Improving Sonar Reconstruction with Sparse Reconstruction Techniques

by

Brian Lindmark

A dissertation submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Auburn, Alabama
December 10, 2016

Keywords: Sparsity, Blind Deconvolution, Sonar

Approved by

Stanley Reeves, Chair, Professor of Electrical and Computer Engineering
Thomas Denney, Professor of Electrical and Computer Engineering
Lloyd Riggs, Professor of Electrical and Computer Engineering
Richard Zalik, Professor of Mathematics and Statistics

Abstract

In many sonar applications, traditional processing can fail to provide adequate information to accomplish the respective goals, because the processed information is either strictly bandlimited to the transmit band, distorted (matched filtering places higher emphasis on larger frequency components), or subject to large error due to noise. By using sparse reconstruction techniques on systems using bandpass chirp transmit signals, we can extrapolate additional bandwidth beyond the transmit signal bandwidth without directly amplifying the noise. This approach could help significantly in areas where more than locational information is desired, such as in target classification. We approach this problem with a basic one-dimensional sparse reconstruction algorithm, extend this to blind deconvolution, and then extend this again by addressing a spatially-varying reconstruction problem.

Matched filtering is used successfully as a way to find targets in sonar applications, but fails to retain details of the targets. The output of matched filtering is strictly bandlimited, limited to the transmit band. The frequency response also prioritizes stronger frequency components, so the output is distorted even in the ideal case. Targets of interest have finite extent and are therefore not bandlimited. In order to get a more accurate representation of the target, it is sometimes preferable to use a regularized deconvolution method. This has a flatter response across the signal bandwidth and has a wider recovery bandwidth, which can lead to improved resolution. Regularized restoration also suffers from a high sensitivity to noise, particularly outside of the strong portions of the transmit band. We propose using a more robust, ℓ_1-based deconvolution method that exploits the finite extent and relative sparsity of targets of interest in order to obtain a more accurate reconstruction. We demonstrate the effectiveness of this technique both on simulated data and on tank data.

In many sonar applications, either an analytic description of the transmitted acoustic signal or a recorded transmit replica is used for processing the return data. Unfortunately, there are multiple processes that occur from the analytic design until receiving the return that could severely degrade its performance in processing. We propose a method to simultaneously estimate the target and the actual in-water transmit signal. The improved estimate on the target can be used directly, while the improved estimate on the in-water transmit can be used as an alternative to the analytic description or recorded replica in additional processing while maintaining a reasonable processing speed for active use.

An acoustic signal changes as it travels through a medium. A simple example of this can be found in attenuation due to spherical or cylindrical spreading. This causes the received signal to have less strength than the one transmitted. The observed changes can be much more complicated, however. The signal degradation can be frequency dependent, such as in a dispersive medium or in transmission loss from absorption. In general, the degradation will be more severe as the distance from the source increases, since these effects compound with each other. Scene estimates obtained from deconvolution-based processing can suffer greatly due to errors in the transmit signal used in processing. We propose a method to adaptively estimate any transmit signal at multiple distances from the source using a semi-blind deconvolution estimation method. This will allow for more accurate processing at longer distances than would otherwise be possible with a single transmit description.

Acknowledgments

I'd like to thank my advisor, Dr. Stanley Reeves for his assistance throughout this process. I'd also like to thank my Master's advisor, Dr. Richard Zalik for encouraging me to pursue a Ph.D. in electrical engineering. Thank you to Dr. Jermaine Kennedy and David Malphurs from the Naval Surface Warfare Center (NSWC) at Panama City Beach, FL, who collected the experimental data and assisted greatly in the interpretation of the results based on target physics. I'd also like to thank Dr. Myoung An and Prometheus, Inc. for introducing me to the sonar work that led me down this path. Finally, I'd like to thank my wife Kirsten for her constant love and support throughout this process.

Table of Contents

List of Figures

Chapter 1

Sonar Background

The typical application of sonar is to determine where things are in the water using an acoustic waveform. There is also an increasing interest in determining not just where things are but what they are. In order to extract that kind of additional information from a sonar return, it is necessary to retain some amount of the structural information of the target in the return waveform. All transmitted waveforms are inherently bandlimited, however, so the return signal necessarily loses some of the frequency content of the interrogated scene, limited to the transmit band. This is because sonar systems are a convolution between a transmit signal and the interrogated scene. When the transmit signal is bandlimited, then the return will necessarily be bandlimited as well. This can be particularly problematic when stealth is desired, as the transmits tend to be narrow in order to minimize the emitted power. In addition to being bandlimited, sonar is also subject to a variety of noise components which can alter both the transmit and return signals away from the ideal case. Electrical component noise, medium inhomogeneities, Doppler frequency shifts, multipath returns, and resonance effects are just some of the noise components that contribute to inaccurate estimates of the underlying desired scene information. One way to mitigate this problem is to increase the bandwidth of the transmit signal. This would partially solve the bandlimited reconstruction case and give increased resolution in the reconstruction, but it also presents its own complications. Increasing the bandwidth requires using transducers that can stably transmit a large range of frequencies. This can be expensive as well as undesirable in some cases, such as the aforementioned stealth. Rather than changing the hardware used, we seek to make processing changes to make it possible to extract additional, accurate information about the interrogated scene.

Discussing a sonar return can be difficult due to the change in coordinate systems involved in the acquisition procedure. In the simplest, single monostatic sensor case, an acoustic waveform is transmitted into three-dimensional space while the recorded return is collected over a single time dimension. Figure 1.1 shows a diagram of three different paths that contribute to the return information of the experimental data. Without knowledge of the setup or additional sensors, we cannot distinguish between a return signal that bounces off of multiple surfaces and a return that comes from a target that has the same path length from a different location. The process of a single sonar return having multiple paths on the sensor recording is known as multipath, and is a source of ongoing study in sonar. Reconstructing this information accurately in multiple dimensions requires either multiple sensors or multiple acquisitions over time (assuming the target information is stationary). Thus, we will refer to the signal in reference to the return information dimension. Thus, a one-dimensional signal is any which is obtained from a single stationary sensor, since the data is obtained as a one-dimensional temporal return. We will refer to the signal as two-dimensional only if the sensor is moved during data acquisition or a sensor array is used to collect the data, as in the case of synthetic aperture sonar collections. The experimental data in this document was obtained from a single stationary sensor, and will therefore be referred to as one-dimensional.

A sonar return is dependent on the geometry of the acquisition and target as well as the composition of the water and the target(s). The geometry of the acquisition procedure includes measurements such as height of the transmitter/receiver relative to the surface of the water, the ground, and the target. This can also include whether the collection is monostatic (transmitter and receiver are co-located) or bistatic (transmitter and receiver are not co-located). This information can determine where multipath returns will appear, as well as where additional features like ground and surface bounces will appear in the return. The experimental data in this document uses a setup that is bistatic but is set up to be processed as a monostatic system by placing the transmitter and receiver near eachother.

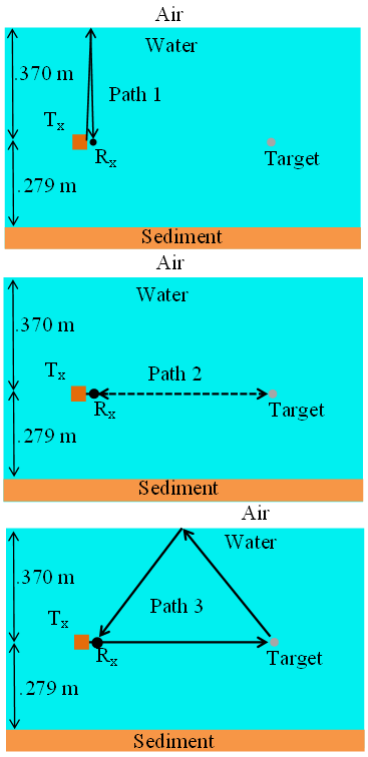

Figure 1.1: Diagram of Some Possible Reflection Paths

The composition of the water determines sound speed, possible presence of resonance effects (bubbles), relaxation and absorption effects, and the possible travel method of the acoustic waveform. In [1], Urick states that absorption of sound in water is caused by shear viscosity, volume viscosity, and ionic relaxation of magnesium sulfate in seawater. The absorption coefficient due to viscosity is directly proportional to the square of the transmitted frequency, while ionic relaxation of magnesium sulfate is the dominant form of absorption below 100kHz transmit frequency. The speed of sound in water is dependent on temperature, salinity, and depth, which are all properties of the medium as well.

The material and geometry of the target determines the target strength and elastic effects of the return. Classically, processing for sonar was performed to obtain the rough measure of the reflection coefficient of the material involved. Objects with a high reflectivity will show up in the processed return and tell you where these objects are. In reality, the signal interacts with objects much more than a simple reflection accounts for. A typical

sonar return can include the incident reflection from the object, lamb-like waves that travel through the front surface of the target and reflect from the back surface, as well as whispering gallery and Rayleigh waves that can couple onto the surface of the target and travel around and contribute to the return signal [2]. The geometry and material of the interrogated target both determine which of these waves can occur as well as where these waves will appear in the processing. Different materials have different sound speeds associated with them, which can change the observed path length (processing the signal assumes a constant sound speed). The strength of these waveforms is also impacted by the type of material and how rigid it is.

1.1 Experimental Data

The experimental data in this document was obtained from a small-scale test bed (SSTB) at the Naval Surface Warfare Center in Panama City Beach, Florida, where the water is assumed to be homogeneous. This facility uses high-frequency transmit signals along with small targets to simulate realistic conditions that correspond in a scaled manner to lower-frequency transmits and larger targets. The acquisition geometry is known, so we can identify the ground and surface bounces separately from the return from the test targets. We used sample targets of a 1" solid stainless steel spherical target and a 2:1, 1" diameter solid stainless steel cylinder in freefield. Figures 1.2 and 1.3 show pictures of the setups for both target types. More information on spherical and cylindrical target scattering effects can be found in [3, 2, 4, 5, 6]. Since the material and geometries are known in advance, we have a reasonable idea of how the target should appear in processing. We would expect to see an initial impulse (or few impulses) from the incident reflection on the front face of the target, followed by additional impulses from the longitudinal waves, whispering gallery waves, and Rayleigh waves. Our particular data collection also had an additional source of waves present due to the source being located nearby the omnidirectional receiver. We noted an additional spike in the data from the specular backscatter that reflects off of the source

and is picked up by the receiver. We specify these locations, along with a breakdown of the visible waves in this data collection in the individual results sections of later chapters.

Figure 1.2: Spherical Target

Figure 1.3: Cylindrical Target

Since we expect the return to be primarily composed of a series of separated impulses or a connected but spatially limited response, we can use sparse signal processing to greatly improve the quality of information received from this data. Sparsity means that, under some basis, we can represent the signal by a small number of nonzero elements. Since the signal is composed of separated impulses, we can represent it as a sparse collection of elements in its native basis. Sparse signal processing will allow us to extrapolate bandwidth outside of the

transmission band. This means that we can attain wider-band estimates from narrowband transmit signals. This can greatly improve the resolution of the system, as well as the overall quality of the reconstruction. These features make this application ideally suited for sparse signal processing.

Chapter 2

Sparse Reconstruction of One-Dimensional Returns

2.1 Introduction

In many sonar applications, target information is obtained by transmitting a narrowband chirp in the water. The return signal is then matched-filtered, and locational information for the major reflectors is obtained. This technique is quick, easy to implement, robust to noise, and effective for most applications. Unfortunately, significant target information is lost outside of the transmit band. This is because matched filtering imposes two bandpass filters on the scene. The first filter occurs when the transmit is convolved with the scene, then the second when the return signal is processed. Typical targets of interest are of finite extent and are not bandlimited, since they are spatially finite. Thus, in order to reconstruct them more accurately, another technique must be used.

Regularized restoration has the benefit of a wider, flatter recovery band. While information outside of the transmit band is still initially suppressed once during convolution with the transmit signal, the reconstruction filter attempts to recover a portion of it with a partial inversion that employs the regularization term. We incorporate information about the targets to keep the reconstruction restricted to solutions we know to be more accurate in our selection of the regularization filter. This means that the information we obtain will retain more recovered content than the matched filtered output. The processing requirement is also minimal, as it can be computed quickly using fast Fourier transforms (FFTs). Unfortunately, it also has significant sensitivity to noise, since the filter can recover information in frequency bands with relatively low transmit strength, which also amplifies the noise in those bands. In addition, the computational complexity is somewhat increased over matched

filtering, and additional user input is required to choose some parameters involved in the reconstruction equation.

Instead of choosing an ℓ_2 regularization term, we can instead reconstruct based on sparsity information of the target. Sparse reconstruction techniques have been used successfully in a variety of applications in recent years. Overviews of some of these methods can be found for image processing [7], radar [8], and a general overview in [9]. Some specific radar applications include through-the-wall radar imaging [10], multi-target detection [11], and synthetic aperture radar imaging [12]. Sonar applications have also used some sparse reconstruction techniques, such as in 3D sonar imaging [13], and synthetic aperture imaging applications [14, 15, 16]. We seek to apply sparse reconstruction techniques to one-dimensional sonar returns collected from high-frequency bandpass chirp transmit signals. Typical sonar targets have finite extent, and the majority of most scanning areas are primarily empty water. This makes it ideally suited for a sparse basis reconstruction. We can take a similar approach as the ℓ_2-based regularized restoration method- but now restrict our solutions to those that are sparse in our chosen basis. While this is similar in concept, the sparsity constraint is a powerful restriction that can significantly improve the noise sensitivity and spatial resolution of our results. This method induces a correlation among surrounding frequency values in the reconstruction, allowing us to reduce the noise variance in each estimated frequency component [17]. This allows us to extrapolate frequency information not contained in the transmit band, which can translate into increased resolution and more accurate target representations. To implement this method, we use a least absolute shrinkage and selection operator (LASSO). This operator has been studied extensively for applications with sparse representations such as in [18, 19]. Unfortunately, the computational complexity of this method increases significantly, and we can no longer solve this in closed-form. Instead, we use an iterative approximation method to find the solution.

This sparse reconstruction algorithm is similar to those used in compressed sensing, which has seen significant research in recent years. While the mathematics involved may be

the same, the goal is different. In compressed sensing, the objective is to improve acquisition speed while maintaining reconstruction quality by exploiting the sparsity of the signal. This is shown in [20, 21, 22], where they develop alternate sampling methods to reduce sampling times using a compressed sensing framework. While this technique has many potential benefits associated with it, it requires a change in hardware and/or acquisition procedure in order to satisfy incoherence between the sensing and sparse matrices. We seek to use the sparsity of the target to improve reconstruction quality without changing hardware or acquisition procedures, which comes at the expense of additional processing time. Although compressed sensing techniques have typically been applied to synthetic aperture data, they serve to further validate the use of sparse modeling in sonar processing.

Another approach to increase target information is to simply widen the bandwidth. This would allow both matched filtering and regularized restoration techniques to increase the recovered bandwidth and obtain more relevant target information. This is frequently restricted by hardware, however, and comes with some additional problems. A wider bandwidth requires transducers capable of outputting accurately over the entirety of the desired transmit bandwidth. This requires either replacing existing transducers with new ones that are able to transmit over the entire band effectively, or figuring out methods to effectively process the data while transmitting with a transducer outside of its optimal operating frequencies. In addition, transmitting a wideband signal also increases the amount of power put into the water. This can be problematic if stealth is important. Wideband signals can also make it more difficult to keep nearby vehicles from overlapping their transmissions, since band separation will be more difficult to implement.

2.2 Reconstruction Algorithm

We assume a high-frequency, discretely sampled sonar model

$$r = Xs + n \tag{2.1}$$

where X is the operator representing linear convolution with the transmit signal x, s is the interrogated scene, n is noise, and r is the recorded return. We seek to recover s from r.

To do this, we define the reconstruction as follows:

$$\hat{s} = \underset{s \in \mathbb{R}^n}{\arg\min} \{ ||Xs - r||_2^2 + \alpha ||As||_1 \} \tag{2.2}$$

This approach can be viewed as the computationally tractable convex relaxation of the underlying sparsifying ℓ_0-norm problem. The use of the ℓ_1 norm prevents us from finding a closed-form solution. Instead, we use a process of majorization as outlined in [23]. This process uses a quadratic of the same curvature at each iteration, which allows for an efficient computation of the minimum.

2.2.1 Majorization

Using majorizing functions, we reduce the problem to a sequence of quadratic problems. We use the method proposed in [23] to choose our majorizers such that the quadratic problem has a closed-form solution at each iteration. To do this, we first choose to use a modified Huber function to replace the ℓ_1 term in (3.2)

$$\rho_T(s) = \begin{cases} \frac{s^2}{2T} + \frac{T}{2}, & |s| \leq T \\ |s|, & |s| > T \end{cases} \tag{2.3}$$

When s is inside of our defined region T, $\rho_T(s)$ is quadratic. This function is differentiable everywhere and approaches $|s|$ as T approaches 0. As long as the majorized function touches the original function at the current estimate and is greater than or equal to the original function everywhere, convergence to the original minimizer is guaranteed [24]. We define this function as

$$\rho_T^M(s; s^k) = \frac{1}{2T} s^2 - d(s^k; T)s + \frac{1}{2T}(s^k)^2 \tag{2.4}$$

where

$$d(s;T) = \begin{cases} \frac{s}{T} + 1, & s < -T \\ 0, & -T \leq s \leq T \\ \frac{s}{T} - 1, & s > T \end{cases} \tag{2.5}$$

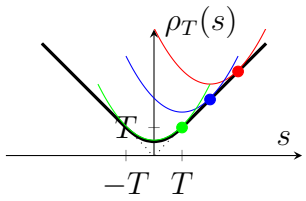

Figure 2.1: Majorized criterion for various operating points

The original function and various majorizing functions are plotted in Figure 2.1. This choice of majorizing function leads to the possibility of a very efficient iterative algorithm. To see this, consider the majorized criterion:

$$e_1^M(s;s^k) = \|r - Xs\|^2 + \alpha R(s;s^k) \tag{2.6}$$

where

$$R(s;s^k) = \sum_i \left[\frac{1}{2T}(As)_i^2 - d(As^k;T)(As)_i + C_i \right] \tag{2.7}$$

The constants C_i can be neglected, since they have no effect on the minimum.

Taking a derivative,

$$\frac{d}{ds} e_1^M(s;s^k) = -2X^H(r - Xs) + \frac{\alpha}{T} A^H As - \alpha d(As^k;T) = 0 \tag{2.8}$$

$$\hat{s} = (X^H X + \frac{\alpha}{2T} A^H A)^{-1} [X^H r + \frac{\alpha}{2} A^H d(As^k;T)] \tag{2.9}$$

If we treat this solution as one step of an iterative algorithm, we obtain:

$$s^{k+1} = (X^H X + \frac{\alpha}{2T} A^H A)^{-1} [X^H r + \frac{\alpha}{2} A^H d(As^k; T)] \tag{2.10}$$

Note that the matrix to be inverted in each iteration is invariant to the current estimate. X represents convolution with the transmit signal. If we also assume that A represents a convolution operator, then we can diagonalize both matrices using the unitary DFT matrix F. So $X = F^H \Sigma_X F$, and $A = F^H \Sigma_A F$. Then

$$(X^H X + \frac{\alpha}{2T} A^H A)^{-1} = F^H (\Sigma_X^H \Sigma_X + \frac{\alpha}{2T} \Sigma_A^H \Sigma_A)^{-1} F \tag{2.11}$$

This matrix can be applied to a vector through fast transforms and pointwise scalings. This approach allows us to easily change both the parameters α and T without recalculating inverse matrices.

2.2.2 Method Summary

1. Find an initial estimate for the scene using the known transmit and return data. A quick scene estimate can be obtained by using the deconvolution formula given in (3.6). An initial scene estimate is not required, but using one can speed up the convergence of the algorithm.

2. Choose the sparsifying basis A for s. In our simulation and experimental data we use $A = I$, so the scene itself is assumed to be sparse.

3. Choose parameter α, which controls the trade-off between sparsity of the scene, and data consistency. We used $\alpha = 0.1$ for the experimental data.

4. Choose parameter T from (2.3). This controls the level of quadratic approximation to the ℓ_1 problem. Smaller T values are better approximations, but will take longer to run. We use $T = 10^{-4}$ with the experimental data.

5. Run the algorithm, which iteratively estimates the minimum of (3.2) using the majorizing functions defined in (2.4).

6. Obtain a new scene estimate as the output to the algorithm.

2.3 Results

2.3.1 Simulation Results

The simulations were generated using a 10-20kHz LFM chirp transmit sampled at 60kHz. We begin by examining the impulse responses for matched filtering, regularized restoration, and our sparse reconstruction algorithm in an ideal, noiseless case. We then examine the results of each technique after adding Gaussian noise to the return at SNR levels of 15 dB and 5 dB.

In the following examples, we use a regularized restoration model for deconvolution given by

$$\min_{s\in\mathbb{R}^n}\{||Xs-r||_2^2+\gamma||Bs||_2^2\} \tag{2.12}$$

where B is a regularization filter. Note that this is very similar to (3.2), but we replace the ℓ_1 norm with an ℓ_2 norm. We can solve this directly with Fourier transforms, which gives the following result:

$$\hat{s}_{\ell_2}=\mathcal{F}^{-1}\Big\{\frac{\overline{X(\omega)}R(\omega)}{|X(\omega)|^2+\gamma|B(\omega)|^2}\Big\} \tag{2.13}$$

where capital letters denote the Fourier transforms of the lowercase counterparts. Observe that when $\gamma=0$ this is an inverse filter. When $B=I$ and γ is large, this is equivalent to matched filtering with a constant scale factor. We use $A=I$ from (3.2) and $B=I$ from (3.5) in the following examples. The parameter values α and γ were chosen to be 10. They were selected from observation to give reasonable results for all the tested noise levels. Better results could be achieved for each noise level by tuning these parameters for each example.

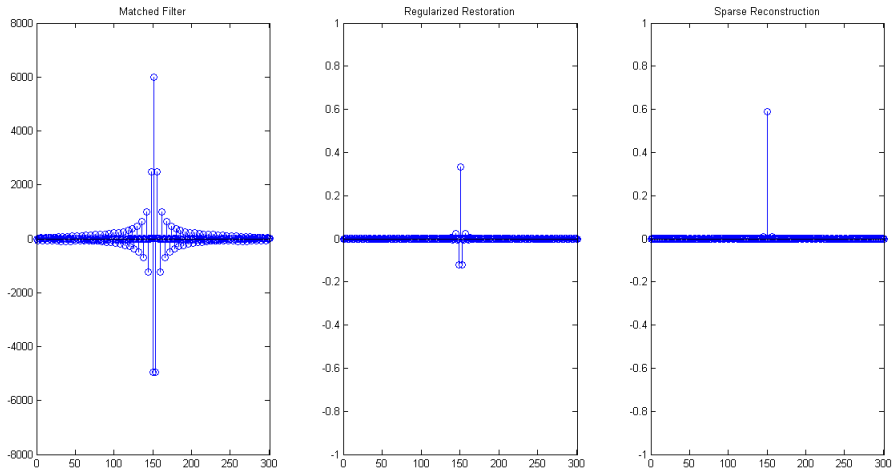

Figure 2.2: Matched Filtered Output (Left), Regularized Restoration (Center), Sparse Reconstruction (Right) of Impulse with No Noise

Since each of these techniques is an attempted recovery of the original signal, they should each resemble a centered impulse with height of 1. The matched-filtered result in Figure 2.2 no longer resembles an impulse because the filter is strictly bandlimited and recovers some frequency components more heavily than others. This leads to a significant spreading of the energy content, which decreases the resolution. Both the regularized restoration and sparse reconstruction methods perform fairly well at reconstructing the original impulse. The sparse reconstruction method is more successful both in concentrating the energy at the single impulse location (resolution) and in the accuracy of the recovered magnitude of the impulse. Note that both the regularized restoration and sparse reconstruction techniques can achieve perfect reconstruction in the noiseless case by setting $\alpha = 0$ and $\beta = 0$, but these values aren't practical selections for realistic applications. We show a comparison of resolution in Figure 2.3. This shows the percentage of energy contained within an interval relative to the total energy in the reconstruction as the interval size increases from the center of the impulse for each technique. A perfectly reconstructed impulse would spike up to 1 at the first value, since all of the signal energy is contained at that single value. The faster each technique reaches high values, the better the resolution. This graph shows that the

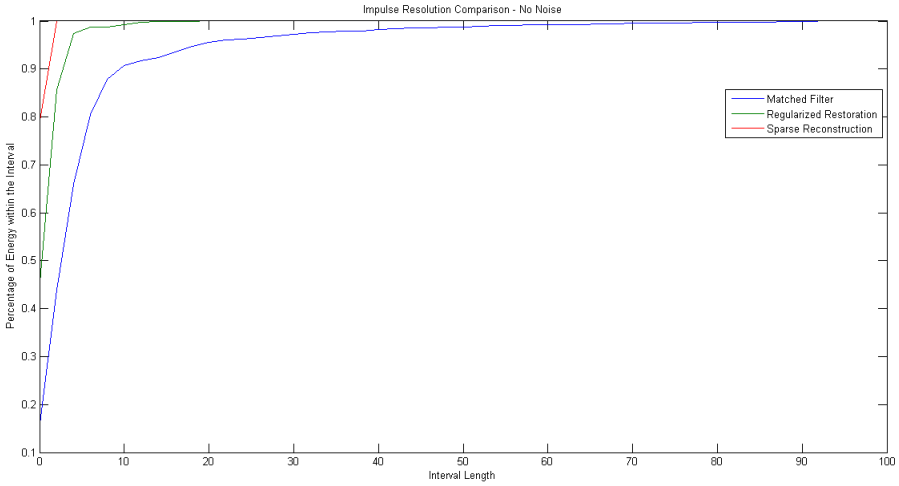

Figure 2.3: Measure of Energy Concentration in Increasing Intervals from the Center, Noiseless Case

sparse reconstruction method is the most impulse-like in its energy concentration, with an initially large value, followed by a single significant jump as the interval size increases. The regularized restoration method is the next best reconstruction for concentrating the energy into a small region, followed by the matched filtered reconstruction.

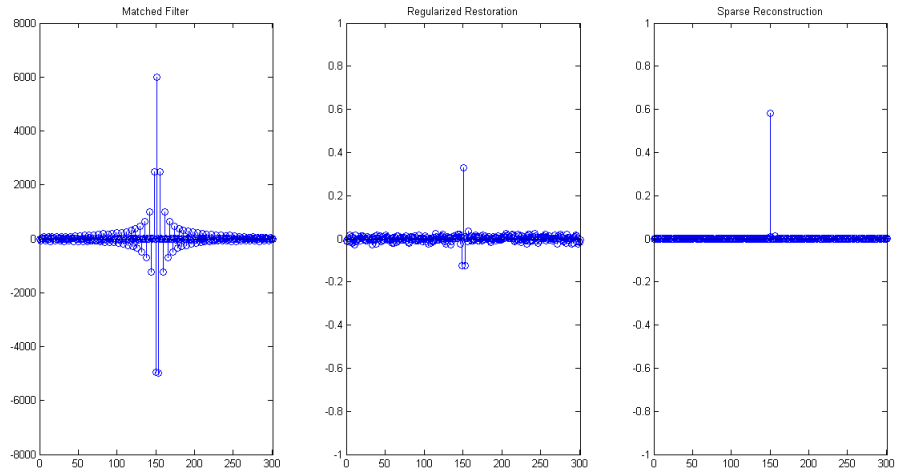

Figure 2.4: Matched Filtered Output (Left), Regularized Restoration (Center), Sparse Reconstruction (Right) of Impulse at 15 dB SNR

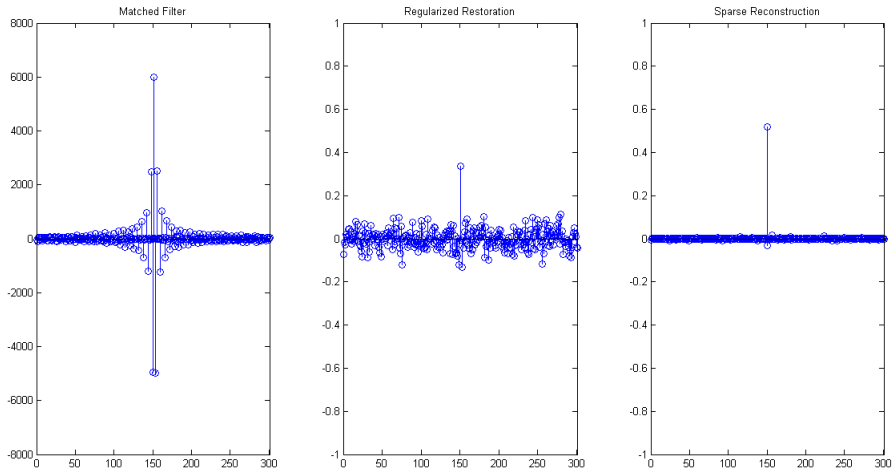

Figure 2.5: Matched Filtered Output (Left), Regularized Restoration (Center), Sparse Reconstruction (Right) of Impulse at 5 dB SNR

As the noise increases, these results become more pronounced, as shown in Figures 2.4 and 2.5. The matched-filtered reconstruction remains similar for both noise levels. The regularized restoration method shows significant errors in the reconstruction as the noise level increases. This is consistent with the technique, as reconstruction is attempted in frequency bands that can contain more noise than signal. The sparse reconstruction technique is shown to be fairly robust to noise for this parameter value. While additional reconstruction errors are created, they are much less significant than those created by regularized restoration. The resolution comparison curves in Figures 2.6 and 2.7 show similar results to the noiseless case for both sparse reconstruction and matched filtering. The regularized restoration method begins to suffer significantly as the noise increases. In the 15 dB SNR case, the energy is still initially more concentrated than the matched-filtered reconstruction, but the noise is consistently higher across the regularized restoration, so the matched-filtered output becomes more concentrated once the interval length reaches 7 pixels. In the 5 dB SNR case, the regularized restoration (at this parameter selection) performs worse than matched filtering at concentrating the total energy over any selected interval beyond the single central point.

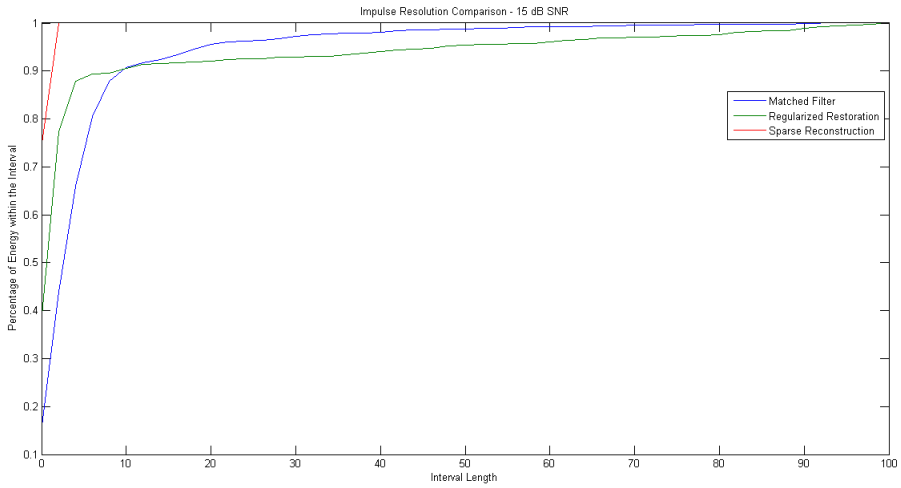

Figure 2.6: Measure of Energy Concentration in Increasing Intervals from the Center at 15 dB SNR

Figures 2.8 and 2.9 show MAE and RMSE values, respectively, for each technique at three different noise levels. The matched-filtered reconstruction was normalized for comparison, since the values are on a much different scale. One hundred trials were performed at both 15 dB SNR and 5 dB SNR with the same parameter selections, and the error calculations were averaged over all trials for each technique. The sparse reconstruction significantly outperforms the regularized restoration and the matched-filtered output at every noise level in both metrics. It is also much more robust to noise than the regularized restoration reconstruction, which performs significantly worse as the SNR decreases. The regularized restoration output outperforms the matched-filtered output at low noise levels for this parameter selection, but matched filtering stays more stable once the noise increases. Although matched filtering is robust to noise, the reconstruction error is relatively high for all tested noise levels. Note that the matched filtered result vastly differs in magnitude scaling from the original before normalizing, whereas regularized restoration and sparse reconstruction techniques recover values close to the original impulse.

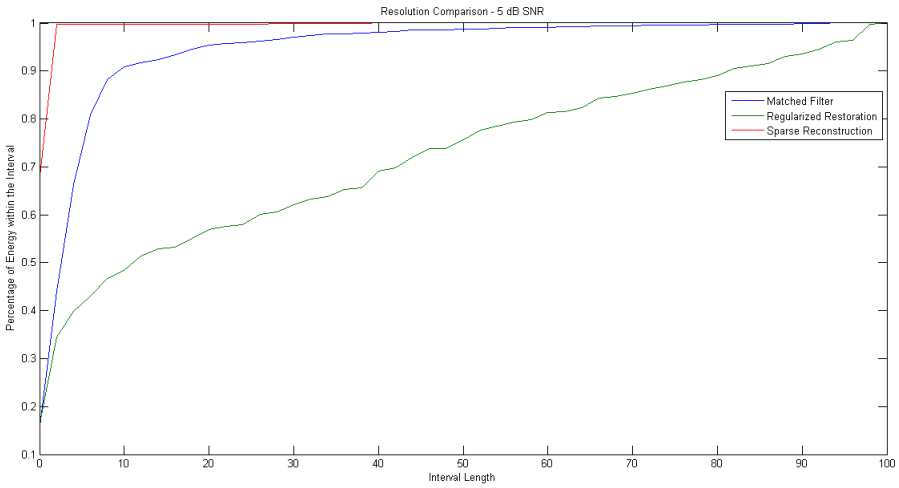

Figure 2.7: Measure of Energy Concentration in Increasing Intervals from the Center at 5 dB SNR

2.3.2 Experimental Results

The following plots were generated from data collected using an LFM chirp at various transmit bandwidths centered at 700kHz from the data collection at the SSTB at the Naval Surface Warfare Center in Panama City Beach, Florida. There was no significant movement during the data collection, so Doppler was not considered in the processing. Since the scene should resemble clusters of point reflectors for both target types, the regularized restorations and sparse reconstructions were performed using $A = I$ and $B = I$ in (3.2) and (3.5), respectively. By using a high-frequency transmit, the scene should appear as point reflectors from the specular return, followed by separated additional returns from acoustic waves travelling around and through the target, with multipath returns farther out. More specifically, we expect to see reflections from the surface of the water near 15", the initial specular reflection near 22.2", an unknown structural wave (whispering gallery or glory wave) near 22.7", the first Rayleigh wave near 22.9", a reflection of the specular return off of the source near 23.2", an unknown structural wave near 23.5", the second Rayleigh wave near 23.6", an unknown structural wave near 23.9", and a multipath reflection near 30". These distances were found by analyzing the known quantities of the target physics and

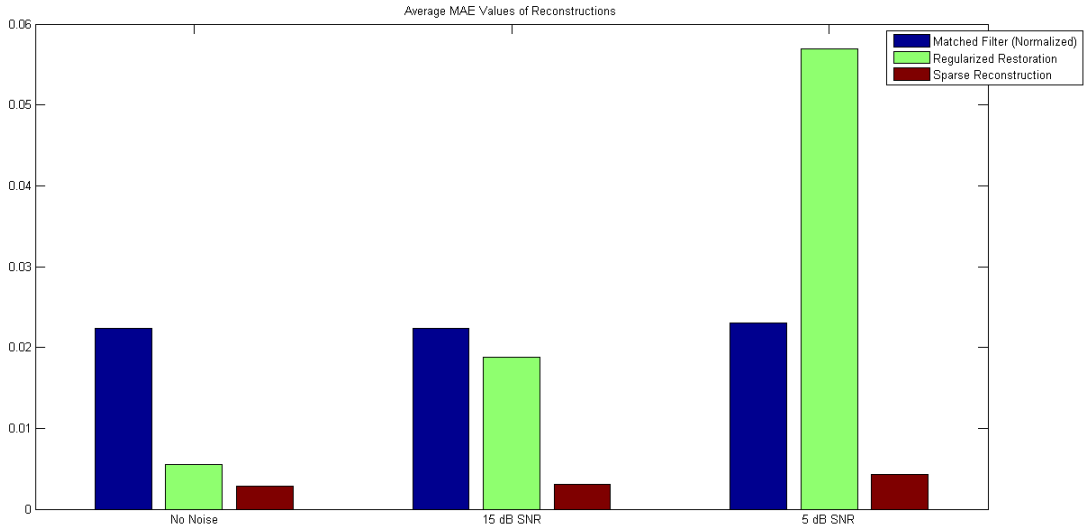

Figure 2.8: Average MAE values for Reconstructions

geometry of the tank setup, as well as by examining the visible returns in the wideband matched filter. Since this is experimental data, we no longer have an example of a perfect reconstruction, so much of the information is extrapolated from the expected physics as well as the behavior of the wideband matched-filtered reconstruction. Matched filtering is a common reconstruction technique that has been used and studied for years, and we consider the results from the wideband case, along with analysis of the target physics, to give the most trustworthy baseline available. We expect that all methods will become more similar as the bandwidth increases, because much less frequency content is allowed to move freely with the reconstruction algorithms. At narrow bandwidths, changing any frequency content of the scene outside of the transmit bandwidth will change the data consistency term very little (convolution with the transmit will suppress all but extreme values outside of the transmit band). All processing in this section used the transmit replicas shown in 2.10 as the given transmit description. These are recordings of the transmit signals recorded on a direct path from the source (without reflections) in the water. Using the replicas rather than analytic descriptions should increase the accuracy of all the techniques.

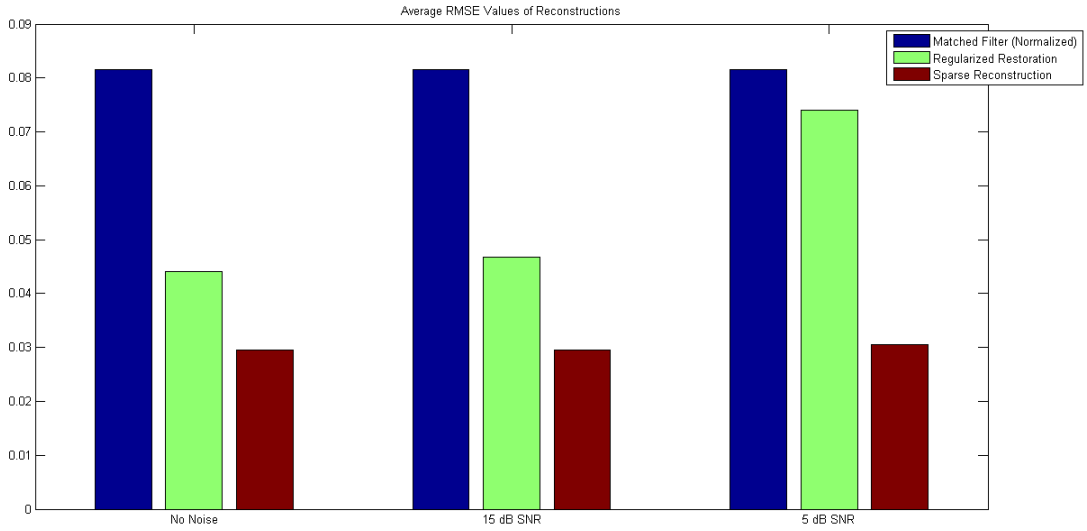

Figure 2.9: Average RMSE values for Reconstructions

Figure 2.11 shows a semilog plot of reconstruction frequencies for each of the three reconstruction techniques using the 600-800kHz LFM transmit replica. The matched filtered output is essentially bandlimited to the transmit band, with an additional DC component. The regularized restoration output has some significant frequency content outside of the transmit band, based on frequency content in the transmit replica outside of the desired transmit band. The sparse reconstruction output extrapolates bandwidth and can therefore have significant frequency content across the entire band.

Figure 2.12 diagrams the dominant paths visible in this data. Figures 2.13-2.18 show comparisons for the three techniques using different transmit bandwidths for a spherical target. Figures 2.13-2.15 show the entire reconstructed scene for different transmit bandwidths. The scaling for the regularized restoration and sparse reconstruction outputs are the same for both the standard and wideband cases, as both of these techniques attempt to recover the actual reflection coefficient values. The scaling was adjusted for the regularized restoration output in Figure 2.14 for visibility, as the reconstructed signal was barely visible above the noise content for this signal type. If we look at these reconstructions combined with the path information from Figure 2.12, we conclude that the information from Path 1, which is

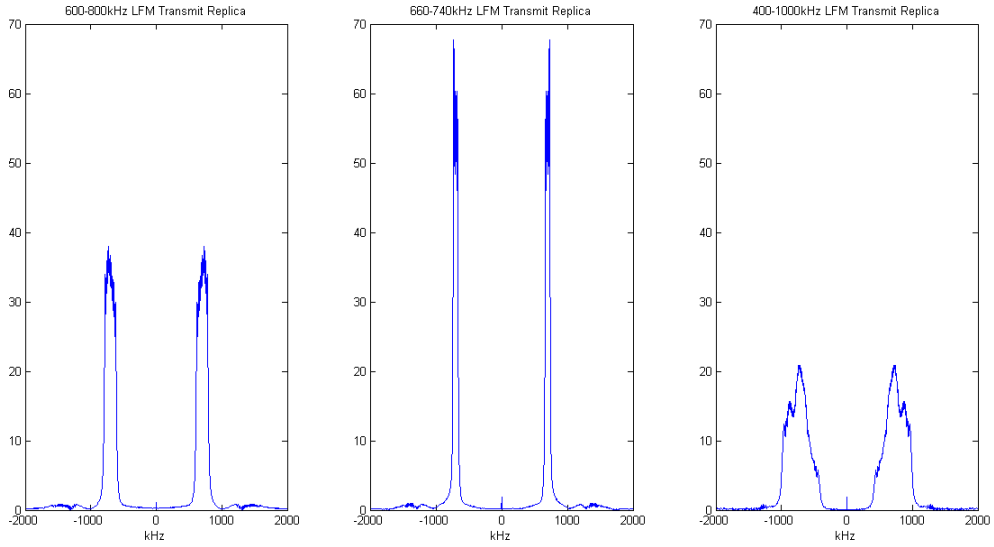

Figure 2.10: Comparison of Transmit Replica Frequency Magnitudes

a reflection off of the surface of the water containing no target information, occurs near 15" in the reconstructions. Path 2 contains the direct path target backscatter, and is present from approximately 21"-25" in the reconstructions. This includes the majority of the target information, including the specular reflection, unknown structural waves, the first and second Rayleigh waves, and a reflection of the specular return off of the source. Information from Path 3 is found near 30" in the reconstructions. This is a multipath reflection, which contains information from the target. We will focus on analysis of the information from Path 2, containing the direct backscatter. To assist in this, we include Figures 2.16-2.18, which are zoomed in on the content from 21"-25" from the same reconstructions. The scaling for the regularized restoration and sparse reconstruction outputs are once again the same for both the standard and wideband cases. The scaling was adjusted for the regularized restoration output in Figure 2.17 for visibility. The regularized restoration method gives an increase in resolution over matched filtering but also amplifies the noise using all three transmit signals. All scenes from the sparse reconstruction algorithm using all three transmit signals saw a significant reduction in noise and improved resolution. This was accomplished without losing most of the structural information defined above. Unfortunately, this parameter selection did

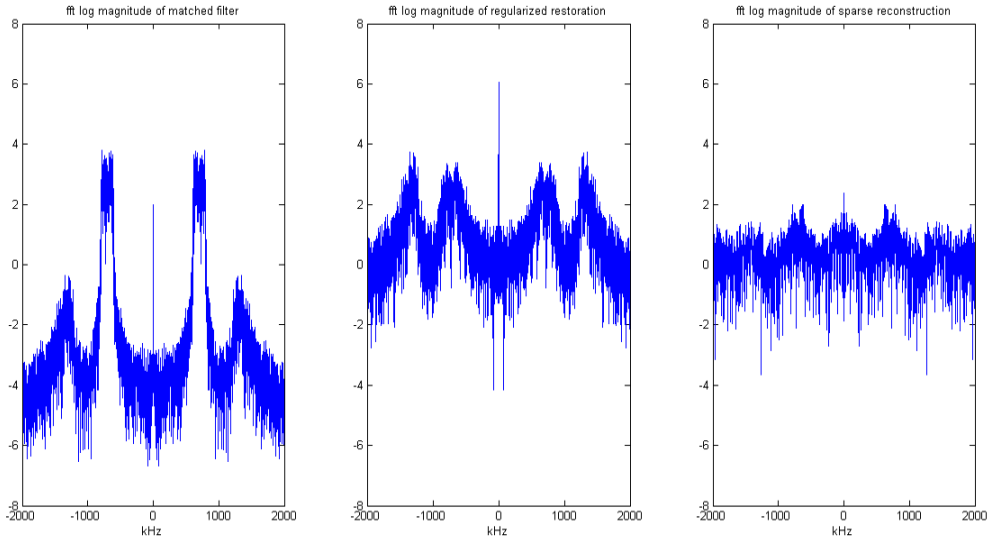

Figure 2.11: Semilog Plot of Frequency for 600-800kHz Transmit, Spherical Target, Matched Filter (Left), Regularized Restoration (Middle), Sparse Reconstruction (Right)

cause a loss of information near 23.5" and 23.6", but some additional fine-tuning should fix that problem. These structural waves were initially identified using the wideband matched-filtered reconstruction, so we still consider it a direct improvement in the 660-740kHz case, where these structural waves are not identifiable in any of the reconstructions. Because we have significantly more bandwidth in the reconstruction, we also see a substantial increase in spatial resolution. This could allow us to better resolve closely placed targets, as well as to detect more detail in the geometry of interrogated targets. This is particularly evident when identifying the return from the unknown structural wave near 22.7" in the spherical target. Some separation between the whispering gallery wave near 22.7" and the first Rayleigh wave near 22.9" is visible in the wideband reconstructions (Figure 2.18), but they overlap significantly in the narrowband matched filtered and regularized restoration reconstructions (Figure 2.17). This makes it incredibly difficult to identify the presence of this wave, and can potentially alter the perceived location of the first Rayleigh wave. In the case of an unknown target, this type of error could change the perceived target material or geometry.

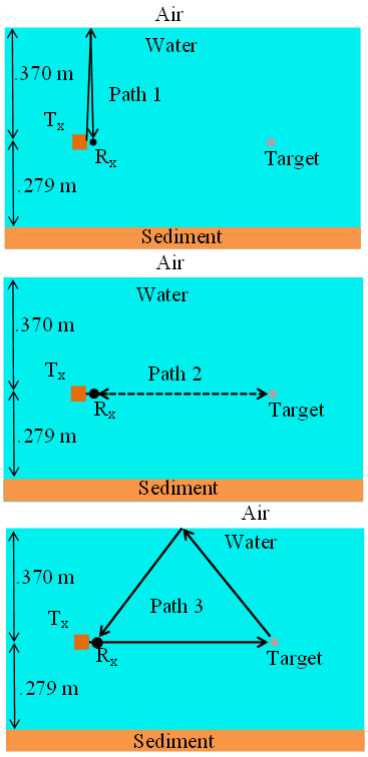

Figure 2.12: Diagram of Reflection Paths

The sparse reconstruction shows a clear separation between these reflectors for all of the transmit bandwidths used in the experiment.

Figures 2.19-2.21 compare regularized restoration and sparse reconstruction results directly, without adjusting the scale in the narrowband case. This shows a considerable improvement for the sparse reconstruction technique, as both an increase in reconstructed signal magnitude and a decrease in noise magnitudes are clearly visible for all transmit bandwidths. The results are most dramatic in the narrowband transmit results but are still significant when using a wideband transmit signal. We show a comparison of resolution in the 600-800kHz transmit case in Figure 2.22. This graph shows a much more rapid increase in the sparse reconstruction towards the maximum energy on the interval than either of the other techniques. This example is consistent with what we expect for this target type at this frequency range over the interval, as we expect it to resemble an impulse (or few impulses) from reflection on the front face of the sphere. Thus the energy should spike at locations

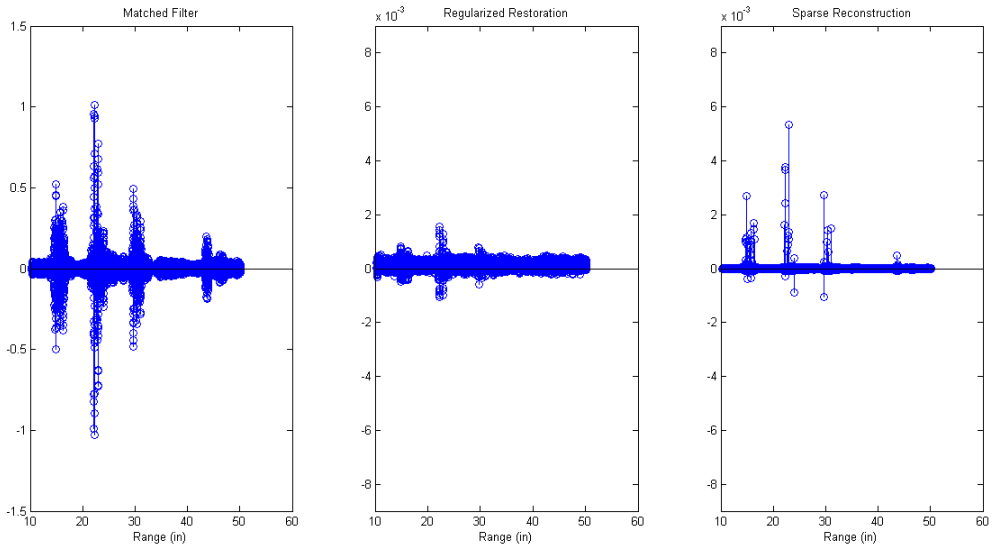

Figure 2.13: 600-800kHz Transmit, Spherical Target

of impulses. Regularized restoration shows an initial increase in the concentration of energy over matched filtering, but the large amount of noise at this parameter selection impacts the concentration of energy as the interval size increases towards the maximum.

Results for a cylindrical target in Figures 2.23-2.31 are similar to the results discussed for the spherical target. In this case, there is a surface reflection near 15" (Path 1), direct target scattering in 23"-27" (Path 2), and multipath returns (Path 3) after that. The axes for Figures 2.24 and 2.27 are adjusted for greater visibility of the regularized restoration. As the bandwidth increases, the three outputs become more similar, as more frequency information is contained within the data, so less needs to be recovered or extrapolated.

Although the sparse reconstruction results have a noticeable increase in the quality of the reconstruction, they also take considerably longer to process. Matched filtering and regularized restoration both took approximately 0.001s to process this experimental data on an Intel Core i5 processor. The sparse reconstruction algorithm took approximately 18.5s to process with the given parameter settings. These values can vary based on the size of the data and optimization of the code, but the sparse reconstruction will always take much longer to complete. The increase in processing time can be somewhat mitigated by relaxing

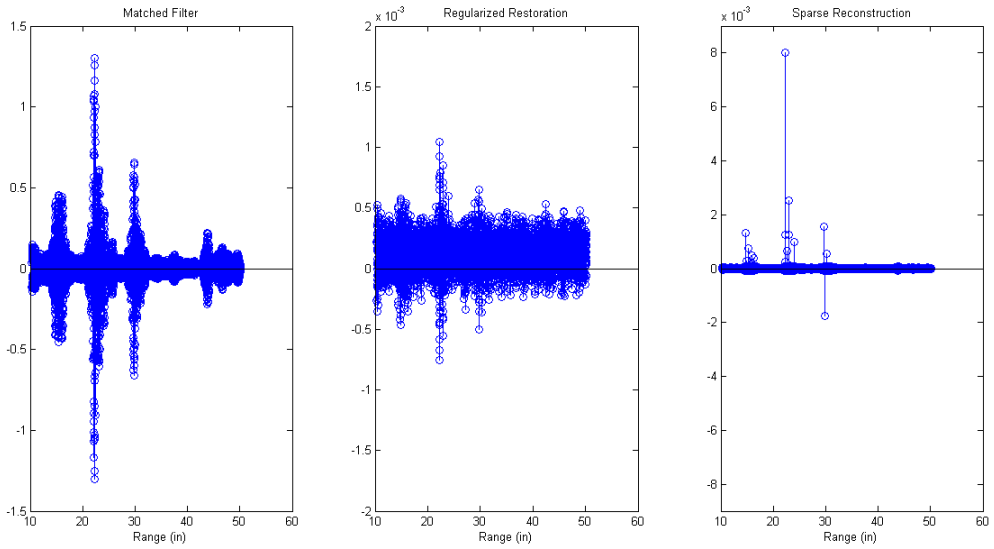

Figure 2.14: 660-740kHz Transmit, Spherical Target (Axes adjusted for Regularized Restoration)

some of the parameter selections (particularly the interval T in the majorization code (2.3)), but this can also lead to a decrease in algorithm performance as it moves the solution away from the ℓ_1 solution.

2.4 Conclusions

We developed and tested a sparse reconstruction algorithm on one-dimensional sonar data using high-frequency chirp transmit signals. Using simulation data, we showed how this method can recover bandwidth that can be lost with other reconstruction techniques without amplifying the noise. We then tested the algorithm on tank data and showed significant noise reduction and improved resolution over both matched filtering and regularized regularization results. This was accomplished with no specifications in the algorithm for a desired target type. Specific knowledge of desired target types could improve the results even further within the same framework. The time required to process the sparse algorithm is much longer than matched filtering or regularized restoration, so it is not a replacement for matched filtering or regularized restoration when rapid detection is the priority. It is well suited for extracting

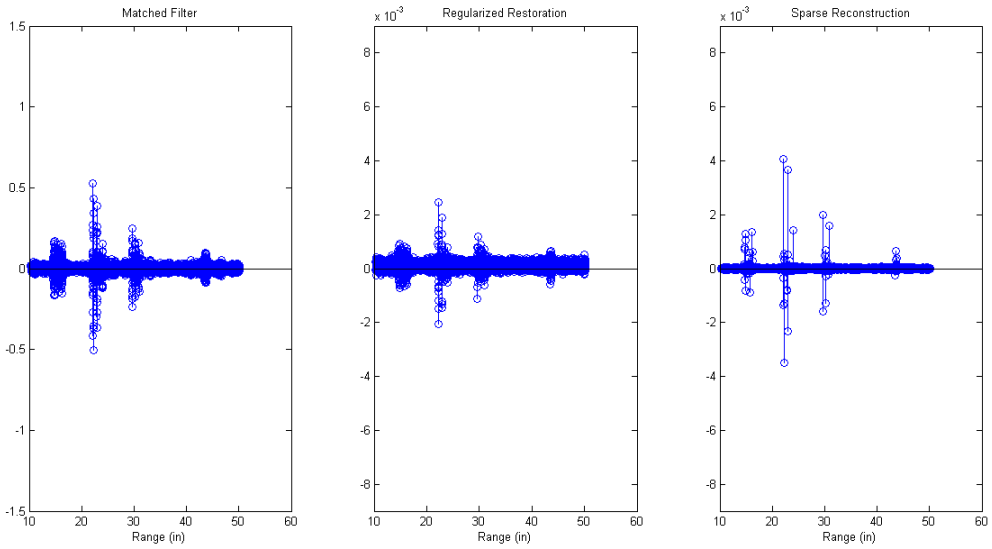

Figure 2.15: 400-1000kHz Transmit, Spherical Target

additional information from previously collected data sets, as well as in improving results in situations where minimizing the number of active transmits or using narrow transmit bandwidths is prioritized over speed.

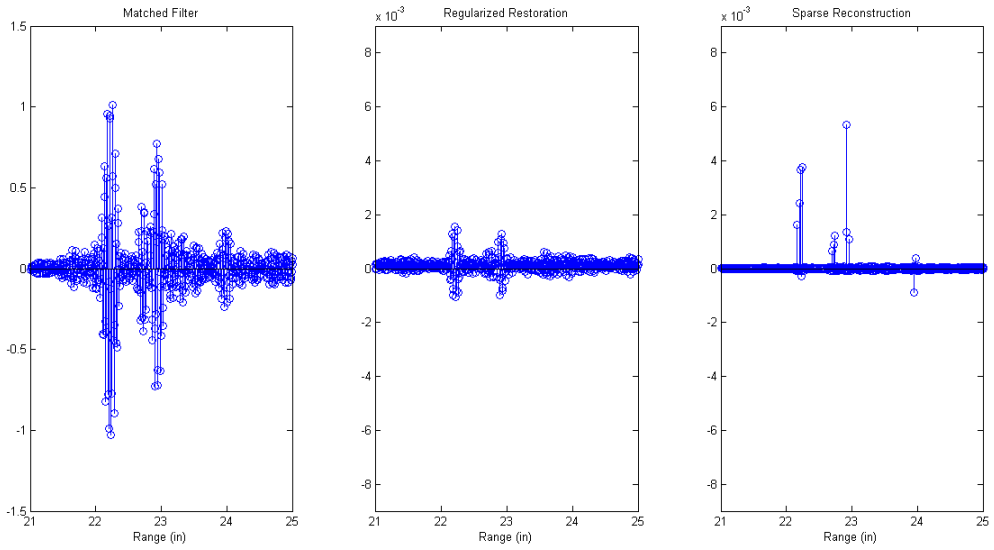

Figure 2.16: 600-800kHz Transmit, Spherical Target, Zoomed on Target

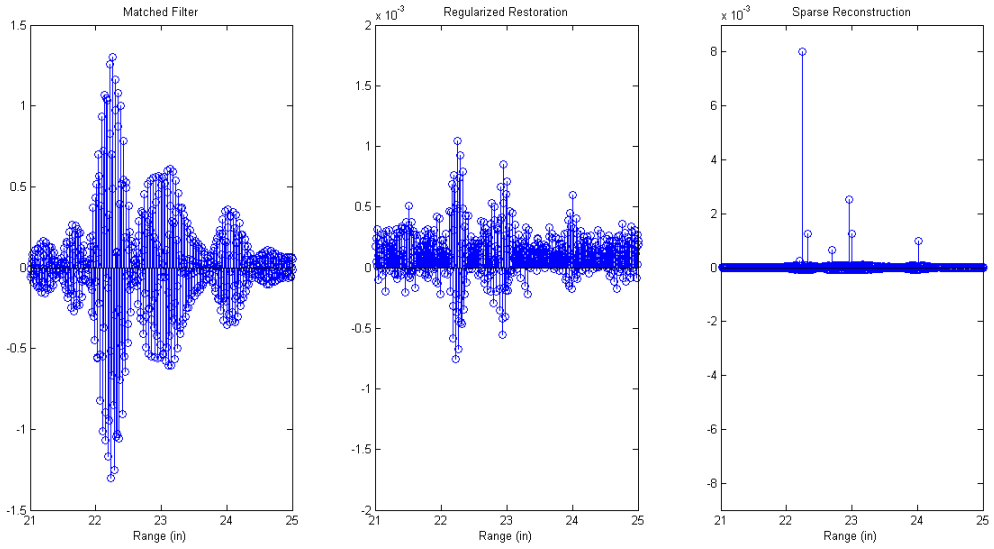

Figure 2.17: 660-740kHz Transmit, Spherical Target, Zoomed on Target (Axes adjusted for Regularized Restoration)

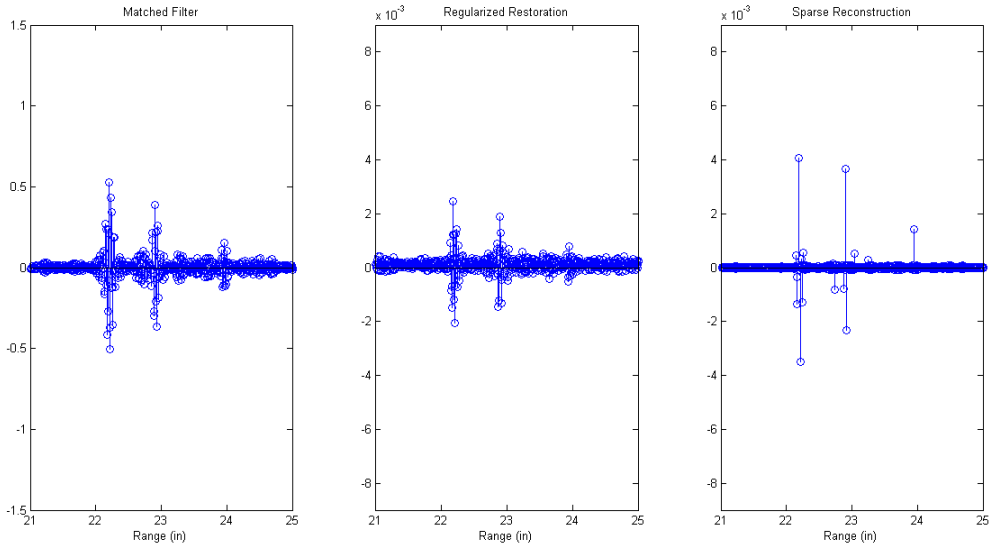

Figure 2.18: 400-1000kHz Transmit, Spherical Target, Zoomed on Target

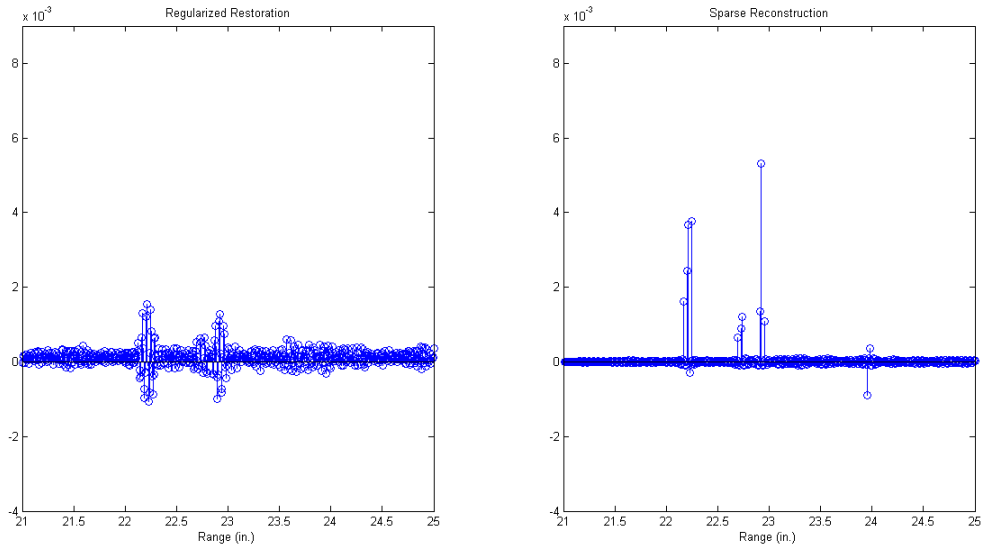

Figure 2.19: 600-800kHz Transmit, Spherical Target, Zoomed on Target, Same Axes

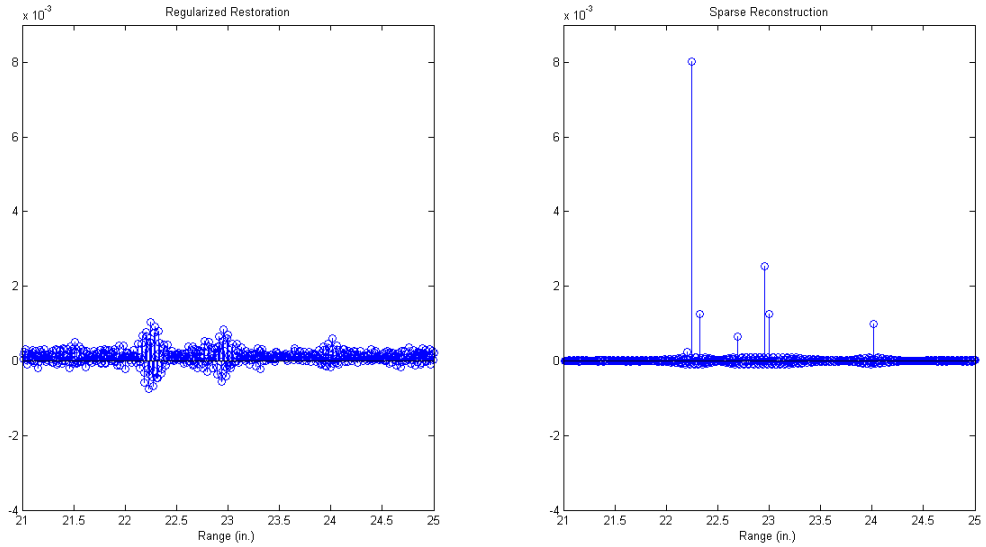

Figure 2.20: 660-740kHz Transmit, Spherical Target, Zoomed on Target, Same Axes

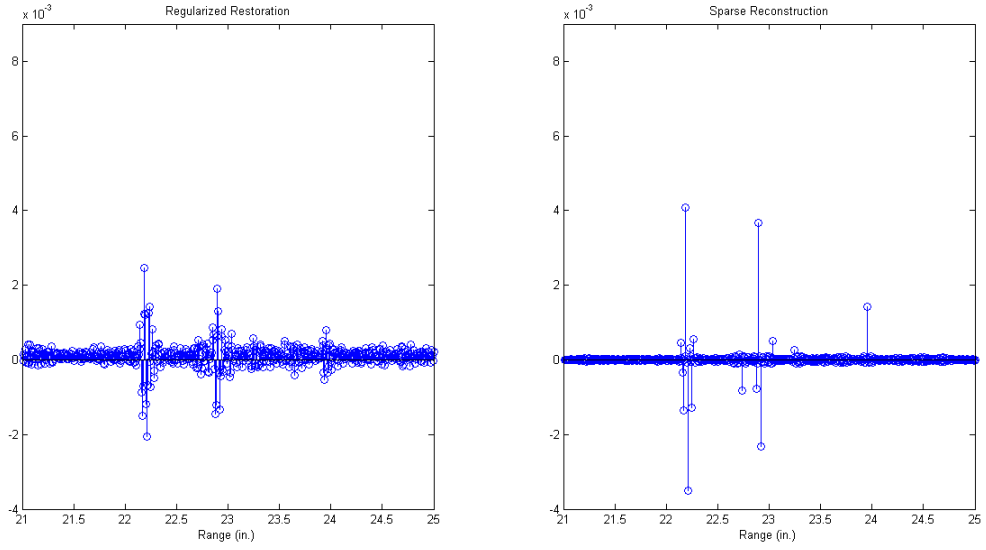

Figure 2.21: 400-1000kHz Transmit, Spherical Target, Zoomed on Target, Same Axes

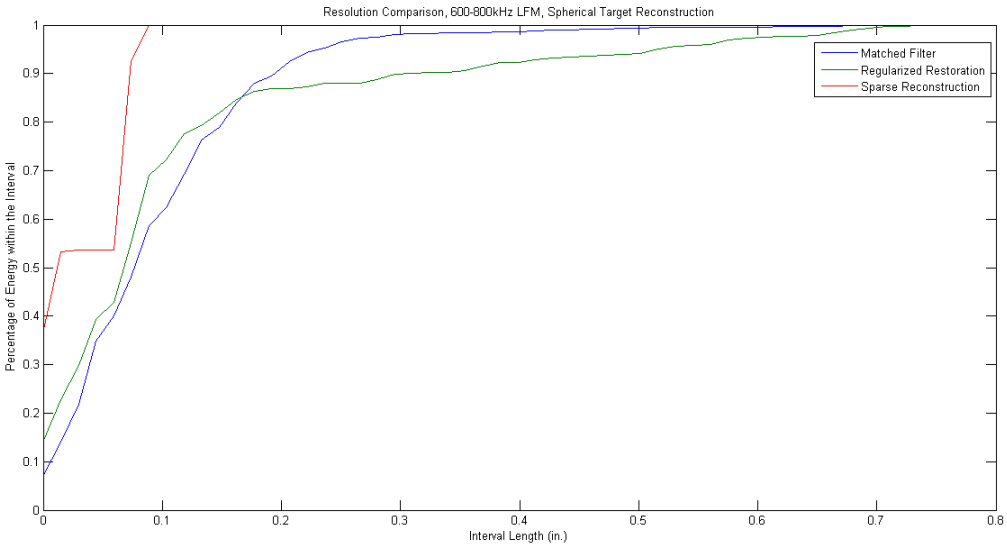

Figure 2.22: Resolution Comparison, 600-800kHz Transmit, Spherical Target

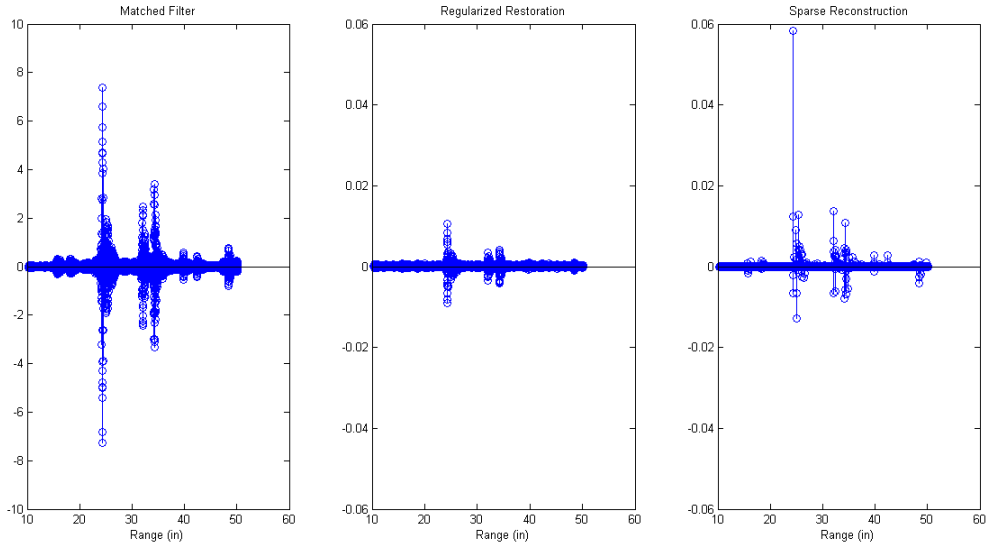

Figure 2.23: 600-800kHz Transmit, Cylindrical Target

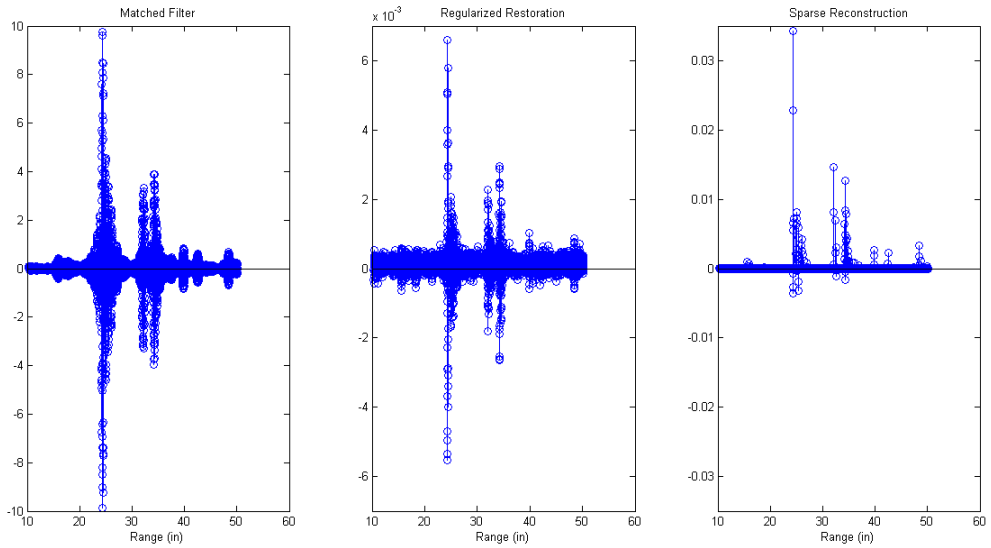

Figure 2.24: 660-740kHz Transmit, Cylindrical Target, (Axes adjusted for Regularized Restoration)

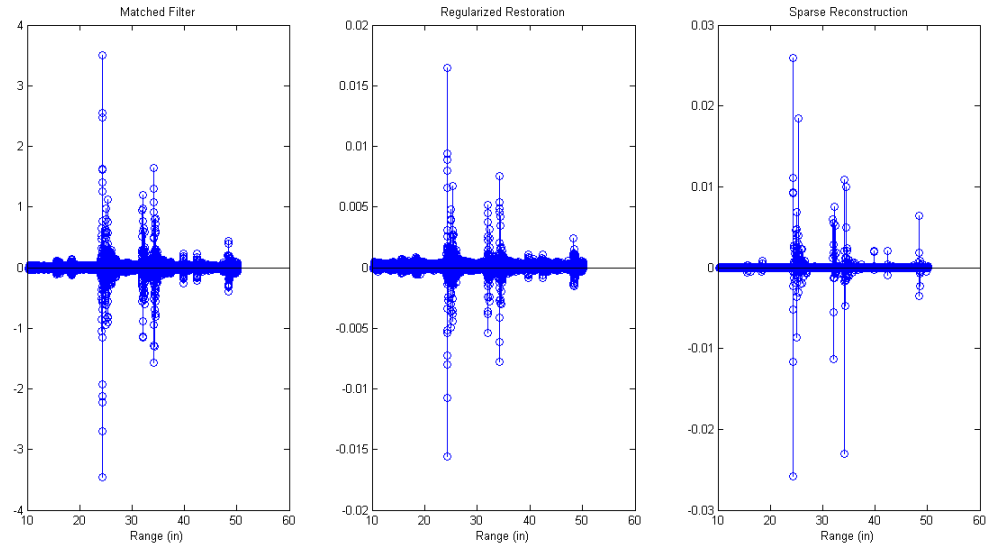

Figure 2.25: 400-1000kHz Transmit, Cylindrical Target

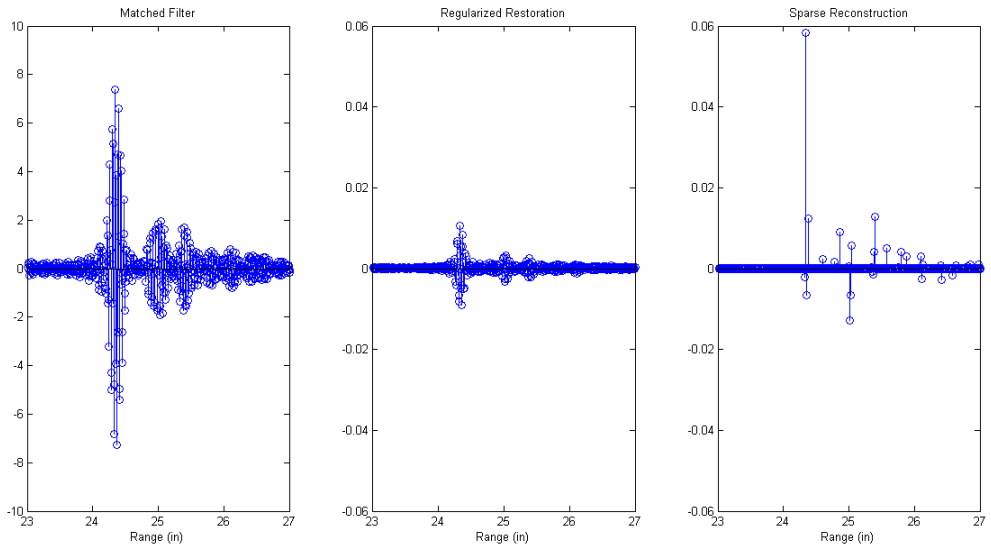

Figure 2.26: 600-800kHz Transmit, Cylindrical Target, Zoomed on Target

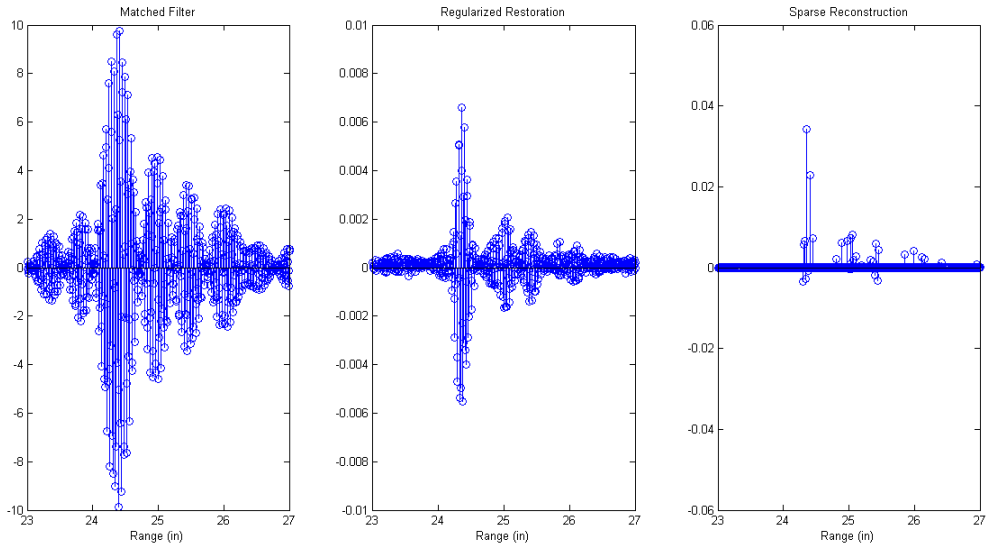

Figure 2.27: 660-740kHz Transmit, Cylindrical Target, Zoomed on Target, (Axes adjusted for Regularized Restoration)

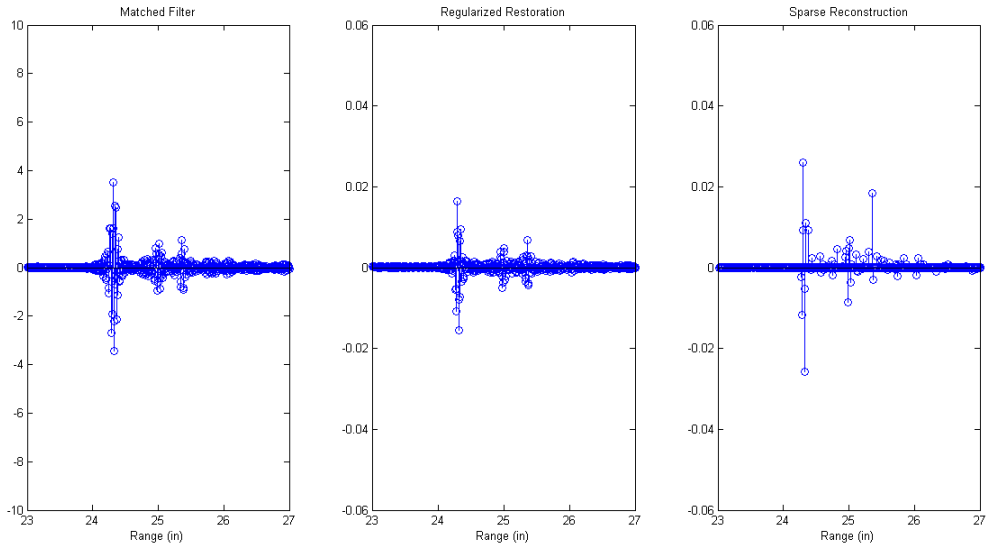

Figure 2.28: 400-1000kHz Transmit, Cylindrical Target, Zoomed on Target

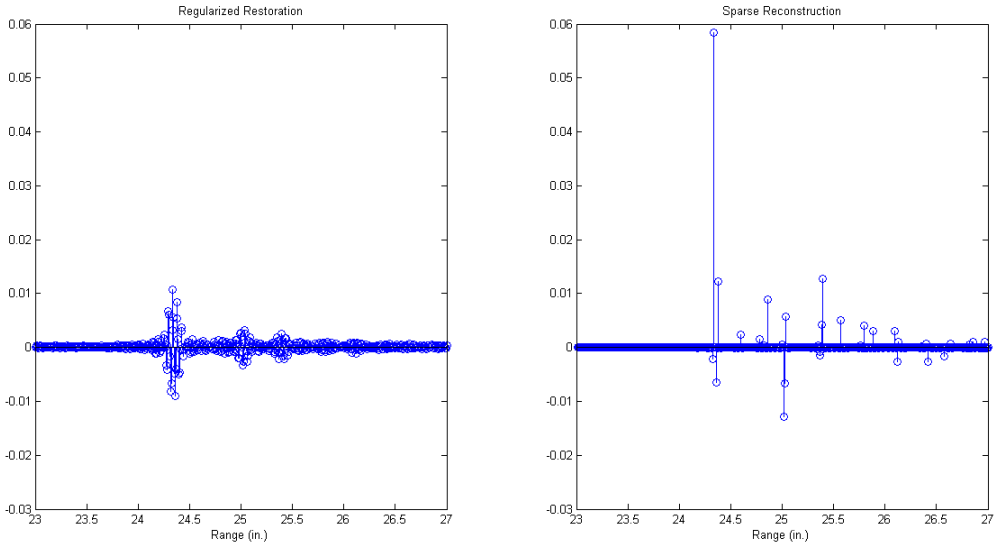

Figure 2.29: 600-800kHz Transmit, Cylindrical Target, Zoomed on Target, Same Axes

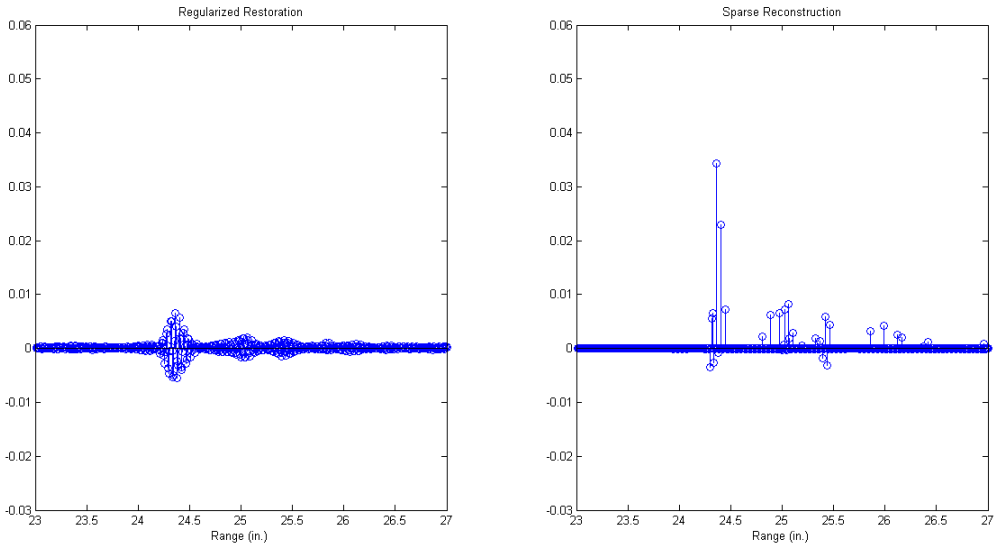

Figure 2.30: 660-740kHz Transmit, Cylindrical Target, Zoomed on Target, Same Axes

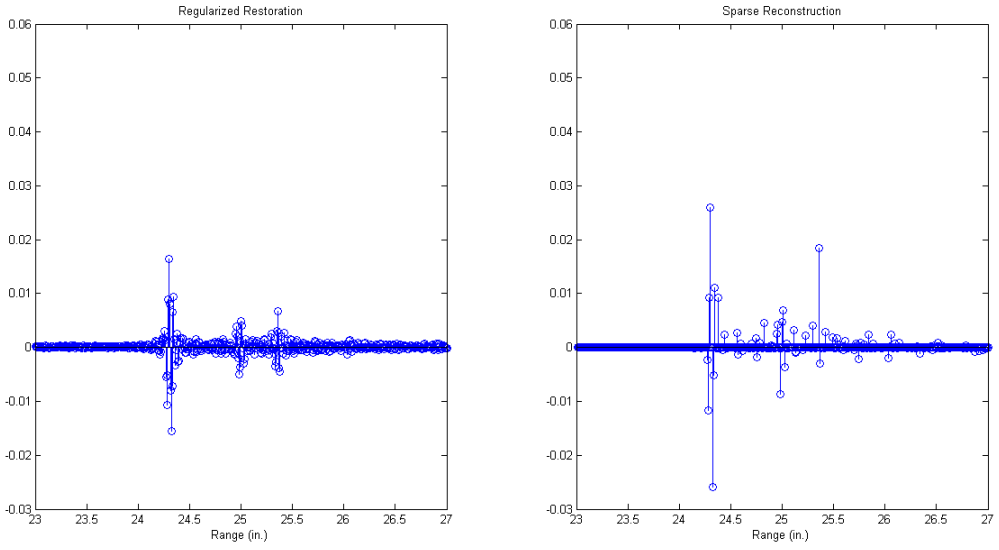

Figure 2.31: 400-1000kHz Transmit, Cylindrical Target, Zoomed on Target, Same Axes

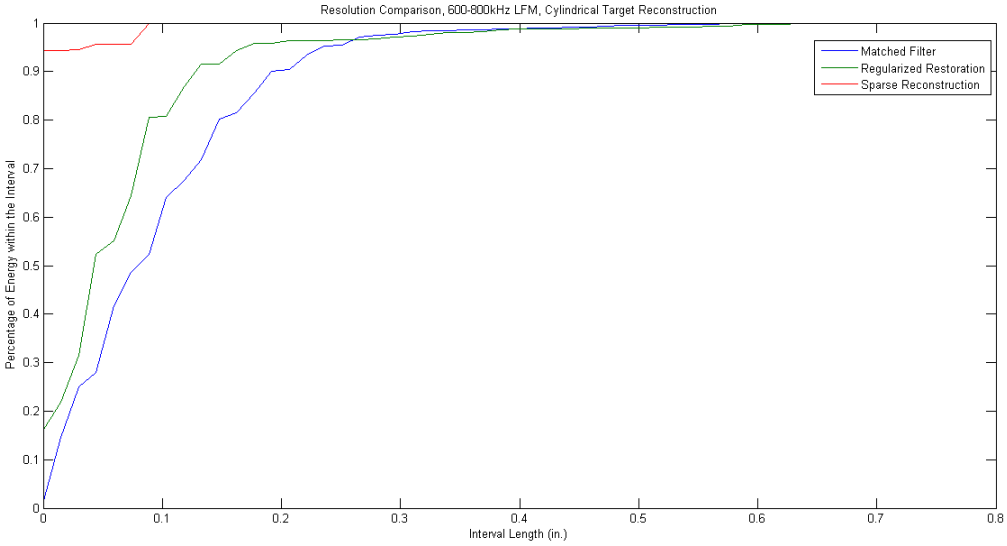

Figure 2.32: Resolution Comparison, 600-800kHz Transmit, Cylindrical Target

Chapter 3

Blind Deconvolution of One-Dimensional Returns

3.1 Introduction

In many sonar applications, the transmit signal is unreliable for accurate processing when using deconvolution-based processing. An unreliable transmit can result from a lack of recorded replica, inaccurately recorded replica, or additional signal degradation before transmission or through the water. When using a matched filter to process, the lack of an accurate replica can be relatively unimportant (in the case of additive Gaussian noise). Matched filtering is robust to this noise in both the transmit and recorded returns, so the lack of accuracy creates very little error in the output. When using regularized restoration, however, inaccuracies in the transmit description can create significant errors in the output. Inconsistencies with the phase can create problems for both reconstructions.

Instead, we seek to use the recorded data to simultaneously estimate the in-water transmit and interrogated scene information. We begin by assuming a sparse representation of the interrogated scene. Sparse reconstruction techniques have gained popularity in recent years for a large range of applications. An overview of some of these methods for image processing and radar can be found in [7] and [8], respectively. The applications are varied, but include through-the-wall radar imaging [10], synthetic aperture radar imaging [12], 3D sonar imaging [13], and synthetic aperture imaging applications [14, 15, 16]. In addition to the sparsity assumption, we will also require either an analytic transmit description or a recorded replica. This will create an additional term in the reconstruction equation corresponding to transmit consistency. This is justified because the in-water signal should bear some resemblance to the given description (either analytic or recorded replica), or none of the processing would be successful. This makes it a problem of semi-blind deconvolution. Blind

deconvolution is the process of simultaneously estimating two unknown, convolved signals. This case is semi-blind because we have a description of one of the signals, which bounds the range of credible solutions for that signal. By using this technique we can obtain a more accurate representation of the interrogated scene, since the transmit is allowed to change to a more accurate description during the processing. We also obtain a final representation of the in-water signal that can then be used in place of a replica for future processing. One can then modify the processing by either continuing to use the semi-blind deconvolution formula with a better initial estimate or by using faster processing methods with the new, more accurate transmit description. This approach allows for a great degree of control over the required processing time and level of accuracy.

Blind deconvolution techniques have been researched primarily in the context of image processing such as in [25, 26, 27, 28, 29]. They have also seen some use in acoustic systems, such as in [30]. Our approach is semi-bind deconvolution because we begin with a description of one of the signals- but believe it to be corrupted from the signal that is transmitted. Other semi-blind deconvolution applications include archaeological magnetic prospecting [31], Glottal flow estimation [32], and magnetic resonance force microscopy [33]. This setup gives us more information to work with than traditional blind deconvolution algorithms, which will restrict the solutions to realistic estimates.

Note that the simplest solution to this problem would be to record accurate replicas with each data collection. However, this can be problematic, as it may require a change in hardware and possibly the data acquisition process. This is also not a perfect solution, as noise in the recording, hydrophone inaccuracies, and environmental changes can still create errors in the replica which are then compounded in deconvolution-based processing.

3.2 Reconstruction Algorithm

We assume a high-frequency, discretely sampled sonar model with a modified transmit signal

$$r = T(x + n_1) * s + n_2 \tag{3.1}$$

where x is the analytic transmit signal, T is some transformation on x (such as from a nonideal transducer frequency response, platform movement/Doppler) $*$ represents linear convolution, s is the interrogated scene, n_1 and n_2 are noise processes, and r is the recorded return. Even in the simple case of $r = x * s$, this problem is ill-posed when assuming both x and s are unknown. In the frequency domain, this is equivalent to solving $X(k)S(k) = R(k)$ for $1 < k \leq M$, when only R is known. This yields M equations, each with an infinite number of solution pairs. We restrict our solutions by assuming that s has a sparse representation in some basis, and that the transmit signal x is a degraded form of a given transmit description x_0.

We begin with a basic sparse reconstruction equation to promote a sparse scene output with a known transmit signal.

$$\hat{s} = \underset{s \in \mathbb{R}^n}{\arg\min} \{ ||Xs - r||_2^2 + \alpha ||As||_1 \} \tag{3.2}$$

This equation can be viewed as the computationally tractable convex relaxation of the underlying sparsifying ℓ_0-norm deconvolution problem. We maintain this underlying sparsity model in the scene so that we can continue to reconstruct s under the same sparsity assumption but add a term to maintain consistency with the given transmit description x_0. This gives the equation

$$(\hat{x}, \hat{s}) = \underset{x \in \mathbb{R}^m, s \in \mathbb{R}^n}{\arg\min} \{ ||x * s - r||_2^2 + \alpha ||As||_1 + \beta ||x - x_0||_2^2 \} \tag{3.3}$$

The scene estimate will maintain consistency with the convolution model and also promote sparsity in the A basis based on parameter selection α. The transmit estimate $\hat{x}$ will also maintain consistency with the convolution model and retain similarities with x_0 based on parameter selection β. The use of the ℓ_2 norm for transmit consistency with x_0 is primarily for speed of computation. With specific knowledge of the way that the transmit signal has degraded, a better predictor model could conceivably be implemented instead of ℓ_2-norm consistency with a given description.

We use a gradient method to find the minimizers $(\hat{x}, \hat{s})$ of (3.3). The gradient in s is the same as that of (3.2), since the transmit consistency term does not depend on s. There are a variety of ways to approximate the solution to this sparse model. We find $\hat{s}$ by using a majorization approximation method outlined in [23]. This problem is then reduced to a sequence of quadratic problems with a closed-form solution at each iteration. This method allows us to compute $\hat{s}$ quickly, which is particularly important since it needs to be computed many times while updating $\hat{x}$.

To find $\hat{x}$, we need to take the gradient in x. The ℓ_1 sparsity term for s is eliminated and the problem is greatly simplified. The estimate $\hat{x}$ can be computed directly with Fourier transforms and is defined as

$$\hat{x} = \mathcal{F}^{-1}\Big\{ \frac{\overline{S(\omega)}R(\omega) + \beta X_0(\omega)}{|S(\omega)|^2 + \beta} \Big\} \tag{3.4}$$

where capital letters denote Fourier transforms of their lowercase counterparts.

The algorithm begins with initial estimates for x and s, then alternately updates each estimate with every iteration until the minimum is achieved. We can exploit the speed of calculating the gradient in x to significantly speed up the computation of the entire method. By running relatively few iterations in the s majorization algorithm, then finding the direct minimum in x, and alternating between these, we can more rapidly approach the global

minimum without adding significant processing time that would be required by a large number of iterations on s.

We also use regularized restoration both as a comparison with the semi-blind scene estimate and to show the improvements possible with using the updated transmit signal for additional processing. This is given by

$$\min_{s\in\mathbb{R}^n}\{||Xs-r||_2^2+\gamma||Bs||_2^2\} \tag{3.5}$$

where B is a regularization filter. Note that this is very similar to (3.2), but we replace the ℓ_1 norm with an ℓ_2 norm. We can solve this directly with Fourier transforms, which gives the following result:

$$\hat{s}_{\ell_2}=\mathcal{F}^{-1}\Big\{\frac{\overline{X(\omega)}R(\omega)}{|X(\omega)|^2+\gamma|B(\omega)|^2}\Big\} \tag{3.6}$$

where capital letters denote the Fourier transforms of the lowercase counterparts. Observe that when $\gamma=0$ this is an inverse filter. When $B=I$ and γ is large, this is equivalent to matched filtering with a constant scale factor.

3.2.1 Method Summary

1. Find initial estimates for both the scene and transmit. We can use either a transmit replica or analytic transmit description as the initial transmit estimate. We can find an initial scene estimate using (3.6) with the known return and initial transmit estimate.

2. Choose the sparsifying basis A for s. We chose a sparse identity for our examples since spherical and cylindrical targets at high frequencies resemble clusters of point reflectors.

3. Choose parameters α and β, which control the trade-offs between data consistency, sparsity of As, and transmit consistency with x_0. We used $\alpha=10^{-1}$ and $\beta=10^{-3}$ for the experimental data.

4. Alternately minimize with respect to each variable. The minimum in x is a single calculation, while the minimum in s requires an iterative process.

5. Obtain a new transmit estimate $\hat{x}$ and a new scene estimate $\hat{s}$ as outputs of the algorithm.

3.3 Results

3.3.1 Simulation Results

We begin by showing reconstruction errors incurred by using incorrect transmit signals with noisy data. Independent, identically distributed (i.i.d.) Gaussian noise was added to both the transmit signal, a 10-20kHz LFM chirp, and the return signal. We use a centered impulse target as an example throughout this section. Reconstructions were performed using the ideal signal description using matched filtering and regularized restoration techniques. We compare matched filtering and regularized restoration at 15 dB total SNR and 5 dB total SNR.

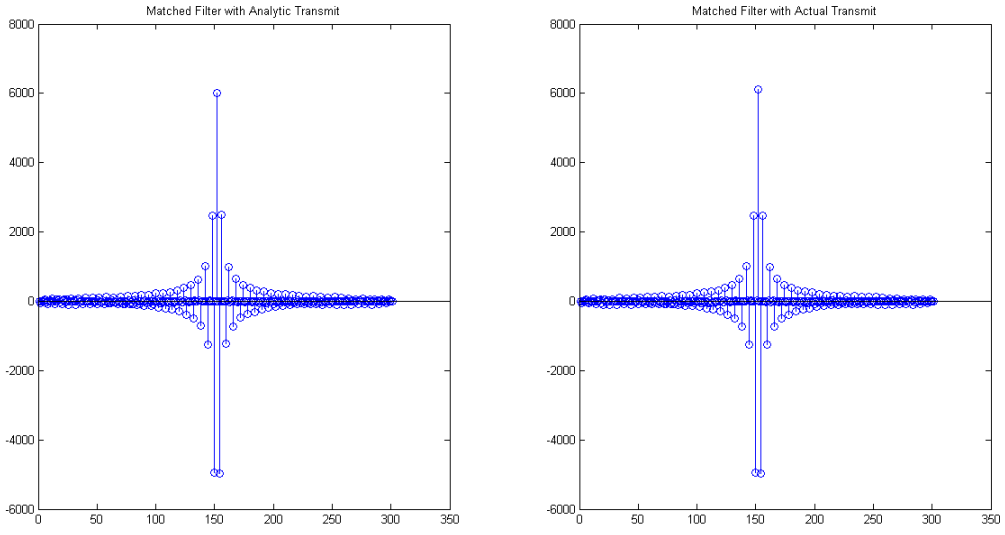

Figure 3.1: Matched Filter mismatch, 15 dB SNR

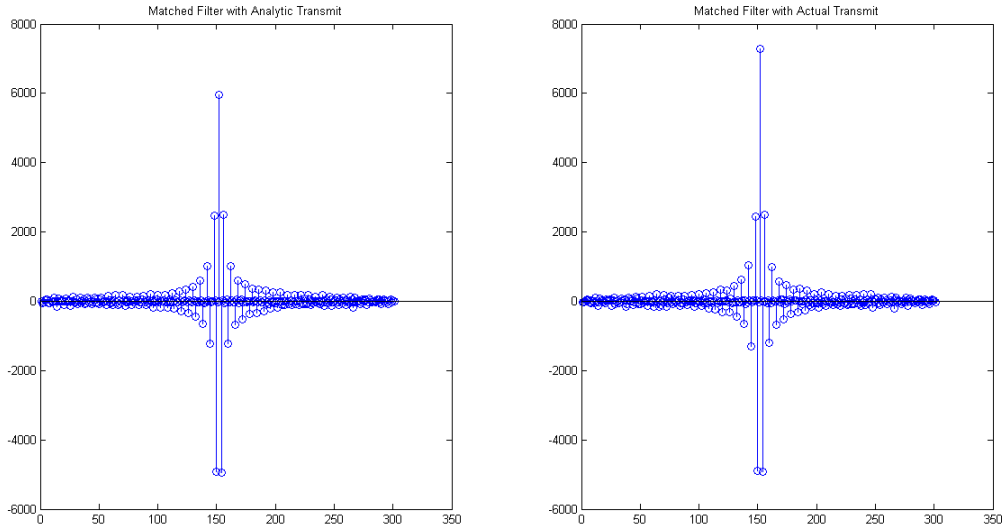

Figure 3.2: Matched Filter mismatch, 5 dB SNR

Figures 3.1 and 3.2 show matched filtered reconstructions using both the actual, noisy transmit and the ideal transmit description that did not include the noise. There is little difference between the reconstructions at both 15 dB SNR and 5 dB SNR. Figures 3.3 and 3.4 show the same experiment, but using regularized restoration instead of matched filtering. There is a much more pronounced difference between using the actual transmit and using the noiseless initial transmit description. There is a significant performance improvement when the correct transmit signal is known in its entirety. This observation motivates our use of semi-blind deconvolution to allow the transmit signal to change to a more accurate representation.

In Figures 3.5-3.9, we took the same noisy simulation data and ran our semi-blind deconvolution algorithm using only the return r and the noiseless transmit description x_0 as inputs. We used $A = I$, $\alpha = 10$, and $\beta = .01$ for (3.3).

Figure 3.5 shows the semi-blind deconvolution scene output, regularized restoration with the analytic (noiseless) transmit description, regularized restoration with the actual (noisy) transmit description, and regularized restoration with the semi-blind deconvolution transmit output at 15 dB SNR. The semi-blind deconvolution scene output does the best

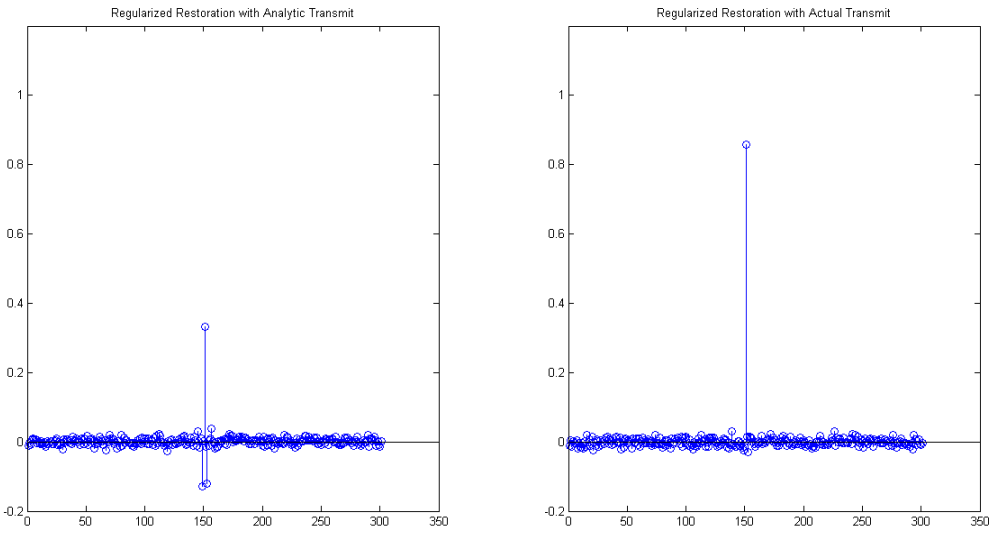

Figure 3.3: Regularized Restoration Mismatch, 15 dB SNR

at reconstructing the original impulse target. Regularized restoration with the analytic transmit has a noticeable decrease in magnitude of the impulse, along with some spreading of the energy and some noticeable noise. Regularized restoration with the actual transmit and the estimated transmit both do very well at reconstructing the impulse with some minor visible noise throughout.

Figure 3.6 compares the frequency magnitudes of the transmit signals. The leftmost plot (analytic transmit) is the input to the semi-blind deconvolution algorithm, while the middle plot (estimated transmit) is the transmit output of the algorithm. The transmit estimate fills out the frequency content that was added to the analytic transmit. This makes it much more accurate when used in additional processing. Figures 3.7 and 3.8 show the same comparisons at 5 dB SNR. Regularized restoration with the analytic transmit suffers even more at this point, while the other methods maintain stable results. The estimated transmit fills out the frequency content to more closely match the actual transmit, as in the 15 dB SNR case.

Figure 3.9 shows another example at 5 dB SNR, but this example was done using a lower regularization parameter in the regularized restoration algorithm. In this case, regularized

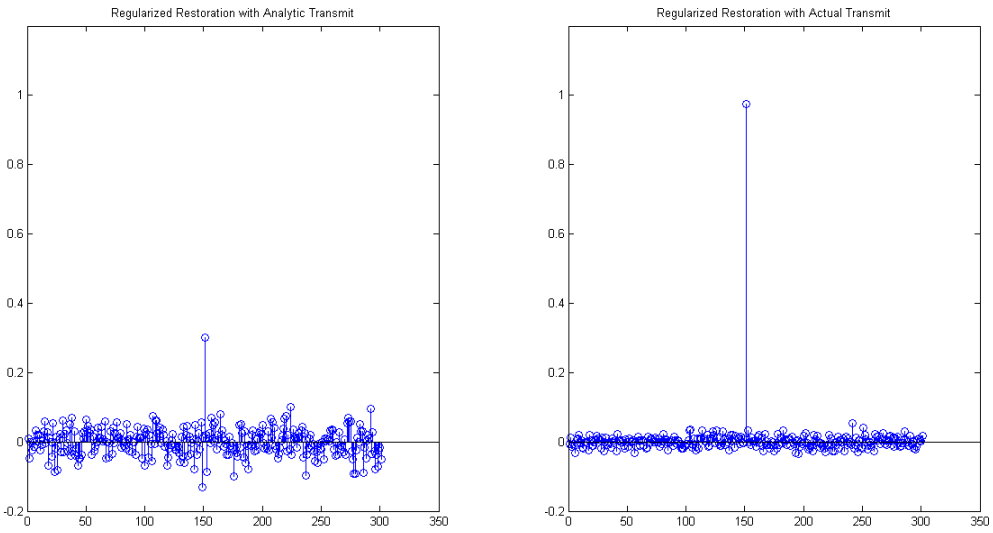

Figure 3.4: Regularized Restoration Mismatch, 5 dB SNR

restoration with the analytic transmit description is unrecognizable as an impulse, while regularized restoration with the actual and estimated transmits are largely unchanged from Figure 3.7. Using the semi-blind deconvolution technique has added stability to regularized restoration processing, since it has filled out the frequency information. Much of the reconstruction error caused by using the analytic transmit for the regularized restoration is due to dividing by very low frequency magnitude values in the transmit, which causes the noise in the return to get amplified. Since the transmit estimate fills in much more of that frequency content, it no longer amplifies the noise in the return as much as in the analytic signal. This is particularly important if the return signal has significant frequency content outside of the analytic transmit band (such as from Doppler, resonance, transducer response, etc.), as these unplanned frequencies can get amplified greatly if the regularization parameter is not carefully tuned to suppress them.

Error comparisons are shown in Figures 3.10 and 3.11 for SNR values of 15 dB, 10 dB, and 5 dB, averaged over 1000 trials. These show a consistent pattern of error behavior for each method. The semi-blind deconvolution scene estimate tends to do the best, followed by

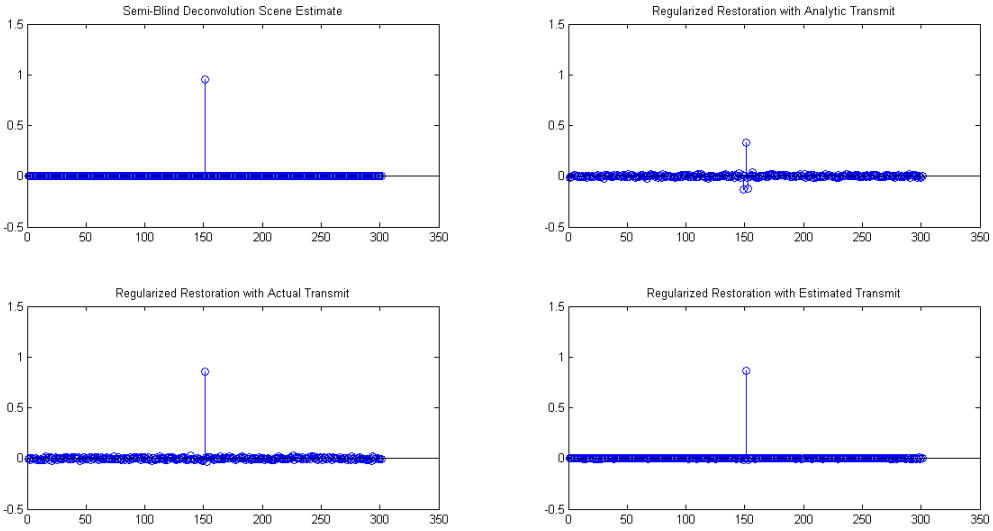

Figure 3.5: Semi-Blind Deconvolution, 15 dB SNR

regularized restoration with the estimated transmit, then the actual transmit, then the analytic transmit. Regularized restoration with the estimated transmit outperforms regularized restoration with the actual transmit because regularized restoration with the actual transmit does not attempt to resolve the additional noise added to the return except through the regularization parameter (which is the same for all the regularized restorations). The sparse modeling involved in finding the estimated transmit encourages the estimated transmit output to have a sparser result, which is more accurate to the scene. It is worth noting that the majority of the error in the scene estimates as the SNR decreases comes from overestimating the magnitude of the impulse. The additional content from the noise causes the algorithm to raise the magnitude of the reconstructed impulse, yielding higher error values. It still succeeds in heavily suppressing values outside of the central impulse.

Figures 3.12 and 3.13 show comparisons of the average resolution of each reconstruction method. Since the scene is an impulse, all of the energy should be contained in the single central point, meaning that the graph should reach 1 immediately. The bigger the interval required to reach a given energy percentage, the worse the resolution of the method. The semi-blind deconvolution scene estimate and regularized restoration with the transmit

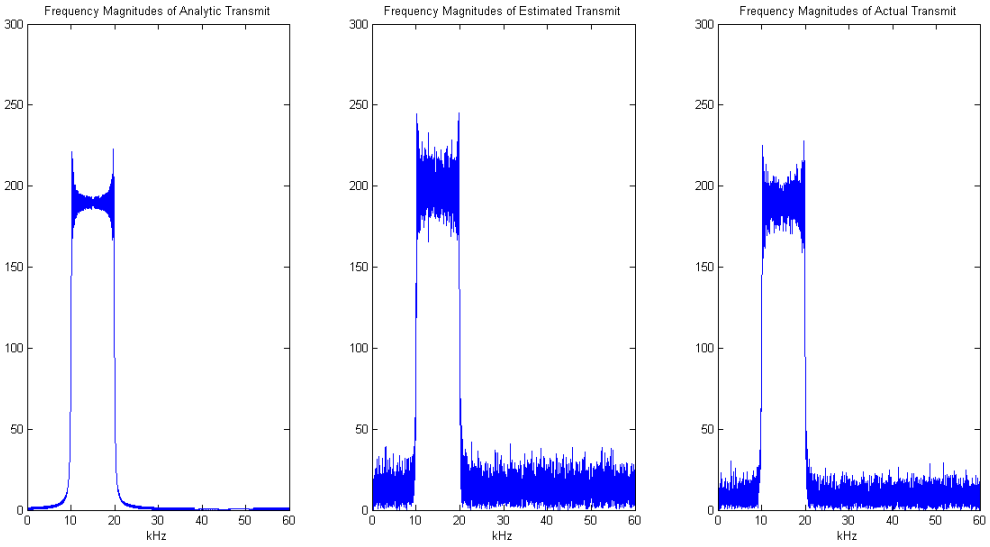

Figure 3.6: Transmit Comparison, 15 dB SNR

estimate have a clear resolution advantage over regularized restoration with the actual and analytic transmits. Regularized restoration with the estimated transmit has a better resolution than regularized restoration with the actual transmitin this example. We expect to see an imperfect reconstruction when using the actual transmit because noise was added to the scene as well as the transmit. This means that some energy should be spread across the reconstructed scene. The estimated transmit is not designed to distinguish between the noise on the signal and the noise on the scene and incorporates both in the estimated transmit in order to promote a sparse scene estimate.

3.3.2 Experimental Results

The following data was collected at the SSTB at the Naval Surface Warfare Center in Panama City Beach, Florida. We compare the blind deconvolution scene output s, regularized restoration with the transmit estimate $\hat{x}$, regularized restoration with an in-water replica, and regularized restoration with the analytic transmit description. The algorithm used only the transmit replica and return as inputs in the following examples. Results are similar when using the ideal description as an input, but the recorded replica is closer to the

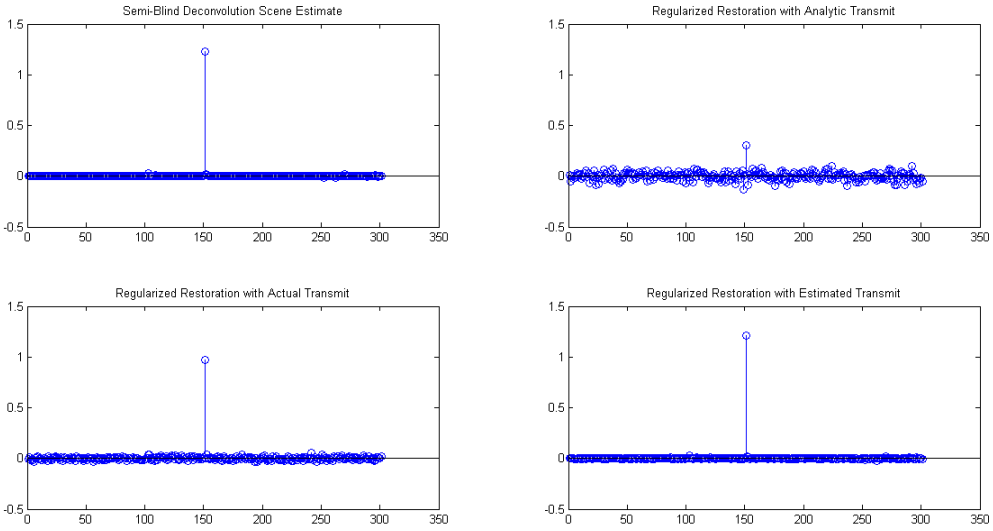

Figure 3.7: Semi-Blind Deconvolution, 5 dB SNR

actual in-water transmit and should thus create fewer errors in the reconstruction. The data is the same as the previous chapter, so the expected target locations are identical. Thus, we expect to see reflections from the surface of the water near 15", the initial specular reflection near 22.2", an unknown structural wave (whispering gallery or glory wave) near 22.7", the first Rayleigh wave near 22.9", a reflection of the specular return off of the source near 23.2", an unknown structural wave near 23.5", the second Rayleigh wave near 23.6", an unknown structural wave near 23.9", and a multipath reflection near 30". Like the previous chapter, we expect the results to become more similar as the bandwidth increases.

We begin with comparison of analytic, estimated, and replica transmit signals in Figures 3.14-3.16. The leftmost images in these figures are the frequency magnitudes of the analytic signals, the center images are the frequency magnitudes of the estimated transmit signals, and the rightmost images are the frequency magnitudes of the recorded replicas of the transmit signals. Our algorithm maintained consistency with the replicas, which is clear in the figures, as much of the content remains unchanged in the transmit bands. All of the estimated transmits have a much higher DC component and larger magnitudes on the higher frequency content contained in the replicas. This is consistent with results which are discussed more

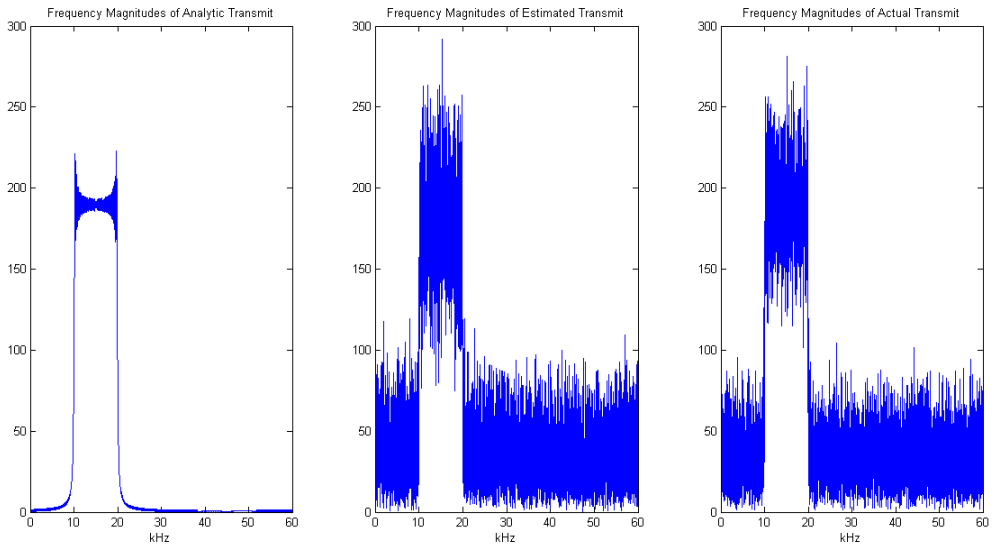

Figure 3.8: Transmit Comparison, 5 dB SNR

in-depth in Chapter 4, for what happened to this transmit replica over a larger distance. The DC component and higher frequency content grows in magnitude relative to the in-band content. So while it is difficult to claim that these estimated transmit are definitively more accurate, the changes that occur are consistent with what what occurs to the transmit replicas over a larger distance.

Figures 3.17-3.19 give the full range reconstructions from the semi-blind scene estimates, regularized restoration with the analytic transmits, regularized restoration with the transmit replicas, and regularized restoration with the semi-blind estimated transmit. Path information from 2.12 is the same as the previous chapter, where information resulting from waves following Path 1 is contained near 15" in range. Information from waves following Path 2 is contained in the 21"-25" in range, and information from Path 3 is located near 30" in range. We focus on the direct target information in Path 2 by focusing on the 21"-25" region in Figures 3.20-3.22. These images show similar improvements in the scene estimate to those generated in Chapter 2. In addition to this, we also generate regularized restoration results using an estimated transmit signal. In order to give a fair comparison of this method, we include Figures 3.23-3.25. These are direct comparisons of all three regularized restorations

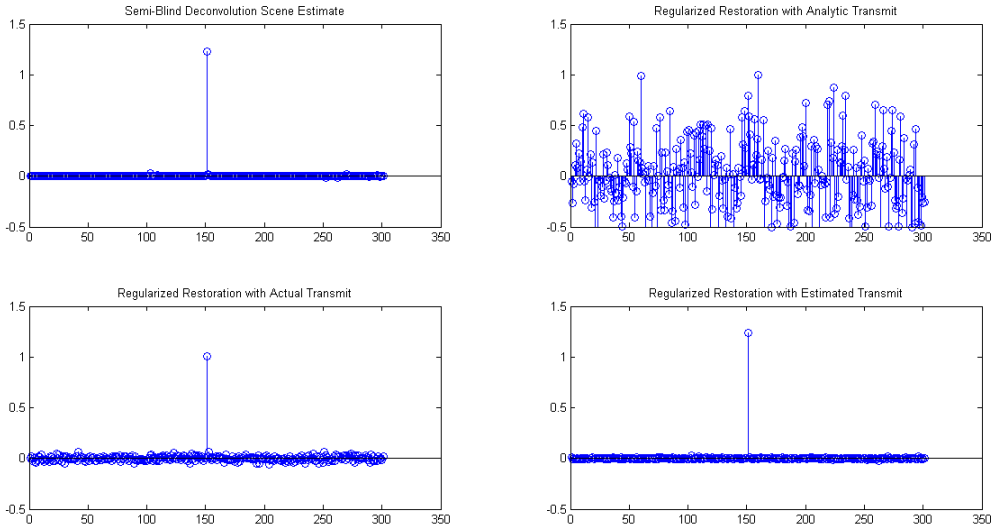

Figure 3.9: Semi-Blind Deconvolution, 5 dB SNR

at all three transmit bandwidths. In each of these examples, the reconstructed reflections are strongest when using the estimated transmit signal in the regularized restoration. This increase comes without a noticeable increase in noise over the replica reconstruction (this only considers the general noise floor present throughout, as we cannot accurately measure inaccuracies in the reconstructed signal arising from this technique). Regularized restoration with the estimated transmit maintains much more structural information and a larger increase in magnitude in the narrowband case. We can see a visible separation between the unknown structural wave near 22.7" and the first Rayleigh wave near 22.9". This separation is much more difficult to spot in both of the other regularized restoration reconstructions when using the narrow bandwidth transmits. Overall, this makes regularized restoration with the estimated transmit more desirable for classification purposes, as target features are more easily identifiable.

Figure 3.26 gives a resolution comparison using the spread of energy as the interval size increases from a central point. This comparison is more difficult to analyze than previous versions because the first point does not contain the majority of the energy. The blue line represents the semi-blind deconvolution scene estimate. While it appears to have lower values

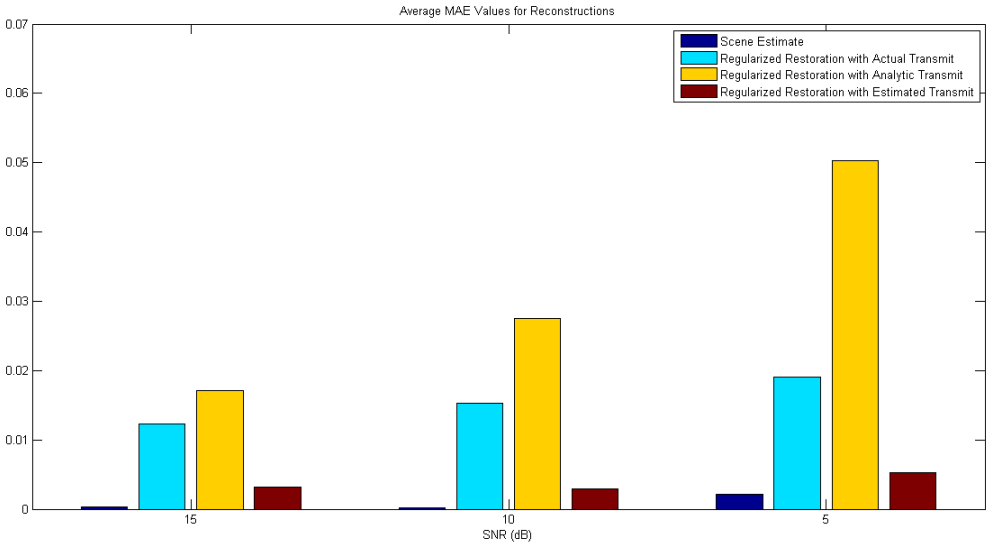

Figure 3.10: Average MAE Values of Reconstructions

initially, we actually look for large jumps to represent impulse-like behavior. Thus, while the energy content at the initial point isn't particularly high relative to the total energy, it has only a few impulse-like jumps towards the maximum. We see some similar behavior from regularized restoration with the estimated and replica transmits. They tend to follow the semi-blind curve, but in a smoother fashion. Regularized restoration with the analytic transmit has higher energy concentration than all techniques near the start, but the steadily sloping behavior indicates a gradual accumulation of energy as the interval increases.

Results for a cylindrical target in Figures 3.27-3.39 are similar to the results discussed for the spherical target. The estimated transmits were altered somewhat more than their spherical counterparts in Figures 3.14-3.16. This is because the blind deconvolution has a dependence on the given return information for the output. The majority of the changes are consistent with the spherical examples, which are justified for the same reasons as above, but there is some error inherent in using a single target for this technique. This indicates that the best method for attaining an accurate in-water estimate would likely come from averaging the estimated transmit signals over several data collections. The reconstructions in Figures 3.30-3.38 have similar results to the spherical counterparts. The estimated transmit

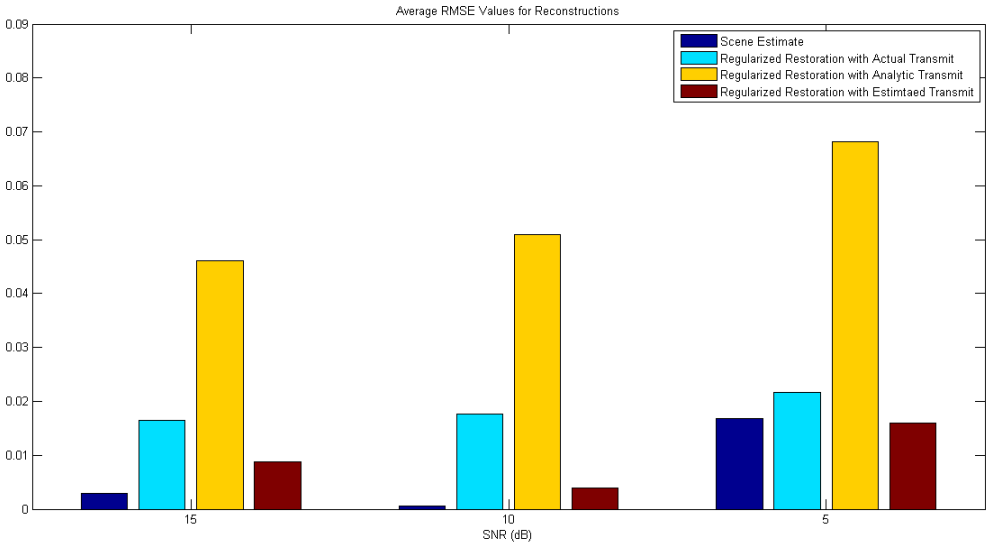

Figure 3.11: Average RMSE Values of Reconstructions

yields higher reflection magnitudes relative to the noise. As the bandwidth increases, the three outputs become more similar. Figure 3.39 is more straightforward than the spherical resolution comparison. The semi-blind scene output has a clear concentration of energy in the first .1", while the other three techniques are all pretty similar using the cylindrical target.

3.4 Conclusions

We developed and tested a semi-blind deconvolution algorithm on one-dimensional sonar data using high-frequency chirp transmit signals. We showed how this method can recover bandwidth with the scene estimate as well as generate a transmit estimate that is more accurate for deconvolution-based processing. We tested this algorithm on tank data and showed resolution improvement over regularized restoration with the analytic description as well as the recorded replica. One of the weaknesses of this method is in the time of computation. Its implementation requires an iterative process nested in another iterative process. This makes it ill-suited for situations where time is the priority. This has some flexibility, however, as it can be run once (or averaged over several times) in advance to find

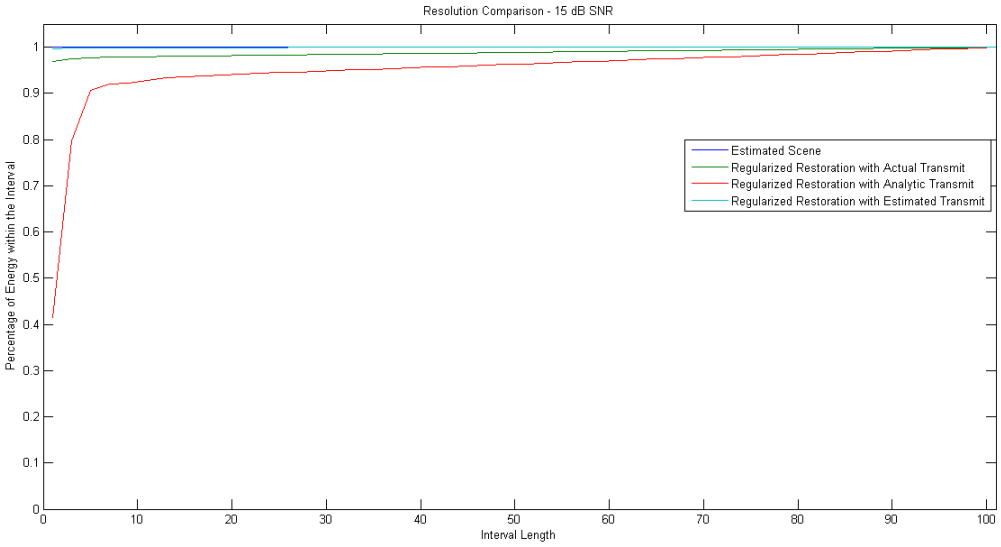

Figure 3.12: Average Energy Concentration in Increasing Intervals from the Center, 15 dB SNR

an accurate estimate of the transmit description. This transmit estimate can then be used in place of a recorded replica or analytic signal in later processing without sacrificing any additional time after the initial processing. This model is also limited in the type of signal degradation it can correct for. Using the $||x - x_0||_2$ term in the reconstruction equation is computationally efficient, but it will severely penalize large deviations from the transmit description. In particular, this will make it difficult for the new transmit estimate to add or remove significant frequency content outside of the given description. This makes this method poorly suited for dealing with Doppler shifts, resonances, or aliasing effects that add or remove frequencies from the in-water transmit. With specific knowledge of the transmit signal degradation in the system, a more specific model could be designed and implemented in a similar framework to improve the results.

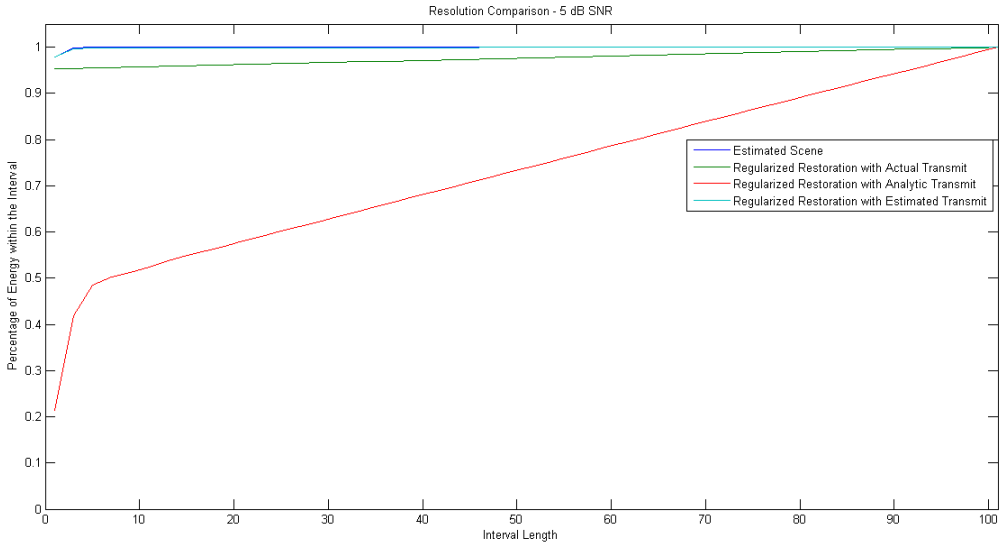

Figure 3.13: Average Energy Concentration in Increasing Intervals from the Center, 5 dB SNR

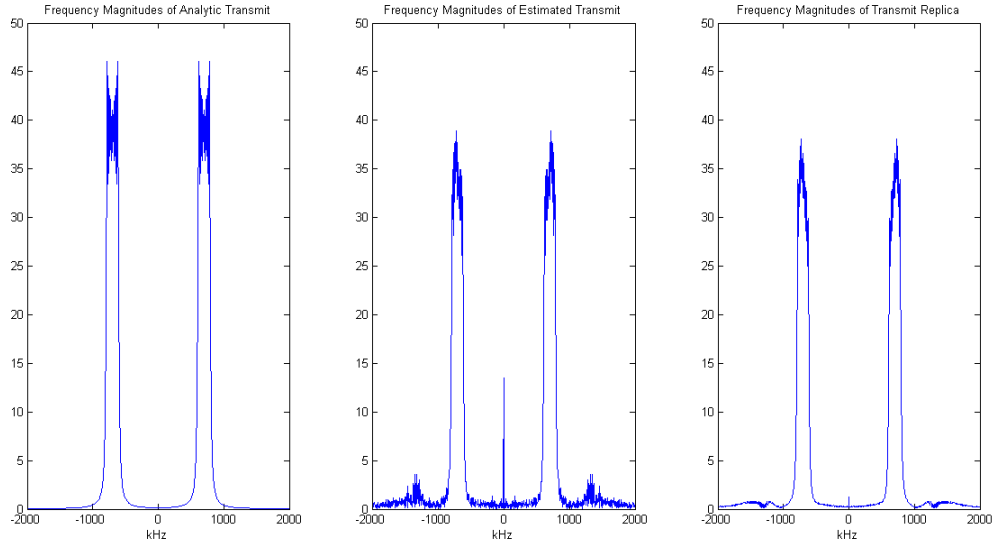

Figure 3.14: 600-800kHz Transmit, Estimate using Spherical Target

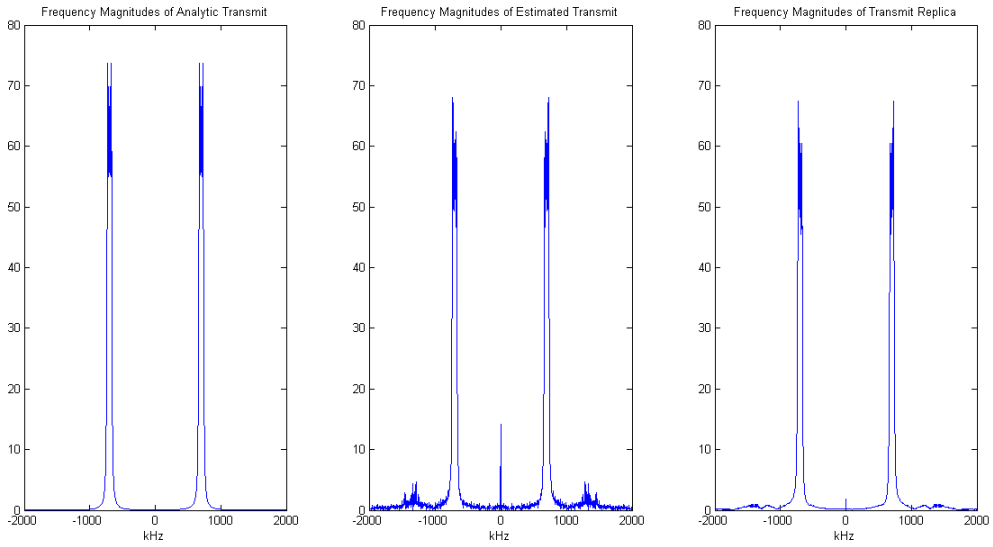

Figure 3.15: 660-740kHz Transmit, Estimate using Spherical Target

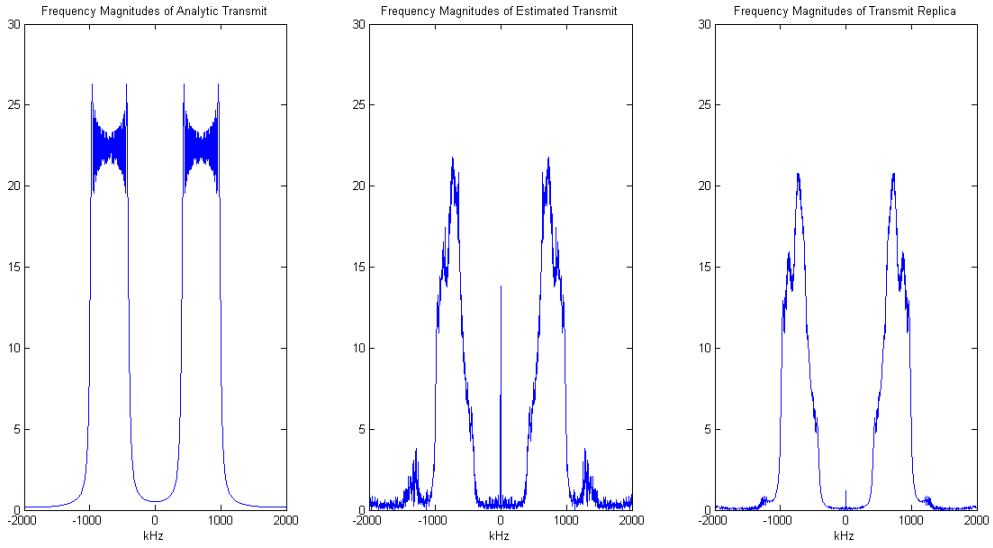

Figure 3.16: 400-1000kHz Transmit, Estimate using Spherical Target

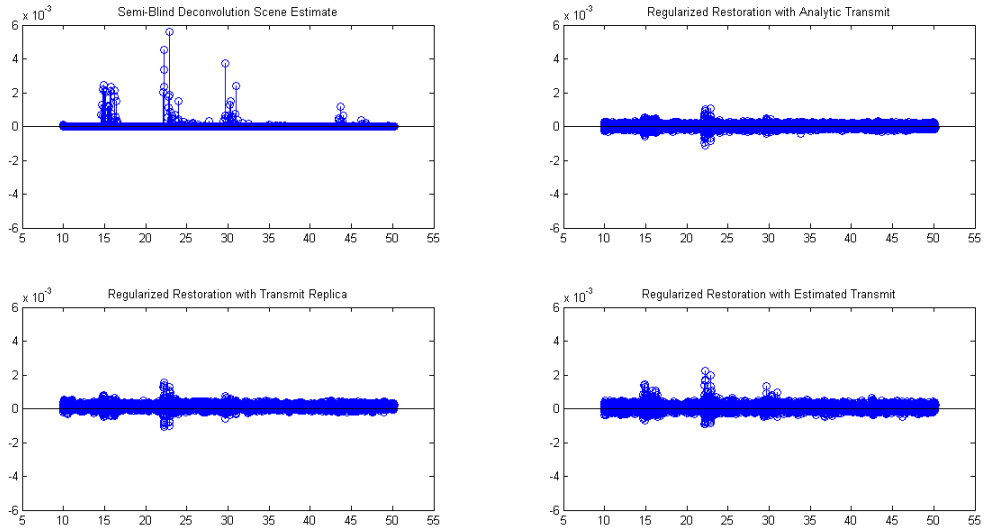

Figure 3.17: 600-800kHz Transmit, Spherical Target

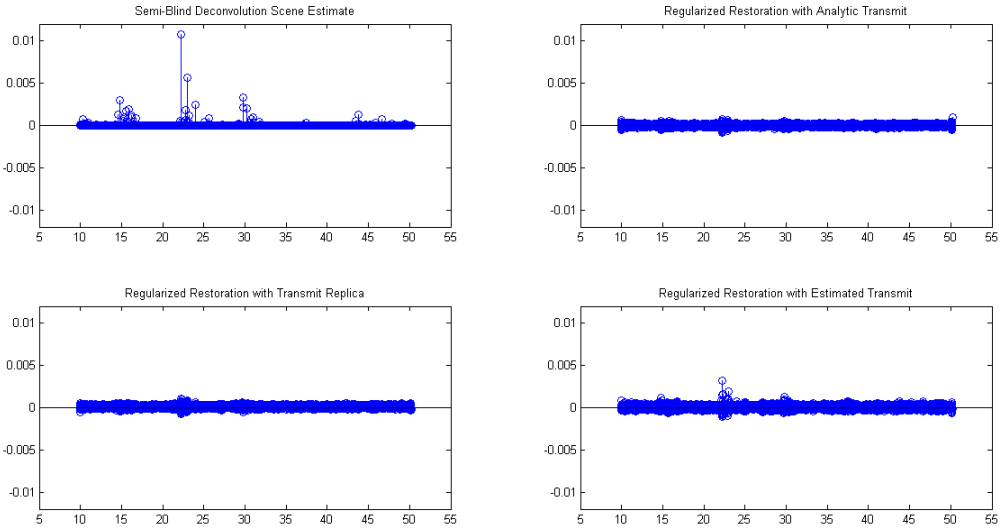

Figure 3.18: 660-740kHz Transmit, Spherical Target

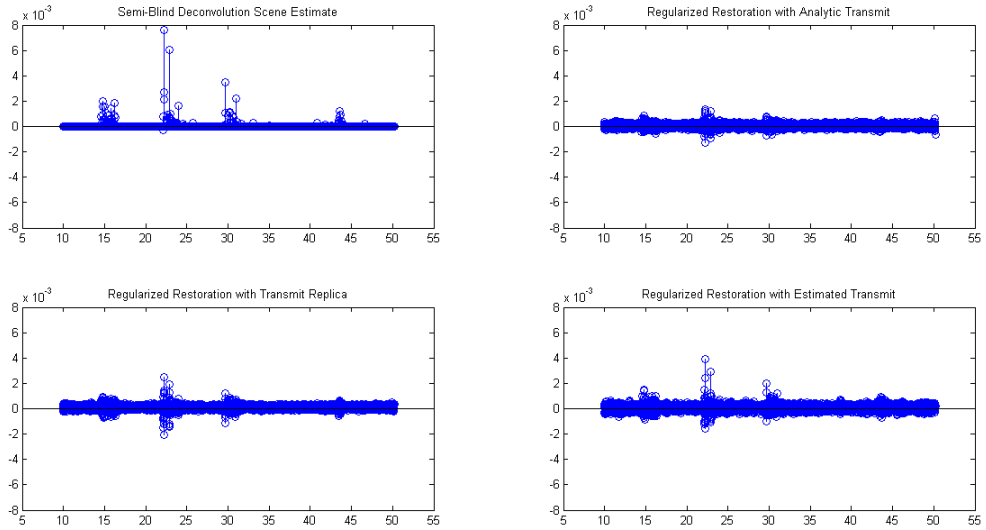

Figure 3.19: 400-1000kHz Transmit, Spherical Target

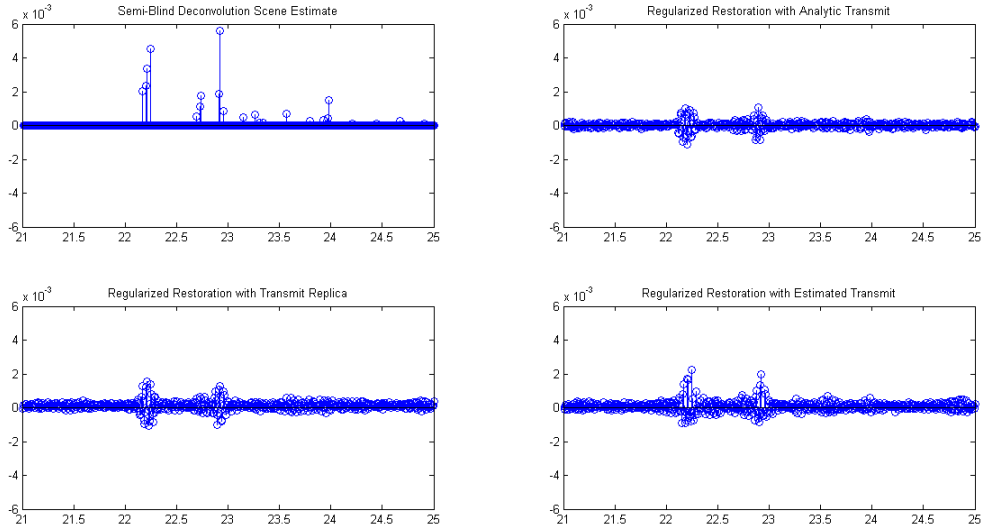

Figure 3.20: 600-800kHz Transmit, Spherical Target, Zoomed on Target

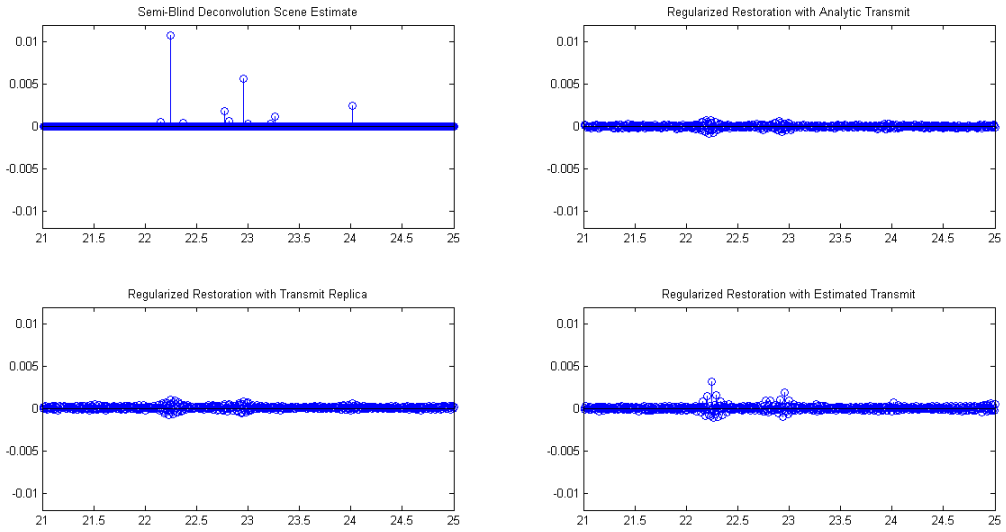

Figure 3.21: 660-740kHz Transmit, Spherical Target, Zoomed on Target

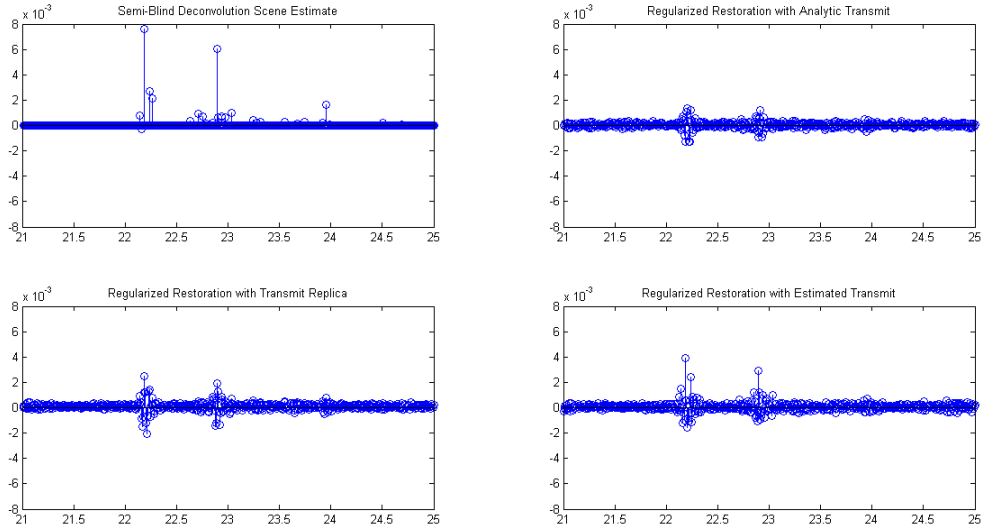

Figure 3.22: 400-1000kHz Transmit, Spherical Target, Zoomed on Target

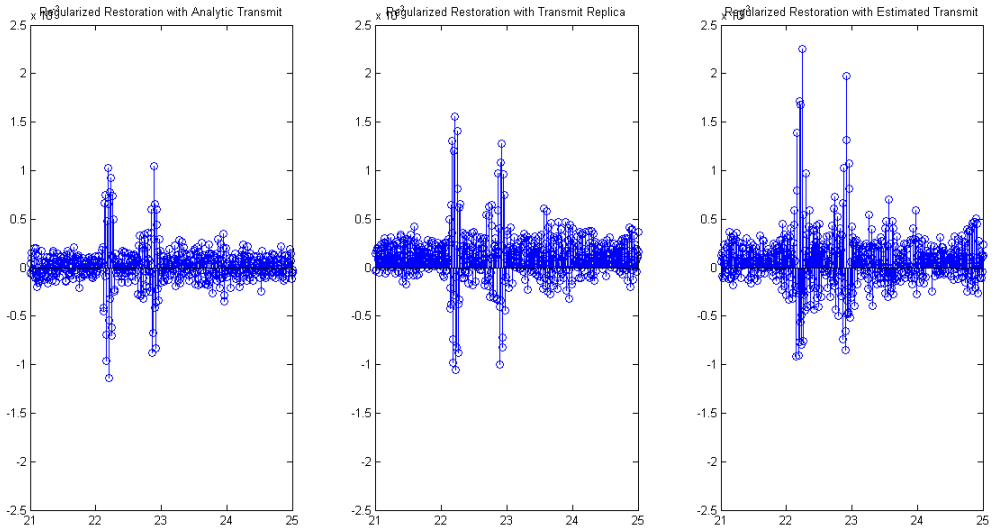

Figure 3.23: 600-800kHz Transmit, Spherical Target, Zoomed on Target, Regularized Restoration Comparison

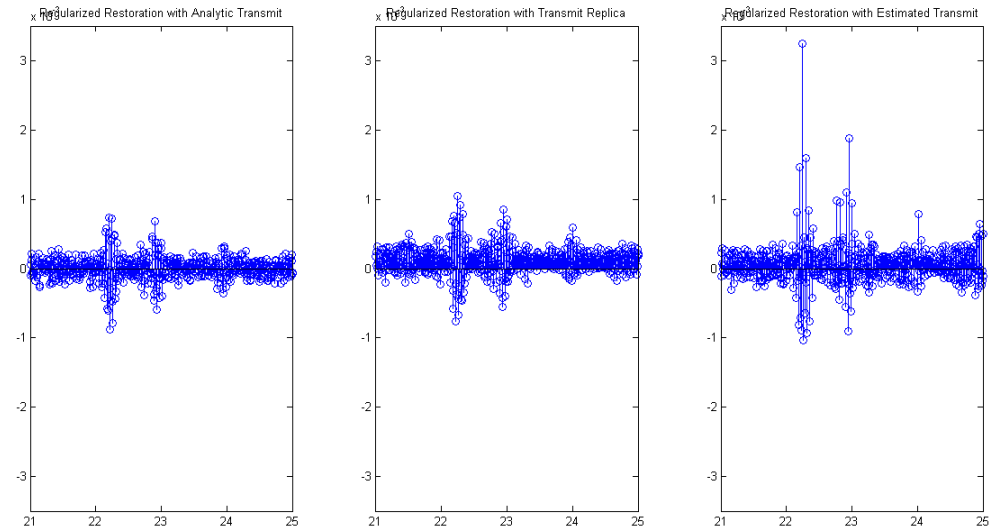

Figure 3.24: 660-740kHz Transmit, Spherical Target, Zoomed on Target, Regularized Restoration Comparison

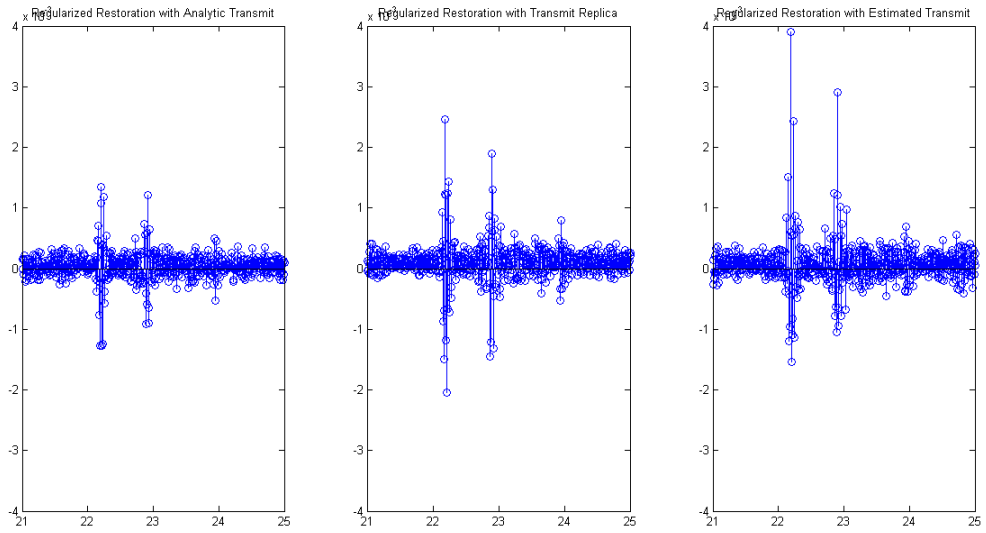

Figure 3.25: 400-1000kHz Transmit, Spherical Target, Zoomed on Target, Regularized Restoration Comparison

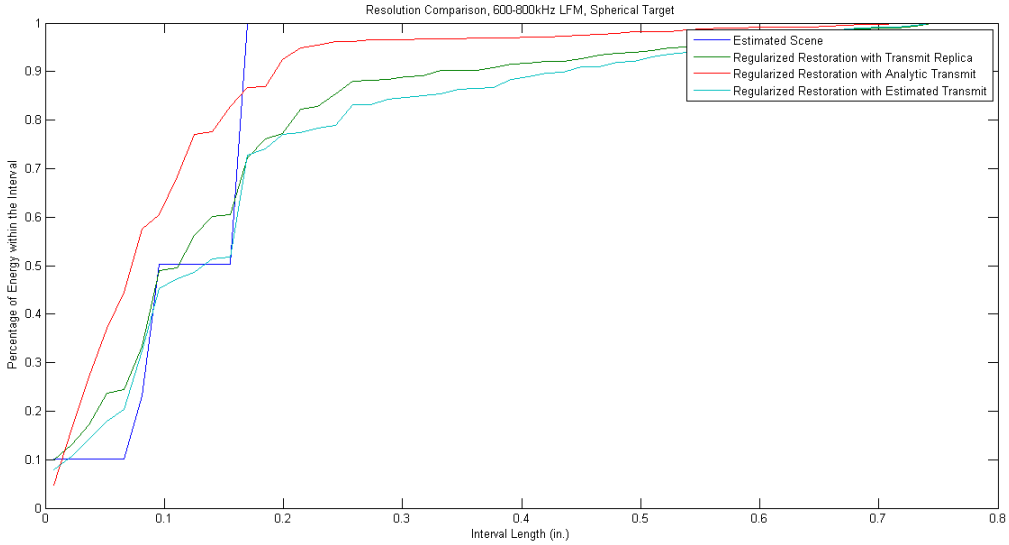

Figure 3.26: Resolution Comparison, 600-800kHz Transmit, Spherical Target

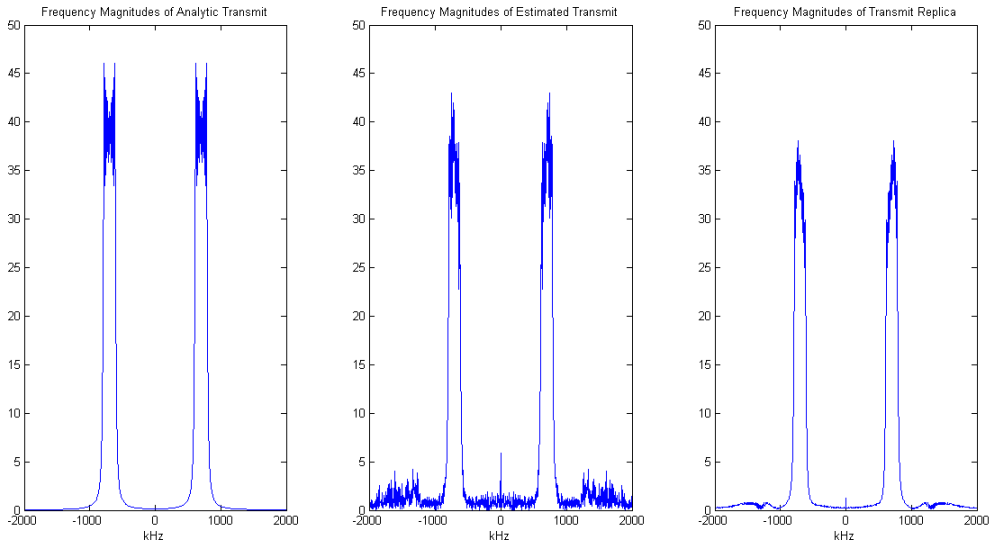

Figure 3.27: 600-800kHz Transmit, Estimate using Cylindrical Target

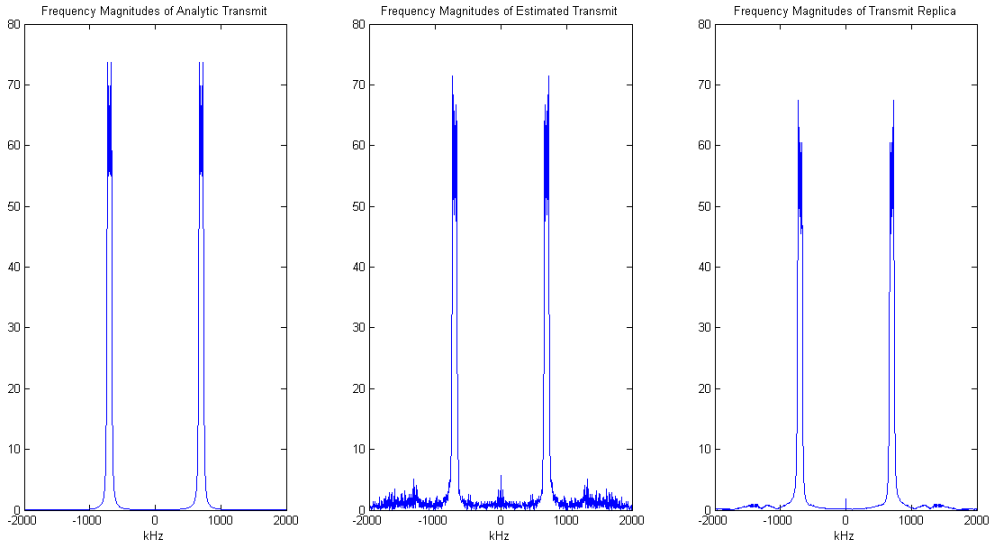

Figure 3.28: 660-740kHz Transmit, Estimate using Cylindrical Target

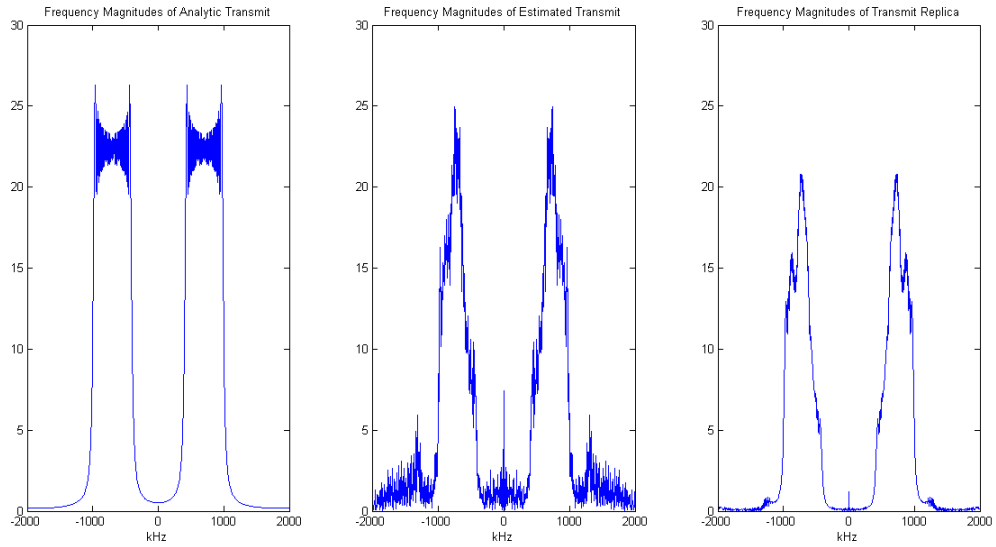

Figure 3.29: 400-1000kHz Transmit, Estimate using Cylindrical Target

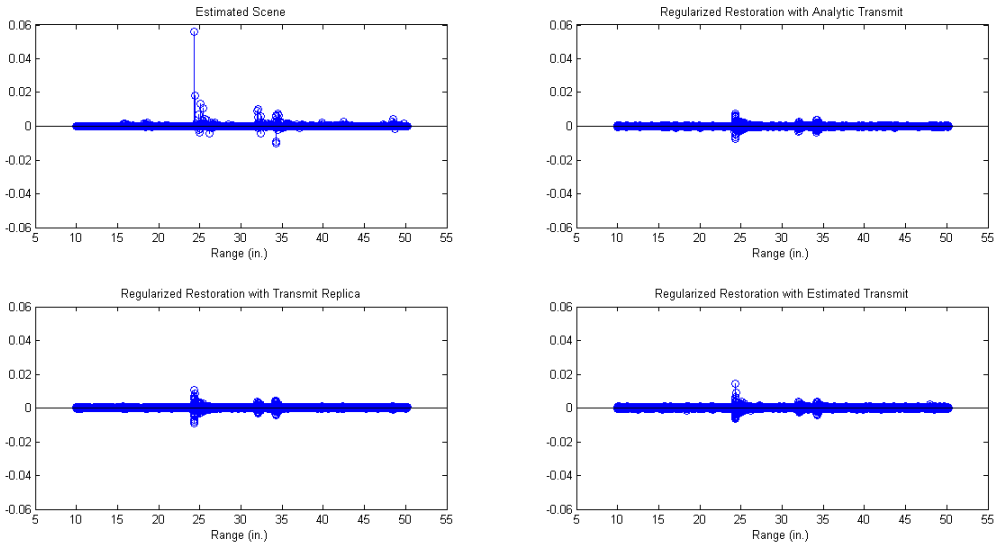

Figure 3.30: 600-800kHz Transmit, Cylindrical Target

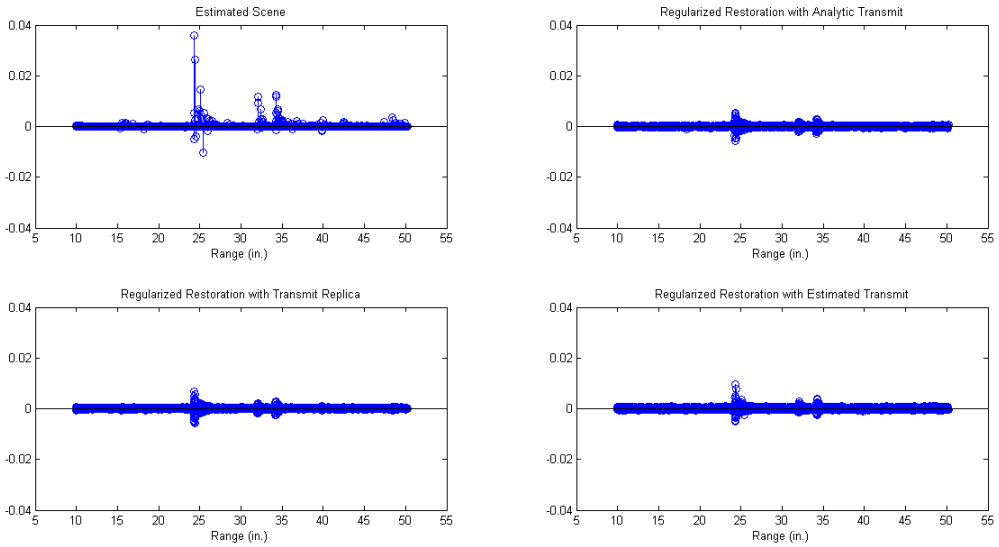

Figure 3.31: 660-740kHz Transmit, Cylindrical Target

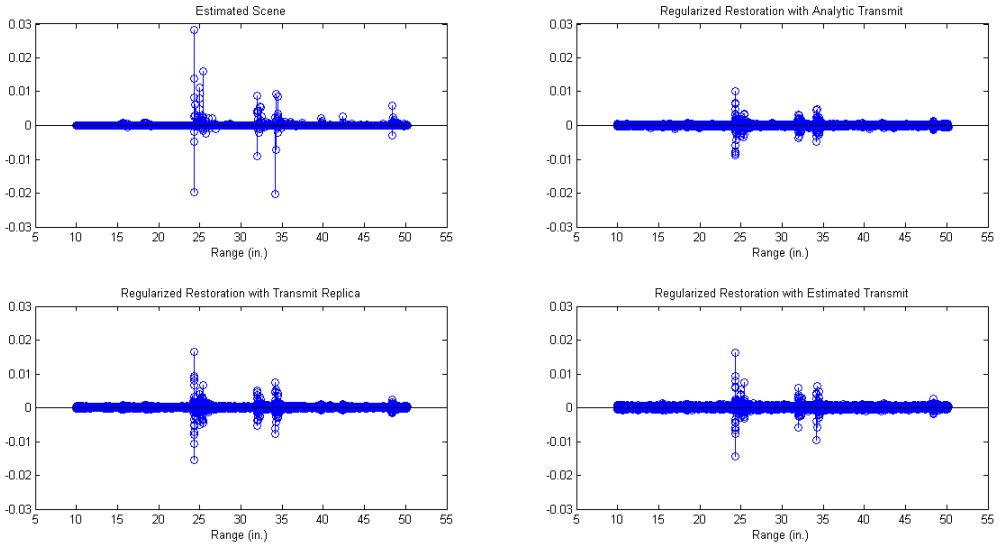

Figure 3.32: 400-1000kHz Transmit, Cylindrical Target

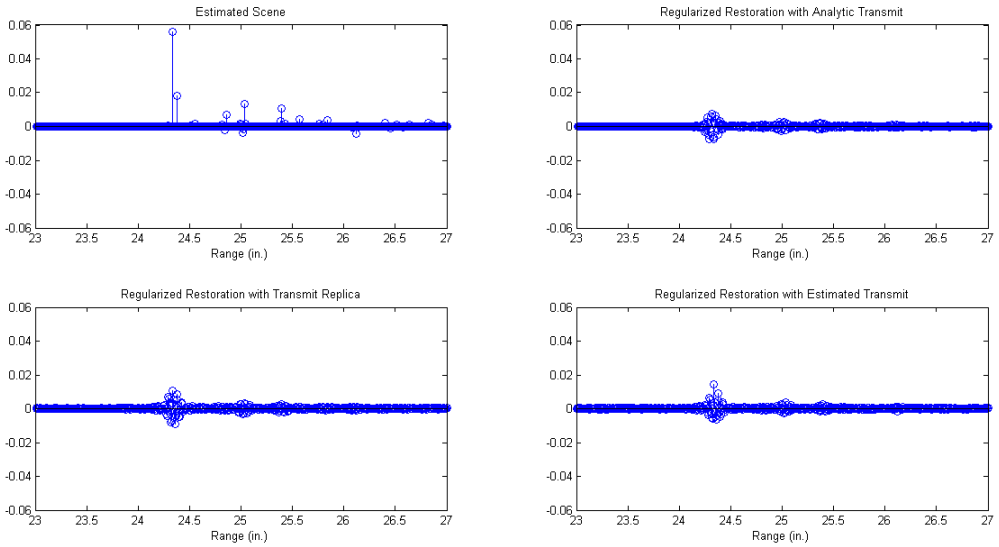

Figure 3.33: 600-800kHz Transmit, Cylindrical Target, Zoomed on Target

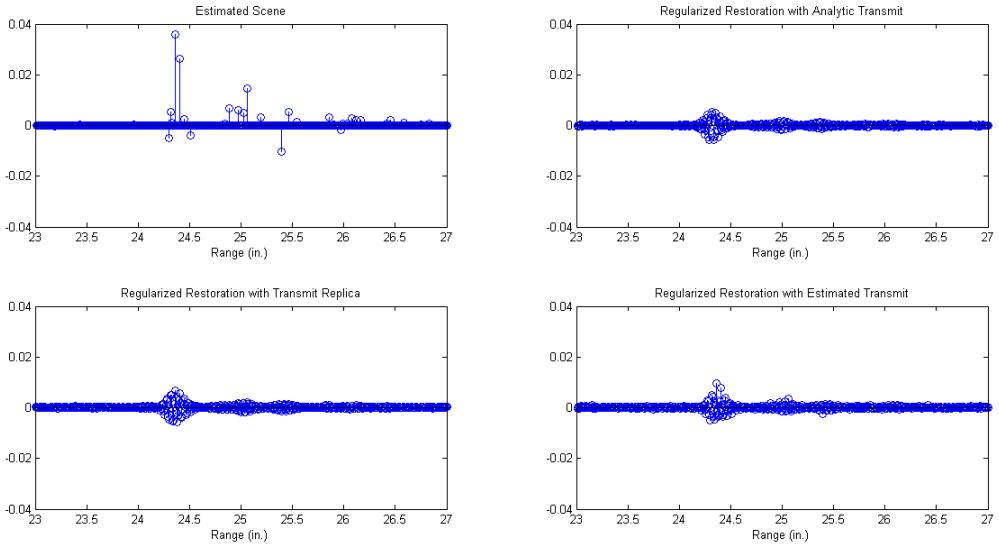

Figure 3.34: 660-740kHz Transmit, Cylindrical Target, Zoomed on Target

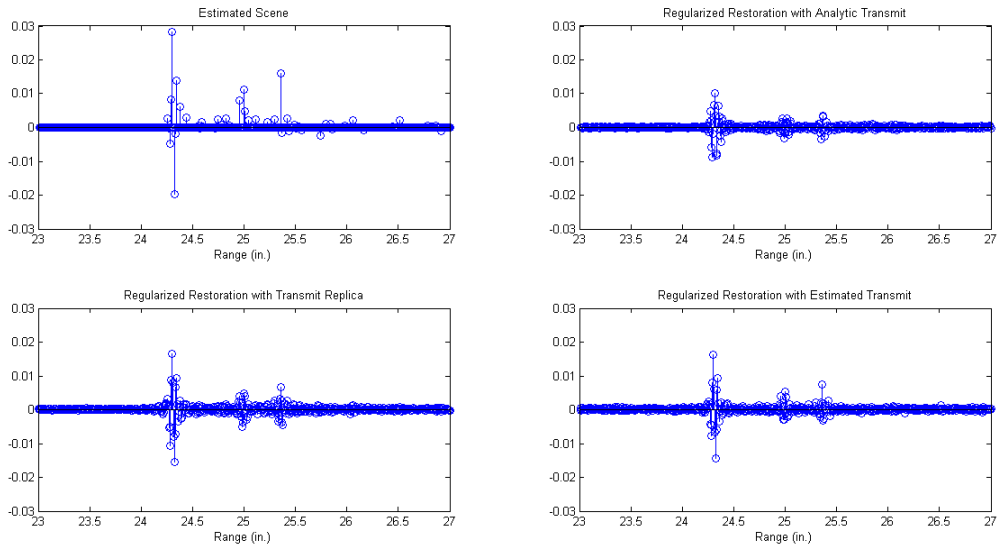

Figure 3.35: 400-1000kHz Transmit, Cylindrical Target, Zoomed on Target

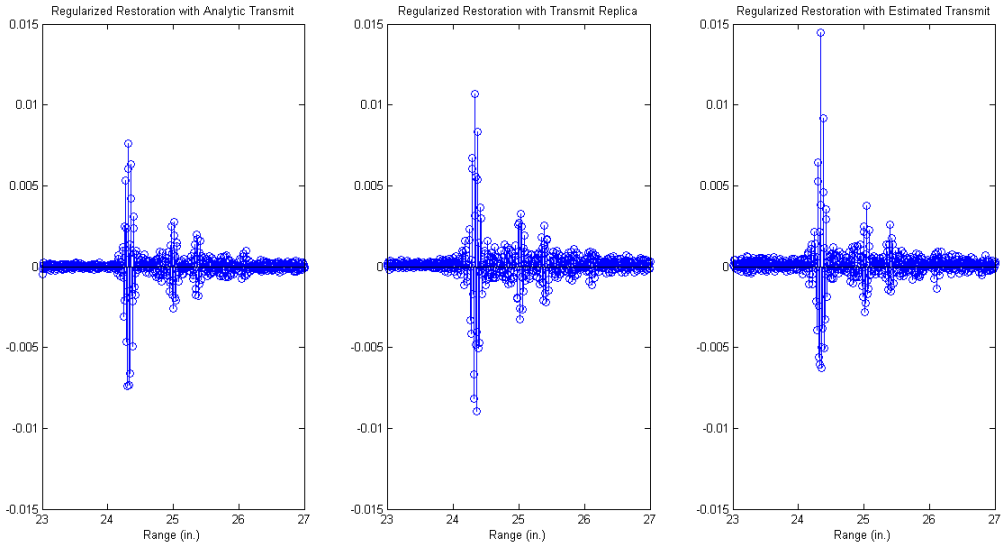

Figure 3.36: 600-800kHz Transmit, Cylindrical Target, Zoomed on Target, Regularized Restoration Comparison

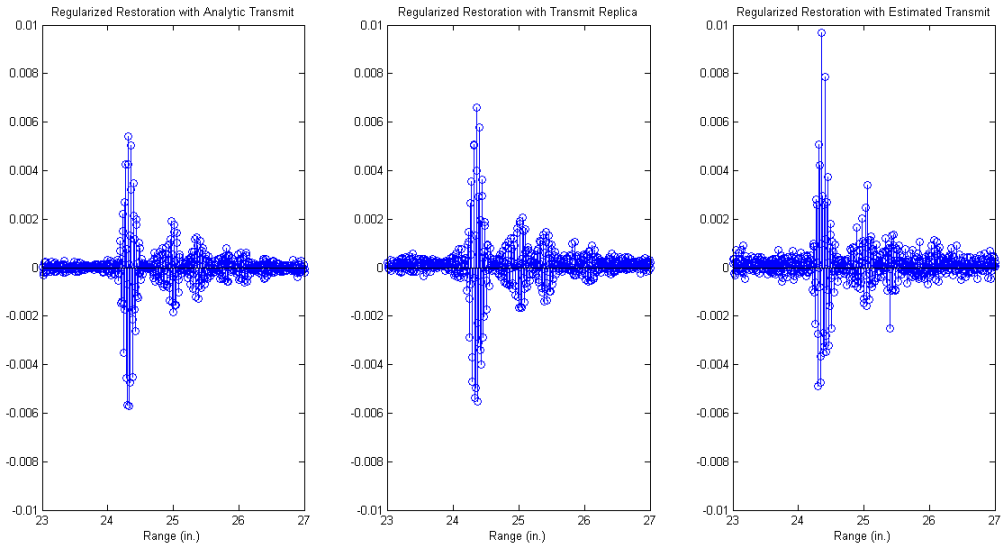

Figure 3.37: 660-740kHz Transmit, Cylindrical Target, Zoomed on Target, Regularized Restoration Comparison

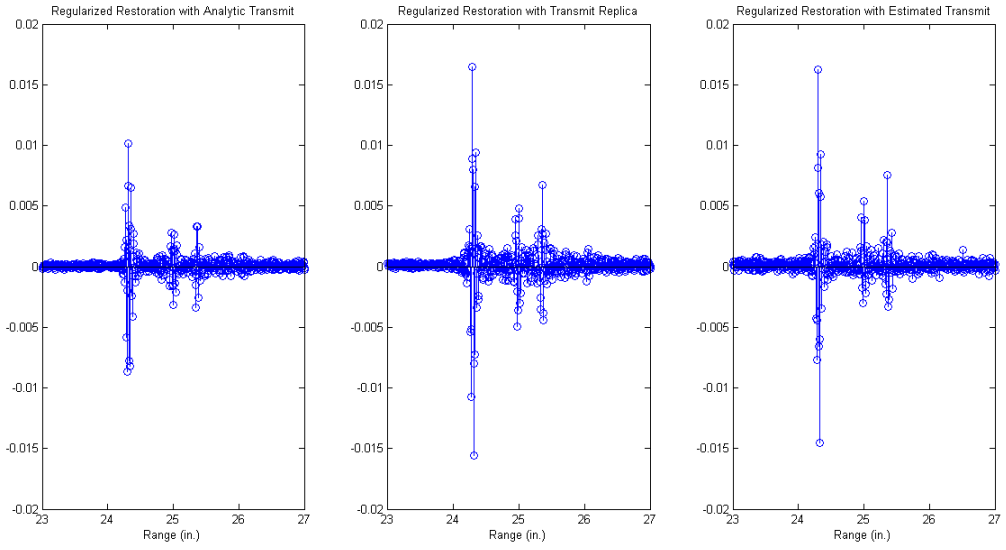

Figure 3.38: 400-1000kHz Transmit, Cylindrical Target, Zoomed on Target, Regularized Restoration Comparison

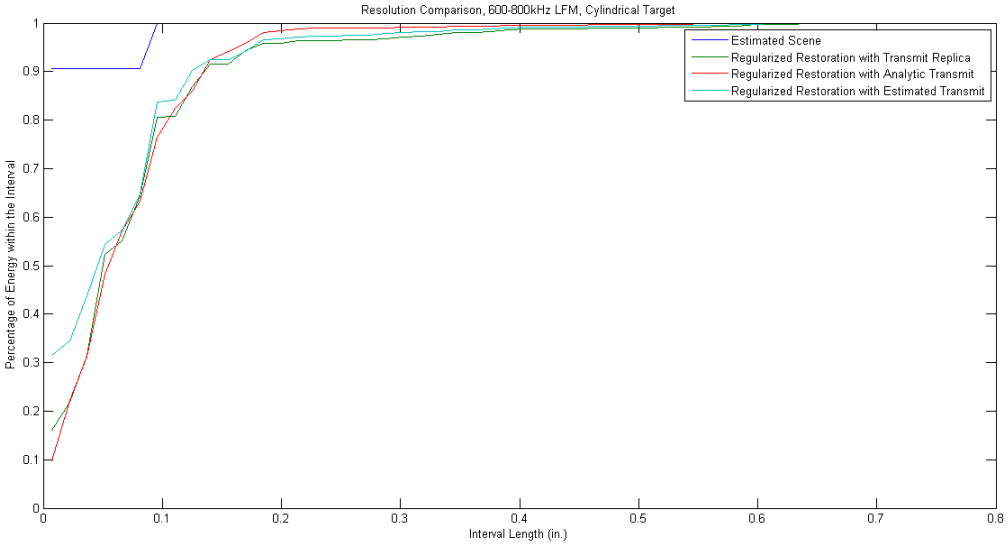

Figure 3.39: Resolution Comparison, 600-800kHz Transmit, Cylindrical Target

Chapter 4

Distance-Based Reconstruction

4.1 Introduction

In any sonar application where information is desired over a range of distances, a single replica will be insufficient to properly process using deconvolution techniques. If the primary processing is matched filtering, then this is likely not a problem. This is because matched filtering can be represented by a single multiplication in frequency. Thus, in the simple case of spreading attenuation, a scale factor on the transmit signal just yields a scale factor on the scene reconstruction. A simple scaling on the reconstructed scene can then compensate for the attenuation over distance of the transmit signal. Even when the degradation is more severe, matched filtering is a very robust method that is only slightly altered by many of the more severe changes that can occur. Matched filtering is inherently limited in the quality of reconstruction, however, as it is strictly bandlimited and distorts the reconstructed frequency content. More information can potentially be obtained by using a deconvolution-based technique, which attempts to restore the frequency content more accurately. While still bandlimited, the effective bandwidth of the filter is wider than that of matched filtering, so more information is retained. The increased bandwidth comes at the cost of increased noise sensitivity, as reconstruction is attempted on lower strength frequency content. These frequency bands can have signal strength much closer to the level of noise in the return signal, which affects the quality of reconstruction.

We will consider the specific method of deconvolution known as regularized restoration. This technique has a term for deconvolution as well as a regularization term which can reduce reconstruction in very low strength frequency bands. In this case, even a simple change in scaling due to spreading can have a larger impact on the reconstructed scene. A

scale factor on the transmit is equivalent to both a scale factor on the output as well as a change in the regularization parameter. This means that even a very basic inaccuracy in the transmit signal can have a significant impact in the reconstructed frequencies of the scene. Frequency-dependent degradation can have an even greater impact on the reconstruction, as the transmit description more directly impacts the reconstructed frequencies of the scene. Changing the transmitted frequencies (such as in the case of Doppler, resonance, etc.) can greatly alter the reconstruction. Having an accurate transmit description can greatly decrease the level of error in the reconstruction using this method.

We chose a method to adaptively estimate the transmit signal at multiple distances using a semi-blind deconvolution method. We begin with a description of the transmit at or near the source. This can either be the analytic description used to generate the transmit or a recorded replica. We then use that description, along with a portion of the return data nearest the source, to generate estimates on both the scene and transmit description in this region. We then move along to a farther distance and repeat the process using the updated transmit estimate from the previous step. This allows us to obtain estimates of the transmit signal at a variety of distances based on the return data. We can then use the transmit signals to rebuild the scene at a variety of distances and combine the results for a more accurate total reconstruction. As an added benefit, this method also outputs sparse representations of the scene at the same distances.

In many sonar applications, the transmit signal is unreliable for accurate processing when using deconvolution-based processing. An unreliable transmit can result from a lack of recorded replica, inaccurately recorded replica, or additional signal degradation before transmission or through the water. When using a matched filter to process, the lack of an accurate replica is relatively unimportant. Matched filtering is robust to noise in both the transmit and recorded returns, so the lack of accuracy creates very little error in the output. When using regularized restoration, however, inaccuracies in the transmit description can create significant errors in the output.

Instead, we seek to use the recorded data to simultaneously estimate the in-water transmit and interrogated scene information. We begin by assuming a sparse representation of the interrogated scene. Sparse reconstruction techniques have gained popularity in recent years for a large range of applications. An overview of some of these methods for image processing and radar can be found in [7] and [8], respectively. The applications are varied, but include through-the-wall radar imaging [10], synthetic aperture radar imaging [12], 3D sonar imaging [13], and synthetic aperture imaging applications [14, 15, 16]. In addition to the sparsity assumption, we will also require either an analytic transmit description or a recorded replica. This will create an additional term in the reconstruction equation corresponding to transmit consistency. This is justified because the in-water signal should bear some resemblance to the given description (either analytic or recorded replica), or none of the processing would be successful. This makes it a problem of semi-blind deconvolution. Blind deconvolution is the process of simultaneously estimating two unknown, convolved signals. This case is semi-blind because we have a description of one of the signals, which bounds the range of credible solutions for that signal. By using this technique we can obtain a more accurate representation of the interrogated scene, since the transmit is allowed to change to a more accurate description during the processing. We also obtain a final representation of the in-water signal that can then be used in place of a replica for future processing. This can modify the processing by either continuing to use the semi-blind deconvolution formula with a better initial estimate, or by using faster processing methods with the new, more accurate transmit description. This allows for a great degree of control over the required processing time and level of accuracy.

Blind deconvolution techniques have been researched primarily in the contexts of image processing [26, 27, 28, 29] and communications [34, 35, 36]. They have also seen some use in sonar applications, primarily in passively detecting unknown signals collected from an unknown scene. Our approach is semi-bind deconvolution because we begin with a description

of one of the signals but believe it to be distorted in some way from the signal that is transmitted. This gives us more information to work with than traditional blind deconvolution algorithms, so our approach restricts the solutions to realistic estimates.

Note that the simplest approach to solve this problem could be to record accurate replicas with each data collection. This can be problematic, as it may require a change in hardware and possibly the data acquisition process. This is also not a perfect solution, as noise in the recording, hydrophone inaccuracies, and environmental changes can still create errors in the replica which can be reflected in deconvolution-based processing. In addition, these changes can vary with distance traveled as the transmitted signal propagates out to the scene and then is reflected back to the receiver.

4.2 Reconstruction Algorithm

We assume a high-frequency, discretely sampled sonar model with a modified transmit signal

$$r = T(x + n_1) * s + n_2 \tag{4.1}$$

where x is the analytic transmit signal, T is some distance-dependent transformation on x (such as Doppler, dispersive medium, absorption effects, etc.), $*$ represents linear convolution, s is the interrogated scene, n_1 and n_2 are noise processes, and r is the recorded return. Even in the simple case of $r = x * s$, this problem is ill-posed when assuming both x and s are unknown. In the frequency domain, this is equivalent to solving $X(k) * S(k) = R(k)$ for $1 < k \leq M$, when only R is known. This yields M equations, each with an infinite number of solution pairs. We restrict our solutions by assuming that s has a sparse representation in some basis and that the transmit signal x is a degraded form of a given transmit description x_0.

We then divide up the return and scene information into m pieces such that

$$r_i = T_i(x + n_{1,i}) * s_i + n_{2,i} \tag{4.2}$$

We do this because we are assuming that a portion of the transformation T is dependent on distance. We seek to estimate the transformed transmit at each successive distance, thus improving the reconstruction as the signal degrades further.

We begin with a basic sparse reconstruction equation to promote a sparse scene output with a known transmit signal.

$$\hat{s_1} = \underset{s_1 \in \mathbb{R}^n}{\arg\min} \{||X s_1 - r_1||_2^2 + \alpha ||A s_1||_1\} \tag{4.3}$$

This equation can be viewed as the computationally tractable convex relaxation of the underlying sparsifying ℓ_0-norm deconvolution problem. We maintain this underlying sparsity model in the scene so that we can continue to reconstruct s_i under the same sparsity assumption, but add a term to maintain consistency with the given transmit description x_0. This gives the equation

$$(\hat{x_i}, \hat{s_i}) = \underset{x_i \in \mathbb{R}^m, s_i \in \mathbb{R}^n}{\arg\min} \{||x_i * s_i - r_i||_2^2 + \alpha ||A s_i||_1 + \beta ||x_i - x_{i-1}||_2^2\} \tag{4.4}$$

Thus, the first iteration begins at $i = 1$, and uses the given transmit description x_0 as the transmit consistency term. Ideally, this should either be defined as a replica recorded within the distance range defined by r_1 or the analytic transmit description. The scene estimate will maintain consistency with the convolution model and also promote sparsity in the A basis based on parameter selection α. The transmit estimate $\hat{x_i}$ will also maintain consistency with the convolution model and retain similarities with x_{i-1} based on parameter selection β. We change the transmit consistency with distance because we assume that the transmit is degrading further with distance. Thus, if we find an estimate for the transmit at one distance, it will be closer to the estimate at the next distance than the original transmit description. The use of the ℓ_2 norm for transmit consistency with x_{i-1} is primarily for speed of computation. With specific knowledge of the way that the transmit signal has degraded,

a better predictor model could be implemented instead of ℓ_2-norm consistency with a given description.

We use a gradient method to find the minimizers $(\hat{x}_i, \hat{s}_i)$ to (4.4). The gradient in s_i is the same as that of (4.3), since the transmit consistency term does not depend on s_i. There are a variety of ways to approximate the solution to this sparse model. We find $\hat{s}_i$ by using a majorization approximation method outlined in [23]. This reduces the problem to a sequence of quadratic problems with a closed-form solution at each iteration. This allows us to compute $\hat{s}_i$ quickly, which is particularly important since it needs to be computed many times while updating $\hat{x}_i$.

To find $\hat{x}_i$, we need to take the gradient in x_i. The ℓ_1 sparsity term for s_i is eliminated and the problem is greatly simplified. The estimate $\hat{x}_i$ can be computed directly with Fourier transforms and is defined as

$$\hat{x}_i = \mathcal{F}^{-1}\left\{\frac{\overline{S_i(\omega)}R_i(\omega) + \beta X_{i-1}(\omega)}{|S_i(\omega)|^2 + \beta}\right\} \tag{4.5}$$

where capital letters denote Fourier transforms of their lowercase counterparts.

The algorithm begins with initial estimates for x_i and s_i, then alternately updates each estimate with every iteration until the minimum is achieved. We can exploit the speed of calculating the gradient in x_i to significantly speed up the computation of the entire method. By running relatively few iterations in the s_i majorization algorithm, then finding the direct minimum in x_i, and alternating between these, we can more rapidly approach the global minimum without adding significant processing time required of a large amount of iterations in s_i.

We complete this process sequentially for all values of i in order to obtain scene and transmit estimates at each desired distance.

We also use regularized restoration as a way to find an initial scene estimate, to compare with the semi-blind scene estimate, and to show the improvements possible with using the

updated transmit signal for additional processing. This is given by

$$\min_{s_i \in \mathbb{R}^n} \{||X_i s_i - r_i||_2^2 + \gamma ||B s_i||_2^2\} \tag{4.6}$$

where B is a regularization filter. Note that this is very similar to (3.2), but we replace the ℓ_1 norm with an ℓ_2 norm. We can solve this directly with Fourier transforms, which gives the following result:

$$\hat{s}_{i\ell_2} = \mathcal{F}^{-1}\Big\{ \frac{\overline{X_i(\omega)} R_i(\omega)}{|X_i(\omega)|^2 + \gamma |B(\omega)|^2} \Big\} \tag{4.7}$$

where capital letters denote the Fourier transforms of the lowercase counterparts. Observe that when $\gamma = 0$ this is an inverse filter. When $B = I$ and γ is large, this is equivalent to matched filtering with a constant scale factor.

4.2.1 Regularized Restoration Amplitude Correction

At the most basic level, this algorithm should compensate for amplitude changes (such as from attenuation) from the source. While the impact on matched filtering would simply be a scale factor, this change can have more pronounced effects on the results of regularized restoration processing. We begin with a simplified form of (4.2), where the transformation is simply a constant scale factor. We consider correcting for the scaling in the regularized restoration by scaling X by a constant α inside the regularized restoration formula.

$$\begin{aligned}
\hat{s}(t) &= \mathcal{F}^{-1}\Big\{ \frac{\alpha X^*(\omega) R(\omega)}{|\alpha X(\omega)|^2 + \gamma} \Big\} \\
&= \mathcal{F}^{-1}\Big\{ \frac{\alpha^2 |X(\omega)|^2 S(\omega)}{\alpha^2 |X(\omega)|^2 + \gamma} \Big\} + \mathcal{F}^{-1}\Big\{ \frac{\alpha X^*(\omega) N(\omega)}{\alpha^2 |X(\omega)|^2 + \gamma} \Big\} \\
&= \mathcal{F}^{-1}\Big\{ \frac{|X(\omega)|^2 S(\omega)}{|X(\omega)|^2 + \frac{\gamma}{\alpha^2}} \Big\} + \frac{1}{\alpha} \mathcal{F}^{-1}\Big\{ \frac{X^*(\omega) N(\omega)}{|X(\omega)|^2 + \frac{\gamma}{\alpha^2}} \Big\}
\end{aligned} \tag{4.8}$$

The first term in (4.8) corresponds to the reconstruction of the scene. It is nearly identical to a regularized restoration without any constant scaling but now has an effective

regularization parameter of $\frac{\gamma}{\alpha^2}$ instead of γ. Since we assume that $0 < \alpha < 1$, this represents an increase in the regularization parameter. This means that at large distances (since an attenuation/absorption factor will decrease with distance), the regularization parameter will increase and the reconstruction will prioritize the frequency content that is strong in the transmit band, while more heavily regularizing outside of this band. The second term corresponds to how the noise is handled in the reconstruction. We can compare this term with the noise term of a regularized restoration reconstruction that is compensated for attenuation after the restoration:

$$\begin{aligned} r(t) &= \mathcal{F}^{-1}\{\alpha X(\omega)S(\omega) + N(\omega)\} \\ s'(t) &= \mathcal{F}^{-1}\left\{\frac{X^*(\omega)R(\omega)}{|X(\omega)|^2+\gamma}\right\} \\ &= \alpha\mathcal{F}^{-1}\left\{\frac{|X(\omega)|^2S(\omega)}{|X(\omega)|^2+\gamma}\right\} + \mathcal{F}^{-1}\left\{\frac{X^*(\omega)N(\omega)}{|X(\omega)|^2+\gamma}\right\} \\ s(t) &= \mathcal{F}^{-1}\left\{\frac{|X(\omega)|^2S(\omega)}{|X(\omega)|^2+\gamma}\right\} + \frac{1}{\alpha}\mathcal{F}^{-1}\left\{\frac{X^*(\omega)N(\omega)}{|X(\omega)|^2+\gamma}\right\} \end{aligned} \tag{4.9}$$

Thus, the noise contribution is lowered if the attenuation factor is done in the processing instead of compensating afterwards. So while preemptively correcting for distance effectively changes the regularization parameter, it also decreases the noise contribution in all cases over correcting in the reconstruction. We can make these processes identical by changing the regularization parameter in the preemptively corrected version to $\gamma' = \alpha^2\gamma$. By not changing this regularization parameter, however, we allow the regularization to increase as a function of distance in the distance-based reconstruction. This means that we decrease the noise contribution and more heavily emphasize the strongest transmit frequencies as the SNR decreases, assuming that the noise is full spectrum and the strength is independent of the signal strength.

4.2.2 Method Summary

1. Find initial estimates for both the scene and transmit at the initial distance. We can use either a transmit replica or analytic transmit description as the initial transmit estimate. We can find an initial scene estimate at the first distance using (3.6) with the known return over the the same distance r_1 and initial transmit estimate x_0.

2. Choose the sparsifying basis A for s. We chose a sparse identity for our examples, since spherical and cylindrical targets at high frequencies resemble clusters of point reflectors.

3. Choose parameters α and β, which control the trade-offs between data consistency, sparsity of As_i, and transmit consistency with x_{i-1}. We used $\alpha = 10^{-1}$ and $\beta = 10^{-3}$ for the experimental data.

4. Alternately minimize with respect to each variable for $i = 1$. The minimum in x_i is a single calculation, while the minimum in s_i requires an iterative process.

5. Repeat Step 4 for $i = 2 : M$, using new initial values $x_i = \hat{x}_{i-1}$, and compute the initial estimate for s_i using (3.6) with r_i and $\hat{x}_{i-1}$.

6. Obtain transmit estimates $\hat{x}_i$ and scene estimates $\hat{s}_i$ for all i as outputs of the algorithm. The individual estimates $\hat{s}_i$ can be combined to form an overall scene estimate s.

4.3 Results

4.3.1 Signal Change over Distance

The following results use data that was collected at the small-scale test bed (SSTB) at the Naval Surface Warfare Center in Panama City Beach, Florida. This facility uses high-frequency transmit signals along with small targets to simulate realistic conditions that correspond in a scaled manner to lower-frequency transmits and larger targets. Transmit replicas were collected using a moving linear rail at a range of distances from the receiver.

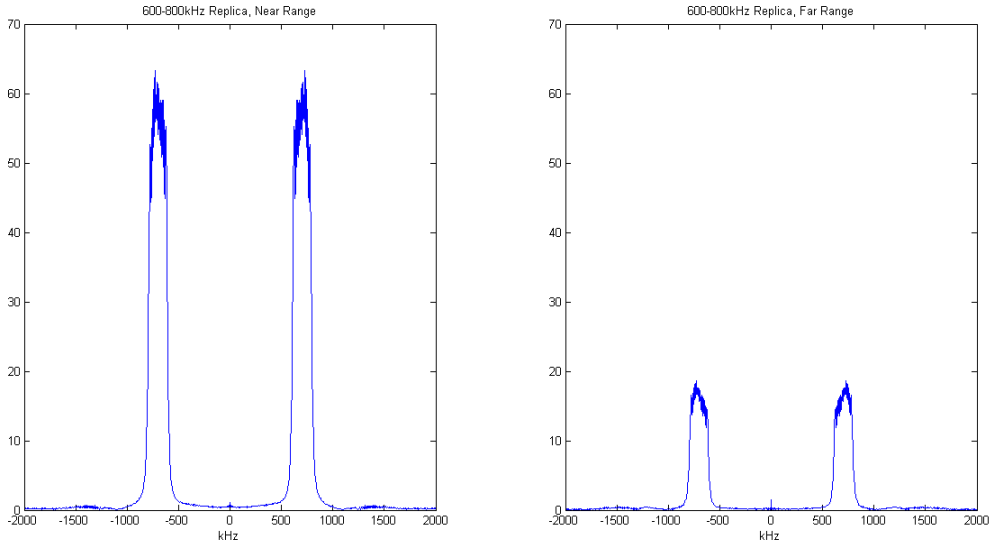

Figure 4.1: 600-800kHz Transmit, Near Recording (Left) and Far Recording (Right)

Figures 4.1-4.3 show comparisons of the recorded replicas at each frequency band at the nearest recording distance and farthest recording distance. The most noticeable difference between them is the attenuation/absorption effects that have significantly reduced the magnitudes of the original values. In order to compare them in a fair manner, Figures 4.4-4.9 use normalized versions of the transmit signal to mitigate the effects of attenuation on the reconstruction. These figures show semilog plots of the frequency values contained in the regularized restoration reconstructions using the near transmit and far transmit on the same return data. It can be difficult to accurately see the differences in these images, so Figures 4.7-4.9 show the differences between values in 4.4-4.6. From these figures, it becomes apparent that while the in-band frequencies are largely unchanged (once attenuation is compensated for), there are higher frequencies that gain in strength as the signal propagates out from the source. This could be because the far field calculations for this setup were based on the desired frequency band. If there were higher frequencies present in the transmit, the far field would be farther than initially calculated. These are represented by a continuous string of primarily negative values in the difference plots. If we examine the relative magnitude changes of both in-band frequencies and these higher frequencies, we can see in

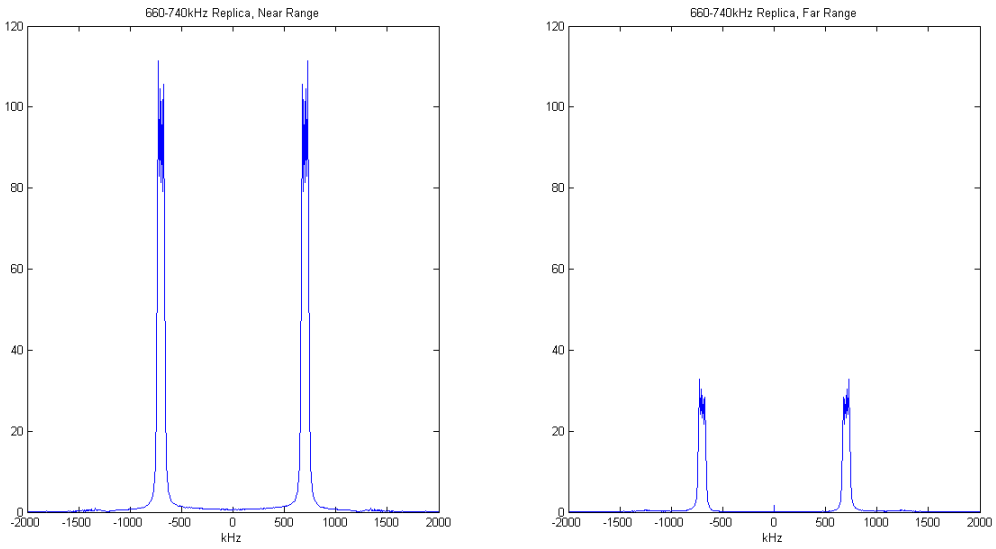

Figure 4.2: 660-740kHz Transmit, Near Recording (Left) and Far Recording (Right)

Figures 4.10-4.12 that the higher frequencies appear to grow for a distance before decaying at roughly the same rate as the in-band frequencies. Note that the higher frequencies are much lower in magnitude, and thus much more strongly impacted by the noise, which makes the values jump around significantly more than the in-band frequencies in the graphs. This behavior is a strong justification for processing signals differently at different distances. The reconstructions can be improved by making use of the additional frequency content that is in the signal at farther distances, but is largely unrecorded nearer to the source.

4.3.2 Experimental Results

The following data was also collected at the small-scale test bed (SSTB) at the Naval Surface Warfare Center in Panama City Beach, Florida. The reconstructions were performed using the distance-based semi-blind deconvolution algorithm using five divisions of the return signal. This in turn generated five different transmit estimates over different distances and five portions of a sparse scene estimate. The sparse scene estimates and regularized restoration portions were combined to form a single image of the entire interrogated range.

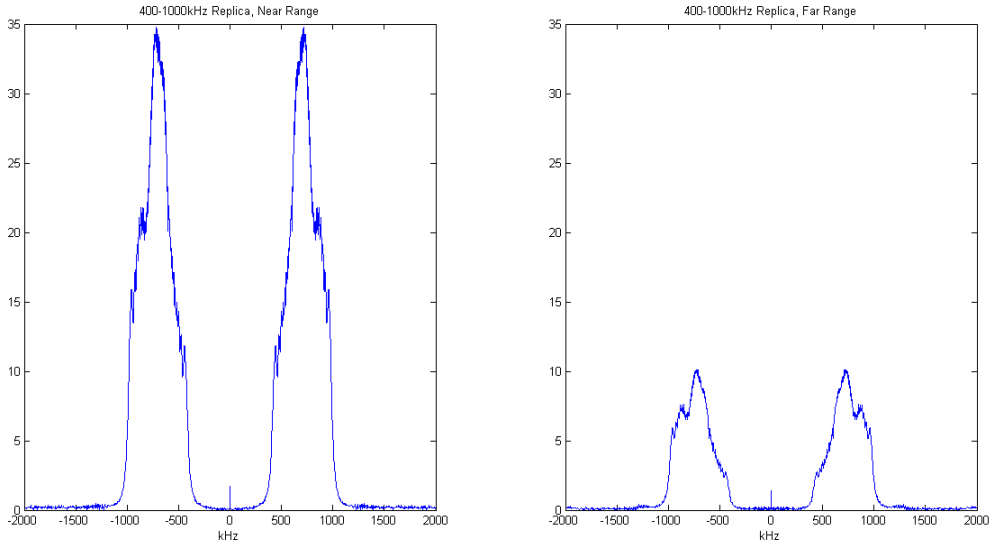

Figure 4.3: 400-1000kHz Transmit, Near Recording (Left) and Far Recording (Right)

There was no significant movement during the data collection, so Doppler was not considered in the processing. By using a high-frequency transmit, the scene should appear as point reflectors from the specular return, followed by separated additional returns from acoustic waves travelling around and through the target, with multipath returns farther out. We use a penalty that promotes a sparse identity basis since the scene should resemble clusters of point reflectors for both target types.

Figures 4.13-4.15 show the reconstructed transmits at the five different distances for each of the three transmit bands. While much of the content is similar, the DC and higher frequency content grows as the distance increases. Figures 4.16-4.36 show the full scene reconstructions (pieced together from five smaller scene reconstructions) for the full distance, while figures 4.19-4.21 show comparisons of normalized reconstructions of regularized restoration with the analytic transmit and regularized restoration with the distance estimated transmits (pieced together to make a single image over the full range). The comparison results were normalized in order to compensate for the change in magnitude of the reconstruction due to changing the scaling on the transmit estimate. The full scene includes information from

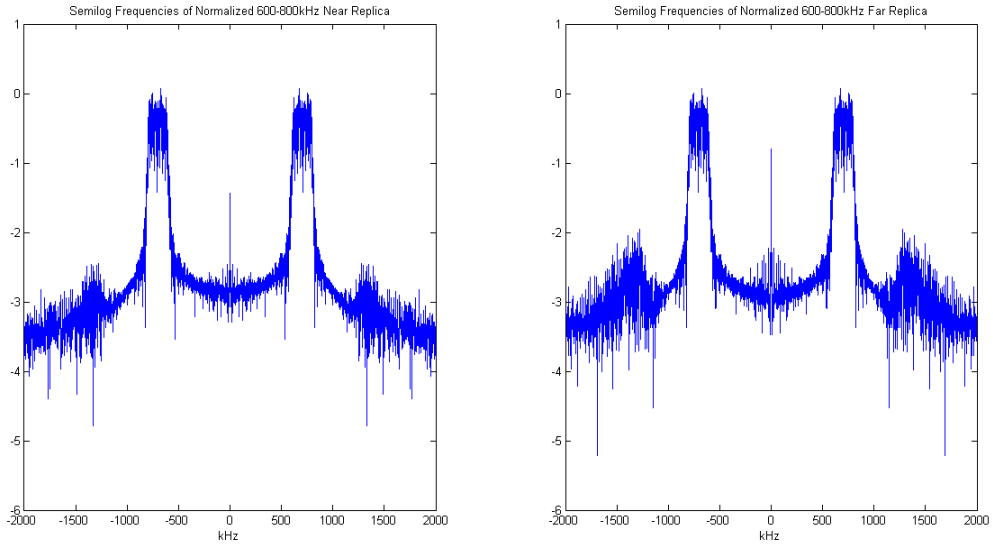

Figure 4.4: 600-800kHz Transmit, Normalized Regularized Restoration Semilog Frequency Comparison of Near Recording (Left) and Far Recording (Right)

a surface bounce, direct reflections from the target and elastic effects, and multipath reflections. Figures 4.22-4.27 show the same sparse and regularized restoration reconstructions but are focused on the target characteristics only (21”-25”).

Figures 4.28-4.30 show resolution comparisons for reconstructions from each transmit band. The sparse scene generated from the distance reconstruction has much better resolution than regularized restoration techniques at 600-800kHz and 660-740kHz transmit bands. It has a more sparse representation than the other two in the wideband case, but the largest component is located at a slightly different location, causing it to appear to have more energy spread. Regularized restoration with the estimated transmits has improved resolution over regularized restoration at 600-800kHz and 660-740kHz and is roughly the same in the 400-1000kHz example.

Figures 4.31-4.48 are similar to the previous figures but are generated using the cylindrical target data. The results in the cylindrical case are similar to those in the spherical case.

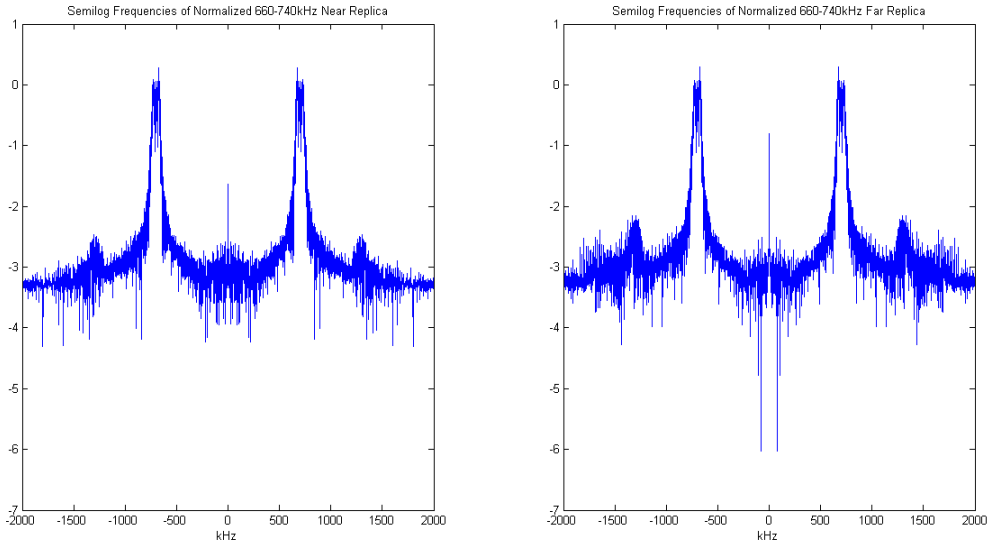

Figure 4.5: 660-740kHz Transmit, Normalized Regularized Restoration Semilog Frequency Comparison of Near Recording (Left) and Far Recording (Right)

4.4 Conclusions

We designed and implemented a distance-varying semi-blind deconvolution technique to simultaneously estimate the scene and transmit signals at various distances. We showed experimental data of replicas at various distances that supports the need for distance-dependent processing. We then tested this technique on experimental tank data and showed improvements in the resolution of the images obtained using distance-varying transmit estimates as well as an increase in the strength of the reconstructed target features relative to the noise. The algorithm is much slower than matched filtering or traditional regularized restoration, so it is more suited to post-processing or for extracting information outside of the effective range of other techniques.

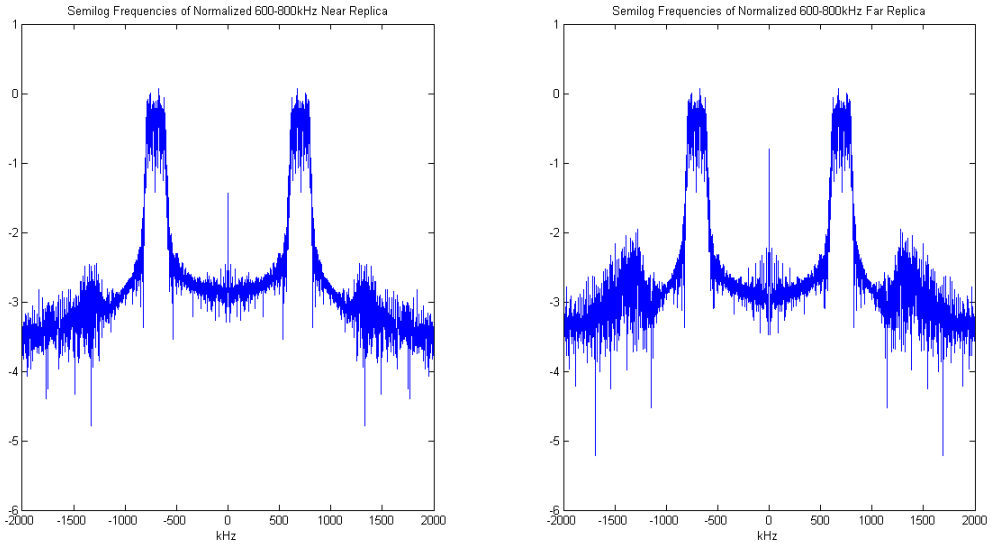

Figure 4.6: 400-1000kHz Transmit, Normalized Regularized Restoration Semilog Frequency Comparison of Near Recording (Left) and Far Recording (Right)

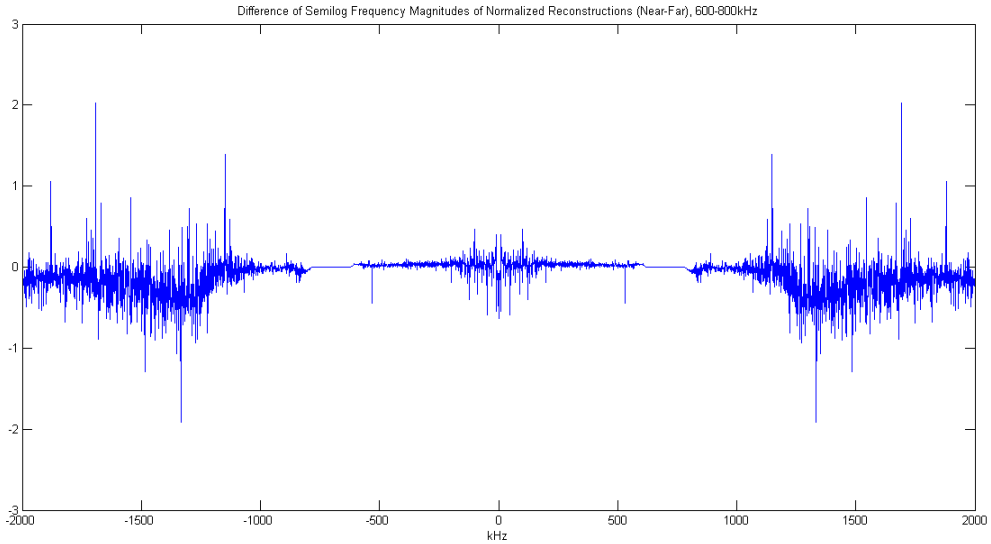

Figure 4.7: 600-800kHz Transmit, Normalized Regularized Restoration Semilog Frequency Difference of Near and Far Recordings

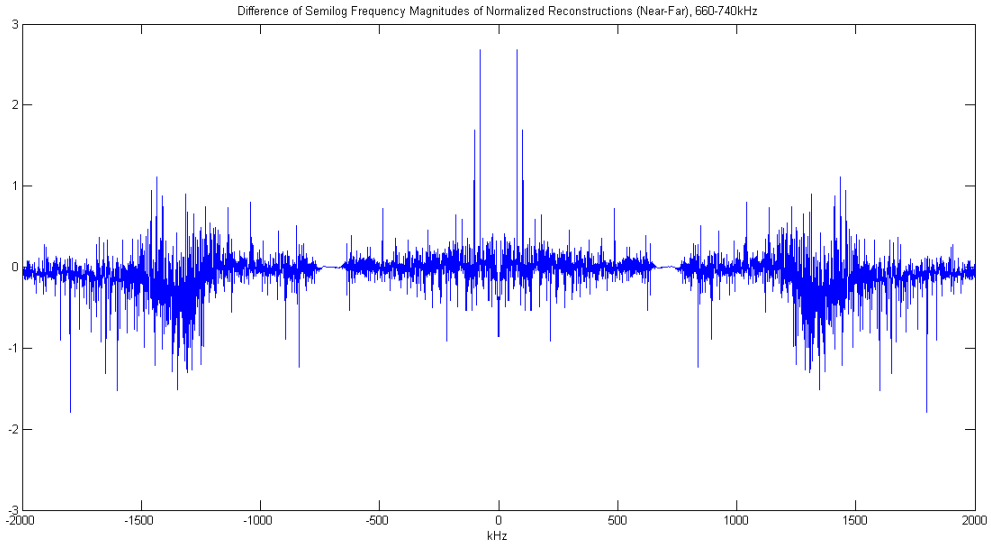

Figure 4.8: 660-740kHz Transmit, Normalized Regularized Restoration Semilog Frequency Difference of Near and Far Recordings

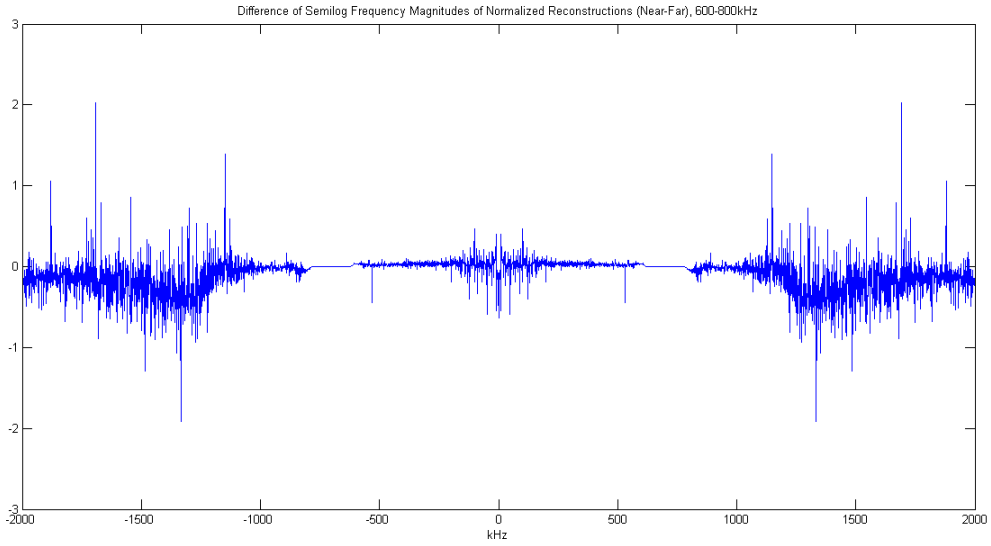

Figure 4.9: 400-1000kHz Transmit, Normalized Regularized Restoration Semilog Frequency Difference of Near and Far Recordings

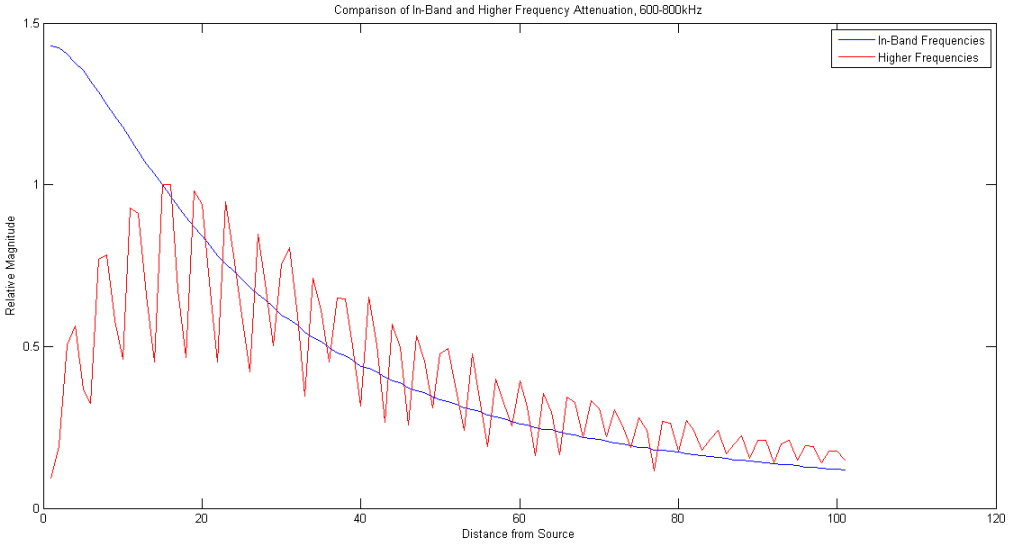

Figure 4.10: 600-800kHz Transmit, Attenuation Comparison

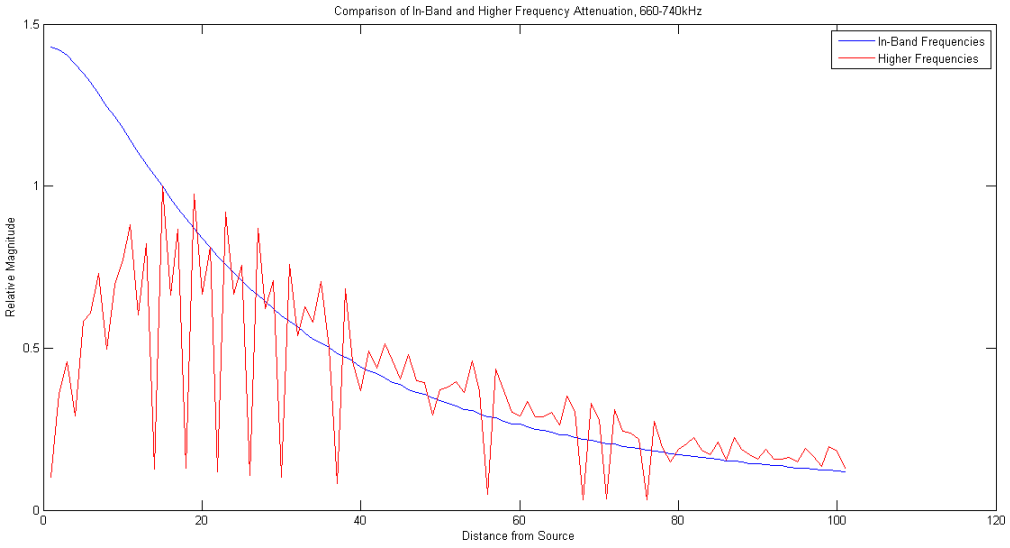

Figure 4.11: 660-740kHz Transmit, Attenuation Comparison

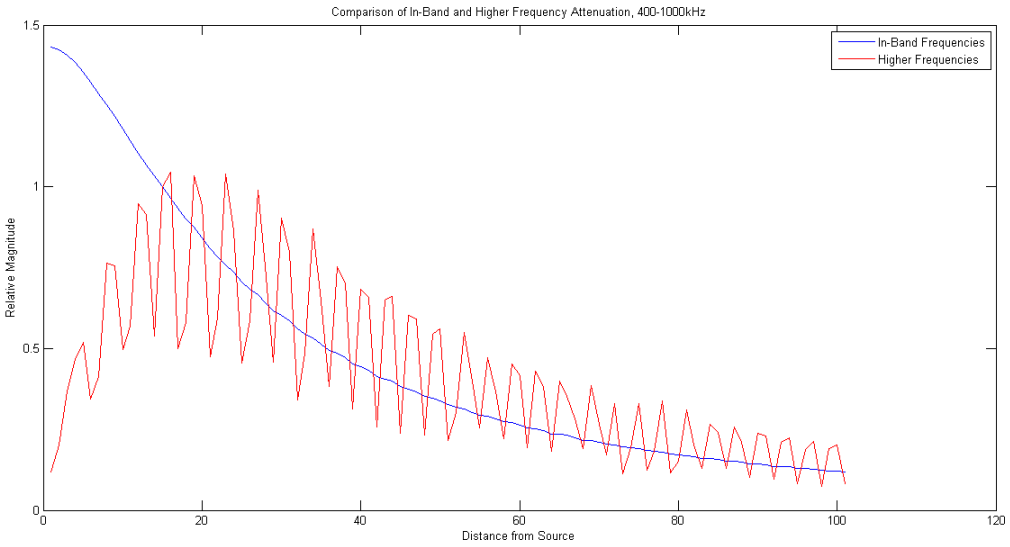

Figure 4.12: 400-1000kHz Transmit, Attenuation Comparison

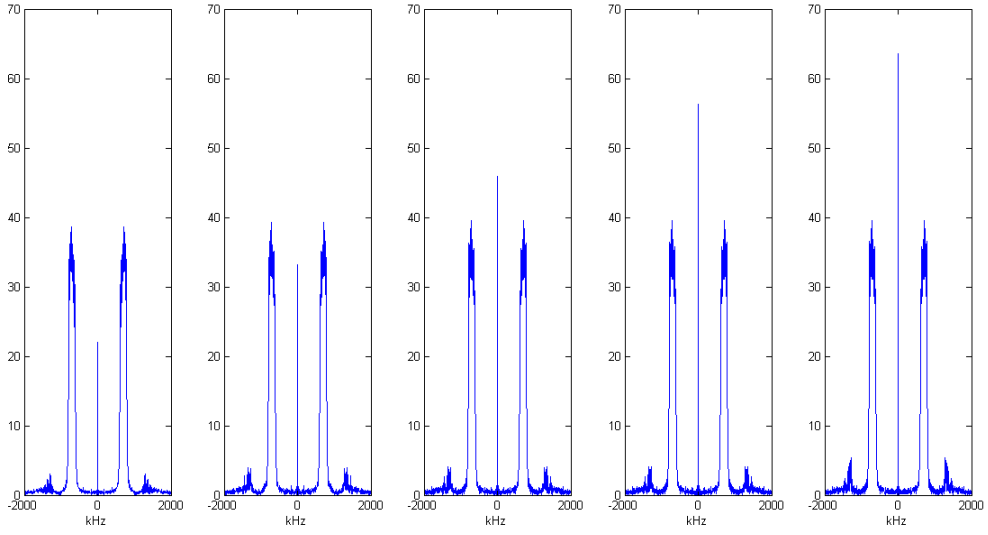

Figure 4.13: 600-800kHz Transmit, Spherical Target, Transmit Estimates Frequencies

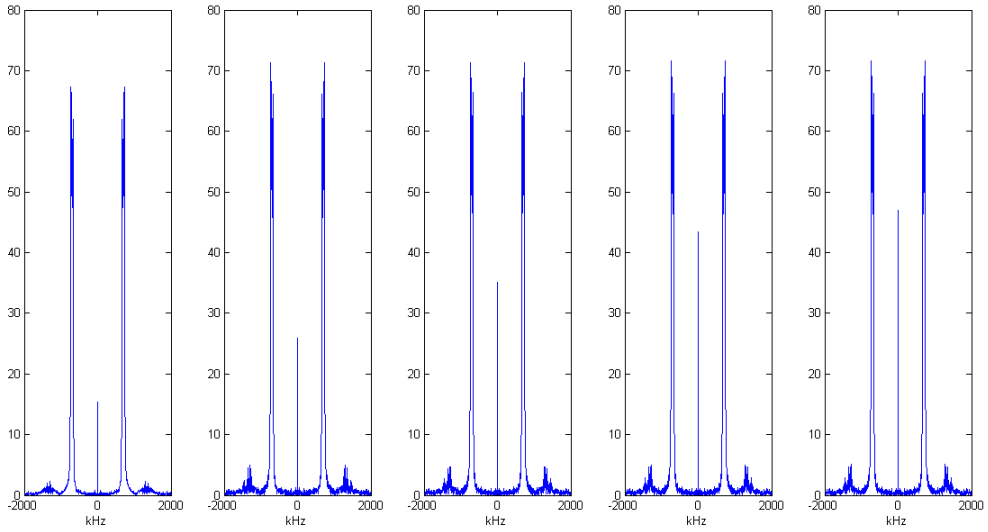

Figure 4.14: 660-740kHz Transmit, Spherical Target, Transmit Estimates Frequencies

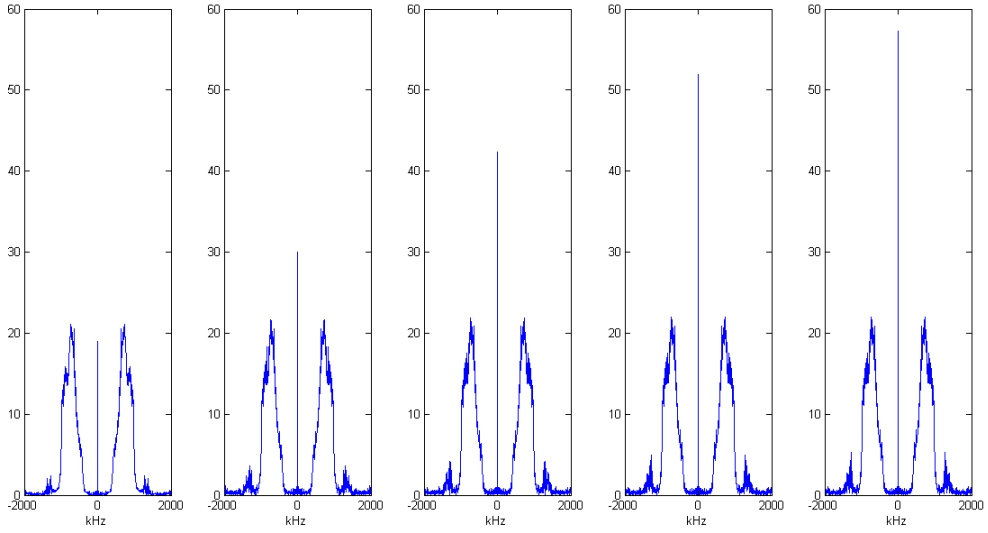

Figure 4.15: 400-1000kHz Transmit, Spherical Target, Transmit Estimates Frequencies

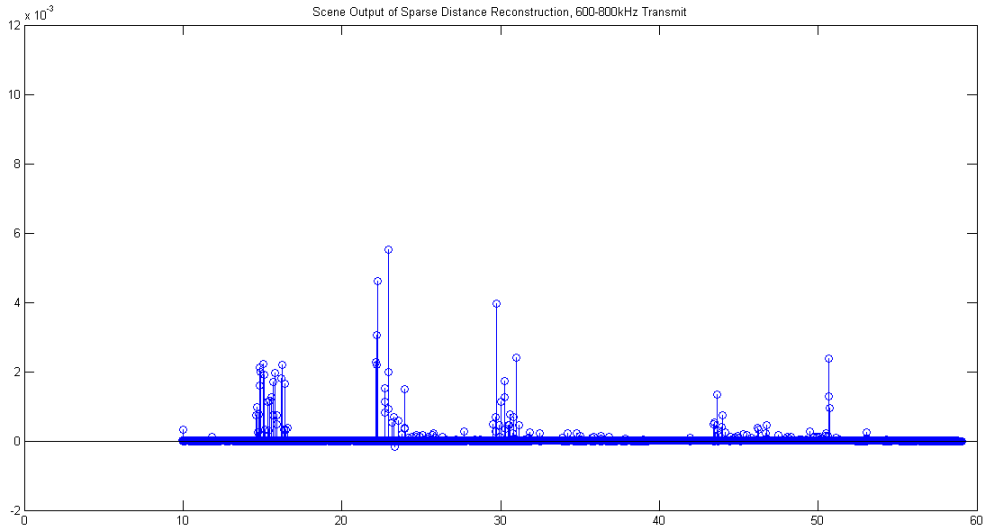

Figure 4.16: 600-800kHz Transmit, Spherical Target, Sparse Distance Reconstruction

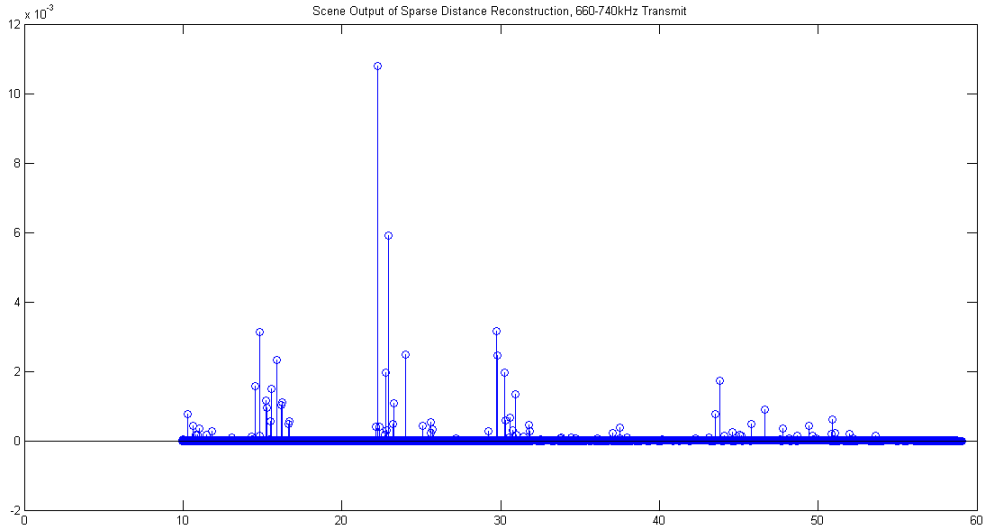

Figure 4.17: 660-740kHz Transmit, Spherical Target, Sparse Distance Reconstruction

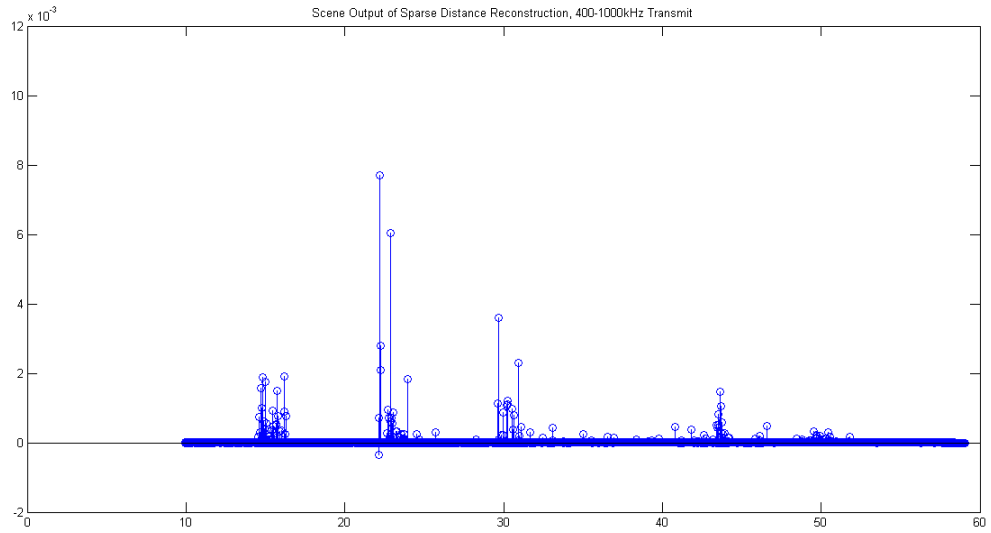

Figure 4.18: 400-1000kHz Transmit, Spherical Target, Sparse Distance Reconstruction

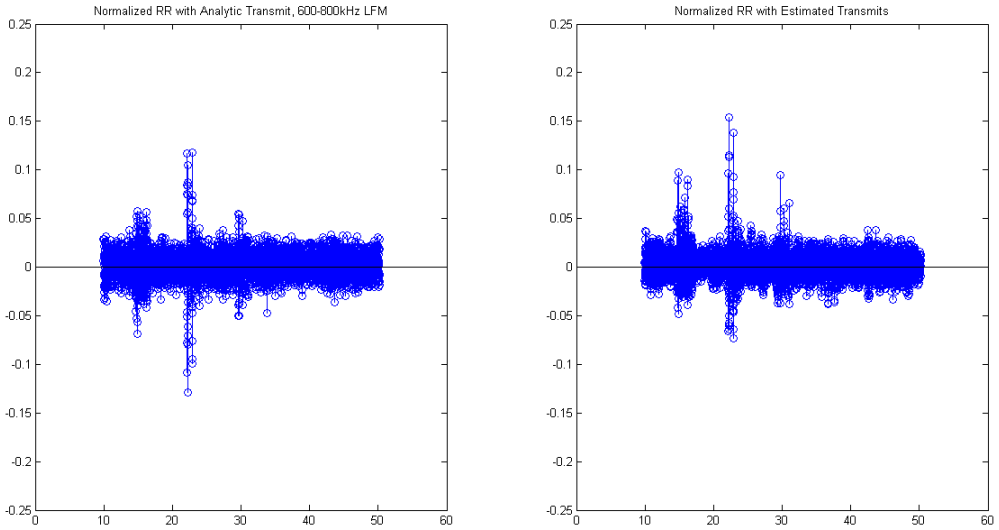

Figure 4.19: 600-800kHz Transmit, Spherical Target, Normalized Regularized Restoration Comparison

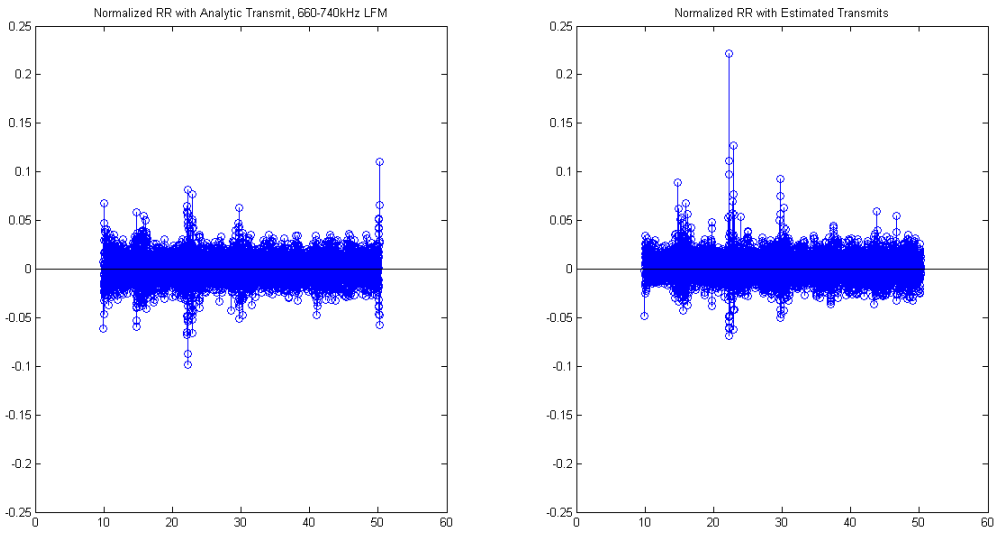

Figure 4.20: 660-740kHz Transmit, Spherical Target, Normalized Regularized Restoration Comparison

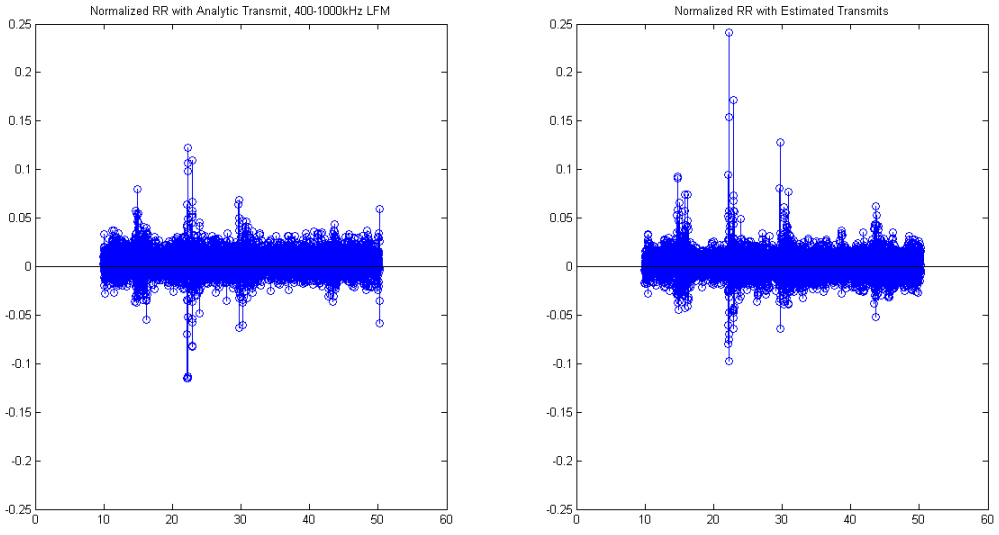

Figure 4.21: 400-1000kHz Transmit, Spherical Target, Normalized Regularized Restoration Comparison

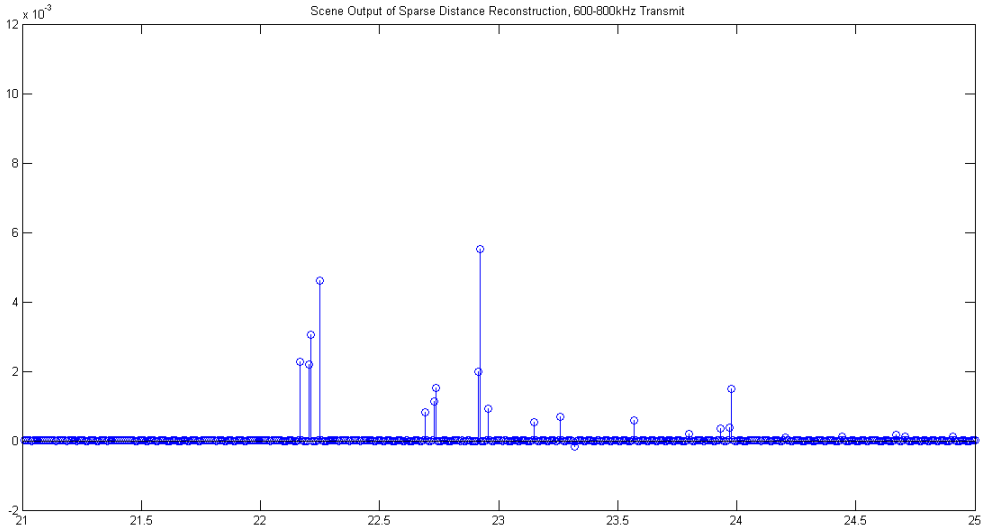

Figure 4.22: 600-800kHz Transmit, Spherical Target, Sparse Distance Reconstruction, Zoomed on Target Features

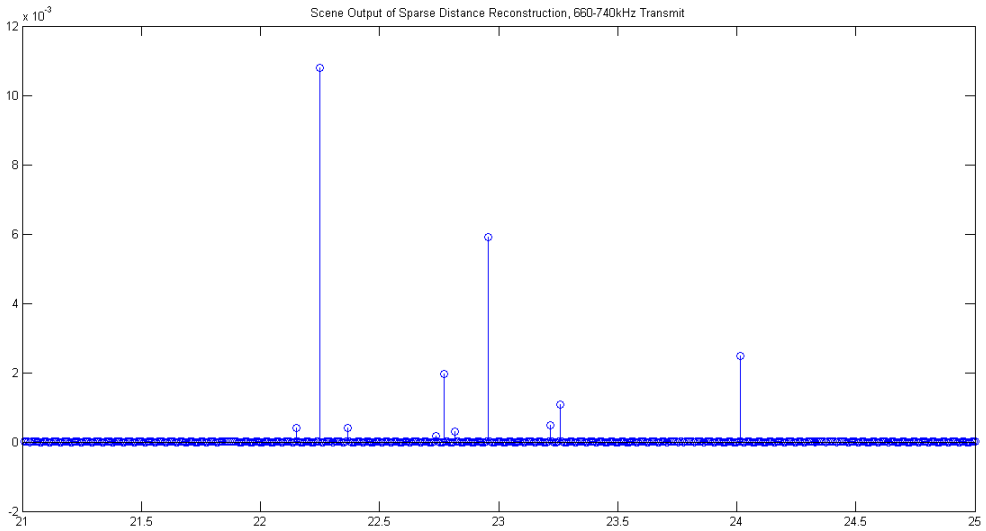

Figure 4.23: 660-740kHz Transmit, Spherical Target, Sparse Distance Reconstruction, Zoomed on Target Features

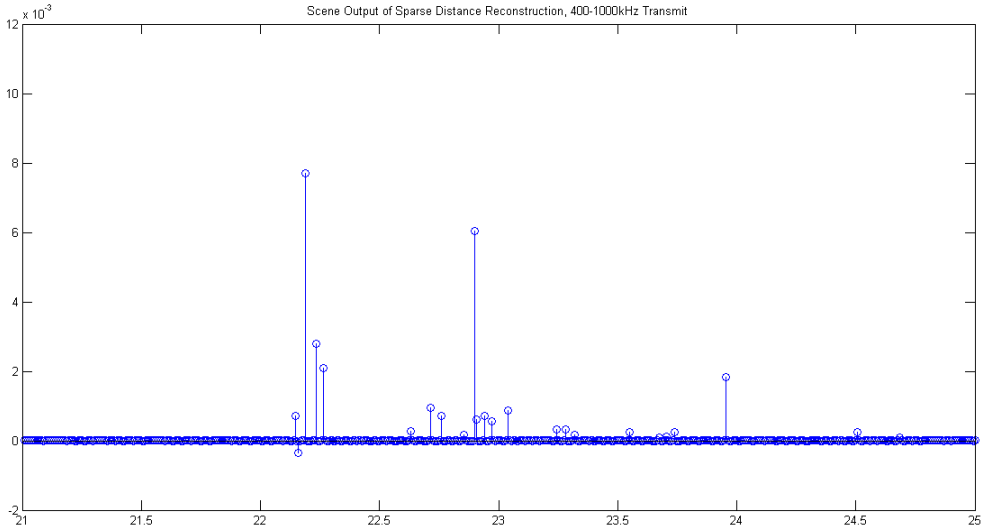

Figure 4.24: 400-1000kHz Transmit, Spherical Target, Sparse Distance Reconstruction, Zoomed on Target Features

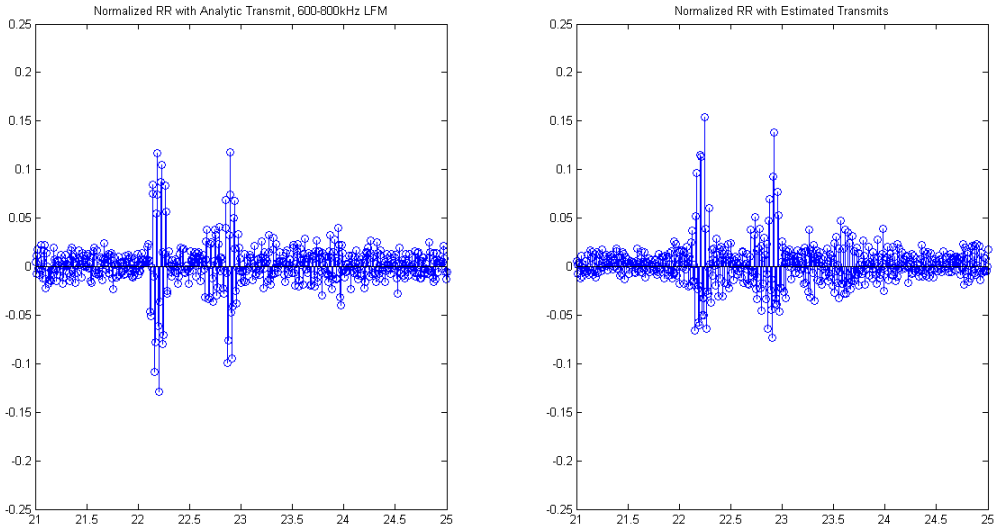

Figure 4.25: 600-800kHz Transmit, Spherical Target, Normalized Regularized Restoration Comparison, Zoomed on Target Features

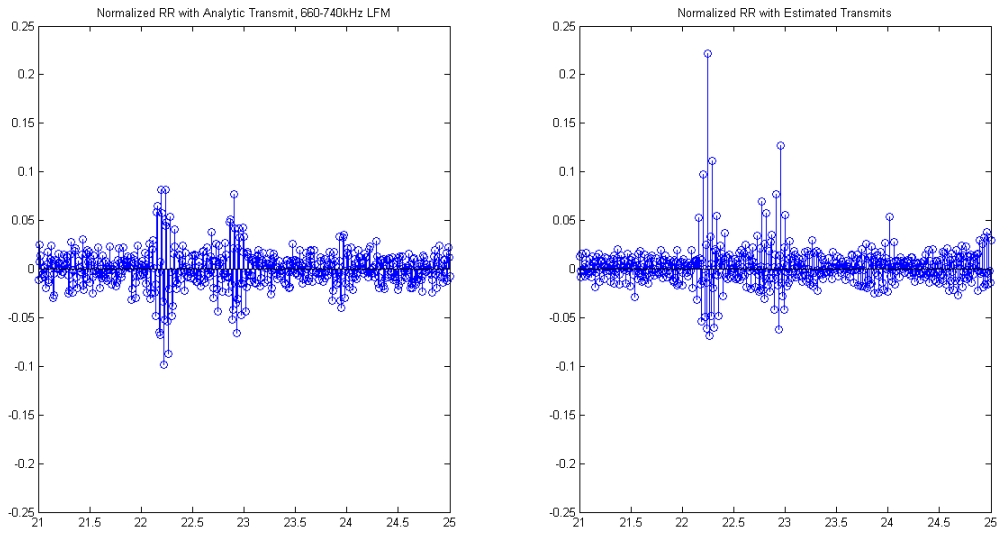

Figure 4.26: 660-740kHz Transmit, Spherical Target, Normalized Regularized Restoration Comparison, Zoomed on Target Features

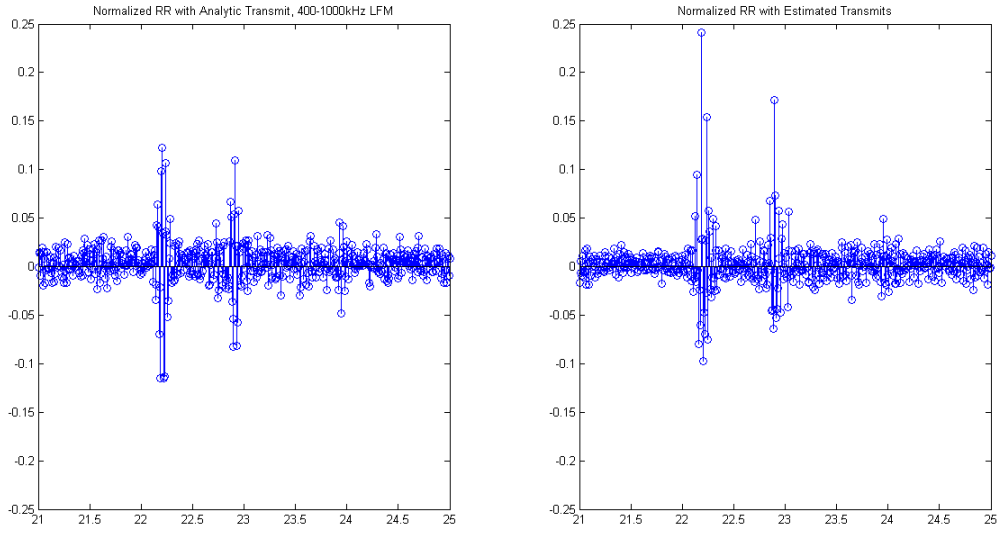

Figure 4.27: 400-1000kHz Transmit, Spherical Target, Normalized Regularized Restoration Comparison, Zoomed on Target Features

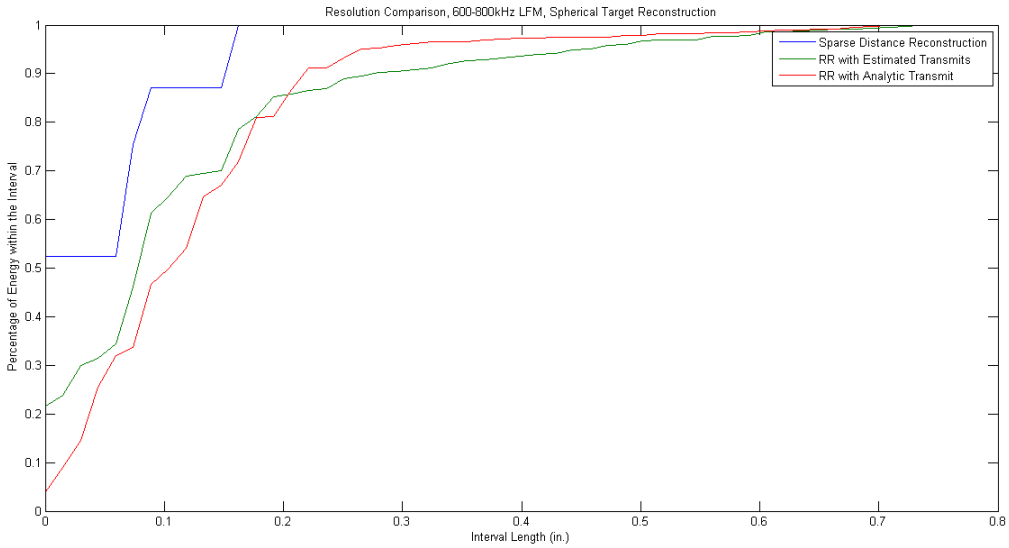

Figure 4.28: 600-800kHz Transmit, Resolution Comparison for Spherical Target Reconstructions

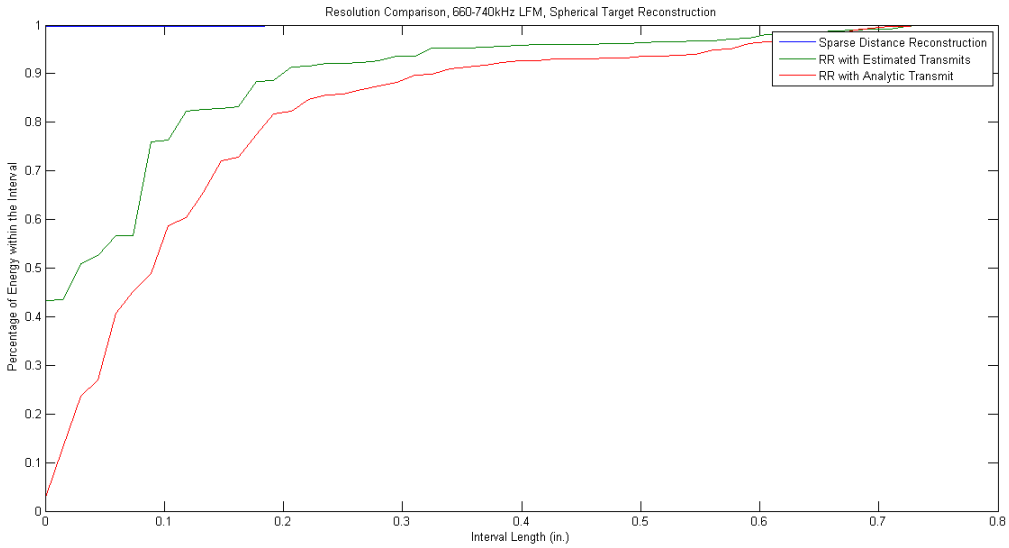

Figure 4.29: 660-740kHz Transmit, Resolution Comparison for Spherical Target Reconstructions

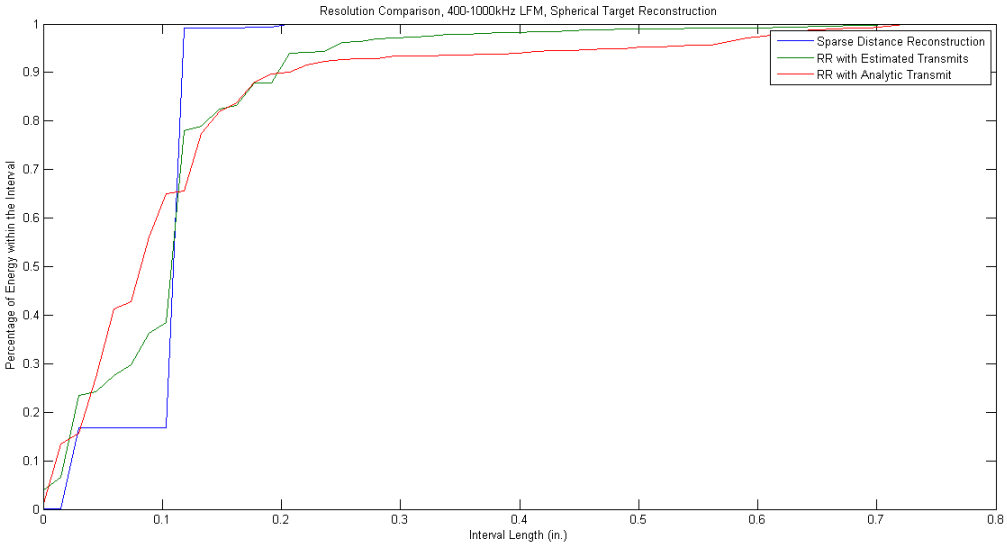

Figure 4.30: 400-1000kHz Transmit, Resolution Comparison for Spherical Target Reconstructions

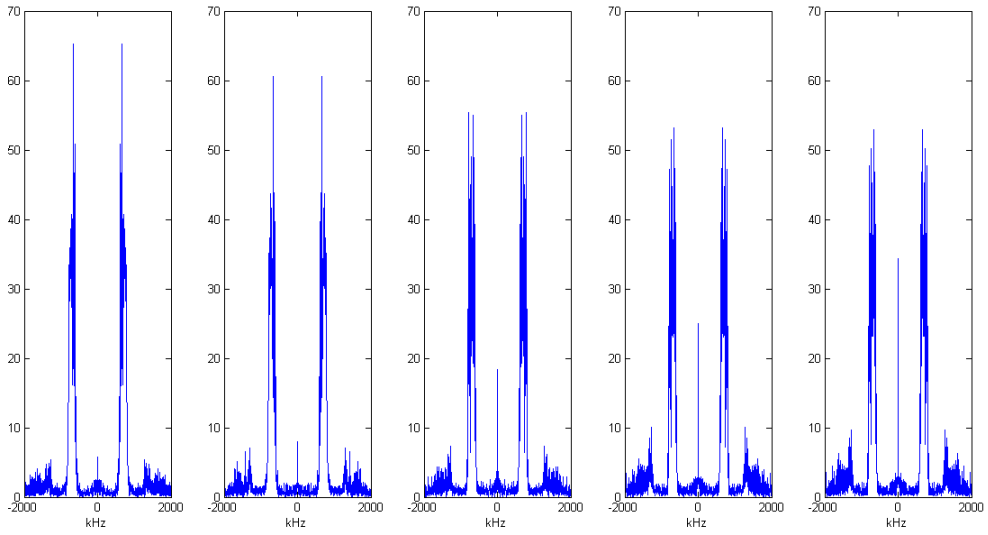

Figure 4.31: 600-800kHz Transmit, Cylindrical Target, Transmit Estimates Frequencies

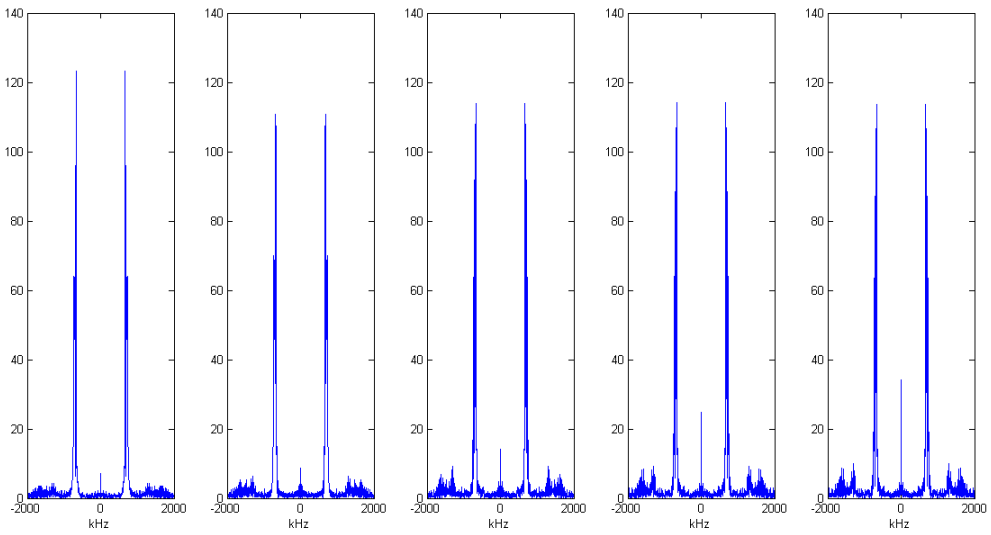

Figure 4.32: 660-740kHz Transmit, Cylindrical Target, Transmit Estimates Frequencies

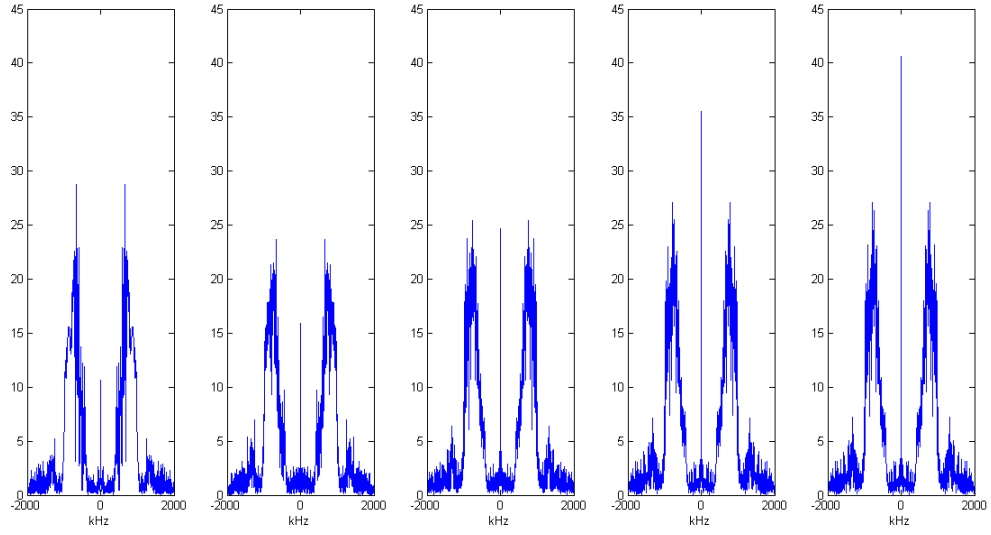

Figure 4.33: 400-1000kHz Transmit, Cylindrical Target, Transmit Estimates Frequencies

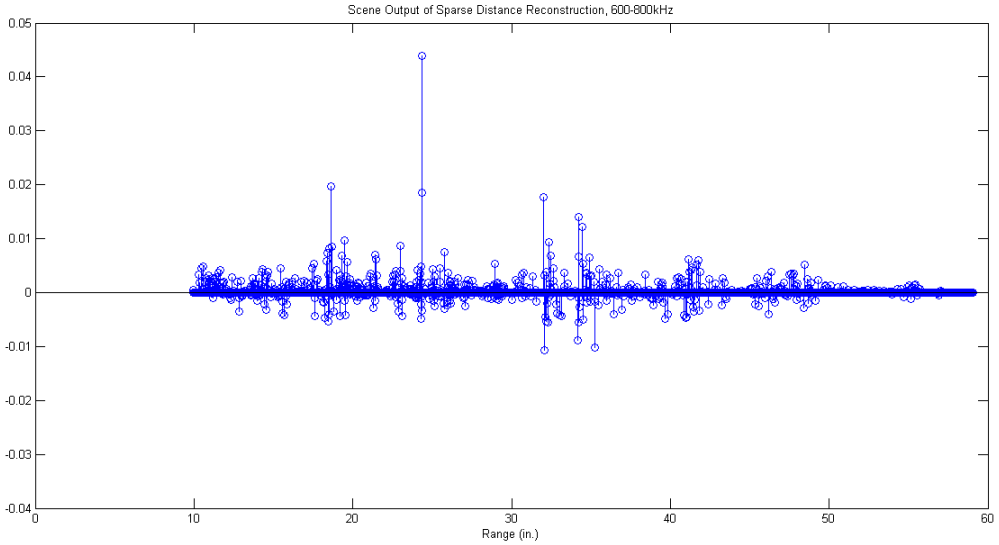

Figure 4.34: 600-800kHz Transmit, Cylindrical Target, Sparse Distance Reconstruction

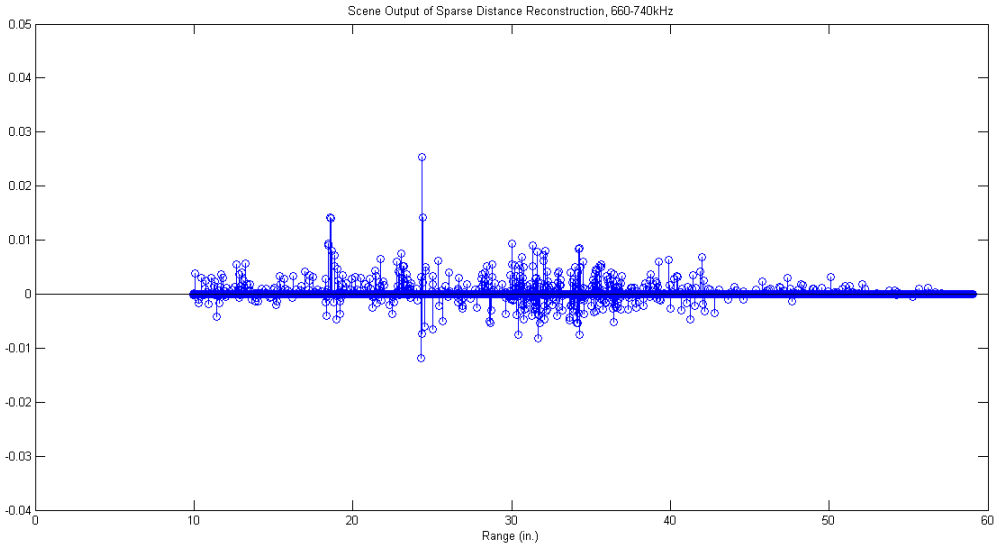

Figure 4.35: 660-740kHz Transmit, Cylindrical Target, Sparse Distance Reconstruction

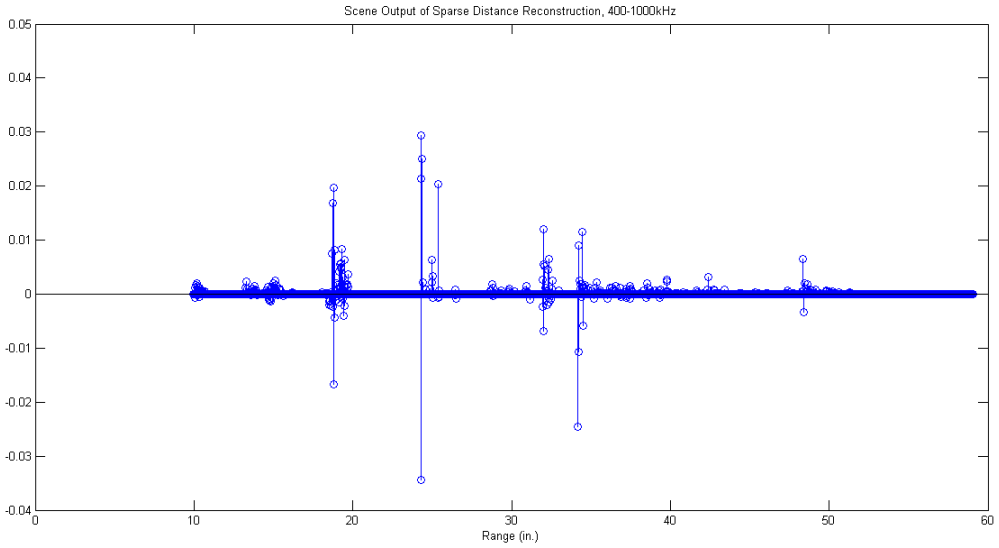

Figure 4.36: 400-1000kHz Transmit, Cylindrical Target, Sparse Distance Reconstruction

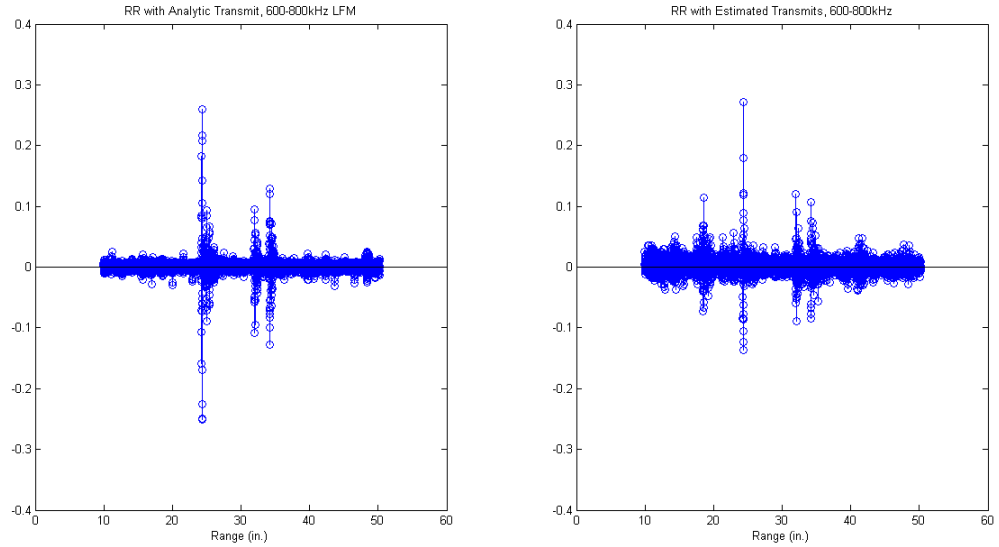

Figure 4.37: 600-800kHz Transmit, Cylindrical Target, Normalized Regularized Restoration Comparison

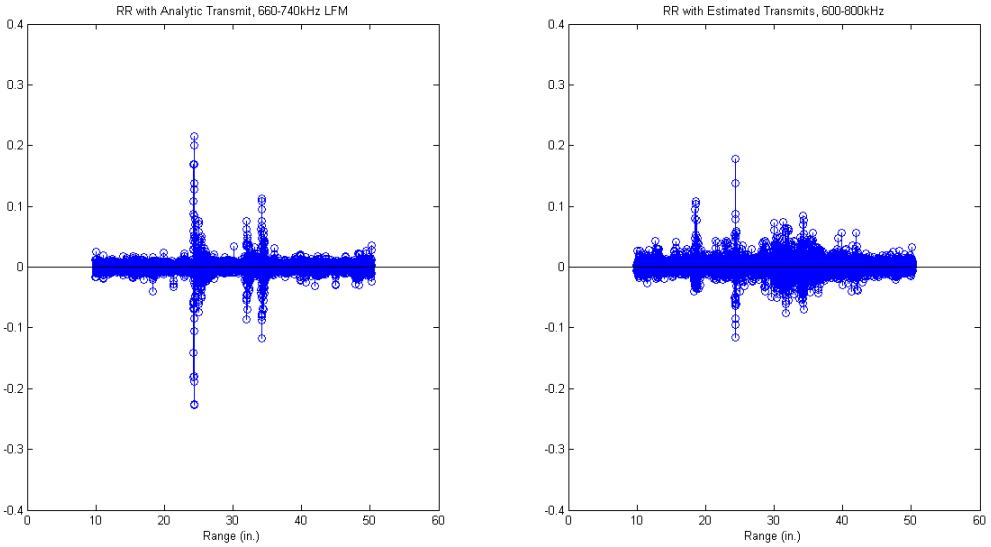

Figure 4.38: 660-740kHz Transmit, Cylindrical Target, Normalized Regularized Restoration Comparison

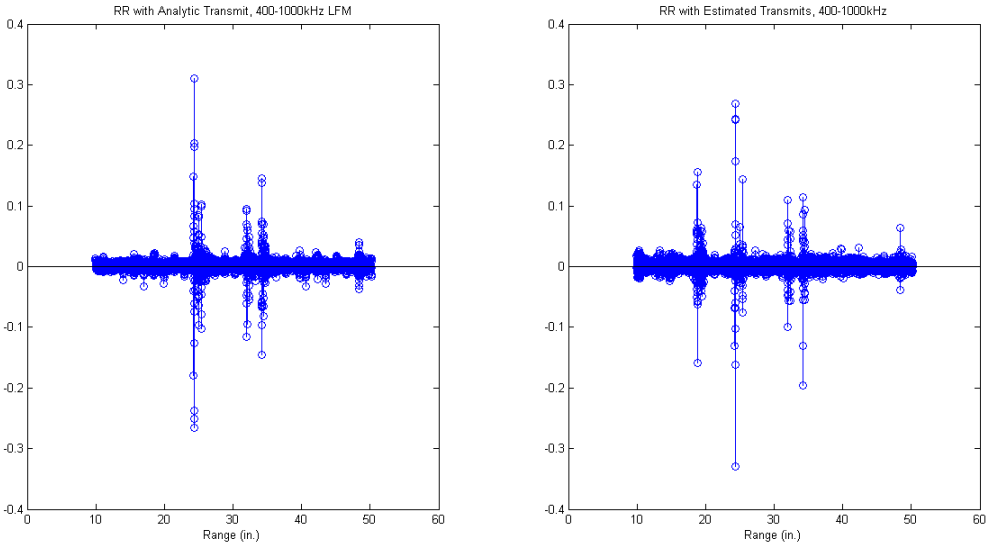

Figure 4.39: 400-1000kHz Transmit, Cylindrical Target, Normalized Regularized Restoration Comparison

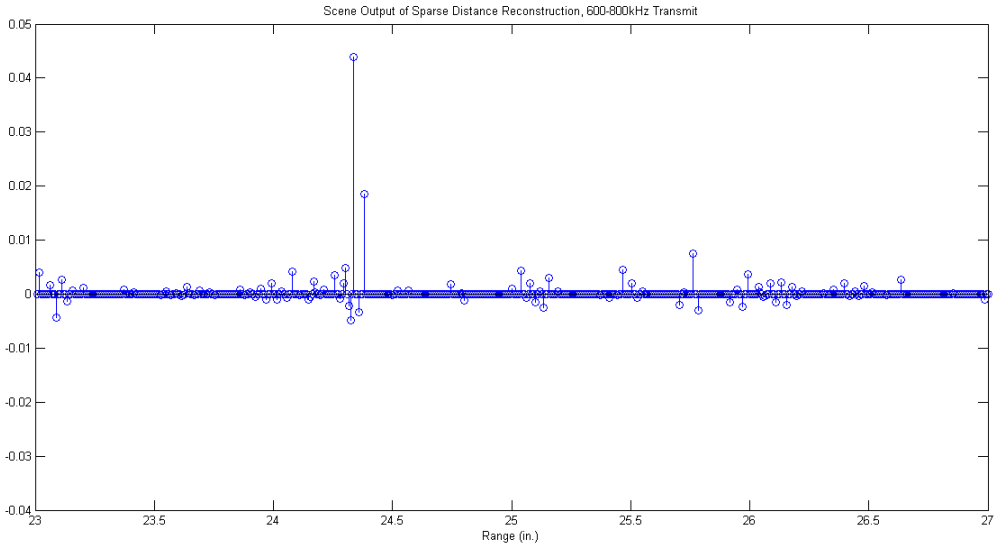

Figure 4.40: 600-800kHz Transmit, Cylindrical Target, Sparse Distance Reconstruction, Zoomed on Target Features

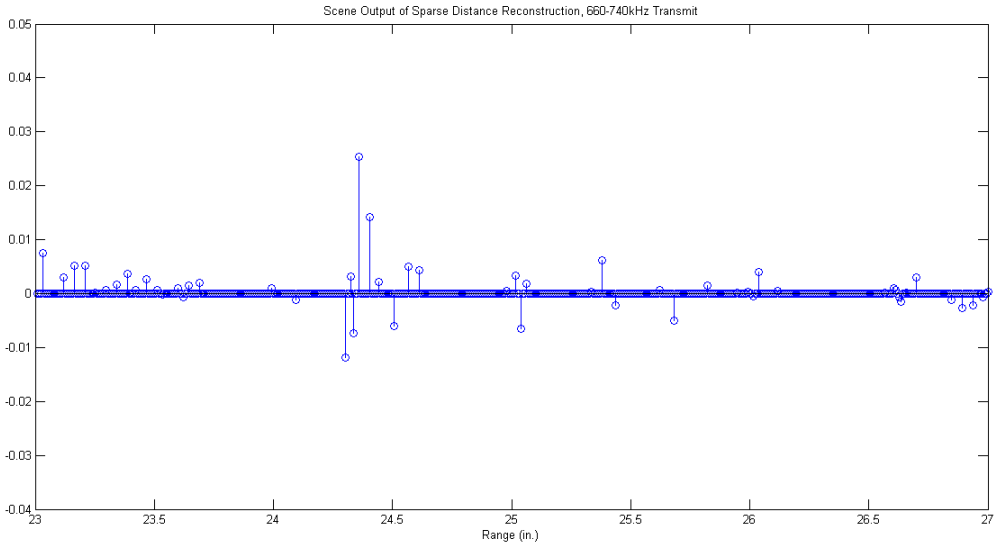

Figure 4.41: 660-740kHz Transmit, Cylindrical Target, Sparse Distance Reconstruction, Zoomed on Target Features

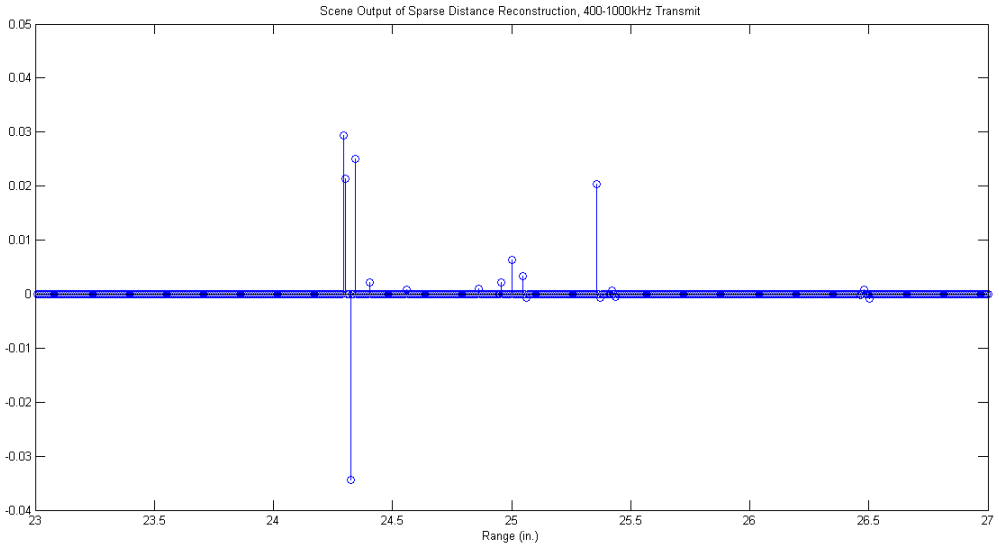

Figure 4.42: 400-1000kHz Transmit, Cylindrical Target, Sparse Distance Reconstruction, Zoomed on Target Features

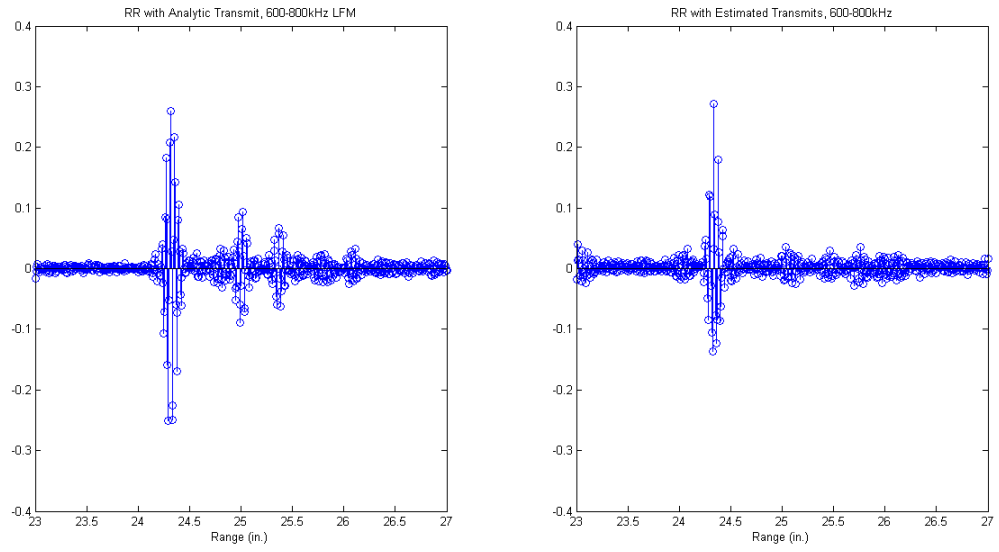

Figure 4.43: 600-800kHz Transmit, Cylindrical Target, Normalized Regularized Restoration Comparison, Zoomed on Target Features

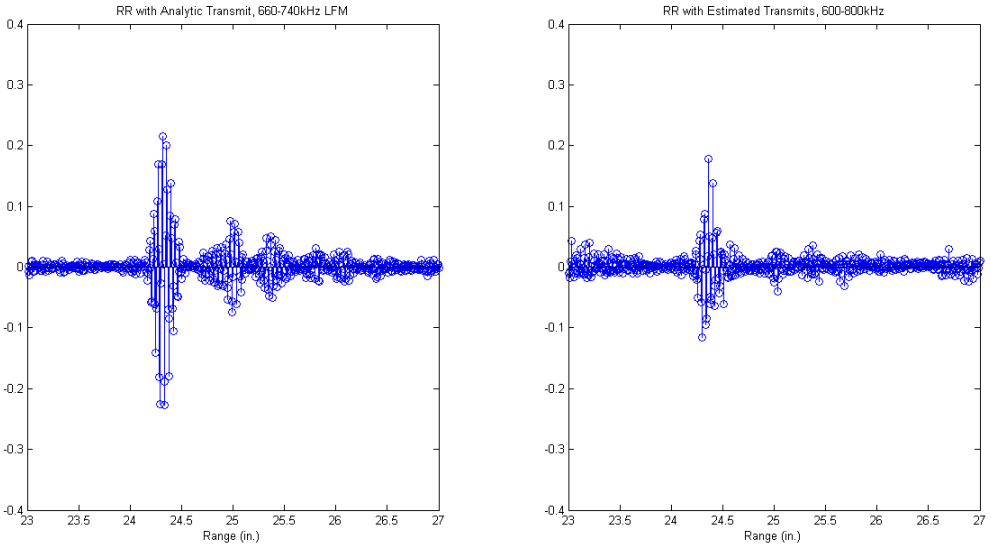

Figure 4.44: 660-740kHz Transmit, Cylindrical Target, Normalized Regularized Restoration Comparison, Zoomed on Target Features

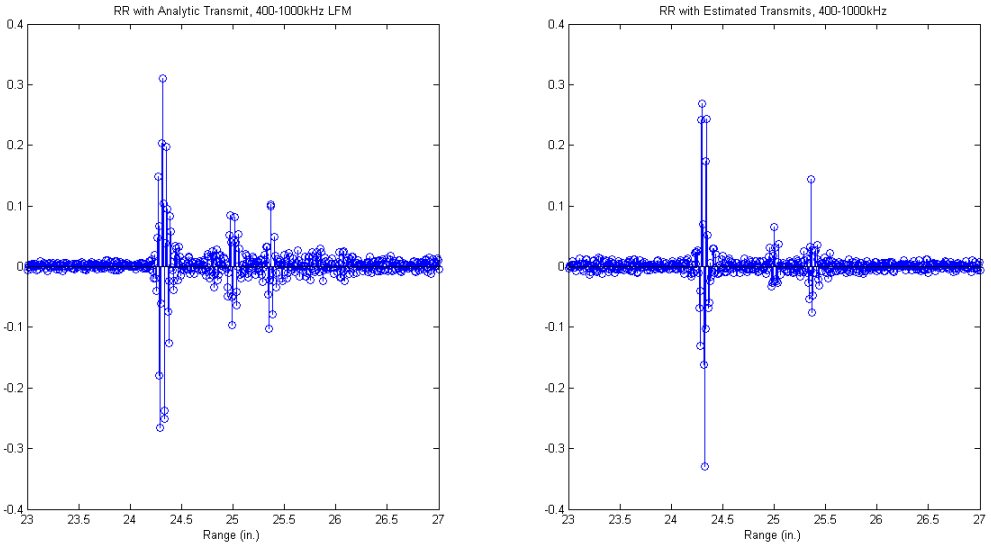

Figure 4.45: 400-1000kHz Transmit, Cylindrical Target, Normalized Regularized Restoration Comparison, Zoomed on Target Features

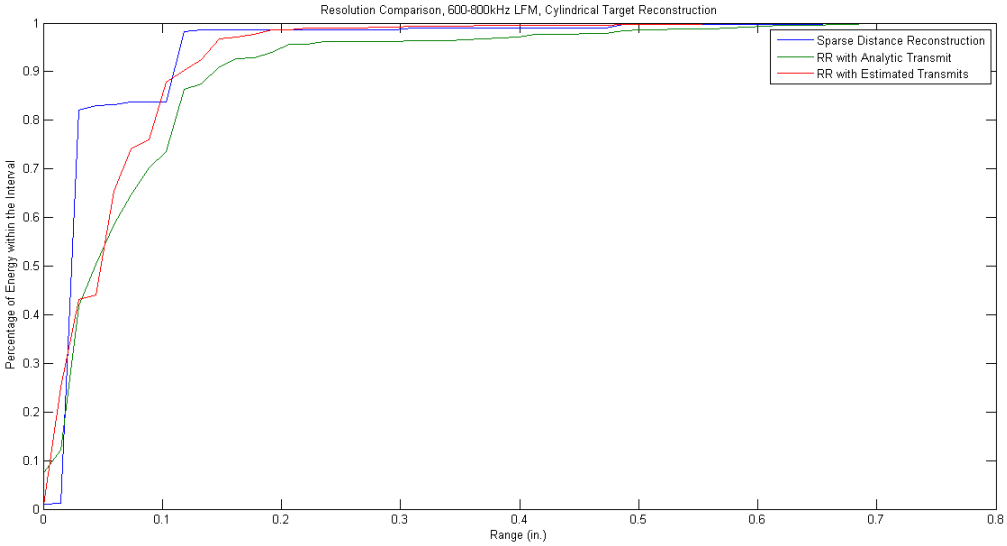

Figure 4.46: 600-800kHz Transmit, Resolution Comparison for Cylindrical Target Reconstructions

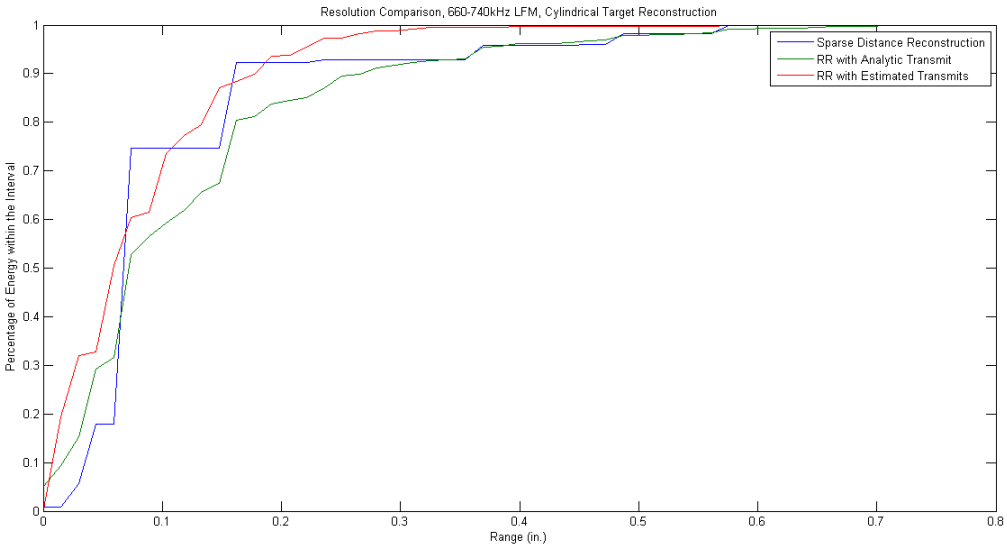

Figure 4.47: 660-740kHz Transmit, Resolution Comparison for Cylindrical Target Reconstructions

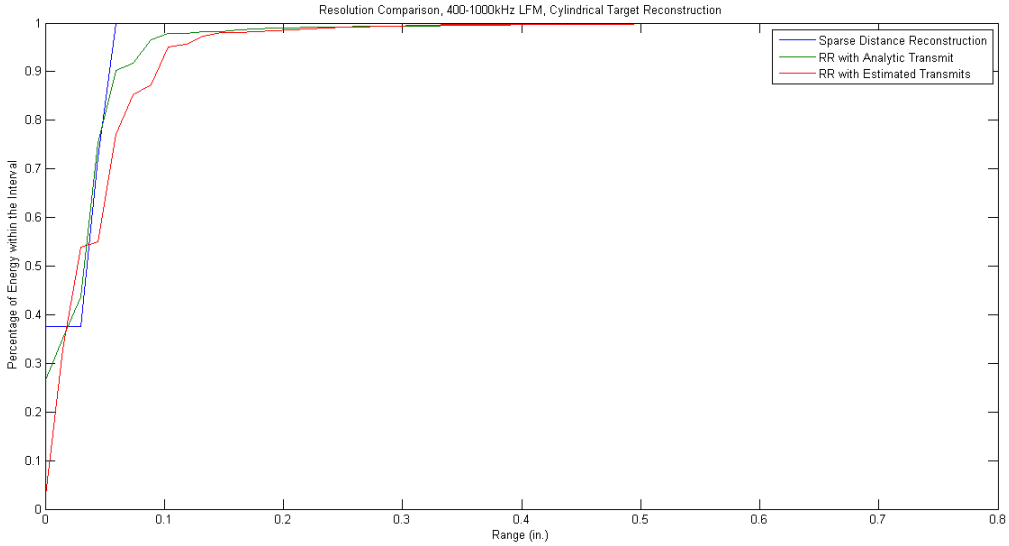

Figure 4.48: 400-1000kHz Transmit, Resolution Comparison for Cylindrical Target Reconstructions

Chapter 5

Conclusion

5.1 Summary of Work

We designed and implemented a variety of sparsity-based reconstruction techniques to one-dimensional sonar systems. We began with a sparse reconstruction algorithm for one-dimensional signals using a LASSO operator that we minimized with a majorization process. We applied the algorithm to simulated impulses at various noise levels and compared errors in the reconstructions with traditional sonar processing. We then applied the algorithm to experimental data obtained from a small-scale test bed (SSTB) and compared the results of our algorithm to expected results based on wideband matched-filtered data and the known target physics. The results were generally favorable for our algorithm in terms of resolution and the ease of identifying notable target features. Unfortunately, our algorithm also suppressed some weaker target features and also took significantly longer to run than traditional processing.

Once this was successfully implemented, we extended the concept to a case of an unreliable transmit description in one-dimensional sonar returns. This led to an algorithm that incorporated the previous work but also allowed for simultaneous modification of the transmit description using a process of semi-blind deconvolution. We showed a justification for some errors that can occur when using deconvolution-based processing on a signal generated using a noisy transmit. We then tested this technique on the same experimental data from the SSTB with positive results. This algorithm also generated a sparse scene reconstruction that was comparable to the sparse reconstruction in the previous section, but it also generated an estimated transmit description available for faster processing speeds. We compared the sparse reconstruction and regularized restoration with the analytic transmit description,

recorded replica of the transmit, and estimated transmit description. Regularized restoration with the estimated transmit description gave slightly better results than regularized restoration with the analytic transmit description and the recorded transmit replica, making it a promising technique for cases where the transmit is known to be unreliable.

Finally, we extended the semi-blind deconvolution concept to a spatially-varying transmit description. We provided evidence using recorded replicas from the SSTB that the transmit description can vary significantly from the initial design and that these variations can vary over distance. We then implemented a spatially-dependent semi-blind deconvolution algorithm that updated the transmit description based on consistency with the previous transmit estimate. We tested this on experimental data from the SSTB with positive results. We showed that this technique can be more effective than regular processing at identifying key target features in the data.

5.2 Future Research Applications

Much of the current research in sonar target identification focuses on the elastic responses of the materials. These algorithms were tuned for a sparse representation of the overall scene, and therefore suppressed some of the elastic effects in the output. A close analysis and re-tuning of any of the above techniques in regards to improving the identification of the elastic target response could be very promising.

In addition, these techniques largely relied on simulations that only incorporated additive noise. Realistic conditions are more complicated and these techniques (particularly the semi-blind deconvolutions) could prove effective in countering more complicated phase-type distortions such as Doppler. Some initial testing with the basic semi-blind deconvolution algorithm showed promise in mitigating errors created from an unknown Doppler effect in the return signal.

The semi-blind reconstruction has an underlying assumption of transmit signal distortion, but the current algorithm makes no assumptions on the type of distortion occurring.

Realistic estimates on known transmit distortions could greatly improve the accuracy of this reconstruction method. It also uses an ℓ_2 penalty on this term for the sake of speed, while a different penalty might be more appropriate for the situation (such as an ℓ_1 penalty). This could also be extended to the distance-dependent reconstruction if more is known about what can occur over various distances.

We could potentially improve on all of the reconstructions by optimizing the regularization parameters for a specific reconstruction goal. This work did not attempt to find optimal values for the parameters in the regularized restorations or the sparse reconstruction algorithms. This could be particularly important in the spatially-varying transmit estimation problem, as we show that the parameter effectively changes at each distance.

Bibliography

[1] R. Urick, *Principles of Underwater Sound*. McGraw-Hill, 1983.

[2] G. C. Gaunaurd and M. F. Werby, "Sound scattering by resonantly excited, fluid-loaded, elastic spherical shells," *The Journal of the Acoustical Society of America*, vol. 90, no. 5, p. 2536, 1991.

[3] K. L. Williams and P. L. Marston, "Axially focused (glory) scattering due to surface waves generated on spheres: Model and experimental confirmation using tungsten carbide spheres," *The Journal of the Acoustical Society of America*, vol. 78, no. 2, p. 722, 1985.

[4] K. L. Williams, S. G. Kargl, E. I. Thorsos, D. S. Burnett, J. L. Lopes, M. Zampolli, and P. L. Marston, "Acoustic scattering from a solid aluminum cylinder in contact with a sand sediment: Measurements, modeling, and interpretation," *The Journal of the Acoustical Society of America*, vol. 127, no. 6, p. 3356, 2010.

[5] A. Tesei, P. Guerrini, and M. Zampolli, "Tank measurements of scattering from a resin-filled fiberglass spherical shell with internal flaws," *The Journal of the Acoustical Society of America*, vol. 124, no. 2, p. 827, 2008.

[6] D. Royer and D. Clorennec, "Theoretical and experimental investigation of rayleigh waves on spherical and cylindrical surfaces," *Proceedings of Laser Ultrasonics, Montral*, p. 5, 2008.

[7] M. Elad, M. A. Figueiredo, and Y. Ma, "On the role of sparse and redundant representations in image processing," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 972–982, 2010.

[8] L. Potter, E. Ertin, J. Parker, and M. Cetin, "Sparsity and Compressed Sensing in Radar Imaging," *Proceedings of the IEEE*, vol. 98, pp. 1006–1020, June 2010.

[9] F. Marvasti, A. Amini, F. Haddadi, M. Soltanolkotabi, B. H. Khalaj, A. Aldroubi, S. Sanei, and J. A. Chambers, "A unified approach to sparse signal processing.," *EURASIP J. Adv. Sig. Proc.*, vol. 2012, p. 44, 2012.

[10] M. Leigsnering, F. Ahmad, M. Amin, and A. Zoubir, "Multipath exploitation in through-the-wall radar imaging using sparse reconstruction," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, pp. 920–939, Apr. 2014.

[11] H. L. Yap and R. Pribi, "Multi-target Radar Detection within a Sparsity Framework," *arXiv preprint arXiv:1311.1259*, 2013.

[12] X. Mao, Y. D. Zhang, and M. G. Amin, "Low-complexity sparse reconstruction for high-resolution multi-static passive SAR imaging*," *EURASIP Journal on Advances in Signal Processing*, vol. 2014, no. 1, pp. 1–12, 2014.

[13] N. Stefanakis, J. Marchal, V. Emiya, N. Bertin, R. Gribonval, and P. Cervenka, "Sparse underwater acoustic imaging: a case study," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pp. 2509–2512, IEEE, 2012.

[14] N. Kumar, Q. F. Tan, and S. S. Narayanan, "Object classification in sidescan sonar images with sparse representation techniques," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pp. 1333–1336, IEEE, 2012.

[15] J. McKay, R. Raj, V. Monga, and J. Isaacs, "Discriminative Sparsity for Sonar ATR," *arXiv preprint arXiv:1601.00119*, 2016.

[16] R. Fandos, L. Sadamori, and A. M. Zoubir, "Sparse Representation based Classification for mine hunting using Synthetic Aperture Sonar," pp. 3393–3396, IEEE, Mar. 2012.

[17] S. J. Reeves, "Analysis of the difficulties and possibilities for superresolution," in *AeroSense'97*, pp. 239–248, International Society for Optics and Photonics, 1997.

[18] N. Meinshausen and B. Yu, "Lasso-type recovery of sparse representations for high-dimensional data," *The Annals of Statistics*, vol. 37, pp. 246–270, Feb. 2009.

[19] Y. Xiang, B. Zhang, and W. Hong, "Lasso Based Performance Evaluation for Sparse One-Dimensional Radar Problem Under Random Sub-Sampling and Gaussian Noise," *Progress In Electromagnetics Research*, vol. 142, pp. 559–578, 2013.

[20] C. Debes, S. Leier, F. Nikolay, and A. Zoubir, "Compressive sensing for synthetic aperture imaging using a sparse basis transform," in *Geoscience and Remote Sensing Symposium (IGARSS), 2012 IEEE International*, pp. 7420–7423, July 2012.

[21] S. Leier and A. M. Zoubir, "Aperture undersampling using compressive sensing for synthetic aperture stripmap imaging," *EURASIP Journal on Advances in Signal Processing*, vol. 2014, p. 156, Oct. 2014.

[22] B. Sun, Y. Cao, J. Chen, C. Li, and Z. Qiao, "Compressive sensing imaging for general synthetic aperture radar echo model based on Maxwells equations," *EURASIP Journal on Advances in Signal Processing*, vol. 2014, p. 153, Oct. 2014.

[23] S. J. Reeves, "Accelerating sparse reconstruction for fast and precomputable system matrix inverses," in *IS&T/SPIE Electronic Imaging*, pp. 78730V–78730V, International Society for Optics and Photonics, 2011.

[24] P. Stoica and Y. Selen, "Cyclic minimizers, majorization techniques, and the expectation-maximization algorithm: a refresher," *IEEE Signal Processing Magazine*, vol. 21, pp. 112–114, Jan. 2004.

[25] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding and evaluating blind deconvolution algorithms," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 1964–1971, IEEE, 2009.

[26] F. Sroubek and P. Milanfar, "Robust Multichannel Blind Deconvolution via Fast Alternating Minimization," *IEEE Transactions on Image Processing*, vol. 21, pp. 1687–1700, Apr. 2012.

[27] S. Bonettini, A. Cornelio, and M. Prato, "A New Semiblind Deconvolution Approach for Fourier-Based Image Restoration: An Application in Astronomy," *SIAM Journal on Imaging Sciences*, vol. 6, pp. 1736–1757, Sept. 2013.

[28] D. Kundur and D. Hatzinakos, "A novel blind deconvolution scheme for image restoration using recursive filtering," *Signal Processing, IEEE Transactions on*, vol. 46, no. 2, pp. 375–390, 1998.

[29] G. Harikumar and Y. Bresler, "Perfect blind restoration of images blurred by multiple filters: Theory and efficient algorithms," *Image Processing, IEEE Transactions on*, vol. 8, no. 2, pp. 202–219, 1999.

[30] M. J. Roan, M. R. Gramann, J. G. Erling, and L. H. Sibul, "Blind deconvolution applied to acoustical systems identification with supporting experimental results," *The Journal of the Acoustical Society of America*, vol. 114, no. 4, p. 1988, 2003.

[31] A. Zunino, F. Benvenuto, E. Armadillo, M. Bertero, and E. Bozzo, "Iterative deconvolution and semiblind deconvolution methods in magnetic archaeological prospecting," *GEOPHYSICS*, vol. 74, pp. L43–L51, July 2009.

[32] M. Lankarany and M. H. Savoji, "Semi Blind Deconvolution; Application to Glottal Flow Estimation," 2010.

[33] S. U. Park, N. Dobigeon, and A. O. Hero, "Variational semi-blind sparse deconvolution with orthogonal kernel bases and its application to MRFM," *Signal Processing*, vol. 94, pp. 386–400, 2014.

[34] G. B. Giannakis and E. Serpedin, "Linear multichannel blind equalizers of nonlinear FIR Volterra channels," *IEEE Transactions on Signal Processing*, vol. 45, no. 1, pp. 67–81, 1997.

[35] S. Choudhary and U. Mitra, "Fundamental limits of blind deconvolution Part I: Ambiguity kernel," *arXiv preprint arXiv:1411.3810*, 2014.

[36] O. Shalvi and E. Weinstein, "New criteria for blind deconvolution of nonminimum phase systems (channels)," *IEEE Transactions on Information Theory*, vol. 36, pp. 312–321, Mar. 1990.