2012

Image reconstruction through polyfiltered variation minimization

Michael Hagen

Image Reconstruction Through Polyfiltered Variation Minimization

by

Michael Hagen

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts in Mathematics
(College of Arts & Sciences)
Eastern Michigan University
2012

Thesis Committee:

Dr. Andrew Ross, Chair
Dr. Jayakumar Ramanathan
Dr. William Sverdlik


EASTERN
MICHIGAN UNIVERSITY

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF SYMBOLS AND NOTATION

TV	Total Variation
PV	Polyfiltered Variation
BP	Basis Pursuit
BPDN	Basis Pursuit Denoise
DFT	Discrete Fourier Transform
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
SAR	Synthetic Aperture Radar
RFI	Radio Frequency Interference
$\mathcal{R}(\mathbf{A})$	Range of the linear transform $\mathbf{A}$
$\mathcal{N}(\mathbf{A})$	Null space of the linear transform $\mathbf{A}$
$\mathcal{F}(\mathbf{x})$	Fourier transform of $\mathbf{x}$
$\mathbf{a} * \mathbf{b}$	Discrete convolution of vectors $\mathbf{a}$ and $\mathbf{b}$
$a \otimes b$	Cartesian product of a and b
$\mathbf{a} \odot \mathbf{b}$	Hadamard product of vectors $\mathbf{a}$ and $\mathbf{b}$

ABSTRACT

Image Reconstruction Through Polyfiltered Variation Minimization

by

Michael Hagen

Chair: Dr. Andrew Ross

There has been considerable interest in reconstruction of remotely sensed imagery from incomplete frequency measurements for some time now. Given the nature of the collection process, it may be that portions of the spectrum are either missing or corrupted such that one is left with an incomplete representation of the original image. The advances in both the theory and available software for sparse signal reconstruction through function minimization make it an attractive approach for recreating the missing frequency data. It is the aim of this thesis to generalize the reconstruction technique known as Total Variation (TV) minimization from a signal processing perspective and to show that it is but one instance of a more general class of multi-filter operators. The approach will be demonstrated using freely available third-party software, and the reconstruction accuracy of TV minimization will be compared to that of several of the developed alternative operators. Last, the relationship between these operators and the frequencies to be reconstructed will be examined.

CHAPTER I

Theory And Background

1.1 Total Variation Minimization

It is well known that a sparse signal can be reconstructed exactly from incomplete frequency information using the ℓ_1 norm [1]. The most common formulations of the problem are usually referred to as Basis Pursuit (BP) and Basis Pursuit Denoise (BPDN) [2] and are given by

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{A}\mathbf{x} = \mathbf{b} \tag{1.1}$$

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \sigma, \tag{1.2}$$

respectively, where $\mathbf{b} \in \mathbb{C}^m$ is the vector of known frequency measurements and $\mathbf{A} \in \mathbb{C}^{m \times n}$, $m < n$, is the incomplete DFT matrix the rows of which correspond to the known frequencies. $\mathbf{A}$ and $\mathbf{b}$ may be some other basis in which the signal can be represented and need not be frequency. For this thesis, however, we will solely use orthogonal frequency components in the constraints. A requirement for the above reconstruction to be successful is that the signal meet some threshold of sparseness with respect to the ℓ_1 norm [3]; otherwise, there will not be enough information in the incomplete frequency data and, in the limiting case as the signal becomes more and more dense, the entire set of orthogonal frequency measurements

will be required. In this last case, $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ and no minimization is necessary. For completeness of the above formulation, we note that for the Cartesian product $\mathbf{x} \in \mathbb{C}^K \otimes \mathbb{C}^K \otimes \cdots \otimes \mathbb{C}^K$, the definition of the ℓ_1 norm is

$$\|\mathbf{x}\|_1 = \sum_{k=0}^{K-1} \|x_k\|_2, \tag{1.3}$$

which for the most common cases of $\mathbf{x} \in \mathbb{R}^K$ and $\mathbf{x} \in \mathbb{C}^K$ works out to be the sum of absolute values and the sum of the moduli, respectively.

For cases where the signal may not be sparse in the ℓ_1 norm, other measures are used as the objective function, measures for which the signal is expected to have a greater degree of sparsity. One of these alternate measures is the Total Variation (TV) of the signal. The TV operator is generally intended for 2-D signals and can be thought of as the ℓ_1 norm of an estimate of the gradient:

$$\mathrm{TV}(\mathbf{X}) = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \sqrt{\left(\mathrm{X}_{i,j+1} - \mathrm{X}_{i,j}\right)^2 + \left(\mathrm{X}_{i+1,j} - \mathrm{X}_{i,j}\right)^2}, \tag{1.4}$$

where $\mathbf{X} \in \mathbb{R}^{M \times N}$. Thus, the general TV problem is similar to the BP problems above:

$$\min_{\mathbf{X}} \mathrm{TV}(\mathbf{X}) \text{ s.t. } \mathbf{A}\mathbf{x} = \mathbf{b} \tag{1.5}$$

$$\min_{\mathbf{X}} \mathrm{TV}(\mathbf{X}) \text{ s.t. } \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \sigma \tag{1.6}$$

where $\mathbf{x} = \mathrm{vec}(\mathbf{X})$ is a vectorized version of the image and the DFT constraint matrix is modified accordingly such that $\mathbf{A} \in \mathbb{R}^{K \times MN}$, $K < MN$, and $\mathbf{b} \in \mathbb{R}^K$.

1.2 Total Variation As A Filter

Another way to view the TV operator is as the ℓ_1 norm of the vectorized output of two FIR filters applied to the image. If we consider the first-order filter $\mathbf{H} =$

$[1; -1]$, then

$$\mathrm{TV}(\mathbf{X}) = \|\mathrm{vec}(\mathbf{H}^{\mathrm{T}} * \mathbf{X}) \otimes \mathrm{vec}(\mathbf{H} * \mathbf{X})\|_1, \tag{1.7}$$

where $\otimes$ indicates a Cartesian product. (We will often omit the vec operator for readability though we include it here since the ℓ_1 norm has a different definition when applied to matrices than it does to vectors.) Elementary filter theory tells us that we can rewrite this in terms of the Fourier transform, $\mathcal{F}$:

$$\mathbf{H} * \mathbf{X} = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{H}) \odot \mathcal{F}(\mathbf{X})), \tag{1.8}$$

so that we may analyze what it is that TV is measuring by looking at the product of the filter's spectrum and the spectrum of the input image. Since the filter is one-dimensional, it suffices to plot the frequency response of $\mathbf{H}$ knowing that $\mathbf{H}^{\mathrm{T}}$ has the same effect but along the rows of the image. This is done in Figure 1.1, where the high-pass nature of the TV operation can be seen. This high-pass nature illustrates how using TV rather than a straight ℓ_1 can give a sparse representation to an otherwise dense signal, viz. by filtering out the signal's lower frequency components. It also raises the question of the effectiveness of using other filters that have different passbands and transition regions than the first-order TV filter given above.

1.3 Polyfiltered Variation

We would like to expand the TV problem to a more general operation, and this is easy to do in light of the definition used in (1.7). All that is required is to exchange the filter $\mathbf{H}$ for whatever filters are of interest. Moreover, there is no reason to limit the operation to two filters. We refer to this generalization as *polyfiltering* and de-

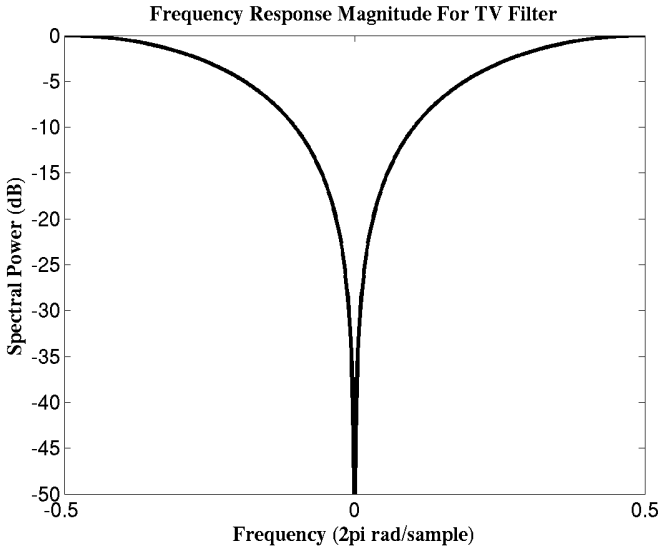

Figure 1.1: The magnitude of the frequency response for the TV filter.

fine it as follows

$$\mathrm{P}(\mathbf{X}) = (\mathbf{H}_1 * \mathbf{X}) \otimes (\mathbf{H}_2 * \mathbf{X}) \otimes \cdots \otimes (\mathbf{H}_L * \mathbf{X}). \tag{1.9}$$

We then define *Polyfiltered Variation* (PV) as the ℓ_1 norm of the vectorized version of the above L-dimensional Cartesian product:

$$\begin{aligned} \mathrm{PV}(\mathbf{X}) &= \|\mathrm{P}(\mathbf{X})\|_1 \\ &= \sum_i \sqrt{p_{i1}^2 + p_{i2}^2 + \cdots + p_{iL}^2}, \end{aligned} \tag{1.10}$$

where $p_{i1} \otimes p_{i2} \otimes \cdots \otimes p_{iL}$ is the i^{th} element of $\mathrm{P}(\mathbf{X})$.

An implicit assumption made in the above definition of (1.9) is that $\dim \mathbf{H}_i = \dim \mathbf{H}_j$ for all individual filters. If this is not the case, then the filter with smaller dimension may be padded with zeros so that the dimensions match up. We will make the assumption from here on that this is the case. To give an example of this zero padding as well as a specific instance of polyfiltering, if we set $L = 2$ and

choose

$$\mathbf{H}_1 = \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix} \tag{1.11}$$

$$\mathbf{H}_2 = \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix}, \tag{1.12}$$

then the PV reduces to the original TV formulation.

1.4 Significance Of The PV Interpretation

With the above interpretation, we can reformulate the TV minimization problem to a more general case of which TV is one specific instance:

$$\min_{\mathbf{x}} \mathrm{PV}(\mathbf{x}) \ \text{ s.t. } \ \mathbf{A}\mathbf{x} = \mathbf{b} \tag{1.13}$$

$$\min_{\mathbf{x}} \mathrm{PV}(\mathbf{x}) \ \text{ s.t. } \ \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \sigma. \tag{1.14}$$

The significance of developing the problem as a series of FIR filters is that our constraints are composed of incomplete frequency measurements. Since the effect of an FIR filter is determined by its frequency response, it stands to reason that we would expect different filters to affect the problem differently. Moreover, we would expect there to be a relation between the frequency response of the filters used in the PV minimization and the set of incomplete frequency measurements.

CHAPTER II

Application

2.1 Problem Application

The problem to which we would like to apply PV minimization is estimating missing frequency samples from 2-D imagery. We estimate the missing frequency samples by estimating the complete image, that is, the image formed from the full set of orthogonal Fourier basis vectors. Starting with the image shown in Figure 2.1, we simulate the corruption or unavailability of frequency samples by zeroing out a subset of them. This is shown in Figure 2.2 where the noticeably distorted incomplete image is shown along with the portions of the spectrum that were zeroed. To get an estimate of the complete image from the incomplete set of frequency samples, we solve

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \mathrm{PV}(\mathbf{x}) \text{ s.t. } \mathbf{A}\mathbf{x} = \mathbf{b}, \tag{2.1}$$

where we have used equality constraints rather than the more general case of $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 < \sigma$ for reasons stemming from the software used to solve the problem (see Section 3.3).

In the above formulation, $\mathbf{x} \in \mathbb{C}^{MN}$ is a vectorized representation of an $M \times N$ image. The rows of the matrix $\mathbf{A}$ are the Fourier vectors corresponding to the known frequency samples, and the vector $\mathbf{b}$ is the vector of these known frequency sam-

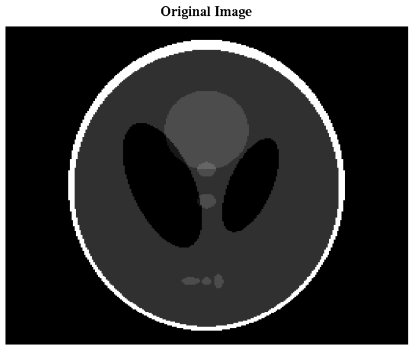

(a) Original Image

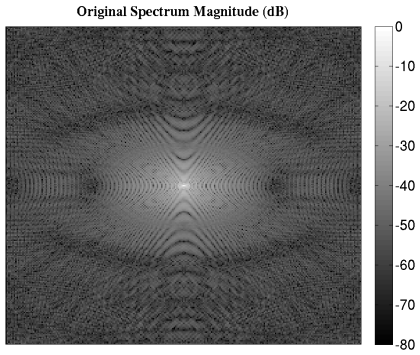

(b) Original Spectrum

Figure 2.1: Original (complete) image and the magnitude of its frequency representation.

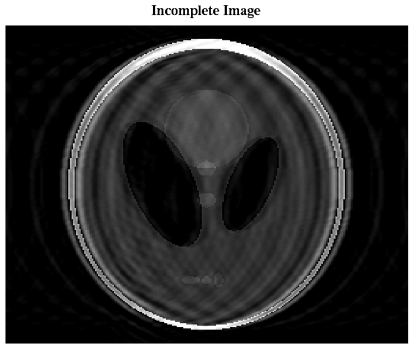

(a) Incomplete Image

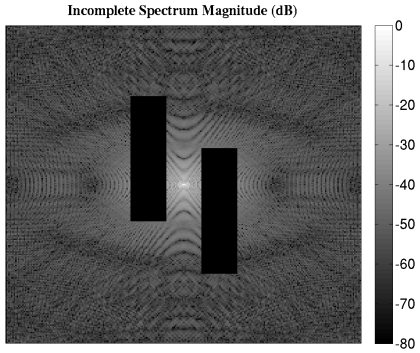

(b) Incomplete Spectrum

Figure 2.2: Incomplete image and the magnitude of its frequency representation. The black regions of the spectrum signify the missing frequency measurements.

ples, viz. all non-blacked-out samples in Figure 2.1(b). The solution to 2.1, $\mathbf{x}^*$, is what is used as an estimate of the complete image. For the current example, this estimate and its spectrum are shown in Figure 2.3.

Theory shows that the above approach is valid for l_0, ℓ_1, and TV minimization if the complete image has a required degree of sparsity or is sparse after applying the TV transformation [1], [3], [4]. We will not address this theory here. To the contrary, we will operate under the conjecture that these theoretical results extend to PV minimization.

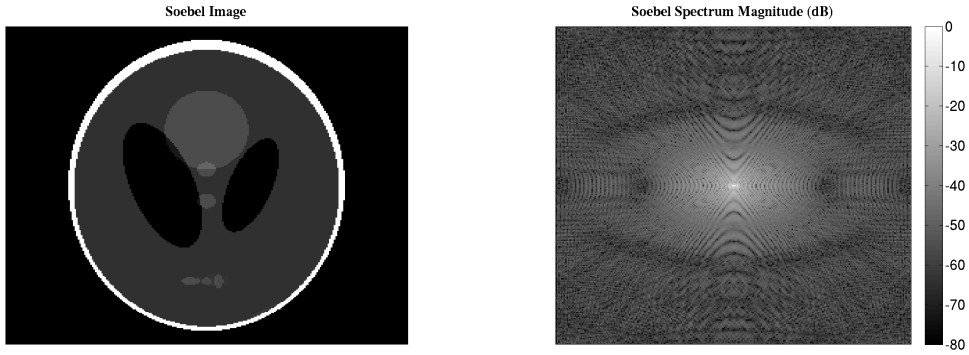

(a) Reconstructed Image (b) Reconstructed Spectrum

Figure 2.3: Estimate of the complete image and the magnitude of its frequency representation. PV minimization was performed using Soebel filters in this case.

2.2 Optical Imagery

Fourier optics shows that various lenses, mirrors, and other apparatuses such as slits can be described by Fourier analysis, which is to say that their effects can be defined by how they act on individual frequency components. If an optical sensor uses one of these components that is defective for some reason (say, a dirty lens or improperly manufactured mirror), then it can result in imagery formed from a mix of uncorrupted and corrupted frequency samples. If this sensor is irretrievable (e.g., on an orbiting platform) or if it is cost-prohibitive to replace the defective components, then the most attractive option may be to try to reconstruct the corrupted frequency samples in software using these ℓ_1 minimization techniques.

2.3 Radar Imagery

Synthetic Aperture Radar (SAR) imagery can likewise suffer from corrupted frequency measurements. The causes may be hardware related, similar to the above examples for an optical sensor; however, SAR sensors tend to operate at frequencies that are heavily used for other purposes such as electronic communications, thereby

making them susceptible to Radio Frequency Interference (RFI). A common effect of RFI is to saturate certain portions of the image's spectrum with what is essentially noise, which can render the image unusable unless it is somehow mitigated. Looking at the SAR image and its spectrum shown in Figure 2.4, we can illustrate a common effect of RFI by adding high-powered noise to its spectrum, as is done in Figure 2.5 where the white bands show the location of the added noise. One can see from the latter that the effect of the RFI is to make the image so noisy that only the brightest of targets are still visible. Overall, the image has lost all practical interpretability. The simplest approach to take in order to mitigate the RFI is to zero out the affected frequency samples and form the image from the incomplete set. The results of doing this are shown in Figure 2.6. It appears that zeroing out the RFI has successfully ameliorated its effects, and indeed it has with respect to the added noise. However, in doing so, a large distortion has been introduced into the image. In Figure 2.7, a close-up of a region of the image is shown for both the complete (original) and incomplete cases. The RFI has been removed but at a cost of a large amount of image distortion. This is one area where PV minimization could be applied. Namely, after zeroing out the RFI, the remaining frequency samples can be used in a PV minimization problem in order to estimate the complete image and remove the introduced distortion.

(a) Original Image

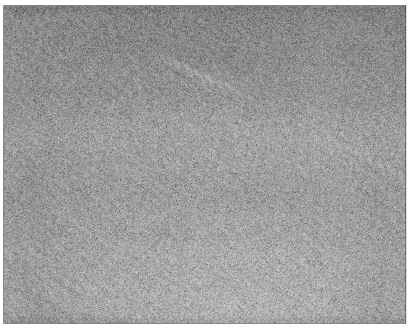

(b) Original Spectrum

Figure 2.4: Original (complete) image and the magnitude of its frequency representation.

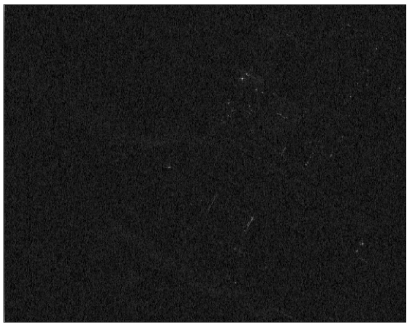

(a) Image With Synthetic RFI Added

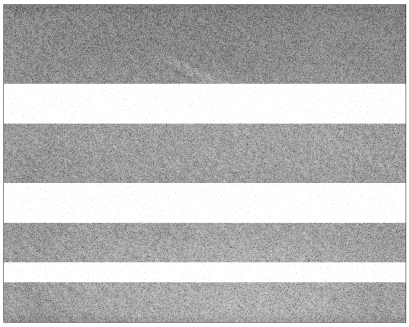

(b) Spectrum With Synthetic RFI Added

Figure 2.5: Result of adding synthetic high-power RFI to the image in Figure 2.4. The white bands in the spectrum show the added noise.

(a) Incomplete Image With RFI Zeroed

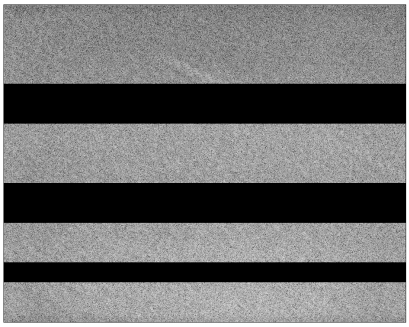

(b) Incomplete Spectrum With RFI Zeroed

Figure 2.6: Result of zeroing out synthetic high-power RFI in Figure 2.5(a). The zeroes in the spectrum are visible as the black bands.

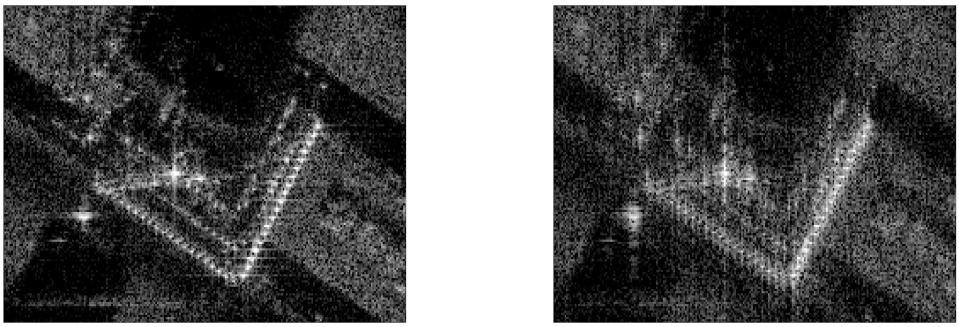

(a) Original Image (b) Incomplete Image

Figure 2.7: Close-up comparison of the original image and the incomplete image that results from zeroing out the RFI. Forming the image from the incomplete set of frequency samples has introduced noticeable distortion.

CHAPTER III

Software Implementation

3.1 Polyfiltering As A Single Linear Transformation

Convolution, being a linear transformation, has a matrix representation. The construction of such a matrix is straightforward and is little more than the elements of one of the vectors, usually the smaller of the two, repeated but shifted on each row. For our case, it would be a matrix the rows of which are shifted copies of the filter coefficients padded with zeros to match the dimensions. As to notation, we will use the same symbol $\mathbf{H}$ when referring to either the filter or its matrix representation and rely on context to make clear which is implied. Thus, $\mathbf{H} * \mathbf{X}$ would imply that $\mathbf{H}$ is the array of filter coefficients. Likewise, $\mathbf{HX}$ would indicate that $\mathbf{H}$ is the matrix representation of such.

We can represent the polyfiltering operation as a single matrix $\mathbf{P}$ by concatenating the L dimensions of $\mathrm{P}(\mathbf{X})$ together. To do so, all that is required is to stack the filtering matrices on top of one another:

$$\mathrm{P}(\mathbf{X}) = \begin{bmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \vdots \\ \mathbf{H}_L \end{bmatrix} \mathbf{X} = \mathbf{PX}. \tag{3.1}$$

Moreover, we can transform the 2-D input array $\mathbf{X} \in \mathbb{R}^{M\times N}$ to a 1-D signal by stacking its columns on top of one another to form $\mathbf{x} = \text{vec}(\mathbf{X}) \in \mathbb{R}^{MN}$. We will take for granted the well known fact that given the matrix-matrix product $\mathbf{AX}, \mathbf{A} \in \mathbb{R}^{L\times M}, \mathbf{X} \in \mathbb{R}^{M\times N}$, there is an equivalent matrix $\mathbf{A}_{\text{vec}} \in \mathbb{R}^{MN\times MN}$ such that $\text{vec}(\mathbf{AX}) = \mathbf{A}_{\text{vec}}\mathbf{x}$. We will use the lowercase $\mathbf{x}$ when referring to the vectorized form of an image and will omit the vec subscript from the matrix operator, relying instead on whether the right-hand element in the product is upper- or lowercase to make clear the form of the matrix. With this in mind, the polyfiltered signal will often be written in vectorized form as

$$\text{P}(\mathbf{x}) = \mathbf{Px}. \tag{3.2}$$

This vectorized representation allows for an alternate expression for the polyfiltered variation defined previously in (1.10). Defining the selection matrix $\mathbf{S}_i$ as an $L \times MN$ matrix of all zeros except for ones in elements $(0, i), (1, i + MN), (2, i + 2MN), \cdots, (L - 1, i + (L - 1)MN)$ allows us to represent the i^{th} element of the Cartesian product $\text{P}(\mathbf{x})$ as the $L \times 1$ vector $\mathbf{S}_i\mathbf{Px}$. From this point it is a simple extension to arrive at

$$\text{PV}(\mathbf{x}) = \sum_i \|\mathbf{S}_i\mathbf{Px}\|_2. \tag{3.3}$$

We thereby have an alternate expression for the PV of an image as a general sum of ℓ_2 norms.

3.2 Convex Optimization Using PV

Referring to the form of polyfiltering given in (3.2), we have that $\text{P} : \mathbb{R}^{MN} \rightarrow \mathbb{R}^{LMN}$, or, equivalently, $\mathbf{P} \in \mathbb{R}^{LMN\times MN}$, which means that $\mathbf{P}$ has a non-trivial nullspace $\mathcal{N}$. We mention this because for any $\mathbf{x} \in \mathcal{N}(\mathbf{P})$ we have that $\text{PV}(\mathbf{x}) = 0 \nLeftrightarrow \mathbf{x} = \mathbf{0}$, thereby making PV a semi-norm rather than a true norm. However,

if we restrict the input to a subspace such that $\mathrm{PV}(\mathbf{x}) = 0 \Leftrightarrow \mathbf{x} = \mathbf{0}$, then we can treat $\mathrm{PV}(\mathbf{x})$ as a true norm. From the fundamental theorem of linear algebra we have that $\mathcal{N}(\mathbf{P}) \perp \mathcal{R}(\mathbf{P}^{\mathrm{T}})$. Therefore, if we restrict $\mathbf{x}$ to the range of the adjoint of $\mathbf{P}$, then $\mathrm{PV}(\mathbf{x})$ is a norm. This is significant because it gives us a convex function (the norm) over a convex set (any vector subspace is a convex set), which then means that the problem

$$\min_{\mathbf{x} \in \mathcal{R}(\mathbf{P}^{\mathrm{T}})} \mathrm{PV}(\mathbf{x}) \text{ s.t. } \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \sigma \tag{3.4}$$

has a unique global minimum [5]. Given the interpretation of PV as set of FIR filters, the nullspace of $\mathbf{P}$ is simply composed of the nulls of the filters. We can thus restate the above as requiring that $\mathbf{x}$ not contain any frequency components for which the frequency response of all the used filters is zero. However, it is expected that frequencies which, while not identically nulled to zero value, are attenuated a significant amount (the definition of 'significant' purposely being left open-ended) will play a role numerically due to either algorithmic implementations or finite precision arithmetic.

3.3 TVAL3

In order to solve the PV minimization problem, we have modified the *TVAL3* MATLAB software written by Cheng Bo Li [6]. We have done this by augmenting it to implement any arbitrary linear transform rather than only the TV operation for which it was written. We represent each specific PV implementation as a single matrix as in (3.1). This gives us the sum of l_2 norms cost function in (3.3). This in turn allows for a simple, closed-form gradient that we can then substitute into the augmented Lagrangian equations used by the *TVAL3* solver regardless of the linear transformation. This allows us to use the same solver for ℓ_1, TV, and all the other

PV implementations discussed herein.

The main drawback to the *TVAL3* software is that the augmented Lagrangian method it uses to solve the problem does not allow us to solve the more general problem of

$$\min_{\mathbf{x}} \mathrm{PV}(\mathbf{x}) \;\; \text{s.t.} \;\; \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \sigma. \tag{3.5}$$

Rather, we are restricted to equality constraints such that all results herein are found from the following formulation:

$$\min_{\mathbf{x}} \mathrm{PV}(\mathbf{x}) \;\; \text{s.t.} \;\; \mathbf{A}\mathbf{x} = \mathbf{b}. \tag{3.6}$$

Even with this limitation, *TVAL3* was deemed superior to other packages such as [7], [8], and [9] due to the inability of the others to accurately process images larger than 128×128 pixels without excessive memory requirements or numerical instability problems.

CHAPTER IV

Example PV Implementations

4.1 General Purpose

4.1.1 ℓ_1

The most basic of PV implementations is to use a single identity filter, $\mathbf{H} = [1]$. In this case, the PV problem reduces to the standard ℓ_1 minimization. The filter is an all-pass filter as it leaves all frequency components unchanged.

4.1.2 TV

To implement the TV problem using the PV framework, we use the following two previously mentioned filters:

$$\mathbf{H}_1 = \begin{bmatrix} 1 & 0 \\ -1 & 0 \end{bmatrix} \tag{4.1}$$

$$\mathbf{H}_2 = \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix}, \tag{4.2}$$

The frequency responses of the two filters are shown in Figure 4.1.

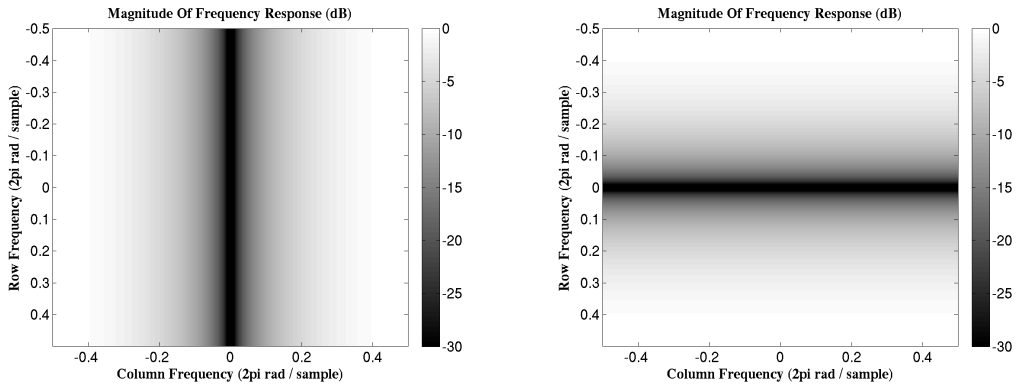

(a) Filter 1 (b) Filter 2

Figure 4.1: Frequency responses of the two TV filters.

4.1.3 Soebel

The Soebel filters see applications most often in edge detection and gradient estimation [10]. The filters have the following form:

$$\mathbf{H}_1 = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix} \tag{4.3}$$

$$\mathbf{H}_2 = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}. \tag{4.4}$$

The frequency responses of the two filters are shown in Figure 4.2.

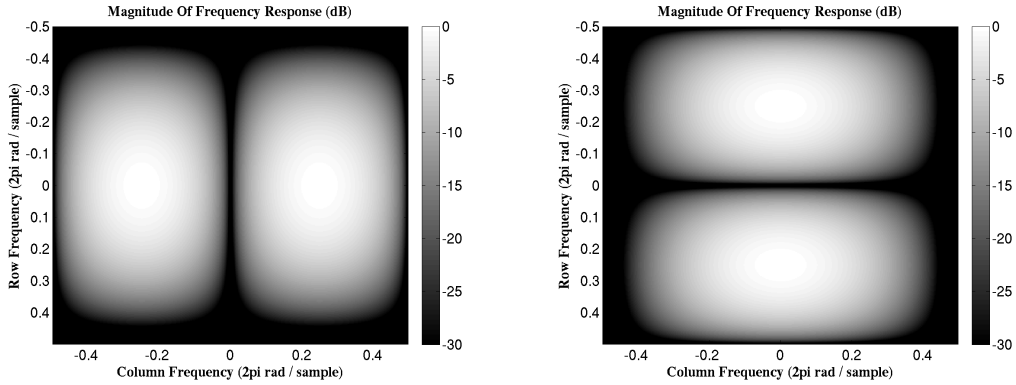

(a) Filter 1 (b) Filter 2

Figure 4.2: Frequency responses of the two Soebel filters.

4.1.4 Prewitt

Similar to the Soebel filters, the Prewitt filters see applications most often in edge detection and gradient estimation [10]. The filters have the following form:

$$\mathbf{H}_1 = \begin{bmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{bmatrix} \tag{4.5}$$

$$\mathbf{H}_2 = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}. \tag{4.6}$$

The frequency responses of the two filters are shown in Figure 4.3. Moreover, though we do not implement it here, given the general form of PV and how the software is set up, we could easily add a third and fourth filter, each meant to act along the diagonals of an image rather than the rows and columns. These diagonal Prewitt

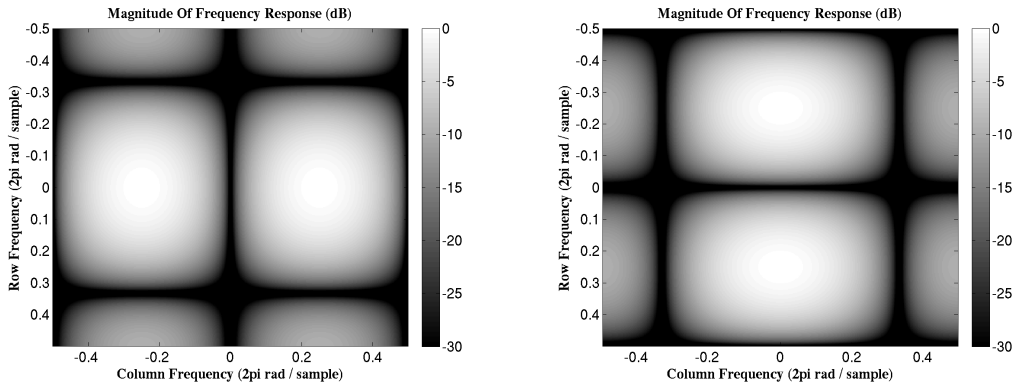

(a) Filter 1 (b) Filter 2

Figure 4.3: Frequency responses of the two Prewitt filters.

filters are given by [10]

$$\mathbf{H}_3 = \begin{bmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{bmatrix} \tag{4.7}$$

$$\mathbf{H}_4 = \begin{bmatrix} -1 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}. \tag{4.8}$$

4.1.5 2nd Derivative

The last general purpose PV implementation we will describe is a set of three filters meant to estimate the 2nd derivative of an image:

$$\mathbf{H}_1 = \begin{bmatrix} 1 & -2 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{4.9}$$

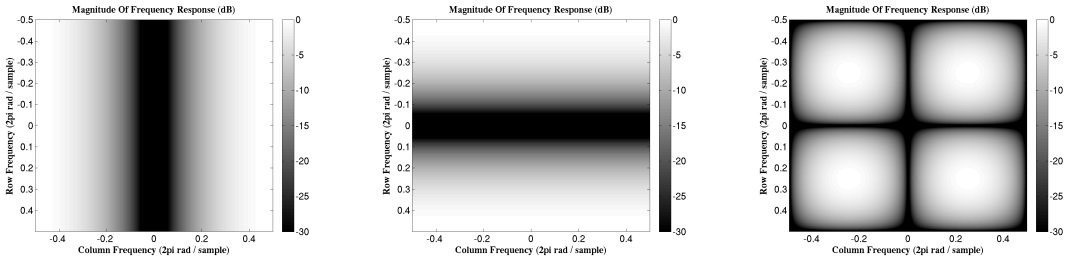

(a) Filter 1 (b) Filter 2 (c) Filter 3

Figure 4.4: Frequency responses of the filters used to estimate the 2nd derivative of the image.

$$\mathbf{H}_2 = \begin{bmatrix} -1/4 & 0 & 1/4 \\ 0 & 0 & 0 \\ 1/4 & 0 & -1/4 \end{bmatrix} \tag{4.10}$$

$$\mathbf{H}_3 = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \tag{4.11}$$

The frequency responses of these three filters are shown in Figure 4.4.

4.2 Problem Specific

4.2.1 TV + Notch Filter

Given the flexibility of the PV solver to work with any linear transformation, we can also use filters that are tailored to specific problems. Here we will combine the same TV filters described above with a notch filter meant to remove energy known to cause poor performance. More specifically, consider the input optical image shown in Figure 4.5. We have modified the spectrum of the original image such that it has two small areas of concentrated energy, visible as the two bright circles in Figure 4.5(b). In Figure 4.6 is shown the missing frequency samples that we will be estimating and the resulting incomplete image.

(a) Complete Image

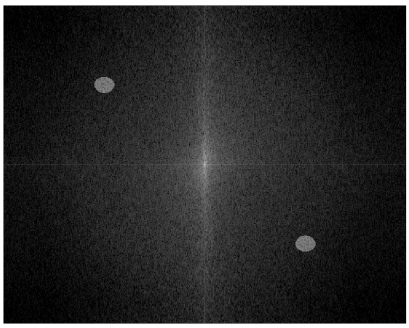

(b) Complete Spectrum

Figure 4.5: Input image used in the TV + Notch example and its spectrum.

(a) Incomplete Image

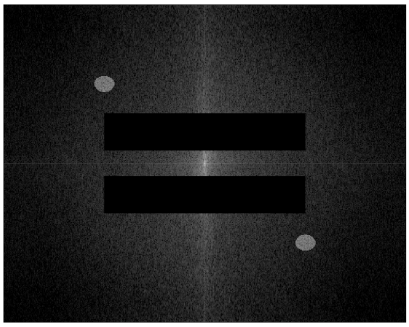

(b) Incomplete Spectrum

Figure 4.6: The incomplete image and its spectrum used in the TV + Notch example.

We know from the basic properties of the Fourier transform that we expect the energy in these small areas of the frequency domain to be diffuse in the image domain. Furthermore, when we look at the frequency responses of the TV filters shown previously in Figure 4.1, we know that this concentrated energy will be passed by both of them nearly unchanged. This tells us that the usual approach to reconstructing the missing data (i.e., TV minimization) will not perform well because TV($\mathbf{x}$) will not give a sparse representation. To address this we can augment the TV filters to include notches where the concentrated energy is located. An example of this TV + Notch design is shown in Figure 4.7.

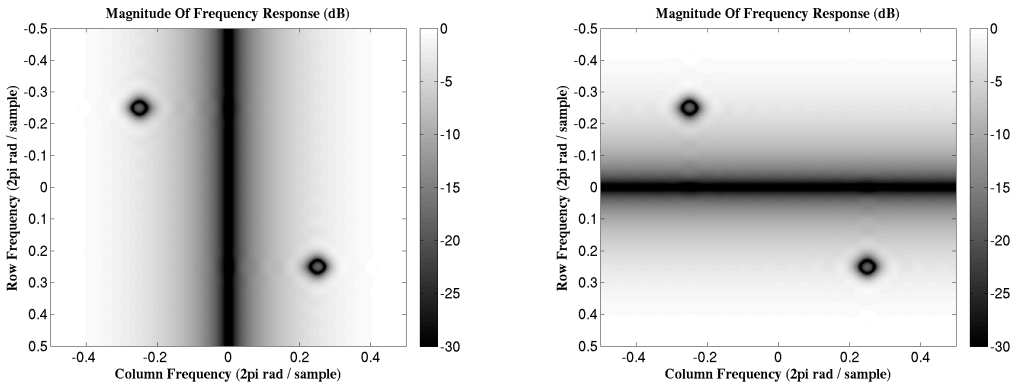

(a) Filter 1 (b) Filter 2

Figure 4.7: Frequency responses of the TV + Notch filters.

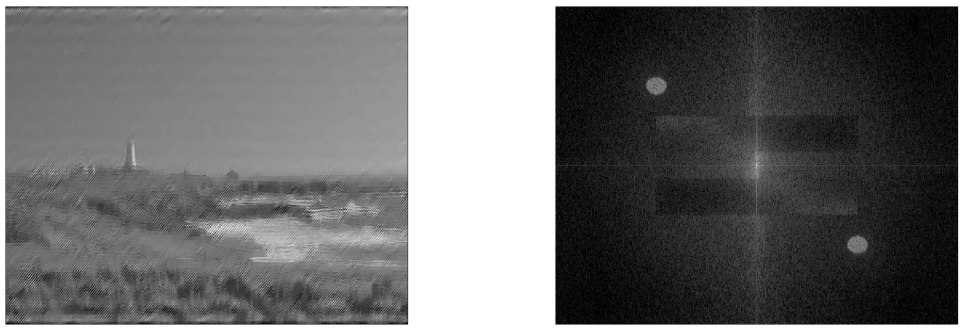

(a) Reconstructed Image (b) Reconstructed Spectrum

Figure 4.8: The estimated complete image and its spectrum as found using TV minimization.

To conclude this example we show the results of reconstructing the above incomplete image using the usual TV minimization in Figure 4.8 and the results from using the TV + Notch implementation in Figure 4.9. Though neither method reconstructs the image without error, by using the TV + Notch rather than the usual TV, the reconstruction error (defined in Section 5.1) was reduced 27% from 1.1512 to 0.8447. Moreover, the correlation coefficient between the original and reconstructed frequency samples was increased from 0.2341 to 0.5384, which is an improvement of 130%.

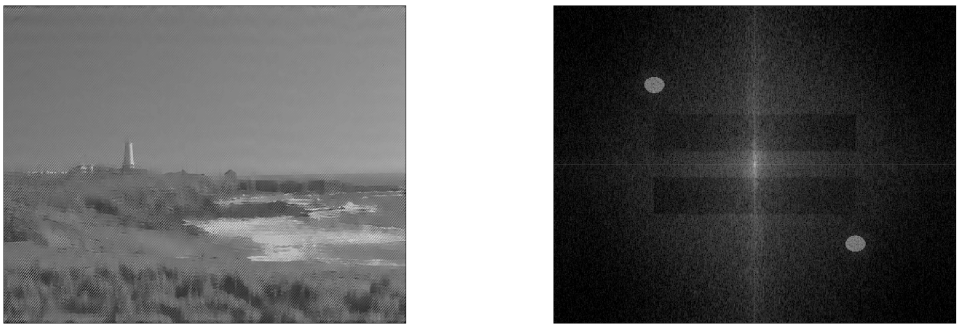

(a) Reconstructed Image (b) Reconstructed Spectrum

Figure 4.9: The estimated complete image and its spectrum as found using TV + Notch minimization.

4.2.2 ℓ_1 + Matched Filter

Another instance of a problem specific filter is a matched filter used to compress a signal with broad support. If we consider a quadratic phase complex exponential with uniform amplitude, then we have the following general form for it and its Fourier transform:

$$\text{rect}\left(\frac{t}{T}\right) e^{j\pi\alpha t^2} \overset{\mathcal{F}}{\longleftrightarrow} \text{rect}\left(\frac{f}{\alpha T}\right) e^{-j\pi f^2/\alpha}, \tag{4.12}$$

where $j = \sqrt{-1}$, f is frequency, T is the signal duration, and α is known as the chirprate and has units of Hz/s^2. Note that the term αT on the right hand side is the bandwidth of the signal, that is, the difference between the maximum and minimum frequencies. To implement a matched filter, we convolve this signal with a time-reversed and conjugated version of itself. This gives

$$\text{rect}\left(\frac{t}{T}\right) e^{j\pi\alpha t^2} * \text{rect}\left(\frac{-t}{T}\right) e^{-j\pi\alpha(-t)^2} \overset{\mathcal{F}}{\longleftrightarrow} \text{rect}\left(\frac{f}{\alpha T}\right) e^{-j\pi f^2/\alpha} e^{j\pi f^2/\alpha}. \tag{4.13}$$

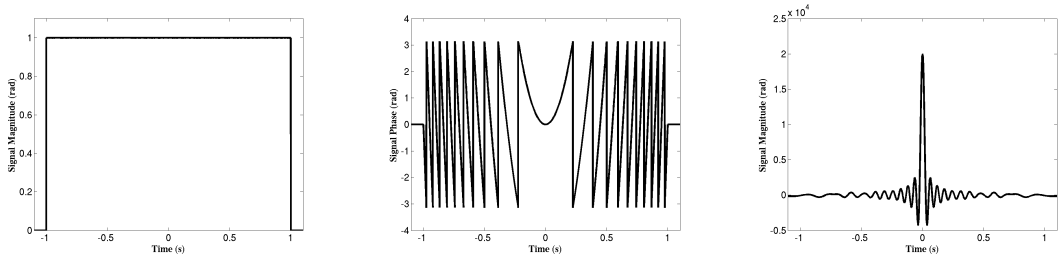

(a) Amplitude (b) Phase (c) Matched Filter Output

Figure 4.10: The phase and amplitude of a quadratic phase signal along with its output from a matched filter.

The two complex exponentials on the right cancel out, and we are left with just the rect function in the frequency domain, the inverse Fourier transform of which is

$$\text{rect}\left(\frac{f}{\alpha T}\right) \overset{\mathcal{F}^{-1}}{\longleftrightarrow} \text{sinc}\left(\alpha T t\right). \tag{4.14}$$

Thus, the result of matched filtering the quadratic phase signal is that it yields a sinc function, the width of which is inversely proportional to αT, the bandwidth of the signal. The foregoing is illustrated in Figure 4.10

The above is easily extended to 2-D such that we can form an image from a quadratic phase signal (also called a target) that has wide support in the image domain but is compressible with a matched filter. Taking this further, we know that the usual ℓ_1 reconstruction will perform poorly on such a signal since it is not sparse; however, since matched filtering is a convolution (more generally it is a linear transformation), we can incorporate the matched filter into the PV problem formulation and, according to our conjecture, improve the performance. To work through this we begin with the image of a single target in Figure 4.11, where the 2-D quadratic phase is also shown. We then form an image with 20 such targets located randomly throughout the image. This is shown in Figure 4.12 along with the result of applying a matched filter to the image, which shows clearly the sparsity it induces.

To show the difference in reconstruction performance between ℓ_1 reconstruction and

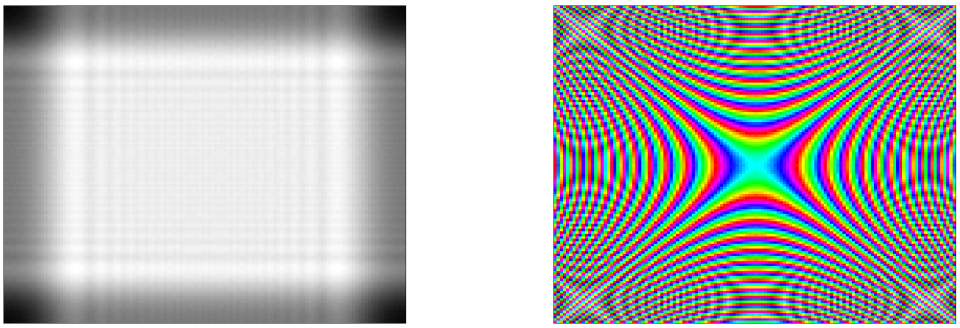

(a) Single Target Image (b) Single Target Phase

Figure 4.11: Amplitude of single quadratic phase target with wide support in the image domain. On the right is its 2-D quadratic phase.

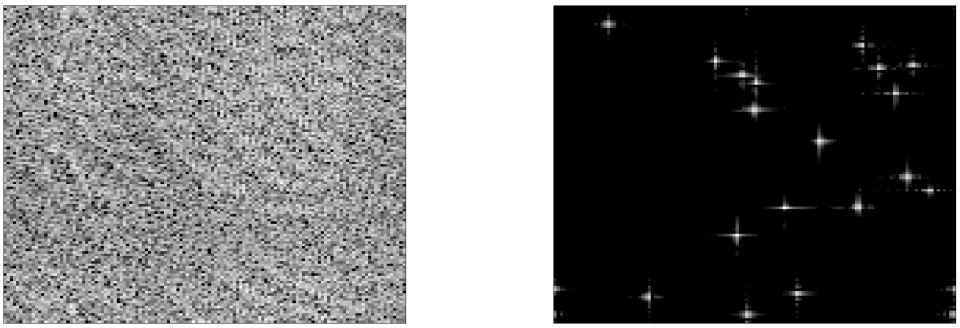

(a) 20 Target Image (b) Matched Filter Output

Figure 4.12: Amplitude of 20 quadratic phase signals summed together and the output after matched filtering.

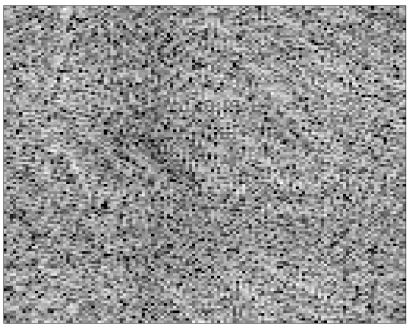

(a) Incomplete 20 Target Image

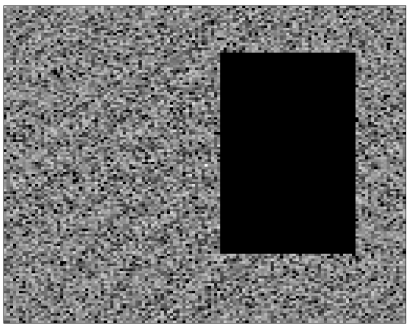

(b) Incomplete 20 Target Spectrum

Figure 4.13: The incomplete image used in the ℓ_1 and ℓ_1 + matched filter reconstructions.

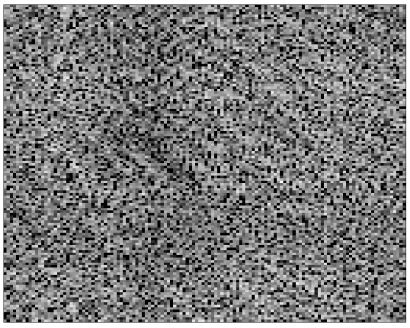

(a) ℓ_1 Reconstructed Image

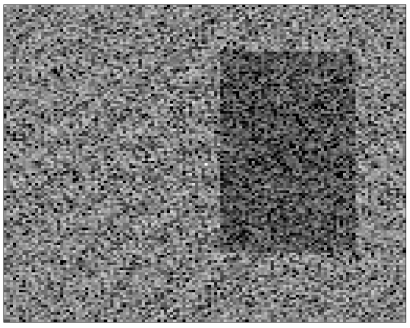

(b) ℓ_1 Reconstructed Spectrum

Figure 4.14: The estimated complete image as found using ℓ_1 minimization.

a PV implementation that combines ℓ_1 with a matched filter, we use the incomplete image illustrated in Figure 4.13. The results of the ℓ_1 reconstruction are pictured in Figure 4.14 and the results of combining the ℓ_1 with a matched filter in Figure 4.15. Quantitatively, the reconstruction error (see Section 5.1) was reduced from 1.1267 to 0.6387 by using the matched filter, and the correlation coefficient increased by two orders of magnitude from 0.0129 to 0.7808. We thus see a significant improvement by tailoring the PV implementation to fit the data rather than using the typical ℓ_1 reconstruction method.

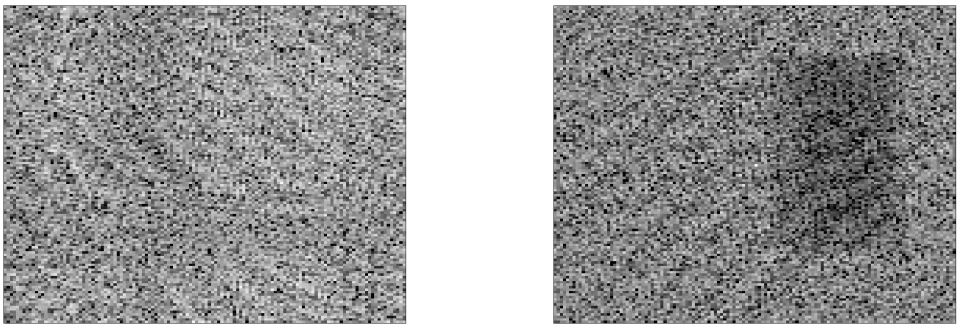

(a) ℓ_1 + Matched Filter Reconstructed Image (b) ℓ_1 + Matched Filter Reconstructed Spectrum

Figure 4.15: The estimated complete image as found using ℓ_1 + matched filter minimization.

CHAPTER V

Experimental Results

5.1 Reconstruction Metrics

We use two metrics to evaluate the performance of the various reconstructions. The first is a weighted two-norm of the difference between the original and missing frequency samples; the second is the correlation coefficient between these same two sets. Specifically, if $\mathbf{X}_m$ is the set of missing frequency samples and $\mathbf{X}_r$ the set of the reconstructed frequency samples, then the first type of error, which we call the reconstruction error, is given by

$$\epsilon_r = \frac{\|\mathbf{X}_m - \mathbf{X}_r\|_2}{\|\mathbf{X}_m\|_2}. \tag{5.1}$$

The second metric is a measure of the similarity between the missing components of the image and their reconstructed counterparts. Since the correlation coefficient is invariant to Fourier transforms, we can calculate this relation over the same sets as above:

$$\sigma_r = \mathrm{corr}\left(\mathbf{X}_m, \mathbf{X}_r\right). \tag{5.2}$$

The reconstruction error will give an idea of the relative distance between the original and the estimated frequency samples. The correlation coefficient, on the other hand, is meant to provide an idea of how similar the reconstructed image is to the

(a) Optical Image 1

(b) Optical Image 2

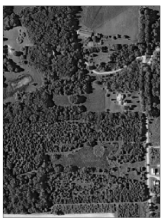

(c) Optical Image 3

Figure 5.1: The three rural optical test images. The three images comprise a contiguous area of rural Michigan west of Ann Arbor.

original regardless of the distance between them.

5.2 Optical

We will run two experiments on six overhead optical images: three of a rural area and three of an urban area. The images are publicly available from ESRI's ArcGIS Web Map Server [11] and are shown in Figures 5.1 and 5.2. Each image is 556 by 413 pixels with 1 meter spacing, which covers 0.005 degrees of both latitude and longitude. (Note that they have been displayed with square pixels so as not to distort them, while their spectrums have not [e.g., Figure 5.3]. Hence, the spectrums will appear square rather than rectangular like the input images do. We have made this choice to minimize the un-used whitespace in the images.)

We will use four PV implementations for each case and each image: TV, Soebel, Prewitt, and 2^{nd} derivative. The definition and frequency responses of each were given in Section 4.1. We do not use ℓ_1 as it is known to perform poorly in this situation due to the lack of sparsity in optical images of the earth's surface (thus transforms are used to induce sparsity).

(a) Optical Image 4

(b) Optical Image 5

(c) Optical Image 6

Figure 5.2: The three urban optical test images. The three images comprise a contiguous area of downtown Ypsilanti, Michigan.

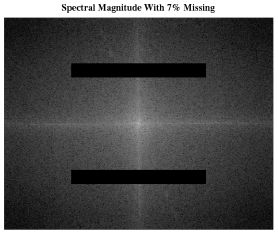

(a) 7% Block Size

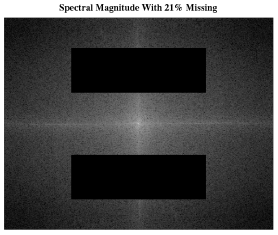

(b) 21% Block Size

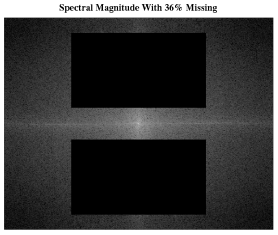

(c) 36% Block Size

Figure 5.3: Example optical spectrums with various block sizes removed.

5.2.1 Reconstruction Vs. Percent Missing

For the first experiment with the overhead optical images, we will remove an increasingly large contiguous block of frequency samples and record the reconstruction error and the correlation. The block of missing frequency samples is illustrated in Figure 5.3 for the cases of 7%, 21%, and 36% of the data missing. The centers of the blocks remain stationary while the lower and upper extents are increased in a step-wise manner. The percentage of missing samples ranges from 0% to slightly less than 50%. The purpose of this experiment is to see how TV minimization, which is what one will find in the current literature, performs relative to the other three filters. Moreover, we would like to see if there is a difference in performance between rural images composed of natural targets and urban images where one would expect sharper edges from the man-made objects.

In Figures 5.4 through 5.9, the results of the above are displayed. It is clear that

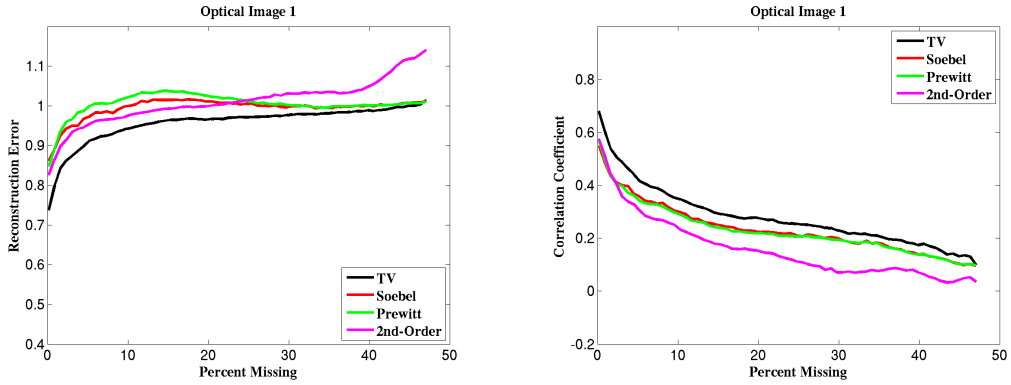

(a) Reconstruction Error

(b) Correlation Coefficient

Figure 5.4: Reconstruction vs. percent missing experimental results for optical test image 1 (rural).

TV outperforms the other three filters in all six cases for both metrics. It was thought that the Soebel and Prewitt filters might outperform the TV filters in the rural case, but this is not supported by the data. There are also two distinct shapes for the reconstruction error curves according to the scene type (i.e., rural or urban). The rural images all have reconstruction error curves that increase sharply and then level out quickly at a high error value. This tells us that none of these filters performs particularly well for this experiment, though, as already mentioned, TV is the best suited of the four. For the urban images we see that the error curves eventually reach the same error level but take more than twice as long to get there, which suggests that one can expect much better performance out of the urban-type images, at least up to a certain point.

We note that the 2^{nd} derivative estimate appears to have failed to converge once for the 5^{th} image and twice for the 6^{th}. Theses failures, however, do not impact the behavior visible in the plots.

5.2.2 Reconstruction Vs. Noise Level

The second optical experiment will fix the missing frequency samples to be a contiguous block that accounts for 12.8% of the frequency samples. We will then add

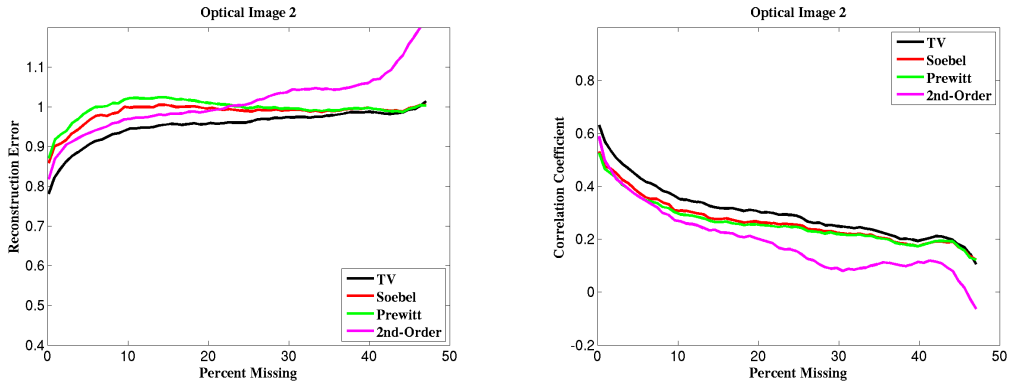

(a) Reconstruction Error (b) Correlation Coefficient

Figure 5.5: Reconstruction vs. percent missing experimental results for optical test image 2 (rural).

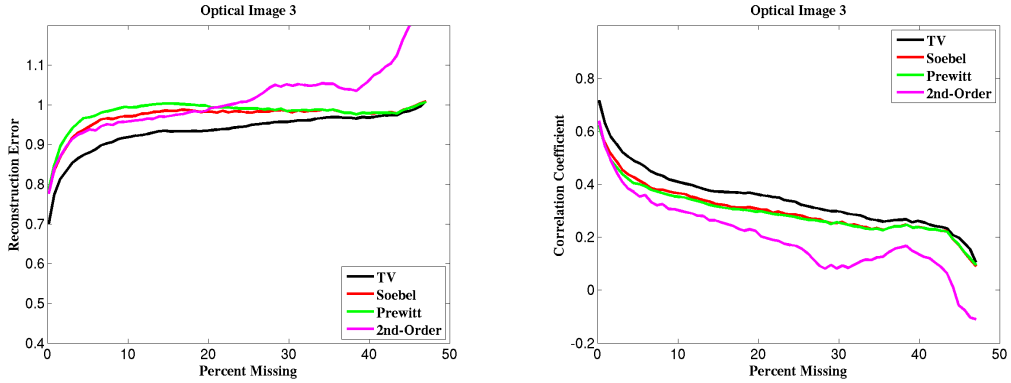

(a) Reconstruction Error (b) Correlation Coefficient

Figure 5.6: Reconstruction vs. percent missing experimental results for optical test image 3 (rural).

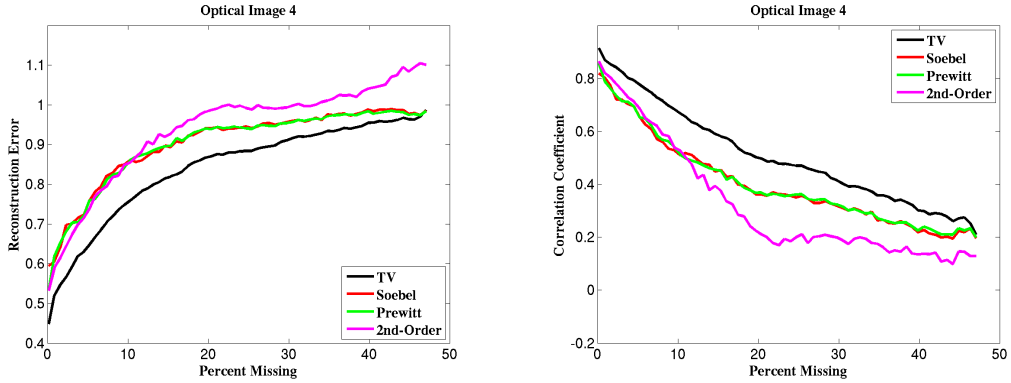

(a) Reconstruction Error (b) Correlation Coefficient

Figure 5.7: Reconstruction vs. percent missing experimental results for optical test image 4 (urban).

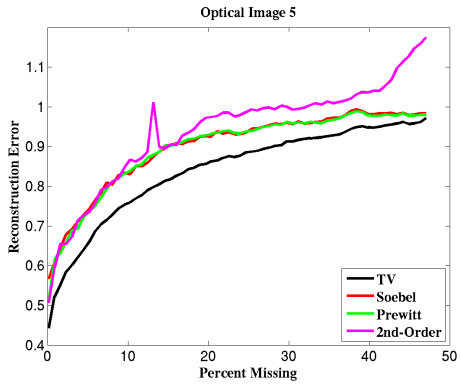

(a) Reconstruction Error

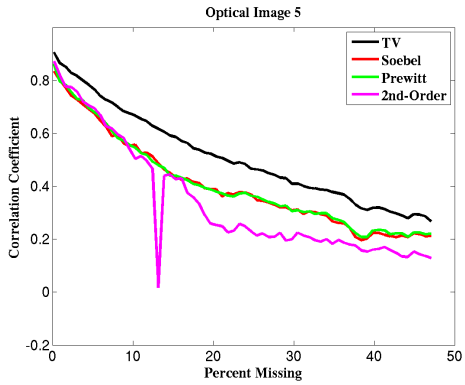

(b) Correlation Coefficient

Figure 5.8: Reconstruction vs. percent missing experimental results for optical test image 5 (urban).

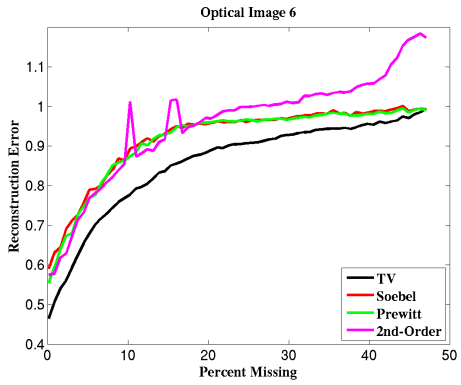

(a) Reconstruction Error

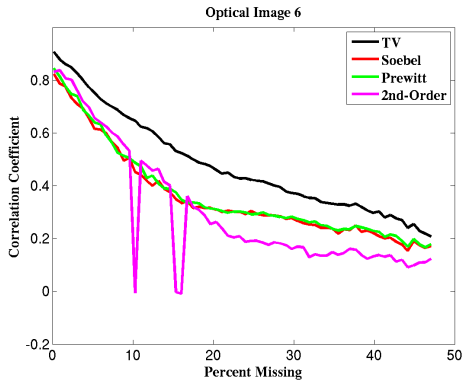

(b) Correlation Coefficient

Figure 5.9: Reconstruction vs. percent missing experimental results for optical test image 6 (urban).

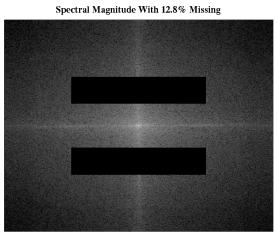

(a) Starting Incomplete Spectrum

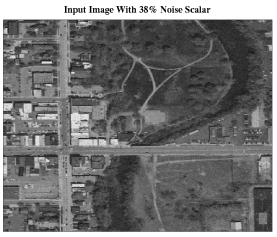

(b) 38% Noise Scalar

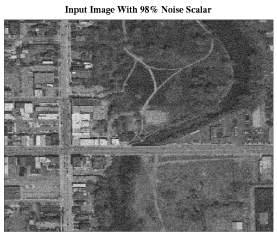

(c) 95% Noise Scalar

Figure 5.10: The incomplete spectrum showing the missing frequency samples for this experiment along with two example noise levels.

an increasing amount of uniformly distributed random noise to the image and record the performance metrics for the same four filters. The scalar used to control the noise level will vary from 0% of the image mean to 100%. The missing frequency samples and two example noise levels are illustrated in Figure 5.10. The motivation for this experiment is that both the Soebel and Prewitt filters are known to give better gradient estimates in the presence of noise [10] than TV, thereby leading to the supposition that they should outperform TV in the reconstruction of noisy data.

The results of this experiment are shown in Figures 5.11 through 5.16. As predicted, both the Soebel and Prewitt filters outperform the TV implementation; however, it is not until the noise scalar reaches 20% in each case that they do so. Moreover, we again see two distinct shapes to the curves according to whether the image was rural or urban.

5.3 Radar

For the experiments with the SAR images we will use the ℓ_1, TV, and 2^{nd} derivative filters. In theory, the ideal error-free SAR image has uniform amplitude in the frequency domain, so we expect the ℓ_1 reconstruction to perform best in such a situation since it should lead to the sparsest representation. We base this expectation

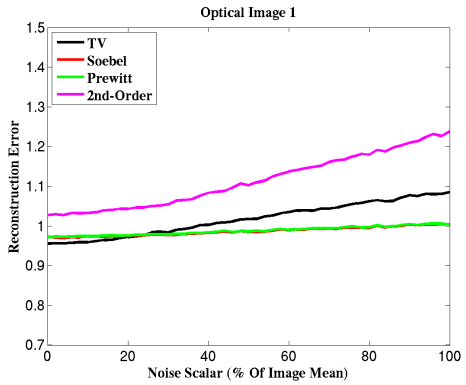

(a) Reconstruction Error

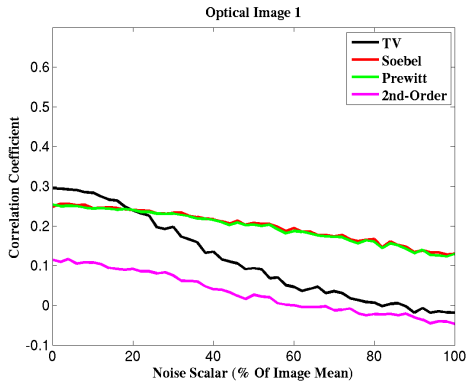

(b) Correlation Coefficient

Figure 5.11: Reconstruction vs. noise level experimental results for optical test image 1 (rural).

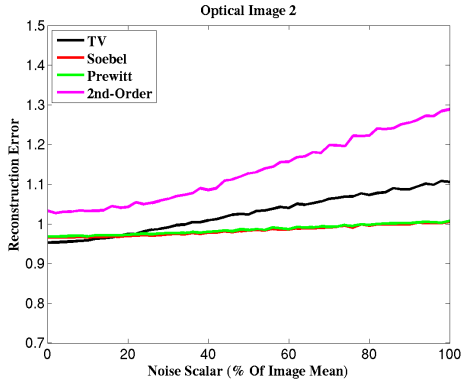

(a) Reconstruction Error

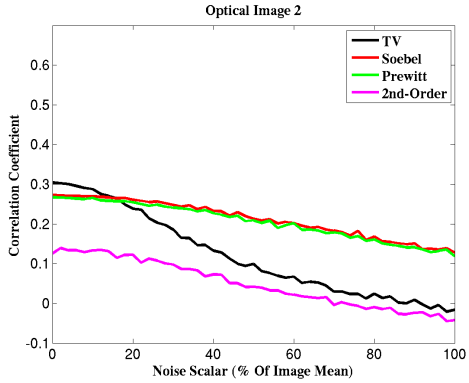

(b) Correlation Coefficient

Figure 5.12: Reconstruction vs. noise level experimental results for optical test image 2 (rural).

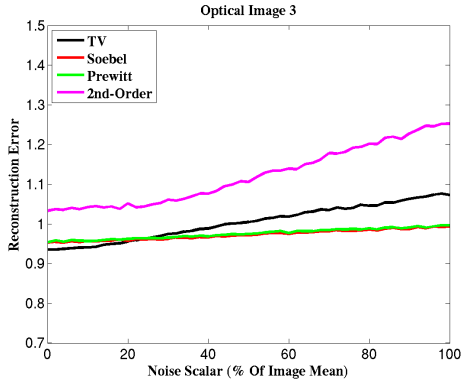

(a) Reconstruction Error

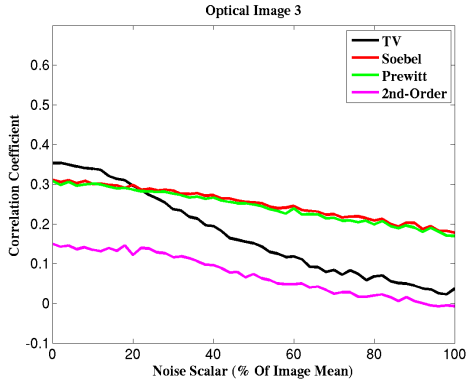

(b) Correlation Coefficient

Figure 5.13: Reconstruction vs. noise level experimental results for optical test image 3 (rural).

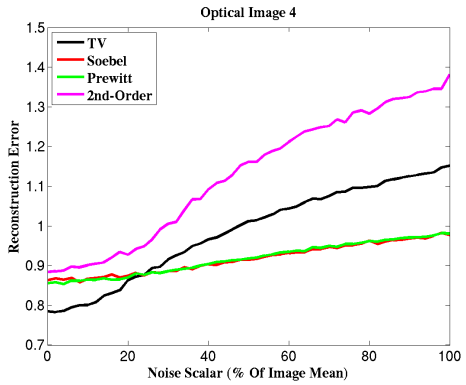

(a) Reconstruction Error

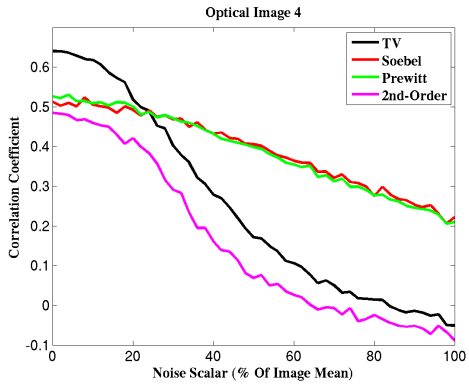

(b) Correlation Coefficient

Figure 5.14: Reconstruction vs. noise level experimental results for optical test image 4 (urban).

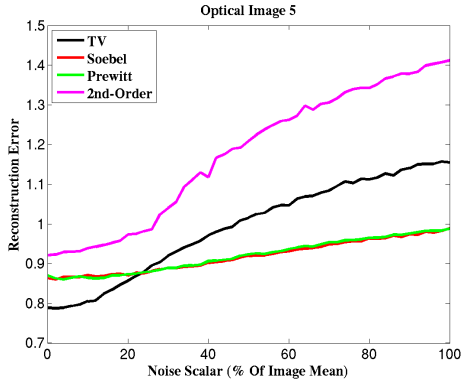

(a) Reconstruction Error

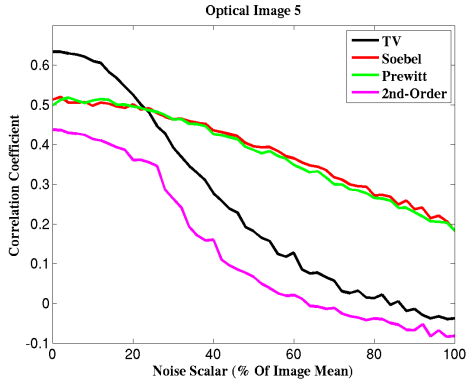

(b) Correlation Coefficient

Figure 5.15: Reconstruction vs. noise level experimental results for optical test image 5 (urban).

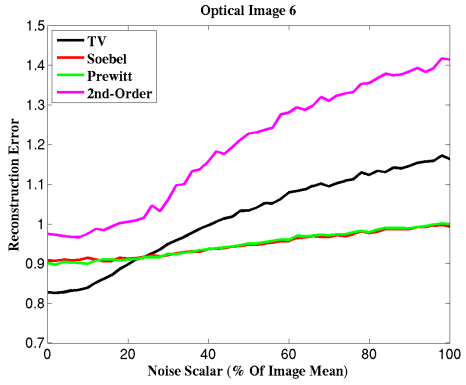

(a) Reconstruction Error

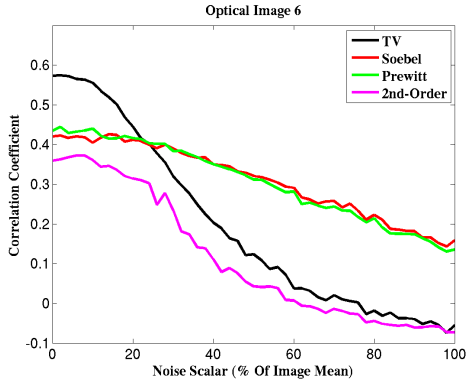

(b) Correlation Coefficient

Figure 5.16: Reconstruction vs. noise level experimental results for optical test image 6 (urban).

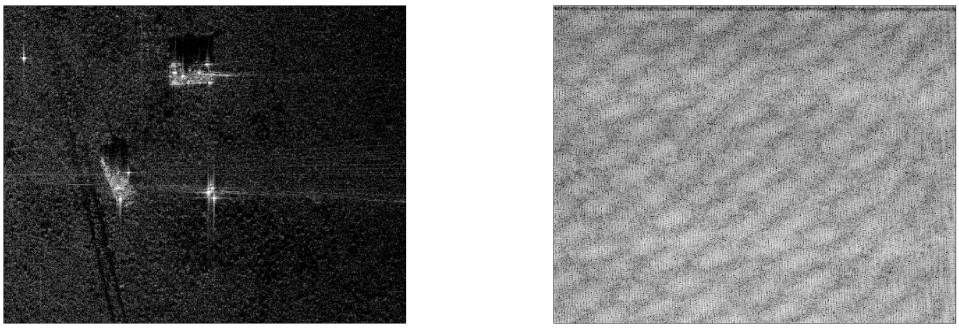

(a) SAR Image 1 (b) SAR Image 1 Spectrum

Figure 5.17: The first SAR test image and the magnitude of its spectrum.

on the conjecture that a sinc function, which is what an error-free SAR spectrum leads to for each target in the image, is the sparsest possible image domain function for a finite length signal. It might be possible to prove this using the Gini coefficient [12], but we do not pursue that here.

The images used herein are taken from the publicly available images provided by Sandia National Laboratories' MiniSAR system [13]. The system operates in the Ku-band and provides resolution on the order of 0.10(m) in both range and azimuth.

5.3.1 Reconstruction Vs. Percent Missing - Contiguous

Similar to the first optical experiment above, we will remove an increasingly large contiguous block of frequency samples from two SAR images and record the reconstruction metrics for each. In Figures 5.17 and 5.18 the two test images and the magnitudes of their spectrums are shown. In Figure 5.19 three example spectrums for the first image with various frequency block sizes removed are displayed.

The reconstruction metrics for this experiment are shown in Figures 5.20 and 5.21. The surprising result is that both TV and the 2^{nd} derivative implementations outperform ℓ_1 for both metrics on both images. Moreover, TV appears anywhere from

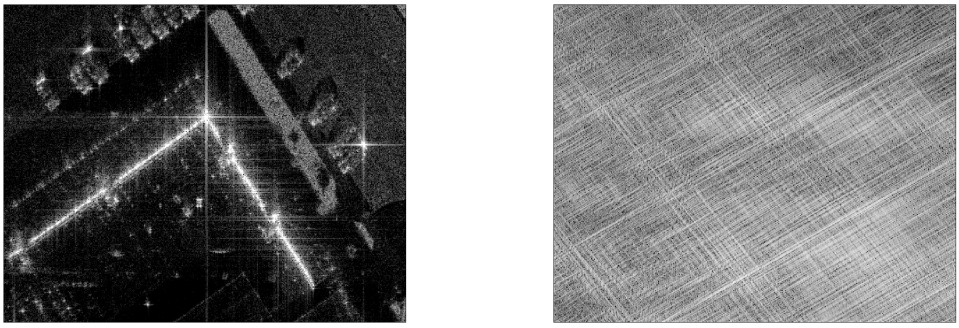

(a) SAR Image 2 (b) SAR Image 2 Spectrum

Figure 5.18: The second SAR test image and the magnitude of its spectrum.

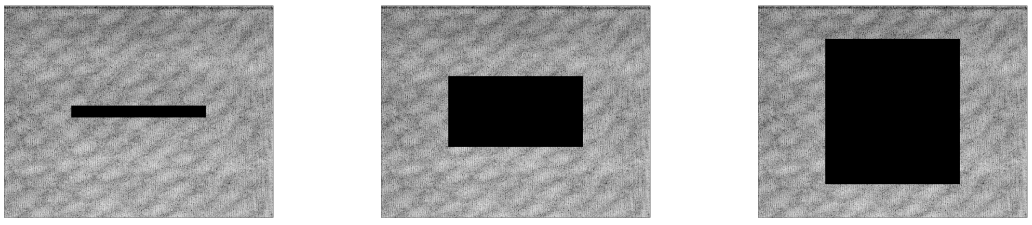

(a) 3% Block Size (b) 16% Block Size (c) 34% Block Size

Figure 5.19: Example SAR spectrums with various block sizes removed.

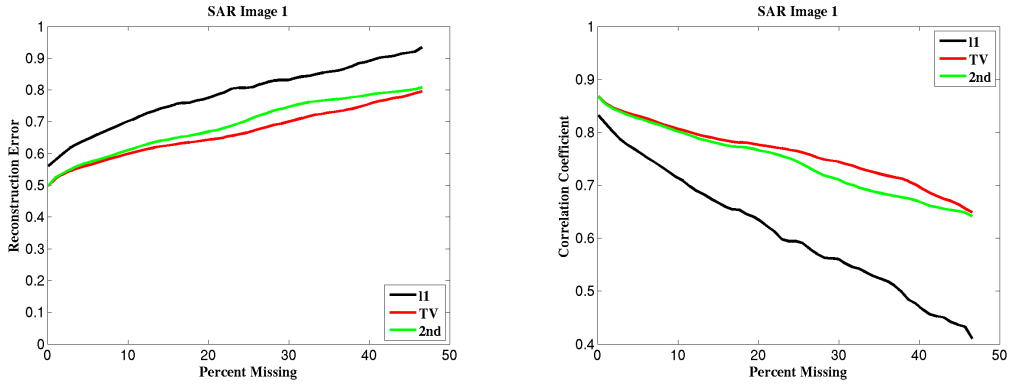

(a) Reconstruction Error

(b) Correlation Coefficient

Figure 5.20: The reconstruction metrics for the first SAR image for the first SAR experiment.

slightly to consistently better than the 2nd derivative. To get an idea of the different results, we show the incomplete spectrum for the case of 20% of the frequency samples missing and the reconstructed spectrums for all three implementations in Figures 5.22 through 5.25. Based on these spectrums, it appears that the ℓ_1 implementation is estimating a denser spectrum than the other two, which appear to be mostly reconstructing higher-energy patterns only. The manifestation of this difference in the image domain is suspected to be a reduction in the noise and/or low-energy returns. We try to measure this using the area outlined in red in Figure 5.26, where we also plot the average power in this area for each reconstructed image. Ignoring the downward spikes in the TV and 2nd derivative curves (which represent cases where the minimization hit the iteration limit), we see that as the percentage of missing samples increases, the TV and 2nd derivative implementations leave out more and more of the energy from this low-return area.

5.3.2 Reconstruction Vs. Percent Missing - Random

Since the amplitude of SAR spectrums tends to be fairly flat, we can expect the image domain contribution of each frequency sample to be roughly equal. This is in contrast to optical imagery, where the vast majority of the frequency domain

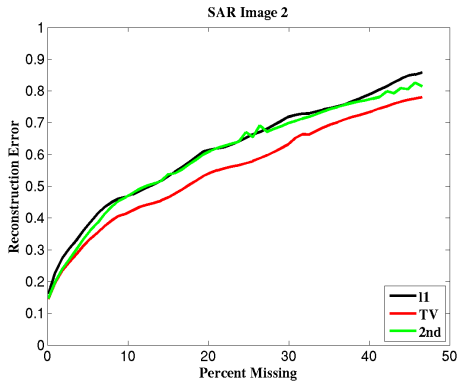

(a) Reconstruction Error

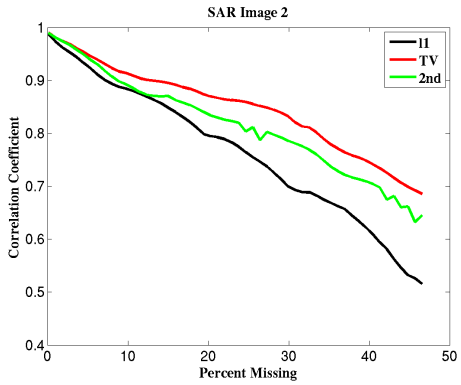

(b) Correlation Coefficient

Figure 5.21: The reconstruction metrics for the second SAR image for the first SAR experiment.

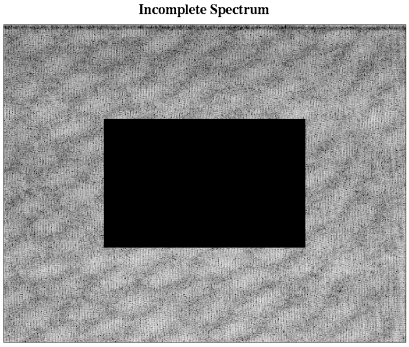

(a) Incomplete Spectrum

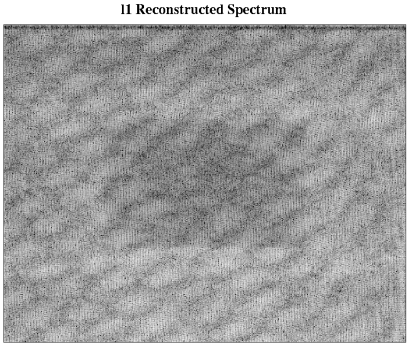

(b) ℓ_1 Reconstructed Spectrum

Figure 5.22: The incomplete and ℓ_1 reconstructed spectrums for the first SAR image with 20% of the frequency data removed.

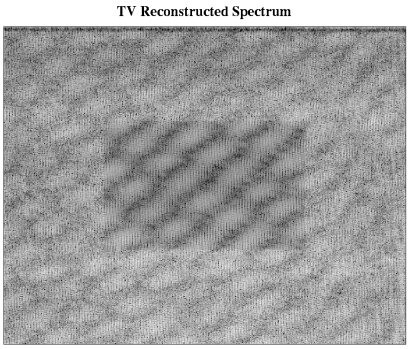

(a) TV Reconstructed Spectrum

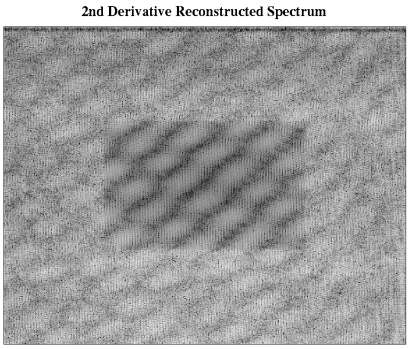

(b) 2^{nd} Deriv. Reconstructed Spectrum

Figure 5.23: The TV and 2^{nd} derivative reconstructed spectrums for the first SAR image with 20% of the frequency data removed.

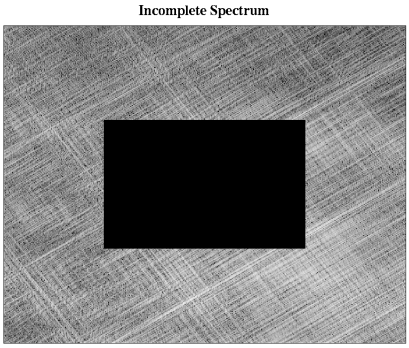

(a) Incomplete Spectrum

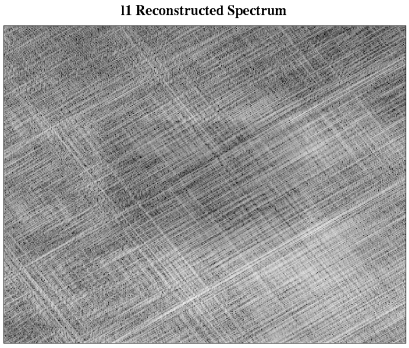

(b) ℓ_1 Reconstructed Spectrum

Figure 5.24: The incomplete and ℓ_1 reconstructed spectrums for the second SAR image with 20% of the frequency data removed.

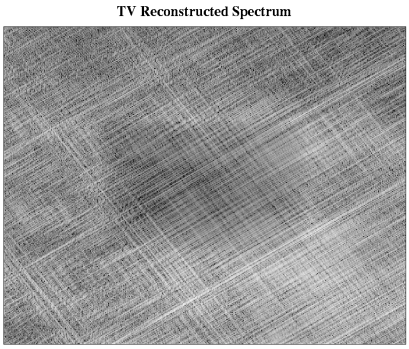

(a) TV Reconstructed Spectrum

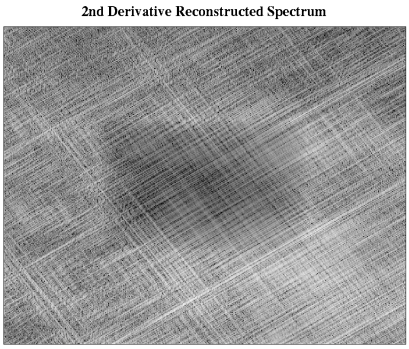

(b) 2^{nd} Deriv. Reconstructed Spectrum

Figure 5.25: The TV and 2^{nd} derivative reconstructed spectrums for the second SAR image with 20% of the frequency data removed.

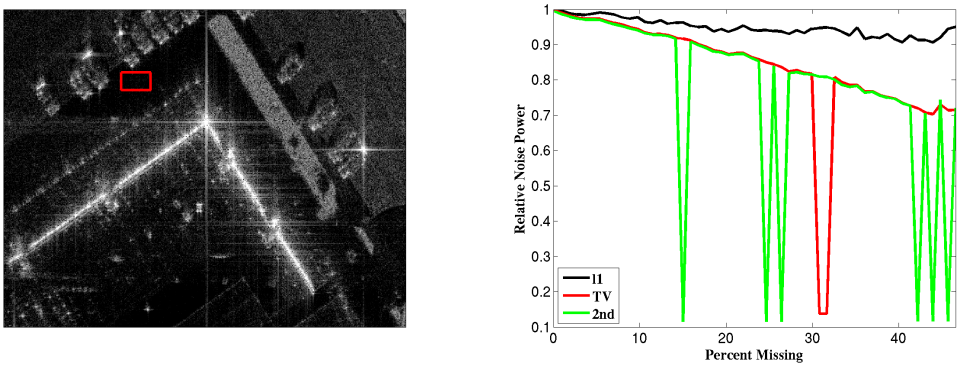

(a) Low-Return Area

(b) Estimate Noise Power

Figure 5.26: Outline of the area used to estimate the noise power and the relative noise power as a function of the amount of missing frequency samples. The downward spikes in the noise power are cases where the minimizer hit the iteration limit.

energy is concentrated in relatively few frequency samples. This characteristic of SAR imagery means that we can expect consistent behavior based not on which samples are removed but on how many. Also, since we are selecting the samples at random, we can evaluate the reconstruction error on a sample-by-sample basis and see, if as supposed in Section 1.4, there is relation between the gain of the filters and the accuracy of the reconstruction – the conjecture being that the lower the gain, the less accurate the estimate. In order to gauge the error for each reconstructed frequency sample individually, we use the same normalized distance as in (5.1) except that we do it for each individual sample, that is, for all $x_m \in \mathbf{X}_m$, $x_r \in \mathbf{X}_r$, we compute $\|x_m - x_r\|_2/\|x_m\|_2$. We will then group these measurements according to the combined gain of the filters used in the reconstruction at that particular frequency. This allows us to then compute the average error as a function of combined filter gain. We define the combined filter gain in a manner similar to how the elements of the polyfiltered image are combined to calculate the ℓ_1 norm: it is the 2-norm of the vector formed by the gain of the individual filters at any one particular frequency. For instance, let $g_1(m,n),\ldots,g_L(m,n)$ be the gains of the L filters at frequency bin (m,n); then we take the combined gain of the filters to be

$\sqrt{g_1^2(m,n) + \cdots + g_L^2(m,n)}$.

We begin illustrating the above in Figure 5.27, where we show the incomplete spectrums for the cases of 8%, 40%, and 80% of the frequency samples being chosen from a uniform random distribution. The reconstruction metrics are shown in Figure 5.27, where the TV implementation again consistently outperforms ℓ_1. (We remark here that we were unable to include the 2nd derivative implementation in this experiment due to both computational and time constraints. [1]) To see if the reconstruction behavior is consistent with the contiguous missing frequency case of the first SAR experiment, we display the incomplete spectrum for the 60% case along with the original for reference in Figure 5.29. The ℓ_1 and TV reconstructed spectrums are then shown in Figure 5.30. We see again that the TV reconstruction appears to give a less noisy spectrum than the ℓ_1. To see if this is born out in the image domain, we show the original image again for reference, the incomplete image, and the ℓ_1 and TV reconstructed images in Figures 5.31 and 5.32. There is a noticeable decrease in the noise level for the TV image as compared to the ℓ_1 image. The last topic we look at is the accuracy of the reconstruction as a function of filter gain. Since the ℓ_1 filter is an all-pass filter (which means it passes all frequency components unchanged), we only look at this for the TV case. In Figure 5.33 we show both the combined filter gain for TV along with the best-fit line [2] of the average error as a function of gain for the same 60% case as in the other images. The expected increase in error as the gain of the filters approaches zero is visibly present.

[1] The 2nd derivative implementation had to be stopped due to the amount of time it was taking to find a solution. It is suspected that a better choice of algorithm parameters could fix this, but in the mean time it had to abandoned in order to allow the ℓ_1 and TV implementations to finish.

[2] The line parameterization used in the fit is $a_0 + a_1 e^{-a_2 g}$, where g is the gain and the a_i are chosen to minimize the l_2 error of the line fit.

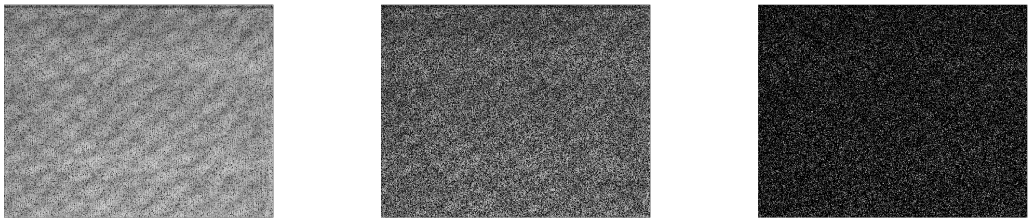

(a) 8% Missing (b) 40% Missing (c) 80% Missing

Figure 5.27: Incomplete spectrums of the first SAR image with various amounts frequency samples removed.

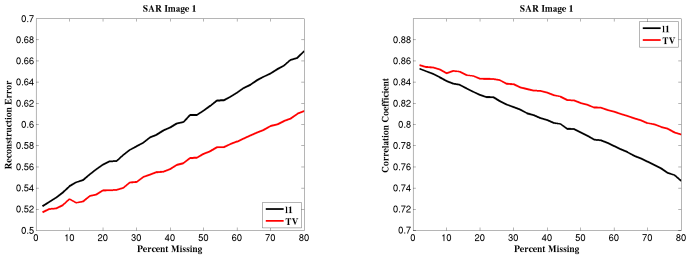

(a) Reconstruction Error (b) Correlation Coefficient

Figure 5.28: The reconstruction metrics of the first SAR image for the second SAR experiment.

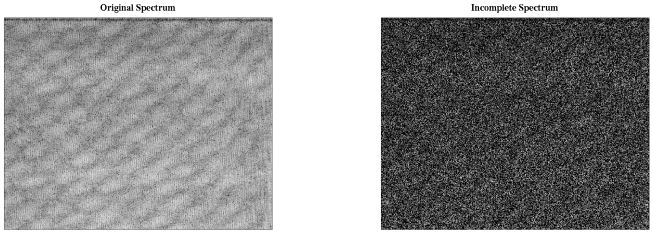

(a) Original Spectrum (b) 60% Missing

Figure 5.29: The original spectrum and the incomplete version of the 60% case for the second SAR experiment.

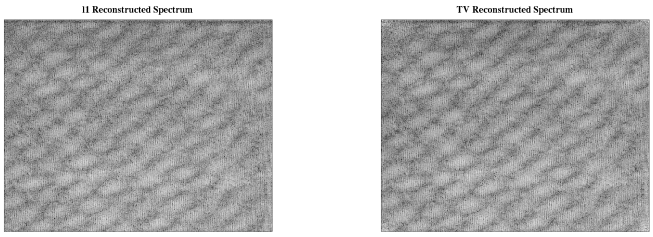

(a) ℓ_1 Reconstructed Spectrum (b) TV Reconstructed Spectrum

Figure 5.30: The ℓ_1 and the TV reconstructed spectrums of the 60% case. for the second SAR experiment.

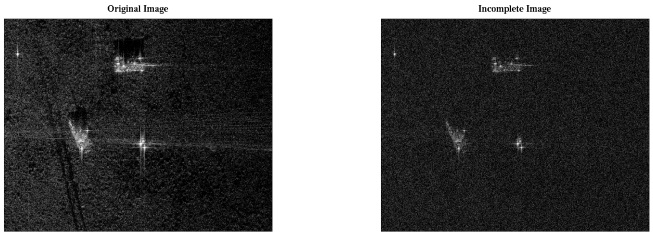

(a) Original Image (b) Incomplete Image

Figure 5.31: The original image and the incomplete image of the 60% case for the second SAR experiment.

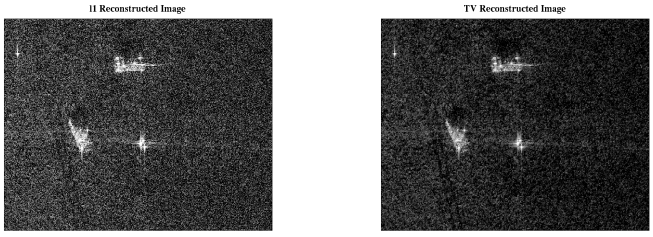

(a) ℓ_1 Image (b) TV Image

Figure 5.32: The ℓ_1 and TV reconstructed images of the 60% case for the second SAR experiment.

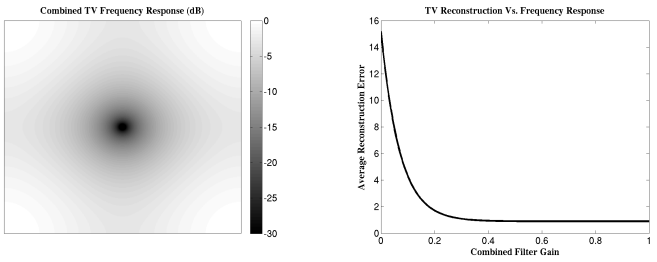

(a) Combined TV Filter Response

(b) Error Vs. Gain

Figure 5.33: The combined filter gain for TV along with the estimated average error as a function of gain.

CHAPTER VI

Conclusion And Future Work

6.1 Conclusion

The main conclusion of the preceding work is that the interpretation of TV minimization as a filtering operation and its generalization to allow for arbitrary filters allows the problem to be developed more robustly from a signal processing point of view. Treatments of TV minimization such as [14] develop the problem from a functional analysis standpoint which leads to proofs that are both prolonged and unfamiliar to many. Moreover, the sole focus is on the first-order difference operation and this leads to other rigid considerations such as the divergence of an image. To stay with this last example, in the PV development, where the operator is a general convolution, what is defined as the divergence in [14] is no more than the adjoint of the filtering operation. Moreover, given the implementation of polyfiltering as a single linear transform, this is available simply as $\mathbf{P}^{\mathrm{H}}$, regardless of the underlying filter.

As important as the preceding point about the generality of the theory, is that the PV development lets one use the same software for a whole class of problems. Software packages such as [7] have separate solvers for both ℓ_1 and TV, where here we have used the same solver for the ℓ_1, TV, Prewitt, Soebel, 2^{nd} derivative, and even two specialty implementations. We will allow that specific problems sometimes re-

quire specific software, however, for general purpose and research-oriented software, the end user should not be limited to a fixed either/or choice.

Experimental conclusions include that the expected result that Soebel and Prewitt filters can improve performance with optical imagery with noise beyond a certain level. We also saw that TV performed better than both of these alternatives for both rural and urban, the former being unexpected.

The most surprising experimental conclusion is that both TV and the 2^{nd} derivative implementations outperform ℓ_1 with SAR imagery. This may prove to be a significant improvement and is something, to the best of our knowledge, that has not been looked at before.

6.2 Future Work

One of the immediate possibilities for future work is to extend the PV implementation to the denoising case,

$$\min_{\mathbf{x}} \mathrm{PV}(\mathbf{x}) \;\text{ s.t. }\; \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \sigma, \tag{6.1}$$

and see if the preceding results among the various implementations hold. For instance, the denoising problem lets us add noise directly to the frequency measurements. This is in contrast to what was done here, where the noise was added to the image prior to the frequency sampling. It is expected that the Soebel and Prewitt filters would still give better performance but at this point that is conjecture. Also, if that is the case, it would be interesting to see if the amount of improvement is consistent.

With respect to the reconstruction of SAR images, we saw that TV gives estimates with significantly less noise, so it would be interesting to see if the denoising formulation could give even better results. Also, the noise level is only one measure of the

quality of SAR images. To give a complete comparison of how ℓ_1 and TV compare, one would need to examine calibration targets and compare things such as the Peak Sidelobe Ratio (PSLR) and the Integrated Sidelobe Ratio (ISLR) among others [15]. There are also outstanding questions as to *why* TV outperforms ℓ_1 in regards to the SAR imagery and, moreover, if there are other filters that may yield even better results (though the answer to the former would presumably suggest answers to the latter).

APPENDICES

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, pp. 489–509, Feb. 2006.

[2] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, pp. 129–159, Jan. 2001.

[3] Y. Zhang, "On theory of compressive sensing via ℓ_1-minimization: Simple derivations and extensions," *Rice CAAM Department Technical Report TR08-011*, pp. 1–21, 2008.

[4] E. Candes, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on pure and applied mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.

[5] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.

[6] C. Li, "An efficient algorithm for total variation regularization with applications to the single pixel camera and compressive sensing," Master's thesis, Rice University, Houston, Texas, 2009.

[7] E. Candes, "L1-Magic: Recovery of Sparse Signals," *http://www.acm.caltech.edu/l1magic/*, 2005.

[8] J. F. Sturm, "Using sedumi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11–12, pp. 625–653, 1999.

[9] CVX Research, Inc., "CVX: Matlab software for disciplined convex programming, version 2.0 beta." `http://cvxr.com/cvx`, Sept. 2012.

[10] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2nd ed., 2001.

[11] ESRI, "ArcGIS Web Map Server." `http://services.arcgisonline.com/ArcGIS/rest/services/World_Imagery/MapServer/`, 2012. [Online; accessed October 2012].

[12] D. Zonoobi, A. A. Kassim, and Y. V. Venkatesh, "Gini index as sparsity measure for signal reconstruction from compressive samples," *J. Sel. Topics Signal Processing*, vol. 5, no. 5, pp. 927–932, 2011.

[13] Sandia National Laboratories, "Minisar: Miniaturized synthetic aperture radar." `http://www.sandia.gov/radar/minisar.html`, 2012.

[14] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imaging Vis.*, vol. 20, pp. 89–97, Jan. 2004.

[15] W. Carrara, R. Goodman, and R. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*. The Artech House Remote Sensing Library, Artech House, 1995.