

RATIONAL APPROXIMATIONS FOR TOMOGRAPHIC RECONSTRUCTIONS

MATTHEW REYNOLDS, GREGORY BEYLKIN AND LUCAS MONZÓN

ABSTRACT. We use optimal rational approximations of projection data collected in X-ray tomography to improve image resolution. Under the assumption that the object of interest is described by functions with jump discontinuities, for each projection we construct its rational approximation with a small (near optimal) number of terms for a given accuracy threshold. This allows us to augment the measured data, i.e., double the number of available samples in each projection or, equivalently, extend (double) the domain of their Fourier transform. We also develop a new, fast, polar coordinate Fourier domain algorithm which uses our nonlinear approximation of projection data in a natural way.

Using augmented projections of the Shepp-Logan phantom, we provide a comparison between the new algorithm and the standard Filtered Back-Projection (FBP) algorithm. We demonstrate that the reconstructed image has improved resolution without additional artifacts near sharp transitions in the image.

1. INTRODUCTION

As perceptively noted in [29], despite the development of many new algorithms for the inversion of the Radon transform, the quality of image reconstruction (in e.g., X-ray tomography) has not improved noticeably when compared with the output of the traditional Filtered Back-Projection (FBP) algorithm (see, for example [28]). This lack of improvement in image quality may be traced to the fact that the signal model for collected data is subject to the Nyquist sampling criterion. Since we are typically interested in reconstructing piece-wise continuous objects, the collected data (within the standard signal model) may be insufficient to resolve the image in the vicinity of discontinuities.

We introduce a different signal model for collected data and, as a consequence, for reconstruction. Assuming that the object of interest is described by functions with a limited number of jump discontinuities, the measured projections have, in the worst case, the same type of discontinuities (for most projections only the first derivative is discontinuous). As shown in [6], this implies that the projections are well approximated by proper rational functions yielding a robust recovery of signals from band-limited data. In this paper, we demonstrate that a rational model of

Key words and phrases. optimal rational approximations, signal models, X-ray tomography, Radon transform, polar grids, Unequally Spaced Fast Fourier Transform, Filtered Back Projection, image resolution.

The authors are with the Department of Applied Mathematics, University of Colorado, Boulder, CO 80309-0526, United States

This research was partially supported by NSF grants DMS-1009951, DGE-0801680, DMS-0602284, and DOE/ORNL grant 4000038129.

Inverse Problems, 29 (2013) 065020, <http://dx.doi.org/10.1088/0266-5611/29/6/065020> .

a (periodic) signal allows us, in principle, to recover its entire Fourier series from a small number of samples. In practice, even though the presence of noise limits such a recovery, a rational model still outperforms models based on the Nyquist sampling criterion. For objects with a limited number of isolated singularities, optimal rational approximations of projection data yield a significant improvement in resolution without introducing artifacts near singularities.

Classical Shannon-Nyquist sampling theory (see, for example, [36]) describes band-limited signals as a linear combination of sinc-functions. This signal model requires at least two samples per period corresponding to the highest frequency present in the signal. In practice, the data must be oversampled to attenuate the impact of noise and allow for local interpolation. The Shannon-Nyquist model may be generalized by replacing the sinc-functions by an alternative basis set, e.g., splines or wavelets [36]. However, in all such models the sampling requirements are directly related to the highest frequency present in the signal.

For band-limited periodic functions (i.e., trigonometric polynomials of degree less or equal to N), it is well-known that they may be recovered from their samples via

$$(1.1) \quad f(x) = \sum_{l=0}^{2N} f\left(\frac{l}{2N+1}\right) D_N\left(x - \frac{l}{2N+1}\right), \quad x \in [0, 1),$$

where

$$D_N(x) = \frac{1}{2N+1} \sum_{|l| \leq N} e^{2\pi i l x} = \frac{1}{2N+1} \frac{\sin(2N+1)\pi x}{\sin \pi x}$$

is the Dirichlet kernel.

In our signal model, a real-valued periodic function is represented by a rational function with $4M+1$ real parameters,

$$(1.2) \quad g(x) = a_0 + 2\operatorname{Re} \sum_{m=1}^M \frac{w_m}{e^{-2\pi i x + \eta_m} - 1}, \quad x \in [0, 1)$$

where $w_m \in \mathbb{C}$, $\eta_m \in \mathbb{C}$, $\operatorname{Re}(\eta_m) > 0$, and $a_0 \in \mathbb{R}$ is a constant. The frequency content of this rational function does not depend on the number of terms M but, rather, on the proximity of nodes $e^{-\eta_m}$, $m = 1, \dots, M$ to the unit circle (the use of the term “nodes” will become clear later). This is fundamentally different from the signal model (1.1), where the sampling rate is controlled by the degree N , i.e., the number of terms in the model. In other words, we have a sparse representation for functions with isolated singularities as described in [6].

One of the goals of this paper is to present a robust reconstruction algorithm for the signal model (1.2) in the presence of noise. Our algorithms are based on representing Fourier values of g via decaying exponentials [6]. Using this representation, we extend (double) the domain of the Fourier transform, or, equivalently, double the number of available samples in each projection. By augmenting the measured data in this fashion, we improve the resolution of reconstruction near singularities. This approach introduces a another mechanism for increasing resolution, different from previous approaches (see, for example, [14, 10, 9, 12, 13, 17, 16]).

We also present a new fast numerical algorithm for inverting the Radon transform. Following [3], we construct a polar grid in the Fourier domain that allows us to mimic the Fourier slice theorem in setting up the image reconstruction. Constructing an image from such grids via the Unequally Spaced Fast Fourier transform

(USFFT) [11, 2, 26] requires $\mathcal{O}(N^2 \log N)$ operations, where N is the number of projection samples and the number of projections is $\mathcal{O}(N)$.

We start in Section 2 by a brief description of our rational model and provide background information in Section 3. In Section 4 we introduce a new approach for the computation of weights in the rational model of projection data. Then, in Section 5, we describe a new, fast, Fourier domain algorithm (dubbed Polar Quadrature Inversion) for tomographic reconstruction using polar grids. Finally, in Section 6, we perform numerical tests demonstrating the improved resolution resulting from the use of nonlinear approximations for the projections. We also compare the performance of the new fast algorithm and the standard FBP algorithm.

2. A RATIONAL MODEL FOR SIGNALS

Let us describe analytic relations between samples of g and the parameters of the rational model (1.2). For this purpose, we derive a relation between the Fourier series coefficients of g and the Discrete Fourier transform (DFT) of its samples $g_n = g(n/N)$, $n = 0, \dots, N-1$, where $N \geq 4M+1$, the total number of parameters in (1.2).

The coefficients $\{a_k\}_{k \in \mathbb{Z}}$ of the Fourier series of g in (1.2) are readily available,

$$(2.1) \quad a_k = \int_0^1 g(x) e^{-2\pi i k x} dx = \sum_{m=1}^M w_m e^{-\eta_m k}, \quad a_{-k} = \bar{a}_k \quad k \in \mathbb{N},$$

and

$$a_0 = \int_0^1 g(x) dx,$$

which coincides with the constant term in (1.2), i.e., the proper rational terms in (1.2) do not contribute to a_0 . We emphasize that all Fourier coefficients are fully described by $4M+1$ real parameters whereas the frequency content of g is controlled by the distance of the nodes, $e^{-\eta_m}$, $m = 1, \dots, M$, from the unit circle, i.e., the real part of the exponents, $\text{Re}(\eta_m)$.

In order to compute the DFT of g_n , we substitute the representation (2.1) into the Fourier series of g . For any integer $N > 0$, summing the geometric series, we have

$$(2.2) \quad \begin{aligned} g(x) - a_0 &= \sum_{k \geq 1} (a_k e^{2\pi i k x} + \bar{a}_k e^{-2\pi i k x}) \\ &= 2\mathcal{R}e \sum_{n \geq 0} \sum_{j=0}^{N-1} \sum_{m=1}^M w_m e^{-\eta_m(j+Nn)} e^{2\pi i(j+Nn)x} - 2\mathcal{R}e \sum_{m=1}^M w_m \\ &= 2\mathcal{R}e \sum_{j=0}^{N-1} \sum_{m=1}^M \frac{w_m}{1 - e^{-\eta_m N} e^{2\pi i N x}} e^{-\eta_m j} e^{2\pi i j x} - 2\mathcal{R}e \sum_{m=1}^M w_m. \end{aligned}$$

Sampling $g(x)$, we obtain

$$(2.3) \quad g\left(\frac{n}{N}\right) = a_0 - 2\mathcal{R}e \sum_{m=1}^M w_m + 2\mathcal{R}e \sum_{j=0}^{N-1} \sum_{m=1}^M \frac{w_m}{1 - e^{-\eta_m N}} e^{-\eta_m j} e^{2\pi i j n/N},$$

where $n = 0, \dots, N - 1$. Computing the DFT of $g\left(\frac{n}{N}\right)$, $n = 0, \dots, N - 1$,

$$\hat{g}_j = \frac{1}{N} \sum_{n=0}^{N-1} g\left(\frac{n}{N}\right) e^{-2\pi i n j / N}, \quad j = 0, \dots, N - 1,$$

we obtain

$$(2.4) \quad \hat{g}_j = \sum_{m=1}^M \frac{w_m}{1 - e^{-\eta_m N}} e^{-\eta_m j} + \sum_{m=1}^M \frac{\bar{w}_m}{1 - e^{-\bar{\eta}_m N}} e^{-\bar{\eta}_m (N-j)}, \quad j = 1, \dots, N - 1,$$

and

$$\hat{g}_0 = a_0 - 2\mathcal{R}e \sum_{m=1}^M w_m + 2\mathcal{R}e \sum_{m=1}^M \frac{w_m}{1 - e^{-\eta_m N}} = a_0 + 2\mathcal{R}e \sum_{m=1}^M w_m \left(\frac{1}{1 - e^{-\eta_m N}} - 1 \right).$$

Note that the DFT coefficients and the Fourier coefficients (2.1) of g share the same nodes $e^{-\eta_m}$, $m = 1, \dots, M$. On the other hand, the DFT coefficients also contain (what we call) companion nodes $e^{\bar{\eta}_m}$, $m = 1, \dots, M$ which lie outside the unit disk. These companion nodes appear due to aliasing, i.e., folding of Fourier series coefficients corresponding to high frequencies onto low frequencies. We note that if the sampling rate is sufficient, then the DFT coefficients \hat{g}_j , $j = 1, \dots, N/2 - 1$ accurately approximate the Fourier series coefficients with the same indices and the companion nodes may be ignored. In our case, due to presence of singularities, we have to deal with the potential influence of the companion nodes. The algorithm for finding nodes given the DFT coefficients is described in Section 3.4 and the details of computing weights are described in Section 4.

In the presence of noise, recovering the $4M + 1$ real parameters in (1.2) requires appropriate oversampling of g . However, the required sampling rate is typically less than the rate dictated by the Nyquist criterion.

Finally, let us emphasize that the exponential representation of the Fourier coefficients in (2.1) defines the rational spatial representation (1.2) and vice versa [6]. We use this correspondence throughout the paper.

3. PRELIMINARY CONSIDERATIONS

3.1. Tomographic reconstruction problem. We consider the classical problem of X-ray tomography, the inversion of the Radon transform. Given a function u defined on the plane, the Radon transform is defined as integral over the lines $\{\mathbf{x} : s = \boldsymbol{\nu} \cdot \mathbf{x}\}$ parametrized by $s \in [-1, 1]$ and $\boldsymbol{\nu} = (\cos \theta, \sin \theta)$ a unit vector,

$$(3.1) \quad v(s, \boldsymbol{\nu}) = (Ru)(s, \boldsymbol{\nu}) = (R_\theta u)(s) = \int_{\mathbb{R}^2} u(\mathbf{x}) \delta(s - \boldsymbol{\nu} \cdot \mathbf{x}) d\mathbf{x}.$$

Note that the points $(s, \boldsymbol{\nu})$ and $(-s, -\boldsymbol{\nu})$ define the same line. The inversion of the Radon transform may be accomplished via the FBP algorithm which formally may be written in operator form as

$$I = R^* K R,$$

where the back-projection operator (the dual transform) is defined as

$$(R^* v)(\mathbf{x}) = \int_{\|\boldsymbol{\nu}\|=1} v(s, \boldsymbol{\nu}) |_{s=\boldsymbol{\nu} \cdot \mathbf{x}} d\boldsymbol{\nu}$$

and the operator K (applied to each projection separately)

$$(Kf)(s) = \frac{1}{2} \int_{\mathbb{R}} |\rho| \hat{f}(\rho) e^{2\pi i \rho s} d\rho.$$

Since multiplication by $|\rho|$ is an unbounded operator, in practice it is replaced by its band-limited approximation,

$$(Kf)(s) = \frac{1}{2} \int_{\mathbb{R}} |\rho| W(\rho) \hat{f}(\rho) e^{2\pi i \rho s} d\rho = \int_{\mathbb{R}} k(s-t) f(t) dt,$$

where

$$k(t) = \frac{1}{2} \int_{\mathbb{R}} |\rho| W(\rho) e^{-2\pi i \rho t} d\rho.$$

In FBP algorithms a typical choice for $|\rho| W(\rho)$ and, therefore $k(t)$, is the Shepp-Logan filter (see e.g., [28] for details). Since any filter modifies the measured data within its bandlimit, our goal is to first extend the bandlimit of the data so that when we apply the filter (see Algorithm 5.1), its impact on the measured data (within its original bandlimit) is reduced.

We also have from (3.1) the Fourier slice theorem,

$$\hat{u}(\rho \boldsymbol{\nu}) = \int_{\mathbb{R}^2} u(\mathbf{x}) e^{-2\pi i \rho \boldsymbol{\nu} \cdot \mathbf{x}} d\mathbf{x} = \int_{\mathbb{R}} (Ru)(s, \boldsymbol{\nu}) e^{-2\pi i \rho s} ds,$$

which we use to build a fast and accurate Fourier domain reconstruction algorithm as an alternative to FBP. In this paper we only examine the so-called parallel beam tomography, i.e. the most simple geometry for the tomographic reconstruction problem. However, the methods used in this paper are applicable to other, more complex geometries, such as fan beam tomography.

3.2. Quadratures for the disk. In order to have an accurate and fast reconstruction algorithm in the Fourier domain, we construct, for any user-specified accuracy $\epsilon > 0$, quadratures for integration (in polar coordinates) in the Fourier domain. Following [3, 4], we discretize integrals using quadratures for band-limited exponentials. In the radial variable these quadratures integrate exponentials against the weight $|\rho|$. In the angular variable we integrate using the trapezoidal rule.

Let us first examine the behavior of the Fourier transform of a real function supported within the box $B = [-1, 1] \times [-1, 1]$.

Proposition 3.1. *Consider a real function f supported in B such that, in polar coordinates, $\hat{f}(c\rho, \theta)$, $\theta \in [0, 2\pi]$, is negligible for $\rho > 1$, where $c > 0$ is the bandlimit. Extending the integration domain to a spatial disk of radius $\sqrt{2}$ containing B , we write*

$$\hat{f}(c\rho, \theta) = \int_0^{2\pi} \int_0^{\sqrt{2}} f(r, \varphi) e^{-icr\rho \cos(\theta-\varphi)} r dr d\varphi = \sum_{l \in \mathbb{Z}} q_l(c\rho) e^{il\theta},$$

where

$$q_l(c\rho) = \frac{1}{2\pi} \int_0^{2\pi} \hat{f}(c\rho, \theta) e^{il\theta} d\theta.$$

Then, for any $\epsilon > 0$, we have

$$(3.2) \quad \left| \hat{f}(c\rho, \theta) - \sum_{l=-L}^L q_l(c\rho) e^{i\theta l} \right| \leq \epsilon, \quad 0 \leq \rho \leq 1$$

where

$$(3.3) \quad L \geq \max \left\{ 7, e\sqrt{2}c, \log_2(2\|f\|_\infty \epsilon^{-1}) \right\}.$$

For the proof see Appendix 8. A similar observation was made in [27, 28] in the context of sampling for tomographic reconstruction algorithms.

Since the function f is real and $(\rho, \theta + \pi)$ and $(-\rho, \theta)$ describe the same point in a disk, we extend $\hat{f}(c\rho, \theta)$ for negative ρ as

$$\hat{f}(-c\rho, \theta) = \hat{f}(c\rho, \theta + \pi) = \overline{\hat{f}(c\rho, \theta)},$$

so that the coefficients $q_l(c\rho)$ in (3.2) for negative ρ are extended as

$$q_l(-c\rho) = \overline{q_{-l}(c\rho)}.$$

Under the assumptions of Proposition 3.1, the number of angular modes to represent \hat{f} is proportional to c , so that we can use polar grid quadratures in [3] to discretize the inverse Fourier transform of \hat{f} ,

$$(3.4) \quad \begin{aligned} f(\mathbf{x}) &\approx c^2 \int_{\{\mathbf{p} \mid \|\mathbf{p}\| \leq 1\}} \hat{f}(c\mathbf{p}) e^{ic\mathbf{p} \cdot \mathbf{x}} d\mathbf{p} \\ &= \frac{c^2}{2} \int_{-\pi}^{\pi} \int_{-1}^1 \hat{f}(c\rho, \theta) e^{ic\rho(x_1 \cos \theta + x_2 \sin \theta)} |\rho| d\rho d\theta, \end{aligned}$$

where $\mathbf{x} = (x_1, x_2)$, $\mathbf{x} \in B$. We obtain

$$(3.5) \quad \left| f(\mathbf{x}) - \frac{\pi c^2}{N_\theta} \sum_{\nu=-N_\rho}^{N_\rho} \sum_{k=0}^{N_\theta-1} w_\nu \hat{f}(c\rho_\nu, \theta_k) e^{ic\rho_\nu(x_1 \cos \theta_k + x_2 \sin \theta_k)} \right| \leq \epsilon,$$

where ρ_ν , $|\rho_\nu| < 1$, are (unequally spaced) quadrature nodes on the diameters, w_ν are corresponding quadrature weights, and $\theta_k = 2\pi k/N_\theta$. The choice of quadratures on the diameters (and the number of nodes) depends on the bandlimit c and the desired accuracy ϵ . The number of diameters, N_θ , is proportional to c . An example of such a grid is illustrated in Figure 3.1. The sum in (3.5) is evaluated on a $N_{image} \times N_{image}$ Cartesian grid using USFFT at the cost of $\mathcal{O}(c^2 \log c)$ since N_{image} , N_ρ and N_θ are all proportional to c .

Remark 3.2. The evaluation of (3.4) follows the approach in [3] which provides a method to obtain quadratures for functions essentially supported on a disk in the Fourier domain. On the diameters we use quadratures for band-limited exponentials for the weight $|\rho|$ [4, 3]. Unlike quadratures for polynomials, the nodes of these quadratures do not concentrate excessively towards the end points of diameters. As a result, even for large bandlimits, we avoid clustering of nodes near such points [8]. We also note that using (3.5) in the new algorithm described in Section 5, we apply the filter (5.3) to enforce the bandlimit.




FIGURE 3.1. A polar grid for integration in the Fourier domain where the number of nodes per diameter is 74 and the number of diameters is 37. This grid (along with appropriate weights) is designed to yield an accuracy of $1.68 \cdot 10^{-7}$ for computing the inverse Fourier transform of functions with bandlimit 31.75π (i.e., supported within the disk).

3.3. Space-limited and band-limited functions in tomography. In a tomography setup the objects of interest are compactly supported and, at the same time, only band-limited measurements are available to solve the inverse problem. The natural class of functions that captures this setup are the eigenfunctions of space-limiting and band-limiting operators defined on appropriate domains. Such integral operators and their analysis were first introduced by Slepian et. al. in a series of papers [35, 24, 25, 31, 32, 33, 34], where they observed that the differential operator for the Prolate Spheroidal Wave Functions commutes with the space-limiting and band-limiting operator and, thus, provides a way to compute its eigenfunctions. Our approach to identifying the class of functions for tomography problems is based on [3], where the Slepian operator mapping a square in space to a disk in the Fourier domain is considered. This differs from the original construction in [31]

of the disk-to-disk mapping since, *inter alia*, for the square-to-disk case no ordinary differential equation is available to compute the eigenfunctions of the integral operator.

We believe that in the tomography setup, the eigenfunctions of the square-to-disk Slepian operator provide a linear space that naturally describes both the measured data and the reconstructed image. The well-localized nature of tomographic projections allow us to approximate them using only the eigenfunctions corresponding to eigenvalues close to 1 since these eigenfunctions are well-localized in the spatial box (slightly away from the boundary) while their Fourier transforms are supported in the disk (see [3, Sections 2.2-2.3] for details). In practice, this property leads to quadratures with controlled accuracy for functions band-limited within a disk (as described in the previous section) and, also, allows us to extend the Fourier transform of the projections without modifying significantly their spatial support (see Section 5).

3.4. Methods of approximation via exponentials. For each real-valued, sampled projection, we compute its DFT and approximate positive frequencies via a linear combination of exponentials with a (near) minimal number of terms for a user-selected accuracy. For this purpose, we use the algorithm in [5, 7], whose key steps are presented below in Algorithm 3.3. Given a user-selected accuracy $\epsilon > 0$ and $2L + 1$ equally spaced samples of a complex-valued function $h(c\xi)$, $\xi \in [0, 1]$, this algorithm yields an approximation

$$(3.6) \quad \left| h\left(c\frac{l}{2L}\right) - \sum_{m=1}^M w_m e^{-c\eta_m l} \right| < \epsilon, \quad \mathcal{R}e(\eta_m) > 0,$$

for $0 \leq l \leq 2L$, where the number of terms, M , is near minimal for the choice of c and ϵ . In this formulation, the constant $c > 0$ scales the problem to the interval $[0, 1]$.

Provided that $h(c\xi)$ is sufficiently sampled, by defining $\tau_m = 2L\eta_m$, and replacing $l/(2L)$ with a continuous variable ξ in (3.6), we obtain

$$(3.7) \quad \left| h(c\xi) - \sum_{m=1}^M w_m e^{-c\tau_m \xi} \right| < \epsilon',$$

for $\xi \in [0, 1]$. The new error ϵ' is only slightly worse than ϵ .

Algorithm 3.3. [5, 7]

The algorithm to produce the approximation (3.6) has the following steps:

- Construct the $(L + 1) \times (L + 1)$ Hankel matrix $\mathbf{H}_{ll'}$ = $h_{l+l'}$, where $h_{l+l'} = h\left(c\frac{l+l'}{2L}\right)$, $0 \leq l, l' \leq L$.
- Find a vector \mathbf{u} satisfying $\mathbf{H}\mathbf{u} = \sigma\bar{\mathbf{u}}$ with positive σ close to ϵ . A solution is guaranteed by Tagaki's factorization [19] and may be reduced to finding the Singular Value Decomposition (SVD) of \mathbf{H} . Given singular values $\sigma_0 \geq \sigma_1 \geq \dots \geq \sigma_M \geq \dots \geq \sigma_L$, we choose M such that $\epsilon \approx \sigma_M/\sigma_0$ and the corresponding singular vector $\mathbf{u} = \{u_l\}_{l=0}^L$.
- Compute the roots γ_m of the polynomial

$$(3.8) \quad u(z) = \sum_{l=0}^L u_l z^l.$$

- The exponents η_m in equation (3.6) are defined by the roots γ_m via $\eta_m = -\log(\gamma_m)/c$, where \log is the principal value of the logarithm. We also compute $\tau_m = 2L\eta_m$.

The usual last step of Algorithm 3.3 is to form a linear Vandermonde system

$$(3.9) \quad h\left(\frac{cl}{2L}\right) = \sum_{m=1}^M w_m \gamma_m^l$$

and solve for the weights w_m as the best fit in the ℓ^2 sense. However, due to the nature of problems in tomography, we depart from this conventional approach. In the next section we provide a new method for calculating weights w_m using the ℓ^1 norm.

Remark 3.4. An important feature of Algorithm 3.3 is that it detects the level of noise in the data. Specifically, the rate of decay of the singular values of the matrix \mathbf{H} changes significantly once the ratio σ_m/σ_0 reaches the level of noise in the signal. Indeed, unlike a coherent signal, noise does not have an efficient representation via exponentials. For this reason, the gain in accuracy of fitting the data becomes negligible as we add additional terms. This directly affects the rate of decay of singular values, which we use as a tool for the selection of the singular value, σ_M , and hence, the number of terms M in the representation. For more details see [5, 6, 7].

Remark 3.5. We want to point out alternatives to Algorithm 3.3 for finding (near) optimal approximations via linear combinations of exponentials. For this purpose, we may use the HSVD [23] or the (equivalent) matrix pencil method [20, 21, 22]. The advantage of such methods is that we avoid the step of finding roots of the polynomial (3.8). On the other hand, Algorithm 3.3 is faster (if properly implemented) since we need to find only a single singular vector of a structured (Hankel) matrix. The alternative methods mentioned above require computing all singular vectors up to the index corresponding to that of the accuracy threshold.

4. CALCULATING WEIGHTS FOR TOMOGRAPHY PROBLEMS

The results of Section 2 show that in order to represent a projection g using a rational signal model it is enough to approximate its DFT values \hat{g}_j , $1 \leq j \leq N/2-1$ as a sum of exponentials. This approximation is achieved via Algorithm 3.3. Since we would like to extrapolate the Fourier data beyond the original bandlimit, we have to ensure that all exponents in the representation of \hat{g}_j have negative real parts. Unfortunately, as mentioned at the end of Section 2, the exponential representation of the DFT coefficients (2.4) differs from that of the Fourier coefficients (2.1) in that it has a second sum,

$$(4.1) \quad \sum_{m=1}^M \frac{\bar{w}_m}{1 - e^{-\bar{\eta}_m N}} e^{-\bar{\eta}_m(N-j)}, \quad \mathcal{R}e(\eta_m) > 0.$$

which typically has a small contribution within the range $1 \leq j \leq N/2-1$, but, as a function of j , it contains growing terms $e^{\eta_m j}$. Hence, when using Algorithm 3.3 on \hat{g}_j , some roots γ_m may occasionally end up just outside the unit disk. Since such roots (typically with small weights) would prevent us from extrapolating the

Fourier data, we discard them at the expense of introducing an additional error. The following approach aims to isolate such errors to the vicinity of singularities, which are responsible for the slow decay of the Fourier data and the mismatch between DFT and Fourier series coefficients.

The proximity of a node to the unit circle controls the frequency contribution of that node while the position of the corresponding pole (of the rational approximation) is directly associated with the location of a singularity [6]. Therefore, the impact of removing nodes just outside the unit disk should be localized to neighborhoods of singularities. However, when we remove nodes from just outside the unit circle and use the ℓ^2 norm to calculate the weights via (3.9), the error spreads out to a large neighborhood of the singularity as illustrated in Figures 4.1 (a) and (c), and Figure 4.2 (a). In order to remedy this situation, we calculate weights by minimizing the ℓ^1 norm of the the residual with respect to the original spatial data. We choose to minimize the ℓ^1 norm due to its well known sparsity properties (see, e.g., [30]). The effect of using the ℓ^1 norm to calculate weights is illustrated in Figures 4.1 (b) and (d), and Figure 4.2 (b).

Let us describe the details. Once we obtain the exponents η_m via Algorithm 3.3 and select all those with $\Re(\eta_m) > 0$, we proceed to compute weights in the space domain. Specifically, using the fact that the function is real we extend the approximation of positive frequencies to negative frequencies and analytically sum the Fourier series to obtain a rational function of the form (1.2). Denoting the grid at which we measure the projection data as $\{x_j\}_{j=0}^{N-1}$, we discretize (1.2) and obtain a Cauchy-like system to solve for the weights w_m .

$$(4.2) \quad g_n = g\left(\frac{n}{N}\right) = \widehat{g}_0 + 2\Re \sum_{m=1}^M w_m \left(\frac{1}{e^{-2\pi i n/N + \eta_m} - 1} - \frac{1}{1 - e^{-\eta_m N}} + 1 \right),$$

where $n = 0, \dots, N - 1$. Let us denote the $N \times 2M$ matrix of this system as \mathbf{C} and the right hand side $\{g_n - \widehat{g}_0\}_{n=0}^{N-1}$ as \mathbf{g} . We note that the contribution of terms $1/(1 - e^{-\eta_m N}) - 1$ is minor since $e^{-\eta_m N}$ is typically small.

We first solve

$$(4.3) \quad \mathbf{C}\mathbf{w} = \mathbf{g}, \quad \text{with } \operatorname{argmin}_{\mathbf{w}} \|\mathbf{C}\mathbf{w} - \mathbf{g}\|_2,$$

where \mathbf{g} is the vector containing the (shifted) data samples. We then verify if the residual is within the error tolerance. If it is not (which is the case for only a few projections), we then solve

$$(4.4) \quad \mathbf{C}\mathbf{w} = \mathbf{g}, \quad \text{with } \operatorname{argmin}_{\mathbf{w}} \|\mathbf{C}\mathbf{w} - \mathbf{g}\|_1.$$

In this approach we rely on an observation that if solving (4.3) satisfies the error tolerance, then the difference in using ℓ^1 or ℓ^2 norms is insignificant.

For solving (4.3) we use the SVD while for (4.4) we use convex optimization. For the latter, since \mathbf{C} is complex valued, we solve via a second order cone program implemented in CVX [15]. An alternative approach for solving (4.4) relies on the iteratively re-weighted least squares algorithm (see, e.g. [30]).

We illustrate the difference between using the ℓ^2 norm vs the ℓ^1 norm in an example. In Figure 4.1 (a), using the ℓ^2 norm, we show a clearly visible artifact near one of the sharp transitions in the projection. This should be compared with Figure 4.1 (b) where, using (4.4), the error is significantly reduced away from the sharp transition. The comparison of the errors for the entire projection is presented in Figures 4.2 (a) and (b).




FIGURE 4.1. Rational approximation of a projection in the vicinity of a sharp transition. The result of using weights obtained by ℓ^2 minimization of the residual (a) and ℓ^1 minimization of the residual (b). The corresponding errors are illustrated in (c) and (d). We observe a significant improvement in error localization using the ℓ^1 norm.

Remark 4.1. If we were to form the $N/2 \times N/2$ Hankel matrix \mathbf{H} using all \hat{g}_j , $j = 1, \dots, N-1$ (for even N) and, thus, insist on the form (2.4), then the resulting matrix is equivalent to a self-adjoint Toeplitz matrix. This may be seen by multiplying \mathbf{H} by the matrix \mathbf{J} with entries $\{\delta_{i, N/2-j-1}\}_{i,j=0, \dots, N/2-1}$. This leads to a construction similar to that in [4] and, in this case the roots (except those on the unit circle) indeed come in pairs, $\gamma_m^{out} = 1/\overline{\gamma_m^{in}}$. However, in this case some nodes may lie on the unit circle which creates problems with the extension. An approach to deal with such nodes may provide an alternative to the method of this paper.

5. POLAR QUADRATURE INVERSION

Current numerical implementations of Fourier methods for inverting the Radon transform, the so-called gridding methods, use interpolation in the Fourier domain to generate values on a rectangular (typically square) equally spaced grid so that inversion can then proceed via the FFT [28, Section 5.2]. In contrast, our new algorithm uses a carefully designed polar grid for inversion (see Section 3.2) which is a natural choice within the setup of the projection slice theorem. This grid allows us to accurately discretize the Fourier domain integral (3.4) to represent the image via the sum in (3.5). We evaluate this sum via USFFT [11, 2, 26]. Note that the only inputs into this algorithm are the values of the Fourier transform at the grid




FIGURE 4.2. Errors in rational approximation of the entire projection in Figure 4.1 using weights obtained by ℓ^2 minimization of the residual (a) and ℓ^1 minimization of the residual (b). We observe a consistent improvement of the error away from sharp transitions.

points. Finding these values is fairly straightforward due to the functional form of our approximation of projections.

We denote projections as $g_n(k)$, $n = 0, \dots, N - 1$ and $k = 0, \dots, N_\theta - 1$, where the index k identifies a projection corresponding to the angle $\theta_k = \pi k / N_\theta$, and the index n identifies a location within a projection. For each projection k we compute the DFT,

$$\hat{g}_j(k) = \frac{1}{N} \sum_{n=0}^{N-1} g_n(k) e^{-2\pi i j n / N},$$

yielding both positive and negative discrete frequencies. For simplicity, we assume that the number of samples N in each projection is a multiple of 4. For each projection k , we use Algorithm 3.3 to construct a (near) optimal approximation of the values $\hat{g}_j(k)$, $j = 1, \dots, N/2 - 1$,

$$(5.1) \quad \left| \hat{g}_j(k) - \sum_{m=1}^{M_k} w_{mk} e^{-c\eta_{mk}j} \right| \leq \epsilon, \quad j = 1, \dots, N/2 - 1, \quad \mathcal{R}e(\eta_{mk}) > 0,$$

where $\epsilon > 0$ is the user-supplied accuracy and the weights w_{mk} are calculated using either (4.3) or (4.4), whichever is appropriate (see considerations in Section 4). Based on (5.1), we then define a continuous representation of the Fourier transform of the projection,

$$(5.2) \quad \hat{g}^{\text{ext}}(\rho, k) = \sum_{m=1}^M w_{mk} e^{-c\tau_{mk}\rho}, \quad \rho \geq 0, \quad \mathcal{R}e(\tau_{mk}) > 0,$$

where $\tau_{mk} = (N/2) \eta_{mk}$ (so that the Nyquist frequency of the sequence $\{g_n(k)\}_{n=0}^{N-1}$ corresponds to $\rho = 1$). While (5.2) only matches the measured data for $\rho \in [0, 1]$, we use its functional form to extend its definition to values $\rho > 1$. We restrict the extension to a finite radius since the rate of decay of \hat{g}^{ext} is faster than that of the Fourier transform of a function with, for example, jump discontinuities. We choose to restrict (5.2) to $\rho \in [0, 2]$, effectively defining the new bandlimit $c_{\text{new}} = 2c$. Additionally, to ensure that the resulting extension is supported within the disk of radius c_{new} , we apply the radial Hann window $W(\rho)$,

$$(5.3) \quad W(\rho) = \sin^2\left(\frac{\pi}{4}\rho - \frac{\pi}{2}\right), \quad \rho \in [0, 2].$$

In comparison, within the FBP algorithm, the bandlimit c (corresponding to the measured data) is enforced by e.g., the Shepp-Logan filter. Since we first extend to the larger bandlimit c_{new} , and then apply the filter (5.3), the impact of the filter on the measured data of bandlimit c is significantly reduced, since the filter mostly modifies the extended part of \widehat{g}^{ext} .

Therefore, we obtain

$$(5.4) \quad \widehat{g}^{bl}(\rho, k) = \widehat{g}^{ext}(\rho, k) W(\rho),$$

where the window is specified in (5.3). Since the image is a real valued function, we define

$$(5.5) \quad \widehat{g}^{bl}(-\rho, k) = \overline{\widehat{g}^{bl}(\rho, k)}$$

for each projection, $k = 0, \dots, N_\theta - 1$ and then sample $\widehat{g}^{bl}(\rho, k)$ on the diameter using quadrature nodes $2\rho_\nu$, $\nu = -N_\rho, \dots, N_\rho$ described in Section (3.2) (we may also choose a quadrature with an even number of nodes on the diameters).

The final step of the algorithm uses the USFFT [11, 2, 26] with input values $\widehat{g}^{bl}(2\rho_\nu, k)$, $\nu = -N_\rho, \dots, N_\rho$ and $k = 0, \dots, N_\theta - 1$ to produce an image. This step (effectively) delegates interpolation in the Fourier domain to the USFFT (which attains any finite user-specified accuracy). We now summarize the Polar Quadrature Inversion (PQI) algorithm,

Algorithm 5.1.

- (1) Compute the FFT of the projection data $\{g_n(k)\}_{n=0}^{N-1}$ for each $k = 0, \dots, N_\theta - 1$.
- (2) Using DFT coefficients $\widehat{g}_j(k)$, $j = 1, \dots, N/2 - 1$ of each projection k , construct its exponential approximation via Algorithm 3.3. The accuracy $\epsilon > 0$ may be estimated by examining the change in the rate of decay of the singular values of the Hankel matrix.
- (3) Find the weights using either (4.3) or (4.4).
- (4) Evaluate the windowed approximations $\widehat{g}^{bl}(\rho, k)$ in (5.4-5.5) at the quadratures nodes $2\rho_\nu$, $\nu = -N_\rho, \dots, N_\rho$, corresponding to the new bandlimit c_{new} for each $k = 0, \dots, N_\theta - 1$. Note that this step includes extrapolation.
- (5) Obtain the image by computing the sum in (3.5) via the USFFT using as function values those computed in the previous step.

Operation count. Step 1 of the algorithm requires $\mathcal{O}(N_\theta N \log N)$ operations. Using Algorithm 3.3 for all projections in Step 2 currently requires $\mathcal{O}(N_\theta N^3)$ but may be reduced to $\mathcal{O}(N_\theta N M^2)$, where M is the number of terms in the exponential approximation (this may be achieved using randomized projections, see the review article [18]). We also note that Step 2 is trivially parallelizable as each projection is treated separately. Step 3 currently requires $\mathcal{O}(N_\theta N M^2)$ when using (4.3) and a similar complexity to solve (4.4) but with a larger constant. However, only a few projections require using (4.4). The remaining steps require $\mathcal{O}(N^2 \log N)$ operations.

6. NUMERICAL EXAMPLES

We now analyze the impact on image resolution of the rational approximation of projections and compare the new PQI algorithm with the FBP algorithm. Following tradition, we demonstrate our reconstruction technique using the Shepp-Logan phantom. To examine the effect of steps (1)-(3) of the PQI algorithm, we apply the standard FBP algorithm to both the original projection data and the augmented projections, where we double the number of samples via rational representations. This choice corresponds to doubling the bandlimit in the Fourier domain as in the PQI algorithm and allows us to isolate the impact of the rational approximation on image formation. Namely, using FBP on the original and augmented projections, we observe the impact of rational approximation; by using augmented projections in FBP we compare its output with that of the PQI algorithm.

The test data consist of 512 projections and 512 samples per projection resulting (via the standard FBP algorithm supplied in MatlabTM) in a 512×512 reconstructed image while the output of PQI is a 1024×1024 reconstructed image. In all experiments we use the radial Hann window (5.3). We perform two types of experiments: noiseless reconstruction and reconstruction after adding Gaussian noise to each projection.

The choice of the number of projections relative to the number of samples follows that in an example in [29] and may not be ideal (see [27, 28]). However, since our goal is to show no additional artifacts near singularities are introduced by our method (while achieving a higher resolution), this choice of parameters is sufficient for the demonstration.

6.1. Noiseless examples. First, we take our original noiseless data set of 512 projections with 512 samples per projection and approximate all of the projections by rational functions. For our experiments in a noiseless setting we use the threshold of $\sigma_M/\sigma_0 \leq 5 * 10^{-4}$ to find rational approximations (see Section 3.4). Using these approximations, we increase the sampling of each projection by a factor of 2. We then take these augmented projections, with 512 projections with 1024 samples per projection, and apply the standard FBP algorithm yielding a 1024×1024 reconstructed image of the Shepp-Logan phantom. The reconstruction from augmented projections is displayed in Figure 6.2 (a) and the reconstruction error is displayed in Figure 6.2 (b). Comparing the output of FBP applied to augmented samples with the output of FBP applied to our original data set (512 projections and 512 samples per projection), see Figure 6.1, we observe better resolution near the sharp boundaries when augmented projections are used as input. In both cases the error highlights the jump discontinuities of the phantom and contains streak artifacts typical in tomographic reconstructions where the angular variable is under-sampled.

We compare the results of this experiment (using augmented projections with 1024 samples as input for FBP) with the output of the PQI algorithm described in Section 5 yielding a 1024×1024 reconstructed image of the Shepp-Logan phantom. Using this algorithm we double the size of the bandlimit in the Fourier domain which corresponds to the increase of sampling in each projection by a factor of two in using FBP. The result of the reconstruction using the PQI algorithm is shown in Figure 6.3 (a) and the error in Figure 6.3 (b). When compared with Figure 6.1 we

observe that the PQI algorithm achieves a significantly improved resolution without introducing additional artifacts.

This also allows us to compare the output of our algorithm with that of the FBP algorithm on augmented projections. We see no significant difference in the reconstructions Figure 6.3 (a) and Figure 6.2 (a). A comparison of the reconstruction errors, Figure 6.3 (b) and Figure 6.2 (b) also shows no significant difference.

6.2. Zooming on details. In order to demonstrate the increased resolution, we zoom in on two areas of the reconstruction. The locations of these areas are shown in Figure 6.3 (a). In Figures 6.4 and 6.5 we compare images near the center of the phantom using the FBP and the PQI algorithms, respectively. Figure 6.4 compares a 32×32 pixel patch (obtained from the original data set of 512 projections with 512 samples each and extracted from 512×512 pixel reconstruction via FBP) and a 64×64 pixel patch extracted from 1024×1024 reconstruction using the FBP on augmented data. Similarly, in Figure 6.5, we compare the same 32×32 pixel patch with a 64×64 patch extracted from 1024×1024 reconstructed image using the PQI algorithm. A similar comparison is shown in Figures 6.6 and 6.7 but zooming into a different section of the Shepp-Logan phantom.

We observe that higher resolution is not accompanied by any additional artifacts and results in visibly sharper images.

6.3. Noisy examples. To test the stability of approximation by rational functions in the presence of noise, we add (to each projection) Gaussian white noise with zero mean and standard deviation of $2.5 * 10^{-4}$. The noise level is of the same order as the smallest features captured by the projections. In all experiments involving noise we used the threshold $\sigma_M/\sigma_0 \leq 2 * 10^{-3}$ to find rational approximations (see Remark 3.4). We construct rational approximations for each noisy projection and increase the sampling rate by a factor of two. We then use the resulting augmented projections as input to the FBP algorithm. The reconstruction error of this experiment is displayed in Figure 6.8 (b) and compared with the reconstruction error of the same experiment performed on noiseless data displayed in Figure 6.8 (a). The Gaussian noise added to the projections appears as a speckle noise in the image in addition to the error of noiseless reconstruction.

Similarly, we display the reconstruction error of the PQI algorithm applied to noisy data in Figure 6.9 (b) next to the error of reconstruction from the noiseless data in Figure 6.9 (a). The qualitative difference between the errors in Figure 6.9 (b) and Figure 6.8 (b) is negligible.

7. DISCUSSION

We present an approach to improve the resolution of X-ray tomography using rational approximation of projections and the new, fast Fourier algorithm (PQI) for this purpose. We also provide a comparison between the PQI algorithm and the standard FBP.

While we use rational approximation to improve the resolution in the classical X-ray tomography setup, a similar approach may be applied in Electron Microscope Tomography and in MRI. However, in both of these modalities of non-destructive evaluation one needs to consider additional issues associated with these techniques. In particular, in Electron Microscope Tomography, one has to address a missing cone of data as well as very high noise level when compared with X-ray tomography.

In MRI the collected data needs to be carefully re-sampled before using the rational model. An interesting potential application in MRI is to use the PQI algorithm to reduce the sampling requirements in the Fourier domain. We plan to address these issues in our future work.

Finally, we note that by modifying the input into the PQI algorithm, it is possible to produce the same image (within a user-supplied accuracy) as that constructed by an Algebraic Reconstruction Technique (ART), i.e., the PQI algorithm can reproduce the ART solution.

8. APPENDIX: PROOF OF PROPOSITION 3.1

Given the assumption that $\hat{f}(c\rho, \theta)$ is negligible for $\rho > 1$, it is enough to show that for $0 \leq \rho \leq 1$ the tail of the Fourier series

$$\hat{f}(c\rho, \theta) = \sum_{l \in \mathbb{Z}} q_l(c\rho) e^{i\theta l},$$

is also negligible

$$\left| \sum_{|l| > L} q_l(c\rho) e^{i\theta l} \right| < \epsilon.$$

Computing

$$\begin{aligned} q_l(c\rho) &= \frac{1}{2\pi} \int_0^{2\pi} \hat{f}(c\rho, \theta) e^{i\theta l} d\theta \\ &= \int_0^{2\pi} \int_0^{\sqrt{2}} f(r, \varphi) \left(\frac{1}{2\pi} \int_0^{2\pi} e^{-icr\rho \cos(\theta-\varphi)} e^{i\theta l} d\theta \right) r dr d\varphi, \end{aligned}$$

we obtain

$$q_l(c\rho) = (-i)^l \int_0^{2\pi} \int_0^{\sqrt{2}} f(r, \varphi) J_l(c\rho) e^{i\varphi l} r dr d\varphi, \quad l \in \mathbb{Z},$$

and, hence,

$$|q_l(c\rho)| \leq 2\pi \|f\|_\infty \max_{x \in [0, \sqrt{2}c]} |J_l(x)|.$$

Using [1, 9.1.62] and [1, 6.1.38], for $x > 0$ and $l \geq 7$, we have

$$|J_l(x)| \leq \frac{1}{l!} \left(\frac{x}{2} \right)^l \leq \frac{1}{2\pi} \left(\frac{ex}{2l} \right)^l.$$

With $l \geq e\sqrt{2}c$, we obtain

$$|q_l(c\rho)| \leq \|f\|_\infty 2^{-l},$$

and arrive at

$$\left| \sum_{|l| > L} q_l(c\rho) e^{i\theta l} \right| \leq 2\|f\|_\infty \sum_{l=L+1}^{\infty} 2^{-l} = \|f\|_\infty 2^{-L+1} < \epsilon$$

for some L ,

$$L \geq \max \left\{ 7, e\sqrt{2}c, \log_2(2\|f\|_\infty \epsilon^{-1}) \right\}.$$

REFERENCES

- [1] M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions*. Dover Publications, 9 edition, 1970.
- [2] G. Beylkin. On the fast Fourier transform of functions with singularities. *Appl. Comput. Harmon. Anal.*, 2(4):363–381, 1995.
- [3] G. Beylkin, C. Kurcz, and L. Monzón. Grids and transforms for band-limited functions in a disk. *Inverse Problems*, 23(5):2059–2088, 2007.
- [4] G. Beylkin and L. Monzón. On generalized Gaussian quadratures for exponentials and their applications. *Appl. Comput. Harmon. Anal.*, 12(3):332–373, 2002.
- [5] G. Beylkin and L. Monzón. On approximation of functions by exponential sums. *Appl. Comput. Harmon. Anal.*, 19(1):17–48, 2005.
- [6] G. Beylkin and L. Monzón. Nonlinear inversion of a band-limited Fourier transform. *Appl. Comput. Harmon. Anal.*, 27(3):351–366, 2009.
- [7] G. Beylkin and L. Monzón. Approximation of functions by exponential sums revisited. *Appl. Comput. Harmon. Anal.*, 28(2):131–149, 2010.
- [8] G. Beylkin and K. Sandberg. Wave propagation using bases for bandlimited functions. *Wave Motion*, 41(3):263–291, 2005.
- [9] E. J. Candès and L. Demanet. The curvelet representation of wave propagators is optimally sparse. *Comm. Pure Appl. Math.*, 58(11):1472–1528, 2005.
- [10] E. J. Candès and D. L. Donoho. New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities. *Comm. Pure Appl. Math.*, 57(2):219–266, 2004.
- [11] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Comput.*, 14(6):1368–1393, 1993.
- [12] S. Engelberg and E. Tadmor. Recovery of edges from spectral data with noise—a new perspective. *SIAM J. Numer. Anal.*, 46(5):2620–2635, 2008.
- [13] A. Gelb and D. Cates. Segmentation of images from Fourier spectral data. *Commun. Comput. Phys.*, 5(2-4):326–349, 2009.
- [14] A. Gelb and E. Tadmor. Detection of edges in spectral data. *Appl. Comput. Harmon. Anal.*, 7(1):101–135, 1999.
- [15] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, February 2011.
- [16] L. Greengard and C. Stucchio. Spectral edge detection in two dimensions using wavefronts. *Appl. Comput. Harmon. Anal.*, 30(1):69–95, 2011.
- [17] K. Guo, D. Labate, and W. Lim. Edge analysis and identification using the continuous shearlet transform. *Appl. Comput. Harmon. Anal.*, 27(1):24–46, 2009.
- [18] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [19] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1990.
- [20] Y. Hua and T.K. Sarkar. Matrix pencil method and its performance. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, 1988.
- [21] Y. Hua and T.K. Sarkar. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(5):814–824, 1990.
- [22] Y. Hua and T.K. Sarkar. On SVD for estimating generalized eigenvalues of singular matrix pencil in noise. *IEEE Transactions on Signal Processing*, 39(4):892–900, 1991.
- [23] S.Y. Kung, K.S. Arun, and D.V. Bhaskar Rao. State-space and singular-value decomposition-based approximation methods for the harmonic retrieval problem. *Journal of the Optical Society of America*, 73(12):1799–1811, 1983.
- [24] H. J. Landau and H. O. Pollak. Prolate spheroidal wave functions, Fourier analysis and uncertainty II. *Bell System Tech. J.*, 40:65–84, 1961.
- [25] H. J. Landau and H. O. Pollak. Prolate spheroidal wave functions, Fourier analysis and uncertainty III. *Bell System Tech. J.*, 41:1295–1336, 1962.
- [26] J.-Y. Lee and L. Greengard. The type 3 nonuniform FFT and its applications. *J. Comput. Phys.*, 206(1):1–5, 2005.
- [27] F. Natterer. *The Mathematics of Computerized Tomography*. Wiley, NY, 1986.

- [28] F. Natterer and F. Wübbeling. *Mathematical Methods in Image Reconstruction*. SIAM, Philadelphia, PA, 2001.
- [29] X. Pan, E. Y. Sidky, and M. Vannier. Why do commercial CT scanner still employ traditional, filtered back-projection for image reconstruction? *Inverse Problems*, 25(12):123009, 36, 2009.
- [30] J. A. Scales and A. Gersztenkorn. Robust methods in inverse theory. *Inverse Problems*, 4(4):1071–1091, 1988.
- [31] D. Slepian. Prolate spheroidal wave functions, Fourier analysis and uncertainty IV. Extensions to many dimensions; generalized prolate spheroidal functions. *Bell System Tech. J.*, 43:3009–3057, 1964.
- [32] D. Slepian. Some asymptotic expansions for prolate spheroidal wave functions. *J. Math. and Phys.*, 44:99–140, 1965.
- [33] D. Slepian. Prolate spheroidal wave functions, Fourier analysis and uncertainty V. The discrete case. *Bell System Tech. J.*, 57:1371–1430, 1978.
- [34] D. Slepian. Some comments on Fourier analysis, uncertainty and modeling. *SIAM Review*, 25(3):379–393, 1983.
- [35] D. Slepian and H. O. Pollak. Prolate spheroidal wave functions, Fourier analysis and uncertainty I. *Bell System Tech. J.*, 40:43–63, 1961.
- [36] M. Unser. Sampling - 50 Years After Shannon. *Proc. IEEE*, 88(4):569–587, 2000.




FIGURE 6.1. A 512×512 reconstruction of the Shepp-Logan phantom using FBP algorithm (a) (the gray scale is $[-0.05, 1.05]$) and the corresponding error (b) (the gray scale is $[-.05, .05]$). A combination of filtering errors and under-sampling in angle produces streak artifacts associated with sharp transitions in the phantom.



FIGURE 6.2. A 1024×1024 reconstructed image of the Shepp-Logan phantom via the FBP algorithm using projections (with twice as many samples) generated by near optimal rational approximation (a) and the corresponding error (b). Gray scales are the same as in Figure 6.1 which should be used for comparison.

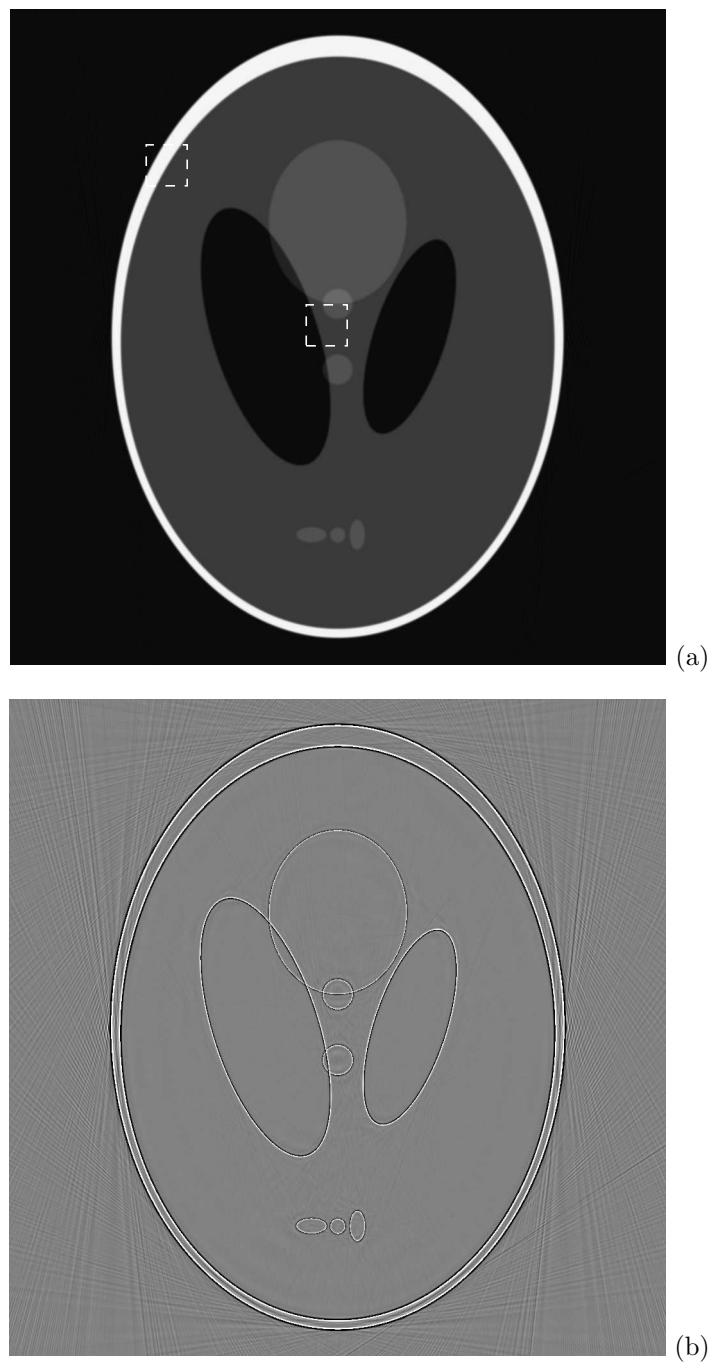


FIGURE 6.3. A 1024×1024 reconstructed image via the PQI algorithm of Section 5 (a) and the corresponding error (b). The gray scales are the same as in Figures 6.1 and 6.2, which should be used for comparison. The two boxes in (a) outline areas of the reconstructed image on which we zoom to examine the reconstruction at a pixel level.




FIGURE 6.4. A zoom of a small area near the center of the phantom indicated in Figure 6.3. In (a) we show a 32×32 square (out of a 512×512 reconstruction via the FBP algorithm) and in (b) the same area with 64×64 samples taken from a 1024×1024 reconstruction via the FBP algorithm using as input augmented projections.




FIGURE 6.5. A different comparison for the area in Figure 6.4. In (a) we show a 32×32 square (out of a 512×512 via the FBP algorithm) and in (b) the same area with 64×64 samples taken from a 1024×1024 reconstruction via the PQI algorithm.




FIGURE 6.6. A zoom of a small area near the sharpest transition in the phantom indicated in Figure 6.3. In (a) we show a 32×32 square (out of a 512×512 reconstruction via the FBP algorithm) and in (b) the same area with 64×64 samples taken from a 1024×1024 reconstruction via the FBP algorithm using as input augmented projections.




FIGURE 6.7. A different comparison for the area in Figure 6.6. In (a) we show a 32×32 square (out of a 512×512 via the FBP algorithm) and in (b) the same area with 64×64 samples taken from a 1024×1024 reconstruction via the PQI algorithm.




FIGURE 6.8. Comparison of errors of 1024×1024 reconstructions in Figure 6.2 (via the standard FBP applied to noiseless data) (a) and the same data with added Gaussian noise (b). We observe that the Gaussian noise creates a speckle component in the error.




FIGURE 6.9. The same comparison as in Figure 6.8 but using the PQI algorithm of Section 5. The effect of introducing Gaussian noise is qualitatively and quantitatively the same as that illustrated in Figure 6.8.