

SPECTRAL SAMPLING AND DISCONTINUITY DETECTION METHODS WITH APPLICATION TO
MAGNETIC RESONANCE IMAGING

by

Adityavikram Viswanathan

A Thesis Presented in Partial Fulfillment
of the Requirements for the Degree
Master of Science

ARIZONA STATE UNIVERSITY

August 2008

ABSTRACT

Typical implementations of Fourier methods demand acquisition and evaluation of input and output data on a Cartesian grid. This presents a problem in applications such as magnetic resonance imaging (MRI), where non-Cartesian input measurements are available. The reconstruction of support-limited functions from such non-Cartesian spectral data constitutes the subject matter of this thesis.

A survey of prevalent reconstruction methods is provided, along with a comparison of their computational cost, implementation steps and reconstruction quality. Several illustrative examples and sampling schemes are considered, with emphasis toward those encountered in MRI. Additionally, analysis of one such scheme - uniform re-sampling (URS), which relates non-Cartesian measurements to Cartesian spectral samples is provided. It is shown that for sampling schemes typically encountered in MRI, high frequency coefficients are recovered with large error, resulting in reconstruction artifacts. Use of spectral reprojection methods such as Gegenbauer reconstruction are suggested to overcome this problem. Empirical results are provided which show that equivalent, or better reconstruction quality can be achieved using only accurate, low frequency spectral coefficients in conjunction with Gegenbauer reconstruction, as opposed to traditional Fourier reconstruction employing all spectral samples.

Since high-resolution reconstruction methods require discontinuity data, an elaborate treatment of discontinuity detection from spectral data is given. The concentration method of jump detection is employed and its statistical formulation developed. A noise robust detector is designed and its associated performance metric provided. Further, an improved jump detection method using combinations of concentration factors is developed.

TABLE OF CONTENTS

| | Page |
|---|-----------|
| LIST OF TABLES | vii |
| LIST OF FIGURES | viii |
| 1 INTRODUCTION | 1 |
| 1.1. Problem Formulation | 1 |
| 1.2. The Need for Spectral Interpolation | 3 |
| 1.3. Fourier Reconstruction and Piecewise-Smooth Functions | 4 |
| 1.4. Discontinuity Detection | 5 |
| 1.5. Outline | 6 |
| 2 RECONSTRUCTION METHODS FOR IRREGULARLY SAMPLED SPECTRAL DATA | 7 |
| 2.1. Non-equispaced “FFTs” | 7 |
| 2.2. Convolutional Gridding | 12 |
| 2.2.1. Accuracy of the Gridding Solution | 14 |
| 2.2.2. Implementation Steps | 16 |
| 2.3. Uniform Re-Sampling (URS) | 19 |
| 2.3.1. Implementation Notes | 21 |
| 2.4. Iterative Methods | 24 |
| 3 ANALYSIS OF UNIFORM RE-SAMPLING AND HIGH RESOLUTION RECONSTRUCTION | 29 |
| 3.1. Error Analysis of URS Reconstruction | 30 |
| 3.2. High Resolution Reconstruction | 33 |
| 3.2.1. Gegenbauer Reconstruction | 35 |
| 3.2.2. Results using Gegenbauer Reconstruction | 36 |
| 4 DISCONTINUITY DETECTION FROM SPECTRAL DATA | 38 |
| 4.1. Concentration Method | 38 |
| 4.2. Statistical Analysis of the Concentration Method | 42 |
| 4.3. Detector Design | 44 |
| 4.4. Performance of the Detector | 46 |

| CHAPTER | Page |
|---|------|
| 4.5. Choice of Detector Parameters | 47 |
| 4.6. Results | 48 |
| 4.7. Extensions and Remarks | 51 |
| 5 SUMMARY | 54 |
| 5.1. Avenues for Further Work | 55 |
| 5.1.1. Spectral Sampling | 55 |
| 5.1.2. Discontinuity Detection from Spectral Data | 55 |
| REFERENCES | 57 |
| APPENDIX A THE SHEPP-LOGAN PHANTOM AND SAMPLING SCHEMES | 60 |
| A.1. Shepp-Logan Phantom | 61 |
| A.2. Sampling Schemes | 61 |

LIST OF TABLES

| Table | Page |
|---|------|
| 2.1 Comparison of reconstruction methods | 28 |
| 4.1 Popular concentration factors | 41 |
| 4.2 Choice of locations x_i and effect on performance using only the trigonometric concentration factor ($L = 5, N = 32, \rho^2 = 7.5$) | 48 |
| 4.3 Choice of concentration factors and effect on performance using a fixed set of locations x_i ($L = 5, N = 32, \rho^2 = 7.5$) | 48 |

LIST OF FIGURES

| Figure | Page |
|--|------|
| 1.1 Non-harmonic kernels ($N = 128$) | 4 |
| 1.2 Fourier reconstruction of piecewise-smooth functions | 5 |
| 2.1 Test function and sampling scheme | 7 |
| 2.2 Trapezoidal quadrature rule | 13 |
| 2.3 Comparison of reconstruction with and without density compensation, $N = 128$ | 13 |
| 2.4 Voronoi tessellation for spiral and radial sampling trajectories, $N = 16$ | 15 |
| 2.5 Variation of Δ_p with N | 17 |
| 2.6 Convergence behavior of gridding reconstruction for different sampling schemes | 18 |
| 2.7 Illustration of convolutional gridding, $N = 128$ | 19 |
| 2.8 Comparison of reconstruction with and without oversampling, $N = 128$ | 20 |
| 2.9 Reconstruction using URS, $N = 256$ | 21 |
| 2.10 Illustration: block uniform re-sampling (BURS) | 23 |
| 2.11 Iterative reconstruction computed using the conjugate-gradient method | 26 |
| 2.12 Shepp-Logan phantom reconstruction from spiral k -space samples using BURS | 27 |
| 3.1 Need for regularization | 30 |
| 3.2 Plot of σ_p and $u_p^H b$ for log sampling, $N = 256$ | 31 |
| 3.3 Error in recovered coefficients, $N = 256$ | 33 |
| 3.4 Error in recovered Fourier coefficients of the Shepp-Logan phantom | 33 |
| 3.5 Test function reconstruction (log sampling), $N = 256$ | 34 |
| 3.6 Comparison of filtered Fourier ($N = 256$) and Gegenbauer reconstructions ($N = 64, m =$ $\lambda = 2$) | 36 |
| 3.7 Comparison of filtered Fourier ($N = 256$) and Gegenbauer reconstructions ($N = 64, m =$ $4, \lambda = 4$) with jump locations estimated using the concentration method | 37 |
| 4.1 Example of the concentration method | 40 |
| 4.2 Step response of different concentration factors ($N = 32$) | 42 |
| 4.3 Concentration factors plotted in Fourier space | 43 |

| Figure | Page |
|--|------|
| 4.4 Example of a signal vector ($N = 32$, trigonometric factor) | 45 |
| 4.5 ROC curve ($L = 3, N = 32$, exponential factor) | 49 |
| 4.6 Single concentration factor versus a combination of factors ($L = 3, N = 32, \rho^2 = 7.5$) | 50 |
| 4.7 Performance on a test function ($L = 3, N = 128, \rho^2 = 7.5$) | 51 |
| 4.8 Edge detection on the Shepp-Logan phantom (trigonometric factor, $L = 7, N = 128, \rho^2 =$ 25.5) | 53 |
| A.1 Shepp-Logan phantom and its Fourier transform, $N = 128^2$ | 61 |
| A.2 Sampling schemes | 63 |

1. INTRODUCTION

Fourier analysis and synthesis methods find widespread application in several branches of science, mathematics and engineering. For the class of compactly supported functions, application of the Whittaker-Kotelnikov-Shannon (WKS) sampling theorem permits computation of Fourier integrals using discrete sums, thereby allowing computer implementations on discrete measurements. Further, use of the fast Fourier transform (FFT) algorithm, [1], provides significant savings in computational cost. However, this framework poses a stringent constraint on all input and output data, requiring them to be acquired (or evaluated) on a uniform grid. In many applications, this is either not physically possible or is computationally costly. For example, in magnetic resonance imaging (MRI), the MR scanner acquires samples in “ k -space”, with the physical image related to k -space data by a two-dimensional Fourier transform. If k -space samples are acquired on a Cartesian grid, the two-dimensional discrete inverse Fourier transform, computed using the FFT algorithm, would provide the desired image reconstruction. However, for several reasons, including robustness to motion artifacts, reduction of field inhomogeneities, and mechanical constraints, acquiring Cartesian data is not popular. Instead, samples are acquired along spiral, radial and other non-Cartesian trajectories. There is no consensus on what constitutes an accurate, yet computationally feasible reconstruction method under these circumstances.

The theory and results presented in this thesis partially address these issues. In particular, we are interested in the reconstruction of support-limited functions from their non-uniform spectral samples. Although MRI constitutes a motivating application, the methods and results presented here are general, and can be adapted to other problems. For example, these methods may be used to compute the spectral coefficients of a bandlimited function from non-uniform physical space samples, or to perform edge detection from irregular spectral data.

A secondary objective of this thesis is the detection of jump discontinuities in such functions, from equispaced but noisy Fourier spectral data. Jump detection is pursued both as an independent signal processing task, as well as part of a reconstruction scheme for piecewise-smooth functions.

1.1. Problem Formulation

For ease of analysis, this development assumes a one-dimensional function. However, remarks and examples on its extension to two dimensions are provided. Let $f \in L^2(\mathbb{R})$ be supported in $[-\pi, \pi]$; i.e., $f(x) = 0, \forall |x| > \pi$ with $f(-\pi) = f(\pi)$.

We shall assume that f has a Fourier transform, \hat{f} , defined by

$$\begin{aligned}\hat{f}(\omega) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} f(x)e^{-i\omega x} dx, \quad \omega \in \mathbb{R} \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x)e^{-i\omega x} dx\end{aligned}\tag{1.1}$$

The compact support also permits a Fourier series representation of f (albeit, a 2π -repeated function $\sum_{p=-\infty}^{\infty} f(x - 2\pi p)$, $p \in \mathbb{Z}$) from the discrete, Fourier series coefficients,

$$\hat{f}(k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x)e^{-ikx} dx, \quad k \in \mathbb{Z}\tag{1.2}$$

With a slight abuse of notation, we shall refer to both the function, and its 2π -repeated version as f .

Therefore, given a finite number of arbitrary spectral samples, $\hat{f}(\omega_k)$, $k = 0, \dots, N - 1$, with ω_k not necessarily integers, our objective is two-fold:

1. To obtain f . This can include recovery of the function's values
 - at equispaced grid-points, $f(x_k)$, where $x_k = -\pi + \frac{2\pi k}{N}$, $k = 0, \dots, N - 1$
 - or preferably, over the entire domain, $f(x)$, $\forall x \in [-\pi, \pi]$
2. To relate $\hat{f}(\omega_k)$ to the equispaced Fourier coefficients $\hat{f}(k)$. In addition to enabling function recovery over the entire domain $[-\pi, \pi]$, knowledge of equispaced coefficients is essential to perform certain other tasks such as high-resolution reconstruction, discontinuity detection from spectral data, etc.

Related to the objective of obtaining the underlying piecewise-smooth function f , is the ability to detect jump discontinuities, including their locations, magnitudes and signs. In fact, many reconstruction methods for recovering f require a-priori jump discontinuity location information. We will consider the problem of obtaining the jump discontinuities of f from a finite number of its equispaced, but noisy Fourier coefficients, with the assumption that noise is additive, zero-mean and complex Gaussian. The formulation and corresponding notation for this problem is provided in the concerned chapter.

1.2. The Need for Spectral Interpolation

Given the equispaced samples $\hat{f}(k)$, we reconstruct f by computing the partial Fourier sum¹,

$$S_N f(x) = \sum_{k=-N/2}^{N/2-1} \hat{f}(k) e^{ikx} \quad (1.3)$$

Since we have equispaced data, this sum constitutes a good quadrature, [2], to the inverse Fourier integral, thereby resulting in a good approximation. Alternatively, it is a simple exercise to show that $S_N f$ may be written as

$$S_N f(x) = (f * D_N)(x), \quad D_N(\eta) = \sum_{k=-N/2}^{N/2-1} e^{ik\eta} \quad \text{is the Dirichlet kernel} \quad (1.4)$$

Since the Dirichlet kernel is the Fourier partial sum approximation of the Dirac delta measure, the error component is purely due to truncation, with $S_N f \rightarrow f$ as N increases. However, the same does not hold when we are given non-equispaced spectral data $\hat{f}(\omega_k)$. Consider reconstruction using the “non-equispaced sum”

$$\tilde{f}(x) \approx \sum_{k=-N/2}^{N/2-1} \hat{f}(\omega_k) e^{i\omega_k x} \quad (1.5)$$

With $\hat{f}(\omega_k)$ non-equispaced, the above sum does not constitute a good quadrature, [2]. It is simple to show that

$$\tilde{f}(x) = (f * A_N)(x), \quad A_N(\eta) = \sum_{k=0}^{N-1} e^{i\omega_k \eta} \quad (1.6)$$

Here, $A_N(\eta)$ is a non-harmonic kernel plotted in Figure 1.1 for two sampling distributions

- ω_k jittered randomly about equispaced nodes (Figure 1.1(a))
- ω_k distributed logarithmically with more samples near the origin and fewer away from the origin (Figure 1.1(b))

Details of these and other sampling schemes can be found in Appendix A. In either case, we note significant deviation from the Dirichlet kernel, with $A_N(\eta) \not\rightarrow D_N(\eta)$ as N increases. This holds for any sampling distribution deviating from equispaced nodes, with the consequence being non-convergence of \tilde{f} to the true function, i.e.,

$$(f * A_N)(x) \not\rightarrow f(x) \text{ as } N \text{ increases}$$

¹In this particular passage, we shall ignore presence of the Gibbs phenomenon. While reduced order of convergence is of concern to us, a bigger issue to address with arbitrary spectral samples is that of non-convergence of the “non-equispaced sum”.

A representative reconstruction of the boxcar function using jittered sampling is shown in Figure 1.1(d). In

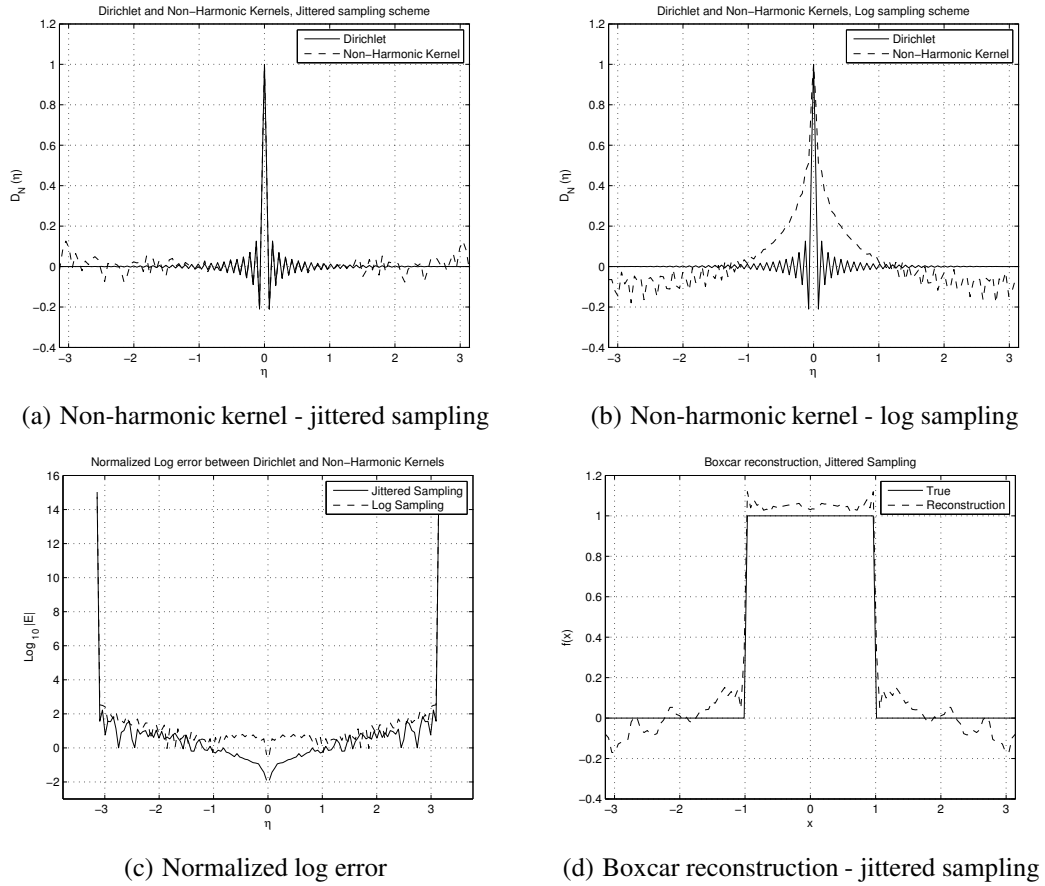



Fig. 1.1. Non-harmonic kernels ($N = 128$)

order to employ a Fourier partial sum reconstruction, we therefore need to interpolate $\hat{f}(\omega_k)$ to equispaced nodes. Further, to obtain a good reconstruction, we require the interpolated coefficients to closely approximate $\hat{f}(k)$.

1.3. Fourier Reconstruction and Piecewise-Smooth Functions

It is well known that Fourier reconstruction offers spectral convergence only for smooth, periodic functions, [3]. Specifically, the approximation error for Fourier reconstruction depends on the number of smooth derivatives in the function. For piecewise-smooth functions, one encounters the Gibbs phenomenon, [4], that is, reduced overall convergence to first order and the introduction of non-physical oscillations in the vicinity of jump discontinuities. Figure 1.2 illustrates this behaviour. Also plotted is the log error of the reconstruc-

tion for increasing number of coefficients, N . It can be observed that accuracy does not improve significantly by increasing the number of coefficients. Further, oscillations at the discontinuity still persist. Filtering is the conventional post-processing method to improve the quality of reconstruction. This is, however, not a solution, since the resulting reconstruction suffers from smeared jump discontinuities, while the new convergence rate is still not exponential. Methods to overcome the Gibbs phenomenon exist, e.g., [5, 6], and it is conceivable that their incorporation in our scheme of reconstruction will provide better results.



(a) Reconstruction illustrating Gibbs phenomenon

(b) Log error in reconstruction

Fig. 1.2. Fourier reconstruction of piecewise-smooth functions

1.4. Discontinuity Detection

Detection of jump discontinuities in piecewise-smooth functions is an important signal processing task. While discontinuity data may be useful in itself, discontinuity detection also constitutes an important pre-processing step in other tasks such as segmentation and feature extraction in image processing, and spectral reprojection for high-resolution function reconstruction. For example, accurate jump discontinuity information is necessary for an alternative reconstruction procedure outlined in Chapter 3.

The recovery of jump discontinuities from spectral coefficients is complicated due to the fact that spectral data are global in nature, while a jump is a local phenomenon. This problem is further exacerbated by the addition of noise to spectral data. Besides, traditional jump detection methods require use of a threshold which is data dependent and empirically determined. Therefore, there is need to develop a noise-robust jump detection scheme, with a clearly defined performance metric and ability to determine the threshold based on desired performance and noise level.

1.5. Outline

The remainder of this thesis is organized as follows: Chapter 2 provides a survey of popular methods which address the reconstruction of support-limited functions from non-Cartesian spectral data. These include convolutional “gridding” reconstruction, uniform-resampling (URS) and its variants, and a class of iterative methods. Since the non-uniform FFT is extensively utilized in many of these methods, a section is devoted to its explanation. Remarks about reconstruction error and computational cost for these methods are provided. Chapter 3 contains an analysis of the uniform-resampling method, which establishes the relationship between $\hat{f}(\omega_k)$ and $\hat{f}(k)$. The quantum of deviation from regular node spacing is related to the accuracy of the interpolated coefficients, and illustrative examples for different types of sampling patterns are provided. Further, variants of traditional Fourier reconstruction and an alternative high-resolution reconstruction scheme are suggested. Chapter 4 is concerned with determining the location, magnitude and sign of jump discontinuities from spectral data. The *concentration method* for jump detection is employed and its statistical characterization provided. The detection problem is formulated and solved, resulting in a simple design and performance metric. Using this metric, an improved edge detection method is developed. Several illustrative examples are included. Finally, Chapter 5 provides some concluding comments and lists avenues for future work.

2. RECONSTRUCTION METHODS FOR IRREGULARLY SAMPLED SPECTRAL DATA

This chapter begins with an introduction to the non-uniform or non-equispaced FFT algorithm, which is widely used in computing non-harmonic Fourier partial sums. A brief survey of methods to recover support-limited functions from non-Cartesian Fourier samples follows, along with illustrative examples. Throughout this chapter, we shall use the one-dimensional example displayed in Figure 2.1 to illustrate different reconstruction methods.

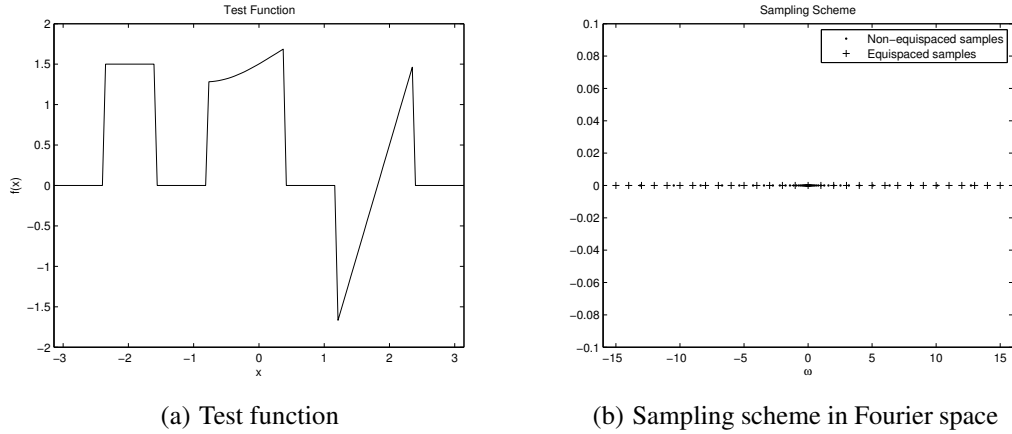


Fig. 2.1. Test function and sampling scheme

Samples in Fourier space are acquired at logarithmic intervals, with more samples acquired in lower frequencies. If ω_k are the nodes at which measurements are acquired, $|\omega_k|$ is logarithmically distributed between 10^{-v} and $N/2$, where N is the total number of samples. A value of $v = 1.5$ was used. This sampling scheme will henceforth be referred to as “log sampling”. Motivation for this sampling scheme can be found in MRI, where typical data acquisition schemes oversample the low frequencies of k -space, while undersampling the high frequencies.

2.1. Non-equispaced “FFTs”

Fast Fourier transforms for non-equispaced data (NEFFT), as described in [7–10], are a family of efficient methods to compute trigonometric sums of the form,

$$\hat{f}(\omega_p) = \sum_{k=0}^{N-1} f(x_k) e^{-i\omega_p x_k}, \quad p = 0, \dots, N-1 \quad (2.1)$$

Here, x_k denote the nodes at which input data are available, and ω_p denote the nodes at which the output is evaluated. The NEFFT allows for the following formulations:

Formulation (1): equispaced input data, with output evaluated at non-equispaced nodes

Formulation (2): non-equispaced input data, with output evaluated at equispaced nodes

Formulation (3): non-equispaced input and output nodes

It is important to remember that Formulations (2) and (3) do not, in general, constitute valid evaluations of the discrete Fourier transform. Further, (2.1) computes spectral coefficients from physical space data; however, the NEFFT is equally applicable for computing function values from spectral data.

Non-equispaced FFT schemes (for example, [8]) result from replacing the non-harmonic exponential $e^{-i\omega_k x}$ with its Fourier series expansion; i.e.,

$$e^{-i\omega_k x} = \sum_{m=-\infty}^{\infty} a_m e^{-imx}, \quad x \in [-\pi, \pi]$$

where

$$a_m = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-i\omega_k x} e^{imx} dx = \text{sinc}(m - \omega_k)$$

The coefficient a_m obtains a maximal value at $m = [\omega_k]$ and decays as $|m - [\omega_k]|$ increases¹. In a practical computation, one would use a $2q + 1$ term truncated sum of the form

$$e^{-i\omega_k x} \approx \sum_{m=[\omega_k]-q}^{[\omega_k]+q} a_m e^{-imx} \quad (2.2)$$

If e is the truncation error, we have

$$\begin{aligned} \|e\|_2 &= \left\| \sum_{|m-[\omega_k]|>q} a_m e^{-imx} \right\|_2 \\ &= 2\pi \left(\sum_{|m-[\omega_k]|>q} |a_m|^2 \right)^{1/2} \\ &= 2\pi \left(\sum_{|m-[\omega_k]|>q} |\text{sinc}(m - \omega_k)|^2 \right)^{1/2} \end{aligned} \quad (2.3)$$

where the second step is obtained using Parseval's identity. Since the sinc function decays as $1/m$, a large number of terms are required in the sum to permit a good approximation. This is counter-productive to our objective of reducing computational cost. A better solution is obtained by computing the Fourier series

¹ $[y]$ denotes the nearest integer to y .

expansion of a composite function $\phi(x)e^{-i\omega_k x}$, with $\phi(\cdot)$ satisfying specific properties. This approach yields

$$\phi(x)e^{-i\omega_k x} = \sum_{m=-\infty}^{\infty} b_m e^{-imx}, \quad x \in [-\pi, \pi]$$

where

$$b_m = \frac{1}{2\pi} \int_{-\pi}^{\pi} \phi(x) e^{-i\omega_k x} e^{imx} dx = \hat{\phi}(m - \omega_k)$$

Therefore,

$$e^{-i\omega_k x} = \frac{1}{\phi(x)} \sum_{m=-\infty}^{\infty} \hat{\phi}(m - \omega_k) e^{-imx}, \quad x \in [-\pi, \pi]$$

Using a $2q + 1$ term truncated series, we have

$$e^{-i\omega_k x} \approx \frac{1}{\phi(x)} \sum_{m=[\omega_k]-q}^{[\omega_k]+q} \hat{\phi}(m - \omega_k) e^{-imx}, \quad x \in [-\pi, \pi] \quad (2.4)$$

with truncation error

$$\|e\|_2 = \left\| \sum_{|m-[\omega_k]|>q} b_m e^{-imx} \right\|_2 = 2\pi \left(\sum_{|m-[\omega_k]|>q} |\hat{\phi}(m - \omega_k)|^2 \right)^{1/2} \quad (2.5)$$

For computational cost to be low, we require $\hat{\phi}(\cdot)$ to have bounded support in Fourier space. Additionally, ϕ should be 2π -periodic and have a Fourier series representation. Further, since we divide by $\phi(x)$, we cannot have $\phi(x) \approx 0$ in $[-\pi, \pi]$. Since these are contradictory statements, we settle for a compromise solution with the energy of $\phi(\cdot)$ concentrated in a small band in either domain. Some examples of popular interpolating functions are given below, while more detailed treatments can be found in [9, 11].

- Periodized Gaussian Bell

$$\phi(x) = \frac{1}{\sqrt{\pi\beta}} \sum_{r \in \mathbb{Z}} e^{-\frac{\eta^2(x+r)^2}{\beta}}, \quad \hat{\phi}(k) = \frac{1}{2\pi\eta} e^{-\frac{k^2\beta}{4\eta^2}}$$

- Kaiser-Bessel²

$$\phi(x) = \begin{cases} I_0(\beta\sqrt{\eta^2 - x^2}) & |x| \leq \eta \\ 0 & |x| > \eta \end{cases}, \quad \hat{\phi}(k) = \frac{1}{\pi} \frac{\sinh(\eta\sqrt{\beta^2 - k^2})}{\sqrt{\beta^2 - k^2}}$$

²The Kaiser-Bessel function is a computationally simple approximation to a prolate spheroidal wave function [12] (PSWF).

Here, β and η are parameters which control the decay of $\phi(\cdot)$ in either domain. Let us assume that x_k are equispaced in (2.1). Substituting (2.4) in (2.1), we have

$$\begin{aligned} \hat{f}(\omega_p) &\approx \sum_{k=0}^{N-1} f(x_k) \left(\frac{1}{\phi(x_k)} \sum_{m=[\omega_p]-q}^{[\omega_p]+q} \hat{\phi}(m - \omega_p) e^{-imx_k} \right) \\ &\approx \sum_{m=[\omega_p]-q}^{[\omega_p]+q} \hat{\phi}(m - \omega_p) \left(\sum_{k=0}^{N-1} \frac{f(x_k)}{\phi(x_k)} e^{-imx_k} \right) \end{aligned} \quad (2.6)$$

Therefore, the non-equispaced FFT is computed as follows,

Algorithm 1 Computing the NEFFT

Input:

f at the equispaced grid points (x_k)

Output evaluation points ω_p

Truncation parameter q

Interpolating function ϕ

Pre-computation: $\phi(x_k)$ and $\hat{\phi}(m - \omega_k)$, $m = [[\omega_k] - q, [\omega_k] + q]$

1: Compensation for use of the interpolating function: $g(x_k) = \frac{f(x_k)}{\phi(x_k)}$

2: Standard FFT computation: $\hat{g}(p) = \sum_{k=0}^{N-1} g(x_k) e^{ipx_k}$, $p = -N/2, \dots, N/2 - 1$

3: Interpolation to the desired grid points: $\hat{f}(\omega_p) = \sum_{m=[\omega_p]-q}^{[\omega_p]+q} \hat{\phi}(m - \omega_p) \hat{g}(p)$, $p = 0, \dots, N - 1$

While the N -point FFT entails a computational cost of $\mathcal{O}(N \log N)$, the additional interpolation and compensation in the NEFFT only require $N(2q + 1)$ and N computations respectively. Since q is small (results in [8] use the Kaiser-Bessel function with $q = 7$), the additional computational burden is insignificant³. Therefore, evaluation of the NEFFT requires $\mathcal{O}(N \log N + \mathcal{M})$ computations, where \mathcal{M} is a small multiple of N . Modifications to account for oversampling are straightforward, and can be found, for example in [8]. Equation (2.6) is an example of the NEFFT which takes equispaced input data and evaluates the output at non-equispaced nodes. Formulation (2), or the adjoint NEFFT is an efficient method to compute

³This cost does not include precomputation of the interpolation function $\hat{\phi}(m - \omega_k)$.

trigonometric sums of the form

$$\tilde{f}(x_p) = \sum_{k=0}^{N-1} \hat{f}(\omega_k) e^{i\omega_k x_p}, \quad x_p = -\pi + \frac{2\pi p}{N}, \quad p = 0, \dots, N-1 \quad (2.7)$$

Using arguments similar to that in Formulation (1), it is possible to show (eg., [8]) that (2.7) can be computed efficiently as

$$\tilde{f}(x_p) \approx \frac{1}{\phi(x_p)} \sum_{k=-N/2}^{N/2-1} \left(\sum_{m \text{ st. } |k-\omega_m| \leq q} \hat{f}(\omega_m) \hat{\phi}(k-\omega_m) \right) e^{ikx_p}, \quad x_p = -\pi + \frac{2\pi p}{N}, \quad p = 0, \dots, N-1 \quad (2.8)$$

Compared to (2.6), it can be observed that compensation, FFT evaluation and interpolation are now performed in reverse order.

Algorithm 2 Computing the adjoint NEFFT

Input:

Non-equispaced measurement nodes ω_k

\hat{f} at the non-equispaced nodes

Truncation parameter q

Interpolating function ϕ

Pre-computation: $\phi(x_p)$ and $\hat{\phi}(k-\omega_m)$, $m \text{ st. } |k-\omega_m| \leq q$

- 1: Interpolation to equispaced nodes: $\hat{g}(\omega_k) = \sum_{m \text{ st. } |k-\omega_m| \leq q} \hat{f}(\omega_m) \hat{\phi}(k-\omega_m)$, $k = -N/2, \dots, N/2-1$
 - 2: Standard FFT computation: $g(x_p) = \sum_{k=-N/2}^{N/2-1} \hat{g}(\omega_k) e^{ikx_p}$, $x_p = -\pi + \frac{2\pi p}{N}$, $p = 0, \dots, N-1$
 - 3: Compensation for use of the interpolating function: $f(x_p) = \frac{g(x_p)}{\phi(x_p)}$
-

It is important to emphasize that the NEFFT is a tool to compute trigonometric sums with non-harmonic exponentials efficiently. Consequently, all error metrics are with respect to the corresponding direct sum evaluations. For example, error measurements for the adjoint NEFFT are with respect to the direct computation in (2.7) (or \tilde{f}). The NEFFT algorithm makes no guarantees on the relative error between \tilde{f} and the true function f . That said, accuracy of the NEFFT algorithm (with respect to direct evaluation) is well studied in literature. For example, [9] provides bounds for approximation error for different ϕ , besides highlighting the relation between number of terms in the sum, q and approximation error. A variation of this algorithm can

be found in [10], where min-max interpolation techniques are used to compute “optimum” interpolation coefficients. Several open source software implementations of non-uniform FFT algorithms are also available ([13, 14]).

2.2. Convolutional Gridding

Convolutional gridding has its origins in astronomical image reconstruction, [15, 16], with O’Sullivan, [17], introducing the method in medical imaging. It is by far the most popular method for computing approximate Fourier transforms on arbitrary grids. The method is based on the discretization of the inverse Fourier integral,

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(\omega) e^{i\omega x} d\omega \quad (2.9)$$

Given N arbitrary spectral samples, $\hat{f}(\omega_k)$, $k = 0, \dots, N - 1$ with ω_k not necessarily integers, the simplest discretization is given by

$$f(x) \approx \sum_{k=0}^{N-1} \alpha_k \hat{f}(\omega_k) e^{i\omega_k x} \quad (2.10)$$

Here, α_k are known as density compensation factors (DCFs) which are data independent weighting terms used to define an appropriate quadrature rule, [2], for approximating the inverse Fourier integral. Further, it is clear that the standard FFT algorithm cannot be directly applied to evaluate the partial sum. However, Formulation (2) of the NEFFT algorithm (§2.1) can be employed, which computes (2.10) at cost $\mathcal{O}(N \log N + \mathcal{M})$, where \mathcal{M} is a small multiple of N . The only additional computation required over the NEFFT is pre-multiplication of $\hat{f}(\omega_k)$ by the DCFs α_k .

To understand the need for additional weighting factors, consider evaluation of the non-equispaced sum as opposed to the DFT. With equispaced data, the DFT describes a trapezoidal quadrature rule with equal sub-intervals (Figure 2.2(a)), i.e., the integral is approximated as the sum of areas of individual trapeziums of constant width. Hence, a single weighting factor (equal to the sampling interval) suffices. Moreover, for smooth, periodic functions, this quadrature rule guarantees exponentially accurate approximations, [3]. However, this is not the case with the non-equispaced sum, which describes a trapezoidal quadrature rule with unequal sub-intervals (Figure 2.2(b)). To satisfy the quadrature, it is therefore necessary to have additional weighting factors, α_k which are equal to the corresponding sub-interval lengths. i.e.,

$$\alpha_k = \omega_{k+1} - \omega_k \quad (2.11)$$

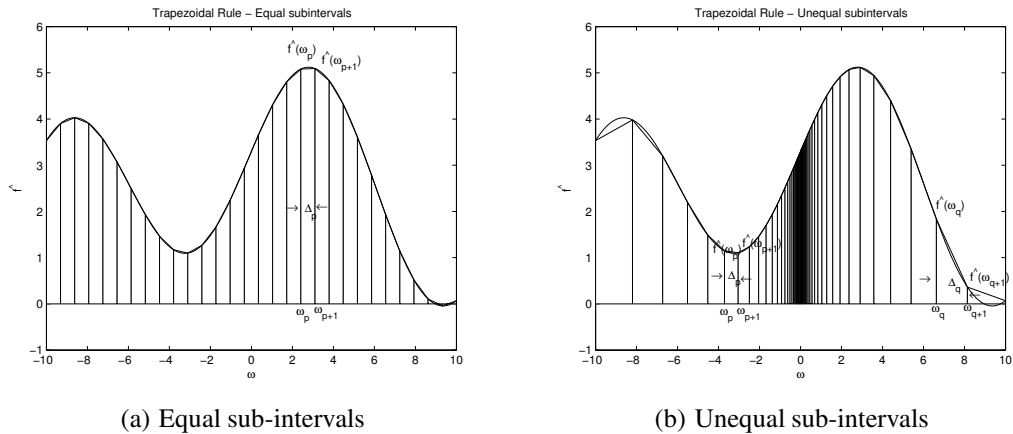


Fig. 2.2. Trapezoidal quadrature rule

As an illustration of the importance of density compensation, consider Figure 2.3, which shows reconstruction of the test function using the log sampling scheme. Figure 2.3(a) shows the poor reconstruction as a result of not using any density compensation, while Figure 2.3(b) shows the reconstruction after using the DCFs in (2.11). Several other density compensation factors and methods exist, including

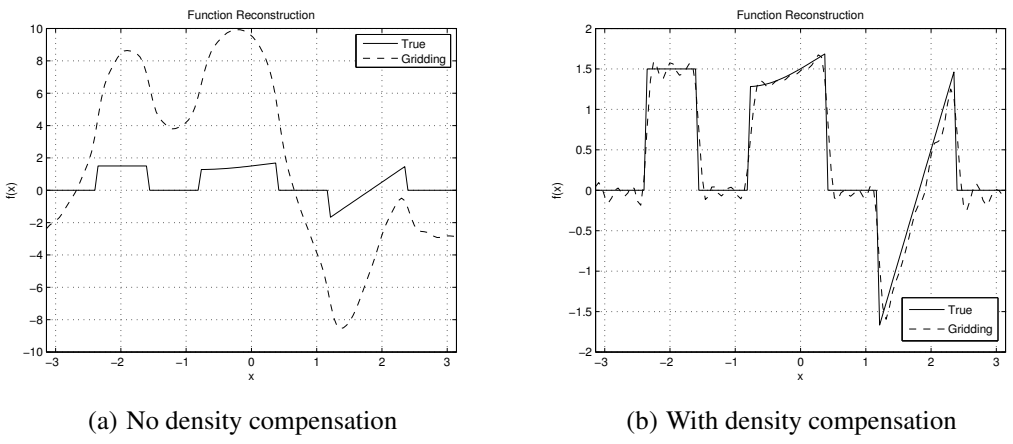


Fig. 2.3. Comparison of reconstruction with and without density compensation, $N = 128$

- Voronoi compensation [18]: This method is applicable to two-dimensional problems where Voronoi tessellation of the non-Cartesian nodes is performed, and each sample is weighted by the area of its Voronoi cell. This can be interpreted as a direct extension of the DCFs introduced in (2.11). An illustration of Voronoi tessellation is provided in Figure 2.4 for spiral and radial sampling trajectories.

- DCF of Jackson et al. [11]: An expression for the DCF is

$$\alpha_k = \frac{1}{(\check{f} * \hat{\phi})(\omega_k)}$$

Here, $\check{f} := \sum_{k=0}^{N-1} \hat{f}(\omega) \delta(\omega - \omega_k)$ is the set of discrete measurements, while $\hat{\phi}$ is the interpolating window function (§2.1, §2.2.2). In regions of oversampling, $\check{f} * \hat{\phi}$ takes large values (due to the large number of measurements contributing to the convolution sum), with relatively smaller values in regions of undersampling due to the smaller number of measurements contributing to the convolution sum. Therefore, as in the case of Voronoi compensation, α_k assumes small values in regions of oversampling and large values in regions of undersampling.

- Pipe-Menon density compensation [19]: From Chapter 1, we know that use of a “non-equispaced sum” of the form (2.10) results in the reconstruction being corrupted or blurred by a non-harmonic kernel. In two dimensions, this kernel is referred to as a point spread function (PSF). The Pipe-Menon method computes weights such that the weighted kernel or PSF is a close approximation to the Dirac delta measure. Let $W(\omega) = \sum_{k=0}^{N-1} \alpha_k \delta(\omega - \omega_k)$ denote the weights and $\hat{\phi}$ denote the interpolating window function. The Pipe-Menon DCFs are the solution to

$$\phi(x) \mathcal{F}^{-1}\{W(\omega)\} \approx \delta(x) \quad (2.12)$$

with an iterative formulation as follows

$$W(\omega)|_{i+1} = \frac{W(\omega)|_i}{(\hat{\phi} * W)(\omega)} \quad (2.13)$$

Details and examples of the method can be found in [19, 20]. This approach has the additional advantage of not requiring the compensation step (§2.1, §2.2.2) while reconstructing the gridding solution.

Other notable density compensation methods include Jacobian based methods [21] and the Sedarat-Nishimura DCFs [22].

2.2.1. Accuracy of the Gridding Solution

A comprehensive error analysis for the gridding method is beyond the scope of this thesis due to the wide variety of density compensation methods in use. However, it is instructive to work out the nature of error

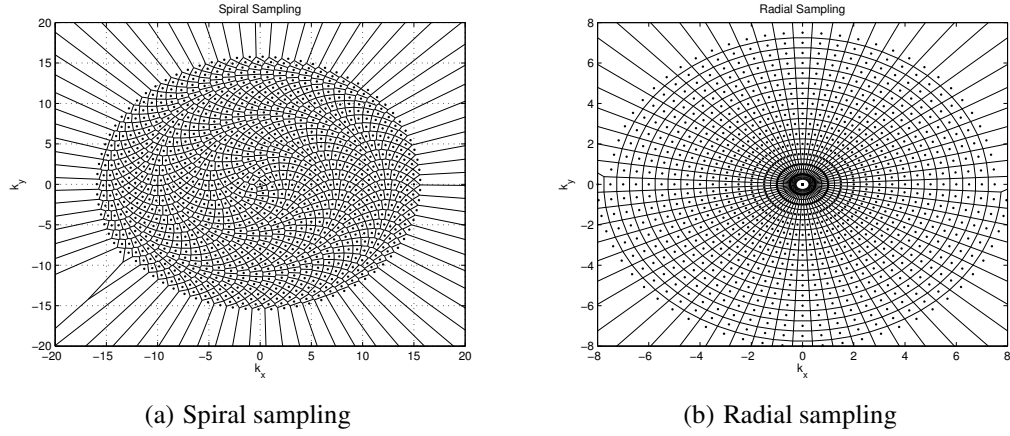


Fig. 2.4. Voronoi tessellation for spiral and radial sampling trajectories, $N = 16$

for the simplest of cases: a one-dimensional function with DCFs as in (2.11). Consider a single sub-interval of Fourier space, $[\omega_p, \omega_{p+1}]$. If e_p denotes the approximation error in this sub-interval, we have

$$e_p = \int_{\omega_p}^{\omega_{p+1}} \hat{f}(\omega) e^{i\omega x} d\omega - \frac{1}{2} (\omega_{p+1} - \omega_p) \left(\hat{f}(\omega_p) e^{i\omega_p x} + \hat{f}(\omega_{p+1}) e^{i\omega_{p+1} x} \right)$$

Note that e_p is a function of x , and the sub-interval bounds, ω_p, ω_{p+1} . From error analysis of the trapezoidal quadrature rule, [2], we know that the (leading order) error in each sub-interval obeys

$$e_p \leq \frac{(\omega_{p+1} - \omega_p)^3 v_p}{12}, \quad p = 0, \dots, N - 2 \quad (2.14)$$

where

$$\begin{aligned} v_p &= \left. \frac{d^2(\hat{f}(\omega) e^{i\omega x})}{d\omega^2} \right|_{\omega=\zeta}, \quad \zeta \in [\omega_p, \omega_{p+1}] \\ &= -x^2 \hat{f}(\zeta) e^{i\zeta x} + 2ix \left. \frac{\partial \hat{f}}{\partial \omega} e^{i\omega x} \right|_{\omega=\zeta} + \left. \frac{\partial^2 \hat{f}}{\partial \omega^2} e^{i\omega x} \right|_{\omega=\zeta} \end{aligned}$$

Consequently, the total approximation error is the sum of errors incurred in each sub-interval, i.e.,

$$e \leq \sum_{p=0}^{N-2} e_p \leq \sum_{p=0}^{N-2} \frac{(\omega_{p+1} - \omega_p)^3 v_p}{12}$$

Let $\Delta_p = \omega_{p+1} - \omega_p$. Then,

$$e \leq \sum_{p=0}^{N-2} \frac{\Delta_p^3 v_p}{12} \quad (2.15)$$

Therefore, the error depends on two factors – the inter-sample intervals Δ_p , and the data-dependent term v_p .

We may make the following observations:

- The data dependent term, v_p , is a function of the Fourier transform \hat{f} and ω . Due to the decay of \hat{f} (\hat{f} decays as $\mathcal{O}\left(\frac{1}{\omega}\right)$ for piecewise-smooth f), $v_p \rightarrow 0$ as $|\omega| \rightarrow \infty$. Therefore, the dominant contribution to e is from terms corresponding to the low frequency samples, i.e.,

$$e \approx \sum_{p \text{ st. } |\omega_p| < \eta} \frac{\Delta_p^3 v_p}{12}$$

where η is a parameter which depends on the decay rate of Fourier coefficients.

- For a given function, convergence of the gridding approximation therefore depends on Δ_p , and consequently, the sampling scheme employed. Consider Δ_p corresponding to the low frequency coefficients. In the case of jittered sampling (Appendix A), in which acquired samples are jittered about the equispaced nodes, Δ_p remains approximately constant as N increases (Figure 2.5), while for the log sampling scheme, which has more low frequency samples than high frequency samples, $\Delta_p \rightarrow 0$ as $N \rightarrow \infty$ (quadrature convergence). Unsurprisingly, empirical results in Figure 2.6 suggest that the gridding approximation with DCFs of (2.11) shows poor convergence properties for jittered sampling, while better convergence behavior can be expected for the log sampling case. Similar arguments can be made for the popular Voronoi compensation scheme in two dimensions.

Gridding reconstruction is particularly popular in MRI, where sampling schemes similar to the above mentioned log sampling are encountered. Oversampling of the lower frequencies in k -space is typical, with undersampling of high frequency regions. For such schemes, as noted previously, we can expect good convergence behavior.

2.2.2. Implementation Steps

Although computation of the gridding solution shares several steps with the adjoint NEFFT algorithm, investigation of the constituent steps provides additional intuition behind the method. Further, an alternative interpretation of the adjoint NEFFT is given. We start with the N arbitrary spectral samples, $\hat{f}(\omega_k)$, $k = 0, \dots, N - 1$ with ω_k not necessarily integers.

1. Density Compensation: The first step (Figure 2.7(a)) involves weighting the measured coefficients by precomputed density compensation factors yielding

$$\hat{f}_\alpha(\omega_k) = \hat{f}(\omega) \cdot \alpha_k \cdot \delta(\omega - \omega_k), \quad k = 0, \dots, N - 1$$




Fig. 2.5. Variation of Δ_p with N

2. Gridding: To be able to use the FFT algorithm for reconstruction, we need to have data on equispaced nodes; however, the above function is a discrete set of scaled Dirac delta measures at non-equispaced nodes. To overcome this difficulty, we convolve \hat{f}_α with a specifically chosen interpolating function $\hat{\phi}$ to yield a continuous function \hat{g} , which can then be sampled at the desired equispaced nodes,

$$\hat{g}(k) = (\hat{f}_\alpha * \hat{\phi})(k) \approx \sum_{m \text{ st. } |k-\omega_m| \leq q} \hat{f}_\alpha(\omega_m) \hat{\phi}(k - \omega_m), \quad k \in \mathbb{Z}$$

The properties of $\hat{\phi}$ for efficient computation of the above sum have already been enumerated while discussing the NEFFT (§2.1). Figure 2.7(b) illustrates gridding with a Gaussian interpolating function, and shows evaluation of $(\hat{f}_\alpha * \hat{\phi})(k) \Big|_{k=-10}$. The discrete measurements $\hat{f}_\alpha(\omega_k)$ and the continuous interpolating function $\hat{\phi}$ can be observed. From the decay of $\hat{\phi}$, it is seen that a small number of terms is sufficient to approximate the convolution sum.

3. FFT Reconstruction: This is followed by reconstruction of g using the standard FFT algorithm⁴.

$$g(x_p) = \sum_{k=-N/2}^{N/2-1} \hat{g}(k) e^{ikx_p}, \quad x_p = -\pi + \frac{2\pi p}{N}, p = 0, \dots, N-1$$

4. Compensation: Finally, to recover f , we compensate for the Fourier space convolution by dividing out the corresponding transform (Figure 2.7(d)). We note that for the Pipe-Menon density compensation

⁴ g can be theoretically recovered $\forall x \in [-\pi, \pi]$; however, use of the FFT algorithm necessitates reconstruction at discrete grid points x_p .

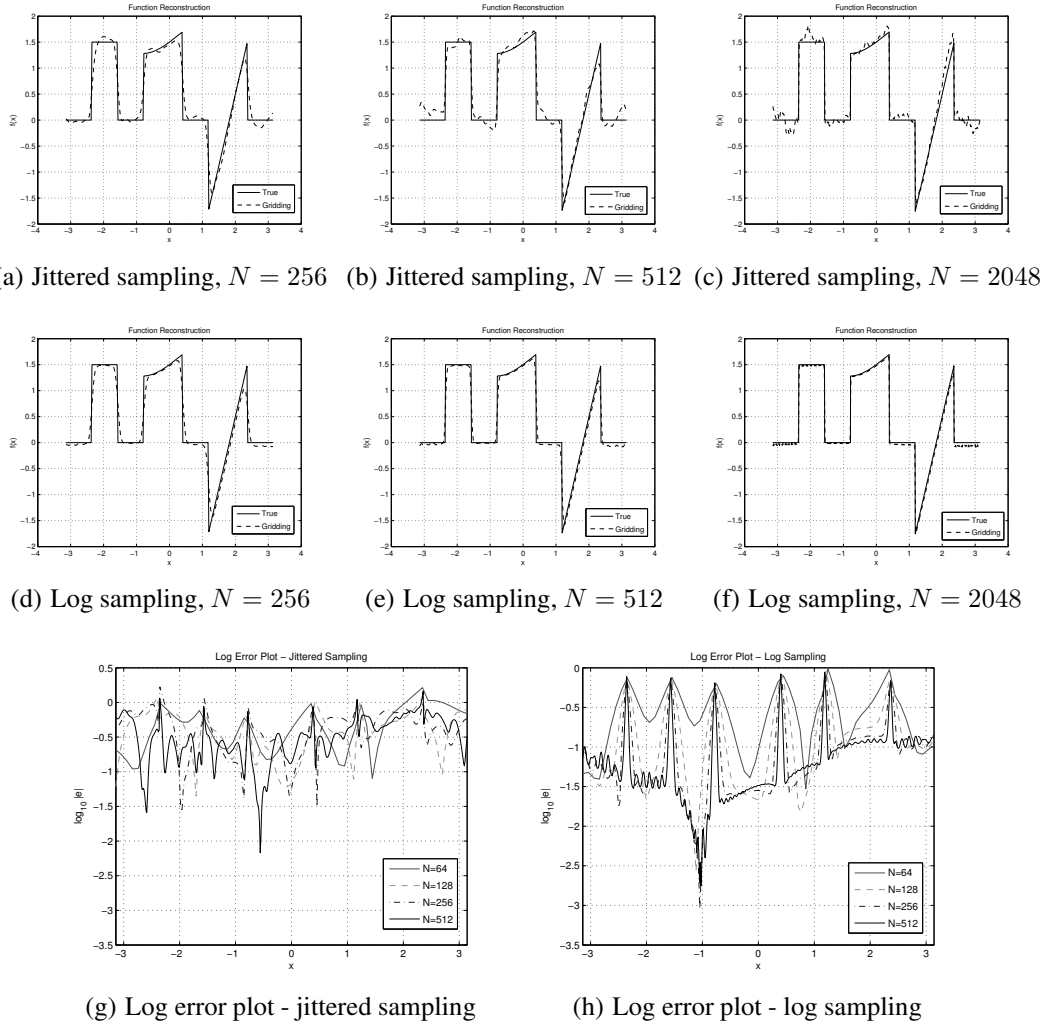


Fig. 2.6. Convergence behavior of gridding reconstruction for different sampling schemes

method, this step is not required.

$$f(x_p) \approx \frac{g(x_p)}{\phi(x_p)}, \quad x_p = -\pi + \frac{2\pi p}{N}, p = 0, \dots, N-1$$

Figure 2.7(d) also highlights a possible issue with compensation. Given that $\phi(x) \rightarrow 0$ as $|x| \rightarrow \pi$, it is possible for the reconstruction to suffer from artifacts in the vicinity of $|x| = \pi$. These artifacts result from dividing by a very small number (ϕ can be of the order of machine precision at $|x| \approx \pi$). The traditional solution to this is oversampling, wherein g is reconstructed in the domain $x = [-c\pi, c\pi)$, where $c > 1$ is an oversampling factor. In this case, $\phi(x/c) \rightarrow 0$ as $|x/c| \rightarrow \pi$. No artifacts due to compensation occur in $[-\pi, \pi]$, where our function exists, and the reconstruction for $|x| > \pi$ can be safely ignored. Figure

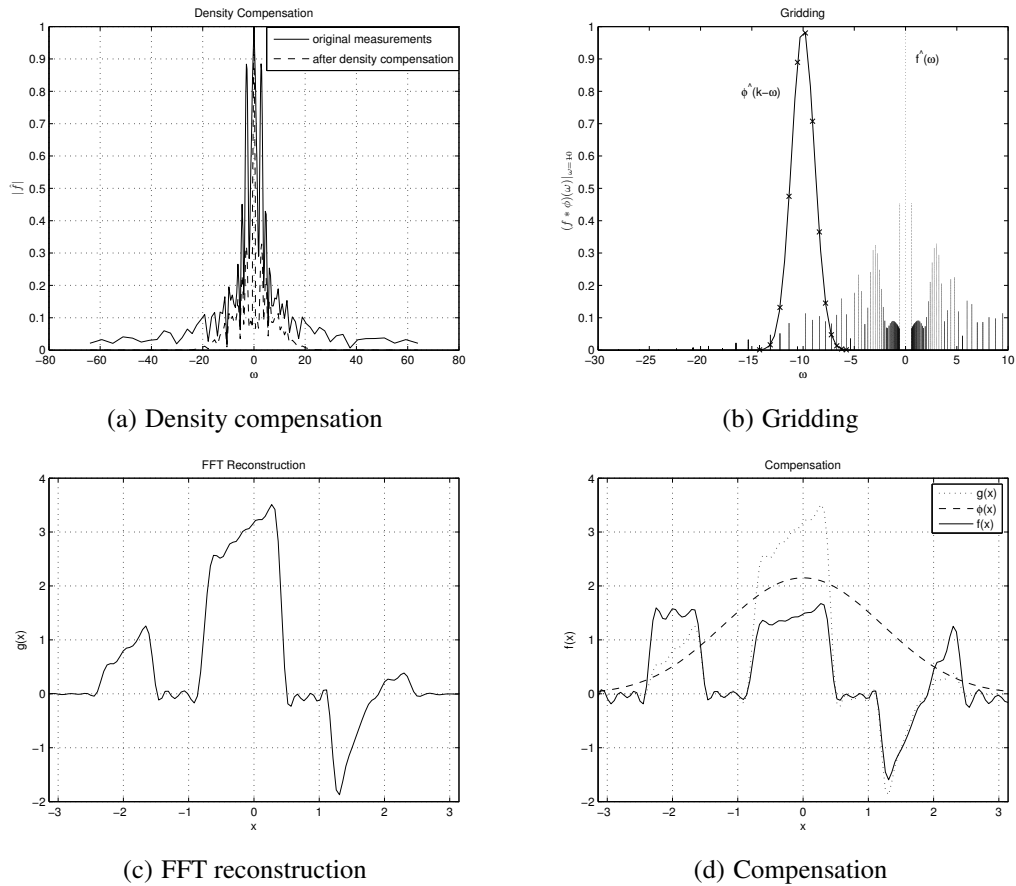


Fig. 2.7. Illustration of convolutional gridding, $N = 128$

2.8 provides a comparison of reconstructions with and without oversampling. The reconstruction without oversampling can be seen to have “winging” artifacts in the vicinity of $|x| = \pi$. While oversampling does improve the quality of reconstruction, the price paid is an increase in computational cost to $\mathcal{O}(cN \log cN + c\mathcal{M})$.

2.3. Uniform Re-Sampling (URS)

The uniform re-sampling algorithm, [23], was first introduced by Rosenfeld in 1998, and improved in further articles (e.g. [24]) by Rosenfeld and others. Reconstruction by this method involves two steps:

1. recovery of Fourier coefficients at equispaced nodes
2. partial Fourier sum reconstruction using the FFT algorithm

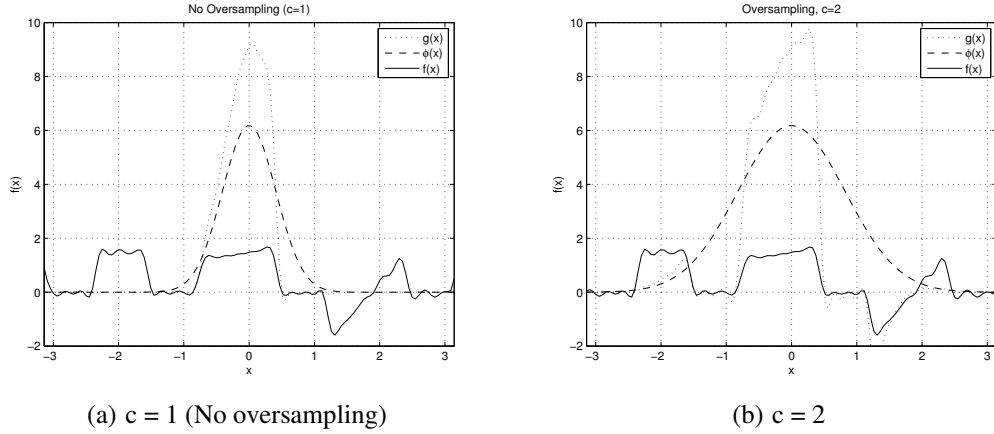


Fig. 2.8. Comparison of reconstruction with and without oversampling, $N = 128$

To recover the equispaced Fourier coefficients, we employ the WKS sampling theorem. Since $f(x)$ is compactly supported in $[-\pi, \pi]$, the Fourier transform $\hat{f}(\omega)$ may be reconstructed at any point $\omega \in \mathbb{R}$ from the discrete, equispaced samples $\hat{f}(k)$ using

$$\hat{f}(\omega) = \sum_{k=-\infty}^{\infty} \hat{f}(k) \text{sinc}(\omega - k), \quad \omega \in \mathbb{R} \quad (2.16)$$

Let $y = \left(\hat{f}(m) \right)_{m=-N/2}^{N/2-1}$ denote the equispaced Fourier coefficients we seek. If $b = \left(\hat{f}(\omega_k) \right)_{k=0}^{M-1}$ are the set of measurements and A is the matrix with elements $a_{ij} = \text{sinc}(\omega_i - j)$, $i = 0, \dots, M-1$, $j = -N/2, \dots, N/2-1$, we may write

$$Ay = b \quad (2.17)$$

Since y is a finite-length vector, the inverse problem associated with (2.17) is inherently ill-posed. We therefore proceed by computing the least-squares, minimum norm solution \tilde{y} , given by

$$\tilde{y} = A^+ b \quad (2.18)$$

where A^+ is the Moore-Penrose pseudoinverse, [25], computed using the singular value decomposition. Further, depending on the sampling scheme, A may have a high condition number, necessitating use of regularization, [26], in computing the solution. Depending on the problem requirements, truncated SVD, Tikhonov regularization, or an iterative method may be used⁵. Having computed the equispaced Fourier

⁵Error analysis for this method is provided in Chapter 3.

coefficients, we may then proceed to reconstruct f using a Fourier partial sum. Reconstruction using this method is illustrated in Figure 2.9. Figure 2.9(a) plots the recovered, equispaced Fourier coefficients, while Figure 2.9(b) plots the corresponding error in recovering these coefficients. The function reconstruction, after post-processing using a second-order exponential filter, is shown in Figure 2.9(c).

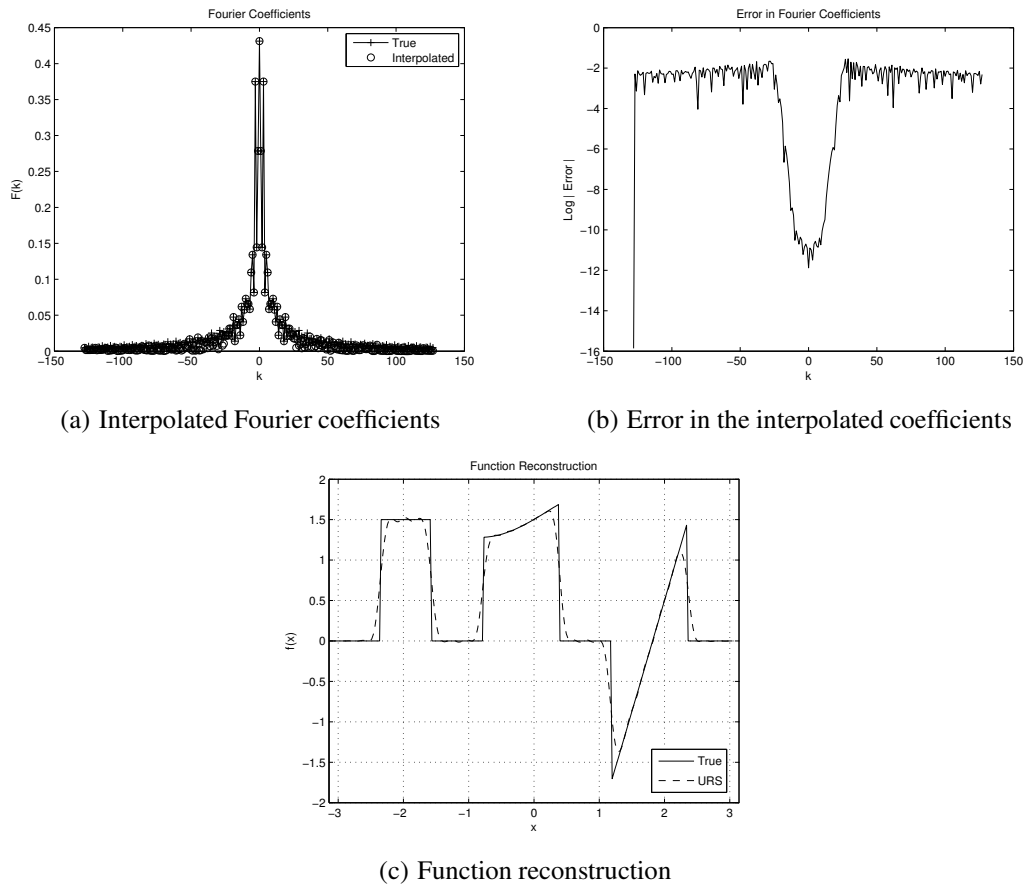


Fig. 2.9. Reconstruction using URS, $N = 256$

2.3.1. Implementation Notes

For large N and in multi-dimensional problems, computing the pseudoinverse of a large matrix becomes infeasible. Block uniform re-sampling (BURS), [23], was designed to use a block based algorithm with smaller system size and reduced computational cost. The underlying principle is as follows: Due to decay of the sinc function in (2.16), the value of $\hat{f}(\omega_i)$ is largely determined by the nearest equispaced

coefficients in the vicinity of ω_i , i.e.,

$$\hat{f}(\omega_i) \approx \sum_{k \text{ st. } |\omega_i - k| < \zeta} \hat{f}(k) \text{sinc}(\omega - k)$$

where ζ is a neighbourhood parameter. Bounds on the error caused by ignoring other terms (truncation error) can be found, for example, in [27]. A set of measurements $\hat{f}(\omega_i), i = 0, \dots, P - 1$, in a local region of Fourier space can therefore be modelled using a smaller system of equations which can be inverted at reduced computational cost. The global solution is then obtained by integrating all the local solutions. It is to be emphasized that use of this block implementation is inherently sub-optimal, because we replace the global approximation in (2.16) by a local approximation. However, experimental results reveal that the error incurred due to block implementation is smaller, or, of the same order as that of the regularization error. Further, ease of implementation and drastic savings in computational cost provide compelling arguments to adopt this implementation.

For completeness, specifics of the algorithm are given below. We shall assume that we are given an M -vector $\mathbf{b} = \left(\hat{f}(\omega_j) \right)_{j=0}^{M-1}$ of non-equispaced Fourier samples and that we seek an N -vector y of equispaced coefficients at the grid points $\mathbf{k} = (k_l)_{l=0}^{N-1}$. For multi-dimensional problems, we shall assume that vectorization of data has been performed. We are looking to construct an $N \times M$ matrix A^+ , with its entries representing interpolation coefficients. For each equispaced node $k_l, l = 0, \dots, N - 1$, the following is performed,

1. The P closest measurements to node k_l , denoted by β are located

$$\{\beta = \mathbf{b} \mid |\omega_j - k_l| < \delta\}$$

2. The Q closest equispaced nodes to k_l , denoted by κ are located

$$\{\kappa = \mathbf{k} \mid |\mathbf{k} - k_l| < \Delta\}$$

3. The $P \times Q$ reduced sinc matrix \mathcal{A} is constructed from the P non-equispaced nodes and Q equispaced nodes

$$\mathcal{A} = (\text{sinc}(\omega_j - \kappa))$$

4. Let γ represent the equispaced Fourier coefficients. Then $\mathcal{A}\gamma \approx \beta$ models a local region of Fourier space. The corresponding solution is $\gamma \approx \mathcal{A}^+\beta$, where \mathcal{A}^+ is the pseudoinverse of \mathcal{A} .
5. The row in \mathcal{A}^+ which yields the equispaced coefficient at node k_l is transferred to A^+ , ensuring that appropriate columns of A^+ are filled.

Here, δ and Δ are neighborhood parameters selected empirically. From formulation of the sampling theorem, it is obvious that Δ has to be greater than δ . Results in [23] indicate that $\Delta \approx 1.5\delta$ provides the best results, with no significant improvement in performance observed by increasing Δ beyond 1.5δ . Figure 2.10 illustrates the selection of β and κ for a two dimensional problem with spiral sampling. Equispaced coefficient recovery at the node $(-2, 0)$ (denoted by a \times) is highlighted. The circles denote elements of the non-equispaced measurement vector β , while the squares denote elements of the equispaced node vector κ . Here values of $P = 5$ and $Q = 10$ were used.

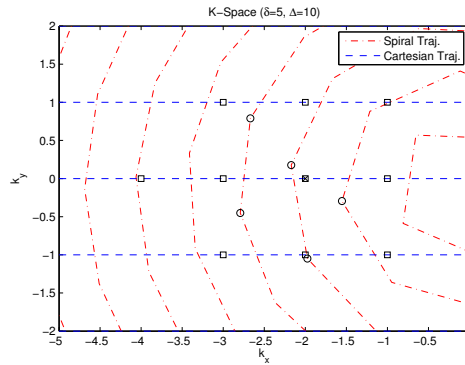


Fig. 2.10. Illustration: block uniform re-sampling (BURS)

After constructing the matrix A^+ , we may compute the approximate solution as

$$y^* = A^+b$$

We note that computation of A^+ only depends on the sampling trajectory, and can therefore be performed offline. Further, the matrix A^+ has only P non zero entries in each row, resulting in an interpolation cost of NP . Adding the cost of standard FFT reconstruction, we require a total of $N \log N + NP$ computations to perform BURS reconstruction.

2.4. Iterative Methods

The class of iterative methods [28, 29] is based on discretization of the Fourier integral (1.1). These methods recover the function at equispaced grid points $x_k = -\pi + \frac{2\pi k}{N}, k = 0, \dots, N-1$. Given M arbitrary spectral samples, $\hat{f}(\omega_k), k = 0, \dots, M-1$ with ω_k not necessarily integers, we may write

$$\hat{f}(\omega_p) \approx \sum_{k=0}^{N-1} f(x_k) e^{-i\omega_p x_k}, \quad p = 0, \dots, M-1 \quad (2.19)$$

We note that x_k are equispaced grid points, and hence this is a valid quadrature for the Fourier integral, [3]. Further, since $\hat{f}(\omega)$ is a smooth function, the partial sum approximates $\hat{f}(\omega_p)$ with exponential accuracy. However, to recover f , we have to invert the system of equations described by (2.19). For this reason, this formulation is sometimes referred to as the ‘‘inverse’’ method or implicit discretization. In matrix notation, we may write

$$b \approx Ay \quad (2.20)$$

where,

$$b = \left(\hat{f}(\omega_j) \right)_{j=0}^{M-1}$$

$$y = \left(f(x_k) \right)_{k=0}^{N-1}, \quad x_k = -\pi + \frac{2\pi k}{N}$$

$$A = \left(e^{-i\omega_j x_k} \right), \quad j = 0, \dots, M-1; \quad x_k = -\pi + \frac{2\pi k}{N}, \quad k = 0, \dots, N-1$$

As a result of using finite N , this system is inherently ill-posed and we compute a least-squares, minimum norm solution. Further, it is typical to have more non-equispaced samples than equispaced nodes, resulting in an overdetermined system. Since the system matrix, A can be badly conditioned depending on sampling trajectory, the above system of equations is solved using an iterative method such as conjugate gradient [25] (CG) or one of its many variants. Since CG requires a symmetric system matrix, it is typically applied on the normal form

$$A^T b \approx A^T A y$$

Each iteration of the CG method requires a matrix-vector multiplication, resulting in prohibitively large computational cost. Since ω_j are not equispaced, the FFT algorithm cannot be applied; however, Formulation (1) of the NEFFT, (§2.1), which is only marginally more expensive than a standard FFT, can be employed. Function recovery using this method is illustrated in Figure 2.11. All reconstructions have been processed with a

second-order exponential filter. In general, the reconstruction is seen to improve as the number of iterations increase. However, beyond a certain number of iterations (8 in the example), no significant improvement in reconstruction quality is seen.

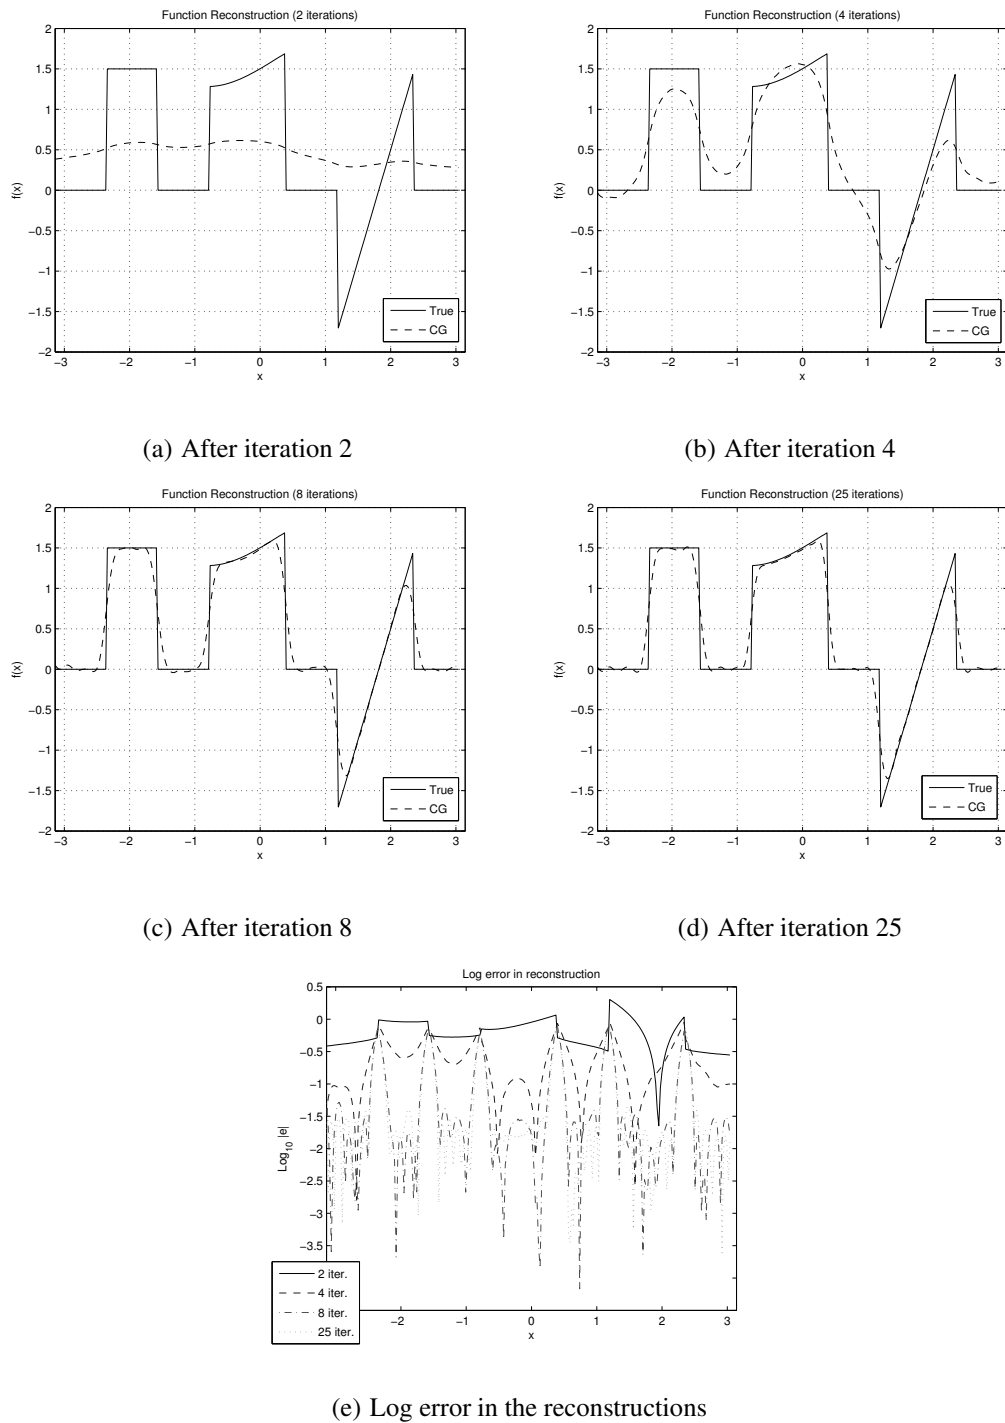


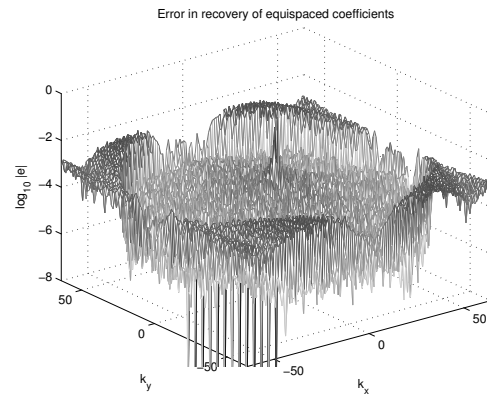
Fig. 2.11. Iterative reconstruction computed using the conjugate-gradient method

To illustrate the applicability to MR image reconstruction, Shepp-Logan phantom reconstruction using BURS from spectral data acquired along a spiral k -space trajectory is shown in Figure 2.12. Spectral samples

of the Shepp-Logan phantom were computed analytically using the method outlined in [30]. The spiral sampling scheme used to obtain the samples is detailed in Appendix A. 65536 spiral samples were acquired and reconstruction was performed on a 128×128 grid. The BURS reconstruction was computed using a precomputed pseudoinverse (calculated using the TSVD) and the following parameters: $\delta = 2$, $\Delta = 5$, $P = 20$.



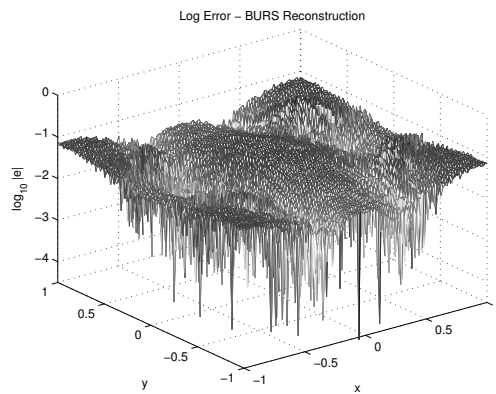
(a) Fourier reconstruction from Cartesian samples



(b) Error in recovery of equispaced k -space samples



(c) BURS reconstruction



(d) Error in BURS reconstruction

Fig. 2.12. Shepp-Logan phantom reconstruction from spiral k -space samples using BURS

TABLE 2.1
COMPARISON OF RECONSTRUCTION METHODS

| | Governing Relation | Cost | Remarks |
|-------------------------------|--|---|---|
| Gridding | $f(x) \approx \sum_{k=0}^{N-1} \alpha_k \hat{f}(\omega_k) e^{i\omega_k x}$ <p>Approximates the inverse Fourier integral α_k are DCFs Compute sum using the NEFFT</p> | $\mathcal{O}(cN \log cN + cM)$ c is an oversampling factor M is a small multiple of N | Accuracy of reconstruction depends on sampling scheme and α_k |
| URS & BURS | <p>Obtain equispaced coefficients $\hat{f}(k)$ as the least-squares solution of</p> $\hat{f}(\omega_p) = \sum_{k=-N/2}^{N/2-1} \hat{f}(k) \text{sinc}(\omega_p - k)$ <p>Reconstruct f using standard FFT For large problems, obtain global solution as union of local solutions, with each local system being of (reduced) order $P \times Q$, ($Q > P$)</p> | $\mathcal{O}(N \log N + NP)$ | System can be solved using iterative methods or computing the pseudoinverse Pseudoinverse methods with block implementation allow for reduced computational cost on account of precomputation of the pseudoinverse Use of regularization is recommended |
| Iterative Solution | $\hat{f}(\omega_p) = \sum_{k=0}^{N-1} f(x_k) e^{-i\omega_p x_k}$ <p>Obtain f as the least-squares solution to the above system Typically solved using conjugate gradient</p> | $\mathcal{O}(\kappa(N \log N + M))$ κ is the iteration count | Only recovers f at grid points |

3. ANALYSIS OF UNIFORM RE-SAMPLING AND HIGH RESOLUTION RECONSTRUCTION

The uniform re-sampling scheme (URS) provides a theoretical framework to relate non-Cartesian measurements $\hat{f}(\omega_k)$ to the equispaced Fourier coefficients $\hat{f}(k)$. URS reconstruction is typically carried out in two stages,

1. recovery of equispaced Fourier coefficients $\hat{f}(k)$
2. recovery of f using partial Fourier reconstruction

To summarize, we have the following governing relation as a result of the WKS sampling theorem,

$$\hat{f}(\omega) = \sum_{k=-\infty}^{\infty} \hat{f}(k) \text{sinc}(\omega - k), \quad \omega \in \mathbb{R} \quad (3.1)$$

Here, $\hat{f}(\omega)$ is the Fourier transform and $\hat{f}(k)$ are the Fourier series coefficients. Let $y = \left(\hat{f}(m) \right)_{m=-N/2}^{N/2-1}$ denote the equispaced Fourier coefficients we seek. If $b = \left(\hat{f}(\omega_k) \right)_{k=0}^{M-1}$ is the set of measurements and A is the matrix with elements $a_{ij} = \text{sinc}(\omega_i - j)$, $i = 0, \dots, M-1$, $j = -N/2, \dots, N/2-1$, we may write

$$Ay \approx b \quad (3.2)$$

The least-squares minimum norm solution y to (3.2) is given by

$$y_{LS} = A^+ b \quad (3.3)$$

where A^+ is the Moore-Penrose pseudoinverse, [25]. Using the singular value expansion, [25], we may write

$$A^+ = \sum_{p=0}^{r-1} \frac{v_p u_p^H}{\sigma_p}$$

where u_p and v_p are the left singular vectors and right singular vectors respectively of A , σ_p are the singular values of A and r is the rank of A . Therefore, the least-squares minimum norm solution is given by

$$y_{LS} = \sum_{p=0}^{r-1} \left(\frac{u_p^H b}{\sigma_p} \right) v_p \quad (3.4)$$

Further, it is possible to show (e.g., [25]) that the norm of the error between the true solution and the least-squares solution is given by,

$$\|Ay_{LS} - b\|_2^2 = \sum_{p=r}^N |u_p^H b|^2 \quad (3.5)$$

URS allows us to reconstruct f in the entire domain $x \in [-\pi, \pi)$ and allows for efficient implementation in the case of precomputed interpolation coefficients and block implementation.


3.1. Error Analysis of URS Reconstruction

There are two principal error components in URS reconstruction,

1. error in recovering equispaced Fourier coefficients
2. error in partial Fourier reconstruction

The latter error component is common to all reconstruction schemes and is treated in (§3.2). To account for the former, we shall assume that the system of equations in (3.2) is inverted using the truncated SVD (TSVD), [25], or Tikhonov regularization, [31]. Regularized solutions such as those mentioned above are necessary because the system described by (3.2) is numerically ill-posed for many sampling schemes of interest. Susceptibility to finite precision effects is accentuated in the presence of measurement noise, leading to spurious solutions.

As an illustration of the need for regularization, consider the system matrix for the “log sampling” scheme, which results in a rank deficient matrix with singular values as shown in Figure 3.1(a). It can be observed that roughly half the singular values are of the order of machine precision. A solution computed using no regularization is shown in Figure 3.1(b). The true function is the test example of Chapter §2 (Figure 2.1(a)).



(a) Singular values of A for the log sampling scheme, $N = 256$

(b) Solution with no regularization

Fig. 3.1. Need for regularization

In such cases, it is preferable to compute the solution using a regularized pseudoinverse, taking the following form

$$A_{reg}^+ = \sum_{p=0}^{r-1} \phi_p \left(\frac{v_p u_p^H}{\sigma_p} \right) \quad (3.6)$$

Here ϕ is a regularization function specific to each method. For the truncated SVD, ϕ takes the form

$$\phi_p = \begin{cases} 1 & p = 0, \dots, r-1 \\ 0 & \text{otherwise} \end{cases} \quad (3.7)$$

Ideally, r is the true rank of the matrix A . For numerically ill-posed systems, determining r is a considerable challenge. One solution is to choose r such that $u_p^H b$ is of the same order as σ_p . This bounds the ratio $\frac{v_p u_p^H}{\sigma_p}$ and ensures that spurious solutions do not result. Figure 3.2 plots σ_p and $u_p^H b$ for the log sampling scheme. The plot indicates that choosing $r \approx 110$, which corresponds to a threshold of 10^{-3} would prevent solution blowup.

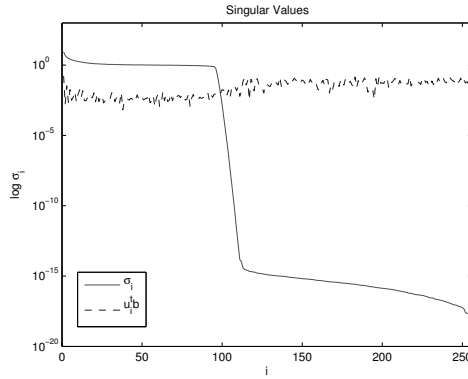


Fig. 3.2. Plot of σ_p and $u_p^H b$ for log sampling, $N = 256$

For Tikhonov regularization [31], we have

$$\phi_p = \frac{\sigma_p^2}{\sigma_p^2 + \alpha^2} \quad (3.8)$$

where, α is a regularization parameter. Selection of optimum α is an open research problem, with several popular methods including the discrepancy principle [26], L-curve [32] and generalized cross validation (GCV) [33].

Let μ be the measurement error in acquiring spectral samples, and b^* the acquired measurements. Then,

$$Ay = b^* = b + \mu \quad (3.9)$$

Consequently, the computed solution is

$$\begin{aligned}
\tilde{y} &= A_{reg}^+ b^* \\
&= \sum_{p=0}^{r-1} \phi_p \left(\frac{v_p u_p^H}{\sigma_p} \right) (b + \mu) \\
\tilde{y} &= \sum_{p=0}^{r-1} \phi_p \left(\frac{u_p^H b}{\sigma_p} \right) v_p + \sum_{p=0}^{r-1} \phi_p \left(\frac{u_p^H \mu}{\sigma_p} \right) v_p
\end{aligned} \tag{3.10}$$

Let e be the error between the computed solution and the least-squares solution of (3.3). Then,

$$\begin{aligned}
e &= y - \tilde{y} \\
&= \sum_{p=0}^{r-1} \left(\frac{u_p^H b}{\sigma_p} \right) v_p - \sum_{p=0}^{r-1} \phi_p \left(\frac{u_p^H b}{\sigma_p} \right) v_p - \sum_{p=0}^{r-1} \phi_p \left(\frac{u_p^H \mu}{\sigma_p} \right) v_p \\
e &= \sum_{p=0}^{r-1} (1 - \phi_p) \left(\frac{u_p^H b}{\sigma_p} \right) v_p - \sum_{p=0}^{r-1} \phi_p \left(\frac{u_p^H \mu}{\sigma_p} \right) v_p
\end{aligned} \tag{3.11}$$

The first term is called the *regularization error*, while the second term is called the *perturbation error*. Although performance in the presence of measurement error is important, we are primarily concerned with the implicit error introduced by non-Cartesian spectral sampling. This manifests as the regularization error in (3.11). As an illustration, consider reconstruction with no measurement error ($\mu = 0$) using the following schemes¹: jittered sampling and log sampling. Jittered sampling results in a system matrix which is full rank and well conditioned ($\kappa = \mathcal{O}(1)$). In this case, minimal or no regularization is required. Log sampling, however, results in a rank-deficient, highly ill-conditioned system matrix ($\kappa = \mathcal{O}(10^{19})$). Figure 3.3 shows error in recovering the solution² for either case, when computed using the TSVD. For jittered sampling, a singular value cut-off equal to machine precision was used, while for log sampling, the singular value cut-off was set to 10^{-4} . It is observed that for the case of log sampling, high frequency coefficients are recovered with large error, in agreement with the low sampling density in these regions.

The recovery of Fourier coefficients with differing error orders is a characteristic of all sampling schemes with variable sampling density. Moreover, regions with undersampling yield interpolated coefficients with large error. For example, consider recovery of equispaced k -space coefficients from spiral and radial sampling trajectories. Figure 3.4 illustrates the error in recovered coefficients when the underlying

¹These schemes are enumerated in Appendix A.

²We note that these error plots show the error between the regularized solution and the true solution (as opposed to the least-squares solution).

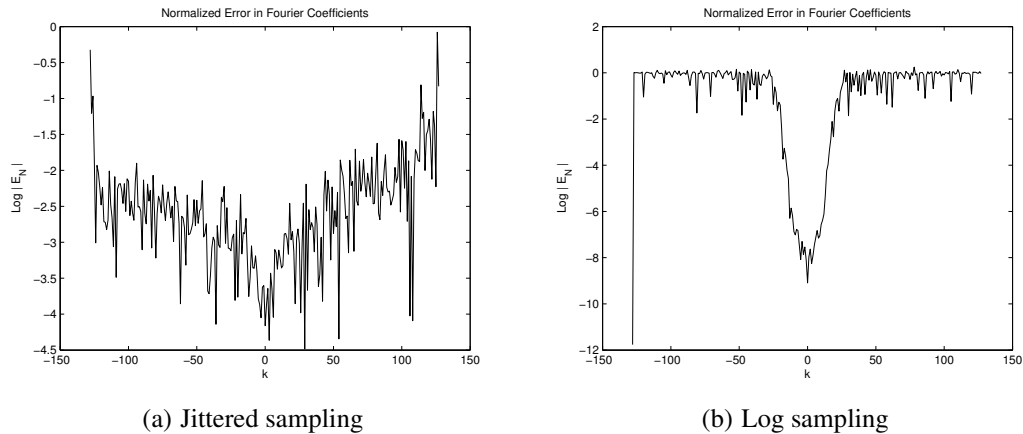


Fig. 3.3. Error in recovered coefficients, $N = 256$

function is a Shepp-Logan phantom. Here, 128×128 Cartesian coefficients were recovered from 65536 non-Cartesian measurements. It is conceivable that by formulating a reconstruction method which takes into

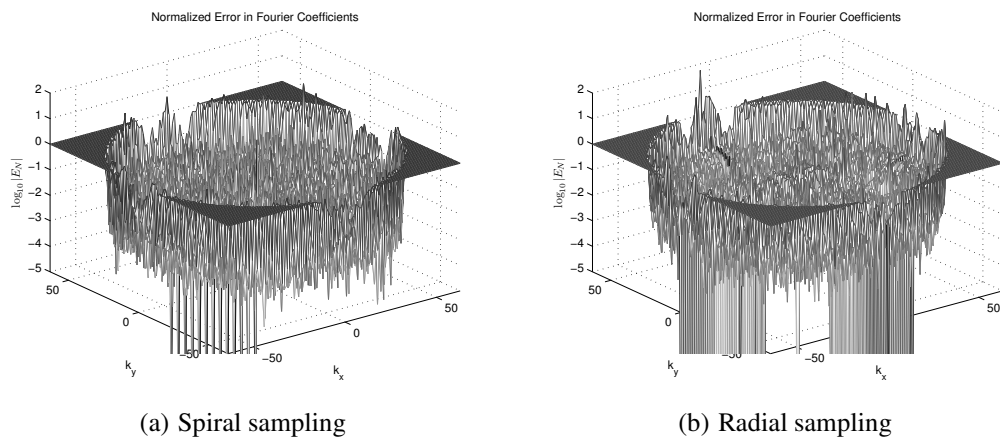


Fig. 3.4. Error in recovered Fourier coefficients of the Shepp-Logan phantom

account the differing accuracy among spectral coefficients, better reconstruction quality can be achieved.

3.2. High Resolution Reconstruction

Given the popularity of sampling schemes such as radial and spiral sampling, it is typical to perform reconstruction with a set of Fourier coefficients characterized by high accuracy at lower frequencies and large error at higher frequencies. Partial Fourier reconstruction without any post-processing yields inaccurate reconstructions characterized by high-frequency “noise-like” artifacts. The conventional solution is to perform

filtering which damps the high frequency coefficients. Figure 3.5 shows the resulting reconstruction of our test function whose Fourier samples have been acquired using the log sampling scheme. The use of large N in a filtered Fourier reconstruction is therefore open to question, since a large number of those coefficients are essentially damped by use of the filter. However, use of small N results in poor order of accuracy, which is expected when Fourier and filtered Fourier methods are used on piecewise-smooth functions.

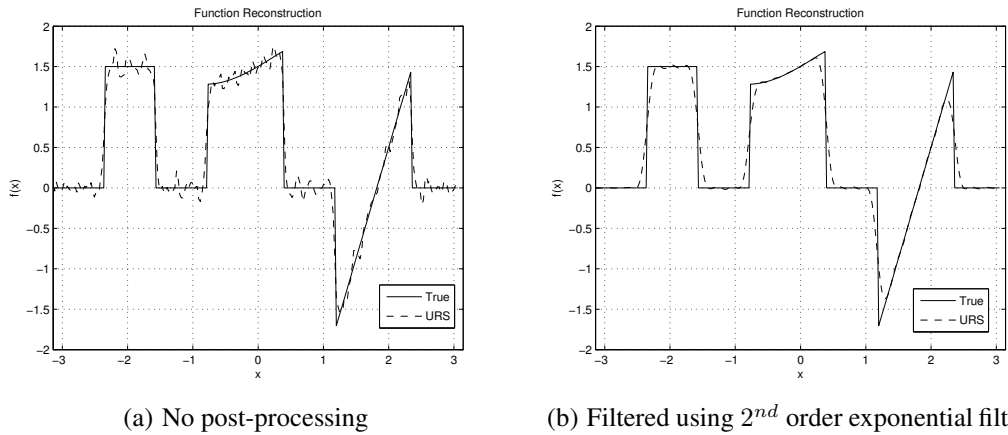


Fig. 3.5. Test function reconstruction (log sampling), $N = 256$

To resolve this apparent quandary, use of a spectral reconstruction scheme such as Gegenbauer reconstruction [5] is suggested. By formulation, these schemes require accurate discontinuity data. Therefore, the detection of signal discontinuities is elaborately treated in Chapter 4. A brief outline of the intuition behind use of Gegenbauer reconstruction is provided here, while the reader is referred to [5, 34] for detailed treatment and implementation notes. Spectral reconstruction schemes were primarily introduced to overcome the Gibbs phenomenon, and involve approximating the function in an alternate basis, known as a Gibbs complementary basis [5]. Approximation is performed in each smooth sub-interval of the function. A Gibbs complementary basis is orthonormal, supports exponential reconstruction accuracy, and has the property that high modes of the Fourier basis have exponentially small projections on low modes of the reconstruction basis. This is exactly our requirement, since we want to minimize the impact of error-filled high frequency Fourier coefficients on our reconstruction. The Gegenbauer reconstruction is an example of a spectral reconstruction scheme, which has been well studied and possesses computationally attractive properties.

3.2.1. Gegenbauer Reconstruction

The Gegenbauer polynomials $C_n^\lambda(x)$ are orthogonal polynomials in $[-1, 1]$ which satisfy

$$\int_{-1}^1 (1-x^2)^{\lambda-\frac{1}{2}} C_k^\lambda(x) C_n^\lambda(x) dx = \begin{cases} h_n^\lambda & k = n \\ 0 & k \neq n \end{cases} \quad (3.12)$$

where (for $\lambda \geq 0$)

$$h_n^\lambda = \frac{\sqrt{\pi} C_n^\lambda(1) \Gamma(\lambda + \frac{1}{2})}{\Gamma(\lambda)(n + \lambda)}$$

with

$$C_n^\lambda(1) = \frac{\Gamma(n + 2\lambda)}{n! \Gamma(2\lambda)}$$

Gegenbauer polynomials obey the recursive relation

$$(n + 1)C_{n+1}^\lambda(x) = 2(n + \lambda)x C_n^\lambda(x) - (n + 2\lambda - 1)C_{n-1}^\lambda(x) \quad (3.13)$$

with $C_0^\lambda(x) = 1$ and $C_1^\lambda(x) = 2\lambda x$.

The Gegenbauer approximation of a function f is given by

$$g_m^\lambda(x) = \sum_{l=0}^m \hat{g}^\lambda(l) C_l^\lambda(x), \quad x \in [-1, 1] \quad (3.14)$$

where \hat{g}^λ are the Gegenbauer expansion coefficients, obtained using the weighted inner product,

$$\tilde{g}^\lambda(l) = \frac{1}{h_l^\lambda} \int_{-1}^1 (1-\eta^2)^{\lambda-\frac{1}{2}} f(\eta) C_l^\lambda(\eta) d\eta \quad (3.15)$$

It has been shown in [5] that f can be replaced by its Fourier expansion in the above expression, while still retaining exponential reconstruction accuracy.

Consider reconstruction of $f(x)$ in a smooth interval $[a, b] \subset [-1, 1]$. We define $\eta \in [-1, 1]$ and $x(\eta) = \epsilon\eta + \delta$ where $\epsilon = \frac{b-a}{2}$ and $\delta = \frac{b+a}{2}$. The Gegenbauer reconstruction may be computed as

$$g_m^\lambda(x(\eta)) = \sum_{l=0}^m \hat{g}_\epsilon^\lambda(l) C_l^\lambda(\eta) \quad (3.16)$$

where the approximate Gegenbauer coefficients $\hat{g}_\epsilon^\lambda(l)$ are given by

$$\hat{g}_\epsilon^\lambda(l) = \frac{1}{h_l^\lambda} \int_{-1}^1 (1-\eta^2)^{\lambda-\frac{1}{2}} f_N(x(\eta)) C_l^\lambda(\eta) d\eta \quad (3.17)$$

Here, $f_N(\cdot)$ is the partial Fourier reconstruction of $f(x)$. [34] shows that these coefficients may be computed using the following Fourier-Bessel sum

$$\hat{g}_\epsilon^\lambda(l) = \delta_{0l} \tilde{f}(0) + \Gamma(\lambda) i^l (l + \lambda) \sum_{0 < |k| \leq N} J_{(l+\lambda)}(\pi k \epsilon) \left(\frac{2}{\pi k \epsilon} \right)^\lambda e^{i k \pi \delta} \tilde{f}(k) \quad (3.18)$$

Further, (3.18) can be computed using the FFT algorithm, leading to considerable computational savings. Selection of m and λ to achieve stable reconstruction is addressed in [35].

3.2.2. Results using Gegenbauer Reconstruction

A comparison of filtered Fourier and Gegenbauer reconstructions is shown in Figure 3.6. The test function was reconstructed from Fourier samples acquired using log sampling. The URS method was used to recover equispaced spectral coefficients. In addition to supporting greater reconstruction accuracy in smooth regions, the Gegenbauer reconstruction is free of smearing artifacts in the neighborhood of jump discontinuities. This is important in MR reconstruction since important features may be contained at the edges. Further, we note that the Gegenbauer reconstruction in Figure 3.6 uses a quarter of the coefficients of filtered Fourier reconstruction.

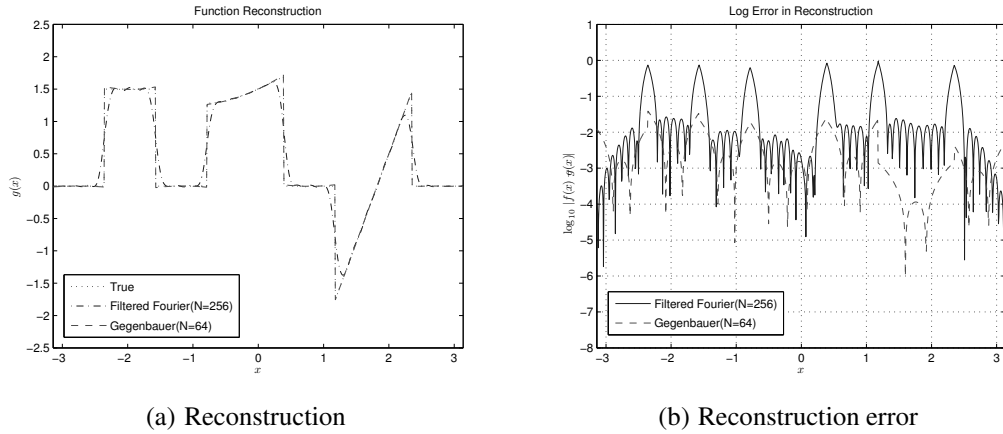


Fig. 3.6. Comparison of filtered Fourier ($N = 256$) and Gegenbauer reconstructions ($N = 64$, $m = \lambda = 2$)

As mentioned in (§3.2), Gegenbauer reconstruction requires information about the location of jump discontinuities in f . In Figure 3.6, it was assumed that this information was available. Figure 3.7 shows the resulting reconstruction when jump locations are estimated from the available spectral data. In this example, $N = 64$ equispaced coefficients recovered using the URS method were used to compute the locations of jump

discontinuities in f , using variants of the concentration method detailed in [36,37]. Compared to Figure 3.6, reconstruction using estimated jump locations shows a small reduction in accuracy in the immediate vicinity of a jump discontinuity. However, the performance is still superior to that of filtered Fourier reconstruction.

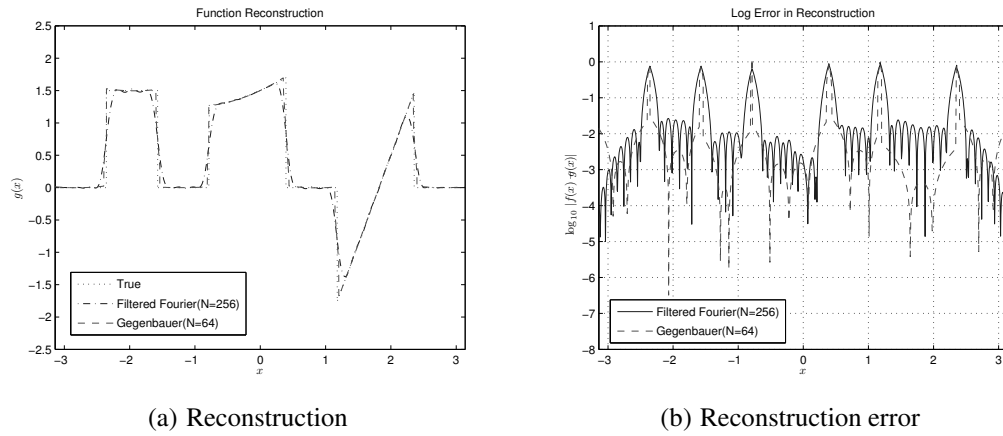


Fig. 3.7: Comparison of filtered Fourier ($N = 256$) and Gegenbauer reconstructions ($N = 64$, $m = 4$, $\lambda = 4$) with jump locations estimated using the concentration method

4. DISCONTINUITY DETECTION FROM SPECTRAL DATA

This chapter addresses the detection of jump discontinuities in a piecewise-smooth function from noisy Fourier data. While detecting jump discontinuities is an important task in its own right, it has been shown in Chapter §3 that knowledge of jump discontinuities is essential for high-resolution function recovery. The concentration method [36, 38] of detecting jump discontinuities from spectral data is employed here. After providing an introduction to this method, its statistical behaviour is studied under the assumption that zero mean additive complex white Gaussian noise corrupts the spectral samples. A method to improve detector performance using a combination of *concentration factors*, first introduced in [39], is further refined, along with the formulation of a performance measure. Examples and receiver operating characteristic (ROC) curves are provided to demonstrate the efficacy of this method. While the following development assumes a one-dimensional function, remarks and examples on its extension to multiple dimensions are provided.

Let f be a 2π -periodic, piecewise-smooth function in $[-\pi, \pi)$. If the right-hand and left-hand limits of f , denoted by $f(x^+)$ and $f(x^-)$ respectively, are defined at every point x , we may define the corresponding *jump function* $[f]$ as

$$[f](x) := f(x^+) - f(x^-) \quad (4.1)$$

We shall concern ourselves with finding an approximation to $[f]$ using a finite number of Fourier coefficients $\hat{f}(k)$ defined by

$$\hat{f}(k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx} dx, \quad k \in [-N, N] \quad (4.2)$$

From (4.1), we see that a jump discontinuity is an essentially local feature; i.e., the jump function at any point x only depends on the values of f at x^+ and x^- . However, \hat{f} is a global representation; i.e., $\hat{f}(k)$ are obtained using values of f over the entire domain $[-\pi, \pi)$. This contrast in definitions makes obtaining a good approximation to $[f]$ from \hat{f} particularly challenging.

4.1. Concentration Method

To understand how jump information can be extracted from Fourier data, consider, without loss of generality, a single jump in f at $x = \zeta^1$. We start by integrating (4.2) by parts.

$$\hat{f}(k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx} dx$$

¹For ease of analysis, we will assume that ζ does not lie near $\pm\pi$, and that $f(\pm\pi)$ is continuous.

$$\begin{aligned}
&= \frac{1}{2\pi} \left(\int_{-\pi}^{\zeta^-} f(x)e^{-ikx} dx + \int_{\zeta^+}^{\pi} f(x)e^{-ikx} dx \right) \\
&= \frac{1}{2\pi} \left(f(x) \frac{e^{-ikx}}{-ik} \Big|_{-\pi}^{\zeta^-} - \int_{-\pi}^{\zeta^-} f'(x) \frac{e^{-ikx}}{-ik} dx + f(x) \frac{e^{-ikx}}{-ik} \Big|_{\zeta^+}^{\pi} - \int_{\zeta^+}^{\pi} f'(x) \frac{e^{-ikx}}{-ik} dx \right) \\
&= \frac{1}{2\pi} \left(\frac{f(\zeta^-)e^{-ik\zeta} - f(-\pi)e^{ik\pi}}{-ik} - \int_{-\pi}^{\zeta^-} f'(x) \frac{e^{-ikx}}{-ik} dx \right. \\
&\quad \left. + \frac{f(\pi)e^{-ik\pi} - f(\zeta^+)e^{-ik\zeta}}{-ik} - \int_{\zeta^+}^{\pi} f'(x) \frac{e^{-ikx}}{-ik} dx \right) \\
&= (f(\zeta^+) - f(\zeta^-)) \frac{e^{-ik\zeta}}{2\pi ik} + \frac{f(-\pi)e^{ik\pi} - f(\pi)e^{-ik\pi}}{2\pi ik} + \mathcal{O}\left(\frac{1}{k^2}\right)
\end{aligned}$$

Since f is periodic, $f(-\pi) = f(\pi)$ and the second term vanishes. Using (4.1), we obtain

$$\hat{f}(k) = [f](\zeta) \frac{e^{-ik\zeta}}{2\pi ik} + \mathcal{O}\left(\frac{1}{k^2}\right) \quad (4.3)$$

As can be observed, location of the jump discontinuity is contained in the phase of the leading term, while the same term also contains the jump magnitude and sign. By suitable design, we can therefore expect to arrive at an approximation to the jump function. The concentration method for jump detection, [36, 38], provides us with exactly this framework, by obtaining jump location, magnitude and sign from a finite set of Fourier coefficients. This method is quite general in formulation – in its most abstract form, we may interpret it as the use of specifically designed convolutional kernels to extract jump data. However, for our particular problem formulation, it is useful to interpret this method as a form of spectral filtering, with the jump function reconstructed using a partial Fourier sum. While the reader is referred to [36,38] for details, it is illustrative to provide a simple example of the theory behind this method. Consider a partial Fourier sum of the form,

$$S_N[f](x) = \sum_{k=-N}^N \left(\frac{i\pi k}{N}\right) \hat{f}(k) e^{ikx} \quad (4.4)$$

Substituting for $\hat{f}(k)$ from (4.3), we obtain

$$\begin{aligned}
S_N[f](x) &= \sum_{k=-N}^N \left(\frac{i\pi k}{N}\right) \hat{f}(k) e^{ikx} \\
&= \sum_{k=-N}^N \left(\frac{i\pi k}{N}\right) \left[[f](\zeta) \frac{e^{-ik\zeta}}{2\pi ik} + \mathcal{O}\left(\frac{1}{k^2}\right) \right] e^{ikx} \\
&= [f](\zeta) \frac{1}{2N} \sum_{k=-N}^N e^{ik(x-\zeta)} + \sum_{k=-N}^N \mathcal{O}\left(\frac{1}{k}\right) e^{ik(x-\zeta)}
\end{aligned} \quad (4.5)$$

The first term is a scaled Dirichlet kernel, localized at $x = \zeta$, i.e., it approximates a scaled Dirac delta distribution at the jump location $x = \zeta$. The second term is a manifestation of the global nature of spectral data. The concentration method gives us a family of spectral filters specifically designed to minimize this contribution. It is not difficult to show [36] that the following holds,

$$S_N[f](x) = [f](x) + \mathcal{O}\left(\frac{\log N}{N}\right) \quad (4.6)$$

Hence, a sum of the form (4.4) provides an approximation to the jump function. As $N \rightarrow \infty$, $S_N[f](x) \rightarrow [f](x)$ with a convergence rate $\mathcal{O}\left(\frac{\log N}{N}\right)$. The spectral filter used in the above example, $\left(\frac{i\pi k}{N}\right)$, is just one example among a class of filters. In its more general form, (4.4) can be written as

$$S_N^\sigma[f](x) = i \sum_{k=-N}^N \hat{f}_k \operatorname{sgn}(k) \sigma\left(\frac{|k|}{N}\right) e^{ikx} \quad (4.7)$$

This is known as the *generalized conjugate Fourier partial sum* with $\sigma_{k,N}(\eta) = \sigma\left(\frac{|k|}{N}\right)$ known as *concentration factors*. These factors determine convergence of the approximation to $[f]$. An example of the concentration method is shown in Figure 4.1, where the jump function of the sawtooth signal is approximated. This function has a single jump at $x = 0$. This plot also sheds light on the nomenclature of this method, which arises from the fact that the jump function approximation “concentrates” at the singular support of f .

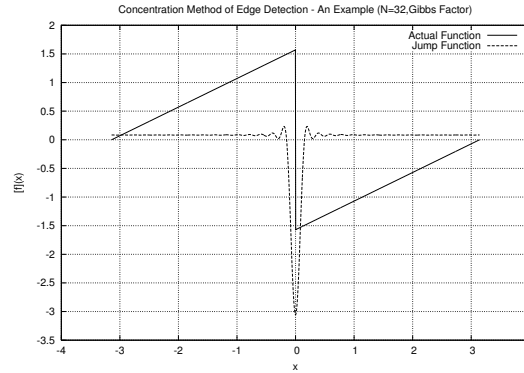


Fig. 4.1. Example of the concentration method

Concentration factors have to satisfy certain properties in order to be admissible. These include,

1. $\sum_{k=1}^N \sigma\left(\frac{k}{N}\right) \sin(kx)$ be odd

2. $\frac{\sigma(\eta)}{\eta} \in C^2(0, 1)$
3. $\int_{\epsilon}^1 \frac{\sigma(\eta)}{\eta} \rightarrow -\pi$, $\epsilon = \epsilon(N) > 0$ being small

Under these conditions, we have the following relation (concentration property, [38])

$$S_N^\sigma[f](x) = [f](x) + \mathcal{O}(\epsilon), \quad \epsilon = \epsilon(N) > 0 \text{ being small} \quad (4.8)$$

While we may synthesize any number of concentration factors satisfying the required conditions, there are several well studied families available for use, each with specific design principles and convergence results [38]. Table 4.1 gives expressions for three such concentration factors: trigonometric, polynomial and exponential. The trigonometric factor emulates physical space divided differences while the p^{th} order polynomial factor is equivalent to an order p differentiated Fourier partial sum. The exponential factor is specifically designed to have exponentially small moments, thereby driving $S_N^\sigma[f]$ to infinitesimally small values in regions of smoothness. Fig 4.2 provides visual indication of the performance of each of these factors when applied to a unit step function, while Figure 4.3 plots $\sigma_{k,N}(\eta)$ in Fourier space. Once the jump function approximation is computed, decisions on absence or presence of a discontinuity are made by comparing against an empirically determined, data-dependent threshold.

TABLE 4.1
POPULAR CONCENTRATION FACTORS

| Factor | Expression | Remarks |
|---------------|---|--|
| Trigonometric | $\sigma_G(\eta) = \frac{\pi \sin(\pi\eta)}{\text{Si}(\pi)}$ | $\text{Si}(\pi) = \int_0^\pi \frac{\sin(x)}{x} dx$ |
| Polynomial | $\sigma_P(\eta) = -p\pi\eta^p$ | p is the order of the factor |
| Exponential | $\sigma_E(\eta) = C \cdot \eta \cdot \exp\left(\frac{1}{\alpha\eta(\eta-1)}\right)$ | α is the order C is a normalizing constant $C = \frac{\pi}{\int_{\frac{1}{N}}^{1-\frac{1}{N}} \exp\left(\frac{1}{\alpha\tau(\tau-1)}\right) d\tau}$ |

In the following section, we provide a statistical analysis and quantitative performance measure of the concentration method in the presence of noise. We also present a linear framework of combining concentration factors for obtaining improved results, comparable to nonlinear variants described in [36, 37].

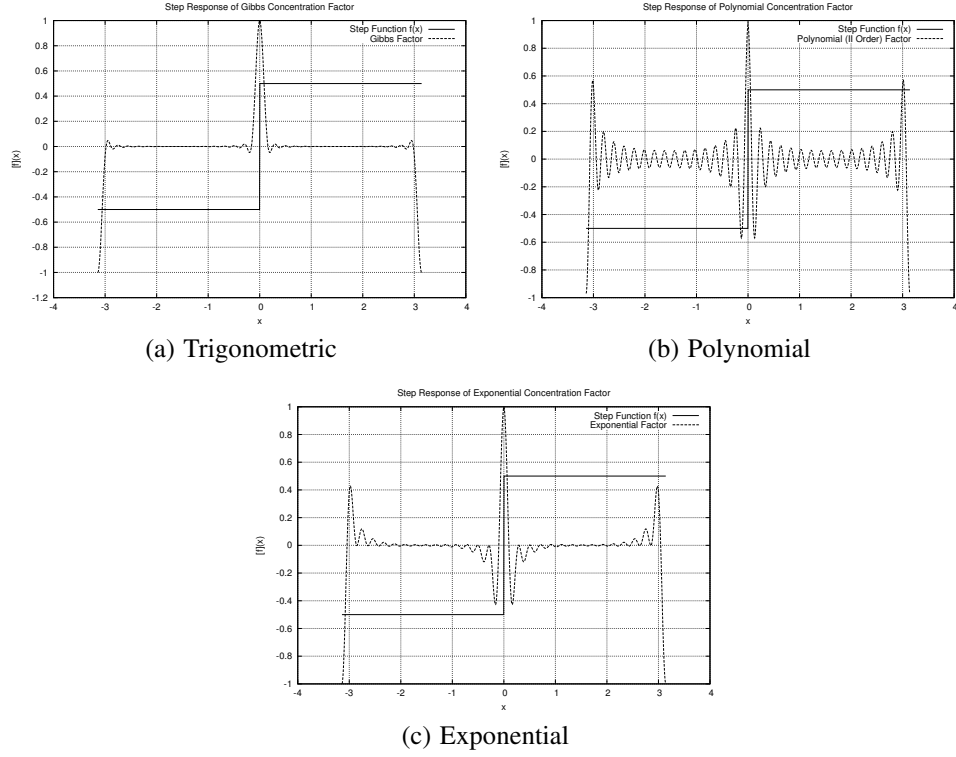


Fig. 4.2. Step response of different concentration factors ($N = 32$)

4.2. Statistical Analysis of the Concentration Method

This section provides a statistical analysis of the concentration method, with the intention of designing a noise robust discontinuity detector. To begin with, we shall assume that the Fourier coefficients $\hat{f}(k)$ are corrupted by zero-mean, additive complex white Gaussian noise of variance ρ^2 . The restriction on correlation is purely for ease of development; extension of the ensuing analysis to incorporate correlated noise is straightforward. Hence,

$$\hat{\mathbf{g}}(k) = \hat{f}(k) + \hat{\mathbf{v}}(k) \quad k \in [-N, N], \quad \hat{\mathbf{v}}(k) \sim \mathcal{CN}[0, \rho^2] \quad (4.9)$$

where $\hat{\mathbf{g}}$ denote the observed coefficients and $\hat{\mathbf{v}}$, the noise contribution. Since the concentration method is a linear operator, we have

$$\begin{aligned} S_N^\sigma[\mathbf{g}](x) &= i \sum_{k=-N}^N (\hat{f}(k) + \hat{\mathbf{v}}(k)) \operatorname{sgn}(k) \sigma\left(\frac{|k|}{N}\right) e^{ikx} \\ &= i \sum_{k=-N}^N \hat{f}(k) \operatorname{sgn}(k) \sigma\left(\frac{|k|}{N}\right) e^{ikx} + i \sum_{k=-N}^N \hat{\mathbf{v}}(k) \operatorname{sgn}(k) \sigma\left(\frac{|k|}{N}\right) e^{ikx} \end{aligned}$$




Fig. 4.3. Concentration factors plotted in Fourier space

$$S_N^\sigma[\mathbf{g}](x) = S_N^\sigma[f](x) + S_N^\sigma[\mathbf{v}](x)$$

Further,

$$\begin{aligned} E[S_N^\sigma[\mathbf{g}](x)] &= E[S_N^\sigma[f](x) + S_N^\sigma[\mathbf{v}](x)] \\ &= S_N^\sigma[f](x) + E[S_N^\sigma[\mathbf{v}](x)] \\ E[S_N^\sigma[\mathbf{g}](x)] &= S_N^\sigma[f](x) \end{aligned} \quad (4.10)$$

Here, we use $E[S_N^\sigma[\mathbf{v}](x)] = 0$, since $\hat{\mathbf{v}}$ is a zero-mean random variable. In summary, the noise component does not bias the jump function approximation.

Even though the corrupting noise is white, use of concentration factors (which are Fourier space filters) result in the jump function approximation acquiring a covariance structure. Entries of the covariance matrix which describe this structure can be obtained as follows. Let $S_N^{\sigma_p}[\mathbf{g}](x_a)$ denote the jump function approximation computed at location x_a using the concentration factor σ_p , and $S_N^{\sigma_q}[\mathbf{g}](x_b)$ denote the jump function approximation computed at location x_b using the concentration factor σ_q . Then

$$\begin{aligned} (C_{\mathbf{v}})_{p,q}^{x_a,x_b} &= E \left[(S_N^{\sigma_p}[\mathbf{g}](x_a) - E[S_N^{\sigma_p}[\mathbf{g}](x_a)]) (S_N^{\sigma_q}[\mathbf{g}](x_b) - E[S_N^{\sigma_q}[\mathbf{g}](x_b)])^* \right] \\ &= E \left[(S_N^{\sigma_p}[\mathbf{v}](x_a)) (S_N^{\sigma_q}[\mathbf{v}](x_b))^* \right] \\ &= E \left[\left(i \sum_{l=-N}^N \hat{\mathbf{v}}(l) \operatorname{sgn}(l) \sigma_p\left(\frac{|l|}{N}\right) e^{ilx_a} \right) \left(-i \sum_{m=-N}^N \hat{\mathbf{v}}(m)^* \operatorname{sgn}(m) \sigma_q\left(\frac{|m|}{N}\right) e^{-imx_b} \right) \right] \end{aligned}$$

$$\begin{aligned}
&= E \left[\sum_{l=-N}^N |\hat{\mathbf{v}}(l)|^2 \sigma_p\left(\frac{|l|}{N}\right) \sigma_q\left(\frac{|l|}{N}\right) e^{il(x_a - x_b)} \right] \\
&\quad + E \left[\sum_{m=-N}^N \sum_{l=-N}^N \hat{\mathbf{v}}(m)^* \hat{\mathbf{v}}(l) \operatorname{sgn}(m) \operatorname{sgn}(l) \sigma_q\left(\frac{|m|}{N}\right) \sigma_p\left(\frac{|l|}{N}\right) e^{-imx_b} e^{ilx_a} \right] \\
&= \sum_{l=-N}^N E[|\hat{\mathbf{v}}(l)|^2] \sigma_p\left(\frac{|l|}{N}\right) \sigma_q\left(\frac{|l|}{N}\right) e^{il(x_a - x_b)} \\
&\quad + \sum_{m=-N}^N \sum_{l=-N}^N E[\hat{\mathbf{v}}(m)^* \hat{\mathbf{v}}(l)] \operatorname{sgn}(m) \operatorname{sgn}(l) \sigma_q\left(\frac{|m|}{N}\right) \sigma_p\left(\frac{|l|}{N}\right) e^{-imx_b} e^{ilx_a}
\end{aligned}$$

where the second equation results from using (4.10). Since $\mathbf{v}(l)$, $l \in [-N, N]$, are independent and zero mean variates, $E[\hat{\mathbf{v}}(m)^* \hat{\mathbf{v}}(l)] = E[\hat{\mathbf{v}}(m)^*] E[\hat{\mathbf{v}}(l)] = 0$. Hence the second term in the above expression goes to zero. Further, $E[|\hat{\mathbf{v}}(l)|^2] = \rho^2$ and

$$(C_{\mathbf{v}})_{p,q}^{x_a, x_b} = \rho^2 \sum_{l=-N}^N \sigma_p\left(\frac{|l|}{N}\right) \sigma_q\left(\frac{|l|}{N}\right) e^{il(x_a - x_b)} \quad (4.11)$$

describes a Hermitian symmetric matrix which is Toeplitz only for identical σ_p and σ_q .

4.3. Detector Design

We are now in a position to formulate the detection problem. For analysis purposes, we will use the unit step function with jump at $x = 0$ (and a 1-sparse jump function) as a template for jumps. Our objective will be to test at $x = 0$ and/or in its neighbouring region, and decide between the presence of a jump (hypothesis \mathcal{H}_1) or its absence (hypothesis \mathcal{H}_0). Let M be the L -length vector constructed by stacking the jump function approximations at points $(x_l)_{l=1}^L$ obtained using concentration factors σ_l ,

$$M = (S_N^{\sigma_1}[f](x_1), \dots, S_N^{\sigma_L}[f](x_L))^T \quad (4.12)$$

This is the true signal vector and is pre-computed from the true Fourier coefficients of the unit step function. An example signal vector is shown in Figure 4.4, where a single concentration factor (trigonometric) has been used. Similarly, let \mathbf{Y} represent the L -length vector of jump function approximations computed from the observed (noisy) Fourier coefficients.

$$\mathbf{Y} = (S_N^{\sigma_1}[\mathbf{g}](x_1), \dots, (S_N^{\sigma_L}[\mathbf{g}](x_L))^T \quad (4.13)$$

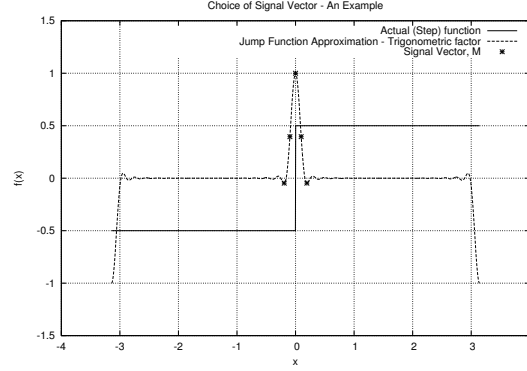


Fig. 4.4. Example of a signal vector ($N = 32$, trigonometric factor)

We note that $E[\mathbf{Y}] = M$ (from 4.10). Further, the covariance structure of \mathbf{Y} is given by (4.11). From formulation of the concentration method, (4.8), we may write the detection problem on data \mathbf{Y} as

$$\mathcal{H}_0 : \mathbf{Y} = \mathbf{V} \sim \mathcal{CN}[0, C_{\mathbf{V}}]$$

$$\mathcal{H}_1 : \mathbf{Y} = M + \mathbf{V} \sim \mathcal{CN}[M, C_{\mathbf{V}}]$$

where $\mathcal{CN}[A, C]$ represents a multivariate Gaussian distribution with mean A and covariance matrix C . To maximize the probability of detection, P_d for a given false alarm rate P_{fa} , we apply the Neyman-Pearson criterion² [40]

$$\begin{aligned} \rightarrow \mathcal{H}_1 : \frac{\Pr(\mathbf{Y}|\mathcal{H}_1)}{\Pr(\mathbf{Y}|\mathcal{H}_0)} &> \gamma \\ \frac{\frac{1}{(2\pi C_{\mathbf{V}})^{L/2}} \exp\left[-\frac{1}{2}(\mathbf{Y} - M)^T C_{\mathbf{V}}^{-1}(\mathbf{Y} - M)\right]}{\frac{1}{(2\pi C_{\mathbf{V}})^{L/2}} \exp\left[-\frac{1}{2}\mathbf{Y}^T C_{\mathbf{V}}^{-1}\mathbf{Y}\right]} &> \gamma \\ \exp\left[-\frac{1}{2}(\mathbf{Y} - M)^T C_{\mathbf{V}}^{-1}(\mathbf{Y} - M)\right] \exp\left[\frac{1}{2}\mathbf{Y}^T C_{\mathbf{V}}^{-1}\mathbf{Y}\right] &> \gamma \\ \exp\left[-\frac{1}{2}(\mathbf{Y}^T C_{\mathbf{V}}^{-1}\mathbf{Y} - M^T C_{\mathbf{V}}^{-1}\mathbf{Y} - \mathbf{Y}^T C_{\mathbf{V}}^{-1}M + M^T C_{\mathbf{V}}^{-1}M - \mathbf{Y}^T C_{\mathbf{V}}^{-1}\mathbf{Y})\right] &> \gamma \\ \exp\left[M^T C_{\mathbf{V}}^{-1}\mathbf{Y} - \frac{1}{2}M^T C_{\mathbf{V}}^{-1}M\right] &> \gamma \end{aligned}$$

Incorporating the data independent term $\exp\left[-\frac{1}{2}M^T C_{\mathbf{V}}^{-1}M\right]$ into the threshold γ , we have

$$\begin{aligned} \exp\left[M^T C_{\mathbf{V}}^{-1}\mathbf{Y}\right] &> \bar{\gamma} \\ \rightarrow \mathcal{H}_1 : M^T C_{\mathbf{V}}^{-1}\mathbf{Y} &> \dot{\gamma} \end{aligned} \quad (4.14)$$

² $\Pr(\cdot)$ denotes the probability density function and γ is a threshold. The notation $\rightarrow \mathcal{H}_1$ stands for “choose hypothesis \mathcal{H}_1 ”.

where $\hat{\gamma} = \ln \bar{\gamma}$. The above formulation is referred to as a generalized matched filter. To make a decision, we compute the weighted scalar product described by (4.14) and compare it to a threshold. The term $M^T C_{\mathbf{V}}^{-1} \mathbf{Y}$ is also referred to as the test statistic.

4.4. Performance of the Detector

Let us determine the statistical properties of the test statistic, $T(\mathbf{Y}) = M^T C_{\mathbf{V}}^{-1} \mathbf{Y}$. The test statistic is a linear combination of independent, identically distributed Gaussian variates, and hence remains Gaussian.

The mean under either hypothesis is

$$\mathbb{E}[T|\mathcal{H}_0] = \mathbb{E}[M^T C_{\mathbf{V}}^{-1} \mathbf{V}] = 0$$

and

$$\mathbb{E}[T|\mathcal{H}_1] = \mathbb{E}[M^T C_{\mathbf{V}}^{-1} (M + \mathbf{V})] = \mathbb{E}[M^T C_{\mathbf{V}}^{-1} M] + \mathbb{E}[M^T C_{\mathbf{V}}^{-1} \mathbf{V}] = M^T C_{\mathbf{V}}^{-1} M$$

Similarly, the variance under either hypothesis can be computed as³

$$\begin{aligned} \text{var}[T|\mathcal{H}_0] &= \mathbb{E} \left[(M^T C_{\mathbf{V}}^{-1} \mathbf{V}) (M^T C_{\mathbf{V}}^{-1} \mathbf{V})^* \right] \\ &= \mathbb{E} [M^T C_{\mathbf{V}}^{-1} \mathbf{V} \mathbf{V}^* C_{\mathbf{V}}^{-1} M] \\ &= M^T C_{\mathbf{V}}^{-1} \mathbb{E}[\mathbf{V} \mathbf{V}^*] C_{\mathbf{V}}^{-1} M \\ &= M^T C_{\mathbf{V}}^{-1} M \end{aligned}$$

and

$$\begin{aligned} \text{var}[T|\mathcal{H}_1] &= \mathbb{E} \left[(M^T C_{\mathbf{V}}^{-1} \mathbf{Y} - \mathbb{E}[M^T C_{\mathbf{V}}^{-1} \mathbf{Y}]) (M^T C_{\mathbf{V}}^{-1} \mathbf{Y} - \mathbb{E}[M^T C_{\mathbf{V}}^{-1} \mathbf{Y}])^* \right] \\ &= \mathbb{E} \left[(M^T C_{\mathbf{V}}^{-1} \mathbf{V}) (M^T C_{\mathbf{V}}^{-1} \mathbf{V})^* \right] \\ &= \text{var}[T|\mathcal{H}_0] = M^T C_{\mathbf{V}}^{-1} M \end{aligned}$$

Therefore, we have

$$\begin{aligned} \mathcal{H}_0 : \quad T(\mathbf{Y}) &\sim \mathcal{N}[0, M^T C_{\mathbf{V}}^{-1} M] = \mathcal{N}[0, d^2] \\ \mathcal{H}_1 : \quad T(\mathbf{Y}) &\sim \mathcal{N}[M^T C_{\mathbf{V}}^{-1} M, M^T C_{\mathbf{V}}^{-1} M] = \mathcal{N}[d^2, d^2] \end{aligned} \quad (4.15)$$

³Note that $\mathbb{E}[\mathbf{V} \mathbf{V}^*] = C_{\mathbf{V}}$.

where, d^2 is termed as the signal-to-noise ratio (SNR). The false alarm and correct detection probabilities may now be computed as the right tail probabilities of the corresponding Gaussian probability distributions:

$$\begin{aligned}
 P_{FA} &= \int_{\dot{\gamma}}^{\infty} Pr(T|\mathcal{H}_0)(\eta) d\eta \\
 &= \int_{\dot{\gamma}}^{\infty} \frac{1}{\sqrt{2\pi d^2}} \exp\left[-\frac{\eta^2}{2d^2}\right] d\eta \\
 &= Q\left(\frac{\dot{\gamma}}{d}\right)
 \end{aligned} \tag{4.16}$$

and

$$\begin{aligned}
 P_D &= \int_{\dot{\gamma}}^{\infty} Pr(T|\mathcal{H}_1)(\eta) d\eta \\
 &= \int_{\dot{\gamma}}^{\infty} \frac{1}{\sqrt{2\pi d^2}} \exp\left[-\frac{(\eta - d^2)^2}{2d^2}\right] d\eta \\
 &= Q\left(\frac{\dot{\gamma} - d^2}{d}\right)
 \end{aligned} \tag{4.17}$$

where $Q(\cdot)$ is the complementary (Gaussian) cumulative distribution function, which is available in tabular form. Using (4.16) and (4.17), we may write

$$\begin{aligned}
 P_D &= Q\left(\frac{Q^{-1}(P_{FA})d - d^2}{d}\right) \\
 &= Q\left(Q^{-1}(P_{FA}) - \sqrt{d^2}\right)
 \end{aligned} \tag{4.18}$$

The above equation completely describes the performance of the detector, since it relates the probability of detection to the false alarm probability. We note that, for a given false alarm probability, the performance only depends on d^2 , the SNR of the detector. Further, we may use (4.16) to set the optimum threshold, $\dot{\gamma}$ in (4.14), to ensure operation of the detector at the desired false alarm rate; i.e.,

$$\dot{\gamma} = Q^{-1}(P_{FA}) \cdot d$$

4.5. Choice of Detector Parameters

For a fixed noise variance ρ^2 , the SNR of the detector depends on the choice of the signal vector M , and consequently, the choice of measurement points and concentration factors, (x_i, σ_i) . Figure 4.2 indicates that each concentration factor produces a distinct response, differing in aspects such as jump resolution (or

mainlobe width) and spurious oscillations away from the jump. If we think of each response as a template indicative of a jump, appropriately chosen combinations of these responses would provide a stronger indication of the jump. This is the rationale behind use of multiple concentration factors in a detector. The strength of each combination is governed by the covariance structure defined by the (x_i, σ_i) pairs, and a numerical measure of this strength is provided by the SNR d^2 . As an example, Table 4.2 shows the variation in the SNR as x_i is varied. Here, a single concentration factor (trigonometric) was used. Table 4.3 illustrates the variation in the SNR as σ_i is varied while keeping x_i constant. In this particular example, $x_i = (-0.193, -0.096, 0, 0.096, 0.193)$ were chosen, which are equispaced grid points spaced $\frac{1}{2N+1}$ apart. The notation used in this table is as follows: T - Trigonometric factor, P - Polynomial factor, E - Exponential factor.

TABLE 4.2

CHOICE OF LOCATIONS x_i AND EFFECT ON PERFORMANCE USING ONLY THE TRIGONOMETRIC CONCENTRATION FACTOR ($L = 5, N = 32, \rho^2 = 7.5$)

| Choice of x_i | Significance | $M^T C_V^{-1} M$ (dB) |
|-----------------------------------|---|-----------------------|
| (-0.294, -0.196, 0, 0.196, 0.294) | maxima/minima locations in $S_N^\sigma[f](x)$ | 13.498 |
| (-0.193, -0.096, 0, 0.096, 0.193) | grid points spaced $\frac{1}{2N+1}$ apart | 14.656 |
| (-0.348, -0.173, 0, 0.173, 0.348) | numerically computed to maximize the SNR | 16.396 |

TABLE 4.3

CHOICE OF CONCENTRATION FACTORS AND EFFECT ON PERFORMANCE USING A FIXED SET OF LOCATIONS x_i ($L = 5, N = 32, \rho^2 = 7.5$)

| Choice of concentration factors ($\sigma_1 \sigma_2 \sigma_3 \sigma_4 \sigma_5$) | $M^T C_V^{-1} M$ (dB) |
|--|-----------------------|
| (T P E P T) | 11.285 |
| (T P E T P) | 13.123 |
| (T P P P T) | 14.727 |

4.6. Results

Receiver operating characteristic (ROC) curves, which plot the probability of detection against false alarm probability, are shown in Figure 4.5 for two different values of noise variance. The predicted perfor-

mance was plotted using relation (4.18), while simulation results were obtained from Monte Carlo trials. For the Monte Carlo simulations, 2000 i.i.d. Gaussian variates were generated with parameters as per (4.15), and the number of false alarms and correct detects were counted for a range of thresholds. After converting these numbers to percentages, the ROC curve was obtained by plotting respective (P_{FA}, P_D) points. For both plots, a three point signal vector utilizing the exponential factor was used with $N = 32$. The set of measurement points x_i was chosen to be $\{-\delta \ 0 \ \delta\}$ with $\delta = \frac{1}{2N+1}$. The SNR metrics for the two cases were 9.0283 dB and 12.338 dB corresponding to $\rho^2 = 7.5$ and $\rho^2 = 3.5$ respectively.

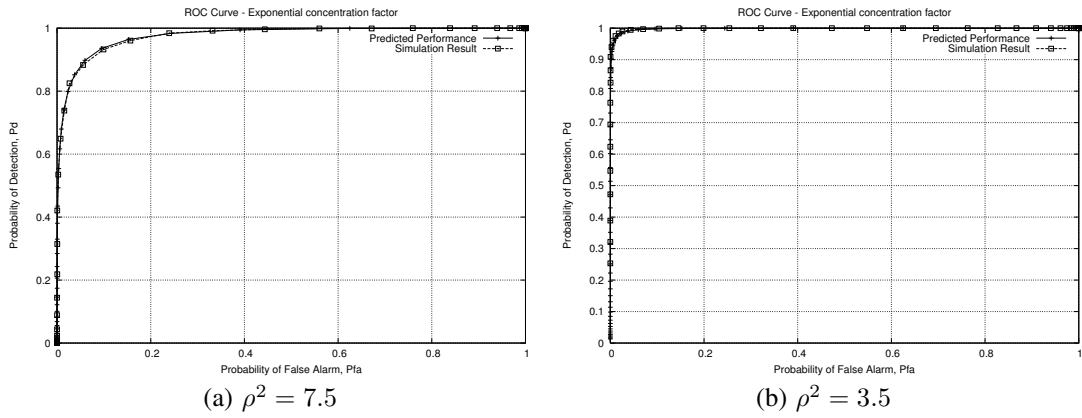


Fig. 4.5. ROC curve ($L = 3, N = 32$, exponential factor)

Figure 4.6 shows the gain in performance from using a combination of concentration factors. Both curves represent three-point signal vector detectors with $N = 32$, operating at a noise level of $\rho^2 = 7.5$. The dashed curve was generated using the exponential concentration factor at points $x_i = (-\delta \ 0 \ \delta)$ with $\delta = \frac{1}{2N+1}$, resulting in a SNR of 9.0283 dB. The solid curve was generated using all three concentration factors and $x_i = 0$, resulting in a SNR of 13.434 dB.

Until now, we have analyzed the performance of our detector on a unit step function, with our only concern being the presence or absence of a jump at $x = 0$. Not only did this permit easy development and analysis of our method, this constitutes a good template since the unit step function has a 1-sparse jump function. The performance of this method on a synthetic test function is shown in Figure 4.7. To provide a visual indication of the level of noise ($\rho^2 = 7.5$), the Fourier approximation computed using the noisy




Fig. 4.6. Single concentration factor versus a combination of factors ($L = 3, N = 32, \rho^2 = 7.5$)

Fourier coefficients is shown in Figure 4.7(a). The function⁴ $f(x)$ shown in the plot is

$$f(x) = \begin{cases} 0 & -\pi \leq x < -\frac{3\pi}{4} \\ \frac{3}{2} & -\frac{3\pi}{4} \leq x < -\frac{\pi}{2} \\ 0 & -\frac{\pi}{2} \leq x < -\frac{\pi}{4} \\ \frac{7}{4} - \frac{x}{2} + \sin(x - \frac{1}{4}) & -\frac{\pi}{4} \leq x < \frac{\pi}{8} \\ 0 & \frac{\pi}{8} \leq x < \frac{3\pi}{8} \\ \frac{11}{4}x - 5 & \frac{3\pi}{8} \leq x < -\frac{3\pi}{4} \\ 0 & \frac{3\pi}{4} \leq x < \pi \end{cases} \quad (4.19)$$

with a jump function

$$[f](x) = \begin{cases} \frac{3}{2} & x = -\frac{3\pi}{4} \\ -\frac{3}{2} & x = -\frac{\pi}{2} \\ \frac{14+\pi}{8} - \sin\left(\frac{\pi+1}{4}\right) \approx 1.28 & x = -\frac{\pi}{4} \\ \sin\left(\frac{2-\pi}{8}\right) - \frac{28-\pi}{16} \approx -1.70 & x = \frac{\pi}{8} \\ \frac{33\pi}{32} - 5 \approx -1.76 & x = \frac{3\pi}{8} \\ 5 - \frac{33\pi}{16} \approx -1.48 & x = \frac{3\pi}{4} \\ 0 & \text{elsewhere} \end{cases} \quad (4.20)$$

A three-point detector using the trigonometric concentration factor with x_i chosen to be $(-\delta \ 0 \ \delta)$ where $\delta = \frac{1}{2N+1}$ was used. Some additional processing was required for this example - prior to applying the

⁴To compare results, the test function used in [39] is used here.

detector, negative valued jumps were inverted. Further, if the detector called a detect at x_i , $S_N^\sigma[g](x_i)$ was assigned as the jump magnitude. The detected jumps are shown in Figure 4.7(b).

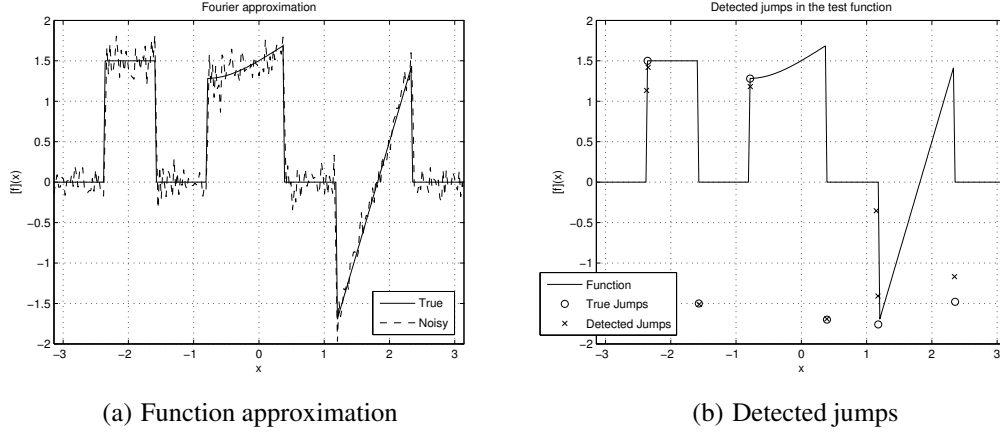


Fig. 4.7. Performance on a test function ($L = 3$, $N = 128$, $\rho^2 = 7.5$)

From the above results, we see that this is a promising method of detecting jump discontinuities in the presence of noise. The only notable shortcoming is the presence of false alarms in the immediate vicinity of a discontinuity. This is expected behaviour, as the \mathcal{H}_0 hypothesis does not apply to regions in the immediate vicinity of a jump (these regions are not $\mathcal{O}(\epsilon)$). Sidelobe mitigation procedures are needed, and are part of future work on this topic.

4.7. Extensions and Remarks

This method is equally applicable to multi-dimensional functions [41]. For example, in the case of a two dimensional function, the concentration method is applied to one dimension, holding the other constant.

$$S_N^\sigma[f](x(\bar{y})) = i \sum_{l=-N}^N \text{sgn}(l) \sigma\left(\frac{|l|}{N}\right) \sum_{k=-N}^N \hat{f}_{k,l} e^{i(kx+l\bar{y})} \quad (4.21)$$

The overbar represents the dimension held constant. It is obvious that overlapping information will be obtained by applying the method to the other dimension. For completeness, the result of edge detection on a Shepp-Logan phantom is shown in Figure 4.8. As in previous chapters, the analytic (Cartesian) Fourier coefficients of the Shepp-Logan phantom were computed using the method outlined in [30]. For a visual illustration of the amount of noise added, the Fourier reconstruction of the phantom is also shown. The trigonometric concentration factor was used in conjunction with a seven-point detector.

The concentration method is also applicable to other spectral expansions such as Chebyshev and Legendre. However, there are no popular applications outside of solving numerical partial differential equations which demand the detection of discontinuities from Chebyshev or Legendre spectral data. Hence, we have restricted our discussion to Fourier methods. A similar development for these spectral expansions is possible, and the interested reader is referred to [36] for details.




(a) True phantom



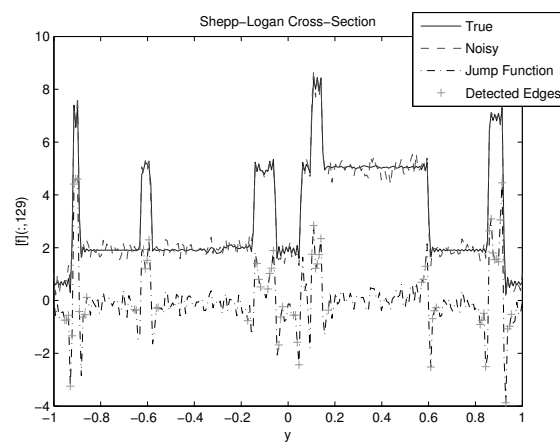
(b) Noisy Fourier reconstruction



(c) Jump function



(d) Edge map



(e) A cross-section

Fig. 4.8. Edge detection on the Shepp-Logan phantom (trigonometric factor, $L = 7$, $N = 128$, $\rho^2 = 25.5$)

5. SUMMARY

The problem of reconstructing support-limited functions from non-Cartesian spectral samples was discussed. Failure of traditional partial Fourier sums in the case of irregularly sampled data was demonstrated. Prevalent reconstruction schemes, including convolutional gridding, uniform re-sampling and iterative methods were surveyed and compared on computational cost, reconstruction quality and implementation steps. Illustrative examples were given, using different sampling schemes, including those prevalent in magnetic resonance imaging (MRI).

For variable sampling density schemes, typically encountered in MRI, it was observed that high frequency spectral coefficients were recovered with large error, leading to reduced accuracy in reconstruction. An alternate reconstruction scheme, based on spectral reprojection, formulated to use only accurate low frequency coefficients was suggested. Empirical results were provided, showing superior performance compared to traditional filtered Fourier reconstruction, despite use of far fewer coefficients.

Additionally, the recovery of discontinuity data, which is essential for spectral reprojection as well as other signal processing tasks was discussed. The concentration method was employed to detect magnitude, location and sign of jump discontinuities in piecewise-smooth functions from a finite number of noisy Fourier coefficients. Statistical formulation of the concentration method was developed under the assumption that corrupting noise is zero mean, complex-additive and white. Expressions for mean and covariance of the resulting jump function approximation were provided and a detection problem formulated. To maximize detection probability for fixed false alarms, a Neyman-Pearson detection problem was solved, resulting in a generalized matched filter. A SNR metric was also derived with the measurement point - concentration factor pair (x_i, σ_i) serving as parameters.

To improve detection performance, a method involving use of multiple concentration factors was developed. Examples showing improvement in SNR for various choices of (x_i, σ_i) were given. To prove the efficacy of this method, test results using a template step function as well as a synthetic test function were provided. Test results for the step function included ROC curves at different noise levels for various choices of (x_i, σ_i) . Finally, the extension of this method to two dimensions was illustrated with the detection of edges in a noisy Shepp-Logan phantom.

5.1. Avenues for Further Work

Both spectral sampling and discontinuity detection are topics rich in avenues for further research. Some of the problems which constitute our immediate research goals are enumerated below,

5.1.1. Spectral Sampling

- Non-uniform FFT methods provide efficient methods to compute non-equispaced sums of the form

$$\tilde{f}(x) = \sum_{k=0}^{N-1} \left(\alpha_k \hat{f}(\omega_k) \right) e^{i\omega_k x}$$

If $\check{f}(x)$ represents the NEFFT approximation, the error is usually defined as $\|\tilde{f} - \check{f}\|$ in a suitable norm.

Window design is performed with the objective of minimizing this error. However, we know that the Fourier basis gives a slowly converging approximation for piecewise-smooth functions. Hence, better results can be obtained if window design is performed with the objective of minimizing $\|f - \check{f}\|$

- A hybrid reconstruction method, suitably combining filtered Fourier reconstruction and Gegenbauer reconstruction is to be implemented. This method promises similar order of accuracy to traditional Gegenbauer reconstruction, but at reduced computational cost.
- It is conceivable that by suitable window design and density compensation, jump discontinuities can be approximated using convolutional gridding. This is in contrast to the present method, where jump detection is performed using coefficients recovered by URS/BURS.
- Use of adaptive mollifiers [42], which promise exponentially accurate reconstructions without the need for spectral reprojecton has to be investigated.

5.1.2. Discontinuity Detection from Spectral Data

- Experimental results with the detector reveal a cluster of false alarms in the immediate neighbourhood of a true discontinuity. Sidelobe mitigation techniques to arrest their presence would further improve detector performance.
- Non-linear variants of the concentration method exist such as enhancement of scales [36] and the minmod algorithm [37]. Although complicated, a statistical analysis of these methods and comparison with the current method would provide more insight towards design of a robust discontinuity detector.

- Performance comparison with corresponding physical space discontinuity detection methods would provide an understanding of the performance gap and relative strengths of global and local discontinuity detection methods.
- The concentration method exploits decay in Fourier coefficients to arrive at a jump function approximation. However, it is reasonable to expect improved performance if structural features of the underlying function are exploited. This is particularly true for multi-dimensional functions.

REFERENCES

- [1] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Math. of Comput.*, vol. 19, no. 90, pp. 297–301, 1965.
- [2] K. Atkinson, *An Introduction to Numerical Analysis*. John Wiley & Sons, 1978.
- [3] J. Hesthaven, S. Gottlieb, and D. Gottlieb, *Spectral Methods for Time-Dependent Problems*. Cambridge University Press, 2007.
- [4] A. J. Jerri, *The Gibbs Phenomenon in Fourier Analysis, Splines and Wavelet Approximations*. Springer, 1998.
- [5] D. Gottlieb and C. W. Shu, "On the Gibbs phenomenon and its resolution," *SIAM review(Print)*, vol. 39, no. 4, pp. 644–668, 1997.
- [6] E. Tadmor, "Adaptive mollifiers for high resolution recovery of piecewise smooth data from its spectral information," *Found. Comput. Math.*, vol. 2, no. 2, pp. 155–189, 2002.
- [7] A. Dutt and V. Rokhlin, "Fast Fourier transforms for nonequispaced data," *SIAM J. Sci. Comput.*, vol. 14, no. 6, pp. 1368–1393, 1993.
- [8] K. Fourmont, "Non-equispaced fast Fourier transforms with applications to tomography," *J. of Fourier Anal. and Applicat.*, vol. 9, no. 5, pp. 431–450, 2003.
- [9] G. Steidl, "A note on fast Fourier transforms for nonequispaced grids," *Advances in Computational Math.*, vol. 9, no. 3, pp. 337–352, 1998.
- [10] J. A. Fessler and B. P. Sutton, "Nonuniform fast Fourier transforms using min-max interpolation," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 560–574, 2003.
- [11] 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. Imag.*, vol. 10, no. 3, pp. 473–478, 1991.
- [12] D. Slepian and H. O. Pollak, "Prolate spheroidal wave functions, Fourier analysis and uncertainty I," *Bell Syst. Tech. J.*, vol. 40, pp. 43–64, 1961.
- [13] J. A. Fessler and B. P. Sutton, "NUFFT-nonuniform FFT toolbox for Matlab," <http://www.eecs.umich.edu/~fessler/code/index.html>.
- [14] S. Kunis and D. Potts, "NFFT, software package," <http://www-user.tu-chemnitz.de/~potts/nfft>.
- [15] W. N. Brouw, "Aperture synthesis," *Methods in Comput. Physics. Volume 14-Radio astronomy*, vol. 76, no. 22804, pp. 09–89, 1975.
- [16] F. R. Schwab, "Optimal gridding of visibility data in radio interferometry," *Indirect Imaging*, pp. 333–343.

- [17] J. D. O'Sullivan, "Fast sinc function gridding algorithm for Fourier inversion in computer tomography," *IEEE Trans. Med. Imag.*, vol. 4, no. 4, 1985.
- [18] J. H. Zwaga, "MR image reconstruction from nonuniform samples using convolution gridding," Master's thesis, Delft Univ., NL, 1997.
- [19] J. G. Pipe and P. Menon, "Sampling density compensation in MRI: Rationale and an iterative numerical solution," *Magnetic Resonance in Medicine*, vol. 41, no. 1, pp. 179–186, 1999.
- [20] F. Wajer, R. Lethmate, R. A. J. de Jong, L. T. Martinez, D. Graveron-Demilly, and D. van Ormondt, "MR image reconstruction from sparse and corrupted k-space data," J. P. Veen, Ed. ProRISC, Mierlo, The Netherlands: IEEE Benelux, Nov. 1999, pp. 577–584.
- [21] R. D. Hoge, R. K. Kwan, and G. B. Pike, "Density compensation functions for spiral MRI," *Magnetic Resonance in Medicine*, vol. 38, no. 1, pp. 117–28, 1997.
- [22] H. Sedarat and D. G. Nishimura, "On the optimality of the gridding reconstruction algorithm," *IEEE Trans. Med. Imag.*, vol. 19, no. 4, pp. 306–317, 2000.
- [23] D. Rosenfeld, "An optimal and efficient new gridding algorithm using singular value decomposition," *Magnetic Resonance in Medicine*, vol. 40, no. 1, pp. 14–23, 1998.
- [24] D. Rosenfeld, "New approach to gridding using regularization and estimation theory," *Magnetic Resonance in Medicine*, vol. 48, no. 1, pp. 193–202, 2002.
- [25] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Johns Hopkins University Press, 1996.
- [26] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*. Polyteknisk Forlag, 1996.
- [27] A. I. Zayed, *Advances in Shannon's Sampling Theory*. CRC Press, 1993.
- [28] T. Knopp, S. Kunis, and D. Potts, "Fast iterative reconstruction for MRI from nonuniform k-space data," *revised Preprint A-05-10, Universitat zu Lubeck*, 2005.
- [29] B. P. Sutton, D. C. Noll, and J. A. Fessler, "Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities," *IEEE Trans. Med. Imag.*, vol. 22, no. 2, pp. 178–188, 2003.
- [30] R. Van De Walle, H. H. Barrett, K. J. Myers, M. I. Aitbach, B. Desplanques, A. Gmitro, J. Cornelis, and I. Lemahieu, "Reconstruction of MR images from data acquired on a general nonregular grid by pseudoinverse calculation," *IEEE Trans. Med. Imag.*, vol. 19, no. 12, pp. 1160–1167, 2000.
- [31] A. N. Tikhonov, *Solutions of Ill-Posed Problems*. Winston Publishing, 1977.
- [32] P. C. Hansen and D. P. OLeary, "The use of the L-curve in the regularization of discrete ill-posed problems," *SIAM J. Sci. Comput.*, vol. 14, p. 1487, 1993.

- [33] G. Wahba, "Practical approximate solutions to linear operator equations when the data are noisy," *SIAM J. Numerical Anal.*, vol. 14, no. 4, pp. 651–667, 1977.
- [34] A. Gelb, "A hybrid approach to spectral reconstruction of piecewise smooth functions," *J. Sci. Comput.*, vol. 15, no. 3, pp. 293–322, 2000.
- [35] A. Gelb, "Parameter optimization and reduction of round off error for the Gegenbauer reconstruction method," *J. Sci. Comput.*, vol. 20, no. 3, pp. 433–459, 2004.
- [36] A. Gelb and E. Tadmor, "Detection of edges in spectral data II - Nonlinear Enhancement," *SIAM J. Numerical Anal.*, vol. 38, no. 4, pp. 1389–1408, Sep.-Oct. 2000.
- [37] A. Gelb and E. Tadmor, "Adaptive edge detectors for piecewise smooth data based on the minmod limiter," *J. Sci. Comput.*, 2006.
- [38] A. Gelb and E. Tadmor, "Detection of edges in spectral data," *Appl. Comput. Harmonic Anal.*, vol. 7, pp. 101–135, 1999.
- [39] D. M. Cates, "Edge detection using Fourier data with applications," Dept. Math. Stat., Ph.D. dissertation, Arizona State Univ., Tempe, Arizona, 2007.
- [40] S. M. Kay, *Fundamentals of Statistical Signal Processing - Detection Theory*. Prentice Hall Signal Processing Series, 1993, ch. 3.
- [41] R. Archibald and A. Gelb, "A method to reduce the Gibbs ringing artifact in MRI scans while keeping tissue boundary integrity," *IEEE Trans. Med. Imag.*, vol. 21, no. 4, pp. 305–319, 2002.
- [42] J. Tanner, "Optimal filter and mollifier for piecewise smooth spectral data," *Math. Comp.*, vol. 75, no. 254, pp. 767–790, 2006.
- [43] L. A. Shepp and B. F. Logan, "The Fourier reconstruction of a head section," *IEEE Trans. Nucl. Sci.*, vol. 21, no. 1, pp. 21–43, 1974.

APPENDIX A

THE SHEPP-LOGAN PHANTOM AND SAMPLING SCHEMES

A.1. Shepp-Logan Phantom

The Shepp-Logan phantom [43] is an analytic test image, commonly used in medical imaging to compare image processing algorithms. It is composed of ten ellipses, each with different major/minor axes, intensity and orientation. Since it is analytic, its Fourier transform at any point $\omega \in \mathbb{R}^2$ can be easily computed [30].

Let (x_p, y_p) denote center of the p^{th} ellipse, ν_p its intensity, α_p the angle between its major axis and the x -axis, and A_p, B_p the major and minor axis lengths respectively. Then, the Fourier transform of the phantom at polar co-ordinate $\omega_k = k e^{i\theta}$ is given by

$$\hat{f}(\omega_k) = \sum_{p=1}^{10} \left(e^{-i\pi k t_p \cos(\gamma_p - \theta)} \nu_p A_p B_p \frac{J_1(a_p(\theta)k)}{a_p(\theta)k} \right)$$

where,

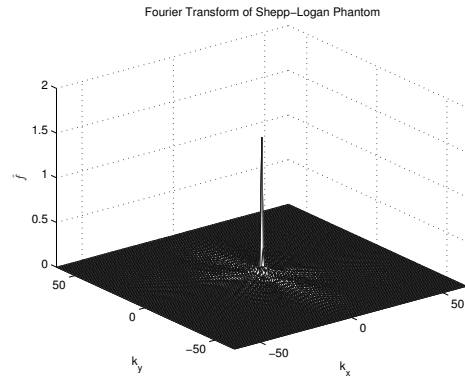
$$t_p = \sqrt{x_p^2 + y_p^2}, \quad \gamma_p = \arctan\left(\frac{y_p}{x_p}\right), \quad a_p(\theta) = \sqrt{A_p^2 \cos^2(\theta - \alpha_p) + B_p^2 \sin^2(\theta - \alpha_p)}$$

and $J_1(\cdot)$ is the first-order Bessel function of the first kind.

The Shepp-Logan phantom and its Fourier transform on a Cartesian grid are plotted in Figure A.1.



(a) Shepp-Logan phantom



(b) Fourier transform

Fig. A.1. Shepp-Logan phantom and its Fourier transform, $N = 128^2$

A.2. Sampling Schemes

Expressions for the different sampling schemes referred to in this thesis are provided below. The jittered sampling and “log sampling” schemes are formulated in one dimension, while radial and spiral sampling

schemes are formulated in two dimensions. In all cases, M denotes the total number of non-Cartesian measurements.

1. Jittered Sampling:

$$\omega_k = k \pm \tau_k, \quad \tau_k \sim U[0, .5], \quad k = 0, \dots, M - 1$$

where, $U[a, b]$ denotes a uniform distribution on the interval $[a, b]$. The τ_k 's are independent, identically distributed (iid) random variables, and represent a uniform jitter about the equispaced nodes with a maximal jitter of .5. 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.

2. Log Sampling:

If ω_k be the nodes at which measurements are acquired, $|\omega_k|$ is logarithmically distributed between 10^{-p} and $M/2$, where M is the total number of samples. A value of $p = 1.5$ was used.

3. Spiral Sampling:

$$\omega_p = C \frac{\sqrt{p}}{2\sqrt{M}} (\cos \theta_p, \sin \theta_p), \quad p = 0, \dots, M - 1$$

where, $\theta_p = \frac{8\pi}{5} \sqrt{p}$ and C is a constant.

4. Radial Sampling:

$$\omega_{r,\theta} = C(-1)^r \left(\frac{r}{R} - .5 \right) \left(\cos \frac{\pi\theta}{\Theta}, \sin \frac{\pi\theta}{\Theta} \right), \quad r = 0, \dots, R - 1; \quad \theta = 0, \dots, \Theta - 1$$

where, Θ is the total number of rays and R is the number of samples in each ray. Therefore, $R \times \Theta = M$, the total number of samples acquired. C is a constant.

For the two-dimensional schemes, C is a scaling constant which determines the span of measurements in Fourier space. Typically, C is chosen such that $\max |\omega_{\mathbf{k}}| = \max |\mathbf{k}|$, with \mathbf{k} denoting Cartesian sampling points. Illustrations of all sampling schemes can be found in Figure A.2.


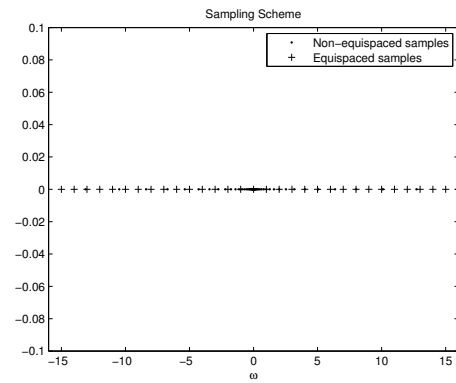
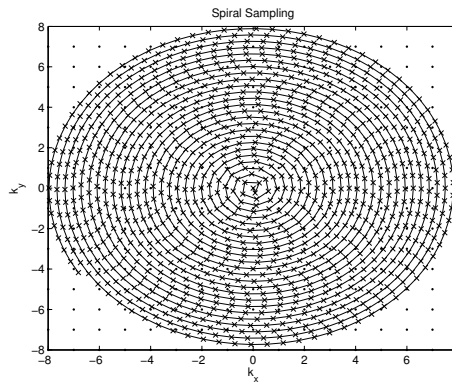

(a) Jittered sampling, $M = 32$ (b) Log sampling, $M = 32$ (c) Spiral sampling, $M = 1024$ (d) Radial sampling, $M = 1024$

Fig. A.2. Sampling schemes