

Synthetic Aperture Radar Image Formation

Via Sparse Decomposition

by

Nicholas Werth

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

Approved February 2011 by the
Graduate Supervisory Committee:

Lina Karam, Chair
Antonia Papandreou-Suppappola
Andreas Spanias

ARIZONA STATE UNIVERSITY

May 2011

ABSTRACT

Spotlight mode synthetic aperture radar (SAR) imaging involves a tomographic reconstruction from projections, necessitating acquisition of large amounts of data in order to form a moderately sized image. Since typical SAR sensors are hosted on mobile platforms, it is common to have limitations on SAR data acquisition, storage and communication that can lead to data corruption and a resulting degradation of image quality. It is convenient to consider corrupted samples as missing, creating a sparsely sampled aperture. A sparse aperture would also result from compressive sensing, which is a very attractive concept for data intensive sensors such as SAR. Recent developments in sparse decomposition algorithms can be applied to the problem of SAR image formation from a sparsely sampled aperture. Two modified sparse decomposition algorithms are developed, based on well known existing algorithms, modified to be practical in application on modest computational resources. The two algorithms are demonstrated on real-world SAR images. Algorithm performance with respect to super-resolution, noise, coherent speckle and target/clutter decomposition is explored. These algorithms yield more accurate image reconstruction from sparsely sampled apertures than classical spectral estimators. At the current state of development, sparse image reconstruction using these two algorithms require about two orders of magnitude greater processing time than classical SAR image formation.

To

My wife Cameron and son Jack

ACKNOWLEDGMENTS

I would like to acknowledge the support of my wife and family. They have provided unwavering encouragement and support, without which I would not have been able to complete this work. My wife has endured the long nights and periodic frustrations alongside me, and provided the motivation to initiate this work and see it through to completion.

Thank you to the thesis committee chair for feedback, suggestions and revisions contributing to this work. Thank you also to the other thesis committee members, who provided helpful comments and discussion.

I would also like to acknowledge the support of my friends and co-workers, who provided countless discussions and insights that helped shape the course of this work. This work was made possible through the support of my employer.

TABLE OF CONTENTS

| | Page |
|--|------|
| LIST OF FIGURES | vi |
| LIST OF ALGORITHMS | ix |
| CHAPTER | |
| 1 INTRODUCTION | 1 |
| 2 SAR BACKGROUND | 6 |
| 2.1. Spotlight mode SAR | 6 |
| 2.2. Tomographic model of spotlight mode SAR | 11 |
| 2.3. Processing spotlight mode SAR | 14 |
| 2.4. Spectral estimation and SAR image formation | 20 |
| 2.5. Problem statement | 21 |
| 2.6. Summary | 26 |
| 3 SPECTRAL ESTIMATORS AND SAR | 27 |
| 3.1. Classical estimators | 28 |
| 3.2. Nonlinear windowing | 29 |
| 3.3. Modern parametric estimators | 29 |
| 3.4. Modern non-parametric estimators | 33 |
| 3.5. Compressive sensing | 36 |
| 3.6. Analytic Comparison | 48 |
| 3.7. Summary | 50 |

| CHAPTER | Page |
|--|------|
| 4 PROPOSED SPARSE RECONSTRUCTION ALGORITHMS FOR SAR | |
| IMAGE FORMATION | 52 |
| 4.1. Iterative Mean Square Extrapolation | 52 |
| 4.2. Stagewise Gradient Pursuit | 58 |
| 4.3. Analytic Comparison | 66 |
| 4.4. Summary | 67 |
| 5 SIMULATIONS AND RESULTS | 68 |
| 5.1. Analytic Assumptions | 68 |
| 5.2. Super-Resolution Properties | 71 |
| 5.3. Point Target Detection | 73 |
| 5.4. Sparse Reconstruction of Clutter in Coherent Diffraction Limited Imaging | 77 |
| 5.5. Sparse Decomposition of Actual Spotlight SAR Imagery | 83 |
| 5.6. Summary | 92 |
| 6 CONCLUSION | 97 |
| 6.1. Contributions | 97 |
| 6.2. Conclusions | 97 |
| 6.3. Future Work | 98 |
| REFERENCES | 100 |

LIST OF FIGURES

| Figure | Page |
|---|------|
| 2.1 Radar azimuth resolution limit. | 7 |
| 2.2 Mapping ground points via range/Doppler method with SAR. | 8 |
| 2.3 Stripmap versus spotlight SAR modes. | 10 |
| 2.4 Collection of spotlight mode SAR data. | 11 |
| 2.5 Projection-slice principle for a single SAR transmit/receive pulse. | 12 |
| 2.6 Sampling in the Fourier domain according to the tomographic model of spotlight SAR. U_0 is the radial center frequency of the image Fourier domain collected in polar coordinates. | 13 |
| 2.7 Phase history sampled as a polar annulus in the Fourier domain. U_0 is the radial center frequency of the image Fourier domain collected in polar coordinates. | 14 |
| 2.8 Time-Frequency plot of a transmit and receive pulse through deramp and deskew processing. | 15 |
| 2.9 Differing range rates (profiles) for targets offset in azimuth. | 17 |
| 2.10 Differing phase history profiles for different points in the image. | 18 |
| 2.11 Effect of motion compensation processing on phase history targets. (a) Pulse echo phase delay of points at varying positions in the scene versus sensor offset position over a synthetic aperture; (b) Motion compensation applies a phase function to each return pulse such that a point at the scene center has zero phase variance over the synthetic aperture. | 19 |

| Figure | Page |
|---|------|
| 2.12 Phase history support and resulting 2D sinc IPR. | 23 |
| 2.13 Effect of irregular sampling on IPR in image. | 25 |
| 4.1 Distinct point targets versus diffuse scattering clutter in a SAR im- age. | 57 |
| 5.1 Resolvability of closely spaced scatterers. | 72 |
| 5.2 Simulated p_d and p_{fa} for a rectangular aperture. | 75 |
| 5.3 Simulated p_d and p_{fa} for a 25% sparse periodic aperture. | 77 |
| 5.4 Phantom image restoration: real-valued and non-diffraction limited. (a) Original. (b) Degraded 75% random aperture. (c) StGP restora- tion. (d) IMSE restoration. | 80 |
| 5.5 Phantom image restoration error convergence. Real-valued, non- diffraction limited, degraded 75% random aperture. | 81 |
| 5.6 Phantom image restoration: real-valued and diffraction-limited. (a) Original. (b) Degraded 75% random aperture. (c) StGP restoration. (d) IMSE restoration. | 82 |
| 5.7 Phantom image restoration error convergence. Real-valued, diffraction-limited, degraded 75% random aperture. | 83 |
| 5.8 Phantom image restoration: complex-valued and diffraction-limited. (a) Original. (b) Degraded 75% random aperture. (c) StGP restora- tion. (d) IMSE restoration. | 84 |
| 5.9 Phantom image restoration error convergence. Complex-valued, diffraction-limited, degraded 75% random aperture. | 85 |

| Figure | Page |
|---|------|
| 5.10 Parking lot scene with a 25% sparse periodic aperture. | 88 |
| 5.11 Parking lot scene target/clutter segmentation with a 25% sparse periodic aperture. | 89 |
| 5.12 Parking lot scene error convergence with a 25% sparse periodic aperture. | 90 |
| 5.13 Capitol scene with a 25% sparse periodic aperture. | 91 |
| 5.14 Capitol scene target/clutter segmentation with a 25% sparse periodic aperture. | 92 |
| 5.15 Capitol scene error convergence with a 25% sparse periodic aperture. | 93 |
| 5.16 Airport scene with a 25% sparse periodic aperture. | 94 |
| 5.17 Airport scene target/clutter segmentation with a 25% sparse periodic aperture. | 95 |
| 5.18 Airport scene error convergence with a 25% sparse periodic aperture. | 96 |

LIST OF ALGORITHMS

| Algorithm | Page |
|--|------|
| 3.1 Focal Underdetermined System Solver (FOCUSS) | 43 |
| 3.2 Mean Square Extrapolation (MSE) | 46 |
| 3.3 Iterative Adaptive Approach (IAA) | 48 |
| 4.1 Iterative Mean Square Extrapolation (IMSE) | 58 |
| 4.2 Stagewise Gradient Pursuit (StGP) | 65 |

CHAPTER 1

INTRODUCTION

Image formation from synthetic aperture radar (SAR) data is a problem of deconvolution. The received signal is representative of the convolution of the reflectivity of the area being imaged and the transmitted radar waveform. Some of the most common SAR image formation algorithms reduce the deconvolution problem to a 2D spectral estimation problem. Inherent constraints on the region of support for SAR data lead to image artifacts due to resolution and spectral leakage of the spectral estimation method used. This work explains the relationship between SAR image formation and spectral estimation, outlines some issues that might preclude classical spectral estimation methods, and examines modern spectral estimation techniques that may provide better performance. Many of these modern spectral estimators are aimed at providing more accurate estimates of the signal power spectral density (PSD). Their properties include high resolution, low spectral leakage, and in some cases non-ideal sampling conditions.

To simplify the problem description, discussion of SAR image formation will be limited to spotlight mode SAR, where the sensor antenna is kept centered on a single point on the ground as the synthetic aperture is sampled over time. Strip map and scan mode SAR follow slightly different image formation steps, but still rely on spectral estimation to transform the 2D signal into an image.

Spectral estimation may be applied to SAR for the simple purpose of image formation, or to enhance or restore image quality. Image enhancement and restoration applications include, but are not limited to: interferometry [1], bandwidth extrapolation (i.e., super-resolution) [2–4], reconstruction of corrupt or missing

bandwidth [5, 6], impulse response sidelobe control [7, 8], sparse aperture combination [9], and image segmentation [10, 11].

A number of spectral estimators have been developed for the purpose of sidelobe control and super-resolution. Classical estimators include the sample spectrum, the periodogram and the correlogram [12]. Modern estimators may be grouped as either parametric or non-parametric. Parametric estimators include MUSIC [3], ESPRIT [13–15], matrix pencil methods [16–20], and AR linear prediction [1, 7, 21]. Non-parametric estimators include Capon methods [6, 22], Prony methods [23] and many sparse decomposition [24–33] and compressive sensing estimators [34–37].

Compressive sampling reconstruction algorithms have been a popular topic in the last decade. Much of the early work in compressive sensing was focused on reconstruction from projections, as in tomography and MRI [38, 39]. Application to radar and SAR has been addressed, but usually for the purpose of super resolution [35, 40, 41]. For example, the work in [42] is based on the assumption of specific radar waveforms and only addresses reconstruction of point scattering returns, not clutter or the effects speckle noise. In contrast, the focus of this work is to address the applicability of sparse decomposition for reconstructing radar imagery that has arbitrary irregular sampling, addressing point target returns versus diffuse clutter returns, and testing with real-world SAR data. Instead of dealing with small-scale simulations where computational efficiency and storage requirements are not limited, this work is directed at algorithmic implementations that are feasible for both large high-resolution images and modest computational requirements.

Ultimately, all of the above applications reduce to the problem of image restoration by deconvolution. The point spread function (PSF) to deconvolve is highly ill-conditioned, in that its transfer function has many zeros (i.e., it is sparse). Research in the last decade has shown that such image restoration may be accomplished if the image characteristics obey the criteria of compressive sensing, that is, the image domain itself is sparse and the sampling transform matrix meets restricted isometry principle (RIP) conditions [43, 44]. In general, most compressive sensing algorithms attempt to solve an underdetermined optimization problem with an image sparsity constraint. Much work on reconstruction algorithms in the last couple of years centers around methods rooted in iterative deconvolution via alternating projections, whether they attempt to solve the stabilized inverse problem directly or indirectly.

This thesis contributes to the fields of SAR imaging and sparse decomposition. Two algorithms, stagewise gradient pursuit (StGP) and iterative mean square extrapolation (IMSE), are developed by combining techniques applied to stagewise orthogonal matching pursuit (StOMP [29]), gradient pursuit (GP [45]), mean square extrapolation (MSE [46]), the iterative adaptive algorithm (IAA [33]) and the focal underdetermined system solver (FOCUSS [30]). The StGP and IMSE algorithms leverage elements from these existing algorithms that reduce computational complexity and improve the rate of convergence to a solution, which is a necessity for near real-time processing of high resolution SAR. StOMP enhances convergence by selecting multiple dictionary atoms per iteration, but requires an expensive calculation to orthogonalize the solution at each iteration, while GP approximates the

orthogonalization step in basic OMP by leveraging conjugate gradient search directions, creating lower computational complexity, but with more iterations required for convergence because only one atom is selected at a time. The StGP algorithm goes one step further by combining multi-atom selection with conjugate gradient search directions, yielding an even more efficient algorithm. The IMSE algorithm draws from IAA and IMSE. The IAA algorithm achieves a sparse solution by iteratively updating the estimate signal covariance, but requires expensive matrix inversions to estimate each possible spectral component. The MSE algorithm, a special case of FOCUSS, provides estimates for all possible spectral components with a single matrix inversion. The IMSE algorithm leverages MSE for calculating spectral estimates and iteratively updates the covariance estimate before applying MSE again, similar to IAA. The StGP and IMSE algorithms could enable improved compression of complex valued imagery (i.e., both amplitude and phase information) or sparse aperture collection, allowing for more imaging area at lower data rates.

This thesis also contributes to the application of sparse decomposition in SAR imaging by demonstrating the StGP and IMSE algorithms on real-world SAR data, whereas existing literature is focused on simulated and laboratory data that is not typical of true SAR imagery. Existing sparse decomposition algorithms have been applied to ideal point scattering targets in the absence of natural clutter, where the effects of speckle noise are pronounced. This work takes the additional step of applying sparse decomposition to coherent diffraction limited images with speckle noise, and investigates the appropriateness of sparsity constraints in the presence

of speckle noise. This work also demonstrates that the StGP and IMSE algorithms result in a decomposition of a real-world SAR image into target and clutter components. A target/clutter decomposition can serve as inputs to target identification and speckle reduction algorithms. Image decomposition, as a byproduct of sparse decomposition in SAR, has not been addressed by the current literature.

This thesis is organized as follows. Chapter 2 is a detailed background on spotlight SAR image formation, its relationship to spectral estimation, and a statement of the missing data problem to which certain spectral estimators are later applied. Chapter 3 is a review of several types of spectral estimators, their properties, and how they have been applied to SAR. Chapter 4 presents the proposed StGP and IMSE algorithms for SAR. The performance of the proposed reconstruction algorithms is examined in relation to the missing data problem, in the context of both uniformly downsampled data and a non-uniform compressive sampling framework. Chapter 5 describes the performance of the two implemented algorithms through simulations and application to real-world SAR data. Finally, a conclusion is provided in Chapter 6, along with areas for future work.

CHAPTER 2

SAR BACKGROUND

This chapter presents a background of synthetic aperture radar (SAR) and the processing necessary to form an image from the collected data. It is necessary to have a basic background in how SAR data is collected and the processing necessary to form an image from SAR data in order to understand why SAR imaging is a spectral estimation problem. The issues and limitations that apply to spectral estimation translate directly to SAR image quality. Discussion is limited to spotlight mode SAR due to its popular use in high resolution applications. Spotlight mode is also similar to tomographic imaging, for which much research and algorithm development have been conducted for medical applications. This chapter begins with a description of data acquisition for spotlight mode SAR. A tomographic model for the data is then presented. The chapter continues by explaining the signal processing required to obtain an image from spotlight mode SAR data. It will be shown that the image formation process is a spectral estimation problem, and that SAR data can be modeled as a sum of complex exponentials. The chapter concludes with some of the spectral estimation issues that impact image quality.

2.1. *Spotlight mode SAR*

The goal of SAR is to determine some representation of an area of ground based on the reflection of radio waves. The reflectivity of the ground may be described by the continuous mathematical function $g(x, y)$, where y is the range direction and x is the cross-range direction (or azimuth). In radar applications, the range dimension is the direction of radio wave propagation from a directional antenna, while azimuth is orthogonal to the range dimension. The function $g(x, y)$ is generally considered




Fig. 2.1. Radar azimuth resolution limit.

to be complex, as the reflectivity of objects affect both the amplitude and phase of the incident radio waves. Typically, the radar sensor provides its own illumination of the ground with some radio frequency (RF) waveform.

Early radar imaging systems suffered from poor resolution in the estimate of the ground reflectivity map $g(x, y)$. Resolution in the range dimension y , given by the equation $\delta_r = \frac{c}{2 * BW_{RF}}$, where c is the speed of light, could easily be improved by increasing the RF bandwidth (BW_{RF}), especially when applying pulse compression waveforms [47]. However, resolution in the azimuth direction is determined primarily by the antenna beam width, $\theta_{3dB} \propto \frac{\lambda}{L}$, which is proportional to the wavelength λ and inversely proportional to the antenna length L . Assuming a small angular beam width, the azimuth resolution $\delta_a \simeq R\theta_{3dB}$ is determined by the length of the arc subtended by the beam width, and thus degrades with the stand-off range R between the sensor and the area being imaged (see Figure 2.1).

Achieving high resolution in azimuth would require an unrealistically large

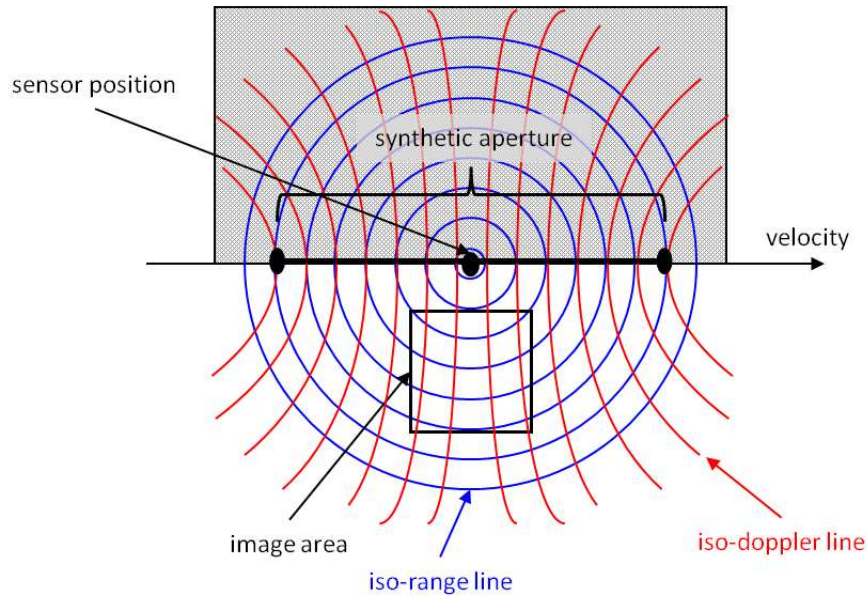


Fig. 2.2. Mapping ground points via range/Doppler method with SAR.

antenna for a mobile platform, which is necessary for SAR data acquisition. One solution to this problem is to sample a much larger synthetic antenna aperture with one or more smaller real antennas. Combining the signals sampled at various positions along the synthetic aperture allows for azimuth resolution that is more consistent with the theoretical beam width of the synthetic aperture. Figure 2.2 shows that, by moving an antenna along a synthetic aperture path, ground points to one side of the sensor platform may be uniquely resolved in two dimensions by their range location or pulse echo time delay, and their Doppler offset relative to the sensor, the determination of which requires pulse integration along the sensor path (i.e., synthetic aperture).

There are two very common modes of collecting SAR data. One mode is stripmap mode, where a fixed side-looking real antenna beam is dragged across a

scene of interest while transmitting and receiving pulses. The other is spotlight mode, where the aim point of the antenna is fixed at the center of the scene of interest as the SAR platform traverses the synthetic aperture. The SAR platform is simply the vehicle that carries the radar payload and support systems (i.e., an airplane). The azimuth resolution in SAR is dependent on the length of the synthetic aperture over which the pulses are integrated, more commonly expressed as the integration angle θ_I . The integration angle for stripmap SAR is limited by the beamwidth of the real antenna as it is dragged over the scene; thus, for a higher azimuth resolution, it is actually necessary to have a smaller antenna, but also requires more transmit power to maintain the desired signal-to-noise ratio (SNR). Power restrictions are very tight on mobile platforms, so spotlight mode SAR is often employed when high resolution is desired. Spotlight mode SAR can dwell on a scene to obtain a larger integration angle at the expense of imaging less scene area, as restricted by the beam footprint. The azimuth resolution of spotlight mode SAR is given by $\delta_a = \frac{\lambda}{2*\theta_I}$ [48]. The difference between stripmap and spotlight mode SAR is shown in Figure 2.3.

The complete SAR data is multi-dimensional in nature, as the SAR data consist of one sampled pulse vector for each transmit/receive pulse that samples the synthetic aperture. Let the index $i = 0, 1, 2, \dots, N - 1$ represent the pulse number along the synthetic aperture. Then, the received SAR signal may be expressed as $r_i(\mathbf{x}(t_2), t_1)$, where t_1 is the intra-pulse time index (fast-time), t_2 is the inter-pulse time index (slow-time). The position $\mathbf{x}(t_2)$ represents the sample locations of the synthetic aperture in space and is dependent on the inter-pulse time, as the real

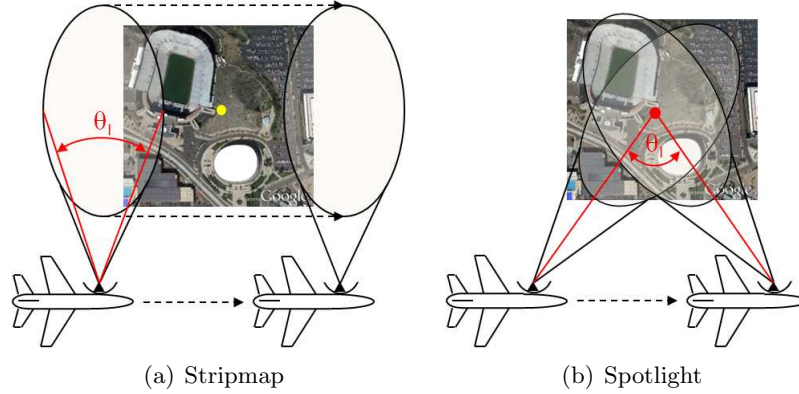


Fig. 2.3. Stripmap versus spotlight SAR modes.

antenna flies along the synthetic aperture. Assuming the sensor is flying a straight path along the cross-range dimension at constant altitude, the received signal becomes $r_i(\mathbf{x}(t_2), t_1)$. It may be noted that this setup resembles a sensor array, except that sampling in space is accomplished with a single sensor, and the signal received at each position in the array represents radar returns from separate transmit pulses $s_i(t_1)$. Each transmit pulse is the same waveform and is a function of t_1 only, but the entire collection of pulses is acquired by a single sensor at different times in t_2 , unlike an array of sensor that would record the same return pulse at different positions. Processing of SAR data requires pulse-to-pulse coherence, or precise time and phase delay knowledge associated with system hardware and propagation effects, in order to properly integrate the received pulses as you would sensor array data. By coherence we mean that the relative phase information between every pulse collected is maintained over the entire aperture. Figure 2.4 depicts the signals received along the synthetic aperture array.

It may now be clearer how high resolution in azimuth (x -axis) may be achieved,




Fig. 2.4. Collection of spotlight mode SAR data.

in light of beamforming theory in array processing. Azimuth resolution can be improved by increasing the length of the synthetic aperture, but the inter-pulse sampling rate (or pulse repetition frequency PRF) must be high enough to avoid aliasing in the array response. The array response may be viewed as the antenna pattern of the synthetic aperture, which is periodic since it is sampled in space. An antenna pattern is generally defined as the radiation pattern of a directional antenna, or rather the gain in radiated power versus angular direction from the antenna. The beamwidth is derived from the antenna pattern.

2.2. Tomographic model of spotlight mode SAR

At this point it, is instructive to introduce a mathematical model for each received pulse. It shall be assumed that the scene is in the far field of the real antenna and that the wave fronts are planar as opposed to spherical. The transmit waveform is demodulated from each received pulse, and each demodulated receive pulse may be represented as the projection of the scene reflectivity along lines perpendicular to the direction of wave propagation, u , and modulated by the transmit waveform. The resulting projection function is given by:




Fig. 2.5. Projection-slice principle for a single SAR transmit/receive pulse.

$p_{\theta}(u) = \int_{-L}^L g(u \cos(\theta) - v \sin(\theta), u \sin(\theta) + v \cos(\theta)) dv$ [49], where $g(x, y)$ is the ground reflectivity. The integration bounds $-L$ to L is derived from the extent of scene illumination due to the beam footprint width in the v direction. This projection observation bears a striking resemblance to tomographic imaging principles, and indeed the projection-slice theorem does apply to SAR data. The reader may refer to the text by Jakowatz for a comprehensive treatment of spotlight SAR from a tomographic perspective [49]. Figure 2.5 depicts the projection-slice principle for SAR data collection.

According to the projection-slice theorem, the projection of target reflectivities on the same iso-range line is equivalent to sampling in the frequency domain. Figure 2.6 shows the Fourier domain of a SAR pulse return according to the projection-slice theorem. Each pulse return is equivalent to a slice of the Fourier transform of the

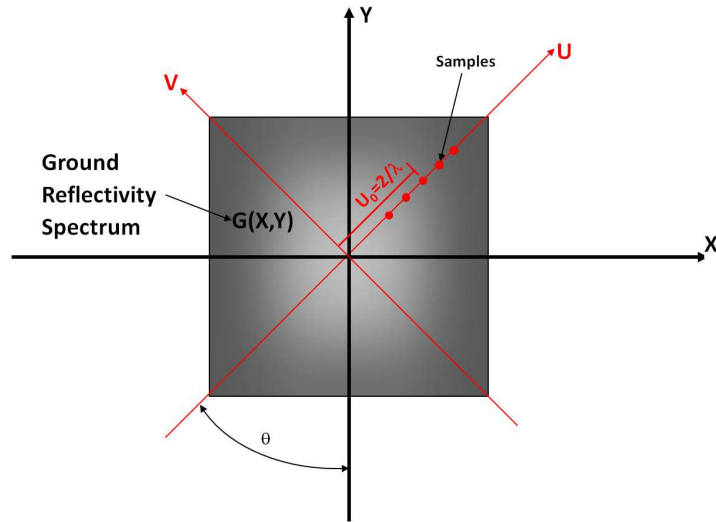


Fig. 2.6. Sampling in the Fourier domain according to the tomographic model of spotlight SAR. U_0 is the radial center frequency of the image Fourier domain collected in polar coordinates.

image area. After collecting a number of pulse returns along the synthetic aperture, we have sampled the Fourier domain via slices at various projection angles. The collection of receive pulses obtained from the SAR sensor is commonly referred to as phase history, in reference to the phase response of targets in the scene over the synthetic aperture. It is important to note that the collected phase history for spotlight mode SAR is equivalent to sampling of the image spatial frequency spectrum, also referred to herein as the Fourier domain, as depicted in Figure 2.6.

Each received pulse is a slice of the ground reflectivity spectrum $G(X, Y)$. The integration angle θ_I is the difference in projection angles between the first and last pulse collected over the synthetic aperture. The accumulation of pulses over an integration angle θ_I results in 2D phase history, which is a sampling of the ground reflectivity spectral content. The SAR system samples the Fourier domain of the

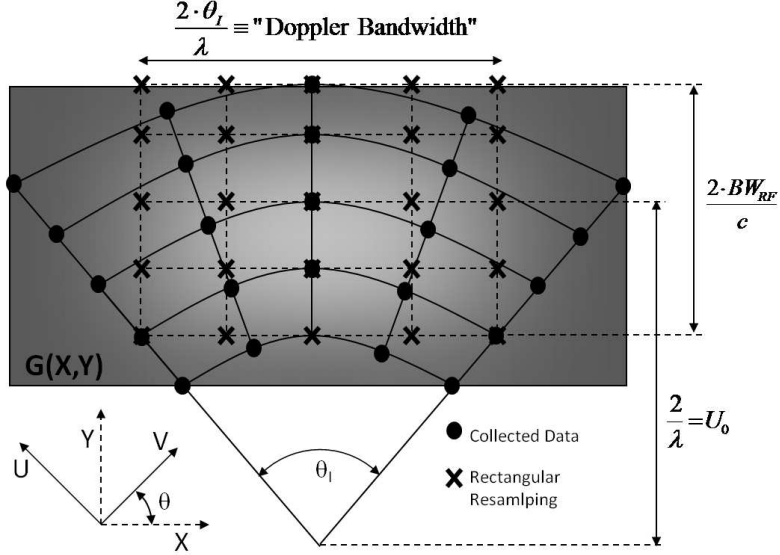


Fig. 2.7. Phase history sampled as a polar annulus in the Fourier domain. U_0 is the radial center frequency of the image Fourier domain collected in polar coordinates.

ground reflectivity $G(X, Y)$ in polar coordinates (U, V) . Figure 2.7 illustrates the resulting polar annulus of data.

2.3. Processing spotlight mode SAR

First, we will examine the processing of individual pulses in the range dimension, along the direction of the pulse propagation. It is common to transmit linear frequency modulated (LFM) waveforms, as depicted in the upper time-frequency plot in Figure 2.8. The equation for the LFM waveform is given by: $s_i(t_1) = e^{[j(\omega_0 t_1 + \alpha_{FM})]}$, where ω_0 is the transmit center frequency, α_{FM} is the FM chirp rate, and t_1 is the fast-time index that spans the transmit pulse width. Transmitting the LFM waveform allows for larger transmit RF bandwidth, and thus finer range resolution, but with longer pulse durations and lower peak power requirements. The receiver may either apply a matched filter to demodulate the




Fig. 2.8. Time-Frequency plot of a transmit and receive pulse through deramp and deskew processing.

received pulse return, or deramp processing may be used, whereby the received pulse is modulated with a reference LFM pulse that has the same chirp rate as the transmit LFM pulse [48]. Depending on the collection parameters, it is common to deramp the received signal to an intermediate frequency (IF) signal that has a smaller bandwidth than the RF bandwidth of the chirped pulse return echoes, as in the middle plot in Figure 2.8. The deramped receive pulses may then be sampled in quadrature, providing complex valued samples from which phase information can be extracted. The range reflectivity projections may then be obtained from the frequency content of the deramped IF signal. The bottom time-frequency plot in Figure 2.8 shows that the reflectivity of target echoes in the range dimension directly corresponds to frequencies in the deramped IF signal.

After deramp processing, the Fourier transform (FT) of each pulse represents the projection $p_{\theta}(u)$. Recall that this projection function is the integration of the

scene reflectivity $g(x, y)$ along iso-range lines (dimension v in Figure 2.5). We will now examine how ground targets lying on a particular iso-range line are resolved by SAR. For this, we must look at sampling of the synthetic aperture from pulse to pulse. Let us consider two points in $g(x, y)$ on the same iso-range line, one at scene center and one at the scene edge. The phase of the chirp echo at the real antenna for a point in $g(x, y)$ is proportional to the round trip time delay. For any single return pulse received, all points on the same iso-range line will share the same propagation phase delay at the receiver. However, those same points will have non-equal phase delays for subsequent pulse returns, because the iso-range dimension v rotates as the collection angle θ rotates from pulse to pulse. Each point on the ground may be uniquely distinguished by the change in its propagation phase delay over the integration angle θ_I . Assuming a straight line sensor flight path and small integration angle θ_I , the phase of a pulse echo from scene center will vary as an approximate quadratic function over the synthetic aperture length. Likewise, other points at the same azimuth coordinate x will have an approximate quadratic phase response, but with an additional unique phase component as well. This added phase component is due to the fact that the rate of change in range from the sensor to targets offset in azimuth varies differently across the synthetic aperture. We shall refer to the change in range from the sensor to a point in $g(x, y)$ as the range profile of that point. Figure 2.9 shows two points in $g(x, 0)$ with different range profiles versus the sensor position in x . Application of a matched filter in the azimuth dimension can compress the pulse-to-pulse samples into ground reflectivity in azimuth. Usually, azimuth compression is done via the




Fig. 2.9. Differing range rates (profiles) for targets offset in azimuth.

Fourier transform, due to intermediate processing that correlates frequency in slow-time t_2 to azimuth in $g(x, y)$.

For a particular range resolution cell, the targets may be resolved in azimuth by matching a filter to the expected phase history response for a desired azimuth offset. The expected phase response is derived from the range profile of a point in $g(x, y)$ and the phase delay function $\phi = 2\pi \frac{\mathbf{R}}{\lambda}$. Every point in the scene has a unique range profile, with approximate linear and quadratic components, as seen in Figure 2.10. The ensemble of pulse echoes collected along the synthetic aperture is referred to as the phase history. A SAR system records the phase history, and an image of the scene is formed via matched filtering, where the filter is matched to the unique phase history expected of a target at a particular position in the scene. Collecting more phase history (i.e., a larger integration angle) allows for a finer resolution of targets in azimuth, as the synthetic aperture length is increased and the beamformer resolution improves. Intuitively, collecting more phase history over a larger range of projection angles θ allows for more variation between range

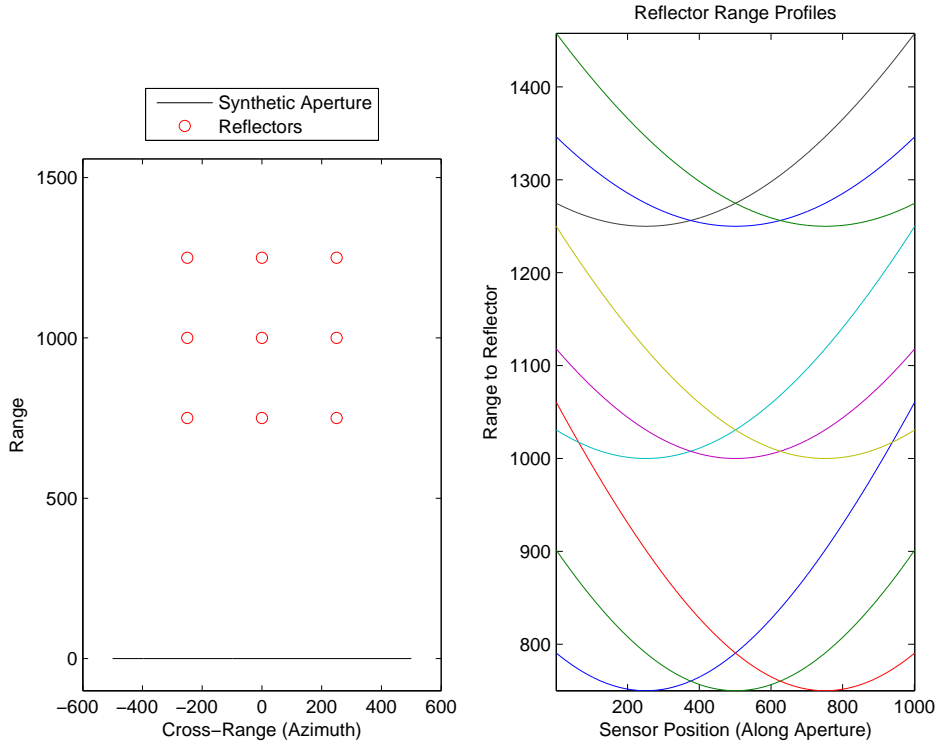


Fig. 2.10. Differing phase history profiles for different points in the image.

profiles of points offset in azimuth, and thus they are easier to differentiate from each other. A shorter wavelength will also yield better azimuth resolution, but there tends to be more restriction on available operating frequencies.

The most direct option for determining the reflectivity response for a point in the scene is to design a unique two-dimensional matched filter for each desired coordinate in the scene being imaged. This method can be extremely inefficient, especially for large images, so it is common to perform some intermediate processing to condition the phase history data such that all points in the imaged scene may be focused simultaneously with a single matched filter. The first step is called motion compensation, which multiplies each pulse response with a different phase function

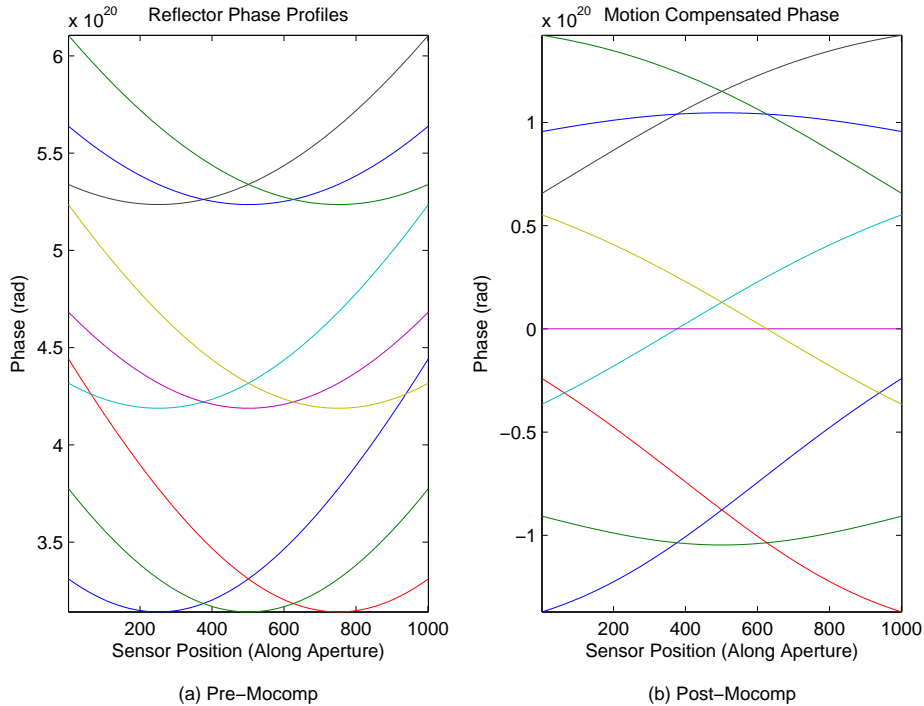


Fig. 2.11. Effect of motion compensation processing on phase history targets. (a) Pulse echo phase delay of points at varying positions in the scene versus sensor offset position over a synthetic aperture; (b) Motion compensation applies a phase function to each return pulse such that a point at the scene center has zero phase variance over the synthetic aperture.

such that a point at the scene center has a zero phase response and corresponds to the DC component of the resulting 2D phase history SAR signal. Figure 2.11 demonstrates the effect of motion compensation on the phase history at different points in the scene.

After motion compensation, the target at scene center will have zero phase response, and targets displaced from the center will have unique residual phase responses, each representing a unique frequency. The residual phases have a linear component, due to azimuth offset, and higher order phases due to migration of the targets through resolution cells.

Note that the received phase history has a natural rectangular window in the azimuth dimension, thus the impulse response (IPR) in azimuth is given by the sinc function. Likewise, the time sample support of each pulse return is finite, so the IPR in range is also given by a sinc function. This is illustrated in Figure 2.12. In fact, the image IPR (point spread function) is a 2D sinc function. After deramp and motion compensation processing, the reflectivity response due to a target at the center of the scene has a zero phase response and is represented by the DC term in the 2D Fourier transform of the phase history. Higher frequency components are due to target returns increasing in distance from the scene center.

2.4. *Spectral estimation and SAR image formation*

In the previous two sections of this chapter, it is shown that data obtained from spotlight mode SAR provide samples in the Fourier domain of the scene being imaged. The image is recovered from the frequency content of the sampled data, where points in the reflectivity map are given by an estimate of the 2D phase history signal spectrum. A popular image formation algorithm for obtaining the spatial frequency estimates is the Polar Reformating Algorithm (PFA) [49], which will now be described.

Referring to Figure 2.7, it is evident that the phase history collected over the synthetic aperture is not uniformly sampled in the Fourier domain $G(X, Y)$. However, a resampling of the data to a uniform rectangular grid in $G(X, Y)$ will allow for the image to be retrieved efficiently with an FFT. Once the phase history is sampled to a uniform grid, the FFT is the matched filter for every output point in the desired image space. Resampling from polar to Cartesian coordinates in the

Fourier domain and then applying a single 2D FFT can be much more efficient than applying a different 2D matched filter to obtain each image sample as an estimated spectral component. The PFA algorithm does assume planar wavefronts impacting the scene, which is only an approximation in reality, but does focus points in the scene that migrate through range cells so that larger areas may be imaged. The PFA algorithm also assumes that polar sampling of the image Fourier domain meets Nyquist criteria. Regardless of the image formation processing used, PFA or brute force matched filtering, the extent of the sampled Fourier domain is limited in both range frequency and azimuth frequency, and thus spectral leakage is to be expected in the image domain. Likewise, the maladies associated with inadequate Nyquist sampling will impact the image domain, irrespective of the image formation algorithm applied.

2.5. *Problem statement*

As mentioned previously, an image obtained from a SAR sensor is given by the spectral content of the raw 2D data processed and formatted in the (X, Y) spatial frequency domain. Several observations about the nature of the collected SAR phase history data can be stated based on the previous discussion.

First, the signal being sampled is bandlimited. The frequency components in the phase history are determined by the spatial extent of the scene being imaged, which is limited by the area that can be illuminated by the SAR sensor. The beam pattern of the real antenna and its footprint on the ground determine how much area is illuminated by the transmit pulse and return echoes that contribute to the received signal. To say that the phase history is bandlimited is equivalent to saying

that we only receive echo energy from a finite swath of ground area. From Figure 2.8, we see that the reflectivity of targets spaced in the range dimension correspond to particular frequencies in the IF bandwidth following deramp processing. Since the receiver IF bandwidth is limited, only a limited swath in the range dimension may be imaged unambiguously, or without aliasing between near range and far range targets. The fast-time (t_1) sample rate, i.e., intra-pulse period or A/D rate, determines how much of the available IF bandwidth is sampled unambiguously. Likewise, the greater the slow-time (t_2) sample rate is, i.e., inter-pulse period or PRF, the more azimuth swath may be imaged unambiguously.

The second important observation is that the sampled SAR signal is finite in both slow-time (t_2) and fast-time (t_1). Since we are dealing with a pulsed radar system, there is a finite receive window between transmit pulses. Also, we can only transmit and receive a finite number of pulses. There is a point of diminishing benefit from collecting more pulses returns, as the scene content is generally non-stationary over time, and integrating over more than 180° provides no additional resolution. To extend the synthetic aperture indefinitely along a straight line is not possible, as the scene being imaged will disappear from view due to the curvature of the Earth. The implication of finite signal support in the SAR phase history is a limit on image resolution in both the range and azimuth dimensions.

One problem resulting from the spectral estimation step is a resolution limit and spectral leakage. Due to the fact that the SAR 2D phase history has a limited spatial frequency support, the IPR in the image domain will have a limited resolution. The IPR will also exhibit sidelobe behavior due to spectral leakage. Less

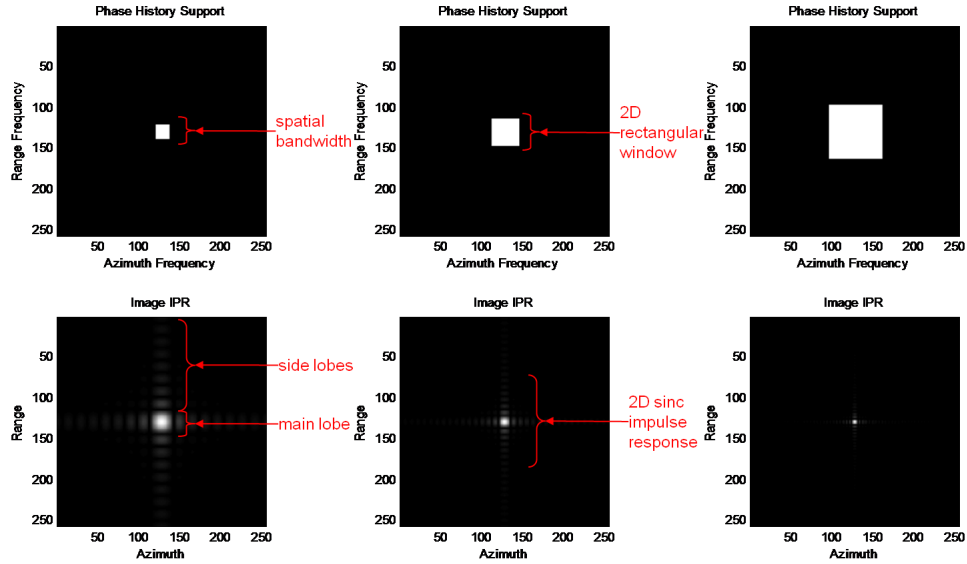


Fig. 2.12. Phase history support and resulting 2D sinc IPR.

resolution means closely spaced scatterers or reflectors in the scene will be indistinguishable in the image. The sidelobes from scatterers with large echo returns (i.e., large radar cross-section or RCS) may occlude the responses from nearby scatterers with smaller RCS. Spectral leakage occurs when the scattering center of a target is between resolution cells in the image, such that the target cannot be represented by a single pixel. Assuming a classical spectral estimator and a rectangular region of phase history support, Figure 2.12 shows the expected 2D sinc IPR response. Notice the dependence of resolution, determined by the main lobe width, on the amount of phase history support. Also, observe the sidelobes spread in a star pattern about the main lobe.

Another problem arises if the phase history is irregularly sampled with respect to Nyquist criteria. If the SAR sampling conditions do not meet Nyquist criteria

for the extent of area being imaged, then artifacts will be introduced into the spectral estimates. For a typical SAR image, these artifacts might cause aliases or an increased multiplicative noise ratio (MNR).

One situation that might lead to irregular sampling is sparse aperture imaging. The designer of a SAR system may want to increase the inter-pulse receive window to collect more range swath. The trade-off is reducing the PRF, which would normally mean less unambiguous azimuth swath for imaging. The PRF distance, or periodicity in the sampled synthetic beam pattern, is reduced and increases the level of aliased energy. With an active electronically scanned array (AESA) antenna, the SAR designer may want to monitor multiple areas of interest simultaneously by alternating dwell (integration) time between different ground scenes and to maintain azimuth resolution by combining non-contiguous synthetic apertures. If operating in a congested RF band, the sensor might need to operate in non-contiguous RF bands. One can either choose a smaller RF bandwidth from which to form an image, resulting in reduced range resolution, or combine non-contiguous RF bands at the risk of introducing artifacts into the image. In general, sparse aperture imaging would involve inadequate sampling at the time of data acquisition, impacting the SAR image formation processing at the earliest stages.

Unexpected errors or noise in the system could result in situations where it is beneficial to ignore some of the phase history samples during image formation processing. Any real-world system is subject to failures of sensor hardware, data transmission or data storage. The signal may deviate from receiver limits resulting in nonlinearities in the sampled signal. Storage devices might fail or experience data

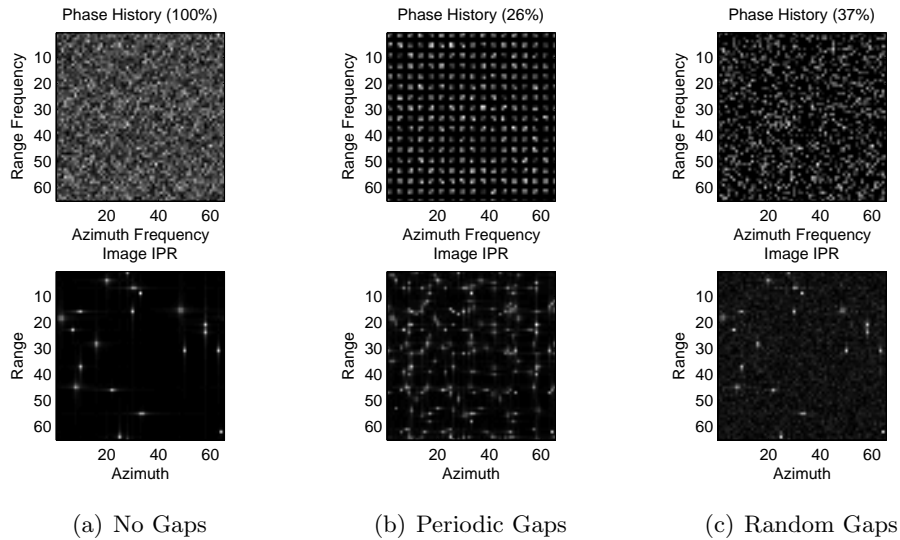


Fig. 2.13. Effect of irregular sampling on IPR in image.

corruption. Data samples are subject to transmission errors imposed by restrictions and limitations of the communication channel. Interference from other RF sources can add unwanted noise to the SAR image. In general, corrupt data samples will decrease the SNR of the system. Excluding the corrupt samples from image formation will get rid of the added noise, but at the expense of introducing spectral estimation artifacts.

Some of the effects of irregularly sampled phase history can be seen in Figure 2.13. In this example, we have twenty point targets spaced randomly in the image, taken as the sample spectrum of the phase history. Periodically deleting samples in range frequency and/or azimuth frequency results in aliases of the targets being created in the associated dimension(s) of the image. Distinguishing between actual targets and their aliases becomes a problem. If random samples of the phase history are deleted, there is an increase of the noise floor, making it difficult to detect weaker

targets in the image. The goal is to find a spectral estimator that can reconstruct the image efficiently and without the artifacts observed in this example.

2.6. *Summary*

In this chapter, the collection and processing was described for a typical spotlight mode SAR system. A tomographic model of spotlight mode SAR reveals that the collected phase history is equivalent to sampling the spatial frequencies of the ground reflectivity map $g(x, y)$ being imaged. The image formation process is shown to be a spectral estimation problem. Sampling conditions that affect resolution, accuracy and quality of spectral estimators translate into SAR image quality. Conditions for insufficient or irregular sampling in SAR are postulated. The same problems that are experienced by spectral estimators due to irregular sampling also impact SAR image quality.

Recent work in modern spectral estimators and sparse sensing may provide answers to some of these problems. Assuming the missing or corrupted samples are known or can be identified, a number of proposed algorithms might be able to reconstruct a SAR image from irregularly sampled phase history.

CHAPTER 3

SPECTRAL ESTIMATORS AND SAR

This chapter presents a more detailed look at known spectral estimators and their applicability to SAR imaging. As mentioned previously, an image obtained from a SAR sensor is given by the spectral content of the raw 2D data processed and formatted in the (X, Y) spatial frequency domain. The simple sample spectrum given by the 2D DFT is by far the most common spectral estimator used to retrieve an image representation of the reflectivity map $g(x, y)$. Other classical estimators, namely the periodogram and variants thereof, have also been utilized to reduce noise at the expense of resolution. However, these classical estimators have some inherent limitations on spectral resolution and susceptibility to side lobes and spectral leakage. They also suffer undesirable artifacts when the data is not ideally sampled according to Nyquist criteria. Modern parametric spectral estimators such as auto-regressive methods, MUSIC, and ESPRIT have super-resolution properties, but require some a priori knowledge of the number of spectral components. Refined filter bank methods are non-parametric modern spectral estimators that also have super-resolution properties, some variations of which can handle irregular sampling conditions, but are very processing intensive. Recent years have seen the rapid development of compressive sensing algorithms, some of which have been utilized to demonstrate super-resolution with synthetic SAR data. The authors in [50] provide a concise overview of classical and modern parametric estimators, and present an iterative hard thresholding based compressive sampling algorithm [51, 52], but only address the 1D case of line spectrum estimation. Many compressive sensing algorithms have good performance, but are processing intensive and are currently

applicable only in a laboratory environment with high performance computers.

3.1. *Classical estimators*

The Fourier method has several advantages making it an attractive spectral estimator, namely a computationally efficient implementation with the FFT and being a separable unitary transform. Key disadvantages are side lobes and point spread of the impulse response (IPR) due to band limiting, as well as speckle noise due to the coherent nature of the collected data. Various windowing methods can be applied with the FFT to reduce side lobes at the expense of spectral resolution. Adaptive windowing techniques, referred to as apodization in the literature, have also been developed with the aim of eliminating side lobes [7, 8]. The limitation stands from the fact that no DFT method alone can provide resolution better than $(\frac{2\pi f_s}{N})$, where f_s is the sampling frequency and N is the signal support in samples in any given dimension.

Yet another major limitation of Fourier spectral estimators is their susceptibility to non-ideal sampling situations. Real-life SAR systems can suffer partial or temporary hardware failures during data acquisition, corruption of data transfers, signal distortion due to receiver limitations, and interference from other sources operating in the same RF band. Or one might be interested in controlled sparse sampling scenarios with the aim to improve imaging area rate or reduce resource requirements without sacrificing resolution [9]. This is where spectral estimation techniques with super-resolution properties can be useful, as they involve implicit extrapolation or interpolation of the SAR signal region of support.

3.2. *Nonlinear windowing*

The apodization technique introduced in [8] has been extended to achieve some sense of super-resolution, and is called the Super-SVA algorithm. Here, a nonlinear apodization technique is applied to the method of energy error reduction [5] for super-resolution. The energy reduction method, or alternating projections, is described by Jain [46]. Nonlinear apodization truncates the IPR side lobes in the image, thus extrapolating the signal support in the frequency domain. This step is followed by a pseudo-inverse filter matched to the truncated IPR response. We may relate this to a deblurring filter, or deconvolving with the truncated IPR, and the end result is sharpened main lobes sampled on a denser grid. The author in [1] demonstrates application of this method, as well as linear prediction (LP), minimum variance, MUSIC, and parametric maximum likelihood estimators to SAR data. This work looks at several other modern spectral estimators not covered therein.

There have been much work on other modern spectral estimators with super-resolution characteristics that may be more accurate and flexible than the Super-SVA algorithm, especially in situations involving missing data in the signal region of support. This will adversely impact the Fourier spectral estimates and cause IPR distortion. This IPR distortion makes apodization techniques in the spatial (image) domain ineffective. We will now look at two groups of modern spectral estimators: parametric subspace methods and non-parametric methods.

3.3. *Modern parametric estimators*

A high-resolution SAR signal $y(m, n)$ can be reasonably modeled as a sum of complex sinusoids with additive noise [49]:

$$y(m, n) = \sum_{k=1}^K \alpha_k e^{j(\omega_{k1}m + \omega_{k2}n)} + \eta(m, n) \quad (3.1)$$

$$y(m, n) = x(m, n) + \eta(m, n) \quad (3.2)$$

where $0 \leq m \leq M - 1$ and $0 \leq n \leq N - 1$, $(\omega_{k1}, \omega_{k2})$ is the 2D frequency of the k^{th} complex sinusoid, α_k is its complex amplitude, and $\eta(m, n)$ is a white Gaussian noise. The complex sinusoid signal model represents a sum of impulses in the image domain. The noise term $\eta(m, n)$ is dominated by returns from low back-scatter reflectors on the ground (e.g., vegetation). The complex sinusoid assumption only applies to returns from targets with a large radar cross-section (RCS) relative to their surroundings.

There exist several parametric estimators that attempt to determine the frequency and amplitude of the complex sinusoids in this model. The reader may refer to [12] for a comprehensive treatment on these methods in the 1D case. Others have proposed 2D extensions of these methods [15, 17]. Parametric spectral estimators based on subspace decomposition have become popular because of their high resolution, accuracy and consistency. The ESPRIT estimator is a favored method, and its 2D extension is described by [15]. The 2D ESPRIT estimator finds frequency components in ω_{k1} and ω_{k2} independently of each other. However, there remains ambiguity as to what combinations of ω_{k1} and ω_{k2} comprise true $(\omega_{k1}, \omega_{k2})$ pairings.

Closely related to, but more general than the ESPRIT method, are matrix pencil methods. The matrix enhancement and matrix pencil (MEMP) method

for estimating 2D complex sinusoid parameters is described in [17]. The authors in [16] develop a modified MEMP (MMEMP) method which they claim to perform more robustly than 2D ESPRIT in the presence of noise. Just like the independent ω_{k1} and ω_{k2} 2D ESPRIT method, the matrix pencil methods need a frequency pairing scheme to estimate the correct $(\omega_{k1}, \omega_{k2})$ pairs [15–17]. A joint $(\omega_{k1}, \omega_{k2})$ 2D ESPRIT method is given in [15].

If frequency estimates can be found for 2D complex sinusoidal components, the corresponding complex amplitudes α_k may be obtained with a least-squares estimate, as described in the nonlinear least-squares estimator presented in [12]. The amplitude estimates $\hat{\alpha}_k$ are given by:

$$\hat{\alpha} = (B^H B)^{-1} B^H \text{Vec}(Y) \quad (3.3)$$

where $B = B_1 \otimes B_2$, and

$$B_1 = \begin{bmatrix} e^{j\omega_{11}} & \dots & e^{j\omega_{K1}} \\ \vdots & & \vdots \\ e^{j\omega_{11}M} & \dots & e^{j\omega_{K1}M} \end{bmatrix}, B_2 = \begin{bmatrix} e^{j\omega_{12}} & \dots & e^{j\omega_{K2}} \\ \vdots & & \vdots \\ e^{j\omega_{12}N} & \dots & e^{j\omega_{K2}N} \end{bmatrix}.$$

In (3.3), the matrix Y corresponds to the sampled 2D signal $y(m, n)$ in (3.1). The above estimate for $\hat{\alpha}_k$ minimizes the function $f = (\text{Vec}(Y) - B\alpha)^H (\text{Vec}(Y) - B\alpha)$, where $\text{Vec}(Y)$ is the stacked column vectorization of the 2D data matrix Y . In the case of missing samples in dimensions M or N , a weighted least-squares estimate form of (3.3) may be appropriate. For example, one optional weighting matrix would give equal weights to observed samples and zero weights to missing samples

when minimizing the error function. The problem with having zero weighting in a least-squares problem is that it creates a very ill-conditioned system, forcing singularities in $(B^H B)$ and thus complicating the calculation of its inverse. A weighting matrix as just described is implemented in the mean square extrapolation algorithm (Algorithm 3.2).

A parametric method for estimating the 2D complex sinusoidal frequencies is nonlinear least-squares (NLS), the 1D description of which can be found in [12]. A multi-modal cost function must be minimized in order to determine the K frequencies, which is a difficult problem. Solving a multi-modal cost function may be avoided if some critical assumptions are made regarding spacing of frequency components in relation to the sampling rate. In particular, if it is assumed that individual sinusoidal components of $y(m, n)$ are spaced no closer than $\frac{1}{M}$ in dimension m , and no closer than $\frac{1}{N}$ in dimension n (i.e., resolvable using the DFT), then the 2D frequencies and their complex amplitudes may be estimated in the image domain using peak detection. This offers super-resolution in some sense, as the IPR response centered at a particular frequency is replaced with an ideal impulse via the nonlinear *argmax* function. However, it is not expected to resolve point scatterers that are not otherwise resolvable in the sample spectrum. The matching pursuit algorithms for sparse decomposition, described later, rely on similar peak detection schemes and rely on the above assumptions to avoid minimization of a multi-modal cost function. However, the matching pursuit algorithms are different from NLS as they depend on a regularization parameter instead of assuming a predetermined number of frequency components. Both techniques should be appli-

cable to interpolation of missing samples, though. Missing samples in the $M \times N$ region of support may be represented as a binary weighting, where a one represents an observed sample and a zero represents a missing sample. This binary weighting may impact the signal to noise ratio in the image domain, but should not impact the location of IPR peaks, since binary weighting has no phase response. This suggests that the peak finding method should be applicable for data interpolation. A binary window in the time/space domain will increase spectral leakage and complicate the peak finding process, but a sanity filtering step implemented in the proposed stage-wise gradient pursuit algorithm (Algorithm 4.2) attempts to mitigate the effect of spectral leakage on peak detection.

While parametric spectral estimators and subspace methods offer arbitrary resolution and can be very accurate, they do require some prior knowledge of how many sinusoidal components K there are in $y(m, n)$. These methods also require computationally expensive SVD calculations, which is currently problematic for near real-time processing. For non-real-time applications, they may be particularly suited for minimizing spectral leakage in the SAR image, target/clutter separation, and signal reconstruction (e.g., extrapolation and interpolation). Section 3.4 describes non-parametric alternatives that do not require a priori knowledge of the number of frequencies K , but still have super-resolution properties.

3.4. *Modern non-parametric estimators*

A modern non-parametric spectral estimator that has been recently developed is called the amplitude and phase estimation (APES) method [6, 22]. It can be summarized as an adaptive filter-bank approach, of which the periodogram and

windowed sample spectrum can be considered special cases. It is adaptive in the sense that, for any particular frequency, an individual filter is designed such that the filtered sequence is a least-squares approximation to a sinusoid signal, yet does not distort the complex spectrum $\alpha(\omega_1, \omega_2)$. For multiple frequencies, a bank of filters is designed, where each filter is refined to each frequency component. Like the parametric estimators discussed in Section 3.3, APES is based on the signal model of Equation (3.1). The analysis behind 2D APES is rather extensive, so only the final result will be presented here, however it is described in full detail in [22].

Let the observed sample sequence be given by $\{y_{n_1, n_2}\}_{n_1=0, n_2=0}^{N_1-1, N_2-1}$. Then, for any particular frequency pair $\alpha(\omega_1, \omega_2)$, we have the following relationship:

$$\begin{aligned}
 y_{n_1, n_2}(\omega_1, \omega_2) &= \alpha(\omega_1, \omega_2) e^{j(\omega_1 n_1 + \omega_2 n_2)} + e_{n_1, n_2}(\omega_1, \omega_2) & (3.4) \\
 n_1 &= 0, \dots, N_1 - 1 \\
 n_2 &= 0, \dots, N_2 - 1 \\
 \omega_1, \omega_2 &\in [0, 2\pi)
 \end{aligned}$$

In (3.4), the amplitude and phase to be estimated are the complex amplitudes $\alpha(\omega_1, \omega_2)$, while $e_{n_1, n_2}(\omega_1, \omega_2)$ represents a residual due to noise and leakage from frequencies other than (ω_1, ω_2) . Let y_{n_1, n_2} be the sampled 2D signal, and \bar{Y}_{l_1, l_2} represent non-overlapping 2D snapshots of y_{n_1, n_2} , such that l_1 and l_2 are the indices corresponding to each snapshot, and M_1 and M_2 is the size of the 2D snapshot in each dimension. The number of snapshots in each dimension is L_1 and L_2 , and the total number of snapshots is $L_1 L_2$. Let $\bar{y}_{l_1, l_2} = \text{Vec}(\bar{Y}_{l_1, l_2})$ be the vectorization of \bar{Y}_{l_1, l_2} , formed by stacking its columns. Also, define $a_{M_1, M_2}(\omega_1, \omega_2)$ to be a column

from the 2D DFT matrix corresponding to the frequency pair (ω_1, ω_2) . The 2D APES estimates $\hat{\alpha}(\omega_1, \omega_2)$ are given by the following equations:

$$a_{M_1, M_2}(\omega_1, \omega_2) = a_{M_2}(\omega_2) \otimes a_{M_1}(\omega_1)$$

$$a_{M_k}(\omega_k) = \begin{bmatrix} 1 & e^{j\omega_k} & \dots & e^{j(M_k-1)\omega_k} \end{bmatrix}^T \quad k = 1, 2$$

$$\hat{\alpha}(\omega_1, \omega_2) = \frac{a_{M_1, M_2}^H(\omega_1, \omega_2) \hat{S}^{-1}(\omega_1, \omega_2) \bar{g}(\omega_1, \omega_2)}{a_{M_1, M_2}^H(\omega_1, \omega_2) \hat{S}^{-1}(\omega_1, \omega_2) a_{M_1, M_2}(\omega_1, \omega_2)} \quad (3.5)$$

$$\bar{g}(\omega_1, \omega_2) = \frac{1}{L_1 L_2} \sum_{l_1=0}^{L_1-1} \sum_{l_2=0}^{L_2-1} \bar{y}_{l_1, l_2} e^{-j(\omega_1 l_1 + \omega_2 l_2)}$$

$$\bar{y}_{l_1, l_2} = \text{Vec}(\bar{Y}_{l_1, l_2}) = \text{Vec} \left(\begin{bmatrix} y_{l_1, l_2} & y_{l_1, l_2+1} & \dots & y_{l_1, l_2+M_2-1} \\ y_{l_1+1, l_2} & y_{l_1+1, l_2+1} & \dots & y_{l_1+1, l_2+M_2-1} \\ \vdots & \vdots & \ddots & \vdots \\ y_{l_1+M_1-1, l_2} & y_{l_1+M_1-1, l_2+1} & \dots & y_{l_1+M_1-1, l_2+M_2-1} \end{bmatrix} \right)$$

$$\hat{R} = \frac{1}{L_1 L_2} \sum_{l_1=0}^{L_1-1} \sum_{l_2=0}^{L_2-1} \bar{y}_{l_1, l_2} \bar{y}_{l_1, l_2}^H \quad (3.6)$$

$$\hat{S}(\omega_1, \omega_2) = \hat{R} - \bar{g}(\omega_1, \omega_2)^H \bar{g}(\omega_1, \omega_2) \quad (3.7)$$

The matrix \hat{R} is an estimate of the covariance of the full sequence $y(m, n)$. The vector $\bar{g}(\omega_1, \omega_2)$ is a filtered snapshot that is the coherent integration of the \bar{y}_{l_1, l_2} , each of which is phase adjusted to maintain phase coherency of a response

due to the frequency pair (ω_1, ω_2) under consideration. However, phase coherency for other (ω_1, ω_2) is not maintained and will tend to average out. The modified covariance estimate $\hat{S}(\omega_1, \omega_2)$ is formed by coherently subtracting the covariance of $\bar{g}(\omega_1, \omega_2)$ from \hat{R} . It is an estimate of the frequency-dependent noise covariance, where any frequency component of $y(m, n)$ not equal to the pair (ω_1, ω_2) under consideration is deemed as noise.

As can be seen in the APES Equation (3.5), spectral amplitude and phase are estimated for each frequency pair (ω_1, ω_2) , one at a time. The APES scheme can be computationally expensive, requiring at least one matrix inversion per frequency pair. However, the authors in [22] have demonstrated that APES has super-resolution qualities, due to its adaptive filter-bank approach. They also develop an extension called Gapped-APES (GAPES) that interpolates gaps of missing observation samples, and another method called Missing Data APES (MAPES) that is able to provide frequency amplitude and phase estimates in cases of arbitrarily spaced missing observation samples. Both GAPES and MAPES are even more computationally intensive than regular APES. The APES methods might be more practical for processing small regions of interest in a SAR image, as points in the ground reflectivity $g(x, y)$ are given by estimates of $\alpha(\omega_1, \omega_2)$, but is not practical for forming complete images that may be thousands of samples in each dimension.

3.5. *Compressive sensing*

Work on compressive sensing has inspired work in recovery of sparsely sampled signals [53]. Part of the original motivation was to recover images from highly undersampled Fourier data in magnetic resonance imaging scenarios, but the authors

in [53] claim that the techniques should apply to tomographic imaging systems as well. As described in the Chapter 2, SAR data acquisition bears some similarities to tomography and obeys the projection-slice theorem. Compressive sensing signal reconstruction methods might be used for signal interpolation or extrapolation and super-resolution if the SAR data meets certain requirements outlined in [53]. These requirements include sparsity in the ground reflectivity function $g(x, y)$ and mutual incoherence of the sample space transform (i.e., the Fourier transform).

The premise behind compressive sensing is that a signal may be reconstructed near perfectly from sparse samples in a dispersive data domain, if that signal has sparse content (i.e., undergoes energy compression) in another transform domain. It is stated in [53] that random frequency samples in the Fourier domain (e.g., SAR signal) adhere to these criteria. Signal reconstruction is performed by searching for a fully sampled signal with minimum $l_1 - norm$ subject to the constraint that the estimated signal matches the observed sparse samples. Optimizing with respect to the $l_1 - norm$ is effective if the compressed signal resembles a line spectrum, or sum of sinusoids. Taking into account observation or measurement noise in the dispersive domain sparse samples, signal reconstruction may be stated as the following optimization problem [43]:

$$\min_x \|x\|_1 \quad \text{subject to} \quad \|Dx - y\|_2 \leq \sigma_e. \quad (3.8)$$

where y is the observed sample sequence, x is the estimate of the sparse transform domain, and D is a decompressing and sub-sampling transform matrix. This implies that the reconstructed signal should have minimum $l_1 - norm$ subject to

the constraint that the reconstruction error among the observed sparse samples is less than the predicted noise level. It is assumed that $\dim(x) \gg \dim(y)$. A condition for signal recovery relates the minimum number of observation samples relative to the number of non-zero components (or impulses) in the compressed signal domain. Let n be the number of observed samples in sequence y , and K be the number of non-zero components in the compressed signal domain $x \in \mathbb{C}^N$, then for some constant C we may state the requirement $n > CK \log N$ [53].

Minimizing the l_1 -norm is actually an approximation to the more ideal solution of minimizing the l_0 -norm, whereby the smallest number of discrete signal components is found such that the reconstruction residual ($\|Dx - y\|_2$) is less than the standard deviation of the noise (σ_e). The l_1 -norm minimization method is popular for its applicability to linear programming. Though much research have been conducted in solutions to the l_1 -norm minimization problem, this method actually is a subset of the more general discipline of sparse decomposition. Sparse signal decomposition is interested in representing the signal as the smallest possible number of discrete components.

Sparse decomposition methods can be categorized under two main objectives: minimizing the l_0 -norm (matching pursuit [24]) or minimizing the l_1 -norm (basis pursuit [28]). Various methods have been developed in each category. The matching pursuit methods are popular for their straightforward implementation via greedy algorithms and include orthogonal matching pursuit (OMP) [24], stage-wise OMP (StOMP) [29], and gradient pursuit (GP) [45], which will be explained in more detail. The basis pursuit algorithms are generally based on linear programming,

such as simplex or interior point methods. The Greedy Basis Pursuit (GBP) algorithm [54] has been proposed for l_1 -norm minimization based on greedy methods. The FOCUSS algorithm [30] minimizes an objective function somewhere between the l_0 - and l_1 -norms.

Matching pursuit algorithms are based on building an approximation to the signal by iteratively combining scaled atoms from a basis dictionary. The atom selected in each iteration is the one with the greatest correlation to the residual signal. The OMP algorithm removes each atom selected from the dictionary, recalculates coefficients for the selected atoms such that the residual is orthogonal to the span of the selected atoms. This orthogonalization step ensures that the sparse approximation to the original signal is optimal in the least-squares sense, with respect to the atomic components selected at the current iteration of the algorithm. The primary fault with matching pursuits is an inherent global sub-optimality, as the greedy approach is only an approximation of the l_0 -norm. Enforcing orthogonality of the residual reduces the sub-optimality of the solution, but adds significant computational complexity. The StOMP algorithm was proposed to speed up convergence to a solution by selecting multiple atoms from the dictionary with each iteration, which reduces the computation time at the expense of yielding a solution that is more sub-optimal. The GP algorithm attempts to approximate the OMP solution by employing a conjugate gradients approach to maintain orthogonality between the selected atoms and the residual, without the need for solving a least-squares inverse problem at each iteration. One of the proposed algorithms used for the simulations in this work combines StOMP and GP in a novel stage-wise

gradient pursuit (StGP) algorithm, which has not previously been demonstrated in the literature. While the sparse decomposition solution resulting from such an algorithm is expected to be more sub-optimal than the straight OMP result, the computational efficiencies inherited from StOMP and GP make it feasible to work with larger images of more than 64x64 on a typical modern PC. The proposed StGP algorithm is described in more detail in Chapter 4.

While matching pursuit algorithms depend on the greedy method, most basis pursuit algorithms are built upon linear programming methods. It has been shown that, under certain conditions on the dictionary and sparsity of the signal being analyzed, the solution to the l_1 - *norm* minimization problem will recover the signal exactly, even if the sampling does not meet the Nyquist criteria [55]. The solution to the l_1 - *norm* minimization problem can be obtained via convex quadratic programming. When noise is present in the signal and regularization is required, it becomes a second-order cone programming (SOCP) problem [56]. When the signal is complex-valued and the problem is large scaled (i.e., there are many free parameters to be estimated), as is the case for SAR images, then SOCP methods can be burdensome. The truncated Newton interior point method (TNIPM) [57] has been proposed for efficiently solving the l_1 - *norm* minimization by employing a conjugate gradient method for approximating the search direction for each step of the algorithm. The authors of [57] state that the definition for the l_1 - *norm* of a complex variable $x = x_r + ix_i$ is often taken as follows:

$$\|x\|_1 = \sum_k (|x_r(k)| + |x_i(k)|) \quad (3.9)$$

which tends to result in a solution that is independently sparse in x_r and x_i . However, if both the dictionary and signal are complex, then defining the complex l_1 - norm as in Equation (3.10) tends to result in a solution that is jointly sparse in x_r and x_i [57]:

$$\|x\|_1 = \sum_k \left[(x_r^2(k) + x_i^2(k))^{\frac{1}{2}} \right] \quad (3.10)$$

The definition for the l_1 - norm of a complex variable according to (3.10) equates to minimizing on the sample amplitudes or, in the case of SAR, minimizing on the amplitude image. Indeed, while pixels in a SAR image tend to be uniform and independently distributed in phase, in magnitude they tend to be sparse with a large dynamic range. SAR imagery tends to be highly compressible when magnitude detected, but the uniform random phase of scatterers in the complex image makes the real and imaginary image components relatively incompressible. In general, magnitude detected SAR images tend to be a combination of large RCS point scatters which are sparse, and clutter that has lower RCS and is more prone to speckle and multiplicative noise, with approximate log-normal distributions. This work assumes that targets in a complex SAR image are statistically independent with uncorrelated phase. The l_1 - norm definition of (3.9) allows the most freedom, however the definition in (3.10) enforces sparsity in the magnitude detected image, where we know the SAR image is compressible. In comparison, we can express the

l_2 – norm as:

$$\|x\|_2 = \left[\sum_k (x_r^2(k) + x_i^2(k)) \right]^{\frac{1}{2}} \quad (3.11)$$

The l_2 – norm is the cost function for mean square extrapolation.

More recent work in compressive sensing has centered on re-weighted norm minimization algorithms for sparse decomposition. The statistical relationship between unweighted norm minimization and re-weighted norm minimization is presented with respect to l_1 – norm minimization in [58]. It has also been shown that re-weighted l_2 – norm minimization can also yield sparse solutions [59]. A popular re-weighted norm minimizer is the Focal Underdetermined System Solver (FOCUSS) algorithm. While matching pursuits address the l_0 – norm and basis pursuits address the l_1 – norm, the FOCUSS algorithm addresses the l_p – norm for $0 < p < 2$. The objective for FOCUSS, as discussed in [30], is a weighted minimum norm problem and is summarized as follows:

$$\begin{aligned} \min_{W^{-1}x} \quad & \|W^{-1}x\|_2 \\ \text{subject to} \quad & \|Dx - y\|_2 \leq \sigma_e \end{aligned} \quad (3.12)$$

The solution to the above optimization problem is given by:

$$x = W (DW)^\dagger y \quad (3.13)$$

where A^\dagger denotes the Penrose-Moore pseudoinverse of A .

The authors in [30] assume that the weight array $W = \text{diag}(x^p)$ for integer $p > 0$, where $\text{diag}(x)$ is an operator that creates a square diagonal matrix with vector x

on the main diagonal. They propose an iterative approach, where the weight array is updated with each iteration. This re-weighted minimum norm method tends to yield a sparser solution after each iteration, and the authors show convergence to a unique solution. The FOCUSS algorithm is summarized in Algorithm 3.1.

Algorithm 3.1 Focal Underdetermined System Solver (FOCUSS)

Let y be the vectorized 2D observation samples

Let D be the sparse sampling transform

Select integer $p > 0$

Set tolerance $\delta > 0$

Initialize $x_0 := y$

$e_0 := x_0$

$error := e^H e$

$k := 1$

while $error < \delta$ **do**

$W_k := diag(x_k^p)$

$q_k := (DW_k)^\dagger y$

$x_k := W_k q_k$

$e_k := x_k - x_{k-1}$

$error := e_k^H e_k$

$k := k + 1$

end while

The M-FOCUSS algorithm was introduced in [25] to take advantage of multiple snapshots of the measured data. This method will work if each snapshot has the same sparse sampling matrix, but that is not guaranteed if we are breaking a sparsely sampled SAR aperture into independent subapertures. The authors in [39] extend M-FOCUSS to large-scale problems by introducing a conjugate gradient solver as part of the algorithm, which they term CG-M-FOCUSS. Similar to the multi-look method of speckle reduction described in [49], the phase history data for spotlight mode SAR may be sub-divided into several independent snapshots due to the holographic nature of SAR data. As in optical holography, SAR directly

measures the spatial Fourier content of the scene being imaged, and any sub-portion of the Fourier bandwidth will form a reduced resolution image of the entire scene. The problem with breaking SAR data into blocks to be used with M-FOCUSS is that we are further limiting our window size (bandwidth) for each snapshot, resulting in less inherent resolution. If our phase history is irregularly sampled, then each snapshot may be subject to a different compressive sampling matrix and the iterative formula for solving the M-FOCUSS problem is broken.

When $p = 2$, we may observe that the FOCUSS algorithm is quite similar to the mean square extrapolation algorithm in [46]. Mean square extrapolation is a super-resolution algorithm that exploits the complex sinusoidal model for the data y_{n_1, n_2} , but does not require prior knowledge of the number of scattering centers in the image. It is a Wiener filtering approach, and as such requires some estimate of the signal-to-noise ratio (SNR) between sinusoidal components, due to point scatterers, and the background clutter in the scene that dominates the noise e_{n_1, n_2} in the model. We may rewrite the signal model in matrix-vector notation as $y = Sx + e = SAs + e$, where the observed signal is $y = \text{Vec}(y_{n_1, n_2}) \in \mathbb{C}^{N_1 N_2}$, the noise subspace is $e = \text{Vec}(e_{n_1, n_2}) \in \mathbb{C}^{N_1 N_2}$, the signal to be estimated is $x = \text{Vec}(x_{n_1, n_2}) \in \mathbb{C}^{N'_1 N'_2}$, the space limiting (pruning) matrix is $S \in \mathbb{R}^{N_1 N_2 \times N'_1 N'_2}$, the transform matrix is $A \in \mathbb{C}^{N'_1 N'_2 \times N'_1 N'_2}$, and $s \in \mathbb{C}^{N'_1 N'_2}$ is the sparse signal in the transform domain. For the underdetermined system $N'_1 > N_1$ and $N'_2 > N_2$. Let \hat{x}_{n_1, n_2} be the estimate of the extrapolated signal, then the least mean square 2D estimator is given by [46]:

$$\hat{x} = R_x S^H [S R_x S^H + R_e]^{-1} y \quad (3.14)$$

where

$$R_x = E [xx^H] = E [(As)(As)^H] = AE[ss^H]A^H = APA^H$$

and

$$R_e = E[ee^H] = \sigma_e^2 I$$

Let matrix $A = F_1 \otimes F_2$ correspond to the 2D DFT, and let $W = S(F_1 \otimes F_2) = SA$ be the space limiting 2D DFT, or rather the compressive sampling matrix when applied to SAR phase history. W is an operator that transforms a vector on $\mathbb{C}^{N'_1 N'_2}$ to $\mathbb{C}^{N_1 N_2}$. We may also write:

$$\hat{s} = PW^H T^{-1} y \tag{3.15}$$

where

$$T = W(P + \sigma_e^2 I)W^H \tag{3.16}$$

The matrix T is block-Toeplitz, and \hat{s} is an estimate of the stacked vector representation of the signal in the complex image domain. We must provide an initial estimate for the signal covariance matrix R_x , often taken as R_y . This method also requires an a priori estimate of the average noise power σ_e^2 . In practice, T is likely to be ill-conditioned and a stabilized inverse must be found. It is suggested that this may be accomplished via conjugate gradient algorithms or Kalman filtering [46].

Algorithm 3.2 Mean Square Extrapolation (MSE)

Let $L = N'_1 N'_2$ be the number of image domain pixels

Let y be the vectorized 2D observation samples

Let A be the 2D DFT matrix

Let S be the Fourier space limiting matrix

Let $W = SA$ be the sparse sampling transform

$$P := \frac{1}{L^2} |W^H y|^2$$

$$\Sigma := \sigma_e^2 I_L$$

$$T := W(P + \Sigma)W^H$$

$$\hat{z} := T^\dagger y$$

$$\hat{x} := APW^H \hat{z}$$

$$\hat{e} := y - Sx$$

The mean square extrapolation method [46] is summarized in Algorithm 3.2. It takes a sampled observation signal of dimension $N_1 \times N_2$ and extrapolates it to a larger $N'_1 \times N'_2$ grid. If the space limiting matrix masks observation samples within the $N_1 \times N_2$ region of support, then the least mean square optimum solution will not take into account the missing samples. The author in [46] states that this super-resolution technique is non-parametric because it does not require estimation of frequency pair locations for complex sinusoids, nor their quantity K , but still results in decomposition of the observed measurements into signal and noise subspaces. Notice that only the signal is extrapolated, while the noise is not because it is assumed to be an uncorrelated random stationary white process.

The weighted least-squares approach has been applied to SAR data in [33] by what the authors call the iterative adaptive approach (IAA). Recall the sparse signal representation for the phase history $y = SA s + e$. The authors in [33] propose avoiding the necessity for polar resampling in SAR data by defining the image transform matrix A such that each column (atom) will focus the 2D phase

function for a desired output grid point in the scene, as given below:

$$\begin{aligned}
b_{k_1, k_2}(n) &= \begin{bmatrix} e^{j4\pi \frac{f_1}{c} \cos\theta [x_{k_1} \cos\phi_n + y_{k_2} \cos\phi_n]} \\ e^{j4\pi \frac{f_2}{c} \cos\theta [x_{k_1} \cos\phi_n + y_{k_2} \cos\phi_n]} \\ \vdots \\ e^{j4\pi \frac{f_M}{c} \cos\theta [x_{k_1} \cos\phi_n + y_{k_2} \cos\phi_n]} \end{bmatrix} \in \mathbb{C}^M \\
a_{k_1, k_2} &= \begin{bmatrix} b_{k_1, k_2}^T(1) & b_{k_1, k_2}^T(2) & \dots & b_{k_1, k_2}^T(N) \end{bmatrix} \in \mathbb{C}^{MN} \\
A &= \begin{bmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,k_2} & a_{2,1} & \dots & a_{k_1, k_2} \end{bmatrix} \in \mathbb{C}^{MN \times K_1 K_2} \quad (3.17)
\end{aligned}$$

In the notation above, M is the number of frequency samples per pulse, N is the number of pulses in the synthetic aperture, and a_{k_1, k_2} is a steering vector that focuses the phase history data to a point (x_{k_1}, y_{k_2}) in the scene. The spatial dimension x_{k_1} is in the direction of line of sight from the synthetic aperture center to the scene center, projected into the focus plane. The spatial dimension y_{k_2} is the focus plane direction orthogonal to x_{k_1} .

The IAA algorithm iteratively settles to a unique solution for a sparse set of scattering centers in the scene. The solution is optimal in the weighted least square sense. Each iteration updates the covariance matrix $R_x = APA^H$, where the power $P = \text{diag}(\hat{s})$ is successively sparser. Each iteration of IAA updates the estimated power at each image pixel according to:

$$P_{k_1, k_2} = \left| \frac{a_{k_1, k_2}^H R_x^{-1} y}{a_{k_1, k_2}^H R_x^{-1} a_{k_1, k_2}} \right|^2 \quad (3.18)$$

where a_{k_1, k_2} is given by (3.17).

Algorithm 3.3 Iterative Adaptive Approach (IAA)

Let the dimensions $M = N = L$
Let y be the vectorized 2D observation samples
Let A be the sparse sampling transform in Equation (3.17)
Initialize $P := \frac{1}{L^2} |A^H y|^2$
Initialize $\Sigma := \sigma_e^2 I_L$
for $iter = 1, 2, \dots, MAXITER$ **do**
 $R := APA^H + \Sigma$
 for $k_1 := 1, 2, \dots, K_1$ **do**
 for $k_2 := 1, 2, \dots, K_2$ **do**
 $P_{k_1, k_2} = \left| \frac{a_{k_1, k_2}^H R^{-1} y}{a_{k_1, k_2}^H R^{-1} a_{k_1, k_2}} \right|^2$
 end for
 end for
 for $l := 1, 2, \dots, L$ **do**
 $\sigma_l^2 = \left| \frac{i_l^H R^{-1} y}{i_l^H R^{-1} i_l} \right|^2$
 end for
 $\Sigma = \text{diag}(\sigma_l^2)$
end for

After a number of iterations, the algorithm will converge. The IAA algorithm is summarized in Algorithm 3.3. Based on the image transform matrix A in (3.17), the computational complexity per an IAA iteration is $\mathcal{O}(M^2 N^2 K_1 K_2)$.

3.6. Analytic Comparison

It is possible to make some comparisons between the compressive sensing algorithms discussed above. First, we note the similarity between the IAA estimates and the APES estimates. Then it shall be shown that the IMSE algorithm is an approximation to IAA. The MSE estimate can be described as a linearly re-weighted least-squares solution, while FOCUSS is a nonlinear re-weighted least-squares solution. Application of certain nonlinear weights tend to result in sparser solutions. While MSE weights enforce the l_2 - norm constraint, the FOCUSS weights enforce the l_p - norm, where $0 < p < 2$. Binary weights in OMP and IHT enforce the

l_0 -norm. For efficient solutions to large-scale problems, both MSE and OMP may utilize the method of conjugate gradients to approximate the solutions.

The similarity between APES and IAA estimates differs in the covariance matrices applied in the inverse filtering step. The IAA algorithm utilizes the covariance matrix \hat{R} as defined in Equation (3.6), and APES uses the covariance matrix $\hat{S}(\omega_1, \omega_2)$ defined in Equation (3.7). Note that APES uses a noise covariance estimate $E[\eta\eta^H]$, and that this estimate is different for each frequency pair. While IAA needs to calculate the inverse $\hat{R}^{-1}y$ only once per iteration, APES must calculate $\hat{S}^{-1}(\omega_1, \omega_2)y$ for each frequency pair at every iteration. Otherwise, the frequency amplitude estimates in Equations (3.5) and (3.18) are the same.

It is possible to show that the MSE algorithm is a special case of FOCUSS. We start by expanding the Penrose-Moore pseudo-inverse in the least-squares FOCUSS solution, Equation (3.13), and setting the weight array $W = \text{diag}(x^p)$ with $p = 1$ as in [30]. With the addition of a regularization term, $\sigma_e^2 I$, we see from the formulation below that the FOCUSS and MSE solutions are the same:

$$\begin{aligned}\hat{s} &= W (AW)^{\dagger} y \\ &= W \left[W^H A^H (A W W^H A^H)^{-1} \right] y \\ &= (W W^H) A^H [A (W W^H) A^H]^{-1} y \\ \hat{s} &= P A^H (A P A^H)^{-1} y\end{aligned}$$

This shows that the MSE algorithm solves a weighted minimum norm problem, and when $p = 1$ FOCUSS optimizes against the l_2 -norm. When $p > 1$ even sparser

solutions are encouraged, as more weight is given to frequencies where the PSD is more dense (i.e., there is less penalty for selecting frequencies where P is large). The downside to selecting $1 < p < 2$ is that the weight array $W = \text{diag}(x^p)$ becomes a nonlinear function with an infinite autocorrelation sequence, thus limiting the ability to efficiently perform 2D convolution with the FFT.

3.7. *Summary*

With the previous chapter showing that SAR image formation is a spectral estimation problem, this chapter presents a review of 2D spectral estimators and how they have been or might be applied to SAR. The chapter starts with a summary of classical Fourier base estimators. Classical estimators are susceptible to spectral leakage and have resolution limits according to Nyquist sampling criteria. Nonlinear and adaptive windowing techniques have been developed that attempt to minimize sidelobe response in SAR imagery, but these techniques are limited by Fourier resolution limits and are not well suited for compressive sampling conditions. The Fourier resolution limit is circumvented by modern line spectral estimators. Popular super-resolution estimators, such as MUSIC and ESPRIT, are not well suited for SAR because their 2D formulations produce ambiguous frequency pairs, their computational complexity make them impractical for large scale problems, and the conditions of compressive sampling provide for a poor estimate of the covariance matrix. However, if Nyquist sampling conditions are met, they would be useful for target/clutter segmentation via subspace decomposition. Given that the problem being addressed in this thesis is that of image formation from a sparsely sampled aperture, recently developed sparse reconstruction algorithms from the field

of compressive sensing are presented. Matching pursuit and basis pursuit methods are described, as well as re-weighted minimum norm methods.

CHAPTER 4

PROPOSED SPARSE RECONSTRUCTION ALGORITHMS FOR SAR IMAGE

FORMATION

This chapter proposes two relatively efficient sparse reconstruction algorithms that may be applied to SAR image formation, iterative mean square extrapolation (IMSE) and stagewise gradient pursuit (StGP), both of which are based on compressive sensing principles. The two proposed algorithms demonstrate super-resolution and other desirable properties. The IMSE algorithm is a combination of the MSE and IAA algorithms discussed in Chapter 3. The primary innovation of the IMSE algorithm is the application of fast 2D transforms and alternating projections to efficiently solve a block Toeplitz matrix inversion problem via conjugate gradients. The StGP algorithm is a combination of the StOMP and GP algorithms, also discussed in Chapter 3. The StGP algorithm adopts the stagewise multiple atom selection method of StOMP in order to speed up convergence, while applying conjugate gradient search directions introduced by GP in order to reduce the computational burden of maintaining orthogonality of the solution. The primary innovation of StGP is to apply a sanity filter step that culls the selected atoms at each iteration, such that no two atoms being added to the solution set exceed a given mutual coherence. The chapter concludes with a comparison of algorithms discussed in Chapters 3 and 4.

4.1. *Iterative Mean Square Extrapolation*

The mean square extrapolation (MSE) algorithm [46] is a particularly attractive method for image formation with missing or sparsely sampled data. This algorithm lends itself to the use of fast separable transforms for large scale

multi-dimensional data. While the algorithm was originally developed for super-resolution, or signal extrapolation, the space limiting matrix S may be defined such that it incorporates any missing samples in the phase history domain. The algorithm may be utilized for signal extrapolation and interpolation, so long as the signal may be modeled as a sparse linear combination of sinusoids in a transform domain that adheres to the mutual incoherence property with the sampling matrix and the restricted isometry property (RIP), as described in detail in [43] and [44]. The major technical difficulty with applying the mean square extrapolation algorithm lies in solving the pseudoinverse problem $\hat{z} := T^\dagger y$ due to the ill-conditioned nature of matrix T . Initial estimates for the signal and noise covariance matrices, R_x and R_e respectively, may have dramatic effects on the final solution as well. This work presents a practical method for calculating a stabilized solution for the $\hat{z} := T^\dagger y$ inverse problem, and also iteratively updating the estimates for R_x and R_e .

A powerful numerical tool for solving a stabilized inverse is the conjugate gradient method, as suggested by Jain in [46]. Jain also mentions the use of discrete prolate spheroidal sequences to obtain a solution, which is similar to the minimum norm solution via SVD when the problem is rank deficient. Such an approach is usually intractable for large scale problems however, in which case the conjugate gradient method may be the only practical solution. The theory behind the conjugate gradient method is well documented, see [60] for example, and the method implemented for this work is derived from a template given in [61]. It is worth repeating that, for a convex solution space, an optimal solution is obtained in at

most $L^2 = (M^2N^2)$ conjugate gradient iterations, and without the need to calculate the pseudoinverse T^\dagger directly. It does require the matrix vector multiplication $T\hat{z}$ a number of times, so an efficient solution to this system is desirable. As we shall see, $T\hat{z}$ may be computed efficiently via the fast 2-D FFT when the signal is composed of complex sinusoids. Substituting Equation (3.16) into the expression $\hat{z} := T^\dagger y$ we get

$$\hat{y} := W(P + \sigma_e^2 I)W^H \hat{z} \quad (4.1)$$

Equation (4.1) can be rewritten as:

$$\hat{y} := W[(P + \sigma_e^2 I)\hat{Z}]$$

where $\hat{Z} = W^H \hat{z}$ is the image domain residual following inverse filtering.

The \hat{z} term is an estimate of the residual noise due to clutter after filtering of the estimated complex sinusoid components, and $W = S(F_1 \otimes F_2) = SA$ is the space limiting 2D DFT. The matrix F_1 is the 1D DFT matrix sized to match the first dimension of the 2D observation signal, and F_2 is the 1D DFT matrix sized to match the second dimension. Then $\hat{Z} = W^H \hat{z}$ is an estimate of the clutter image. The term $(P + \sigma_e^2 I)$ is an estimate of the target image, with the addition of a regularization constant proportional to the clutter power. The multiplication $(P + \sigma_e^2 I)\hat{Z}$ is done at the pixel level in the image domain, and the signal \hat{y} is an estimate of the observed clutter-plus-noise phase history. These operations can be done using the 2D FFT transform and simple pixel multiplication, as opposed to full matrix/vector

multiplication on the vectorized phase history data. This is possible because of the assumption that, although the covariance matrix $R_x = WE[ss^H]W^H$ may not be sparse, it is very sparse (e.g., diagonal) in the amplitude image domain. The vector s is simply the vectorized complex image pixels. It is assumed that the expected covariance $P = E[ss^H]$ between complex image pixels is zero due to the uniformly distributed random phase of the radar returns, and is well represented by just the main diagonal elements. It was mentioned previously that the matrix T defined in Equation (3.16) is block-toeplitz, but it is also a block circulant convolution matrix completely defined by the first row of blocks. Thus, the computation of $T\hat{z}$ as just described is in essence calculating the convolution via multiplication in the Fourier domain. It would seem logical that a highly overcomplete (zero padded) DFT matrix would give a more closely spaced grid in the spectral estimate and thus a higher probability of compressing to the true scattering center frequency (i.e., higher spectral resolution in the sparse estimate), but it can be shown that the RIP property of the overcomplete DFT degrades as zero padding is increased [50].

The conjugate gradient method for solving $\hat{z} := T^\dagger y$ requires an initial estimate for the solution \hat{z} . It is typical to initialize this estimate to zero. Each iteration of the conjugate gradient method refines the estimate of \hat{z} . The conjugate gradient iterations may be halted when the residual $(y - \hat{y})$ drops below some tolerance, where y is the original observed signal and \hat{y} is defined as in Equation 4.1, or when stability of the solution is lost. After a stabilized solution of $\hat{z} := T^\dagger y$ has been found, an estimate of the extrapolated/interpolated target signal $\hat{x} = APW^H\hat{z}$ may be computed. This calculation of \hat{x} from \hat{z} is again equivalent to performing

a convolution via multiplication in the Fourier domain. The method is efficient because the transform matrix A is a dictionary of complex sinusoids that can be calculated using the 2D FFT. Other dictionaries with different atoms may not have a fast transform computation that can be exploited.

The mean square extrapolation algorithm requires an initial estimate of the signal covariance matrix $R_x = WE[ss^H]W^H$, or rather an initial estimate of $E[ss^H]$. Jain suggests in [46] that initializing $R_x = \sigma^2 I$ is appropriate if a better estimate does not exist. For this work, the initial estimate is determined as $R_x = APA^H$, where $P = E[ss^H] = \text{diag}(|W^H y|^2)$, which is the upsampled power image given the observed phase history samples y . For the iterative mean square extrapolation algorithm, successive iterations of mean square extrapolation start with the previous solution \hat{s} , so that $R_x = WE[\hat{s}\hat{s}^H]W^H$. As the extrapolated/interpolated signal estimate is refined at each iteration of mean square extrapolation, the estimated target image \hat{s} tends to become more sparse until it eventually settles to an optimal solution of the regularized least-squares problem. The mean square extrapolation algorithm is analogous to the IAA algorithm, except that the IAA algorithm utilizes a pseudo-spectrum estimate for P at each iteration, which adds significant complexity.

The last input necessary for the mean square extrapolation algorithm is an estimate of the noise covariance matrix R_e . For simplified notation, let us define the total number of image domain pixels as $L = N'_1 N'_2$. Assuming the noise to be stationary and white, we may express the noise covariance as $R_e = \sigma_e^2 I_L$, where σ_e^2 is an estimate of the average noise power. As mentioned previously, the noise in the




Fig. 4.1. Distinct point targets versus diffuse scattering clutter in a SAR image.

signal model is dominated by diffuse scattering clutter where there is no dominant point scatterer in the resolution cell. It is difficult to estimate the average noise power in the phase history (Fourier) domain, where the noise and signal are not well separated. However, point targets tend to be sparse in the image domain and the average clutter power, which dominates σ_e , may be estimated from the image. Figure 4.1 shows an actual SAR image, with highlighted regions of point targets, that as well modeled by complex sinusoids, and clutter that does not fit the sinusoid model. Since the point targets are sparse, most of the image is clutter, and it has been determined experimentally that $\sigma_e^2 = (2 * median(|s|))^2$, where s are the complex image amplitudes, usually yields a stable solution. Setting the noise level too low risks instability and noise in the solution, whereas setting it too high risks not reconstructing lower magnitude point targets.

The sparsity of the solution depends on the estimated noise power σ_e^2 . The greater the noise regularization term, the more sparse the solution will be, since greater reconstruction error will be allowed at the expense of reconstructing fewer sinusoidal components. The trivial solution (all zeros) should result if σ_e^2 is greater than $\max_{k_1, k_2}(P)$, as all image pixels will be treated as noise. This is consistent with the assertion that the trivial solution results for the TNIPM l_1 -norm minimization if the Lagrange multiplier λ is greater than $\max_{k_1, k_2}(P)$ [57]. The proposed iterative mean square extrapolation is summarized in Algorithm 4.1.

Algorithm 4.1 Iterative Mean Square Extrapolation (IMSE)

Let $L = N_1'N_2'$ be the number of image domain pixels
Let y be the vectorized 2D observation samples
Let A be the 2D DFT matrix
Let S be the Fourier space limiting matrix
Let $W = SA$ be the sparse sampling transform in Equation (4.1)
Initialize $P := \frac{1}{L^2} |A^H y|^2$
Initialize $\sigma_e^2 := [2 * \text{median}(P)]^2$
Initialize $\Sigma := \sigma_e^2 I_L$
for $k = 1, 2, \dots, \text{MAXITER}$ **do**
 Given: $y, A,$ and S
 $P := \frac{1}{L^2} |W^H \hat{x}|^2$
 $T := W(P + \Sigma)W^H$
 Solve $\hat{z} := T^\dagger y := \text{cg}(P, \Sigma, W, \hat{x}, y)$
 $\hat{x} := APW^H \hat{z}$
 $\hat{e} := y - S\hat{x}$
end for

4.2. Stagewise Gradient Pursuit

Another sparse decomposition algorithm developed as part of this work is based on the StOMP and gradient pursuit algorithms [29, 45]. This algorithm will be referred to as stagewise gradient pursuit (StGP). The StGP algorithm is a greedy matching pursuit method that incorporates the simultaneous selection of multiple

atoms, as with StOMP, but applies an approximation to the orthogonalization constraint in a manner similar to gradient pursuit. The primary benefit of this hybrid approach is enhanced computational efficiency, by removing the orthogonalization filter in StOMP and allowing the selection of multiple atoms at each gradient pursuit iteration. The order of the efficiency gain is dependent of how many atoms are selected at each iteration, which is in turn dependent of the noise regularization level and the sanity filtering threshold, but can be expected to be a couple orders of magnitude more efficient. The trade-off for increased computational efficiency is a less optimal solution than either StOMP or gradient pursuit, with respect to sparsity, but this may be an acceptable trade when considering near real-time imaging.

The dictionary of atoms to select from is the overcomplete DFT matrix, representing a set of complex sinusoids whose frequency spacing is less than the Fourier resolution supported by the observed data. The goal is to select a sparse set of these atoms that linearly combine to approximate the observed samples with minimal error. The matching pursuit method selects the atoms that correlate the best with the observed samples at each iteration. The StGP algorithm computes these correlations coefficients by zero padding the missing data samples and getting the DFT coefficients on an upsampled grid, such that the frequency spacing of the DFT coefficients is at least several times greater than the Fourier resolution obtainable from the sparsely sampled signal. The atoms selected are ones with the largest DFT coefficients that also correspond to a maximum in the expected IPR response, in accordance with a sanity filter. Recall from Chapter 2 that the IPR is the point

spread function in the image domain due to the shape of the 2D aperture in the phase history domain. For the sanity filter, an atom under consideration is rejected if it does not have the largest magnitude DFT coefficient among all other atoms within its IPR response. This method of selecting atoms based on peak detection in the image domain means that super-resolution properties will break down for closely spaced point targets. Well isolated point targets with non-overlapping main lobes should be super-resolved, in that their IPR responses are replaced by ideal impulses, so that sidelobes should no longer occlude neighboring targets. The ability to accurately reconstruct an image from the sparsely sampled phase history is determined by the ability to differentiate between a DFT sample at the center of an IPR versus a DFT sample in an IPR sidelobe. An assumption required for StGP to work is that the pattern of missing samples does not result in an IPR whose maximum amplitude is not at the center of the IPR, a situation the sanity filter attempt to prevent.

The sanity filtering step requires a mask that defines the IPR in the image domain. The IPR may be obtained by applying the system transfer function to an ideal impulse as input. Given the pattern of missing samples in the phase history signal, we may generate an IPR mask in the following manner. First, initialize an image of zeros such that the correlation coefficient associated with each atom in the overcomplete DFT dictionary is zero. Then, set the center sample in the image equal to one and apply the IDFT to obtain the full set of phase history samples associated with the impulse at the center of the image. The phase history samples should have constant amplitude and zero phase. Next force the

phase history samples associated with the missing samples equal to zero and apply the DFT to get the filtered image. The IPR can be considered as the spectral leakage emanating from the ideal impulse input at the center of the image. Stated simply, the IPR is given by the SAR system response to an ideal impulse input, which is in turn determined by the region support in the sampled phase history. With the expected IPR in hand, the condition of a DFT coefficient being the maximum within an IPR centered on it may be checked. If the check fails, that is an indication that a more dominant atom exists within its neighborhood, so it is rejected from consideration until the next iteration. If the condition is met, that is an indication that the atom under consideration is the most dominant in its neighborhood and should be selected for the current iteration. The sanity filtering step encourages selection of well isolated atomic components of the dictionary, with respect to the residual image. This sanity filtering step can be compared to enforcement of a structural sparsity model, as in the coherence-inhibiting structured sparse approximation algorithm presented in [50].

One method of selecting candidate atoms is to choose those whose DFT coefficients have a power that is greater than a certain dB above the average clutter power. As long as there are atoms with correlation coefficients that exceed this threshold, the iterative matching pursuit algorithm continues. The neighboring atoms within the IPR mask centered on each candidate atom are checked to make sure their DFT coefficients are smaller. Candidate atoms that do not pass the sanity filter are rejects. Multiple candidate atoms may be selected during each iteration. Only the DFT coefficients of the selected atoms are retained, while all

others are set to zero, and the IDFT is applied to obtain the full phase history signal associated with the selected atoms. The window associated with the sparsely sampled phase history is applied to force the sparsity pattern on the phase history of the selected atoms. Then, the DFT is applied to get the image IPR responses of the selected atoms after the sampling sparsity constraint has been applied. This filtered image of the selected atoms is then subtracted from the original image, leaving a residual image of the clutter minus the selected atoms. The subsequent iteration of matching pursuit operates on the residual image of the previous iteration. The next set of candidate atoms is selected from the residual image and the process is repeated. The iterations continue until there are no more atoms whose correlation coefficients exceed the threshold above the average clutter power. The atom selection function is highly non-linear and may be compared to energy error reduction methods, such as SVA and Super-SVA, for side lobe control and super-resolution [46, 62, 63]. Whereas SVA applies a non-linear function that suppresses side lobes and retains the sinc IPR main lobe, StGP replaces the entire IPR with an ideal impulse. This means that Super-SVA must apply the pseudoinverse response of a sinc main lobe in the phase history domain, which acts as a whitening filter to sharpen the sinc main lobes in the image following SVA, resulting in super-resolution properties. Since the StGP non-linear function replaces the entire IPR with an ideal impulse, the whitening filter step is unnecessary, but at the expense of not super-resolving closely spaced targets.

The classic OMP algorithm performs an orthogonalization step after selecting new atoms for the sparse decomposition [24, 29]. The coefficients on the selected

basis of atoms are determined by the following least-squares estimate:

$$\hat{s} = [\Phi_k^T \Phi_k]^{-1} \Phi_k^T y \quad (4.2)$$

In (4.2), the $N \times k$ matrix Φ is composed of the k basis atoms selected from the overcomplete dictionary D through the current iteration. The $k \times 1$ vector \hat{s} is the estimate of the coefficients on the k atoms selected, based on the $L \times 1$ vector of observed samples y . As mentioned above, for the case of SAR image reconstruction from a sparsely sampled phase history, the dictionary is assumed to be the complex sinusoids of an oversampled DFT. The orthogonalization step may be rewritten as

$$\hat{s} = [(A^T S^T)(SA)I_{k \in L}]^{-1}(A^T S^T y)$$

where S is the Fourier space limiting transform, A is the 2D DFT matrix, and $I_{k \in L}$ is the $L \times L$ matrix with ones on the diagonal elements corresponding to the indices of the k selected atoms and zeros elsewhere, representing locations of targets in the target image. Then $[(A^T S^T)(SA)I_{k \in L}]$ is the resulting amplitudes of k unit impulses centered on the selected complex sinusoid atoms, after filtering through the sparse sampling matrix. The term $(A^T S^T y)$ can be interpreted as the amplitude of the original image pixels corresponding to the selected k complex sinusoid atoms. Thus, \hat{s} is simply the original DFT coefficients associated with the k selected atoms, weighted by the inverse of their IPR response due to sparse sampling. If it is assumed that the selected atoms adhere to a random uniform phase distribution, then the inverse weighting can be reasonably approximated as a constant gain factor

due to signal loss from reduced region of support in the phase history. This inverse weighting procedure is necessary so that the k sparse components of the observed signal y may be properly removed to yield the residual $\hat{r} = y - \Phi s$. Weighting the DFT coefficients of the residual search direction by a constant that is the inverse of the compression gain loss due to sparse sampling is a decent approximation to the orthogonalization step of OMP if the selected atoms are mutually incoherent. However, if the cumulative coherence between selected atoms is large enough, this approximation can break down. The sanity filtering step is intended to limit the mutual coherence between atoms selected at each iteration. The compression gain loss is given by the ratio of observed phase history samples to the total number of phase history samples resulting from interpolation/extrapolation.

Algorithm 4.2 summarizes the proposed StGP algorithm. Notice that the search direction $d^{(k)}$ is updated after selecting the sanity filtered atoms. The authors in [45] propose a conjugate gradient search direction update in order to approximate the orthogonalization procedure in Equation (4.2). Let $\Gamma^{(k)}$ be the set of all atoms selected through iteration $k = 1, 2, \dots$, then $\hat{r}_{\Gamma^{(k)}}$ is the residual $\hat{r} = y - W\hat{s}$ for atoms in $\Gamma^{(k)}$. Note that \hat{r} is the gradient of the linear equations $y = W\hat{s}$ and its direction would be the search direction in a steepest descent solver. In general, matching pursuit solvers select a search direction $d^{(k)} = \hat{r}_{\Gamma^{(k)}}$ that is a sparse subsets of \hat{r} . To approximate orthogonal matching pursuit, the GP algorithm iteratively updates the search direction $d^{(k)} = \hat{r}_{\Gamma^{(k)}} + b_1 d^{(k-1)}$, where $d^{(k-1)}$ is the search direction from the previous iteration and the initial search direction is the direction of steepest descent, or $d^{(1)} = \hat{r}_{\Gamma^{(1)}}$. The direction update coefficient b_1 is given by [45]:

Algorithm 4.2 Stagewise Gradient Pursuit (StGP)

Let $L = N'_1 N'_2$ be the number of image domain pixels
 Let y be the vectorized 2D observation samples
 Let A be the 2D DFT matrix
 Let S be the Fourier space limiting matrix
 Let $W = SA$ be the sparse sampling matrix
 Initialize $P := \frac{1}{L^2} |A^H y|^2$
 Initialize $\sigma_e^2 := [2 * \text{median}(P)]^2$
 Calculate system IPR $H := \frac{1}{L^2} |A^H \text{diag}(S^T S)|^2$
 Define IPR mask $(i, j) \in \{H_{i,j} > \alpha(\max_{i,j}(H_{i,j}))\}$
 Initialize $\hat{s}^{(0)} := \mathbf{0}$
 Initialize $\Gamma^{(0)} := \emptyset$
 Initialize $d^{(0)} := \mathbf{0}$
 Initialize $b_1 := 0$
for $k = 1, 2, \dots, \text{MAXITER}$ **do**
 Given: $y, A,$ and S
 $\hat{r}^{(k)} := y - W\hat{s}^{(k)}$
 $P^{(k)} := \frac{1}{L^2} |W^H \hat{r}^{(k)}|^2$
 Hard threshold: $(m, n) \in \{P_{m,n}^{(k)} > \sigma_e^2\}$
 Sanity filter $(m, n) \in \{P_{m,n}^{(k)} > \beta(\max_{i,j}(P_{(m-\frac{i}{2}), (n-\frac{j}{2})}^{(k)}))\}$
 Stop if $(m, n) = \emptyset$
 $\Gamma^{(k)} := \Gamma^{(k-1)} \cup (m, n)$
 If $(k > 1)$ then $b_1 := -\frac{\langle c, \hat{r}^{(k)} \rangle}{\|c\|_2^2}$
 Update direction: $d^{(k)} = W^H \hat{r}^{(k)} + b_1 d^{(k-1)}$
 $c := W d^{(k)}$
 $\mu := \frac{\langle \hat{r}^{(k)}, c \rangle}{\|c\|_2^2}$
 $\hat{s}^{(k)} = \hat{s}^{(k-1)} + \mu d^{(k)}$
end for

$$b_1 := -\frac{\langle Wd^{(i-1)}, \hat{r} \rangle}{\|Wd^{(i-1)}\|_2^2} \quad (4.3)$$

4.3. Analytic Comparison

We may deduce similarities between IAA and IMSE. In particular, we may relate the per iteration frequency amplitude estimates given in Equations (3.18) and (3.15). Let us rewrite Equation (3.18) as

$$P_{k_1, k_2} = \left| \frac{1}{a_{k_1, k_2}^H R^{-1} a_{k_1, k_2}} (a_{k_1, k_2}^H R^{-1} y) \right|^2 = \left| \hat{P} (a_{k_1, k_2}^H R^{-1} y) \right|^2$$

or in matrix form as:

$$P = \left| \hat{P} (A^H R^{-1} y) \right|^2 \quad (4.4)$$

The term \hat{P} in Equation (4.4) is a pseudo-spectrum estimate derived as $\hat{P} = \frac{1}{A^H R^{-1} A}$. Note that \hat{P} is the minimum variance distortionless response (MVDR) solution in relation to beamforming. In the MSE algorithm, this pseudo-spectrum term is replaced with P that is derived from the solution \hat{s} from the previous iteration. Using the reduced rank spectrum, as in IAA, should provide for faster convergence to a sparse solution, but requires additional inverse filtering per iteration. As mentioned previously, calculating the stabilized inverse is the most expensive step in the MSE algorithm, which is also true for the other least-squares sparse decomposition methods discussed in this thesis.

Finally, we may compare the StGP and IMSE algorithms. Both solve an approximation problem $y = W\hat{s}$ against sparsity constraints. Both are solved for

large-scale systems by employing the method of conjugate gradients. The initial search direction for both algorithms is that of steepest descent, and conjugate gradient updates applied for subsequent search directions. The difference between StGP and IMSE is in how the sparsity constraint is applied. In StGP, sparsity is enforced by applying a nonlinear sparsifying function to the gradient search directions, so that the search direction itself is sparse. In IMSE, sparsity is enforced via a reweighted least-squares objective, the effect of which is to sparsify the search direction via linear weightings. Whereas the spectral resolution of StGP is restricted by that of the used DFT dictionary, IMSE should have some super-resolution properties associated with iterative deconvolution.

4.4. *Summary*

Two modified sparse decomposition algorithms are proposed, IMSE in Section 4.1 and StGP in Section 4.2, developed for computational efficiency. The StGP method is a matching pursuit type algorithm that has elements taken from stagewise OMP, gradient pursuit and iterative hard thresholding. The IMSE method is a reweighted least-mean-square algorithm that has similarity to the FOCUSS and IAA algorithms, which are close approximations of the APES refined filterbank methods. Both StGP and IMSE are able to handle sparse sampling conditions. Given phase history data that is formatted to a uniform rectangular grid in the spatial frequency domain, these sparse decomposition algorithms may exploit the FFT so that the overcomplete dictionary of atoms does not need to be stored and the transform may be done efficiently.

CHAPTER 5

SIMULATIONS AND RESULTS

The proposed IMSE and StGP algorithms, as described in Chapter 4, are to be assessed quantitatively for their performance against simulated data, and qualitative examples given for application to real-world SAR data. Monte-Carlo simulations are conducted to study the detection performance for point scatterers at various sparsity levels and signal-to-noise ratios (SNR). Another set of simulations investigates the impact of diffraction limiting and coherent speckle when attempting sparse reconstruction of clutter like responses in SAR imagery. Qualitative assessment demonstrates the performance of sparse decomposition algorithms as applied to actual SAR data, with various sparse sampling constraints forced in the phase history domain.

5.1. *Analytic Assumptions*

The purpose of this work is to improve computational efficiency of sparse reconstruction by proposing the StGP and IMSE algorithms for application to large-scale complex SAR imagery. As explained in Chapter 4, these proposed algorithms are not expected to provide more accurate estimates than the sparse reconstruction algorithms discussed in Chapter 3, but they are more efficient. The goal is to realize computational efficiency for practical application to SAR imaging in a production environment, while demonstrating that the proposed algorithms are still able reconstruct a faithful representation of the image given sparse sampling conditions. As such, a direct comparison in performance between the algorithms of Chapters 3 and 4 is not necessary and is not pursued in this work.

Unless otherwise noted for the simulations in this chapter, the maximum iter-

ations (MAXITER) for IMSE is four, and 300 for StGP. In general, four iterations of IMSE involve many more image transformations than 300 iterations of StGP. For clarity of the convergence plots in Section 5.4, the IMSE error is displayed only for the number of transformations computed during the equivalent StGP run. The IMSE algorithm, in its current state of development, does not have an error dependent stopping criteria. The StGP algorithm terminates early if no more peaks are found in the image for selection of atoms, where peaks must both exceed the noise regularization level σ_e^2 and pass the sanity filter parameterized by β . The parameter β is set to 0.1 for all simulations in this work. The noise regularization level σ_e^2 is calculated once, prior to the first iteration, and held constant thereafter. The only exception is for the phantom image reconstructions in Section 5.4, where σ_e^2 is recalculated from the residual image following each iteration, which is necessary for a more complete image reconstruction in those simulations.

It should be noted that phase errors in the signal (phase history) domain are not modeled in Equation (3.15). The simulations presented do not include phase errors, and the provided image examples start with full aperture imagery that has had phase gradient autofocus (PGA) applied prior to enforcing the sparse aperture constraint in the phase history. Uncorrected phase errors would alter the IPR response in the Fourier domain, and, thus, the atoms in the overcomplete DFT dictionary would be mismatched for point target returns and the sparse decomposition performance would be degraded. Two possible solutions exist to handle a sparsely sampled aperture with phase error. One approach is to decouple the autofocus problem from the sparse decomposition problem and develop an autofocus algo-

rithm that takes into account missing phase history samples. The sparse aperture autofocus algorithm would be applied prior to sparse decomposition. The second option would be to combine autofocus and sparse decomposition by incorporating unknown phase error parameters into the model in equation (3.15) and reformulating a solution that also estimates the phase error. Applying a spatially variant sparse aperture autofocus could be quite difficult, but a global quadratic autofocus might be sufficient under most conditions, except when scene terrain height is widely varied. Most SAR phase error can be attributed to non-deterministic errors in ephemeris data over the synthetic aperture, resulting in global phaser errors. Development of a sparse aperture autofocus algorithm is beyond the scope of this thesis.

The application of sparse decomposition in the simulations to follow assume that in-scene content of the imaged area is stationary. That is, the algorithms presented in the previous chapter do not model moving targets within the scene, which would be manifest as wandering spatial frequency components in the collected phase history data. Treating the entire phase history as a single snapshot of data, a moving target will appear defocused in the image due to Doppler shift and/or range migration. Treating sub-apertures of the phase history as independent data snapshots could conceivably allow for tracking frequency migration of distinct scatterers with simple motion, but defocus of vegetation due to wind motion is often very complex because of spatially variant aspects of wind and the many moving parts of vegetation within a resolution cell. Motion-induced defocus can make the image less sparse, but is a practical limitation that must be noted. As

applied here, these sparse decomposition algorithms would treat a non-stationary target as a collection of individual stationary targets distributed in the image.

5.2. *Super-Resolution Properties*

It was mentioned previously that both the StGP and IMSE algorithms are capable of super-resolution, in that they result in an extrapolation or interpolation of the signal domain. The StGP algorithm achieves this by applying a nonlinear function in the image domain, in which a distinct time-windowed IPR is replaced by an ideal impulse with infinite signal support. However, StGP should not be expected to resolve scattering targets that are spaced more closely than the resolution given by the DFT of the signal, as scattering centers are estimated at peaks in the sample spectrum. The IMSE algorithm should be capable of resolving scatterers spaced more closely than resolvable via the DFT, since it performs a sort of iterative deconvolution.

To test the resolution capability of the StGP and IMSE algorithms against these expectations, a test is set up to demonstrate how well these algorithms can correctly identify a group of targets as their spacing increases from below the DFT resolution limit to above the resolution limit. The target pattern consists of four scattering centers lying in the cardinal sidelobes of each other, at the corners of a square in particular. The rectangular region of signal support is limited to one fourth the size of the DFT of the complex image in each dimension. The area of the test image is 32×32 pixels, so the IPR resolution given the rectangular signal support is approximately the -3dB width of a sinc mainlobe, or $(0.6) \left(\frac{2}{8}\right) (32\text{pixels}) \approx 4$ pixels. For normalized target spacings of $\delta = \frac{\{1,2,\dots,6\}}{32}$, Figure 5.1 shows the StGP




Fig. 5.1. Resolvability of closely spaced scatterers.

and IMSE sparse reconstructions. In the spirit of compressive sensing, the signal domain is limited not only in extent of support, but also every other signal sample in both dimensions is zeroed as if it were missing. It should be stated that the targets in the test pattern are complex with random phase.

The first, or leftmost, column of images in Figure 5.1 shows the original test pattern. The second column is the DFT sample spectrum after limiting the signal support. The original target impulses become sinc responses due to the limited signal extent, but are also aliased in both dimensions because of an effective sampling rate reduction due to missing every other sample. This means that the side lobes and aliases of the multiple targets will interfere with each other and cause a reduc-

tion in resolution. The degraded DFT images in the second column demonstrate how the resulting target responses become indistinct, especially as their spacing decreases. The third column of images show the StGP sparse reconstruction, while the fourth column shows the IMSE sparse reconstruction. Each row of plots represents an increased spacing between targets. The target spacing in the fourth row is equal to the theoretical DFT resolution limit given the restricted signal extent. The images in the first three rows have a target spacing smaller than the DFT resolution, while rows five and six have target spacings greater than the DFT resolution. Notice that the IMSE algorithm reconstructs the true target patterns regardless of their spacing, as long as they are centered on grid points in the output image. The StGP algorithm on the other hand is just barely able to resolve the target test pattern at the DFT resolution limit of $\delta = \frac{4}{32}$. At a smaller target spacing, StGP is unable to reconstruct the original target pattern. For $\delta \geq \frac{4}{32}$, StGP is able to resolve the targets. These results are consistent with the aforementioned expectations of super-resolution properties of StGP and IMSE.

5.3. Point Target Detection

An important measure of performance for any radar system, including SAR imaging, is the ability to detect a target in the presence of noise. The performance of a detector can be summarized by its detection and false alarm rates, p_d and p_{fa} respectively. Let (k_1, k_2) be the set of 2D frequency pairs corresponding to true point target locations, and (\hat{k}_1, \hat{k}_2) be the estimated target locations from the sparse decomposition algorithms under review. Then p_d is the probability that the estimates (\hat{k}_1, \hat{k}_2) correspond to the true target locations (k_1, k_2) , and p_{fa} is the

probability that the estimated locations (\hat{k}_1, \hat{k}_2) are in resolution cells that do not contain a target. The probability of detection p_d is calculated as the percent of targets (k_1, k_2) that are found in the estimated locations (\hat{k}_1, \hat{k}_2) . The probability of false alarm p_{fa} is calculated as the percent of estimates (\hat{k}_1, \hat{k}_2) that are not found in the set (k_1, k_2) .

The quantitative metrics are obtained via Monte-Carlo simulation, so that the true target locations are known. Unit amplitude ideal impulse targets with uniform random phase are placed randomly in an $M \times N$ output image grid, which are by definition the true target locations (k_1, k_2) . Random complex Gaussian noise is added and the noisy image is back transformed to the phase history domain via 2D DFT, where a reduced or sparse aperture function is imposed by selectively zeroing samples according to the sparsity mask $S^T(SI)$. The transform S deletes rows from matrix I , while S^T expands (SI) by zero padding the deleted rows of I , where the diagonal elements of I represented the vectorized 2D phase history. The degraded simulated image is obtained by transforming back into the image/spatial domain via 2D IDFT. The noise level is estimated based on the median amplitude in the degraded image, and point target locations and reflectivity estimates (\hat{k}_1, \hat{k}_2) are obtained by applying the StGP and IMSE algorithms, which result in a separation of point targets and noise/clutter subspaces. The location of non-zero elements in the sparse target subspace are taken to be the estimates (\hat{k}_1, \hat{k}_2) , which are evaluated against the true (k_1, k_2) to obtain the aforementioned quantitative metrics. For comparison, detection is performed on the classic sample spectrum (i.e., the degraded simulated image) by labeling any image pixel with power greater than

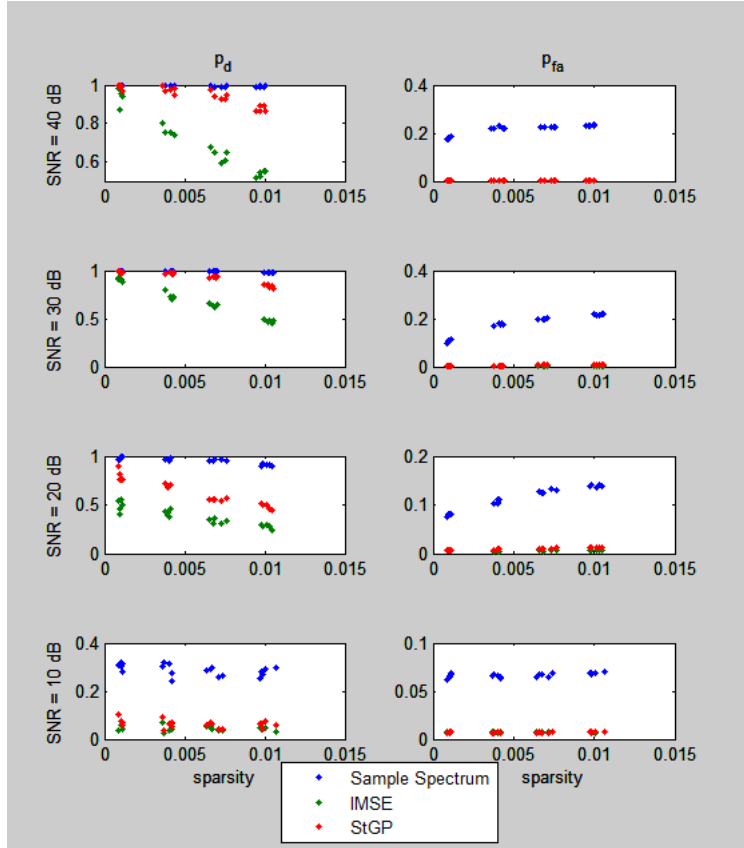


Fig. 5.2. Simulated p_d and p_{fa} for a rectangular aperture.

the estimated noise level σ^2 as a target, and then calculating the same metrics.

Monte-Carlo simulations are performed at various levels of sparsity in the true target locations, and at various sampling (phase history) domain noise levels. The sparsity is defined here as the ratio of the number of image pixels populated by targets to the total number of image pixels. The noise level is given as the signal domain PSNR in dB.

The simulation results in Figure 5.2 show a series of plots of the detection metrics. The left column plots show p_d versus target density, the right column shows the corresponding p_{fa} versus target density, and each row of plots represents

decreasing PSNR. The green data points in the plots represent IMSE results, the red data point represent StGP results, and the blue data points represent target detection via hard threshold detection of the sample spectrum.

The sample spectrum detection finds true target locations, but at the expense of detecting many pixels not at true target locations. This can be attributed to spectral leakage in the sample spectrum estimate, which results in detecting much sidelobe energy as targets when applying a hard threshold. This is characteristic of non-sparse solutions. Applying a similar detection to any of the classical spectral estimators will tend to have many false alarms due to spectral leakage.

The primary advantage of Iterative Mean Square Extrapolation (IMSE) and Stagewise Gradient Pursuit (StGP) over the classical estimators is a reduction in false alarms. The sparse solutions obtained from these algorithms are less likely to identify true target locations, but approach the sample spectrum detector p_d performance when the true target density is very sparse (less than 0.1%) and the peak SNR is sufficiently large (greater than 20 dB). The sparse solutions yield a much smaller false alarm rate (nearly an order magnitude less) because the sparse decomposition estimators suffer less spectral leakage than the classical estimators. In general, StGP performs better than IMSE in terms of p_d , especially at high SNR. Even in the case of a 25% sparse aperture, as in Figure 5.3 where every other sample is missing in both k_1 and k_2 , the relative detection performance stays the same as when a full rectangular aperture is observed.

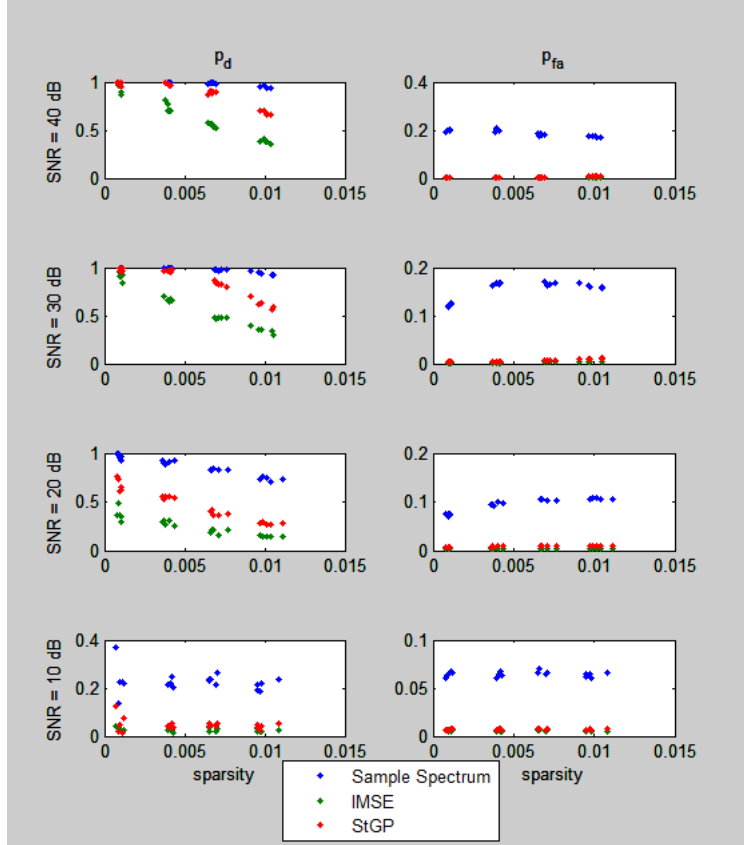


Fig. 5.3. Simulated p_d and p_{fa} for a 25% sparse periodic aperture.

5.4. Sparse Reconstruction of Clutter in Coherent Diffraction Limited Imaging

Much of the literature on compressive sensing is focused on incoherent non-diffraction limited systems. The literature that does address coherent diffraction limited systems, such as SAR, are mostly concerned only with super-resolution of distinct point target returns. Speckle noise reduction in clutter has been attempted through the application of the K-SVD sparse decomposition method [64] with limited success, but there has been no literature to suggest that super-resolution of clutter returns in SAR imagery is feasible. The reality is that clutter returns tend to be severely impacted by constructive and destructive interference from neigh-

boring clutter returns. This is the cause of speckle in a coherent system, when the surface roughness within a resolution cell is on the order of a wavelength. With no dominant scattering centers in a resolution cell, closely spaced clutter returns with approximately the same RCS amplitude but slightly different phases can interfere with each other. Similarly, a coherent diffraction limited system can experience constructive and destructive interference from neighboring resolution cells. The impulse response of a diffraction limited system, like SAR, is not restricted to a single resolution cell. In a non-coherent system, this may result in constructive interference, but a coherent system may also experience destructive interference. The end result is that the problem of reconstructing the clutter from an overcomplete dictionary is very ill-conditioned and the sparsest solution may not be the best, as demonstrated in Figures 5.5 through 5.9.

Figure 5.4 demonstrates sparse reconstruction of the phantom image, commonly used to simulate image reconstruction from projections, as in tomography or MRI, but is being used here to represent clutter and shadows. Figure 5.4(a) is the original real-valued phantom image, while Figure 5.4(b) is degraded with a random 25% of samples missing from its Fourier domain. This simulation is intended to mimic a non-coherent system that is not diffraction limited. Notice the significantly increased noise that obscures the definition of low amplitude areas of the image and a general reduction in contrast, as evidenced by the loss of definition between mid-level amplitude targets within the phantom image. Figures 5.4(c) and 5.4(d) show the sparse reconstructions with StGP and IMSE, respectively. Both StGP and IMSE results are an improvement over the degraded image, restoring def-

initiation between objects of differing amplitudes; however the StGP algorithm results in less residual noise. This may be expected according to [55], as matching pursuit tends toward a sparser solutions than weighted least-squares. The error convergence plot in Figure 5.5 shows that IMSE converges with fewer transformations, but StGP is capable of a much lower reconstruction error under these conditions. The error is plotted against the number of transforms performed, as just a single transform is required to construct the sample spectrum image for comparison, and the fact the each iteration of IMSE requires numerous transform for the conjugate gradients method, while StGP only requires a few transforms per iteration. This simulation also demonstrates that the StGP and IMSE algorithms presented here are capable of reconstructing the special case of a non-coherent image that is not diffraction limited.

Figure 5.6 demonstrates the sparse reconstruction of the phantom image, where Figure 5.6(b) is degraded with a random 25% of samples missing and half the frequency support in either dimension of the Fourier domain (i.e., it has been low-pass filtered). Real-valued amplitudes are retained in the image domain. The reduced frequency support is typical of SAR imagery and yields an impulse response representative of a diffraction-limited imaging system. The degraded image still suffers increased noise and reduced contrast due to the randomly missing samples. Banding artifacts are also introduced by spectral leakage of the IPR. Figures 5.6(c) and 5.6(d) show the StGP and IMSE reconstructions of the degraded, real-valued and diffraction-limited image simulation. Notice that the sparser StGP reconstruction bears less resemblance to the original than the IMSE reconstruction, even though




Fig. 5.4. Phantom image restoration: real-valued and non-diffraction limited. (a) Original. (b) Degraded 75% random aperture. (c) StGP restoration. (d) IMSE restoration.

the StGP signal reconstruction in the Fourier domain has less error (see Figure 5.7). This simulation clearly demonstrates that the sparsest decomposition is not necessarily preferable for closely spaced scatterers in a diffraction-limited system. The IMSE solution is more interpretable when compared to the original phantom image, which appears mostly due to reduced noise from the random missing samples, yet the banding artifacts remain.

The simulation in Figure 5.8 is similar to that in Figure 5.6, except that the phantom image is given complex amplitudes by applying uniform random phase to each pixel. The image in Figure 5.8(b) is degraded with a random 25% of samples missing and half the frequency support in either dimension of the Fourier domain.




Fig. 5.5. Phantom image restoration error convergence. Real-valued, non-diffraction limited, degraded 75% random aperture.

The reduced frequency support and complex-valued reflectivity map is typical of SAR imagery, being both a diffraction-limited and coherent sensor system. The degraded image still suffers increased noise and reduced contrast due to the randomly missing samples, and the same banding artifacts due to spectral leakage of the IPR. In addition, there is speckle from constructive and destructive interference among IPR responses centered at different locations in the image. Figures 5.6(c) and 5.6(d) show the StGP and IMSE reconstructions of the degraded, complex-valued and diffraction-limited image simulation. Notice that the IMSE reconstruction is sparser than it was for the prior simulation without speckle degradation. Both the StGP and IMSE solutions are unable to deal with speckle noise. While the




Fig. 5.6. Phantom image restoration: real-valued and diffraction-limited. (a) Original. (b) Degraded 75% random aperture. (c) StGP restoration. (d) IMSE restoration.

global noise induced by the randomly missing samples is improved, both the loss of contrast and resolution among smaller objects and those with mid-level amplitudes are not recovered. The reduced global noise does appear to help somewhat with the definition of low amplitude areas, but not significantly. Figure 5.9 shows that the StGP solution has less error than the IMSE solution. In this case, StGP yields a result that is comparable, if not preferable to IMSE, if for no other reason than that the performance of IMSE degraded rapidly with the introduction of speckle noise.




Fig. 5.7. Phantom image restoration error convergence. Real-valued, diffraction-limited, degraded 75% random aperture.

5.5. Sparse Decomposition of Actual Spotlight SAR Imagery

An actual SAR image is not usually represented by a random field of point scatterers. Point target returns are likely to be clustered and are often related to man-made structures with dihedral and trihedral cross section components, but are typically sparse relative to the area being mapped, especially for high resolution SAR imagery. With this in mind, it is important to demonstrate the IMSE and StGP sparse decomposition algorithms as applied to actual SAR data.

The data used in this work was collected by Sandia National Laboratories' Otter II platform and was provided by Charles V. Jackowatz as part of a presentation package given to students attending his short course on SAR [65]. The data




Fig. 5.8. Phantom image restoration: complex-valued and diffraction-limited. (a) Original. (b) Degraded 75% random aperture. (c) StGP restoration. (d) IMSE restoration.

is spotlight-mode collection of three different scenes: a parking lot, an airport, and the Capitol building in Washington, D.C. Each data set has been motion compensated, range deskewed, and reformatted from a polar to a rectangular grid. The phase gradient autofocus algorithm [49] is applied to yield a well focused complex image with full Fourier support in the synthetic aperture. The well focused SAR images is the starting point for each of the simulations to follow. Sparse sampling of the apertures is enforced by decompressing the complex image via 2D FFT, and then zeroing selected samples in the Fourier (phase history) domain. The IMSE and StGP algorithms are applied to the sparse apertures in an attempt to reconstruct the original complex image.




Fig. 5.9. Phantom image restoration error convergence. Complex-valued, diffraction-limited, degraded 75% random aperture.

The three images of the three different scenes is typical of real-world SAR imagery, with varying mixtures of point-like impulse responses and diffusely scattering clutter. The parking lot scene has many point like returns from numerous vehicles on top of uniformly scattering clutter from pavement in the background. The airport scene has a few point like returns associated with several airplanes and support equipment, with low backscatter clutter from the tarmac and higher backscatter clutter from unpaved land around the tarmac. The Capitol scene has point like returns from building features, varying levels of clutter from trees and vegetation, and low return areas from specular reflectors such as water and paved walkways, as seen in Figure 4.1.

Both the IMSE and StGP algorithms are applied below to the three SAR images, where a 25% sparse periodic sampling is enforced in the phase history domain. The sparse sampling matrix includes 2x zero padding around the original phase history for super-resolution in the sparse reconstruction, while every other sample in range frequency k_1 and azimuth frequency k_2 is zeroed to represent missing samples. The resulting periodic sparse aperture pattern is similar to that seen in Figure 2.13(b). Imagery formed with classical spectral estimators are characterized by 2D sinc impulse responses, due to the 2x zero padding in the Fourier domain, and aliasing with two-fold symmetry because of periodic missing Fourier samples in both dimensions. For example, a point target in the upper left corner of the image will have periodic replicas in both dimensions. Well known Fourier transform properties tell us that the spacing between target aliases is inversely proportional to the period of the missing samples in the phase history domain. The following demonstrations show that target aliases imposed by periodic sparsely sampled aperture can greatly confuse interpretation of the SAR image.

Figures 5.10, 5.13 and 5.16 demonstrate the application of the IMSE and StGP sparse reconstruction algorithms to three scenes. The first scene is a parking lot with vehicles surrounding a building. The second scene is of the area surrounding the US Capitol building. The third scene is an airport. The subimage in the upper left is the original SAR image formed via the sample spectrum (DFT) with 2x zero padding enforced in the phase history domain. The subimage in the upper right is the DFT image after enforcing the 25% sparse periodic sampling pattern. The subimage in the lower left is the reconstruction from the sparse samples using the

StGP algorithm, while the subimage in the lower right is the reconstruction using the IMSE algorithm.

The StGP and IMSE algorithms result in a decomposition targets and clutter in the image. Point scattering targets are modeled as complex sinusoidal components in the phase history, which is the reconstructed signal. The clutter cannot be modeled as complex sinusoids and factored into the error component, which is not reconstructed. This decomposition of target and clutter components of the image can be useful as inputs to target identification and classification algorithms, or speckle reduction algorithms for the clutter. The reconstructed StGP and IMSE images in Figures 5.10, 5.13 and 5.16 represent the coherent sum of the target and clutter components because, even though the clutter is not reconstructed, clutter can still add image context for the interpretation of reconstructed point targets.

Figures 5.11, 5.14 and 5.17 demonstrate the target and clutter decomposition for the three image examples. Recall that the target image is the sparse signal component of the model in equation (3.2), while the clutter image is the residual noise in the model that cannot be interpolated nor extrapolated. The subimage in the upper left is the target image generated from the StGP signal space solution. The subimage in the upper right is the StGP residual clutter image. The subimage in the lower left is the target image generated from the IMSE signal space solution, while the subimage in the lower right is the associated residual clutter image.

In Figures 5.10, 5.13 and 5.16 we see that the sample spectrum image formed from the sparse aperture is degraded in the form of periodic replicas of targets in both dimensions. All return energy from the scene, both point targets and diffuse

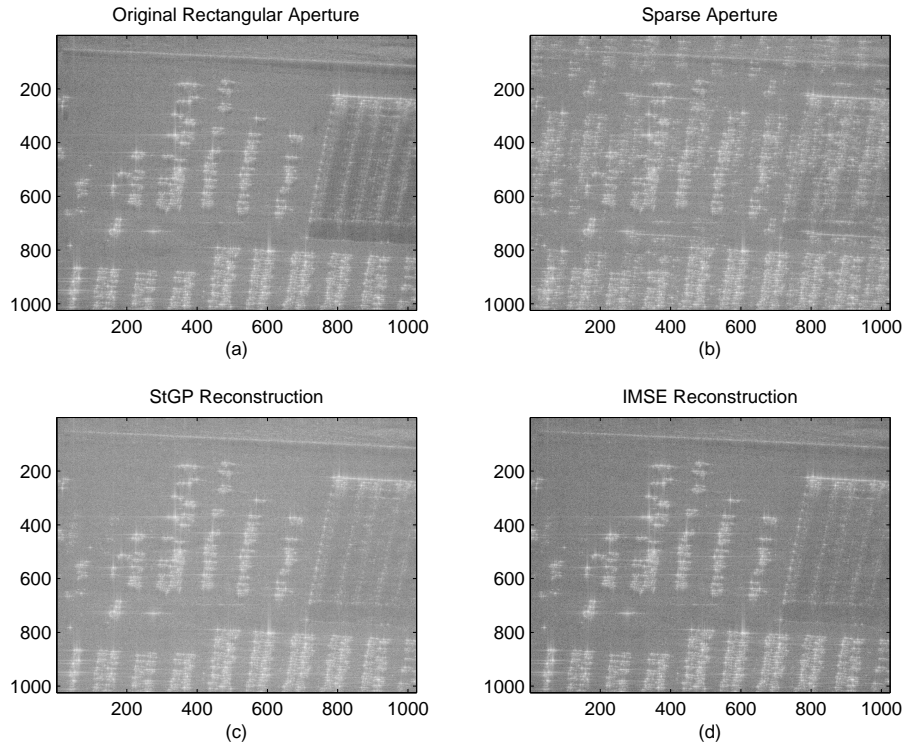


Fig. 5.10. Parking lot scene with a 25% sparse periodic aperture.

clutter, are replicated due to the periodic sparse aperture. The StGP and IMSE reconstructions at the bottom of the figures are effectively able to distinguish between true point-like targets and their replicas, as compared to the original full aperture images in the upper left of the figures. The parking lot scene is particularly confused, as it is difficult to identify the buildings and distinguish between empty or occupied space in the surrounding parking lot. The reconstructed images restore the ability to distinguish these targets. Notice that the noise levels in the shadow areas of the reconstructed images is elevated in comparison to the full aperture image. This is a result of the inability of the sparse decomposition algorithms to extrapolate or interpolate clutter return energy, so the replicated clutter

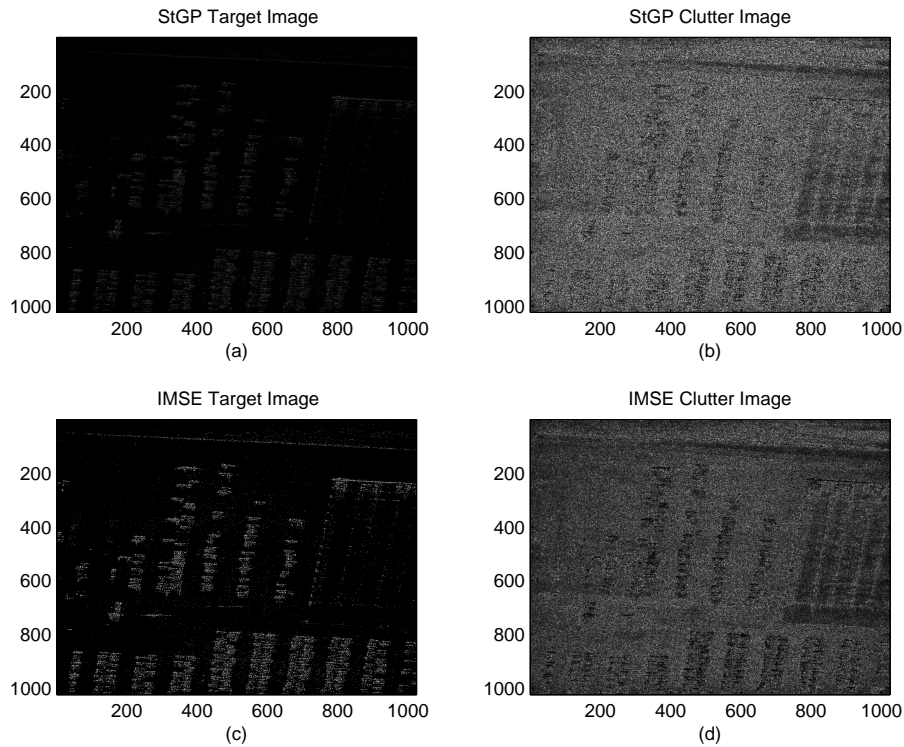


Fig. 5.11. Parking lot scene target/clutter segmentation with a 25% sparse periodic aperture.

tends to raise the noise floor in the reconstructed image. The unrecoverable loss in interpretability due to the elevated noise floor is evident in the airport image. The shadows of the two airplanes in the original image clearly show the shape of their wings, but these shadows are lost in the reconstructed images. It is not expected that compressive sensing algorithms would be able to resolve clutter, as clutter tends not to be sparse in the spectral domain. Furthermore, diffuse clutter is more severely impacted by speckle noise than point target returns, and we can see from Figure 5.2 that the performance of the sparse decomposition algorithm degrades quickly as the noise level approaches the signal level.

A close inspection of Figures 5.11, 5.14 and 5.17 reveals that the IMSE al-

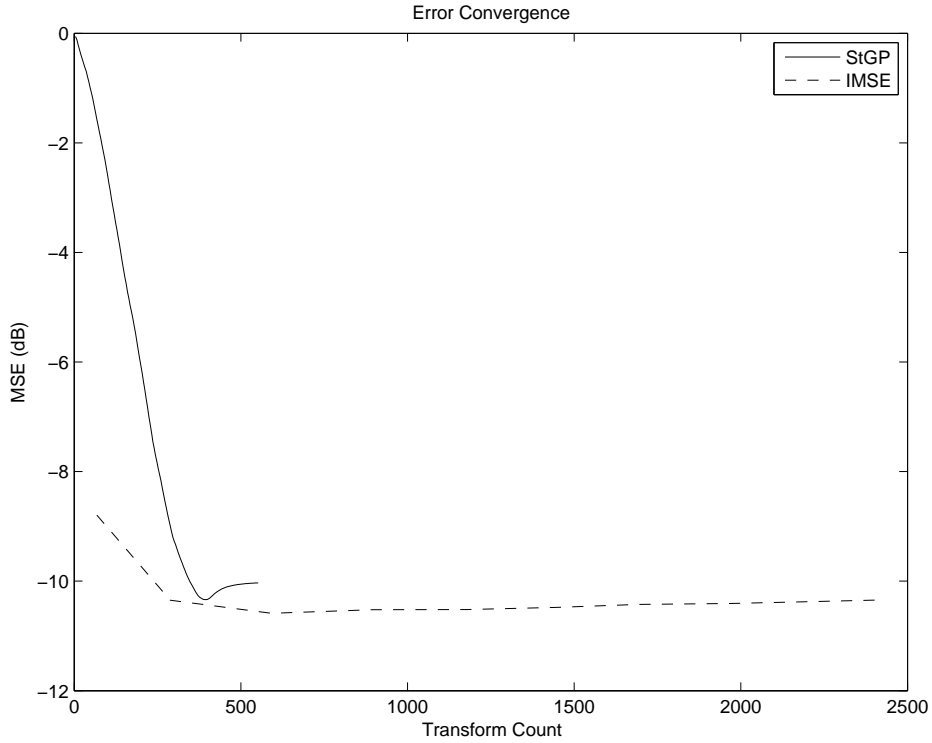


Fig. 5.12. Parking lot scene error convergence with a 25% sparse periodic aperture.

gorithm results in distinct nulls in the place of strong point targets in the clutter subspace image. This behavior is not observed in the StGP clutter image. The presence of nulls in the IMSE clutter image can be related to beamforming techniques. Indeed, the IMSE solution in equation (3.15) is much like the minimum mean square error (MMSE) solution to beamforming, in which beam pattern nulls are steered toward interfering sources, while the main beam is steered toward the desired source. Similarly, the IMSE algorithm steers weighted nulls in the spatial dimension during the process of estimating the signal (target) subspace signal. In contrast, the StGP clutter image tends to leave a residual resembling the surrounding clutter in locations where point target returns are removed. One can imagine

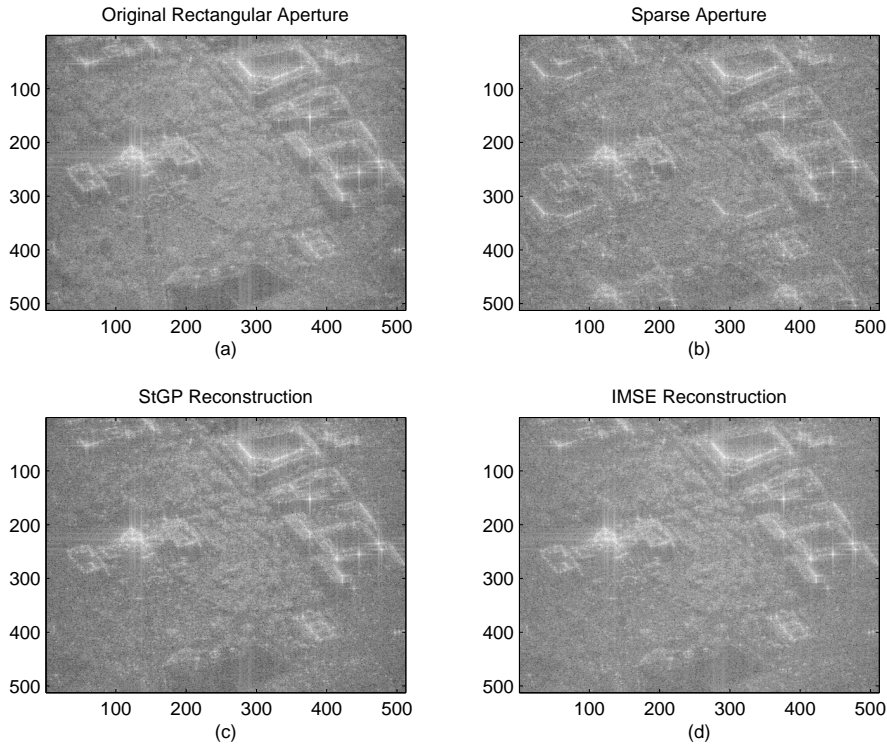


Fig. 5.13. Capitol scene with a 25% sparse periodic aperture.

some post-processing being performed on the clutter image prior to reconstructing the combined image. The preferred behavior of nulls in the clutter image would depend on the goal of the post-processing. For example, nulls may be undesirable for a smoothing speckle reduction filter. For the purpose of image interpretability, the desired clutter image for reconstruction would have the least noise in the radar shadows. Unfortunately, the residual noise in the shadows is most likely due to spectral leakage from surrounding clutter, which neither IMSE or StGP can address. In the three image examples, the IMSE and StGP clutter images both demonstrate similar noise levels in the shadows, given the same noise power estimate for regularization. The IMSE algorithm might demonstrate better noise performance in

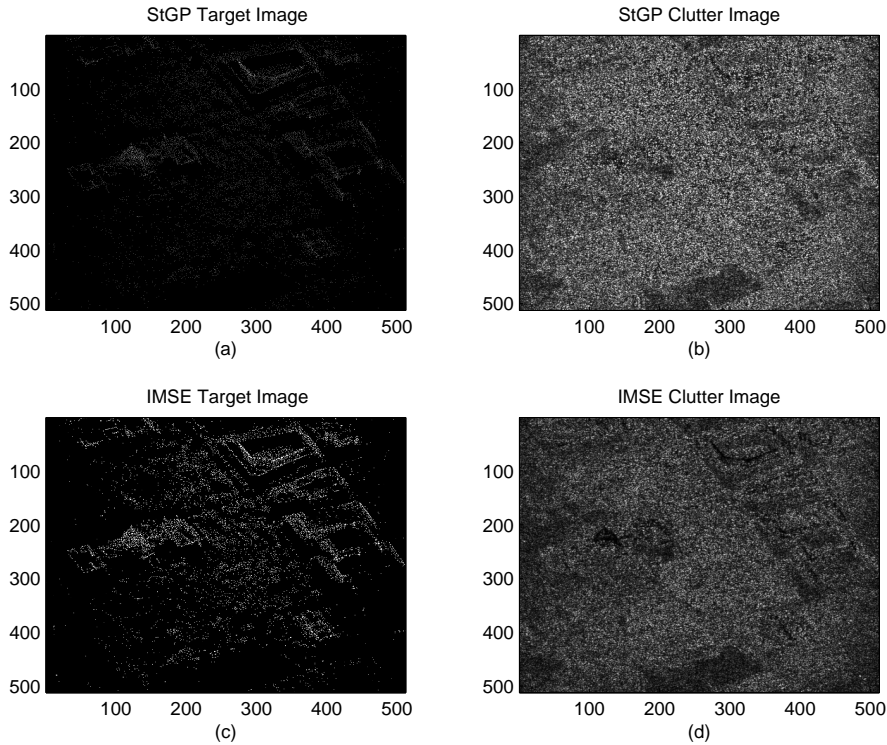


Fig. 5.14. Capitol scene target/clutter segmentation with a 25% sparse periodic aperture.

the clutter image in cases where the clutter to speckle noise ratio is large and the point target density is less sparse, as suggested in the simulation results of Figures 5.2 and 5.3, by lowering the noise power estimate σ_e^2 .

5.6. Summary

The simulations presented in this chapter demonstrate that, when the image content is suitably sparse and the SNR is not too high, the StGP and IMSE sparse decomposition algorithms are able to identify the location of 2D line spectral components accurately with high probability, even when the phase history is sampled suboptimally according to Nyquist criteria. Compared to classical estimation, sparse reconstruction minimizes spectral leakage due to missing samples




Fig. 5.15. Capitol scene error convergence with a 25% sparse periodic aperture.

in the phase history domain, making the resulting image less susceptible to false targets or aliases that would otherwise degrade reliable exploitation.

It is shown that the StGP and IMSE algorithms are able to reconstruct an accurate representation of the phantom image in Figure 5.4, even when samples are missing from its discrete Fourier transform. However, the performance of these algorithms is degraded when the bandwidth of the phantom image is limited by a rectangular frequency aperture to simulate a sinc impulse response in the image domain, which is characteristic of a diffraction-limited system such as high resolution SAR. The algorithms are still capable of reducing noise that is induced by randomly missing frequency samples, but they are unable to deal with point spread

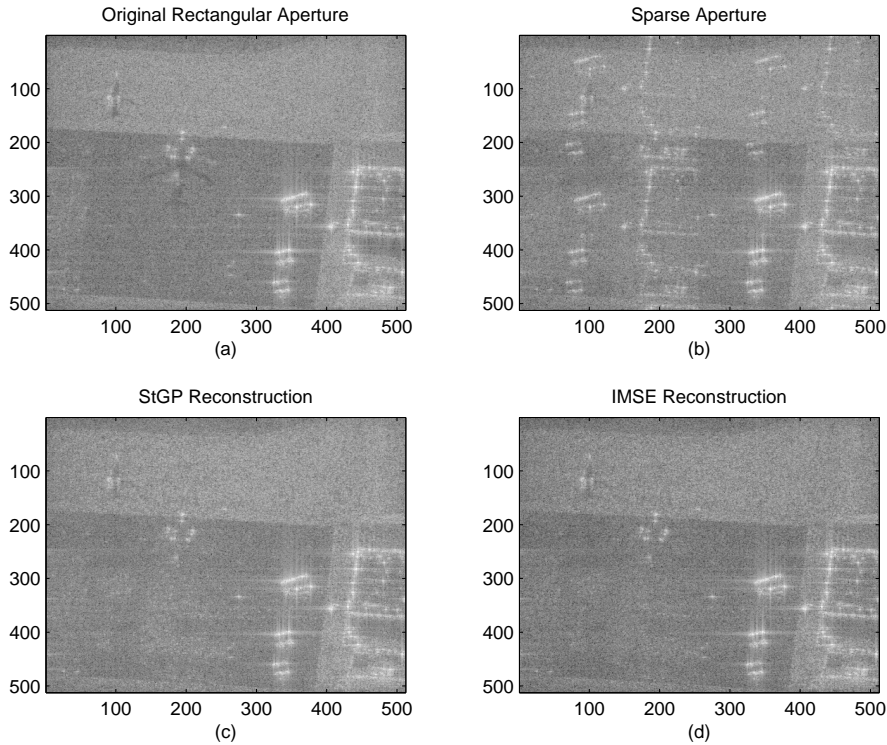


Fig. 5.16. Airport scene with a 25% sparse periodic aperture.

function interaction between neighboring targets of similar magnitude. This limitation is caused by the ill-conditioned nature imposed by coherence between atoms in an overcomplete or zero-padded DFT dictionary. When random phase is added to the image pixels in combination with a sinc IPR, speckle like noise is added to the image due to constructive and destructive interference between neighboring targets. This leads to further ill-conditioning, and the StGP and IMSE algorithms are unable to reconstruct information lost due to coherent interference between frequency components. In fact, the simulated phantom image reconstructions in Figure 5.8 suggest that the sparsest solution to the underdetermined system may not be the most preferable given the presence of speckle noise in the background

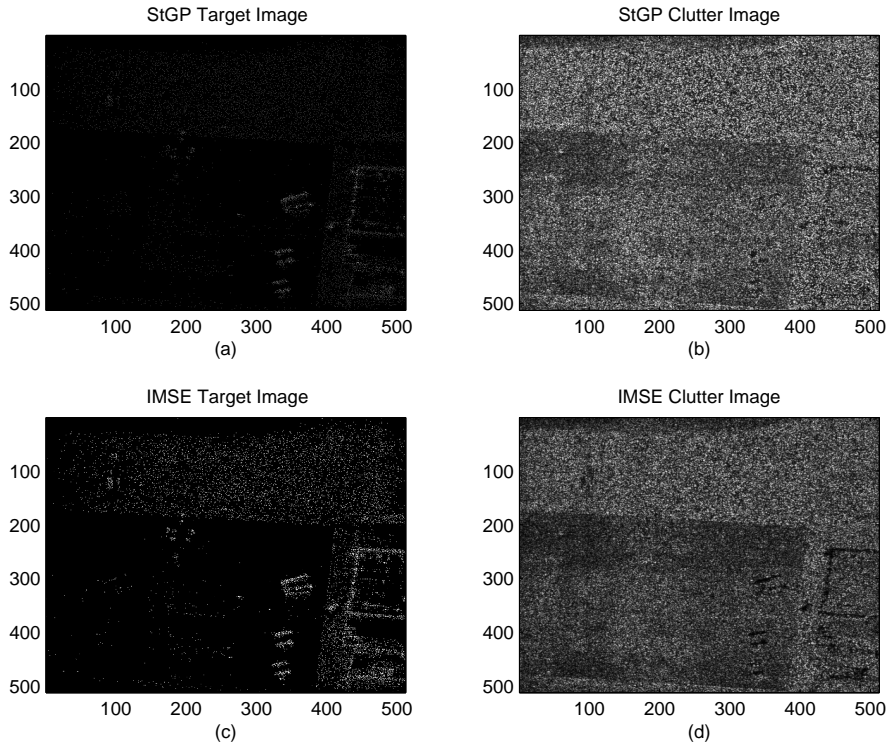


Fig. 5.17. Airport scene target/clutter segmentation with a 25% sparse periodic aperture.

image clutter.

Setting the noise regularization level above the clutter power level, as estimated from image domain statistics, allows for reconstruction of distinct point scattering targets via the StGP and IMSE sparse decomposition algorithms, while leaving the background clutter untouched. This results in a subspace decomposition of point targets from background clutter. Target responses are restored from the sparse phase history aperture, whereas the clutter responses are not restored and still suffer spectral leakage impacts of irregular aperture windowing. Examples are given of target/clutter segmentation on actual high resolution SAR images, which could be useful for target recognition applications or further filtering of image clutter,




Fig. 5.18. Airport scene error convergence with a 25% sparse periodic aperture.

such as speckle filtering.

CHAPTER 6

CONCLUSION

This chapter summarizes the contributions of this thesis to the field of sparse decomposition in relation to spotlight-mode SAR image formation. Also discussed in this chapter are ideas for continued work.

6.1. *Contributions*

The contributions of this work include: 1) The exploration of the relationship between spotlight-mode SAR image formation and spectral estimation; 2) the resulting impact of missing observation samples on SAR imagery for both uniformly downsampled data and in a non-uniform compressive sampling framework; 3) the implementation of two proposed sparse decomposition estimation algorithms, StGP and IMSE, that could be applied in near real-time with modest compute hardware; 4) identifying the limitation of sparsity constraints when reconstructing image clutter in the presence of coherent speckle; 5) demonstrating the application of sparse reconstruction to real-world SAR data, and the resulting target/clutter image segmentation.

6.2. *Conclusions*

It has been shown in this work that it is feasible to apply sparse decomposition algorithms to sparse aperture spotlight-mode SAR data and reconstruct an image that would otherwise be severely degraded if using classical spectral estimation. The proposed StGP and IMSE algorithms have similar effectiveness, though simulations suggest that StGP is more robust in the presence of AWGN, and as target density in the image increases. However, IMSE is capable of super-resolving closely spaced scattering centers, whereas StGP is unable to do the same. These

algorithms can be adept at mitigating spectral leakage and aliasing of point target impulse responses in SAR images but, as with any estimator of complex sinusoids, reconstructing the clutter subspace is beyond their capability, as clutter in SAR imagery is not well modeled by sinusoids. The lack of clutter reconstruction implies that the definition of radar shadows will be difficult to recover. There is no theory to suggest that clutter resolution can be enhanced; so these algorithms are most useful for point target detection and identification, and less useful for exploitation of the clutter map. However, these sparse decomposition algorithms do work well for target/clutter segmentation and, when applied to fully sampled phase history, could be useful for target recognition and other applications that rely on the clutter map. For example, work on speckle noise filtering using K-SVD has been attempted [64]. The application of K-SVD based speckle reduction looks promising, but the results show a severe attenuation of clutter surrounding point targets with strong reflectivity. Applying the K-SVD filter to the segmented clutter image might mitigate such artifacts.

6.3. *Future Work*

There are opportunities to continue work in the area of sparse decomposition and compressive sensing in SAR imaging. Algorithmic robustness to non-AWGN noise, such as speckle, has not been studied in detail. The stopping criteria and regularization levels presented in this work have been selected experimentally to provide acceptable results, but have not been optimized. The StGP and IMSE algorithms can also be extended to address missing phase history data prior to

polar resampling, though this may preclude application of fast transforms for a polar sampling grid.

Nondeterministic azimuthal (slow-time) phase errors are common in SAR imaging. As mentioned in Chapter 5, phase errors distort the IPR in the image domain, such that complex sinusoid atoms would yield a less sparse representation of the image. Future work could look at adapting or implementing an autofocus algorithm that takes into account missing phase history samples, then apply it prior to sparse decomposition. Another option is to integrate phase error correction vectors into the sparse decomposition problem. A logical starting point would be to tackle global quadratic phase errors, given a sparsely sampled aperture.

REFERENCES

- [1] S.R. DeGraaf, "SAR imaging via modern 2-D spectral estimation methods," *IEEE Transactions on Image Processing*, vol. 7, no. 5, pp. 729–761, 1998.
- [2] H.C. Stankwitz and M.R. Kosek, "Super-resolution for SAR/ISAR RCS measurement using spatially variant apodization," in *Proc. AMTA Symp., Williamsburg, VA*, 1995, pp. 251–256.
- [3] J.W. Odendaal, E. Barnard, and C.W.I. Pistorius, "Two-dimensional super-resolution radar imaging using the MUSIC algorithm," *IEEE Transactions on Antennas and Propagation*, vol. 42, no. 10, pp. 1386–1391, 1994.
- [4] A. Quinquis, E. Radoi, and F.C. Totir, "Some radar imagery results using superresolution techniques," *IEEE Transactions on Antennas and Propagation*, vol. 52, no. 5, pp. 1230–1244, 2004.
- [5] H.C. Stankwitz and M.R. Kosek, "Sparse aperture fill for SAR using super-SVA," in *Proceedings of the IEEE National Radar Conference*, 1996, pp. 70–75.
- [6] Y. Wang, J. Li, and P. Stoica, "Two-dimensional nonparametric spectral analysis in the missing data case," in *Proc. 30th IEEE Int. Conf. Acoustics, Speech, Signal Processing (ICASSP)*, 2005, pp. 397–400.
- [7] S.R. DeGraaf, "Sidelobe reduction via adaptive FIR filtering in SAR imagery," *IEEE Transactions on Image Processing*, vol. 3, no. 3, pp. 292–301, 1994.
- [8] H.C. Stankwitz, R.J. Dallaire, and J.R. Fienup, "Nonlinear apodization for sidelobe control in SAR imagery," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 31, no. 1, pp. 267–279, 1995.
- [9] H.C. Stankwitz and S.P. Taylor, "Advances in non-linear apodization," *IEEE Aerospace and Electronic Systems Magazine*, vol. 21, no. 1, pp. 3–8, 2006.
- [10] W. Phillips, S. DeGraaf, and R. Chellappa, "Enhanced segmentation of SAR images using non-Fourier imaging," in *Proc. of the IEEE International Conference on Image Processing (ICIP 98)*, 1998, vol. 1, pp. 583 – 586.
- [11] K.T. Kim and H.T. Kim, "Two-dimensional scattering center extraction based on multiple elastic modules network," *IEEE Transactions on Antennas and Propagation*, vol. 51, no. 4, pp. 848–861, 2003.
- [12] P. Stoica and R. Moses, *Spectral Analysis of Signals*, Pearson Prentice Hall, 2005.

- [13] M. Haardt and J.A. Nossék, “Unitary ESPRIT: How to obtain increased estimation accuracy with a reduced computational burden,” *IEEE Transactions on Signal Processing*, vol. 43, no. 5, pp. 1232–1242, 1995.
- [14] M. Haardt and F. Romer, “Enhancements of unitary ESPRIT for non-circular sources,” in *Proc. of the IEEE Int. Conf. Acoustics, Speech, Signal Processing (ICASSP)*, 2004, vol. 2, pp. 101–104.
- [15] M. Haardt, M.D. Zoltowski, C.P. Mathews, and J.A. Nossék, “2D Unitary ESPRIT for efficient 2D parameter estimation,” in *Proc. of the IEEE International Conference on Acoustics Speech and Signal Processing*, 1995, vol. 3, pp. 2096–2096.
- [16] F. Chen, C.C. Fung, C. Kok, and S. Kwong, “Estimation of two-dimensional frequencies using modified matrix pencil method,” *IEEE Transactions on Signal Processing*, vol. 55, no. 2, pp. 718–7224, 2007.
- [17] Y. Hua, “Estimating two-dimensional frequencies by matrix enhancement and matrix pencil,” *IEEE Transactions on Signal Processing*, vol. 40, no. 9, pp. 2267–2280, 1992.
- [18] Y. Hua and T.K. Sarkar, “Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise,” *IEEE Transactions on Acoustics Speech and Signal Processing*, vol. 38, no. 5, pp. 814–824, 1990.
- [19] S. Ozsoy and A.A. Ergin, “Pencil back-projection method for SAR imaging,” *IEEE Transactions on Image Processing*, vol. 18, no. 3, pp. 573–581, 2009.
- [20] T.K. Sarkar and O. Pereira, “Using the matrix pencil method to estimate the parameters of a sum of complex exponentials,” *IEEE Antennas and Propagation Magazine*, vol. 37, no. 1, pp. 48–55, 1995.
- [21] S.R. DeGraaf, “Parametric estimation of complex 2-D sinusoids,” in *Proc. of the Fourth Annual ASSP Workshop on Spectrum Estimation and Modeling*, 1988, pp. 391–396.
- [22] Y. Wang, J. Li, and P. Stoica, *Spectral analysis of signals: the missing data case*, vol. 1, Morgan & Claypool Publishers, 2006.
- [23] L. Marple, “Spectral line analysis by Pisarenko and Prony methods,” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1979, vol. 4,, pp. 159 – 161.

- [24] J.A. Tropp, “Greed is good: Algorithmic results for sparse approximation,” *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [25] S.F. Cotter, B.D. Rao, K. Engan, and K. Kreutz-Delgado, “Sparse solutions to linear inverse problems with multiple measurement vectors,” *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2477–2488, 2005.
- [26] J. Casar-Corredera, J.M. Alcázar-Fernández, and L.A. Hernández-Gómez, “On 2-D Prony methods,” *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 10, pp. 796 – 799, 1985.
- [27] R. Carriere and R.L. Moses, “High resolution radar target modeling using a modified Prony estimator,” *IEEE Transactions on Antennas and Propagation*, vol. 40, no. 1, pp. 13–18, 1992.
- [28] D.L. Donoho, M. Elad, and V.N. Temlyakov, “Stable recovery of sparse over-complete representations in the presence of noise,” *IEEE Transactions on Information Theory*, vol. 52, no. 1, pp. 6–18, 2006.
- [29] D.L. Donoho, Y. Tsaig, I. Drori, and J.L. Starck, “Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit,” *Submitted for publication*, 2006.
- [30] I.F. Gorodnitsky and B.D. Rao, “Sparse signal reconstruction from limited data using FOCUSS: A re-weighted norm minimization algorithm,” *IEEE Trans. Signal Processing*, vol. 45, no. 3, pp. 600–616, 1997.
- [31] M.D. Plumbley, “On Polar Polytopes and the Recovery of Sparse Representations,” *IEEE Transactions on Information Theory*, vol. 53, pp. 3188 – 3195, 2005.
- [32] G. Rath and C. Guillemot, “Sparse approximation with an orthogonal complementary matching pursuit algorithm,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2009, pp. 3325–3328.
- [33] M. Xue, E. Santiago, M. Sedehi, X. Tan, and J. Li, “SAR imaging via iterative adaptive approach and sparse Bayesian learning,” in *Proceedings of SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery XVI*, 2009, vol. 7337.

- [34] C.R. Berger, J. Areta, K. Pattipati, and P. Willett, “Compressed sensing—a look beyond linear programming,” *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3857–3860, 2008.
- [35] I. Bilik, “Spatial compressive sensing approach for field directionality estimation,” *IEEE Radar Conference*, pp. 1–5, 2009.
- [36] K. Egiazarian, A. Foi, and V. Katkovnik, “Compressed sensing image reconstruction via recursive spatially adaptive filtering,” in *Proc. of the IEEE International Conference on Image Processing (ICIP)*, 2007, vol. 1, pp. 549–452.
- [37] Y. Yu, A.P. Petropulu, and H.V. Poor, “Compressive sensing for MIMO radar,” in *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2009, pp. 3017–3020.
- [38] M. Lustig, D. Donoho, and J.M. Pauly, “Sparse MRI: The application of compressed sensing for rapid MR imaging,” *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [39] Z. He, A. Cichocki, R. Zdunek, and J. Cao, “CG-M-FOCUSS and Its Application to Distributed Compressed Sensing,” in *Proceedings of the 5th International Symposium on Neural Networks: Advances in Neural Networks*. Springer, 2008, pp. 237–245.
- [40] M. Tello, P. López-Dekker, and J.J. Mallorqui, “A Novel Strategy for Radar Imaging Based on Compressive Sensing,” in *Proc. of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2008, vol. 2, pp. 213–216.
- [41] J.L. Ponce-Davalos and Y.V. Shkvarko, “Superresolution of targets on the multi-grade scene: a spectral positional invariance-based approach,” in *Proc. of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2003, vol. 6, pp. 4047 – 4049.
- [42] M. Herman and T. Strohmer, “High-resolution radar via compressed sensing,” *IEEE Trans. Signal Processing*, vol. 57, no. 6, pp. 2275 – 2284, 2007.
- [43] E.J. Candès, J.K. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207, 2006.

- [44] E.J. Candès, “Compressive sampling,” *Proceedings of the International Congress of Mathematicians*, vol. 3, pp. 1433–1452, 2006.
- [45] T. Blumensath and M.E. Davies, “Gradient pursuits,” *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2370–2382, 2008.
- [46] A.K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Inc. Upper Saddle River, NJ, USA, 1989.
- [47] D.R. Wehner, *High resolution radar*, 1987.
- [48] W.G. Carrara, R.S. Goodman, and R.M. Majewski, “Spotlight synthetic aperture radar- Signal processing algorithms(Book),” *Norwood, MA: Artech House, 1995.*, 1995.
- [49] C.V. Jakowatz, D.E. Wahl, P.H. Eichel, D.C. Ghiglia, and P.A. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*, Kluwer Academic Pub, 1996.
- [50] M.F. Duarte and R.G. Baraniuk, “Recovery of frequency-sparse signals from compressive measurements,” *48th Annual Allerton Conference on Communication, Control, and Computing*, pp. 599 – 606, 2010.
- [51] T. Blumensath and M.E. Davies, “Iterative hard thresholding for compressed sensing,” *Journal of Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, 2009.
- [52] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [53] E.J. Candès and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?,” *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [54] P.S. Huggins and S.W. Zucker, “Greedy basis pursuit,” *Signal Processing, IEEE Transactions on*, vol. 55, no. 7, pp. 3760–3772, 2007.
- [55] E.J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.

- [56] S.P. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge Univ Pr, 2004.
- [57] S.J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky, “An Interior-Point Method for Large-Scale l_1 -Regularized Least Squares,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 606–617, 2007.
- [58] E.J. Candès, M.B. Wakin, and S.P. Boyd, “Enhancing sparsity by reweighted L_1 minimization,” *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 877–905, 2008.
- [59] I. Daubechies, R. DeVore, M. Fornasier, and C.S. Güntürk, “Iteratively reweighted least squares minimization for sparse recovery,” *Communications on Pure and Applied Mathematics*, vol. 63, no. 1, pp. 1–38, 2010.
- [60] G.H. Golub and C.F. Van Loan, *Matrix Computations*, Johns Hopkins Univ Pr, 1996.
- [61] V. Frayssé and L. Giraud, “A set of conjugate gradient routines for real and complex arithmetics,” Tech. Rep., CERFACS, 2000, Retrieved from www.cerfacs.fr.
- [62] D.E. Dudgeon and R.M. Mersereau, *Multidimensional Digital Signal Processing*, Prentice Hall Professional Technical Reference, 1990.
- [63] H.C. Stankwitz and M.R. Kosek, “Super spatially variant apodization (Super-SVA),” Nov. 11 1997, US Patent 5,686,922.
- [64] S. Foucher, “SAR Image Filtering Via Learned Dictionaries and Sparse Representations,” in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2008, vol. 1, pp. 229–232.
- [65] C. Jakowatz, *Signal Processing for Spotlight Mode SAR*, presented at Lockheed Martin, Goodyear, Arizona, 2004.