reason, it is important to specify the solver used. In the majority of this paper, as has been previously considered in [4, 5, 10], we solve (2.6) by using truncated singular value decompositions (SVDs). If $U\Sigma V^*$ denotes the SVD of $A^{(T)}$, where Σ is the diagonal matrix of singular values $\sigma_1 \geq \sigma_2 \geq \dots$, then we correspondingly define

$$L_{N,M}^{(T,\epsilon)}(\mathbf{b}) = V \Sigma^{(\epsilon)} U^* \mathbf{b},$$

where $\Sigma^{(\epsilon)}$ is the diagonal matrix with n^{th} entry $1/\sigma_n$ if $\sigma_n > \epsilon$ and 0 otherwise. Here $\epsilon > 0$ is the truncation parameter, which we take to be 10^{-13} unless specified otherwise. We denote the corresponding FE by $F_{N,M}^{(T,\epsilon)}(f)$.

Having said this, we note that the quantities introduced below for studying equispaced FEs – namely, the condition number and numerical defect constant – are not specific to the SVD algorithm. In particular, one can compute such quantities for each different numerical solver and thereby directly compare the effectiveness of an equispaced FE resulting from an SVD with an equispaced FE computed using *Matlab's* \code{\backslash} or *Mathematica's* `LeastSquares` commands, for example. We return to this briefly in §3.3.

Remark 2.3 Regardless of the solver used, it is important that the least-squares (2.6) is regularized when solved numerically. This is done by the parameter ϵ with the SVD approach, or automatically when using an blackbox least-squares solver such as *Matlab's* \code{\backslash} or *Mathematica's* `LeastSquares`.

As shown in [4], the ‘exact’ FE mapping $f \mapsto F_{N,M}^{(T)}(f)$, i.e. that obtained by solving (2.6) in infinite precision, is ill-conditioned and suffers from a Runge phenomenon unless the number of equispaced points M scales quadratically with N . Such severe scaling is undesirable, and is due solely to the behaviour of the Fourier series corresponding to singular vectors with small singular values. Fortunately, when the system (2.6) is regularized and $F_{N,M}^{(T,\epsilon)}$ is computed, this scaling drops

to linear in N . Moreover, the Fourier series of the excluded singular values are precisely those which are small on the domain $[-1, 1]$ but large on $[-T, T] \setminus [-1, 1]$. Thus, their exclusion has little effect on the approximation of f . Note that a similar behaviour is also witnessed when different solvers are used for (2.6), such as those listed above.

2.3 Condition number and error bounds for equispaced FE approximations

We now provide estimates for the accuracy and stability of $F_{N,M}^{(T)}(f)$. The key point is that these formulae involve constants which can be computed numerically. This will be discussed in the next section. First, however, we require the following assumption:

Assumption. The operator $L_{N,M}^{(T)}$ defined by solving (2.6) with a standard numerical solver (e.g. truncated SVDs) is approximately a linear operator.

Note that the exact, i.e. infinite precision, version of $L_{N,M}^{(T)}$ is of course a linear operator. Hence it is not unreasonable that its finite precision counterpart acts in the same way. Observe also that this assumption implies that the overall numerical FE operator $F_{N,M}^{(T)}$ is also linear.

With this in hand, we can now define the condition number in the usual way:

Definition 2.4 (Condition number). *The (absolute) condition number of the equispaced FE approximation $F_{N,M}^{(T)}$ is given by*

$$\kappa_{N,M}^{(T)} = \max_{\substack{\mathbf{b} \in \mathbb{C}^{2M+1} \\ \mathbf{b} \neq 0}} \left\{ \frac{\|R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{b})\|_\infty}{|\mathbf{b}|_\infty} \right\}. \quad (2.7)$$

Here $\|g\|_\infty = \sup_{x \in [-1, 1]} |g(x)|$ is the uniform norm on $[-1, 1]$ for $g \in L^\infty(-1, 1)$ and $|\mathbf{b}|_\infty = \max_{|m| \leq M} |b_m|$ for $\mathbf{b} = (b_m)_{|m| \leq M} \in \mathbb{C}^{2M+1}$.

We remark that $\kappa_{N,M}^{(T)}$ is the absolute condition number, as opposed to the more standard relative condition number [15]. It measures the absolute sensitivity of the FE to perturbations in the samples of f , and transpires to be substantially easier to compute in practice. Note also that the definition implicitly assumes linearity of the mapping $L_{N,M}^{(T)}$.

We now consider the approximation error. For this we require the following definition:

Definition 2.5 (Numerical defect constant). *The numerical defect constant of the equispaced FE is given by*

$$\lambda_{N,M}^{(T)} = \max_{\substack{\mathbf{a} \in \mathbb{C}^{2N+1} \\ \mathbf{a} \neq 0}} \left\{ \frac{\|R_N^{(T)}(\mathbf{a} - L_{N,M}^{(T)} \circ S_M \circ R_N^{(T)}(\mathbf{a}))\|_\infty}{|\mathbf{a}|_\infty} \right\}. \quad (2.8)$$

Before showing the relevance of this constant to error bounds, let us first consider its meaning. Recall that each vector \mathbf{a} corresponds uniquely to a function $\phi \in \mathcal{G}_N^{(T)}$ given by $\phi = R_N^{(T)} \mathbf{a}$. Thus the numerator in (2.8) reads

$$\left\| \phi - F_{N,M}^{(T)}(\phi) \right\|_\infty.$$

In infinite precision, the FE operator $F_{N,M}^{(T)}$ satisfies $F_{N,M}^{(T)}(\phi) = \phi$ for $\phi \in \mathcal{G}_N^{(T)}$. In other words, it is a projection. Hence the numerical defect constant measures how close the numerical, i.e. finite precision, FE operator is to possessing this property.

We are now able to provide an error bound for $F_{N,M}^{(T)}(f)$:

Lemma 2.6. *Let $f \in L^\infty(-1, 1)$ and suppose that $F_{N,M}^{(T)}(f)$ is given by (2.4). Then*

$$\|f - F_{N,M}^{(T)}(f)\|_\infty \leq \inf_{\mathbf{a} \in \mathbb{C}^{2N+1}} \left\{ \left(1 + \kappa_{N,M}^{(T)}\right) \|f - R_N^{(T)}(\mathbf{a})\|_\infty + \lambda_{N,M}^{(T)} |\mathbf{a}|_\infty \right\}, \quad (2.9)$$

where $\lambda_{N,M}^{(T)}$ is as in (2.8).

Proof. Let $\mathbf{a} \in \mathbb{C}^{2N+1}$ be arbitrary and write $\phi = R_N^{(T)}(\mathbf{a}) \in \mathcal{G}_N^{(T)}$. Then, using linearity of $F_{N,M}^{(T)}$ (Assumption 1), we obtain

$$\|f - F_{N,M}^{(T)}(f)\|_\infty \leq \|f - \phi\|_\infty + \|F_{N,M}^{(T)}(f - \phi)\|_\infty + \|\phi - F_{N,M}^{(T)}(\phi)\|_\infty.$$

We consider the latter two terms separately. For the first, note that

$$\|F_{N,M}^{(T)}(f - \phi)\|_\infty = \|R_N^{(T)} \circ L_{N,M}^{(T)} \circ S_M(f - \phi)\|_\infty \leq \kappa_{N,M}^{(T)} |S_M(f - \phi)|_\infty \leq \kappa_{N,M}^{(T)} \|f - \phi\|_\infty.$$

This gives the corresponding second term in (2.9). We now consider the other term. We have

$$\left\| \phi - F_{N,M}^{(T)}(\phi) \right\|_\infty = \left\| R_N^{(T)} \left(\mathbf{a} - L_{N,M}^{(T)} \circ S_M \circ R_M^{(T)}(\mathbf{a}) \right) \right\|_\infty \leq \lambda_{N,M}^{(T)} |\mathbf{a}|_\infty,$$

as required. \square

Let us now interpret this error bound. In §2.4 we shall observe numerically that

$$\kappa_{M/\eta, M}^{(T)} \sim \tilde{\kappa}_{T, \eta} \log M, \quad \lambda_{M/\eta, M}^{(T)} \sim \tilde{\lambda}_{T, \eta} M, \quad M \rightarrow \infty, \quad (2.10)$$

where $\tilde{\kappa}_{T, \eta}$ and $\tilde{\lambda}_{T, \eta}$ are independent of M . Furthermore, the ratio $\mu = \tilde{\lambda}_{T, \eta} / \tilde{\kappa}_{T, \eta}$ is roughly 10^{-13} in magnitude, regardless of the choice of T or η . Hence, one has the estimate

$$\|f - F_{N,M}^{(T)}(f)\|_\infty \leq \tilde{\kappa}_{\eta, T} M \inf_{\mathbf{a} \in \mathbb{C}^{2N+1}} \left\{ \|f - R_N^{(T)}(\mathbf{a})\|_\infty + \mu |\mathbf{a}|_\infty \right\}, \quad N = M/\eta. \quad (2.11)$$

The key aspect of this bound is that it separates the error into two component. The first, namely, $\tilde{\kappa}_{\eta, T}$, is determined by the parameters $\eta = M/N$ and T , and is independent of the function f . Moreover, as we see next, it can be computed numerically. The second, i.e. the term

$$E_N(f) := \inf_{\mathbf{a} \in \mathbb{C}^{2N+1}} \left\{ \|f - R_N^{(T)}(\mathbf{a})\|_\infty + \mu |\mathbf{a}|_\infty \right\}, \quad (2.12)$$

is crucially independent of η and depends only on the intrinsic approximation properties of the subspace $\mathcal{G}_N^{(T)}$ and the smoothness of f . In particular, combining Lemma 2.6 with Theorems 2.1 and 2.2, we immediately obtain the following:

Corollary 2.7. *Let $E_N(f)$ be given by (2.12). If $f \in \mathbf{H}^k(-1, 1)$ then*

$$E_N(f) \leq \min_{0 \leq l \leq k} \left\{ C(l, T_0) \|f\|_{\mathbf{H}^l(-1, 1)} \left(\left(\frac{N\pi}{T} \right)^{-l} + \mu \right) \right\}, \quad (2.13)$$

for each $T \geq T_0$, where T_0 and $C(l, T_0)$ are as in Theorem 2.1. Moreover if f is analytic in $\mathcal{D}(\rho)$ and continuous on its boundary, then one has

$$E_N(f) \leq c_f(T) \rho^{-N} (1 + \mu E(T)^N), \quad (2.14)$$

where $c_f(T)$ is as in Theorem 2.2.

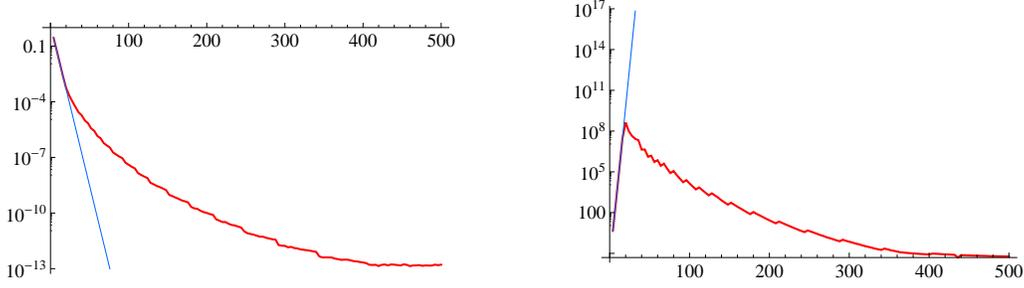


Figure 1: Best approximation error (left) and coefficient norm (right) for the function $f(x) = \frac{1}{40-39x}$ in finite precision (thicker line) and infinite precision (thinner line). The latter was computed in *Mathematica* using additional precision.

This corollary explains the behaviour of $E_N(f)$ in both finite and infinite precision. In infinite precision, where $\mu = 0$, the bound (2.14) shows geometric decay of $E_N(f)$ for all N at a rate equal to ρ . In finite precision, however, the small, but nonzero constant μ dramatically alters the convergence. Geometric decay still occurs for small N , when the term $\mu E(T)^N$ in (2.14) is small, but once $N \geq N_0 = -\log \mu / \log E(T)$ the right-hand side of (2.14) begins to increase. For $N \geq N_0$, $E_N(f)$ no longer decays geometrically. Instead, its decay is described by the bound (2.13). Specifically, algebraic decay in N occurs down to a maximal achievable accuracy on the order of μ . Note that the constant term $C(l, T) \|f\|_{H^l(-1,1)}$ usually grows with l , thus as $E_N(f)$ approaches μ the effective rate of decay usually lessens.

This behaviour is illustrated in Figure 1. As we see, geometric convergence in infinite precision requires geometric growth of the coefficient vector \mathbf{a} . Conversely, in finite precision, such convergence is sacrificed for algebraic convergence whilst maintaining a bounded coefficient norm. We refer to [4] for a more detailed discussion.

2.4 Computing the condition number and numerical defect constant

Whilst Corollary 2.7 explains the decay of $E_N(f)$, in order to understand the error of the equispaced FE $F_{N,M}^{(T)}(f)$ we need to determine the magnitudes of $\kappa_{N,M}^{(T)}$ and $\lambda_{N,M}^{(T)}$. The former also determines the stability of $F_{N,M}^{(T)}$. As we discuss further in §6, it is as of yet unknown how to do this analytically, hence we now resort to numerical investigations. For this we need a means of computing $\kappa_{N,M}^{(T)}$ and $\lambda_{N,M}^{(T)}$. This follows from the next two lemmas:

Lemma 2.8. *Let \mathbf{e}_m , $|m| \leq M$, be the canonical basis for \mathbb{C}^{2M+1} . Then*

$$\kappa_{N,M}^{(T)} = \sup_{x \in [-1,1]} \sum_{|m| \leq M} \left| R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m)(x) \right|.$$

Proof. If $\mathbf{b} = (b_m)_{|m| \leq M} \in \mathbb{C}^{2M+1}$, we may write $\mathbf{b} = \sum_{|m| \leq M} b_m \mathbf{e}_m$. By linearity

$$R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{b}) = \sum_{|m| \leq M} b_m R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m),$$

and therefore

$$\kappa_{N,M}^{(T)} \leq \sup_{x \in [-1,1]} \sum_{|m| \leq M} |R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m)|.$$

Conversely,

$$\kappa_{N,M}^{(T)} = \sup_{x \in [-1,1]} \max_{\substack{\mathbf{b} \in \mathbb{C}^{2M+1} \\ \mathbf{b} \neq \mathbf{0}}} \frac{\left| \sum_{|m| \leq M} b_m R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m)(x) \right|}{\|\mathbf{b}\|_\infty}.$$

We now set b_m equal to the complex sign of $R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m)(x)$ to deduce the lower bound. \square

This lemma allows for approximate computation of $\kappa_{N,M}^{(T)}$. Let $K \in \mathbb{N}$ be given and define

$$x_k = \frac{T(k-1)}{K} - 1, \quad k = 1, \dots, K_T,$$

where

$$K_T = \left\lfloor \frac{2K}{T} + 1 \right\rfloor.$$

Note that $\{x_k\}_{k=1}^{K_T}$ is a set of K_T equispaced nodes in $[-1, 1]$. Therefore

$$\kappa_{N,M}^{(T)} = \lim_{K \rightarrow \infty} \kappa_{N,M,K}^{(T)},$$

where

$$\kappa_{N,M,K}^{(T)} = \max_{k=1, \dots, K_T} \sum_{|m| \leq M} \left| R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m)(x_k) \right|,$$

is a computable quantity. We remark also $\kappa_{N,M,K}^{(T)}$ can be computed efficiently using Fast Fourier Transforms (FFTs), since the functions $R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{e}_m)(x)$ are Fourier series and $\{x_k\}_{k=1}^{K_T}$ are appropriately constructed equispaced nodes. Throughout this paper we shall consistently use the value $K = 2^{15}$ in our numerical experiments.

We use a similar approach in order to compute the numerical defect constant. Analogously to Lemma 2.8, we have the following:

Lemma 2.9. *Let \mathbf{e}_n , $|n| \leq N$, be the canonical basis for \mathbb{C}^{2N+1} . Then*

$$\lambda_{N,M}^{(T)} = \sup_{x \in [-1,1]} \sum_{|n| \leq N} \left| W_{N,M}^{(T)}(\mathbf{e}_n)(x) \right|,$$

where $W_{N,M}^{(T)} = R_N^{(T)} - R_N^{(T)} \circ L_{N,M}^{(T)} \circ S_M \circ R_N^{(T)}$.

Much as before, we may now write

$$\lambda_{N,M}^{(T)} = \lim_{K \rightarrow \infty} \lambda_{N,M,K}^{(T)}, \quad \lambda_{N,M,K}^{(T)} = \max_{k=1, \dots, K_T} \sum_{|n| \leq N} \left| W_{N,M}^{(T)}(\mathbf{e}_n)(x_k) \right|.$$

where the latter can once more be computed efficiently using FFTs.

Having demonstrated how to compute $\kappa_{N,M}^{(T)}$ and $\lambda_{N,M}^{(T)}$, in Figures 2 and 3 we present the result of such computations for the case where the FE is computed using an SVD with tolerance $\epsilon = 10^{-13}$. The results confirm the scaling (2.10) for these quantities. Moreover, these results also show that the quantity

$$\mu_{M/\eta, M}^{(T)} = \left(\frac{\lambda_{M/\eta, M}^{(T)}}{\kappa_{M/\eta, M}^{(T)}} \right) \left(\frac{\log M}{M} \right) \leq \mu, \quad \forall M \gg 1, \quad (2.15)$$

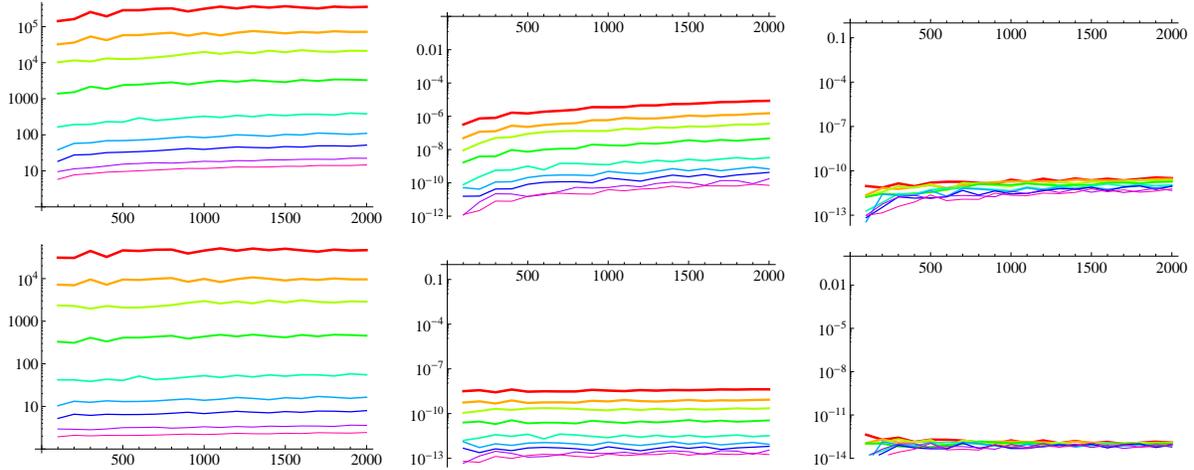


Figure 2: Top row: the quantities $\kappa_{M/\eta, M}^{(T, \epsilon)}$ (left), $\lambda_{M/\eta, M}^{(T, \epsilon)}$ (middle) and $\mu_{M/\eta, M}^{(T, \epsilon)} := \lambda_{M/\eta, M}^{(T)} / \kappa_{M/\eta, M}^{(T)}$ (right) against M for $\eta = 1, 1.125, 1.25, 1.5, 2, 2.5, 3, 4, 5$ (thickest to thinnest), $\epsilon = 10^{-13}$ and $T = 2$. Bottom row: the same quantities scaled by $\log M$, M and $M/\log M$ respectively.

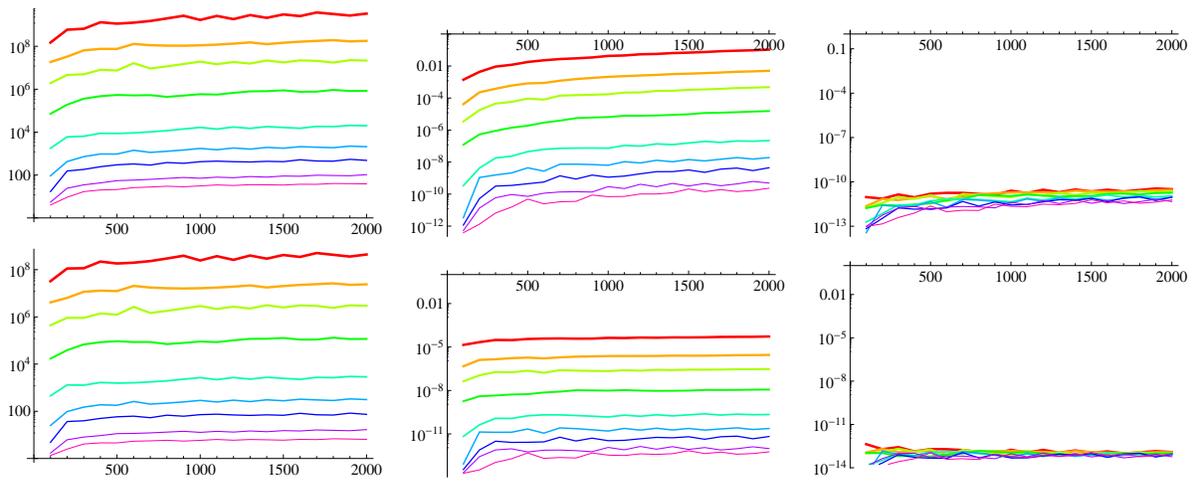


Figure 3: Top row: the quantities $\kappa_{M/\eta, M}^{(T, \epsilon)}$ (left), $\lambda_{M/\eta, M}^{(T, \epsilon)}$ (middle) and $\lambda_{M/\eta, M}^{(T)} / \kappa_{M/\eta, M}^{(T)}$ (right) against M for $\eta = 1, 1.125, 1.25, 1.5, 2, 2.5, 3, 4, 5$ (thickest to thinnest), $\epsilon = 10^{-13}$ and $T = 1.25$. Bottom row: the same quantities scaled by $\log M$, M and $M/\log M$ respectively.

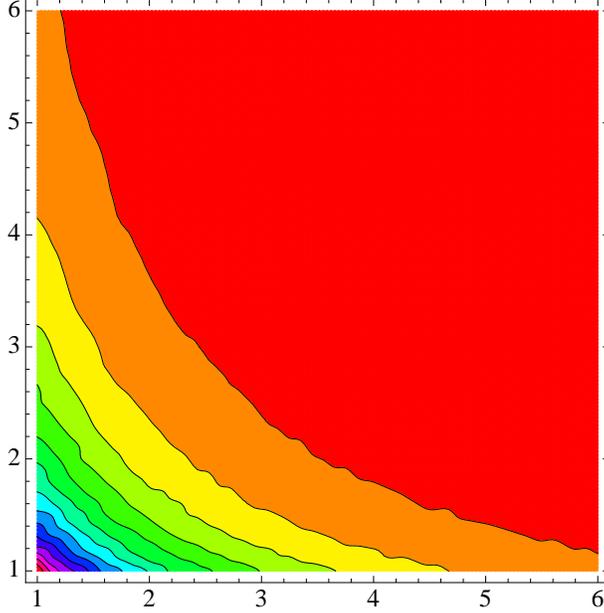


Figure 4: Contour plot of $\kappa_{N,\eta N}^{(T,\epsilon)}$ against $1 < T \leq 6$ (horizontal axis) and $1 \leq \eta \leq 6$ (vertical axis) for $N = 250$ and $\epsilon = 10^{-13}$.

where μ is roughly 10^{-13} in magnitude, regardless of the choice of T and η .

Note that in both cases a larger oversampling ratio η leads to a smaller condition number $\kappa_{M/\eta,M}^{(T,\epsilon)}$ and numerical defect constant $\lambda_{M/\eta,M}^{(T,\epsilon)}$. Moreover, the larger value of T , in this case, $T = 2$, has a smaller condition number for the same oversampling value than the smaller value $T = 1.25$. The purpose of the next section is to investigate the exact nature of these relative scalings.

Remark 2.10 Several previous papers have investigated quantities similar to $\kappa_{N,M}^{(T)}$ and $\lambda_{N,M}^{(T)}$. In [4] and [10], quantities based on the exact singular values and vectors of the matrix $A^{(T)}$ were investigated using high-precision numerical computations. The approach we take above differs from these studies in two aspects. First, $\kappa_{N,M}^{(T)}$ and $\lambda_{N,M}^{(T)}$ can be formulated for any numerical solver, not just truncated SVDs. Second, when truncated SVDs are used, they incorporate the numerical errors in the calculation of the singular values and vectors. Since the matrix $A^{(T)}$ is ill-conditioned, these errors cannot be assumed to be insignificant.

3 Numerical investigation

We now suppose that the FE is computed using an SVD as described in §2.2. Our aim is to examine the behaviour of $\kappa_{N,M}^{(T,\epsilon)}$, and later $\lambda_{N,M}^{(T,\epsilon)}$, with respect to T and the *oversampling* ratio $\eta = M/N$, and how this affects the accuracy of the corresponding FE approximation.

In Figure 4 we give a contour plot of $\kappa_{N,\eta N}^{(T,\epsilon)}$ as function of T and η . As is evident, increasing either η or T leads to a smaller condition number. This suggests that in practice, a balance must be struck between η and T so as to get a good condition number whilst retaining a good approximation properties (recall that larger values of T possess worse resolution power [4], whereas larger values of η yield worse approximations for a fixed budget of M equispaced data points).

Before investigating this interplay further, let us briefly explain why the condition number behaves in this way. Clearly, increasing η results in a smaller value of $N = M/\eta$, and therefore an approximation space $\mathcal{G}_N^{(T)}$ of smaller dimension. The condition number, the maximum taken over this space, therefore decreases. Conversely, when T is decreased, this means that the Fourier basis is defined over a smaller domain. But the Fourier series $\phi = R_N^{(T)} \circ L_{N,M}^{(T)}(\mathbf{b})$ must fit the nonperiodic data \mathbf{b} in a least-squares sense on the original domain $[-1, 1]$ and must be periodic on the extended domain. To do this, ϕ will be required to take increasingly larger values between data points as $T \rightarrow 1^+$, giving it a bigger uniform norm in comparison to the data norm $\|\mathbf{b}\|_\infty$.

3.1 Setup

In order to make a comparison, for each different value of T used we shall choose the ratio $\eta = M/N$ in such a way that the condition number is the same. Specifically, let $\kappa^* > 1$ be fixed. Then for each T and each M in some specified range, we numerically compute the maximum N such that the condition number is no more than $\kappa^* \log M$. In other words, we compute the function

$$\Theta^{(T)}(M; \kappa^*) = \max \left\{ N : \kappa_{N,M}^{(T)} \leq \kappa^* \log M \right\}, \quad M \in \mathbb{N}. \quad (3.1)$$

Note that we allow $\log M$ factor here since $\kappa_{M/\eta, M}^{(T)}$ grows like $\log M$ as $M \rightarrow \infty$. With this scaling, $\Theta^{(T)}(M; \kappa^*)$ will be linear in M (see later).

The function (3.1) can be computed numerically for each M . This follows from the fact that $\kappa_{N,M}^{(T)}$ can be computed (see §2.4). Note also that $\Theta^{(T)}(M; \kappa^*)$ can be computed for any particular numerical solver used to solve the least squares (2.6), and thus allows a comparison between different methods. We consider this further in §3.3.

Having computed (3.1) for each value of $T = T_1, \dots, T_r$ and some range of M , we next use these values to compare approximation properties of the corresponding equispaced FEs

$$F_{N,M}^{(T_j)}(\cdot), \quad \text{where } N = \Theta^{(T_j)}(M; \kappa^*), \quad j = 1, \dots, r.$$

When doing this, we shall consider a suite of different test functions, described further below. Observe that (3.1) determines the largest value of N for which the condition number is at most $\kappa^* \log M$. Thus, setting $N = \Theta^{(T)}(M; \kappa^*)$ when computing the FE $F_{N,M}^{(T)}(f)$ ensures the best approximation properties for each value of T (since N is maximal) whilst retaining the same condition number for the different choices $T = T_1, \dots, T_r$. Thus a comparison between these different values of T can be made, using the condition number as the common fixed point.

We now require appropriate test functions. Our first three functions are as follows:

$$f_1(x) = e^{230\sqrt{2}i\pi x} \quad f_2(x) = \sin(400x^2) \quad f_3(x) = \text{Ai}(-66 - 70x)$$

These functions all exhibit oscillations, which make them challenging to approximate from equispaced data. Plots of f_2 and f_3 are given in Figure 5. Our next collection of test functions feature singularities in the complex plane near $[-1, 1]$, again making them difficult to approximate:

$$f_4(x) = \frac{1}{1 + 1500x^2} \quad f_5(x) = \frac{1}{60 - 59x} \quad f_6(x) = \frac{1}{1 + 25 \sin^2 8x}.$$

A plot of f_6 is shown in Figure 5. Our final collection of functions, also displayed in Figure 5, is

$$f_7(x) = e^{\sin(21.6\pi x - 10.8\pi) - \cos 8\pi x} \quad f_8(x) = e^{-1/(8x)^2} \quad f_9(x) = s(x).$$

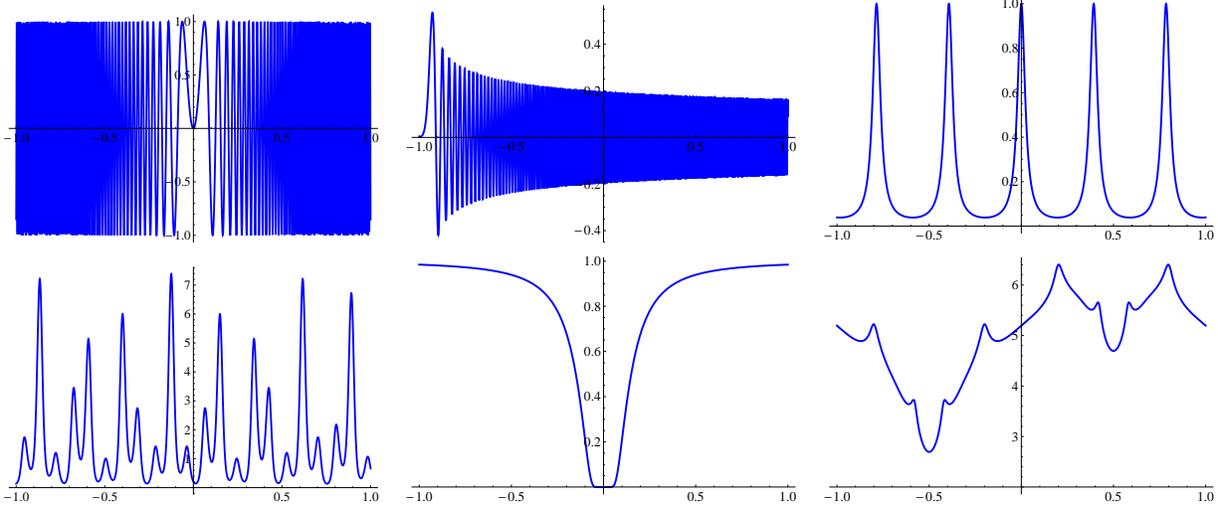


Figure 5: Top row: the functions f_2 , f_3 and f_6 . Bottom row: the functions f_7 , f_8 and f_9 .

T	1.125	1.25	1.5	2.0	2.5	3.0	4.0	5.0	6.0
$\kappa^* = 10$	0.21	0.23	0.28	0.37	0.46	0.55	0.73	0.91	1.00
$\kappa^* = 25$	0.25	0.28	0.33	0.45	0.55	0.66	0.89	1.00	1.00
$\kappa^* = 100$	0.31	0.35	0.42	0.55	0.69	0.82	1.00	1.00	1.00

Table 1: The approximate linear scaling of $\Theta^{(T)}(M; \kappa^*)$ with M . Values were computed using linear regression on the data obtained in Figure 6.

Note that f_7 is often used in testing algorithms for recovering functions to high accuracy from equispaced data, and f_8 is made challenging by its lack of analyticity and the flat region near $x = 0$. The function f_9 is similar to that introduced in [16]. It is obtained by the following iteration:

- $s(x) = \sin \pi x$, $f(x) = s(x)$
- For $j = 1, 2, \dots, 10$, $s(x) = 3/4(1 - 2s(x)^4)$, $f(x) = f(x) + s(x)$.

3.2 Numerical results

In Figure 6 we plot the function $\Theta^{(T)}(M; \kappa^*)$ against M for various values of T and κ^* . Note that this function is approximately linear in M . Moreover, its gradient is larger for bigger values of T , as expected from the results given in Figure 4. The approximate linear rate of growth of $\Theta^{(T)}(M; \kappa^*)$ is shown in Table 1.

Next, in Figure 7 we consider the approximation of the functions f_1, \dots, f_9 using the derived values for $\Theta^{(T)}(M; \kappa^*)$. These results point towards a surprising phenomenon. Besides the choices $T = 5.0$ and $T = 6.0$, all values of T used lead to near-identical approximation errors, regardless of the function considered. Thus, seemingly the value of T , unless taken to be either 5.0 or 6.0 in this case, makes little or no difference to the FE approximation.

Let us explain why the two values $T = 5.0$ and $T = 6.0$ lead to worse approximations. This is due to a phenomenon we refer to as *saturation*:

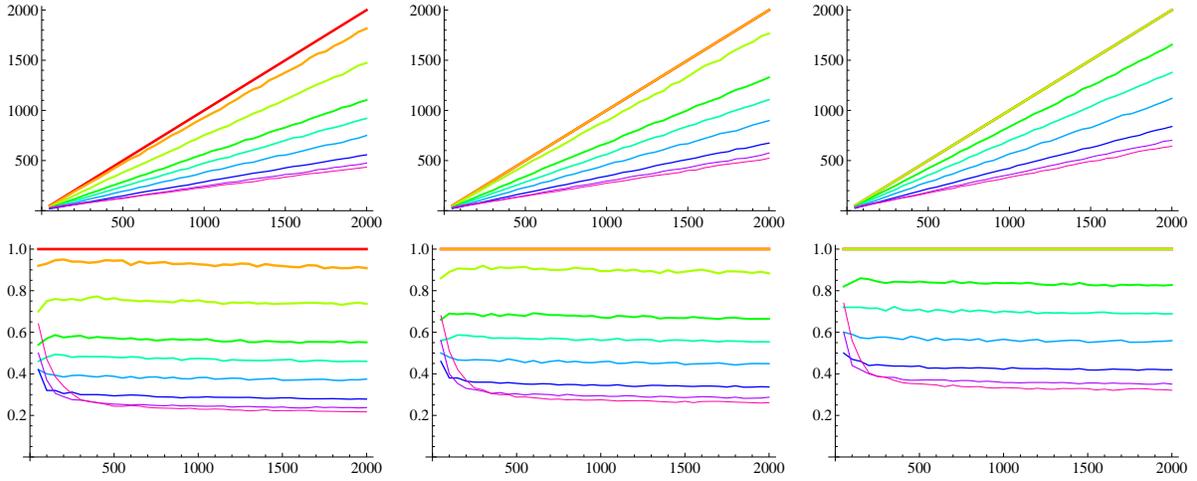


Figure 6: Plots of $\Theta^{(T,\epsilon)}(M; \kappa^*)$ (top) and $\Theta^{(T,\epsilon)}(M; \kappa^*)/M$ (bottom) against M for $\kappa^* = 10$ (left), $\kappa^* = 25$ (middle) and $\kappa^* = 100$ (right) using $\epsilon = 10^{-13}$. The values of T used (in order of increasing thickness) were $T = 1.125, 1.25, 1.5, 2, 2.5, 3.0, 4.0, 5.0, 6.0$. Note that in the middle plots the $T = 5.0$ line is identical to the $T = 6.0$ line, and for the right plots the $T = 4.0, T = 5.0$ and $T = 6.0$ lines are identical.

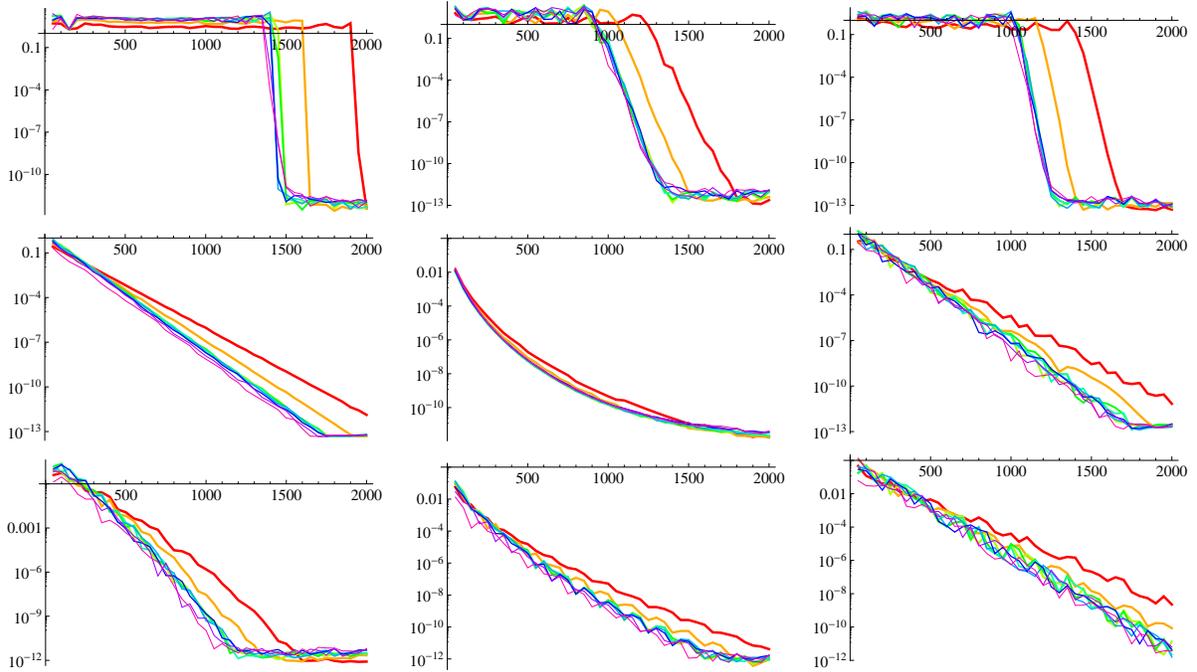


Figure 7: Approximation errors for the functions $f_i, i = 1, \dots, 9$ using the values $\Theta^{(T,\epsilon)}(M; \kappa^*)$ computed in Figure 6 for $\kappa^* = 25$ and $\epsilon = 10^{-13}$.

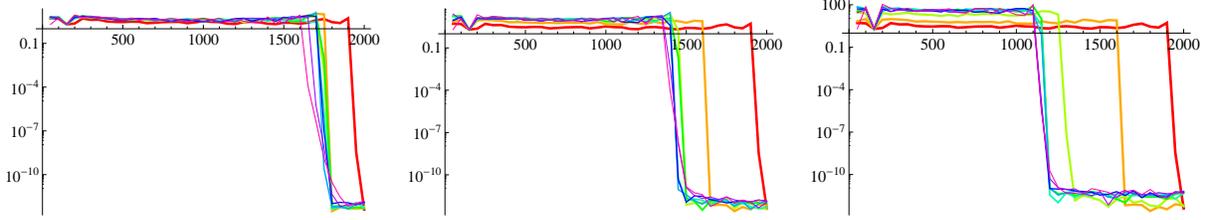


Figure 8: Approximation errors for the function f_1 using the values $\Theta^{(T)}(M; \kappa^*)$ computed in Figure 6 for $\kappa^* = 10$ (left), $\kappa^* = 25$ (middle) and $\kappa^* = 100$ (right).

Definition 3.1. For a given $T > 1$ and $\kappa^* > 1$, saturation occurs if

$$\limsup_{M \rightarrow \infty} \kappa_{M,M}^{(T)} / \log M < \kappa^*.$$

Saturation means that the corresponding FE is too stable to take advantage of the allowed condition number κ^* . In other words, the maximal value of N permitted is limited by the fact that $N \leq M$ in the FE approximation, and not by the condition number constraint $\kappa_{N,M}^{(T)} \leq \kappa^*$. When saturation occurs, the resulting FE approximation performs worse in terms of approximation than that corresponding to a value of T for which saturation does not occur. Figure 6 illustrates that the two values of T which give worse approximations in Figure 7, i.e. $T = 5.0$ and $T = 6.0$, do indeed saturate. This is further demonstrated in Figure 8, where we consider the approximation of the function f_1 using different values of κ^* . Figure 6 shows that for $\kappa^* = 10$ only $T = 6.0$ saturates, whereas $T = 5.0$ and $T = 6.0$ both saturate for $\kappa^* = 25$, and for $\kappa^* = 100$ the values $T = 4.0$, $T = 5.0$ and $T = 6.0$ all saturate. In Figure 8 the FE approximations with these values of T perform worse than the FEs corresponding to values of T which do not saturate. Note that when $\kappa^* = 10$ the effect of the saturation for $T = 6.0$ has less impact, since $\kappa_{M,M}^{(6,\epsilon)} / \log M \approx 5$ is reasonably close to κ^* in this case. Hence saturation occurs, but to a lesser extent. Similarly, since $\kappa_{M,M}^{(6)} < \kappa_{M,M}^{(5)} < \kappa_{M,M}^{(4)}$ the effect of the saturation on the FE approximation when $\kappa^* = 100$ is less for $T = 4.0$ than it is for $T = 5.0$ and $T = 6.0$.

With this in hand, we are now able to state the main empirical conclusion of this section: *unless saturation occurs, the choice of T makes little difference to the FE approximation.* In particular, one may use the value $T = 2$ provided it does not saturate for given value of κ^* . We note in passing that this phenomenon is of course asymptotic in M , and relies on the fact that the functions under consideration require large numbers of equispaced points to be approximated to any accuracy. For smooth function lacking unpleasant features such as close singularities or oscillations, i.e. functions that can be resolved with small M , there will be slight discrepancies for different values of T .

Let us give some explanation for why this conclusion should hold. Figure 6 shows that the function $\Theta^{(T)}(M; \kappa^*)$ is approximately linear in M , i.e. $\Theta^{(T)}(M; \kappa^*) \approx \nu^{(T)}(\kappa^*)M$ for large M for some constant $\nu^{(T)}(\kappa^*) > 0$. Explicit values of the quantity $\nu^{(T)}(\kappa^*)$ are given in Table 1, and these are plotted against T in Figure 9. As is evident from this figure, for each fixed κ^* , the quantity $\nu^{(T)}(\kappa^*)$ is approximately linear in T up to the point at which saturation occurs. Write $\nu^{(T)}(\kappa^*) \approx \tau(\kappa^*)T$ for some $\tau(\kappa^*) > 0$ and now consider the FE approximation $F_{N,M}^{(T)}$ where $N = \tau(\kappa^*)TM$. By (2.11) and (2.12), we see that the error is determined up to a mildly growing factor in M by the quantity $E_N(f) = E_{\tau(\kappa^*)TM}(f)$. However, the error bound (2.13) depends on the ratio N/T , which in this case is equal to $\tau(\kappa^*)M$ regardless of the choice of T . Hence the bound (2.13) is independent of T in the nonsaturating case for this choice of N , and therefore we expect the same T -independence of the FE approximation whenever N is taken according to $\Theta^{(T)}(M; \kappa^*)$.

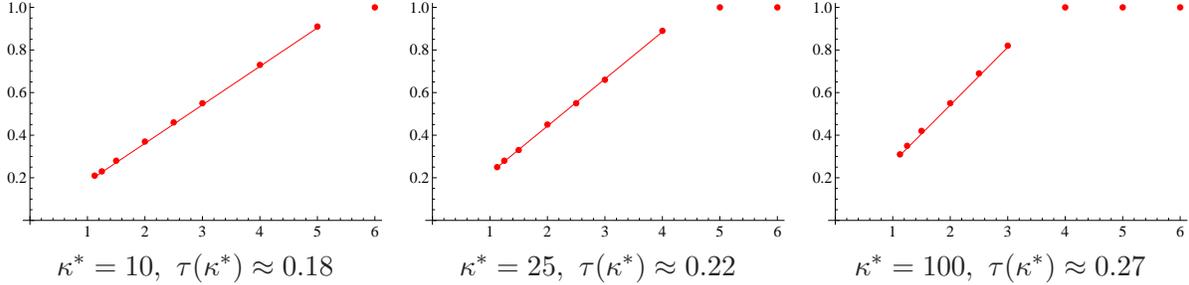


Figure 9: The approximate values $\nu^{(T)}(\kappa^*)$ from the Table 1 plotted against T . The solid curve is the line with slope $\tau(\kappa^*)$, where $\tau(\kappa^*)$ is computed using linear regression on the data $\{T, \nu^{(T)}(\kappa^*)\}$ corresponding to those nonsaturating values of T .

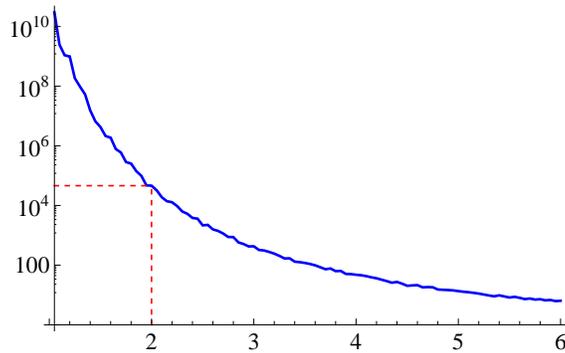


Figure 10: The quantity $\kappa_{M,M}^{(T,\epsilon)}/\log M$ against $1 < T \leq 6$ for $M = 500$ and $\epsilon = 10^{-13}$.

This explanation aside, it is also of interest to numerically determine the so-called *saturation point*: that is to say, the maximal value of T (for a given κ^*) above which saturation occurs. This is shown in Figure 10, where we plot the quantity $\kappa_{M,M}^{(T)}$ against T for $M = 500$. Note that this quantity is decreasing in T (recall Figure 4). Using this figure, we make the following observation. For the case $T = 2$, saturation occurs when the $\kappa^* \approx 5e4$ or greater. Hence, we now conclude the following: *if one implements the equispaced FE with $T = 2$, $\epsilon = 10^{-13}$ and $N = \Theta^{(2,\epsilon)}(M; \kappa^*)$ chosen such that the κ^* is less than $\approx 5e4$, then no other value of T will asymptotically give a better approximation.* This establishes one of the main aims of this paper: namely, determining when the value $T = 2$, and hence the associated fast algorithm, can be used without concern that another value gives better accuracy. Note that this conclusion is very reasonable. In practice, we usually want the condition number to be much smaller than 10^4 in magnitude.

Having now ascertained that one may use $T = 2$ without worry in most cases, we end this subsection by providing numerical values for the constants $\kappa_{M/\eta,M}^{(2,\epsilon)}$ and $\lambda_{M/\eta,M}^{(2,\epsilon)}$. These are shown in Figure 11. As we see, the choices $\eta = 1.5$ or $\eta = 2.0$ seem reasonable in practice. Increasing η beyond this point brings only marginal benefits in stability.

3.3 Other solvers

The phenomenon described above also occurs when a different numerical solver is used. We illustrate this in Figure 12 for *Matlab*'s `\` and *Mathematica*'s `LeastSquares`. Interestingly, the effect is much more pronounced for the latter than for the former, although we believe it would become more

η/M	250	500	750	1000
1.00	2.51e5	2.86e5	2.65e5	3.14e5
1.25	1.01e4	1.25e4	1.72e4	1.99e4
1.50	2.16e3	2.39e3	2.41e3	2.84e3
2.00	1.88e2	2.25e2	2.89e2	3.27e2
3.00	2.70e1	3.29e1	3.94e1	3.94e1
4.00	1.18e1	1.53e1	1.67e1	1.84e1

η/M	250	500	750	1000
1.00	6.16e-7	1.48e-6	2.76e-6	3.49e-6
1.25	3.86e-8	8.82e-8	9.98e-8	1.33e-7
1.50	3.04e-9	7.36e-9	1.56e-8	1.82e-8
2.00	4.26e-10	1.02e-9	1.00e-9	1.23e-9
3.00	2.97e-11	8.20e-11	8.37e-11	1.97e-10
4.00	1.71e-11	1.51e-11	3.59e-11	5.24e-11

Figure 11: Values of $\kappa_{M/\eta, M}^{(T, \epsilon)}$ (left) and $\lambda_{M/\eta, M}^{(T, \epsilon)}$ (right) for $T = 2$ and $\epsilon = 10^{-13}$.

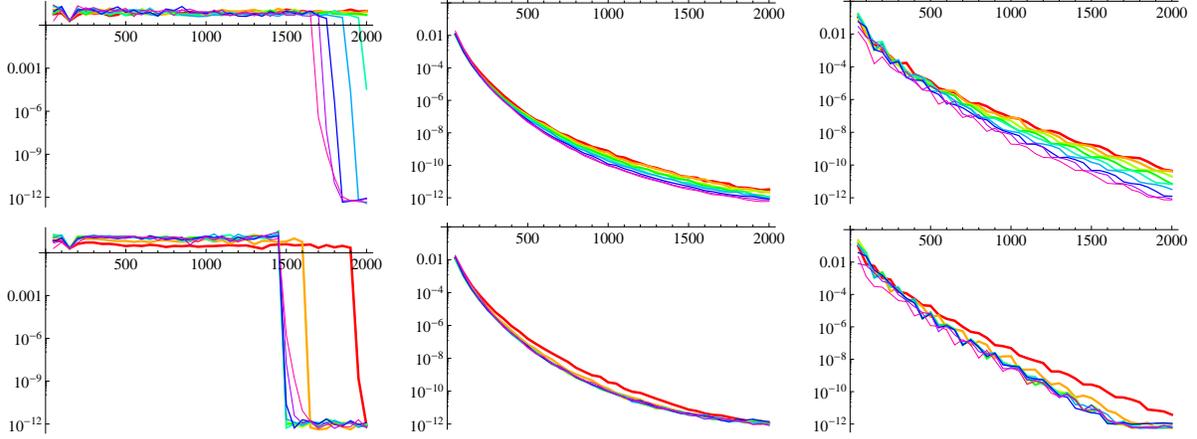


Figure 12: Approximation errors for the functions $f_1(x)$, $f_5(x)$ and $f_8(x)$ (left to right) using *Matlab's * (top row) and *Mathematica's LeastSquares* (bottom row). For each solver, the values $\Theta^{(T)}(M; \kappa^*)$ were computed for $\kappa^* = 25$.

apparent in the former for larger M (recall that this phenomenon is asymptotic in M). The discrepancies in the results are likely due to the different algorithms having somewhat different default tolerances for solving ill-conditioned least-squares problems.

3.4 Influence of the SVD tolerance ϵ

Let us now return to the SVD algorithm. Thus far, we have taken the tolerance ϵ to be equal to 10^{-13} . However, other values are possible, and it is of interest to determine how this affects the approximation. Note that some previous insight in this problem was given in [5], where the effect of ϵ was considered for specific functions.

In Figure 13 we plot $\Theta^{(T, \epsilon)}(M; \kappa^*)$ for different values of ϵ using $T = 2.0$. As is evident, $\Theta^{(T, \epsilon)}(M; \kappa^*)$ grows more rapidly for increasing ϵ . This should come as no surprise, since it is the small singular vectors that cause ill-conditioning. However, the main purpose of this figure is to show the specific improvement that is possible by changing ϵ . For example, when $\epsilon = 10^{-13}$ the function $\Theta^{(T, \epsilon)}(M; \kappa^*)$ is approximately $0.4M$. Conversely, for $\epsilon = 10^{-6}$ it is approximately $0.8M$, i.e. it scales roughly twice as quickly. If δ is some finite tolerance greater than 10^{-6} , this means that the FE approximation with $\epsilon = 10^{-6}$ will approximate a given function to accuracy δ using roughly half the number of equispaced points as is required when $\epsilon = 10^{-13}$. This fact is also confirmed in the right panel of Figure 13, where the oscillatory function $f(x) = \exp(250\sqrt{2}\pi ix)$ is approximated using different values of ϵ . Resolving this function using $\epsilon = 10^{-13}$ requires roughly

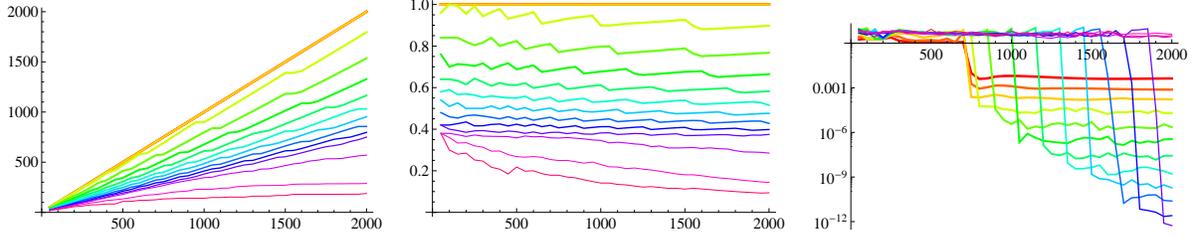


Figure 13: Left and Middle: the functions $\Theta^{(T,\epsilon)}(M; \kappa^*)$ and $\Theta^{(T,\epsilon)}(M; \kappa^*)/M$ against M for $T = 2$, $\kappa^* = 10$ and $\epsilon = 10^{-2}, 10^{-3}, \dots, 10^{-16}$ (thickest to thinnest). Note that the graphs for $\epsilon = 10^{-2}, 10^{-3}, 10^{-4}$ are the same. Right: approximation of $f(x) = \exp(250\sqrt{2}\pi ix)$ using these values.

3700 equispaced data points, whereas when $\epsilon = 10^{-6}$ this value drops to around 1700.

Of course, the downside of a larger ϵ is that the minimal error is limited to approximately ϵ , as can be seen in Figure 13. Nonetheless, the conclusion we draw from this section is that if accuracy close to machine precision is not required – as is typically the case in practice, where three to six digits is often acceptable – then a viable way to increase the performance of the FE approximation whilst maintaining the condition number is to use a larger value of ϵ .

4 Resolution power

In many applications, it is important to have an approximation algorithm with good *resolution power*. Loosely speaking, this means that oscillatory functions are recovered using using a number of measurements that scales linearly with the frequency of oscillation with a constant that is as small as possible. Formally, let $\{F_M\}_{M \in \mathbb{N}}$ be a sequence of approximations such that $F_M(f)$ depends only on the values of f on an equispaced grid of $2M + 1$ points. Let

$$\mathcal{R}(\omega, \delta) = \min \{M \in \mathbb{N} : \|e^{i\pi\omega \cdot} - F_M(e^{i\pi\omega \cdot})\|_\infty < \delta\}, \quad \omega > 0, \quad 0 < \delta < 1,$$

then we say that F_M has resolution constant $0 < r < \infty$ if

$$\mathcal{R}(\omega, \delta) \sim r\omega, \quad \omega \rightarrow \infty, \tag{4.1}$$

for any fixed δ . The approximation F_M has good resolution power if r is small, and bad resolution power if r is large.

For periodic oscillations, the Fourier series approximation on $[-1, 1]$ of degree M has optimal resolution constant $r = 1$. Since the number equispaced points is $2M + 1$, this corresponds to two *points-per-wavelength*. Of course, Fourier series do not converge uniformly for nonperiodic oscillations, which is why we resort to alternative algorithms such as FEs. Naturally, though, it is desirable that the FE resolution constant be as close to the optimal value $r = 1$ as possible.

In Figure 14 we numerically determine the resolution constant for the FE approximation $F_{M/\eta, M}^{(T,\epsilon)}$ by computing the function

$$\mathcal{R}(\omega, \delta) = \mathcal{R}(\omega, \delta; T, \eta).$$

These results suggest that the resolution constant

$$r \approx T\eta, \tag{4.2}$$

is approximately the product of the extension parameter T and the oversampling ratio η .

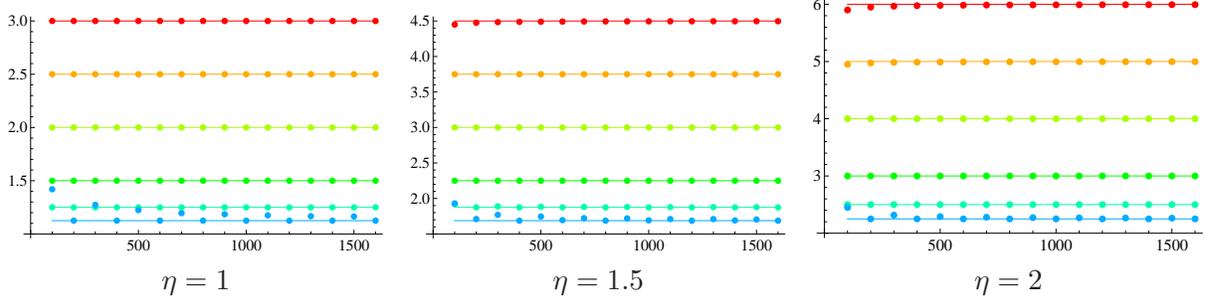


Figure 14: The quantity $\mathcal{R}(\omega, \delta; T, \eta)/\omega$ against ω for $T = 3.0, 2.5, 2.0, 1.5, 1.25, 1.125$ and $\delta = 10^{-3}$. The solid lines indicate the values $T\eta$.

We shall discuss the consequences of this observation in a moment, but we first wish to explain why r should be at most $T\eta$. Recall from (2.11) that the error of the FE approximation is determined (up to a small linear constant in M) by the decay of the factor $E_N(f)$, given by (2.12), where $N = M/\eta$. Now let

$$\mathcal{R}_0(\omega, \delta) = \min \{N \in \mathbb{N} : E_N(e^{i\pi\omega \cdot}) < \delta\}, \quad \omega > 0, \sqrt{2T}\mu < \delta < 1. \quad (4.3)$$

Then to show that $r = T\eta$, it suffices to prove that

$$\mathcal{R}_0(\omega, \delta) \sim T\omega, \quad \omega \rightarrow \infty.$$

We now note the following:

Lemma 4.1. *Let $\mathcal{R}_0(\omega, \delta)$ be as in (4.3). Then*

$$\limsup_{\omega \rightarrow \infty} \mathcal{R}_0(\omega, \delta)/\omega \leq T.$$

Proof. Let $f(x) = e^{i\pi\omega x}$ and write $\omega = P/T + z$ where $0 \leq z < 1/T$ and $P \in \mathbb{N}$. Then

$$f(x) = e^{i\frac{P\pi}{T}x} g(x), \quad g(x) = e^{i\frac{z\pi}{T}x}.$$

Let $K \in \mathbb{N}$ and let $\phi \in \mathcal{G}_{P+K}^{(T)}$ be given by $\phi(x) = e^{i\frac{P\pi}{T}x} \tilde{\phi}(x)$, where $\tilde{\phi} \in \mathcal{G}_K^{(T)}$ is arbitrary. Define $N = P + K$, and let \mathbf{a} be the vector of coefficients of ϕ . Then by definition,

$$E_N(f) \leq \|f - \phi\|_\infty + \mu|\mathbf{a}|_\infty = \|g - \tilde{\phi}\|_\infty + \mu|\tilde{\mathbf{a}}|_\infty,$$

where $\tilde{\mathbf{a}}$ is the vector of coefficients of $\tilde{\phi}$. Since $\tilde{\phi}$ is arbitrary, we deduce that $E_N(f) \leq E_K(g)$. Thus $E_N(f) \leq \delta$ provided $E_K(g) \leq \delta$. Using Corollary 2.7, we see that $E_K(g) \leq \delta$ provided $K \geq K_0(\delta)$, where $K_0(\delta)$ is independent of ω . Thus, $E_N(f) \leq \delta$ whenever $N = P + K_0(\delta) = (\omega - z)T + K_0(\delta)$. Since $K_0(\delta)$ is independent of ω , we obtain the result. \square

From this, we deduce that $\mathcal{R}_0(\omega, \delta) \leq T\omega + \mathcal{O}(1)$ for large ω and hence $r \leq T\omega$. Unfortunately, we have no proof of the lower bound, although it is supported by the results in Figure 14.

Let us now discuss the consequences of (4.2). Suppose we return to the earlier experiment where κ^* is fixed and N is chosen according to $N = \Theta^{(T)}(M; \kappa^*)$. For large M , Figure 6 shows that

$$\Theta^{(T)}(M; \kappa^*) \approx M/\eta^{(T)}, \quad \eta^{(T)} = \eta^{(T)}(\kappa^*), \quad (4.4)$$

κ^*	10	25	100	500	1000	5000	10000
r	5.41	4.44	3.64	2.98	2.77	2.38	2.25

Table 2: The resolution constant $r = \eta^{(T)}T$ for $T = 2$, where $\eta^{(T)} = \eta^{(T)}(\kappa^*)$ is given by (4.4).

for some fixed $\eta^{(T)}$ depending on T and κ^* . In our previous observation, it was found that the resulting FE approximation was independent of T . In particular, this holds for the oscillatory exponential $\exp(i\pi\omega x)$. But since the resolution constant r , which describes the point after which the approximation error begins to decay, is equal to ηT , we therefore deduce the following:

$$\eta^{(T)}T = \eta^{(T')}T',$$

for any T and T' that do not saturate. This implies that the level curves of the condition number $\kappa = \kappa_{M/\eta, M}^{(T)}$ are approximately given by $T\eta = \text{constant}$, for sufficiently large M and non-saturating T , which is in good agreement with the contour plot given in Figure 4.

In addition to this, another important implication of (4.2) is that to get better resolution power one must necessarily worsen the stability of the algorithm. In other words, there is a direct relationship between κ^* and r , regardless of the choice of T made. In Table 2 we give numerical results for the resolution constant when $T = 2$ (and therefore all nonsaturating T) for different values of κ^* . As we see, by allowing κ^* to increase to roughly 500 we get a marked improvement over when $\kappa^* = 10$. Beyond this point, further increases give only marginal gains at the expense of much larger condition numbers.

5 Other data

In this final section, we illustrate that the main phenomena observed for equispaced data are also witnessed for numerous other types of data. To this end, we consider the following four examples:

- *Jittered pointwise data.* Here the measurements of f are pointwise samples at the jittered locations

$$f(x_m), \quad x_m = \frac{m}{M} + z_m,$$

where $z_m \in (-\delta/M, \delta/M)$ and $0 < \delta < 1$. This is a typical example of a nonuniform sampling pattern for scattered data approximation. We shall choose the z_m 's as follows:

$$z_m = \frac{\delta}{M} \sin(M^2/m), \quad m \neq 0, \quad z_0 = 0. \quad (5.1)$$

- *Logarithmic pointwise data.* Here we again sample f pointwise at nodes x_m , but in this case the nodes are logarithmically distributed:

$$x_m = -x_{-m} = 10^{\left(\frac{m-1}{M-1} - 1\right) \log_{10}(cM)}, \quad x_0 = 0, \quad (5.2)$$

where $c > 0$ is a fixed, user-controlled parameter (we take $c = 2$ in our results). This sampling pattern corresponds to a nonuniform sampling scenario where data is collected more densely at the origin.

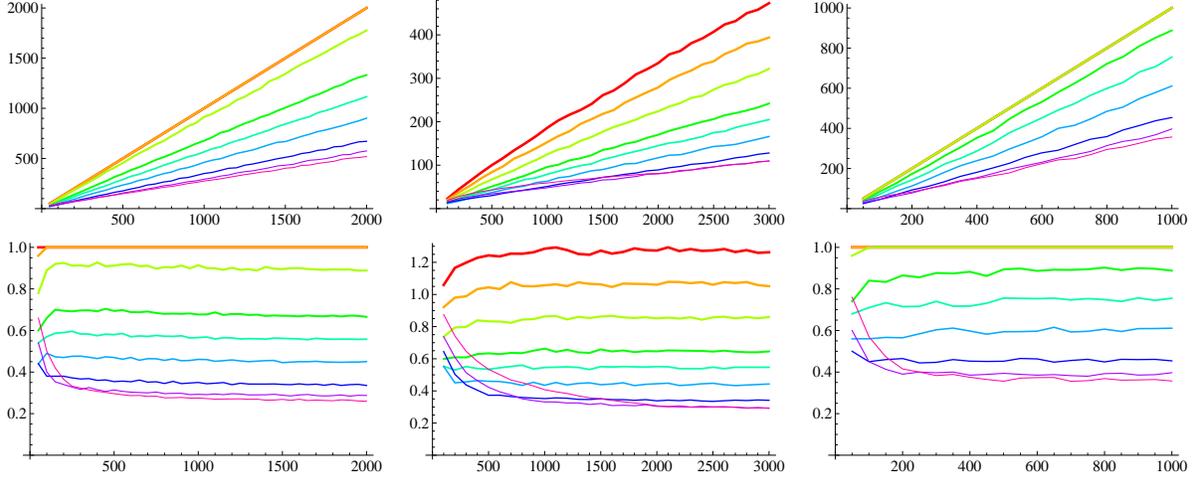


Figure 15: Plots of $\Theta^{(T)}(M; \kappa^*)$ (top row) and $\Theta^{(T)}(M; \kappa^*)/S(M)$ (bottom row) against M for jittered (5.1) (left), logarithmic (5.2) (middle) and Fourier (5.3) (right) data. Here $S(M) = M/\log(M)$ for the logarithmic data and $S(M) = M$ otherwise, and $\kappa^* = 25$ (jittered, logarithmic) or $\kappa^* = 10$ (Fourier).

- *Fourier data.* In some applications, rather than pointwise samples, we may wish to reconstruct f from its Fourier coefficients:

$$\hat{f}(m) = \int_{-1}^1 f(x) e^{-i\pi m x} dx, \quad |m| \leq M. \quad (5.3)$$

As discussed in [1], this can be seen as a continuous analogue of the equispaced data recovery problem. In particular, there is a completely analogous result to that of Platte, Trefethen & Kuijlaars regarding stability and convergence [2]. Note that for this data we replace the uniform norms used in the error estimates and the definitions of $\kappa_{N,M}^{(T)}$ and $\lambda_{N,M}^{(T)}$ by the L^2 and ℓ^2 norms. This is natural in view of Parseval's identity for Fourier coefficients.

- *Optimal pointwise data.* Finally, in order to show that the phenomenon is not witnessed for all data, we consider pointwise samples taken at the so-called mapped symmetric Chebyshev nodes (see [3]):

$$x_m = -x_{-m-1} = m^{-1} \left(\cos \frac{(2m+1)\pi}{2M+2} \right), \quad m = 0, \dots, M, \quad (5.4)$$

where m is given by (2.1). These nodes are derived from the observation that FE approximations correspond to algebraic polynomial approximations in the mapped co-ordinate $z = m(x)$ [4]. Chebyshev nodes provide optimal nodes for polynomial interpolation. Therefore, under the inverse mapping m^{-1} , they provide the optimal nodes (5.4) for FE approximations. Note that since the nodes arise in this way, no oversampling is required in the FE approximation, i.e. we let $\eta = 1$ in this case.

In Figure 15 we give plots of the function $\Theta^{(T)}(M; \kappa^*)$ for the first three data types. For jittered and Fourier data the scaling is linear, whereas for the logarithmic data $\Theta^{(T)}$ scales like $M/\log M$. This scaling is proportional to the reciprocal of the maximal spacing between nodes in the case, and hence is completely expected. Note also that no values of T saturate for the logarithmic data,

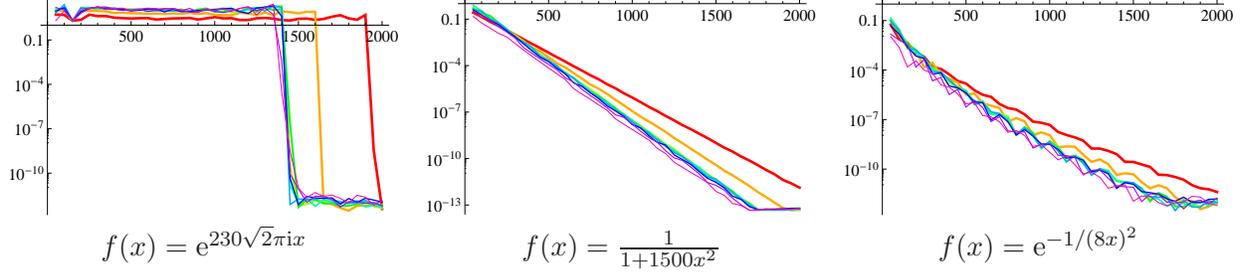


Figure 16: Approximation errors for jittered data (5.1) using the values $\Theta^{(T)}(M; \kappa^*)$ from Figure 15.

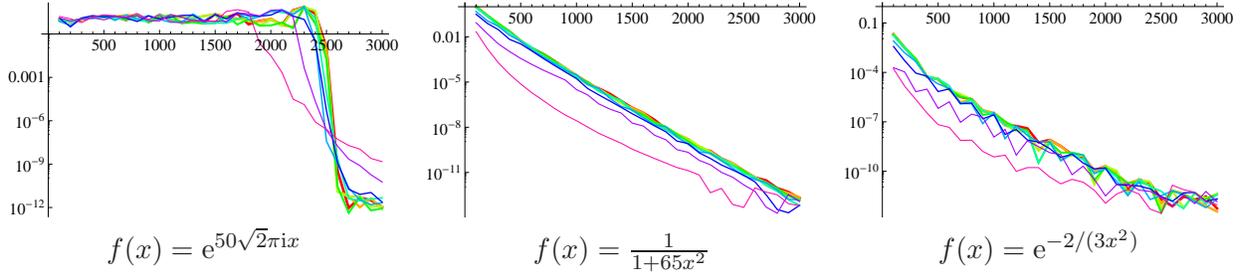


Figure 17: Approximation errors for logarithmic data (5.2) using the values $\Theta^{(T)}(M; \kappa^*)$ from Figure 15.

whereas values $T = 5.0$ and $T = 6.0$ saturate for the jittered data, and for the Fourier data the values $T = 3.0$, $T = 4.0$, $T = 5.0$ and $T = 6.0$ all saturate. The lower saturation point for the latter is due to the fact that the condition number is measured in the weaker L^2 norm.

Next, in Figures 16–18 we compare approximation errors using these values. For the jittered and Fourier data we see exactly the same phenomenon as before: namely, the approximation errors are roughly independent of T . A similar phenomenon is witnessed for the logarithmic data, although it is slightly weaker: $T = 1.125$ gives somewhat better errors than the other values. This is due to the much more severe scaling of $\Theta^{(T)}(M; \kappa^*)$ with M in this case, which means the asymptotic regime takes longer to set in.

Finally, in Figure 19 we display approximation errors for the optimal pointwise data (5.4). As is evident, the phenomenon does not occur in this case. The reason for this is due to the choice of the data, which means that no oversampling is required. Thus the FE approximation error is proportional to $E_N(f)$, where $N = M$, multiplied by a mildly growing factor. Figure 19 therefore

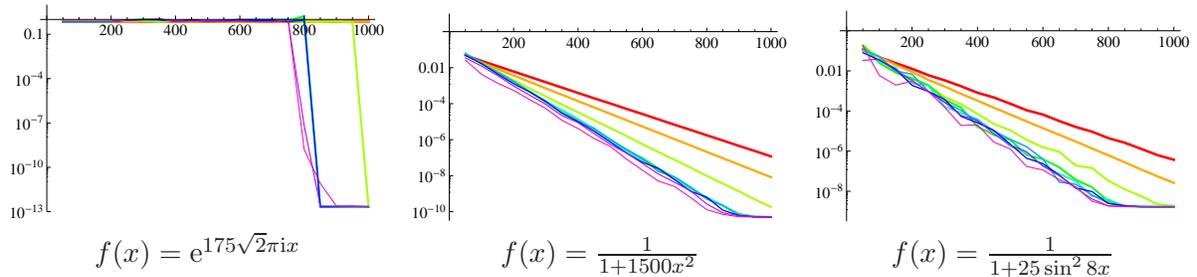


Figure 18: Approximation errors for Fourier data (5.3) using the values $\Theta^{(T)}(M; \kappa^*)$ from Figure 15.

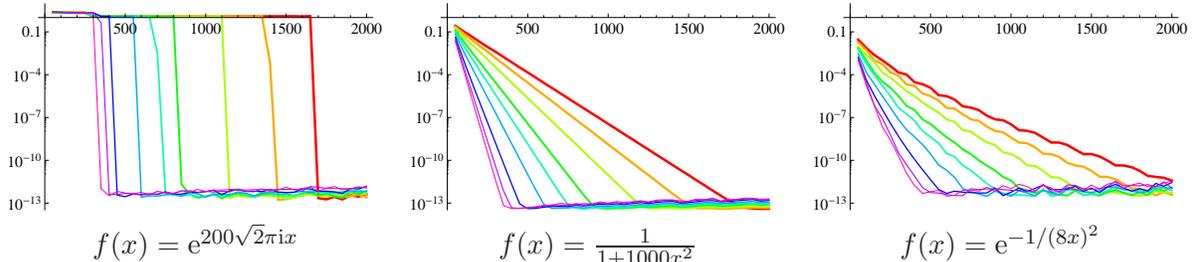


Figure 19: Approximation errors for the optimal pointwise data (5.4) using $M = N$.

serves as a reminder that smaller values of T are intrinsically better than larger values, provided one has freedom to pick ideal data. On the other hand, for nonideal data – such as equispaced, jittered, logarithmic or Fourier data – this effect is nullified by a worse scaling of $\Theta^{(T)}(M; \kappa^*)$.

6 Conclusions and open problems

The purpose of this paper was to document an interesting phenomenon in FE approximations from equispaced data. Namely, when the desired condition number is fixed, the choice of the extension parameter T has no substantial effect on the approximation. This is on the proviso that saturation does not occur, which we have shown to be the case for moderate values of T and κ^* . In particular, one may use $T = 2$, and the associated fast algorithm, without concern that it is suboptimal.

The main open problem is to provide mathematical analysis for the empirical conclusions drawn. We believe this is possible, although not straightforward. One possible approach towards this is to conduct an asymptotic analysis of the singular values and vectors of the matrix $A^{(T)}$. Recall that the normal form $(A^{(T)})^* A^{(T)}$ is a Riemann sum approximation to the Gram matrix G of the FE basis functions $\phi_n(x) = e^{i\frac{n\pi}{T}x}$. As discussed in [4], the matrix G is precisely the prolate matrix. The eigenvalues and eigenvectors of this matrix were analyzed in detail by Slepian [14] (see also [17]). It may be possible to do the same for the discretized version $(A^{(T)})^* A^{(T)}$, and this is an important topic for future work.

Another question raised by this work is that of whether it might be possible to vary T with M to achieve better results; in particular, improved resolution power. We believe this may be the case, the caveat being that there is currently no fast algorithm for $T \neq 2$. Some potential choices for varying T with M were considered previously in [3, 4]. But it may also be possible using the approach of this paper to numerically compute an optimal (in some sense) value of T for each M . We leave this for future work.

Acknowledgments

The authors would like to thank Daan Huybrechs, Mark Lyon and Rodrigo Platte for useful discussions. BA acknowledges support from the NSF DMS grant 1318894.

References

- [1] B. Adcock, A. C. Hansen, and C. Poon. Beyond consistent reconstructions: optimality and sharp bounds for generalized sampling, and application to the uniform resampling problem. *SIAM J. Math. Anal.*, 45(5):3114–3131, 2013.

- [2] B. Adcock, A. C. Hansen, and A. Shadrin. A stability barrier for reconstructions from Fourier samples. *SIAM J. Numer. Anal.*, 52(1):125–139, 2014.
- [3] B. Adcock and D. Huybrechs. On the resolution power of Fourier extensions for oscillatory functions. *J. Comput. Appl. Math.*, 260:312–336, 2014.
- [4] B. Adcock, D. Huybrechs, and J. Martín-Vaquero. On the numerical stability of Fourier extensions. *Found. Comput. Math. (to appear)*, 2013.
- [5] J. P. Boyd. A comparison of numerical algorithms for Fourier Extension of the first, second, and third kinds. *J. Comput. Phys.*, 178:118–160, 2002.
- [6] J. P. Boyd and J. R. Ong. Exponentially-convergent strategies for defeating the Runge phenomenon for the approximation of non-periodic functions. I. Single-interval schemes. *Commun. Comput. Phys.*, 5(2–4):484–497, 2009.
- [7] 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.
- [8] D. Gottlieb and S. A. Orszag. *Numerical Analysis of Spectral Methods: Theory and Applications*. Society for Industrial and Applied Mathematics, 1st edition, 1977.
- [9] D. Huybrechs. On the Fourier extension of non-periodic functions. *SIAM J. Numer. Anal.*, 47(6):4326–4355, 2010.
- [10] M. Lyon. Approximation error in regularized SVD-based Fourier continuations. *Appl. Numer. Math.*, 62:1790–1803, 2012.
- [11] M. Lyon. A fast algorithm for Fourier continuation. *SIAM J. Sci. Comput.*, 33(6):3241–3260, 2012.
- [12] M. Lyon and J. Picard. A fast algorithm for unevenly-spaced non-periodic Fourier approximation. *Preprint*, 2013.
- [13] R. Platte, L. N. Trefethen, and A. Kuijlaars. Impossibility of fast stable approximation of analytic functions from equispaced samples. *SIAM Rev.*, 53(2):308–318, 2011.
- [14] D. Slepian. Prolate spheroidal wave functions. Fourier analysis, and uncertainty V: The discrete case. *Bell System Tech J.*, 57:1371–1430, 1978.
- [15] L. N. Trefethen and D. Bau III. *Numerical Linear Algebra*. SIAM, Philadelphia, 1997.
- [16] L. N. Trefethen. Computing numerically with functions instead of numbers. *Math. Comput. Sci.*, 1:9–19, 2007.
- [17] J. Varah. The prolate matrix. *Linear Algebra Appl.*, 187(1):269–278, 1993.

Symbols

Symbol	Description
T	Extension parameter
N	Number of Fourier modes
M	Number of equispaced nodes
η	The ratio M/N of nodes to modes
$\phi_n(x)$	The exponential $\frac{1}{\sqrt{2T}}e^{i\frac{n\pi}{T}x}$
$\mathcal{G}_N^{(T)}$	The space $\text{span}\{\phi_n : n \leq N\}$
S_M	The sampling operator $f \mapsto \frac{1}{\sqrt{M}}(f(m/M))_{m=-M}^M$
$F_{N,M}^{(T)}(f)$	The FE approximation (2.4)
$\mathbf{a} = (a_n)_{n=-N}^N$	Coefficients of the FE approximation
$A^{(T)}$	Matrix of the FE approximation
$L_{N,M}^{(T)}$	Solution of the least-squares (2.6)
$R_N^{(T)}$	The mapping from coefficients to Fourier series
ϵ	SVD truncation parameter
$\kappa_{N,M}^{(T)}$	Condition number
$\lambda_{N,M}^{(T)}$	Numerical defect constant
$E_N(f)$	Approximation error (2.12)
κ^*	Maximum allowed condition number
$\Theta^{(T)}(M; \kappa^*)$	The modes to nodes ratio (3.1)
$\nu^{(T)}(\kappa^*)$	The approximate linear scaling of $\Theta^{(T)}(M; \kappa^*)$ with M
$\tau(\kappa^*)$	The approximate linear scaling of $\nu^{(T)}(\kappa^*)$ with T
r	The resolution constant