

A FRAME THEORETIC APPROACH TO THE NONUNIFORM FAST FOURIER TRANSFORM*

ANNE GELB[†] AND GUOHUI SONG[‡]

Abstract. Nonuniform Fourier data are routinely collected in applications such as magnetic resonance imaging, synthetic aperture radar, and synthetic imaging in radio astronomy. To acquire a fast reconstruction that does not require an online inverse process, the nonuniform fast Fourier transform (NFFT), also called convolutional gridding, is frequently employed. While various investigations have led to improvements in accuracy, efficiency, and robustness of the NFFT, not much attention has been paid to the fundamental analysis of the scheme, and in particular its convergence properties. This paper analyzes the convergence of the NFFT by casting it as a Fourier frame approximation. In so doing, we are able to design parameters for the method that satisfy conditions for numerical convergence. Our so-called frame theoretic convolutional gridding algorithm can also be applied to detect features (such as edges) from nonuniform Fourier samples of piecewise smooth functions.

Key words. Fourier frames, numerical frame approximation, nonuniform fast Fourier transform, convolutional gridding

AMS subject classifications. 65T50, 65T40, 42C15

DOI. 10.1137/13092160X

1. Introduction. In several imaging applications, such as magnetic resonance imaging (MRI) [3, 29, 32], synthetic aperture radar (SAR), or synthetic imaging in radio astronomy [5], data are finitely sampled in the Fourier domain. Although in many cases the data are presumed to be sampled uniformly, there are some situations for which the data are collected nonuniformly, either by instrument design or by the circumstances surrounding the collection. Regardless of the particular collection protocol, reconstructing the underlying image can be viewed as an approximation to the inverse Fourier transform. As the data often must be accurately analyzed in real time, numerical efficiency and robustness of the numerical algorithm are essential.

One common technique used to reconstruct images from nonuniform finitely sampled Fourier data involves “regridding” the data uniformly, via a convolution, so that the inverse fast Fourier transform (FFT) algorithm can be applied. This technique goes by various names, depending on the particular application. We will use the terms *nonuniform FFT* (NFFT) [14, 15] and *convolutional gridding* [24, 28, 29, 32] interchangeably. Briefly, the method uses practical and heuristic arguments to construct an essentially compactly supported window function, which is also essentially compactly supported in the Fourier domain. Uniform Fourier coefficients are then computed by numerically approximating the convolution of the Fourier transform of the underlying function with the window function from the given nonuniformly sampled Fourier data. Numerical weights, often called *density compensation factors* (DCFs), are introduced to evaluate the convolution integral via summation. The FFT is then used directly to reconstruct the underlying image.

*Received by the editors May 20, 2013; accepted for publication (in revised form) February 26, 2014; published electronically May 22, 2014. The authors were supported in part by grants NSF-DMS 1216559 and AFOSR 12004863.

<http://www.siam.org/journals/sinum/52-3/92160.html>

[†]School of Mathematical and Statistical Sciences, Arizona State University, Tempe, AZ 85287-1804 (anne.gelb@asu.edu).

[‡]Department of Mathematics, Clarkson University, Potsdam, NY 13699 (gsong@clarkson.edu).

Although the convolutional gridding algorithm typically achieves satisfactory results, it is not difficult to generate examples where the method fails [37]. Specifically, if the DCFs are viewed as quadrature weights for the inverse Fourier transform (or convolution integral), then the rate of convergence is tied to the accuracy of the numerical quadrature scheme. Moreover, for sampling patterns that are increasingly sparse in the high frequency domain, such as those found in MRI, increasing the number of data points does nothing to improve the resolution, and hence standard quadrature rules may fail to converge. Thus developing a new rigorous mathematical framework to analyze the convergence properties of the convolutional gridding technique is highly desirable.

There are, of course, other algorithms that recover f from its nonuniform Fourier data. For example, a well studied problem in sampling theory is to recover a band-limited \hat{f} from nonuniform samples (see, e.g., [1, 2, 13, 12] for history and references). Techniques using compressed sensing are also becoming more prevalent, e.g., [26]. While these approaches each have their advantages, they are primarily solved using iterative algorithms, thus making them somewhat difficult to analyze. The purpose of this paper is to establish a theoretical framework for the forward convolutional gridding method. Specifically, instead of viewing convolutional gridding essentially as an interpolation-numerical quadrature technique, i.e., as a means to approximate uniform Fourier coefficients and then perform the (inverse) FFT, we will consider the convolutional gridding algorithm to be a *Fourier frame approximation* of the underlying function.¹ To do this, we incorporate the convolutional gridding window function into the representation of f . That is, we express f in terms of the Fourier basis divided by the window function (rather than the standard Fourier basis). There are several advantages to our approach. First, it completely eliminates the intermediate interpolation step and the errors associated with it. Second, it provides a framework for constructing both the DCFs and the window function using rigorous mathematical arguments while satisfying the need for efficient computation. Finally, it establishes a forward algorithm for the finite (Fourier) frame approximation, which may be useful in other applications where data are sampled in a nonstandard way.

Our paper is organized as follows. In section 2 we review the convolutional gridding algorithm and link it to the Fourier frame approximation. We demonstrate how the DCFs for the convolutional gridding algorithm can be chosen to generate the best finite frame approximation from the given Fourier frame data. In section 3 we illustrate how our frame theoretic convolutional gridding algorithm can also be used to detect edges of a piecewise smooth function from nonuniform Fourier data. Numerical examples are given in section 4, and in section 5 we provide some concluding remarks. Finally, it establishes a forward algorithm for the finite (Fourier) frame approximation, which may be useful in other applications where data are sampled in a nonstandard way. We call our new technique the *frame theoretic convolutional gridding method* (FCG).

2. Convolutional gridding and the finite frame approximation.

2.1. The convolutional gridding algorithm. Let f be an unknown piecewise analytic function supported in $[0, 1]$. Suppose we are given the finite samples

$$(2.1) \quad \hat{f}_j = \langle f, \varphi_j \rangle = \int_0^1 f(x) e^{-2\pi i \lambda_j x} dx, \quad -n \leq j \leq n,$$

¹We note that in [3] it was suggested that certain MRI sampling patterns may indeed be viewed as Fourier frame sequences.




FIG. 1. Nonuniform sampling schemes (right half plane), $N = 16$.

where $\{\lambda_j : j \in \mathbb{Z}\}$ is a (nonuniform) sequence in \mathbb{R} . Our goal is to recover the underlying function f .

For illustrative purposes we will consider the following sampling schemes for $\{\hat{f}(\lambda_j)\}_{j=-n}^n$, depicted in Figure 1:

1. Jittered sampling:

$$(2.2) \quad \lambda_k = k \pm \tau_k, \quad \tau_k \sim U[0, \theta], \quad k = -n, -(n-1), \dots, n.$$

Here $U[0, \theta]$ denotes a uniform distribution on the interval $[0, \theta]$. The τ_k 's are independent, identically distributed random variables and represent a uniform jitter about the equispaced nodes with a maximal jitter of θ . Further, both positive and negative jitters are equiprobable, with the sign of jitter at each node being independent of the sign of jitter at any other node. We will always assume that $|\theta| < 1/2$. One such example occurs is MRI, where Cartesian (equispaced) scans may have stochastic jitters in the scanning trajectory caused by the inaccuracies in generating the magnetic field gradients.

2. Log sampling: The samples are acquired at logarithmic intervals, with more samples acquired in lower frequencies. Specifically, $|\lambda_k|$ is logarithmically distributed between 10^{-v} and n with $v > 0$ and $2n + 1$ being the total number of samples. This sampling pattern is inspired by non-Cartesian sampling patterns in MRI, where typical data acquisition schemes oversample the low frequencies while undersampling the high frequencies.

To motivate the convolutional gridding method, we begin by writing f as its inverse Fourier transform,

$$(2.3) \quad f(x) = \int_{-\infty}^{\infty} \hat{f}(\lambda) e^{2\pi i \lambda x} d\lambda.$$

Since f is assumed to be piecewise analytic in $[0, 1]$, \hat{f} cannot be compactly supported. However, our finite samples lie in the domain $[\lambda_{-n}, \lambda_n]$, so the best we can hope to approximate is

$$(2.4) \quad \tilde{f}(x) = \int_{\lambda_{-n}}^{\lambda_n} \hat{f}(\lambda) e^{2\pi i \lambda x} d\lambda.$$

Unfortunately, finding a good quadrature rule for (2.4) is difficult, due to the oscillatory nature and slow decay of $\hat{f}(\lambda)$ as $\lambda \rightarrow \infty$. Moreover, the set of samples $\{\hat{f}(\lambda_j)\}_{j=-n}^n$ are predetermined, so we are not at liberty to choose them.

The approximation of (2.4) may take the form

$$(2.5) \quad S_n \tilde{f}(x) := \sum_{k=-n}^n \alpha_k \hat{f}(\lambda_k) e^{2\pi i \lambda_k x},$$

where the weights, α_k , are known as the DCFs in the MR imaging community. For example, the DCFs may be generated using the well known trapezoidal rule:

$$(2.6) \quad \alpha_k = \lambda_{k+1} - \lambda_k.$$

In two dimensions, we may compute a Voronoi tessellation of the non-Cartesian nodes and then assign individual cell areas as the density compensation weights. For non-Cartesian sampling patterns which intentionally oversample the low frequencies while undersampling the high frequencies, the approximation (2.5), where the DCFs are chosen by *any* quadrature rule, e.g., (2.6), will fail to converge. That is because as the sampling is increased the space between the samples does not decrease, i.e., the sampling is not refined, and standard quadrature rules are therefore not applicable. Moreover, in approximating (2.4) instead of (2.3), we may incur significant error. As noted above, the underlying function is assumed to be piecewise smooth and supported in $[0, 1]$, and hence it has associated slow decay rate in the Fourier domain. If in addition the function exhibits highly oscillatory behavior, even more energy will be lost in the discarded high frequency modes.

This issue is partially addressed in [18, 37], where spectral reprojection [21] was applied to modified versions of (2.5). In essence, it is possible to reproject (2.5) onto a polynomial basis for which the high frequency information beyond $|\lambda_n|$ contributes only exponentially small values onto the low modes in the reprojection basis. It is important to note, however, that standard filtering algorithms are not effective in reducing the error. This is because classical filters are designed to mollify the high frequency coefficients in the harmonic series, that is, for $\lambda_j = j$, and indeed the numerical convergence estimates rely on the orthogonality of the harmonic basis [36].

Regardless of how the DCFs are chosen, the computational cost associated with evaluating (2.5) directly is significant, $\mathcal{O}(N^2)$, and even more so in multidimensions, since the FFT cannot be directly applied. Hence a speedup mechanism was devised to compute nonuniform sums such as (2.5) efficiently. The procedure goes by several names—the nonuniform FFT [14], the nonequispaced FFT [15], FFTs for nonuniform grids [35], or convolutional gridding in the MR imaging community [24, 28]. We note also that iterative algorithms have become increasingly popular in computing the NFFT [4, 11, 25, 30, 31]. However, the corresponding computational costs may still be too expensive for some applications, such as MRI. Moreover, it is difficult to analyze the convergence properties of most iterative algorithms. Thus we are motivated to both analyze and improve the direct methods of computing NFFT, which is the main purpose of this investigation.

The convolutional gridding procedure involves moving the nonuniform measurements, $\hat{f}(\lambda_j)$, to an equispaced grid (gridding) via a convolution operation. To accomplish this, we define w to be a smooth window function and consider the function $g = fw$. In addition, $w > 0$ should be essentially compactly supported so as to avoid significant aliasing error in the reconstruction of $f = g/w$. Some popular choices for window functions include Gaussian and Kaiser–Bessel functions [24, 28, 29].

The standard Fourier expansion is used to obtain g ,

$$(2.7) \quad g(x) = \sum_{l \in \mathbb{Z}} \hat{g}(l) e^{2\pi i l x},$$

where $\hat{g}(l) = \int_0^1 g(x)e^{-2\pi ilx} dx$, and is determined by the convolution

$$(2.8) \quad \hat{g}(l) = (\hat{f} * \hat{w})(l) = \int_{-\infty}^{\infty} \hat{f}(\lambda)\hat{w}(l - \lambda)d\lambda.$$

Recall that one of the difficulties in approximating (2.3) was the slow decay rate of \hat{f} . Convolutional gridding attempts to address this issue by requiring \hat{w} to be essentially compactly supported. True compact support is of course impossible if the window function w is also strictly compactly supported. One motivation for choosing w as either the Kaiser–Bessel or Gaussian function is that their Fourier coefficients are explicitly known. Any similar mollifier (bump function) should suffice, however, and some may have better convergence properties.

We can now approximate $\hat{g}(l)$ from the given samples $\{\hat{f}(\lambda_j)\}_{j=-n}^n$ using a quadrature rule of the form

$$(2.9) \quad \hat{g}(l) \approx \sum_{|j| \leq n} \alpha_j \hat{f}(\lambda_j) \hat{w}(l - \lambda_j).$$

We note that generating the optimal DCFs, $\{\alpha_j\}_{j=-n}^n$ is an ongoing research problem, since, not surprisingly, using quadrature weights such as (2.6) to solve (2.9) may not be satisfactory for many applications. Iterative algorithms that take into consideration the full construction of (2.10) seem to yield the best accuracy; see, e.g., [23, 29, 32]. As a consequence of our frame theoretic approach to the convolutional gridding algorithm, a new technique to determine DCFs will be presented later in subsection 2.3.

Substituting (2.9) into (2.7) yields

$$(2.10) \quad g(x) \approx \sum_{l \in \mathbb{Z}} \sum_{|j| \leq n} \alpha_j \hat{f}(\lambda_j) \hat{w}(l - \lambda_j) e^{2\pi ilx},$$

and division by the window function w provides an approximation to f as

$$(2.11) \quad f(x) \approx \sum_{l \in \mathbb{Z}} \sum_{|j| \leq n} \alpha_j \hat{f}(\lambda_j) \hat{w}(l - \lambda_j) \frac{e^{2\pi ilx}}{w(x)}.$$

Finally, since \hat{w} is deliberately chosen to have fast decay, (2.11) can be truncated for l , leading to the practical *convolutional gridding* computation for f ,

$$(2.12) \quad A_{cg}(f)(x) = \sum_{|j| \leq n} \sum_{|l - \lambda_j| \leq q} \alpha_j \hat{f}(\lambda_j) \hat{w}(l - \lambda_j) \frac{e^{2\pi ilx}}{w(x)},$$

where q is a truncation threshold to be specified. In practical applications, the reconstruction is often “zero padded” so that the computation of (2.12) near the boundary is not performed. This prevents dividing by very small values of w .

The steps of the convolutional gridding algorithm are enumerated in Algorithm 1 while a graphical illustration of the FFT reconstruction and window compensation is provided in Figure 2.

ALGORITHM 1. *Given \hat{f} at the nonequispaced measurement nodes λ_k , $k = -n, \dots, n$. Choose truncation parameter q , interpolating window function w , and DCFs α_k in (2.12).*




FIG. 2. A graphical illustration of the convolutional gridding method. Note that w is chosen for illustrative purposes only, as the corresponding decay rate of \hat{w} is too slow for practical applications.

(1) Interpolate to equispaced nodes via convolution:

$$(2.13) \quad \hat{g}(l) = \sum_{k \text{ st. } |l-\lambda_k| \leq q} \alpha_k \hat{f}(\lambda_k) \hat{w}(l-w_k) \approx \hat{g}(l), \quad l = -n, \dots, n.$$

(2) Perform standard FFT computation:

$$(2.14) \quad S_n \tilde{g}(x_p) = \sum_{l=-n}^n \hat{g}(l) e^{2\pi i l x_p}, \quad x_p = \frac{p}{n}, \quad p = 0, \dots, n-1.$$

(3) Compensate for using the interpolating window function:

$$A_{cg}(f)(x_p) = \frac{S_n \tilde{g}(x_p)}{w(x_p)}.$$

Since \tilde{g} is also a piecewise smooth function in $[0, 1]$, its Fourier reconstruction will likewise yield the Gibbs phenomenon. Hence the second step is typically filtered, i.e., the coefficients are multiplied by an admissible filter function $\sigma(\eta) = \sigma(\frac{l}{N})$. An in depth discussion on Fourier filters can be found in [36].

2.2. The Fourier frame approximation. Convolutional gridding, as it is described in Algorithm 1, is widely used in a variety of applications. A main advantage is that the process is a forward operation, that is, no inverse problem need be solved in the calculation. Moreover, since we are using the standard Fourier expansion, the fast Fourier transform can be employed to expedite the computation. Thus the process is straightforward to apply and very efficient. However, the error analysis of convolutional gridding is not yet fully understood, and in fact there are some instances for which the method fails to converge [37]. Below we provide a rigorous analysis of convolutional gridding by linking it to a numerical frame approximation for piecewise analytic functions f [7]. By applying the corresponding error analysis from [33], we are able to demonstrate the existence of DCFs and window functions to ensure the numerical convergence of (2.12). To this end, we first review the frame definition and the numerical frame approximation.

Let \mathcal{H} be a separable Hilbert space. We say $\{\varphi_j : j \in \mathbb{Z}\} \subseteq \mathcal{H}$ is a *frame* for \mathcal{H} if there exists some positive constants A, B such that for all $f \in \mathcal{H}$

$$A\|f\|_{\mathcal{H}}^2 \leq \sum_{j \in \mathbb{Z}} |\langle f, \varphi_j \rangle|^2 \leq B\|f\|_{\mathcal{H}}^2.$$

For a frame $\{\varphi_j : j \in \mathbb{Z}\} \subseteq \mathcal{H}$, its frame operator $S : \mathcal{H} \rightarrow \mathcal{H}$ is defined as

$$S(f) := \sum_{j \in \mathbb{Z}} \langle f, \varphi_j \rangle \varphi_j, \quad f \in \mathcal{H}.$$

The standard frame expansion has the form

$$(2.15) \quad f = S^{-1}Sf = \sum_{j \in \mathbb{Z}} \langle f, \varphi_j \rangle \tilde{\varphi}_j = \sum_{j \in \mathbb{Z}} \hat{f}(\lambda_j) \tilde{\varphi}_j,$$

where $\{\tilde{\varphi}_j := S^{-1}\varphi_j : j \in \mathbb{Z}\}$ is called the (*canonical*) *dual frame* and we have defined the frame coefficients as $\hat{f}(\lambda_j) := \langle f, \varphi_j \rangle$. In our investigation we will assume that the sampling sequence $\{\varphi_j : j \in \mathbb{Z}\}$ is a frame in $L^2[0, 1]$.²

Since typically the dual frame is not explicitly known, the (modified) frame algorithm [7] is often used to calculate (2.15). Other inverse frame operator approximation methods have also been developed [6, 8, 9, 22]. Here we choose to use the method of *admissible frames* [33] to obtain an approximation of $\tilde{\varphi}_j$. This method provides a link from (2.15) to (2.12) and therefore relates its corresponding convergence properties.

Briefly, an admissible frame is designed to compute the inverse frame operator when the given frame may only be weakly localized (the off-diagonal decay has order less than 1). Specifically, we define the following [33].

DEFINITION 1. A frame $\{\psi_l : l \in \mathbb{Z}\}$ is *admissible with respect to the frame* $\{\varphi_j : j \in \mathbb{Z}\}$ if the following two conditions hold:

- (i) There exist some positive constants c_0 and $t > 1$ such that

$$|\langle \psi_j, \psi_l \rangle| \leq c_0(1 + |j - l|)^{-t}, \quad j, l \in \mathbb{Z}.$$

²We note that this assumption is validated in [3] for some non-Cartesian sampled MRI data.

(ii) *There exist some positive constants c_1 and $s > \frac{1}{2}$ such that*

$$|\langle \varphi_j, \psi_l \rangle| \leq c_1(1 + |j - l|)^{-s}, \quad j, l \in \mathbb{Z}.$$

Now suppose $\{\psi_l : l \in \mathbb{Z}\}$ is an admissible frame with respect to the frame $\{\varphi_j : j \in \mathbb{Z}\}$. As shown in [33], the dual frame $\tilde{\varphi}_j = S^{-1}\varphi_j$ can be approximated by

$$(2.16) \quad \tilde{S}_n^{-1}\varphi_j := \sum_{|l| \leq k} b_{l,j}\psi_l,$$

where $B := [b_{l,j}]_{|l| \leq k, |j| \leq n}$ is the Moore–Penrose pseudoinverse of the matrix

$$(2.17) \quad \Psi := [\langle \varphi_j, \psi_l \rangle]_{|j| \leq n, |l| \leq k}.$$

The relationship between n and k is given by $n = k + ck^{\frac{1}{2s-1}}$, where c is a constant depending on c_1, s , and the frame bounds of $\{\varphi_j : j \in \mathbb{Z}\}$ [33].

The formula given in (2.16) can be directly inserted into the truncated form of (2.15) as an approximation of f , and, as shown in [33], the rate of convergence is of the same order as that of approximations using the new expansion frame, $\{\psi_l\}$.

As an example, consider the jittered samples (2.2) with $\theta = \frac{1}{4}$ and its corresponding frame $\{\varphi_j(x) = e^{2\pi i\lambda_j} : j \in \mathbb{Z}\}$. The standard Fourier basis $\{\psi_l(x) = e^{2\pi ilx} : l \in \mathbb{Z}\}$ is admissible with respect to φ_j for any positive number t and $s = 1$.

Substituting (2.16) into the frame expansion (2.15) and truncating the infinite summation of j yields the following approximation of f :

$$(2.18) \quad A_{frm}(f) := \sum_{|l| \leq m} \sum_{|j| \leq n} \hat{f}(\lambda_j)b_{l,j}\psi_l.$$

To link (2.18) to the convolutional gridding method, $A_{cg}(f)$ in (2.12), we take $\psi_l(x) = \frac{e^{2\pi ilx}}{w(x)}$.

We next provide the error analysis for the frame approximation method, $A_{frm}(f)$, following the approach of admissible frames proposed in [33]. To this end, we first show that $\{\psi_l : l \in \mathbb{Z}\}$ is an admissible frame with respect to $\{\varphi_j : j \in \mathbb{Z}\}$.

PROPOSITION 2.1. *If there exist some positive constants α_L and α_U such that $\alpha_L \leq w(x) \leq \alpha_U$ for all $x \in [0, 1]$, then $\{\psi_l(x) = \frac{e^{2\pi ilx}}{w(x)} : l \in \mathbb{Z}\}$ is a Riesz basis (and therefore also a frame [7]) for $L^2[0, 1]$ with frame bounds $1/\alpha_U^2$ and $1/\alpha_L^2$.*

Proof. Recall that a Riesz basis is the image of a bounded invertible mapping of an orthonormal basis [7]. Since $\{e^{2\pi ilx} : l \in \mathbb{Z}\}$ is an orthonormal basis for $L^2[0, 1]$ and the bound of $w(x)$ ensures that the mapping $U(f) := f/w$ is bounded invertible in $L^2[0, 1]$, $\{\psi_l(x) = U(e^{2\pi ilx}) : l \in \mathbb{Z}\}$ is a Riesz basis for $L^2[0, 1]$. The frame bounds follow from a direct calculation. \square

We next discuss the conditions for which $\{\psi_l(x) : l \in \mathbb{Z}\}$ is an admissible frame.

PROPOSITION 2.2. *If the 1-periodic extension of $w(x)$ is $C^p(\mathbb{R})$ smooth for some $p \in \mathbb{N}$ and $w(x) > 0$ for any $x \in [0, 1]$, then for the jittered sampling $\{\lambda_j : j \in \mathbb{Z}\}$ given in (2.2), $\{\psi_l(x) = \frac{e^{2\pi ilx}}{w(x)} : l \in \mathbb{Z}\}$ is an admissible frame with respect to $\{\varphi_j(x) = e^{2\pi i\lambda_j x} : j \in \mathbb{Z}\}$ with $t = p + 1$ and $s = 1$.*

Proof. We will check the two admissible conditions in Definition 1 directly. For $j, l \in \mathbb{N}$, we have

$$(2.19) \quad \langle \psi_j, \psi_l \rangle = \int_0^1 \frac{1}{w^2(x)} e^{2\pi i(j-l)x} dx.$$

Since $w \in C^p(\mathbb{R})$ and $w(x) > 0$ for all $x \in [0, 1]$, we have $1/w^2 \in C^p(\mathbb{R})$. Note that (2.19) can be viewed as the $(l - j)$ th Fourier coefficient of $1/w^2$. If a function is in $C^p(\mathbb{R})$, then its Fourier coefficients decay as $p + 1$, [16, 20]. Thus $\{\psi_l(x) : l \in \mathbb{Z}\}$ satisfies the first admissible condition with $t = p + 1$.

To check the second admissible condition, we note that

$$\langle \varphi_j, \psi_l \rangle = \int_0^1 \frac{1}{w(x)} e^{2\pi i(\lambda_j - l)x} dx.$$

By assumption, $1/w$ is bounded on $[0, 1]$. Hence for $j, l \in \mathbb{N}$, it is clear that for some $c > 0$ we have

$$|\langle \varphi_j, \psi_l \rangle| \leq c(1 + |\lambda_j - l|)^{-1}.$$

The second admissibility condition is readily verified for $s = 1$ in the jittered sampling case, (2.2). \square

The convergence results for the frame approximation method $A_{frm}(f)$ can now be given. In what follows, when it is clear from the context, we will use $\|\cdot\|$ to denote all norms. This notation includes the $L^2[0, 1]$ norm $(\int_0^1 |f(x)|^2 dx)^{1/2}$ on a function f , the Euclidean norm of a vector, or the spectral norm of a matrix.

THEOREM 2.3. *Suppose the 1-periodic extension of $w(x)$ is $C^p(\mathbb{R})$ smooth for some $p \in \mathbb{N}$ and there exists a positive constant c_0 such that $|\hat{f}(l)| \leq c_0(1 + |l|)^{-p-1}$ for all $l \in \mathbb{Z}$. If there exist positive constants α_L and α_U such that $\alpha_L \leq w(x) \leq \alpha_U$ for any $x \in [0, 1]$, then for the jittered sampling case, $\{\lambda_j : j \in \mathbb{Z}\}$ with $m = (1 + 2c_1^2 \alpha_U^4)n$ in (2.18), we have*

$$\|f - A_{frm}(f)\| \leq cn^{-p-1/2}$$

for some positive constant $c > 0$.

Proof. From Proposition 2.2, we know $\{\psi_l(x) = \frac{e^{2\pi i l x}}{w(x)} : l \in \mathbb{Z}\}$ is an admissible frame with respect to $\{\varphi_j(x) : j \in \mathbb{Z}\}$ with $t = p + 1$ and $s = 1$. The desired result follows immediately from Theorem 5.1 in [33].³ \square

Note that Theorem 2.3 does not apply to the case of log sampling, since in general its corresponding sequence does not constitute a frame. However, as mentioned in the introduction, in [3] it was suggested that some MRI sampling patterns may yield Fourier frames.

2.3. Linking convolutional gridding and the Fourier frame approximation. We now have two different approaches for approximating piecewise smooth f in $[0, 1]$ from nonuniform Fourier data—the convolutional gridding method, $A_{cg}(f)$ in (2.12), and the Fourier frame approximation, $A_{frm}(f)$ in (2.18). In order to link the two approximations, we first slightly modify the summation in $A_{cg}(f)$ to make the summation index consistent, and define

$$(2.20) \quad \tilde{A}_{cg}(f) := \sum_{|j| \leq n} \sum_{|l| \leq m} \alpha_j \hat{f}(\lambda_j) \hat{w}(l - \lambda_j) \frac{e^{2\pi i l x}}{w(x)}.$$

When \hat{w} decays fast or is compactly supported, the above modification should not significantly change the computational results or the convergence analysis.

³Theorem 5.1 in [33] requires the use of several technical propositions and is therefore omitted here.

In order to compare $A_{frm}(f)$ to the (modified) convolutional gridding method $\tilde{A}_{cg}(f)$, we first rewrite both methods using matrix notation. By defining the vector $\hat{\mathbf{f}} := (\hat{f}(\lambda_j) : |j| \leq n)$, (2.20) yields

$$(2.21) \quad \tilde{A}_{cg}(f) = \sum_{|l| \leq m} c_l \psi_l, \quad \mathbf{c} = (c_l : -m \leq l \leq m) = \Omega \mathbf{D} \hat{\mathbf{f}},$$

where $\Omega = [\hat{w}(l - \lambda_j)]_{|l| \leq m, |j| \leq n}$ and $\mathbf{D} = \text{diag}(\alpha_j : |j| \leq n)$. Likewise, we rewrite A_{frm} in (2.18) as

$$(2.22) \quad A_{frm}(f) = \sum_{|l| \leq m} d_l \psi_l, \quad \mathbf{d} = (d_l : -m \leq l \leq m) = \mathbf{B} \hat{\mathbf{f}} = \Psi^\dagger \hat{\mathbf{f}},$$

where Ψ^\dagger is the Moore–Penrose pseudoinverse of Ψ in (2.17).

Theorem 2.3 provides convergence results for $A_{frm}(f)$. We are now able to establish the difference between the approximations $\tilde{A}_{cg}(f)$ and $A_{frm}(f)$.

THEOREM 2.4. *Suppose there exist some positive constants α_L and α_U such that $\alpha_L \leq w(x) \leq \alpha_U$ for all $x \in [0, 1]$. Then there exists a positive constant c such that*

$$\|\tilde{A}_{cg}(f) - A_{frm}(f)\| \leq \frac{1}{\alpha_L} \|(\Psi^* \Psi)^{-1}\| \| \Psi^* \Psi \Omega \mathbf{D} - \Psi^* \|_F \| \hat{\mathbf{f}} \|,$$

where $\|\cdot\|_F$ is the Frobenius norm and Ψ^* is the conjugate transpose of Ψ .

Proof. From (2.21) and (2.22) we have

$$\|\tilde{A}_{cg}(f) - A_{frm}(f)\| = \left\| \sum_{|l| \leq m} (c_l - d_l) \psi_l \right\| = \left\| \sum_{|l| \leq m} (c_l - d_l) e^{2\pi i l x} \frac{1}{w(x)} \right\|.$$

It follows that

$$\|\tilde{A}_{cg}(f) - A_{frm}(f)\| \leq \|1/w\|_\infty \left\| \sum_{|l| \leq m} (c_l - d_l) e^{2\pi i l x} \right\| \leq \frac{1}{\alpha_L} \|\mathbf{c} - \mathbf{d}\|.$$

To estimate $\|\mathbf{c} - \mathbf{d}\|$, we observe that

$$\mathbf{c} - \mathbf{d} = (\Omega \mathbf{D} - \Psi^\dagger) \hat{\mathbf{f}} = (\Omega \mathbf{D} - (\Psi^* \Psi)^{-1} \Psi^*) \hat{\mathbf{f}} = (\Psi^* \Psi)^{-1} (\Psi^* \Psi \Omega \mathbf{D} - \Psi^*) \hat{\mathbf{f}}.$$

Therefore

$$\|\mathbf{c} - \mathbf{d}\| \leq \|(\Psi^* \Psi)^{-1}\| \| \Psi^* \Psi \Omega \mathbf{D} - \Psi^* \|_F \| \hat{\mathbf{f}} \|,$$

which finishes the proof. \square

Remark 1. Theorem 2.4 depends upon the representation of the dual frame given by (2.16). Various versions of the frame algorithm in [7] do not yield the form needed in (2.22).

Remark 2. Theorem 2.4 actually suggests a way to select optimal DCFs in (2.9). Specifically, for any given sampling sequence $\{\lambda_j\}_{j=-n}^n$ and weight function w , the matrices Ψ and Ω are both fixed. Hence the optimal \mathbf{D} satisfies

$$(2.23) \quad \underset{\alpha \in \mathbb{R}^{2n+1}}{\text{argmin}} \|\Psi^* \Psi \Omega \mathbf{D} - \Psi^*\|_F,$$

where $\|\cdot\|_F$ is the Frobenius norm. The DCFs, $\{\alpha_j\}_{j=-n}^n$, are the diagonal elements in the diagonal matrix \mathbf{D} . In fact, we have the following.

LEMMA 2.5. Let $D = \text{diag}(\alpha_j : -n \leq j \leq n)$, where α_j , $-n \leq j \leq n$ are the DCFs in (2.23). Then the optimal DCFs α'_j are given by

$$\alpha'_j = \frac{\langle \Psi_j^*, \mathsf{T}_j \rangle}{\|\mathsf{T}_j\|^2},$$

where Ψ_j^* and T_j are the j th column of Ψ^* and $\Psi^*\Psi\Omega$, respectively.

Proof. The result follows from calculating

$$\|\Psi^*\Psi\Omega D - \Psi^*\|_F^2 = \sum_{j=-n}^n \|\alpha_j \mathsf{T}_j - \Psi_j^*\|^2$$

as the solution to (2.23). \square

Remark 3. By defining α_j as the diagonal elements of D , Theorem 2.4 allows us to exploit the frame approximation to obtain the optimal DCFs as they appear in the traditional convolutional gridding algorithm. However, the results can be improved by allowing D to be banded instead of diagonal. In particular, in this case we view the DCFs as linear combinations of weights used in approximating the finite frame approximation of f from its sampled nonuniform Fourier data, as opposed to viewing them as numerical quadrature weights that approximate the inverse Fourier transform. For example, suppose D is a $(2r - 1)$ -banded matrix with $r - 1$ subdiagonals above and below the main diagonal, respectively. The optimal banded matrix is determined by solving

$$(2.24) \quad \operatorname{argmin} \{ \|\Psi^*\Psi\Omega D - \Psi^*\|_F : D \text{ is } (2r - 1)\text{-banded} \}.$$

If we define $\mathsf{T} := \Psi^*\Psi\Omega$, it follows that

$$\|\mathsf{T}D - \Psi^*\|_F^2 = \sum_{j=1}^n \|\mathsf{T}d_j - \Psi_j^*\|^2,$$

where d_j are the column vectors of D . The solution to (2.24) is also the minimizer of $\|\mathsf{T}d_j - \Psi_j^*\|^2$ for each j . This is a standard least-squares problem with the explicit solution $d_j = \mathsf{T}^\dagger \Psi_j^*$. Note that the number of nonzero components in each d_j is at most $2r - 1$. That is, we actually need only to calculate the Moore–Penrose pseudoinverse of the corresponding block (at most $(2r - 1) \times (2r - 1)$) instead of the whole matrix T .

The computational cost of calculating D is $O(nr^2)$. Once D is determined, the cost of evaluating f using the FFT is $\mathcal{O}(n \log n)$. Note that the computational cost of determining the DCFs would still be consistent with that of the FFT if we choose $r = \mathcal{O}(\log n)$. In fact, we observe in some of our experiments in section 4 that numerical convergence, especially in the log sampling case, requires r to grow with n . Indeed, the method may not converge for constant r .

Figures 3 and 4 compare the DCFs generated by (A) the trapezoidal rule, (2.6), (B) the iterative procedure proposed in [29], and (C) our new FCG approach for both the jittered (2.2) with $\theta = \frac{1}{4}$ and log sampling cases. As discussed in section 2.1, iterative techniques typically yield more suitable DCFs than the quadrature weights, such as the trapezoidal rule. As is evident here, our new method generates comparable DCFs as those constructed by iterative procedures, although it has less computational




FIG. 3. DCFs for jittered sampling with $n = 128$. We used the window function $w(x) = e^{-a|x-.5|}$, where $a = 5 \times 10^{-5}$ for (B) and (C).




FIG. 4. DCFs for log sampling with $n = 128$. We used the window function $w(x) = e^{-a|x-.5|}$, where $a = 5 \times 10^{-5}$ for (B) and (C).

complexity. Note that although the corresponding sequence for the log sampling scheme does not constitute a frame, the same algorithm used to generate the optimal (frame theoretic) DCFs can still be applied. In section 4, we will also demonstrate that the numerical approximation using the FCG method for log sampling sequences also yield high accuracy for smooth functions.

Remark 4. Figure 3 suggests that interpreting the reconstruction of f as an approximation to its inverse Fourier transform and subsequently applying a quadrature formula (of general order) is suitable when the data are close to uniform. On the other hand, when the given data are highly nonuniform, Figure 4 shows that standard quadrature rules “compensate” by *increasing* the DCF values in the sparse regions. Unfortunately, in many signal processing applications the high frequency data are likely to contain noise which should be mollified rather than enhanced. Both the iterative and frame theoretic approaches yield more evenly weighted DCFs. We point out that the iterative approach involves solving a full $(2n - 1) \times (2n - 1)$ linear system,⁴ while the frame theoretic approach only solves $2n - 1$ much smaller $(2r - 1) \times (2r - 1)$ linear systems. In particular we use the diagonal matrix, that is, $r = 1$, to generate Figures 3 and 4 for the frame theoretic approach.

3. Applications to edge detection. Edge detection is important in many signal processing and imaging applications. For example, in MRI, tissue classification inherently depends on the ability to determine the internal boundaries of an image. In what follows we show how our new FCG algorithm can be applied to locate jump discontinuities in piecewise smooth functions when data are sampled nonuniformly in the Fourier domain.⁵ Determining jump locations from uniformly sampled Fourier data is in itself a difficult problem, mainly because an edge is a local feature while the Fourier data are of a global nature. While there are a variety of algorithms that have been developed to do this, they are in general not easily adapted to the nonuniform case. In particular, any kind of interpolation (regridding) procedure produces oscillations that do not have extractable information. A notable exception can be found in [34], where the the method from [19] was adapted successfully for nonuniform sampling by employing a sparsity enforcing minimization procedure. However, analyzing the corresponding error is difficult, and, moreover, the technique relies on solving an inverse problem which may be computationally intensive or even fail to converge.

A frame theoretic approach to edge detection from nonuniform Fourier data was first suggested in [17]. That method relies on the (modified) frame algorithm, e.g., [7], to compute the dual frame, which may also be computationally intensive or ill-conditioned. Here we extend the ideas from [17] but replace the frame algorithm with the more efficient convolutional gridding algorithm. The link between the Fourier frame approximation and convolutional gridding allows us to do this.

Consider a piecewise continuous function $f : [0, 1] \rightarrow \mathbb{R}$ with $f(0) = f(1) = 0$. We will define the associated *jump function* $[f] : (0, 1) \rightarrow \mathbb{R}$ by

$$(3.1) \quad [f](x) = \lim_{t \downarrow x} f(t) - \lim_{t \uparrow x} f(t).$$

Notice that $[f]$ is well-defined since f is piecewise continuous. Furthermore, $[f] = 0$

⁴In some sense, the iterative DCF approach in [29] attempts to construct a kernel from the nonharmonic series expansion that resembles a regularized Dirac delta function. This is not the main focus of our current investigation but will be analyzed in more detail in future work.

⁵We describe only the one-dimensional case, where an edge is equivalent to a jump discontinuity. Our technique may be adapted to two dimensions if we write (3.4) as a separable function. Results in this direction can be found in [27].

everywhere except at the jump discontinuity locations, where $[f]$ takes on the value of the jump. For simplicity of presentation, we assume that f has a single jump discontinuity (edge) $\xi \in (0, 1)$. In our numerical examples we will demonstrate that the method works to recover multiple jump locations, that is, wherever $[f](x) \neq 0$.

Let us again assume that $\{\varphi_j(x) = e^{2\pi i \lambda_j x} : j \in \mathbb{Z}\}$ is a frame for $L^2[0, 1]$. We seek the jump location, ξ , and its corresponding jump value, $[f](\xi)$, from the given data $\hat{\mathbf{f}} := \{\hat{f}(\lambda_j) = \langle f, \varphi_j \rangle : -n \leq j \leq n\}$. In constructing a numerical algorithm to determine the jump discontinuities of f , it will be more convenient to write (3.1) as

$$(3.2) \quad [f](x) = [f](\xi)\delta_\xi(x),$$

where $\delta_\xi(x)$ is the indicator function taking the value 1 when $x = \xi$, and 0 otherwise. For multiple jumps, (3.2) is correspondingly $[f](x) = \sum_{j=1}^J [f](\xi_j)\delta_{\xi_j}(x)$, where J is the number of jump discontinuities in f .

Since the indicator function is nontrivial only at a single point whose set has zero measure, it is not practical to directly construct a meaningful approximation in $L^2[0, 1]$. Therefore, we regularize (3.2) and instead seek to approximate

$$(3.3) \quad [f](x) \approx [f](\xi)h_n(x),$$

where

$$(3.4) \quad h_n(x) := h\left(\frac{x - \xi}{\epsilon_n}\right), \epsilon_n > 0,$$

with $h(0) = 1$. We choose h to be a compactly supported bump function that approximates $\delta_\xi(x)$ such that h_n is supported in $[\xi - \epsilon_n, \xi + \epsilon_n]$ with $h_n(\xi) = 1$. As $\epsilon_n \downarrow 0$, $h_n(x) \rightarrow \delta_\xi(x)$. Numerical considerations dictate that h is smooth (e.g., bell-shaped) around 0 and $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$. We remark that the choice of ϵ_n is critical in how the method performs. As ϵ_n increases, the approximation is more regularized, but the edges are not as well localized. This trade-off can be decided on a case-by-case basis, depending on other external influences, such as the amount of and corruption in the data; see, e.g., [10, 17].

We now implement the ideas from the FCG algorithm to construct an approximation to (3.3) from the given finite information $\hat{\mathbf{f}} = \{\hat{f}(\lambda_j) : |j| \leq n\}$. To this end, we first establish a relationship between the Fourier coefficients of the regularized indicator function, $\widehat{h_n}(\lambda_j)$, and the given data $\hat{f}(\lambda_j)$. For $\lambda_j \neq 0$, integration by parts on the given Fourier samples (2.1) yields

$$(3.5) \quad \begin{aligned} \hat{f}_j &= \frac{1}{2\pi i \lambda_j} [f](\xi) e^{-2\pi i \lambda_j \xi} + \frac{1}{2\pi i \lambda_j} \int_0^1 f'(x) e^{-2\pi i \lambda_j x} dx, \quad -n \leq j \leq n \\ &= \frac{1}{2\pi i \lambda_j} [f](\xi) e^{-2\pi i \lambda_j \xi} + \mathcal{O}\left(\frac{1}{\lambda_j^2}\right). \end{aligned}$$

Thus, for large λ_j , we have

$$(3.6) \quad 2\pi i \lambda_j \hat{f}_j e^{2\pi i \lambda_j \xi} \approx [f](\xi).$$

On the other hand, direct calculation of the Fourier coefficients in (3.3) gives

$$\begin{aligned} \widehat{h_n}(\lambda_j) &= \int_{-\infty}^{\infty} h\left(\frac{x - \xi}{\epsilon_n}\right) e^{-2\pi i \lambda_j x} dx = \epsilon_n e^{-2\pi i \lambda_j \xi} \int_{-\infty}^{\infty} h(u) e^{-2\pi i \lambda_j u \epsilon_n} du \\ &= \epsilon_n e^{-2\pi i \lambda_j \xi} \hat{h}(\lambda_j \epsilon_n). \end{aligned}$$

Finally, from (3.6) we have

$$(3.7) \quad [f](\xi)\widehat{h}_n(\lambda_j) \approx 2\pi i\lambda_j\epsilon_n\hat{f}_j\hat{h}(\lambda_j\epsilon_n),$$

implying that we can obtain an approximation of $[f](\xi)\widehat{h}_n(\lambda_j)$ directly from the given data \hat{f}_j .

We now formulate our edge detection algorithm as follows: Suppose we are given

$$\hat{\mathbf{h}} := (2\pi i\lambda_j\epsilon_n\hat{f}_j\hat{h}(\lambda_j\epsilon_n) : |j| \leq n),$$

i.e., an approximation to the nonuniform Fourier coefficients of the regularized jump function (3.3). Following the frame theoretic convolutional gridding approach, we use (2.21) to obtain the approximation of $[f](\xi)h_n(x)$ as

$$E_{cg}(f) := \sum_{|l| \leq m} c_l \psi_l, \quad \mathbf{c} = \Omega D \hat{\mathbf{h}}.$$

Likewise, the frame approximation method (2.22) yields an approximation to $[f](\xi)h_n$ as

$$E_{frm}(f) := \sum_{|l| \leq m} d_l \psi_l, \quad \mathbf{d} = \Psi^\dagger \hat{\mathbf{h}}.$$

Analogous results to Theorem 2.4 hold here, although the overall accuracy for recovering the edges strongly depends on the error incurred in approximating (3.5) by (3.6), which clearly increases for multiple jumps. Once again the DCFs can be optimally chosen using (2.24). We demonstrate our results in section 4.

Remark 5. The FCG approach for edge detection given nonuniform Fourier data has two distinct advantages. First, it does not require interpolation to uniform nodes. It is evident from (3.5) that edge information comes from the high frequency part of the Fourier domain. Unfortunately, the sampling patterns in applications such as MRI are typically sparse in the high frequency domain, and the algorithms that interpolate these data to integer coefficients are neither accurate nor robust [37]. By contrast, the proposed FCG approach yields much smaller errors in the high frequency domain. Second, the FCG method is a forward operation that can be efficiently and accurately implemented using the FFT. This is in contrast to the frame based edge detection method introduced in [17], which requires approximating an inverse frame operator.

4. Numerical experiments. The examples below are used to demonstrate the FCG approach for function reconstruction and edge detection. In each case we compare our new method to the traditional convolutional gridding (CG) method, A_{cg} in (2.12), and the frame approximation method, A_{frm} in (2.18).⁶ Before presenting our numerical experiments, it is important to analyze the relationship between these three methods. The traditional CG method shares the same form as the FCG approach with $r = 1$, that is, when we use diagonal approximation in (2.24). Both are linear methods, with the FCG construction providing a new algorithm for choosing optimal DCFs. When $r > 1$, the FFG method can be viewed as a generalization of the traditional

⁶There are no discernible discrepancies when A_{cg} in (2.12) replaces the sum that is analyzed, \tilde{A}_{cg} in (2.20).

CG approach since it allows the use of banded rather than diagonal matrices. Note that since the FCG approach is an approximation of the frame approximation, (2.18), the best convergence we can hope to achieve is that obtained by (2.18). However, the standard frame approximation requires solving a $(2n - 1) \times (2n - 1)$ linear system, while the proposed FCG approach must solve $2n - 1$ smaller $(2r - 1) \times (2r - 1)$ linear systems. As r increases, we get closer to the frame approximation in (2.18). Thus we see the expected trade-off in computational efficiency and numerical convergence.

For convenience, in all of our experiments we use the window function

$$(4.1) \quad w(x) = e^{-a|x-0.5|},$$

where $a = 5 \times 10^{-5}$. We point out that (4.1) is only one such example that can be used in Theorem 2.4 and that the choice of window function may influence the numerical approximation, a topic that will be explored in future work.

Example 1. Consider the target function

$$f(x) = \cos^2(\pi(x - 0.5)^2) \sin(10(x - 0.5)^2), \quad x \in [0, 1].$$

Figure 5 displays the approximation results for jittered sampling, (2.2), with $\theta = \frac{1}{4}$. We compare the frame approximation (FA) method, (2.18), the convolutional gridding (CG) method, (2.12), and the new FCG, (2.21), for $r = 3$. For simplicity we employ the trapezoidal rule, (2.6), to determine the DCFs for the traditional CG method.

Figure 6 displays the analogous results for the log sampling pattern. Note in this case that the corresponding sequence $\{\varphi_j : j \in \mathbb{Z}\}$ may not constitute a frame except for its finite span. Nevertheless, the FCG method is still effective.

Remark 6. When jittered sampling is used, we do not expect to see much difference between the results using the traditional or FCG algorithm. This is because the DCFs look very similar in this case; see Figure 3. On the other hand, the log sampling case yields a substantial difference in the DCFs, and consequently in the numerical results. As noted previously, the iterative DCF algorithm yields qualitatively similar results but is more computationally intensive and is only established for the traditional (linear) convolutional gridding algorithm (e.g., $r = 1$).




FIG. 5. Comparisons of recovered functions (A) and pointwise errors (in log) (B) employing the FA method, the traditional CG method, and the FCG for jittered sampling with $\theta = \frac{1}{4}$ and $n = 200$.




FIG. 6. Comparisons of recovered functions (A) and pointwise errors (in log) (B) of these three methods: the FA method, the traditional CG method, and the FCG for log sampling.




FIG. 7. Approximation errors (in log) for different choices of r .

Figure 7 compares the log approximation error using the FCG algorithm with $r = 1, \mathcal{O}(\log n)$, and the Fourier frame approximation for $n = 16, 32, 64, 128$. Observe that for $r = 1$, that is, traditional convolutional gridding, the method fails to converge as n increases. In fact the method may not converge for any fixed r . The method does converge (assuming no round-off error) for $r = \mathcal{O}(\log n)$, however.

We now demonstrate the FCG as a means of detecting edges from given nonuniform Fourier data. Consider the test function.

Example 2. We consider the piecewise continuous function

$$f(x) = \begin{cases} 3/2, & 1/8 \leq x \leq 1/4, \\ \frac{7}{4} - \frac{x}{2} + \sin(2\pi x - 1/4), & 3/8 \leq x \leq 9/16, \\ \frac{11}{4}x - 5, & 11/16 \leq x \leq 7/8, \\ 0 & \text{otherwise,} \end{cases}$$

for which we are given Fourier coefficients at either the jittered or log samples.




FIG. 8. Example 2 and its corresponding target edge function, $[f](\xi)h_n(x)$ as in (3.3).




FIG. 9. Comparisons of the recovered edge functions and the pointwise approximation errors (in log) to the target edge function for jittered sampling.

As an approximate indicator function h in (3.4) we choose

$$h(x) = e^{-5x^2}, \quad \epsilon_n = 0.02.$$

Figure 8 shows Example 2 and its corresponding target edge function, $[f](\xi)h_n(x)$ as in (3.3).

Figures 9 and 10 compare the recovered edge functions and the approximation errors to the target edge function for given $n = 512$ jittered and log spaced Fourier samples, respectively. Once again we use the window function from (4.1) for the convolutional gridding and FCG, for which we chose $r = 25$. Fewer oscillations in smooth regions are evident using the FCG.

We note that it is not straightforward to analyze the edge detection results. In addition to the flexibility in choosing the window function and truncation parameters associated with convolutional gridding, the indicator function approximation (3.3) and additional error incurred in (3.5) can contribute to the overall approximation. Further, the results displayed in Figures 9 and 10 can be improved by postprocessing. For example, the sparsity enforcing procedure introduced in [34], with some modifications, should be applicable.




FIG. 10. Comparisons of the recovered edge functions and the pointwise approximation errors (in log) to the target edge function for log sampling.

5. Conclusion. By viewing the convolutional gridding (nonuniform FFT) method as a frame approximation, we are able to perform rigorous error analysis. This enables us to choose optimal DCFs for the convolutional gridding method from a frame theoretic rather than heuristic perspective. Furthermore, since the traditional (linear) convolutional gridding method fails to converge, we generalize our approach by defining the DCFs as the elements of a banded matrix instead of a vector, thereby increasing the overall accuracy of the approximation and generating numerical convergence. As the band increases, we get closer to the frame approximation of the underlying function. However, the computational cost for a bandwidth proportional to the log of the number of given Fourier samples remains of the same order as the traditional convolutional gridding algorithm. Our numerical results are consistent with our theoretical results, and even in the case where the underlying sampling pattern does not generate a frame, we are able to see improvement using our method. Finally, it is important to note that the FCG method is a forward algorithm and thus avoids the computational costs necessarily incurred by solving an inverse problem.

There are a large number of parameters in the convolutional gridding algorithm that can be modified for any particular circumstance. For example, greater accuracy is possible if a larger truncation number q is used in Algorithm 1, but this would increase the number of calculations. There also may be better ways to choose the admissible frame, which may well lead to other types of regridding algorithms that are not FFT compatible. In edge detection applications, using different regularizations for the indicator function, (3.4), may improve the results. However, doing so may also mean that a larger number of Fourier samples are needed to resolve the edges. We will explore all of these ideas in future work. Finally, the convolutional gridding algorithm extends analytically and numerically to two dimensions. The parameters of the method, however, are still chosen heuristically. In future work we will investigate extending the use of admissible frames, [33], to two dimensions, which may enable us to choose these parameters from a frame theoretic perspective.

REFERENCES

- [1] A. ALDROUBI AND K. GRÖCHENIG, *Nonuniform sampling and reconstruction in shift-invariant spaces*, SIAM Rev., 43 (2001), pp. 585–620.
- [2] J. J. BENEDETTO AND W. HELLER, *Irregular sampling and the theory of frames. I*, Note Mat., 10 (1992), pp. 103–125.

- [3] J. J. BENEDETTO AND H. C. WU, *Non-uniform sampling and spiral MRI reconstruction*, Proc. SPIE, 4119 (2000), pp. 130–141.
- [4] G. BEYLKIN, *On the fast Fourier transform of functions with singularities*, Appl. Comput. Harmon. Anal., 2 (1995), pp. 363–381.
- [5] D. S. BRIGGS, F. R. SCHWAB, AND R. A. SRAMEK, *Imaging*, in Synthesis Imaging in Radio Astronomy II, Vol. 180, 1999, pp. 127–149.
- [6] P. G. CASAZZA AND O. CHRISTENSEN, *Approximation of the inverse frame operator and applications to Gabor frames*, J. Approx. Theory, 103 (2000), pp. 338–356.
- [7] O. CHRISTENSEN, *An Introduction to Frames and Riesz Bases*, Applied and Numerical Harmonic Analysis, Birkhäuser Boston, Boston, 2003.
- [8] O. CHRISTENSEN AND T. STROHMER, *Methods for approximation of the inverse (Gabor) frame operator*, in Advances in Gabor Analysis, Appl. Numer. Harmon. Anal., Birkhäuser Boston, Boston, 2003, pp. 171–195.
- [9] O. CHRISTENSEN AND T. STROHMER, *The finite section method and problems in frame theory*, J. Approx. Theory, 133 (2005), pp. 221–237.
- [10] D. COCHRAN, A. GELB, AND Y. WANG, *Edge detection from truncated Fourier data using spectral mollifiers*, Adv. Comput. Math., 38 (2013), pp. 737–762.
- [11] A. DUTT AND V. ROKHLIN, *Fast Fourier transforms for nonequispaced data*, SIAM J. Sci. Comput., 14 (1993), pp. 1368–1393.
- [12] H. G. FEICHTINGER AND K. GRÖCHENIG, *Irregular sampling theorems and series expansions of band-limited functions*, J. Math. Anal. Appl., 167 (1992), pp. 530–556.
- [13] H. G. FEICHTINGER, K. GRÖCHENIG, AND T. STROHMER, *Efficient numerical methods in non-uniform sampling theory*, Numer. Math., 69 (1995), pp. 423–440.
- [14] J. A. FESSLER AND B. P. SUTTON, *Nonuniform fast Fourier transforms using min-max interpolation*, IEEE Trans. Signal Process., 51 (2003), pp. 560–574.
- [15] K. FOURMONT, *Non-equispaced fast Fourier transforms with applications to tomography*, J. Fourier Anal. Appl., 9 (2003), pp. 431–450.
- [16] C. GASQUET AND P. WITOMSKI, *Fourier Analysis and Applications: Filtering, Numerical Computation, Wavelets*, Texts Appl. Math. 30, Springer-Verlag, New York, 1999.
- [17] A. GELB AND T. HINES, *Detection of edges from nonuniform Fourier data*, J. Fourier Anal. Appl., 17 (2011), pp. 1152–1179.
- [18] A. GELB AND T. HINES, *Recovering exponential accuracy from non-harmonic Fourier data through spectral reprojection*, J. Sci. Comput., 51 (2012), pp. 158–182.
- [19] A. GELB AND E. TADMOR, *Detection of edges in spectral data. II. Nonlinear enhancement*, SIAM J. Numer. Anal., 38 (2000), pp. 1389–1408.
- [20] D. GOTTLIEB AND S. A. ORSZAG, *Numerical Analysis of Spectral Methods: Theory and Applications*, CBMS-NSF Regional Conference Series in Applied Mathematics 26, SIAM, Philadelphia, 1977.
- [21] D. GOTTLIEB, C.-W. SHU, A. SOLOMONOFF, AND H. VANDEVEN, *On the Gibbs phenomenon. I. Recovering exponential accuracy from the Fourier partial sum of a nonperiodic analytic function*, J. Comput. Appl. Math., 43 (1992), pp. 81–98.
- [22] K. GRÖCHENIG, *Acceleration of the frame algorithm*, IEEE Trans. Signal Process., 41 (1993), pp. 3331–3340.
- [23] R. D. HOGE, R. K. KWAN, AND G. B. PIKE, *Density compensation functions for spiral MRI*, Magn. Reson. Med., 38 (1997), pp. 117–128.
- [24] J. I. JACKSON, C. H. MEYER, D. G. NISHIMURA, AND A. MACOVSKI, *Selection of a convolution function for Fourier inversion using gridding*, IEEE Trans. Med. Imaging, 10 (1991), pp. 473–478.
- [25] J. KEINER, S. KUNIS, AND D. POTTS, *Using NFFT 3—a software library for various nonequispaced fast Fourier transforms*, ACM Trans. Math. Software, 36 (2009).
- [26] M. LUSTIG, D. DONOHO, AND J. M. PAULY, *Sparse MRI: The application of compressed sensing for rapid MRI imaging*, Magn. Reson. Med., 6 (2007), pp. 1182–95.
- [27] A. MARTINEZ, A. GELB, AND A. GUTIERREZ, *Edge detection from non-uniform Fourier data using the convolutional gridding algorithm*, J. Sci. Comput., to appear.
- [28] J. D. O’SULLIVAN, *A fast sinc function gridding algorithm for Fourier inversion in computer tomography*, IEEE Trans. Med. Imaging, 4 (1985), pp. 200–207.
- [29] J. G. PIPE AND P. MENON, *Sampling density compensation in MRI: Rationale and an iterative numerical solution*, Magn. Reson. Med., 41 (1999), pp. 179–186.
- [30] D. POTTS, G. STEIDL, AND M. TASCHE, *Fast Fourier transforms for nonequispaced data: A tutorial*, in Modern Sampling Theory, Appl. Numer. Harmon. Anal., Birkhäuser Boston, Boston, 2001, pp. 247–270.
- [31] D. POTTS AND M. TASCHE, *Numerical stability of nonequispaced fast Fourier transforms*, J. Comput. Appl. Math., 222 (2008), pp. 655–674.

- [32] H. SEDARAT AND D. G. NISHIMURA, *On the optimality of the gridding reconstruction algorithm*, IEEE Trans. Med. Imaging, 19 (2000), pp. 306–317.
- [33] G. SONG AND A. GELB, *Approximating the inverse frame operator from localized frames*, Appl. Comput. Harmon. Anal., 35 (2013), pp. 94–110.
- [34] W. STEFAN, A. VISWANATHAN, A. GELB, AND R. RENAUT, *Sparsity enforcing edge detection method for blurred and noisy Fourier data*, J. Sci. Comput., 50 (2012), pp. 536–556.
- [35] G. STEIDL, *A note on fast Fourier transforms for nonequispaced grids*, Adv. Comput. Math., 9 (1998), pp. 337–352.
- [36] E. TADMOR, *Filters, mollifiers and the computation of the Gibbs phenomenon*, Acta Numer., 16 (2007), pp. 305–378.
- [37] A. VISWANATHAN, A. GELB, D. COCHRAN, AND R. RENAUT, *On reconstruction from non-uniform spectral data*, J. Sci. Comput., 45 (2010), pp. 487–513.