

## Abstract

Three-dimensional diffuse optical tomography (DOT) attempts to map *in vivo* tissue blood oxygenation and blood volume levels by reconstructing the spatial distribution of the optical absorption coefficient from intensity measurements on the surface of the body. This problem is typically ill-posed due to the large attenuation and scattering of the diffuse wave. In addition, certain applications, such as breast mapping from a planar array of optical sources and detectors, increase the ill-posedness of the problem by restricting the views of the sources and detectors. Also, unlike CT or MRI, the interaction of the diffuse wave with the medium can not be restricted to two dimensions. Thus all three spatial dimensions must be considered to accurately account for the diffuse wave propagation. This results in problems that are typically highly underdetermined as well.

We first present a comparison of the currently employed linear model reconstruction techniques, to both identify the most promising class using a linear model approach and to provide a baseline for the comparison of the reconstruction performance of the additionally constrained algorithms. Our approach to improve the reconstruction fidelity of this highly underdetermined and ill-posed problem has been to incorporate *a priori* constraints on the solution. We have developed two types of constrained reconstruction algorithms. The first takes advantage of the necessity of collecting data at two optical wavelengths to implement a constraint on the boundaries and value of the absorption anomaly. The second algorithm is an admissible solution approach where we have examined a number of convex constraint functions on the solution. Using the admissible solution approach

in a simulation scenario with the right constraints we have been able to produce good reconstructions from two-dimensional simulations that approximate the properties of the three-dimensional problem.

# Constrained Reconstruction Techniques for Diffuse Optical

## Tomography

Rick Gaudette

September 8, 2000

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	DOT imaging	3
1.2	Technical challenges of DOT imaging	5
1.3	Background	7
1.3.1	Modeling of photon propagation in highly scattering media	7
1.3.2	Validation of the photon diffusion model	11
1.3.3	Solutions to the diffusion model	12
1.3.4	Boundary modeling and diffuse sources	12
1.3.5	Optical characteristics of biological tissue	14
1.3.6	Imaging using DOT	15
1.3.7	Related work in other disciplines	20
1.4	Contributions of this thesis	21
<b>2</b>	<b>Simulation models</b>	<b>24</b>
2.1	Measurement scenario	25
2.2	Propagation models	27
2.2.1	Discretization of the propagation model	33
2.2.2	Validation of the propagation model	34
2.3	Noise models	41
2.3.1	Detector noise model	43
2.3.2	Shot noise model	44
2.3.3	Combined noise model	44
<b>3</b>	<b>Linear model reconstruction evaluation</b>	<b>48</b>
3.1	Simulation details	49
3.1.1	Computational geometry	50
3.1.2	Source and detector parameters	50
3.1.3	Volume discretization	51
3.1.4	Medium optical parameters	52
3.1.5	Noise model	53
3.1.6	Model mismatch analysis	53
3.2	Reconstruction algorithms	54

3.2.1	Algebraic techniques . . . . .	55
3.2.2	Subspace techniques . . . . .	57
3.2.3	Regularization parameter selection . . . . .	59
3.3	Results and observations . . . . .	60
3.3.1	Example reconstructions . . . . .	61
3.3.2	Performance measures . . . . .	64
3.3.3	Regularization parameter sensitivity . . . . .	70
3.4	Linear model discussion . . . . .	71
<b>4</b>	<b>Two wavelength constrained reconstruction</b>	<b>74</b>
4.1	Constant absorption coefficient ratio (CACR) constraint . . . . .	75
4.2	Applying the CACR constraint . . . . .	77
4.2.1	Independent TSVD reconstruction . . . . .	78
4.2.2	Joint truncated singular value decomposition (J-TSVD) . . . . .	78
4.2.3	Post constrained truncated singular value decomposition (PC-TSVD) . . . . .	79
4.2.4	Joint modified truncated singular value decomposition (JM-TSVD) . . . . .	80
4.3	Simulation scenario . . . . .	85
4.3.1	Two wavelength SV spectra . . . . .	87
4.4	Simulation results . . . . .	89
4.4.1	Qualitative results . . . . .	89
4.4.2	Quantitative results . . . . .	92
4.5	Two wavelength discussion . . . . .	95
<b>5</b>	<b>Admissible solution DOT</b>	<b>97</b>
5.1	The admissible solution approach . . . . .	98
5.2	The Ellipsoid Algorithm . . . . .	100
5.2.1	Subgradients . . . . .	100
5.2.2	Extending the search algorithm to multiple dimensions . . . . .	102
5.2.3	Derivation of subgradients . . . . .	106
5.3	Admissible solution simulation scenario . . . . .	111
5.4	Admissible solution performance . . . . .	114
5.4.1	Constraint comparisons . . . . .	115
5.4.2	Monte Carlo noise test . . . . .	124
5.4.3	Object parameter tests . . . . .	125
5.4.4	Variable background test . . . . .	136
5.5	Computational considerations . . . . .	139
5.6	Admissible solution discussion . . . . .	140
<b>6</b>	<b>Conclusions and future work</b>	<b>144</b>
6.1	Future work . . . . .	145
<b>A</b>	<b>Table of Parameters</b>	<b>148</b>

<b>B Photon Migration Imaging Toolbox</b>	<b>150</b>
B.1 PMI Toolbox Introduction . . . . .	150
B.2 High Level Algorithms . . . . .	151
B.2.1 The Photon Migration Imaging data structure . . . . .	152
<b>Bibliography</b>	<b>164</b>

# List of Figures

1.1	Hemoglobin absorption coefficient . . . . .	15
2.1	Simulation schematic . . . . .	26
2.2	1 cm source detector sensitivity . . . . .	28
2.3	5 cm source detector sensitivity . . . . .	28
2.4	Method of images schematic . . . . .	32
2.5	Born forward model validation scenario . . . . .	36
2.6	Center source scattered fluence rate comparison . . . . .	39
2.7	Left source scattered fluence rate comparison . . . . .	40
2.8	Center source scattered fluence rate ratios . . . . .	41
2.9	Left source scattered fluence rate ratios . . . . .	42
3.1	True absorption perturbation function . . . . .	51
3.2	Singular value spectrum . . . . .	58
3.3	L-Curve for TCG at 20 dB SNR . . . . .	61
3.4	F3D reconstruction examples at 20 dB SNR . . . . .	63
3.5	2DS reconstruction examples at 20 dB SNR . . . . .	64
3.6	MSE for the F3D reconstructions . . . . .	67
3.7	MSE for the 2DS reconstructions . . . . .	68
3.8	OCE for F3D and 2DS reconstructions . . . . .	69
3.9	AE for the F3D and 2DS reconstructions . . . . .	69
3.10	Subspace reconstruction technique regularization parameter sensitivity . . . . .	70
3.11	Algebraic reconstruction technique regularization parameter sensitivity . . . . .	71
4.1	Source and detector positions for collocated source geometry . . . . .	86
4.2	Source and detector positions for separate source geometry . . . . .	86
4.3	Singular value spectra for collocated source geometry . . . . .	88
4.4	Singular value spectra for separate source geometry . . . . .	88
4.5	Collocated source, uniform noise TSVD reconstruction at 30 dB SNR . . . . .	90
4.6	Collocated source, uniform noise J-TSVD reconstruction at 30 dB SNR . . . . .	90
4.7	Collocated source, uniform noise PC-TSVD reconstruction at 30 dB SNR . . . . .	90
4.8	Collocated source, uniform noise JM-TSVD reconstruction at 30 dB SNR . . . . .	90
4.9	Separate source, uniform noise TSVD reconstruction at 30 dB SNR . . . . .	91
4.10	Separate source, uniform noise J-TSVD reconstruction at 30 dB SNR . . . . .	91

4.11	Separate source, uniform noise PC-TSVD reconstruction at 30 dB SNR . . . . .	91
4.12	Separate source, uniform noise JM-TSVD reconstruction at 30 dB SNR . . . . .	91
4.13	Collocated source geometry MSE, uniform noise model . . . . .	93
4.14	Collocated source geometry MSE, simplified shot noise model . . . . .	93
4.15	Separate source geometry MSE, uniform noise model . . . . .	93
4.16	Separate source geometry MSE, simplified shot noise model . . . . .	93
4.17	Collocated source geometry OCE, uniform noise model . . . . .	94
4.18	Collocated source geometry OCE, simplified shot noise model . . . . .	94
4.19	Separate source geometry OCE, uniform noise model . . . . .	94
4.20	Separate source geometry OCE, simplified shot noise model . . . . .	94
5.1	Admissible solution approach illustration . . . . .	99
5.2	One-dimensional subgradient illustration . . . . .	101
5.3	Ellipsoid Algorithm example . . . . .	105
5.4	True absorption perturbation image . . . . .	112
5.5	Singular value spectrum of the forward operator . . . . .	114
5.6	Minimum norm solution with no added noise to the measurements . . . . .	114
5.7	TSVD reconstruction at 30 dB SNR . . . . .	117
5.8	Admissible solution reconstruction with the residual constraint at 30 dB SNR . . .	117
5.9	MSE for the TSVD and residual constraint admissible solution algorithms . . . . .	117
5.10	residual-TV constrained admissible solution reconstruction . . . . .	118
5.11	MSE for the residual-TV constrained admissible solution approach . . . . .	118
5.12	Individual MSE measures for the residual-TV constrained admissible solution . . .	119
5.13	Residual-TV-min/max constrained admissible solution reconstruction . . . . .	122
5.14	True absorption anomaly perturbation . . . . .	122
5.15	MSE for the residual-TV-min/max constrained admissible solution approach . . . .	122
5.16	MSE versus SNR for the residual-TV-min/max admissible solution approach. . . .	122
5.17	Best MSE versus the minimum value constraint . . . . .	123
5.18	Monte Carlo MSE vs. SNR . . . . .	125
5.19	Anomaly position test true anomaly, deep left . . . . .	127
5.20	Anomaly position test true anomaly, center . . . . .	127
5.21	Anomaly position test true anomaly, shallow right . . . . .	127
5.22	The minimum MSE averaged over all nine anomaly positions as a function of SNR	128
5.23	Separate MSE curves for each depth examined . . . . .	128
5.24	Anomaly size test, smallest anomaly . . . . .	129
5.25	Anomaly size test, largest anomaly . . . . .	129
5.26	Anomaly size test MSE averaged over all sizes . . . . .	130
5.27	Anomaly size test MSE separately for each size . . . . .	130
5.28	Anomaly size test example reconstruction at 60 dB SNR . . . . .	131
5.29	Anomaly size test example reconstruction at 20 dB SNR . . . . .	131
5.30	Anomaly contrast test MSE averaged over all contrasts . . . . .	132
5.31	Anomaly contrast test MSE for a 2:1 contrast ratio . . . . .	132
5.32	Anomaly contrast test MSE for a 3:1 contrast ratio . . . . .	133

5.33 Anomaly contrast test MSE for a 4:1 contrast ratio . . . . . 133

5.34 Anomaly contrast test MSE for a 5:1 contrast ratio . . . . . 133

5.35 Multiple anomaly test, shallow anomaly . . . . . 134

5.36 Multiple anomaly test, center anomaly . . . . . 134

5.37 Multiple anomaly test, deep anomaly . . . . . 134

5.38 Multiple absorption anomaly reconstruction example, shallow, 60 dB SNR . . . . . 135

5.39 Multiple absorption anomaly reconstruction example, shallow, 30 dB SNR . . . . . 135

5.40 Multiple absorption anomaly reconstruction example, center, 60 dB SNR . . . . . 135

5.41 Multiple absorption anomaly reconstruction example, center, 30 dB SNR . . . . . 135

5.42 Multiple absorption anomaly reconstruction example, deep, 60 dB SNR . . . . . 136

5.43 Multiple absorption anomaly reconstruction example, deep, 30 dB SNR . . . . . 136

5.44 Multiple anomaly MSE vs. SNR . . . . . 137

5.45 0.1% Variable background MSE vs. SNR . . . . . 138

5.46 0.2% Variable background MSE vs. SNR . . . . . 138

5.47 0.5% Variable background MSE vs. SNR . . . . . 138

5.48 1.0% Variable background MSE vs. SNR . . . . . 138

5.49 2.0% Variable background MSE vs. SNR . . . . . 139

5.50 5.0% Variable background MSE vs. SNR . . . . . 139

5.51 10.0% Variable background MSE vs. SNR . . . . . 139

5.52 Modified range approach MSE contour vs. SNR and background variation . . . . . 140

5.53 Modified model approach MSE contour vs. SNR and background variation . . . . . 140

5.54 Example reconstruction, unmodified approach 30 dB SNR, 1% background variation 141

5.55 Example reconstruction, modified range approach 30 dB SNR, 1% background variation . . . . . 141

5.56 Example reconstruction, modified model approach 30 dB SNR, 1% background variation . . . . . 141

# List of Tables

3.1	Truncated Conjugate Gradient algorithm . . . . .	60
5.1	Admissible solution analysis simulation parameters. . . . .	113
A.1	Table of Parameters . . . . .	149
B.1	The Photon Migration Imaging data structure. This is the top-level data structure that most of the high-level operations take as an argument and return. It contains all of the information and results from a particular experiment. . . . .	153
B.2	The Model data structure, this is used for both the forward and inverse problems. . . . .	154
B.3	The Chromophore data structure. This structure describes the concentration and resultant absorption coefficient of the specified chromophore. . . . .	156
B.4	The Optode data structure. This is used to specify the positions, orientation, and optical parameters of the sources and detectors. . . . .	156
B.5	The Boundary data structure. This describes the type of boundary model to use and any parameters that are needed with the type specified. . . . .	157
B.6	The Method data structure. This describes the method used to compute either the forward problem or forward operator for the inverse problem. . . . .	157
B.7	The Computational Volume data structure. This is used to specify the points at which the reconstruction will be performed and if appropriate where the forward solution will be calculated. . . . .	158
B.8	Object data structure. This data structure specifies the optical heterogeneities present in the medium. . . . .	159
B.9	The Noise data structure. This specifies the noise models to be used along with the weighting coefficients for whitening the system. . . . .	160
B.10	The Reconstruction data structure. This specifies the reconstruction technique and any parameters required. . . . .	162
B.11	The Visualization data structure. This specifies how to visualize the results . . . . .	163

# Chapter 1

## Introduction

Diffuse optical tomography (DOT) imaging is a relatively new technique for imaging inside biological tissues. By using near infrared (NIR) light this technology can potentially map, in three dimensions, the blood volume and oxygenation of biological tissues to depths of up to 10 cm. A number of significant technical hurdles must be solved before DOT imaging is at a level appropriate for widespread use in the medical community. Primary among them are the ill-posedness of DOT imaging in general and the typically underdetermined nature of the three-dimensional reconstruction problem. Both of these properties lead to poor quality reconstructions of optical absorption and ultimately of blood volume and oxygenation. To address these problems we have chosen an *a priori* constrained reconstruction approach, which we have implemented through several different algorithms. Our constrained reconstruction techniques have significantly enhanced the reconstruction fidelity of both the position and magnitude of the optical absorption coefficient in simulation studies.

The use of NIR light to image inside the human body has received significant interest in the last

decade<sup>1</sup>. The interest in DOT has been motivated by several factors. These include:

- the ability of NIR light to penetrate relatively deeply (*i.e.* several centimeters) into tissue without causing harm to the tissue
- the sensitivity of IR light to the physiologically important concentrations of oxygen-carrying components of blood, oxy-hemoglobin and deoxy-hemoglobin
- the low relative cost of the equipment for this type of imaging, compared to Computerized Tomography (CT) or Magnetic Resonance Imaging (MRI) for example.

To expand on these three factors, we note that near IR light (700-900nm wavelength) can propagate significant distances within tissue because of the relatively small rate of absorption by water and hemoglobin at these wavelengths. However, as we will explore in detail later in the thesis, although the absorption of photons in this band is reduced compared to that at surrounding wavelengths, the amount of scattering at these wavelengths is significant. Photons typically only propagate on the order of 10-100 micrometers before being scattered. This highly scattering medium, also known as a *diffuse* medium, requires a different approach to modeling than is typically applied in wave propagation imaging such as radar, sonar, or in other fields of optics.

The sensitivity of DOT imaging to the oxygen carrying components of blood results from two facts. First both oxy-hemoglobin and deoxy-hemoglobin have optical absorption coefficients at these wavelengths which, although low, are higher than that of water in typical concentrations found in tissue. Additionally, around 800nm oxy-hemoglobin and deoxy-hemoglobin exchange

---

<sup>1</sup>As with many new technologies being investigated by many different groups of researchers, this imaging technique is known by several names such as diffuse optical tomography (DOT), photon migration imaging (PMI), and diffuse photon density wave (DPDW) imaging. We will use the first acronym to identify this technology throughout this thesis.

roles as the primary absorber, allowing for the spectroscopic discrimination of the two species of hemoglobin.

The low relative cost of the equipment results from the common availability of inexpensive IR sources and detectors. The relative cost of the optics and electronics may vary considerably, however, depending on whether one uses low-cost, steady-state frequency domain techniques, or relatively high-cost, early-arrival time-domain techniques that require short pulses and fast detectors [1, 2].

## 1.1 DOT imaging

The unique sensitivity of IR light to the two types of hemoglobin provides for several interesting medical applications. The simplest of these devices, in clinical use for more than ten years now, are pulse oximeters. These devices measure the arterial oxygen saturation of the tissue under the measurement region, rejecting the response from other hemoglobin sources, without imaging. In other words, they simply measure an average property of the tissue between transmitter and receiver, using the presence of the pulsatile blood to discriminate arterial from venous and background blood and measurements at two frequencies to discriminate oxy-hemoglobin from deoxy-hemoglobin. By “imaging”, in contrast, we mean the use of multiple sources and receivers to sample a region of tissue in such a way that we can use mathematical algorithms to reconstruct a *localized distribution*, rather than bulk average, of absorption.

The motivating goal of DOT imaging as we apply it here is to generate a three-dimensional, quantitative map of both the fraction and oxygenation of blood within a region of biological tissue.

These quantities can be inferred from the spatial distribution of the optical absorption coefficient provided that measurements are made at at least two wavelengths within the 700-900 nm range. With measurements at a single wavelength, relative or qualitative observations can be made. With imaging, a wide variety of potential new applications includes anatomical and functional brain imaging and detection of carcinomas. Human brain imaging can be grouped into at least two types of applications. In one, DOT imaging is used to assess the size and nature of a stroke. Within this application area, research is also being conducted on using DOT imaging to assess the appropriateness of pharmacological intervention, which for certain drugs need to be administered immediately following a stroke [3, 4]. The other type of application to brain imaging involves localization of cerebral function by imaging hemodynamics in the outer layers of the brain. Brain imaging can also be applied to neonates with better imaging depth due to the thinner skull of a newborn. This is especially useful in the monitoring of brain oxygen levels in premature infants [5].

The other major application area of current interest besides brain imaging is the detection of malignant breast tumors. Breast cancer detection is the application that has been the primary motivation for the investigation presented in this thesis. This motivation derives from the expected difference in the functional parameters of cancerous tissue with respect to normal tissue. Specifically, greater vascularization of tumors, and lower blood oxygenation, can indicate the presence and even perhaps aggressiveness of malignant carcinomas [6, 7, 8]. By taking advantage of its ability to map oxygen saturation and blood volume, researchers hope to use DOT imaging to identify cancerous regions within the human breast.

There are several important aspects of DOT imaging that deserve mentioning. First, unlike other imaging modalities such as computerized tomography (CT), magnetic resonance imaging (MRI),

ultrasound, or positron emission tomography (PET), IR light propagates in a highly scattered manner. This results in an inherently three-dimensional problem even if the sources and detectors are restricted to two dimensions. Second, a typical measurement scenario involves amplitude modulating the IR light source with a radio frequency (RF) sinusoid. This results in a “photon density” wave at the RF modulation frequency that can be modeled as the envelope of the electromagnetic wave. This photon density wave propagation can, in the proper circumstances, be well modeled by the diffusion equation, hence the name diffuse photon density wave. The resulting wavelength of the diffuse photon density wave is dependent upon the optical parameters (absorption and scattering coefficient) of the tissue, the index of refraction, and the modulation frequency, but typical scenarios in tissue result in wavelengths that are much larger than both the objects of interest and the spacing between sources and detectors (*i.e.* wavelengths on the order of tens of centimeters). Thus we are dealing with an effectively near-field imaging problem. A final aspect worth mentioning is that to compute quantitative information about the blood volume and oxygenation from the absorption coefficient, measurements need to be made at multiple wavelengths. This is necessary to discriminate the absorption responses from oxy-hemoglobin and deoxy-hemoglobin, using spectroscopic techniques as mentioned above.

## 1.2 Technical challenges of DOT imaging

Reconstruction of the optical absorption coefficient presents several difficult technical problems which are typical of many inverse scattering problems. The problem is ill-posed due to two factors. First, due to the physics, there is a high degree of attenuation experienced by the diffusive wave as

it propagates through the medium. Second, a typical measurement is affected by a large amount of tissue. A perturbation in the optical absorption coefficient in a region will have similar effects on most measurements made in the vicinity of that region. One way to think of this is that the large amount of photon scattering results in a blurring of the diffuse photon density wave. The general consequence of both factors is that inverse solutions need to amplify small differences in the measurements to recover detail in the reconstructions, and thus are extremely sensitive to noise in the measurements and error in assumed models and parameters. One specific consequence of the ill-posedness is the resulting ill-conditioning of the forward model matrix when a linear approximation is used.

Another significant technical problem is that the linear problem is generally severely underdetermined for three-dimensional reconstructions. Since a typical measurement is affected by a large volume, we must account for any small region in the large volume which may affect the measurements. Thus using a reasonable reconstruction voxel size results in many voxels per measurement. We are limited in techniques that we can apply to directly address this aspect of the measurements. First, there is a practical limit to increasing the number of measurements to reduce the underdetermined nature of the problem. This is a consequence of the ill-posedness of the problem. Specifically, trying to address the underdeterminedness by making a large number of measurements in a small region does not help much because each measurement “sees” a very similar environment. Another approach could be to use a wide range of source and detector locations to help address the underdeterminedness. But since we are limited to placing the sources and detectors on the skin surface, we cannot achieve a large amount of spatial heterogeneity in the sources and detectors.

## 1.3 Background

The development of DOT imaging is based on investigations in a number of areas, some closely related, such as the modeling of photon propagation in diffuse media, and others indirectly related, such as regularization and convex optimization. In the following subsection we explore some of the important contributions in fields relevant to DOT imaging.

### 1.3.1 Modeling of photon propagation in highly scattering media

Direct modeling of electromagnetic propagation in many applications generally follows one of two models, either based on Maxwell's equations describing the electric and magnetic fields, or based on geometric optics. The choice of model depends on the size of the objects the radiation is interacting with, relative to the wavelength. However neither of these methods is computationally feasible when the number of distinct interactions being considered is very large, as is the case with the highly scattering medium we are trying to model. Rather we require a technique that models this large number of interactions by some aggregate approach.

One technique that has been developed to deal with a large number of field-media interactions is linear transport theory [9, 10]. In this approach light is treated as particles propagating through the medium that can be either absorbed or scattered by the particles composing the medium. Light particles are not modeled as interacting with each other, and thus there is no correlation between the fields the particles represent. The implication of this approach is that powers, instead of fields, are additive. This model also does not take into account polarization effects<sup>2</sup>. The equation defining

---

<sup>2</sup>A model that does take into account polarization is available [10] but it is not typically employed in the derivation of the diffusion equation for light propagation in a diffuse medium, our ultimate goal here.

linear transport theory, known synonymously as the linear transport equation, the radiative transport equation, or the Boltzmann transport equation, is a conservation equation of the light particles and is given by

$$\mu_s \int_{4\pi} f(\hat{\Omega}, \hat{\Omega}') L(\mathbf{r}, \hat{\Omega}', t) d\hat{\Omega}' - \mu_t L(\mathbf{r}, \hat{\Omega}, t) + Q(\mathbf{r}, \hat{\Omega}, t), = \frac{1}{v} \frac{\partial L(\mathbf{r}, \hat{\Omega}, t)}{\partial t} + \nabla \cdot L(\mathbf{r}, \hat{\Omega}, t) \hat{\Omega}, \quad (1.1)$$

where  $L(\mathbf{r}, \hat{\Omega}, t)$  is the radiance (the power per unit area and unit solid angle) at position  $\mathbf{r}$  in the direction  $\hat{\Omega}$  at time  $t$ ;  $\mu_t = \mu_s + \mu_a$  are the optical transport, scattering, and absorption coefficients respectively;  $f(\hat{\Omega}, \hat{\Omega}')$  is the scattering phase function;  $Q(\mathbf{r}, \hat{\Omega}, t)$  is the radiant source function; and  $v$  is the electromagnetic propagation speed in the medium. Although the linear transport equation is applicable to a wide range of media, analytical solutions are only available for simple scenarios because of the integro-differential structure of the equation. Numerical solutions to the linear transport equation are computationally intensive due to the dependence on space, angle, and time.

### Photon diffusion equation

We can use the diffusion equation as an approximate solution to model the light propagation described by the linear transport equation if two criteria are met: that the reduced scattering coefficient is much larger than the absorption coefficient and that the reduced scattering coefficient does not have significant spatial variation. The reduced scattering coefficient is given by

$$\mu'_s = (1 - g)\mu_s \quad (1.2)$$

where  $g$  is the mean cosine of the phase function given by

$$g = \int_{4\pi} f(\hat{\Omega}, \hat{\Omega}') \hat{\Omega} \cdot \hat{\Omega}' d\hat{\Omega}'. \quad (1.3)$$

The reduced scattering coefficient specifies the equivalent scattering rate to achieve a uniformly random scattering function.

Detailed derivations of the photon diffusion equation from the linear transport equation are given in Ishimaru [10], Haskell et al. [11], Boas [12], and Arridge [13] among others. We briefly show where these important assumptions apply in developing the diffusion equation below.

If the reduced scattering coefficient is much greater than absorption, the radiance can be expressed by the first two terms of the Taylor series for the radiance given by

$$L(\mathbf{r}, \Omega, t) = \frac{1}{4\pi} \Phi(\mathbf{r}, t) + \frac{3}{4\pi} \mathbf{J}(\mathbf{r}, t) \cdot \hat{\Omega}, \quad (1.4)$$

where  $\Phi(\mathbf{r}, t)$  is the photon fluence rate, which is the integral of the radiance without respect to direction

$$\Phi(\mathbf{r}, t) = \int_{4\pi} L(\mathbf{r}, \hat{\Omega}, t) d\hat{\Omega}, \quad (1.5)$$

and  $\mathbf{J}(\mathbf{r}, t)$  is the photon flux; that is, the first order directional component of the radiance

$$\mathbf{J}(\mathbf{r}, t) = \int_{4\pi} L(\mathbf{r}, \hat{\Omega}, t) \hat{\Omega} d\hat{\Omega}. \quad (1.6)$$

Requiring that the reduced scattering coefficient be much greater than the absorption coefficient results in a radiance that is almost uniform, having only a relatively small flux in the direction  $\hat{\Omega}$ . Expressing the radiance in this form allows for the simplification of the linear transport equation to the form

$$-\nabla \cdot D \nabla \Phi(\mathbf{r}, t) + v \mu_a \Phi(\mathbf{r}, t) + \frac{\partial \Phi(\mathbf{r}, t)}{\partial t} = v S(\mathbf{r}, t). \quad (1.7)$$

where  $S(\mathbf{r}, t)$  is the isotropic source component given by

$$S(\mathbf{r}, t) = \int_{4\pi} Q(\mathbf{r}, \Omega, t) d\Omega, \quad (1.8)$$

and  $D$  is the diffusion coefficient given by

$$D = \frac{v}{3[(1-g)\mu_s]}. \quad (1.9)$$

If we can assume that  $D$  is invariant with respect to space, which allows us to move  $D$  outside of the Laplacian operator, we arrive at the photon diffusion equation

$$-D\nabla^2\Phi(\mathbf{r}, t) + v\mu_a\Phi(\mathbf{r}, t) + \frac{\partial\Phi(\mathbf{r}, t)}{\partial t} = vS(\mathbf{r}, t). \quad (1.10)$$

There is also some discussion in the community as to whether the denominator of the diffusion coefficient should also contain an additive term  $3\mu_a$  [14, 15]. In our work we have chosen to use the expression Eq. (1.9) [14].

Several aspects of this model deserve additional discussion. Near the boundary between a diffuse and non-diffuse medium (*i.e.* the air-tissue interface) the diffusion approximation does not hold, since any photons scattered out of the medium will never have a chance of being scattered back into the medium. This must be handled as a special case as described in Section 1.3.4. Additionally, the distance between sources and detectors should be much greater than the mean transport length  $1/(\mu_s' + \mu_a)$  so that enough scattering events occur to generate a diffuse field. Finally, explicit use of the linear transport equation does become necessary when some region of the medium being interrogated is not diffuse, for example the cerebral spinal fluid in the head [16].

### Frequency domain photon diffusion equation

Taking the Fourier transform of Eq. (1.10) with respect to time gives the frequency domain photon diffusion equation which is in the form of the Helmholtz equation

$$\left[\nabla^2 + k^2\right]\Phi(\mathbf{r}, \omega) = \frac{-v}{D}S(\mathbf{r}, \omega), \quad (1.11)$$

where  $k$  is the complex wavenumber given by,

$$k^2 = \frac{-v\mu_a + j\omega}{D} = 3\mu'_s \left( -\mu_a + j\frac{\omega}{v} \right) \quad (1.12)$$

Note that Eq. (1.12) is derived using the Fourier transform defined as  $F(\omega) = \int_{-\infty}^{\infty} f(t)e^{j\omega t} dt$ .

This is a slightly different form than is commonly used in engineering texts in that the  $j\omega t$  term is positive. This has the consequence that the Fourier transform of the time derivative term becomes  $-j\omega\Phi(r, \omega)$  instead of  $j\omega\Phi(r, \omega)$ .

### 1.3.2 Validation of the photon diffusion model

Early work in investigating the diffusion model of photon propagation focused on validating the model with experimental evidence. O’Leary *et al.* [17] and Boas [12] showed that diffuse photon density waves refract at a boundary just like electromagnetic waves. They conducted an experiment in which a two layer medium of differing optical scattering coefficients was probed using amplitude modulated diffusive light. By measuring the amplitude and phase of the modulated photon density wave throughout the volume they were able to show good agreement with expected refraction at the boundary layer predicted by the Helmholtz equation. Fishkin and Gratton [18] performed a similar experiment showing that diffuse photon density waves diffract at an absorbing straight edge. As in the experiment by O’Leary *et al.* [17], they used an amplitude modulated light source in a highly scattering medium. In the medium they placed an absorbing straight edge and measured the amplitude and phase in the volume. They used both the Helmholtz equation and a Monte Carlo model to predict their results and achieved good agreement with their measurements.

### 1.3.3 Solutions to the diffusion model

Solutions to the diffusion equation for basic geometries have been explored by several groups. Boas *et al.* [12, 19] presented a spherical harmonic solution for the case where a spherical inhomogeneity is present in an infinite medium. In this instance a closed form infinite series solution can be expressed and a reasonable approximation can be calculated numerically. Additionally they compared their calculations to experimental results where a perfect absorber was placed in a large diffuse region. They were able to show good agreement between the experimental measurements and those predicted by their analytical solution. Walker *et al.* [20] extended the spherical results described above to a cylindrical inhomogeneity. They were able to show good agreement between an analytical solution to the diffusion equation and experimental measurements made in an intralipid phantom.

### 1.3.4 Boundary modeling and diffuse sources

Sources used for DOT are typically laser beams incident on the diffuse medium. This creates a complex source function that is typically approximated by an isotropic point source one transport length  $l_{tr} = 1/\mu'_s$  into the medium [11, 12]. Correct modeling of any boundary between diffuse and non-diffuse regions is of great importance to accurately predict the fluence rate distribution in the medium. This is particularly important when the sources and detectors are near a diffuse/non-diffuse interface such as the case in which we are interested. Haskell *et al.* [11] compared three models for the air-tissue boundary: a zero boundary condition, an extrapolated zero boundary condition and a partial current boundary condition. The partial current boundary condition is physically the most

rigorous of the three but difficult to implement computationally. They showed that the extrapolated boundary condition matches the partial current boundary condition up to the quadrupole moments. Thus we have based our simulations on the extrapolated boundary condition. The extrapolated boundary condition implements a zero fluence rate boundary at specified distance from the true air-tissue interface. This distance is a function of the optical parameters and the change in index of refraction between the two media. The expression for the extrapolated boundary distance is

$$z_b = \frac{2}{3} \left( \frac{1}{\mu'_s + \mu_a} \right) \cdot \left( \frac{1 + R_{eff}}{1 - R_{eff}} \right) \quad (1.13)$$

where  $R_{eff}$  is the effective Fresnel reflection at the surface given by

$$R_{eff} = \frac{R_\phi + R_j}{2 - R_\phi + R_j}, \quad (1.14)$$

$$R_\phi = \int_0^{\pi/2} 2 \sin(\theta) \cos(\theta) R_{Fresnel}(\theta) d\theta, \quad (1.15)$$

$$R_j = \int_0^{\pi/2} 3 \sin(\theta) \cos^2(\theta) R_{Fresnel}(\theta) d\theta, \quad (1.16)$$

and  $R_{Fresnel}$  is the Fresnel reflection coefficient for unpolarized light.

Pogue and Patterson [21] examined the effect of inaccurately modeling the shape of the boundary of the diffuse medium. Specifically, they modeled finite volumes consisting of slabs, cylinders, and spheres and then compared the reconstruction of the optical properties from a half space boundary assumption. When the volume of the diffuse medium was larger than 5 cm in width for the slab geometry and 5 cm in diameter for the cylindrical and spherical geometries, they found good agreement between semi-infinite medium assumption and the true solution. In particular, for a medium with a scattering coefficient of  $1.0 \text{ mm}^{-1}$  and absorption coefficient of  $0.01 \text{ mm}^{-1}$  they reported errors in the scattering coefficient estimate of less than  $0.2 \text{ mm}^{-1}$  for the spherical geometry and

much less than that for the slab geometry. The absorption coefficient estimates for the same scenario were within  $0.003 \text{ mm}^{-1}$  for the spherical geometry and much less again for the slab geometry.

### 1.3.5 Optical characteristics of biological tissue

The optical characteristics of biological tissues vary widely depending upon the type of tissue and have been studied in many investigations. The quantities of interest here are the optical scattering coefficient,  $\mu_s$ , the mean cosine of the scattering phase function,  $g$ , (or equivalently  $\mu_s^l$ ), and the absorption coefficient,  $\mu_a$ .

The scattering coefficient varies for both tissue type and as a function of wavelength. Cheong *et al.* [22] compiled a comprehensive set of optical parameter coefficients from a large number of earlier publications. Most of the data compiled at that point was either in the visible range, 500-650 nm, or at infrared wavelengths over 1000 nm. Mitic *et al.* [2] presented *in vivo* measured data of breast tissue from six female volunteers using a time-gated approach at 800 nm. They reported a reduce scattering coefficient of 7.2-10.0  $\text{cm}^{-1}$ . Bevilacqua *et al.* [23] presented *in vivo* results from human brain measurements. They examined several types of brain tissue at 674, 811, 849 and 956 nm. At the two wavelengths nearest the ones we have employed here, 780 and 830 nm, they found reduced scattering coefficients of 8.2-9.5  $\text{cm}^{-1}$  at 811 nm, and 7.6-9.2  $\text{cm}^{-1}$  at 849 nm.

The optical absorption coefficient, at our wavelengths of interest, is primarily affected by the concentration and type of hemoglobin present in the tissue being examined. Mitic *et al.* [2] found bulk optical absorption coefficients of slightly compressed breast tissue in the range of 0.017-0.032  $\text{cm}^{-1}$ . If we assume that the absorption coefficient is entirely due to hemoglobin then the absorption coefficient of tissue for a given blood volume and oxygenation can be computed using the optical

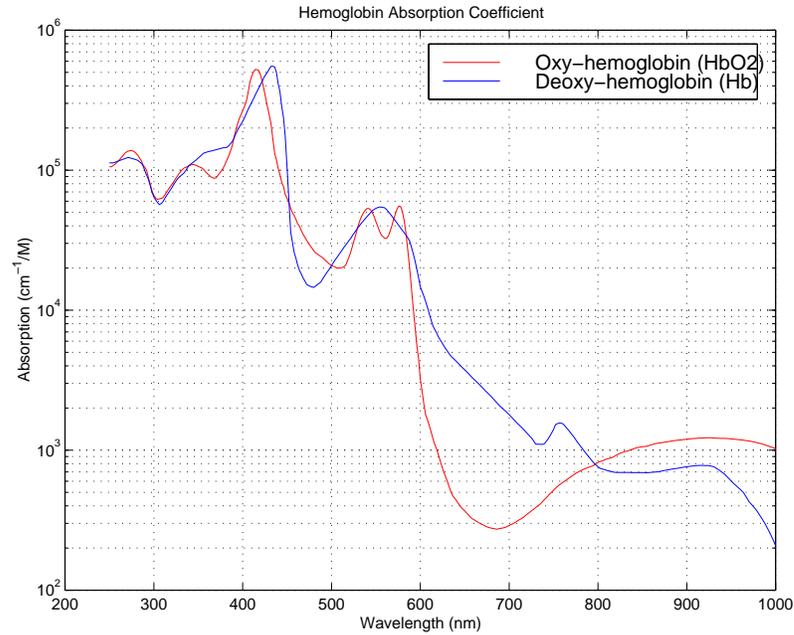


Figure 1.1: Hemoglobin absorption coefficient per mole as a function of wavelength. Note the relatively low absorption between 700 and 1000 nm and the crossover point around 800 nm. Data taken from Prahl [24]

coefficients for pure blood. Prahl [24] compiled data from variety of sources to present the absorption coefficient of whole oxy-hemoglobin and deoxy-hemoglobin versus wavelength, as shown in Figure 1.1.

### 1.3.6 Imaging using DOT

Imaging using DOT consists of computing a spatial map of the optical scattering coefficient, absorption coefficient, or both. From these maps other biological characteristics are often derived, such as a map of blood volume. Imaging of the optical characteristics can be broadly classified into two categories, linear model based imaging and non-linear imaging. In general, linear model

based imaging attempts to use a linearized model approximation of the photon diffusion equation (Eq. 1.10) to create a simpler relationship between the measured radiance and the medium optical parameters. On the other hand, non-linear techniques typically use the photon diffusion equation directly and, through an iterative process, try to find the optical parameter map that best matches the measured data (and possibly meets some other additional criteria).

### **Linear model based imaging techniques**

Linear model based imaging techniques typically employ either the first order Born or Rytov approximations [25] to create a linear relationship between the optical parameter of interest and the measured light. In all of the published studies we have examined, only the absorption coefficient has been imaged using linear model based techniques.

Some of these early attempts include the work by O’Leary *et al.* [26, 27], who examined reconstruction of both the absorption and scattering coefficient. They used both the SIRT algorithm with a positivity constraint and a filtered SVD algorithm to compute their reconstructions. They were able to show that reconstruction of absorption images from frequency domain data is possible.

Another approach used early on was to use interfering diffuse photon density waves to detect anomalies within the medium. This technique was pursued by Chance *et al.* [4] and Knuttel [28]. While this technique provided early confirmation that absorbing inhomogeneities could be detected, it suffers from ambiguity when the *a priori* information about the number of objects is unknown.

A collaboration between groups at Polytechnic University in Brooklyn, NY, SUNY Health Science Center in Brooklyn, and Bioimaging Sciences in New Jersey has been active in developing linear DOT reconstruction techniques. In one series of articles [29, 30, 31, 32, 33, 34, 35] they

explored the use of a Monte Carlo based transport theory forward model and several different linear reconstruction techniques to compute maps of the absorption coefficient for a cylindrical diffuse region. In this work they simulated detector responses for scenarios where there was either a single cylindrical absorbing inhomogeneity or 13 cylindrical absorbing inhomogeneities. In all cases the absorption coefficient of the inhomogeneities was infinite. In Graber *et al.* [29] they explored the effect of modulation frequency on reconstruction quality using the Simultaneous Algebraic Reconstruction Technique (SART) algorithm. They looked at modulation frequencies of DC, 1 KHz, 1 MHz, 100 MHz, 1 GHz, 10 GHz, 100 GHz and 1 THz and found that a 1 GHz modulation frequency provided the most accurate reconstruction of a single off-center rod. Chang *et al.* [31] explored the performance of three different linear reconstruction algorithms, Projection on Convex Sets (POCS), Conjugate Gradient Descent (CGD), and SART. Using these algorithms applied to three-dimensional reconstructions, two-dimensional reconstructions where the axial invariance is known *a priori*, and two-dimensional reconstructions where it is assumed that only those voxels in the plane of the detector contribute to the detector measurements, they showed that the three-dimensional reconstructions provide better images (qualitatively) than the two-dimensional approaches.

In Chang *et al.* [32] they explored a positivity constraint on the CGD algorithm and a rescaling of the forward matrix. In previous work they had applied a positivity constraint by setting any negative value in an intermediate solution to zero. In this work they examined the constrained solution at each stage to detect divergence (an increasing residual). If they detected divergence they would restart the CGD algorithm at the current estimate. Their rescaling of the forward matrix was accomplished by rescaling each column to have a maximum or average of one. Their results

showed that to apply a positivity constraint in this manner (setting negative values to zero) requires the restarting of the CGD algorithm. Without restarting, their images did not converge to anything resembling the true absorption function. With restarting they were able to generate good quality reconstructions. Also, the rescaling of the forward matrix produced better quality images than the CGD algorithm with the unscaled forward matrix. They combined most of the previously described results in the journal article [33].

In Zhu *et al.* [34, 35] they examined the total least squares (TLS) solution to the linear perturbation problem. They looked at simulations of one-dimensional and two-dimensional variations and attempted to reconstruct in two-dimensions using a TLS algorithm. They found that the TLS solution provided better performance than the least squares solution for the one-dimensional variation and no improvement in the case where there was two-dimensional variation. They attributed this result to the fact that they solved an overdetermined system for the one-dimensional variation and an underdetermined system for the two-dimensional variation.

### **Non-linear imaging techniques**

A number of groups have explored non-linear reconstruction techniques. In these approaches the researchers have attempted to retain (or approximate) the exact relationship between the measured fluence rate data and the optical properties of the medium. This approach typically requires a larger amount of computation than an equivalent linear approach. Additionally the solution is typically computed through a minimization algorithm and therefore the solution is not guaranteed to be a global minimum. Jiang *et al.* in [36] presented a detailed review of their group's work in non-linear frequency-domain DOT reconstruction. They used a finite element forward model of a cylindrical

diffuse medium with a cylindrical inhomogeneity. They reconstructed using both simulations and experimental data from a phantom. In both cases they modeled an environment that was invariant in the axial dimension and reconstructed a two-dimensional slice through the volume. The 16 sources and 16 detectors were evenly distributed in a ring around the boundary of the cylinder. They were able to reconstruct quantitatively accurate maps of both the absorption and reduced scattering coefficient using both simulated and measured data. Additionally, they examined the effect of varying the modulation frequency between 50 and 300 MHz and found very little variation in the quality of the reconstructions.

The group based in Brooklyn, previously mentioned in the section on linear techniques, has also explored a non-linear technique based on the Born iterative method. In Yao *et al.* [37] they described their implementation of the Born iterative approach and showed noise free two-dimensional reconstructions that displayed some improvement over a first order Born reconstruction. They also demonstrated (qualitatively) the degradation in the reconstruction resulting from the addition of 10% noise to the simulated data. Finally, they demonstrated the ability to reconstruct absorption and scattering images from a set of sources and detectors that only partially spanned the circumference of the cylinder. In Yao *et al.* [38] they described their implementation of a multigrid forward solver to the diffusion equation and its application to an anatomically accurate geometry derived from MR based data. They simulated optical inclusions into the geometry and showed they were able to qualitatively reconstruct the optical parameter maps. In Zhu *et al.* [39] they described a wavelet based approach to the inversion step of their reconstruction algorithm. They solved the inversion at a coarse scale using the CGD algorithm, then using the result from the coarse scale to initialize the CGD algorithm at the next finer scale they computed the result at that scale. This

process is iterated until the desired precision is reached. The motivation behind this approach is to reduce the computation for the inversion step. They show an example where the total computation is reduced by a factor of seven compared to non-multiresolution reconstruction. In Yao *et al.* [40] they combine all of these results into a single article.

A statistical approach to the non-linear DOT reconstruction is described by Eppstein *et al.* [41]. In this approach they used an extended Kalman filter to estimate the absorption coefficient in a recursive scheme which they originally developed for geohydrology. They also attempted to reduce the number of unknowns at each iteration by grouping like voxels together as a large voxel with the same properties moving the problem from an underdetermined to an overdetermined one.

Parametric approaches have been explored by Kilmer *et al.* [42] and Arridge *et al.* [43]. In [42] the authors take a parametric approach to extracting information about the anomaly present in the medium. The anomaly boundary is modeled by a B-spline curve that is updated through a non-linear optimization approach. In two-dimensional reconstructions they were able to accurately estimate the size and optical absorption coefficient of a 1 cm deep object.

### 1.3.7 Related work in other disciplines

Many of the approaches for imaging are derived from other applications of tomography. A primary reference for our research into tomography has been the monograph Computerized Tomographic Imaging by Kak and Slaney [25] in which they presented the algebraic reconstruction techniques (Algebraic Reconstruction Technique (ART), Simultaneous Iterative Reconstruction Technique (SIRT), etc.) and the first Born approximation, of which we have made extensive use in our

investigation. Many of the matrix techniques that we have employed are well described in *Matrix Computations* by Golub and Van Loan [44]. Specifically, we have referenced this text for the Givens rotation algorithm, the QR formulation of the minimum norm solution to an underdetermined system, the Truncated Singular Value Decomposition (TSVD) and the Conjugate Gradient (CG) algorithm. Shewchuk [45] also presented an accessible introduction to the CG algorithm.

P. C. Hansen has written extensively on the topic of regularization. The L-curve as a method for choosing a regularization parameter is presented in [46] and references therein. His group has also presented a modification to the TSVD algorithm that we have adapted for DOT imaging. The modified (M)TSVD algorithm [47] presents an algorithm to minimize a weighted semi-norm (2-norm) of the solution subject to the solution matching a regularized version of the data.

The Ellipsoid algorithm is used extensively in the latter part of our study. This algorithm was originally presented by Khachiyan [48] as a polynomial time solution for linear programming. Bland *et al.* [49] present a detailed survey of the ellipsoid algorithm, its development, and its application, particularly to linear programming. Boyd and Barratt [50] present the ellipsoid algorithm in the context of convex optimization with applications in control systems. It has also been used in our lab to implement an admissible solution approach for the electrocardiographic inverse problem [51].

## 1.4 Contributions of this thesis

Our approach, described in this thesis, has been to try to understand the limitations and failings of the basic linear model reconstruction approach. With those limitations and failings in mind we have

developed a set of constraints, and algorithms to apply them, that improve reconstruction fidelity of the optical absorption coefficient. The work described in this thesis has been done as part of a collaboration among several different groups here at Northeastern University, Tufts University, and Massachusetts General Hospital. The participants in this group include Dana Brooks, Eric Miller, and Chuck DiMarzio from Northeastern University, Misha Kilmer from Tufts University, and David Boas and Tom Gaudette from Massachusetts General Hospital.

In this thesis we have made three primary contributions to DOT imaging. First, we have provided a detailed analysis of the performance and limits of several three-dimensional linear model reconstruction algorithms. In this analysis we showed that a class of algorithms which has been less commonly used for this particular application, based on subspace techniques, outperformed the more commonly employed algebraic techniques. In addition, in this work we explored several different measures of performance for the reconstruction algorithms and compared them through a thorough subjective analysis of the reconstructions. From this analysis we concluded that both mean square error and position-based object centroid error provide useful although incomplete measures of algorithm performance. This work has been published in Refs. [52, 53]. A second area of contribution has been in the development of a constrained two-wavelength reconstruction algorithm. In this approach we have exploited the existing requirement that multiple wavelengths need to be collected to calculate blood volume and oxygenation. We have developed an efficient algorithm for applying an *a priori* constraint on the solutions at each wavelength that improves the reconstruction performance compared to the individual linear model reconstructions. The final significant area of contribution is the application of more general constraints to linear model reconstructions through an admissible solution approach. By applying these *a priori* constraints, derived from the expected

behavior of the optical absorption coefficient, we have been able to dramatically improve the fidelity of the reconstructions over a wide range of simulation scenarios. In addition we have made two other contributions. We performed a preliminary comparison of several types of forward models for DOT imaging, with the primary goal of testing the validity of linear (first Born approximation) models, and we designed and constructed an extensible, MATLAB-based toolbox for use in DOT imaging research.

This thesis is structured as follows: Chapter 2 provides a description of the simulation models used for this investigation. Chapter 3 presents our evaluation of the linear model reconstruction techniques. Chapter 4 derives the constrained two wavelength approach and presents the results of applying this algorithm. Chapter 5 derives our admissible solution approach and compares the performance of this technique to the previously presented algorithms. In Chapter 6 we present the conclusions of this thesis and our suggestions for future work in the field.

## **Chapter 2**

# **Simulation models**

All of the results presented in this thesis are derived from computer simulations of several different DOT measurement scenarios. In this chapter we present the measurement scenarios and mathematical models we used to examine the DOT reconstruction algorithms we have developed. The measurement scenario consists of the geometric layout of the sources and detectors and their relationship to the diffuse medium. Also in this section we discuss the resulting sensitivity of particular geometries. The mathematical models consist of the diffuse photon density wave propagation equations and approximations, used for both simulation of the forward problem and for constructing a model for the inverse problem. We also describe the noise models we have used to simulate the noise expected in laboratory and clinical measurements.

## 2.1 Measurement scenario

The general measurement geometry explored in our work was a planar array of sources and detectors placed at the air tissue interface. The motivation for studying this geometry was the examination of the human female breast for the detection and characterization of tumors. It is expected that for this scenario the subject would be lying on her back with the source-detector array placed on top of the breast. A simple schematic of this scenario is shown in Figure 2.1. The sources, identified by the circles, are generally interspersed between the detectors. A set of measurements is made by exciting one source and measuring at all of the detectors, and the process is then repeated with a different source turned on. The sources can be excited by either DC current or a sinusoidally modulated current (in the tens to hundreds of MHz range) to produce a spatially and temporally varying photon density wave. If a modulated source is used, both the amplitude and phase are measured through the use of a quadrature demodulator. Additional sources of different wavelengths are simulated in some of the experiments. These can either be collocated with the original sources or in different locations at the air-tissue interface. All of these possibilities allow for a wide range of measurement scenarios, many of which we have explored in this thesis.

With this type of measurement geometry the spatial sensitivity of a particular source-detector pair traces out what are sometimes called “banana patterns.” Two examples of these are shown in Figures 2.2 and 2.3. Each image in these figures shows the sensitivity, over the spatial region being plotted relative to a unit source, of the specified source-detector pairing. Another interpretation of these figures is the relative probability for a photon to travel through a given region. The sensitivities are plotted on a dB scale with the range for each source-detector pair set at 100 dB. The first (upper

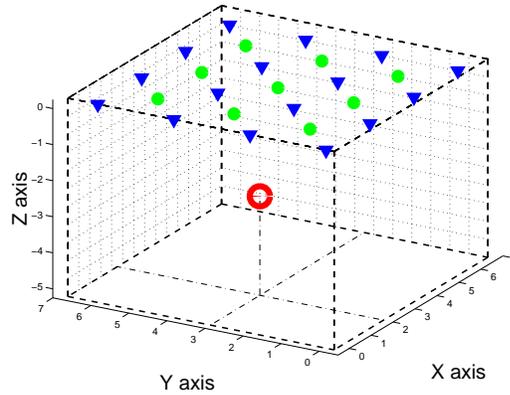


Figure 2.1: A schematic of the geometric layout for some of the simulations in this thesis. The sources (small green circles) and detectors (blue triangles) lie on the tissue-air boundary ( $z=0$  cm). The entire volume was 7 cm x 7 cm x 5.5 cm deep. The absorption anomaly (red sphere) is 2.5 cm into the medium at  $x=2$  cm and  $y=3$  cm with a radius of 1 cm. The grid size used to discretize the volume for some of the inverse computations is illustrated by the grid shown in the figure.

left) image in each figure shows the vertical plane passing through the line connecting the source and detector. Subsequent images show the sensitivity (on the same scale) for planes at increasing distance from the source-detector plane, using increments of 0.25 cm, labeled in the title of each image. The source was simulated at the origin of the coordinate system. Figure 2.2 shows the case where the detector is 1 cm away from the source at (1,0,0). This set of images shows that with a close source-detector separation the region sampled is relatively local to the source and detector. On the other hand, Figure 2.3 shows the case where the detector is 5 cm away from the source at (5,0,0). In this case we can see the volume sampled by this source-detector distance is relatively large. The surface where sensitivity falls to -40 dB relative to peak sensitivity is at about 3 cm deep into the medium in the source detector-plane and extends out to at least 2 cm out of the source-detector plane. This image displays clearly why DOT imaging is a three-dimensional problem. There is a

large diffusive wave interaction outside of the source-detector plane.

## 2.2 Propagation models

To understand how diffuse photon density waves interact with the optical parameters of the medium we needed a propagation model for the density wave. Both the forward problem (*i.e.* simulating measurements) and the inverse problem (*i.e.* creating maps of the optical parameters from measurements) require a model of diffuse photon density wave propagation through a heterogeneous medium. Furthermore, we have explored a linear model which is an approximation to the true non-linear relationship between the measured fluence rate and the spatially varying optical absorption parameter we wish to reconstruct. The motivation for exploring a linear approximation model was to develop a computationally less intensive reconstruction approach than is required by a full non-linear inversion. In this section we first present the details of the techniques we used to model diffuse photon density wave propagation. Following this we show how we have discretized the model so that it could be solved numerically. Finally, we present a validation of the model by comparing it to spherical harmonic and finite difference solutions of the full non-linear model.

The optical absorption function that represents the scenario described in the previous section can be expressed as

$$\mu_a(\mathbf{r}) = \mu_a^o + \Delta\mu_a(\mathbf{r}) \quad (2.1)$$

where  $\mu_a(\mathbf{r})$  describes the optical absorption as a function of space,  $r$ ,  $\mu_a^o$  specifies the background value of the absorption parameter, and  $\Delta\mu_a(\mathbf{r})$  specifies the spatial deviation of the absorption parameter. Substituting Eq. (2.1) into the equation for the spatial wavenumber in the frequency

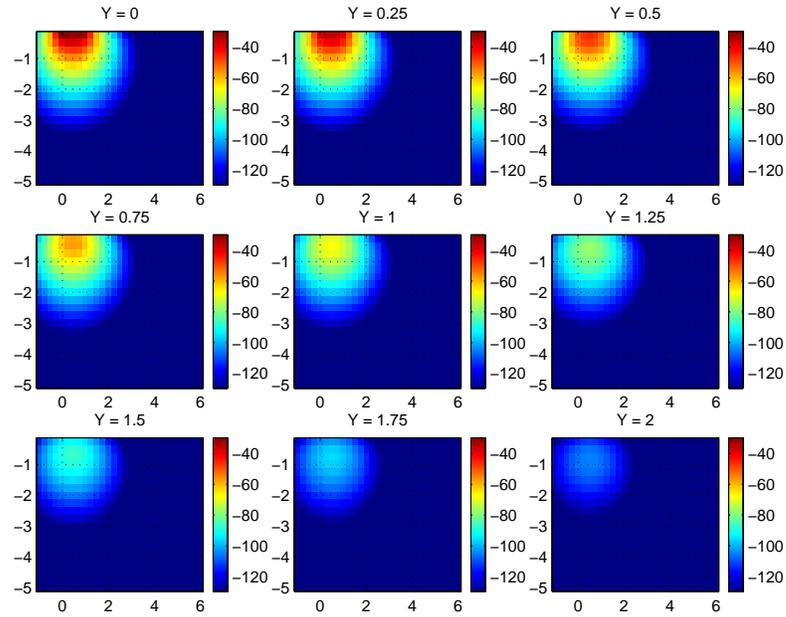


Figure 2.2: Source-detector sensitivity for a 1 cm source-detector separation.

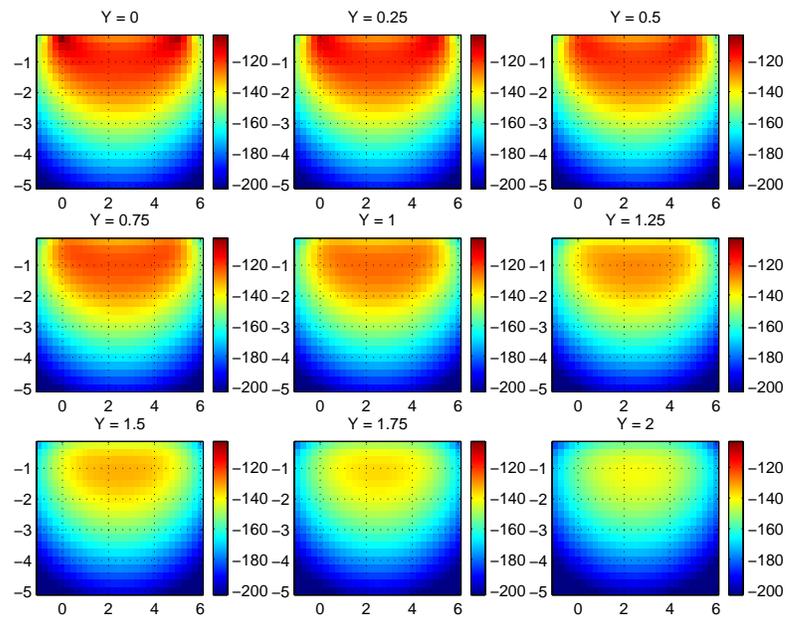


Figure 2.3: Source-detector sensitivity for a 5 cm source-detector separation.

domain photon diffusion equation Eq. (1.12) gives

$$k^2(\mathbf{r}) = k_o^2 + \Delta k^2(\mathbf{r}) = \frac{j\omega - v\mu_a^o}{D} + \frac{-v}{D}\Delta\mu_a(\mathbf{r}) \quad (2.2)$$

where we have expressed the squared wavenumber as the sum of a background component,  $k_o^2$ , due to the background optical parameters and  $\Delta k^2(\mathbf{r})$  due to the absorption deviation. Substituting Eq. (2.2) into the frequency domain photon diffusion equation Eq. (1.11) gives

$$\left[\nabla^2 + k_o^2 + \Delta k^2(\mathbf{r})\right] \Phi(\mathbf{r}) = \frac{-v}{D}S(\mathbf{r}). \quad (2.3)$$

Note that we have also dropped the explicit dependence on the temporal frequency  $\omega$  present in Eq. (1.11). Since our measurement procedure is to collect steady state frequency domain data at one or more discrete frequencies that we specify, we have left in  $\omega$  only as a parameter.

We have already described approaches previous groups have taken to solving the frequency domain diffusion equation for special cases in Section 1.3.3. When the variation in the absorption coefficient is not spherical or cylindrical and the change in spatial wavenumber is not too large compared to its average or background value, we can apply a perturbation approach to solve the Helmholtz equation Eq. (2.3) with the wavenumber given by Eq. (2.2). The perturbation approach that we chose to apply is known as the first Born approximation [25]. We first rewrite Eq. (2.3) as

$$\left[\nabla^2 + k_o^2 + \Delta k^2(\mathbf{r})\right] \left(\Phi^i(\mathbf{r}) + \Phi^s(\mathbf{r})\right) = \frac{-v}{D}S(\mathbf{r}), \quad (2.4)$$

where we have expanded the fluence rate as the sum of the incident fluence rate,  $\Phi^i(\mathbf{r})$  and a scattered fluence rate,  $\Phi^s(\mathbf{r})$ . The incident fluence rate is that resulting from the source acting on the medium with the spatially varying wavenumber parameter set to zero. The scattered fluence rate is the fluence rate remaining to correctly solve the complete equation and can be thought of as the fluence rate that scatters from the inhomogeneities.

If we consider only the source acting on the homogeneous medium we have the equation for the incident fluence rate given by

$$\left[\nabla^2 + k_o^2\right] \Phi^i(\mathbf{r}) = \frac{-v}{D} S(\mathbf{r}). \quad (2.5)$$

Subtracting the homogeneous medium equation Eq. (2.5) from (2.4) we are left with the scattered fluence rate equation

$$\left[\nabla^2 + k_o^2\right] \Phi^s(\mathbf{r}) = -\Delta k^2(\mathbf{r}) \left(\Phi^i(\mathbf{r}) + \Phi^s(\mathbf{r})\right). \quad (2.6)$$

Making the assumption that  $\Phi^i(\mathbf{r}) \gg \Phi^s(\mathbf{r})$  allows us to drop the  $-\Delta k^2(\mathbf{r})\Phi^s(\mathbf{r})$  term from Eq. (2.6). This is the first order Born approximation and results in

$$\left[\nabla^2 + k_o^2\right] \Phi^s(\mathbf{r}) \approx -\Delta k^2(\mathbf{r})\Phi^i(\mathbf{r}). \quad (2.7)$$

Thus we now have a linear relationship between the scattered fluence rate  $\Phi^s(\mathbf{r})$  and the variation in the absorption coefficient through  $\Delta k^2(\mathbf{r})$ . Physically this approximation amounts to treating each point in an inhomogeneity as if it existed in isolation from the rest of the inhomogeneity, ignoring the contributions of perturbations of the scattered fluence rate from one part of an inhomogeneity on the fluence rate incident on another part.

To solve Eq. (2.7) we have used the method of Green's functions [54] which when discretized (as described below) provides us with a system of linear equations relating the scattered fluence rate and optical absorption. The Green's function solution for Eq. (2.7) is given by

$$\Phi^s(\mathbf{r}) = \int_V G(\mathbf{r}, \mathbf{r}') \Phi^i(\mathbf{r}') \Delta k^2(\mathbf{r}') dr', \quad (2.8)$$

where the integral is over the entire volume  $V$ . To calculate the incident fluence rate in Eq. (2.7) we need to solve for the homogeneous medium given in Eq. (2.5). Again we use the Green's function

solution given by

$$\Phi^i(r) = \frac{-v}{D} \int_V G(\mathbf{r}, \mathbf{r}') S(\mathbf{r}') dr'. \quad (2.9)$$

In our simulations one simply computes only the scattered fluence rate; in practice, of course, this is usually impossible. Hence, in practice some approximate homogeneous fluence rate is estimated by first estimating the background optical parameters and then subtracting the homogeneous fluence rate estimate from the measurements.

For a free space geometry the Green's function for the Helmholtz equations Eq. (2.7) and Eq. (2.4) is given by [55]

$$G_f(\mathbf{r}, \mathbf{r}') = \frac{-1}{4\pi|\mathbf{r} - \mathbf{r}'|} e^{jk_o|\mathbf{r} - \mathbf{r}'|}. \quad (2.10)$$

Thus if we assume that the source function is modeled by an isotropic point source at  $\mathbf{r}_{src}$  then equation for the incident fluence rate is

$$\Phi^i(\mathbf{r}) = \frac{v}{D} \frac{1}{4\pi|\mathbf{r} - \mathbf{r}_{src}|} e^{jk_o|\mathbf{r} - \mathbf{r}_{src}|}. \quad (2.11)$$

For the half space geometry described above we can use the method of images to impose the Dirichlet boundary defined by the extrapolated boundary condition when we calculate both the incident and scattered fluence rates. A diagram of the geometry of the method of images under the extrapolated boundary condition is shown in Figure 2.4. This approach results in two sources for each source simulated, the original source and its image which has a negative amplitude with respect to the original source. Any anomalies in the medium also need to be reflected across the extrapolated boundary. For the reflected anomalies the sign of the amplitude of the object parameter is the same as the true object. This symmetry results in zero fluence rate in the extrapolated boundary, the condition we are trying to achieve.

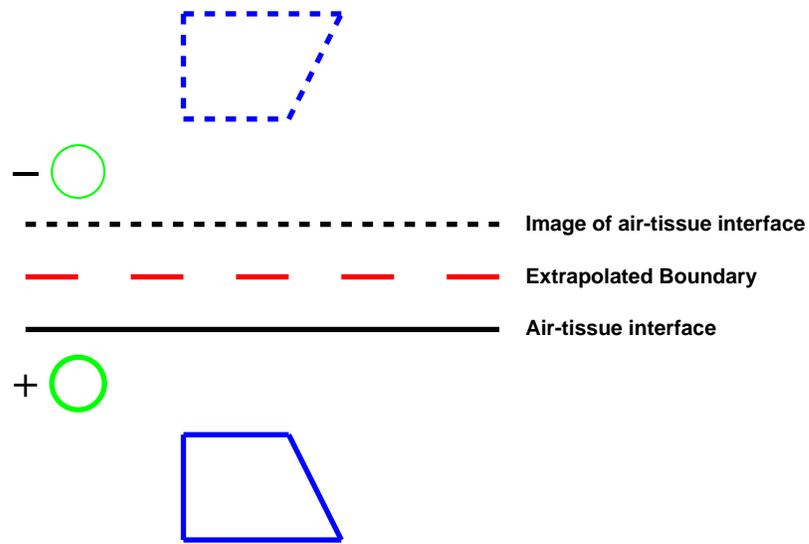


Figure 2.4: Method of images diagram for the extrapolated boundary half space. The true air-tissue interface is represented by solid black line, the dashed red line shows the position of the extrapolated boundary, the green circle represents a source, and the blue polygon represents a anomaly in the diffuse medium. The dotted line objects show the position and orientation of the images of the true object.

The method of images can also be implemented directly in the Green's function by adding a second term that accounts for the image source. This additional term results directly from the application of superposition of the two sources. The equation for the half space Green's function is given by

$$G_h(\mathbf{r}, \mathbf{r}') = \frac{-1}{4\pi|\mathbf{r} - \mathbf{r}'|} e^{jk_o|\mathbf{r} - \mathbf{r}'|} + \frac{1}{4\pi|\mathbf{r} - \mathbf{r}'_i|} e^{jk_o|\mathbf{r} - \mathbf{r}'_i|}, \quad (2.12)$$

where  $\mathbf{r}'_i$  is the image point of  $\mathbf{r}'$  with respect to the extrapolated boundary.

### 2.2.1 Discretization of the propagation model

We now have a linear approximation that can be used to compute the forward problem if  $k^2(\mathbf{r})$  is known and can also be used to construct an inverse problem if  $\Phi(\mathbf{r})$  is known and we wish to estimate  $k^2(\mathbf{r})$ . To solve Eq. (2.8) numerically we need to discretize the integral. We performed this by discretizing the volume into uniform voxels. The integral was approximated by assuming that the Green's function, incident fluence rate and wavenumber perturbation were constant over each voxel, with the value of each function taken at the center of the voxel. Making this assumption leads to a summation approximation to the integral

$$\Phi^s(\mathbf{r}) \approx \sum_{j \in V} G(\mathbf{r}, \mathbf{r}'_j) \Phi^i(\mathbf{r}'_j) \Delta k^2(\mathbf{r}'_j), \quad (2.13)$$

where the summation is over all voxels in  $V$  and  $\mathbf{r}'_j$  is the center of the  $j^{\text{th}}$  voxel. To construct the discrete inverse problem from multiple measurements we stacked the summations of Eq. (2.13) from each measurement. This gave us a matrix vector equation

$$\Phi^s = \mathbf{G}\mathbf{k}, \quad (2.14)$$

where  $\Phi_m^s$  is the scattered fluence rate from  $m^{\text{th}}$  measurement,  $\mathbf{k}_n$  is the  $\Delta k^2(\mathbf{r}_n)$  for the  $n^{\text{th}}$  voxel, and  $\mathbf{G}_{m,n}$  is given by

$$\mathbf{G}_{m,n} = G(\mathbf{r}_m, \mathbf{r}_n) \Phi^i(\mathbf{r}_n) \quad (2.15)$$

where  $\mathbf{r}_m$  is the position of the  $m^{\text{th}}$  measurement and  $\mathbf{r}_n$  is the position of the  $n^{\text{th}}$  voxel.

When sources are amplitude modulated the frequency domain measurements are represented as complex values from a quadrature demodulation at the modulation frequency. In general when computing the inverse solution this could result in a complex valued  $\mathbf{k}$  which would imply a complex valued optical absorption coefficient. To restrict the elements of  $\mathbf{k}$  to the real line we separate the real and imaginary parts of  $\Phi^s$  and  $\mathbf{G}$  and stack them. In matrix-vector notation this is given by

$$\tilde{\Phi}^s = \begin{bmatrix} \text{Re}(\Phi^s) \\ \text{Im}(\Phi^s) \end{bmatrix} = \tilde{\mathbf{G}} \mathbf{k} = \begin{bmatrix} \text{Re}(\mathbf{G}) \\ \text{Im}(\mathbf{G}) \end{bmatrix} \mathbf{k} \quad (2.16)$$

## 2.2.2 Validation of the propagation model

Because the linear propagation model presented in the previous section is an approximation of the frequency domain diffusion equation, we needed to characterize how well the approximation agreed with the full solution. We validated the accuracy of the first Born propagation model by comparing the scattered fluence rate solution from our first Born model and two other forward solutions, the spherical harmonic solution presented in Boas et al. [19] and discussed in 1.3.3, and a finite difference frequency domain (FDFD) solution. We examined the scattered fluence rate under several different scenarios to understand when the first Born model was appropriate and how it deviated from the true solution when it was not accurate.

### Simulation scenarios

All of the scenarios consisted of simulating a spherical absorption anomaly in a diffuse medium. The diffuse medium was modeled with a reduced scattering coefficient of  $10 \text{ cm}^{-1}$ , an absorption coefficient of  $0.041 \text{ cm}^{-1}$  and a index of refraction of 1.37. The absorption anomaly had an absorption coefficient of  $0.18 \text{ cm}^{-1}$  and was centered 2.5 cm into the medium with a radius of 1 cm. Two medium geometries were examined, a free space geometry and a half space geometry. The half space geometry was implemented with the extrapolated boundary condition, where the extrapolated boundary was 0.1832 cm away from the air-tissue interface. A schematic diagram for this scenario is shown in Figure 2.5. A single source was simulated in one of three positions; directly above the anomaly, or -3 cm or 2 cm away from directly above the anomaly. The effective position of the diffuse source was  $1/\mu'_s = 0.1 \text{ cm}$  into the medium as discussed in Section 1.3.4. In the half space medium case the offset direction was normal to the air-tissue interface, simulating fibers that were normal to the interface. In the free space medium simulation, since there was no air-tissue interface, the direction was chosen so that it matched the geometry of the half space medium simulation. Both CW and 200 MHz source modulation were simulated. The latter produced a diffuse photon density wave with a wavelength of 9.441 cm. The scattered fluence rate field was sampled along a line at the air-tissue interface, that coincided with the offset direction of the sources, in 0.5 cm increments.

### First Born details

We implemented the first order Born approximation using a 0.1 cm grid just surrounding the 1 cm diameter sphere. This resulted in a 21x21x21 computational grid. The half space boundary was

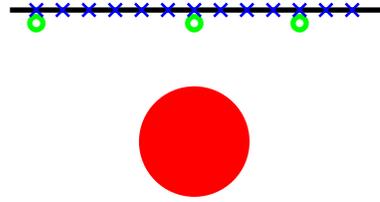


Figure 2.5: Diagram of the half space validation scenario. The air-tissue interface is represented by the black line, the detector position are shown by the blue crosses, the effective source positions are shown by the open green circles and the absorption anomaly position and size is represented by the large red circle.

implemented by first computing the incident fluence rate at the inhomogeneity as the difference of the effective source at 0.1 cm into the medium and the image source at 0.4664 cm outside the medium using the free space Green's function for the diffusion equation. We then calculated the scattered fluence rate at the detectors as the difference of the effective scattering source and its reflection across the boundary. The effective scattering source is due to the incident fluence rate at the inhomogeneity voxel and the difference in absorption from the background.

### **Spherical harmonic details**

The spherical harmonic solution is based on a spherical series expansion solution to the three-dimensional Helmholtz equation in an infinite medium. For this approach we used the code embedded in the PMI program developed by David Boas [56]. The number of terms in the series used by this code is unknown to the user. For the infinite medium case the solution was calculated by using a straight forward PMI script. The only parameter to mention is that the source was specified 0.1 cm into the medium because PMI does not change the position of the diffuse point source relative to the specified source position for an infinite medium.

For the half space medium we implemented a method of images by using PMI to calculate the spherical harmonic solution in an infinite medium for all four combinations of the true source and its image and the true object and its image. This resulted in sources at 0.1 cm and -0.4664 cm relative to the detectors and the sphere center 2.5 cm and -2.8664 cm relative to the detectors.

### **FDFD details**

The FDFD solution [55, 57] solved the scattered fluence rate equation Eq. (2.6). Since this technique does not rely on the simplifying assumptions we have made to derive the first Born approximations it provides a good test of the accuracy of the first Born approximation. We chose to apply the finite difference method to Eq. (2.6), the partial differential equation describing the scattered fluence rate, as opposed to Eq. (2.3), which describes the total fluence rate, in order to more accurately compute the scattered fluence rate in a reasonable amount of time. Applying the FDFD approach to solve the total diffusion equation requires a large amount of computation time to accurately characterize the scattered fluence rate. This is because the scattered fluence rate can often be many orders of magnitude less than the total fluence rate. By calculating the incident fluence rate through the Green's function solution Eq. (2.9) and then plugging the result into the FDFD solution of Eq. (2.3), we were able to accurately calculate the scattered fluence rate in much less time.

The finite difference method was implemented using a large computational grid of 16 x 16 x 16 cm with a grid step size of 0.25 cm. This resulted in 65 x 65 x 65 computation grid. For the free space simulation the grid was centered on the anomaly. For the half space case the grid was placed such that edge of the computational grid coincided with the extrapolated boundary. This implicitly implemented the zero boundary at the correct location. The large computational volume

was used to approximate free space (an unbounded diffuse domain) for all of the other faces of the computational volume. The specific technique used to solve the FDFD problem is known as the Jacobi method [57]. Typically this technique is considered to converge too slowly for practical purposes, but within MATLAB it lends itself to a much larger number of vector operations relative to element operations, compared to the Gauss-Seidel or Successive Overrelaxation (SOR) techniques. This resulted in a Jacobi implementation that, although requiring a larger number of iterations, was faster than Gauss-Seidel or SOR.

### **Validation results**

The results from simulations using the three forward solutions are shown in Figures 2.6 through 2.9. These results were from the 200 MHz source modulation half space scenario for the source directly above the absorption anomaly and for the source at the -3 cm position (left source in Figure 2.5). All of the other simulations showed similar results in comparing the three forward solutions. Figure 2.6 shows the amplitude and phase (both relative to the source) of the scattered fluence rate simulated at the detector positions for the source position directly above the absorption anomaly. From the plots we can see that all three forward solutions produce the same spatial structure in both the amplitude and phase plots. The first Born approximation produces a slightly larger scattered fluence rate. This is expected because the Born approximation does not account for any attenuation in the fluence rate interacting with the lower part of the anomaly after it has traveled through the upper part of the anomaly. Figure 2.6 shows the detected fluence rate for the case where the source was offset -3 cm from absorption anomaly (the left most source position in Figure 2.5). Again we see close agreement in the spatial structure of both the amplitude and phase responses between the

three simulation methods.

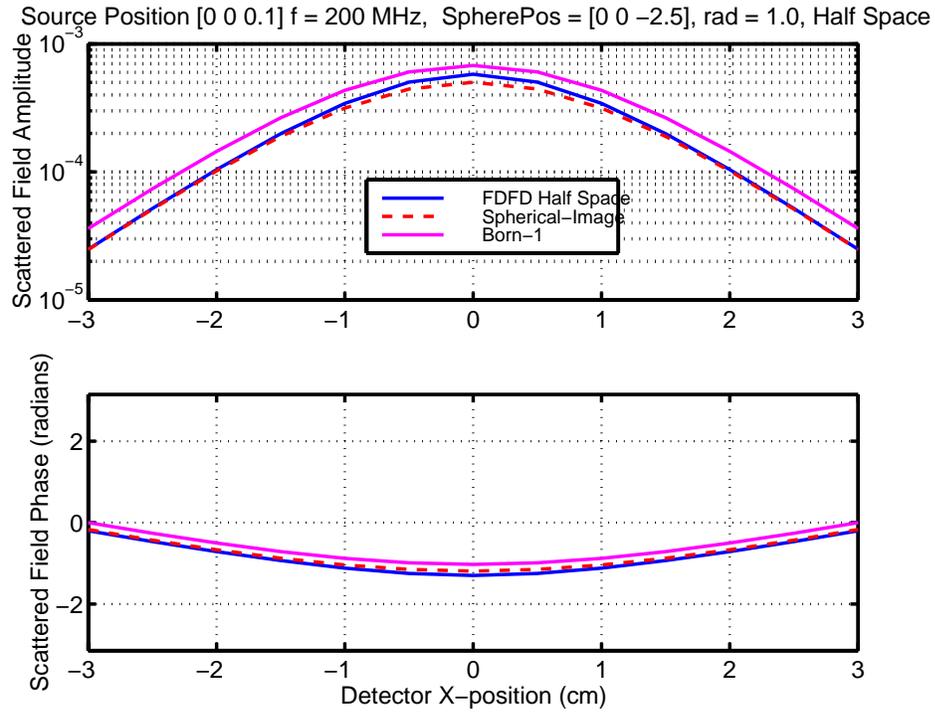


Figure 2.6: The scattered fluence rate amplitude (top) and phase (bottom) relative to the source for the center source position. The simulation scenario consisted of 200 MHz source modulation and a half space medium. The three curves show the scattered fluence rate simulated from three forward solutions, spherical harmonics, FDFD, and first Born.

To get a better understanding of the differences between the first Born simulation and the other methods we have plotted the ratio of the amplitudes relative to the spherical harmonic solution in Figures 2.8 and 2.9. These plots provide a more convenient view in which to examine the overestimation of the first Born approximation. For all three source positions we observed that the first Born approximation overestimates the scattered fluence rate by 50-70%, as is evident in Figures 2.8 and 2.9. Thus when using the first Born approximation as a linear model for reconstructing the optical absorption coefficient we can expect an underestimation of the coefficient. But, from these figures it

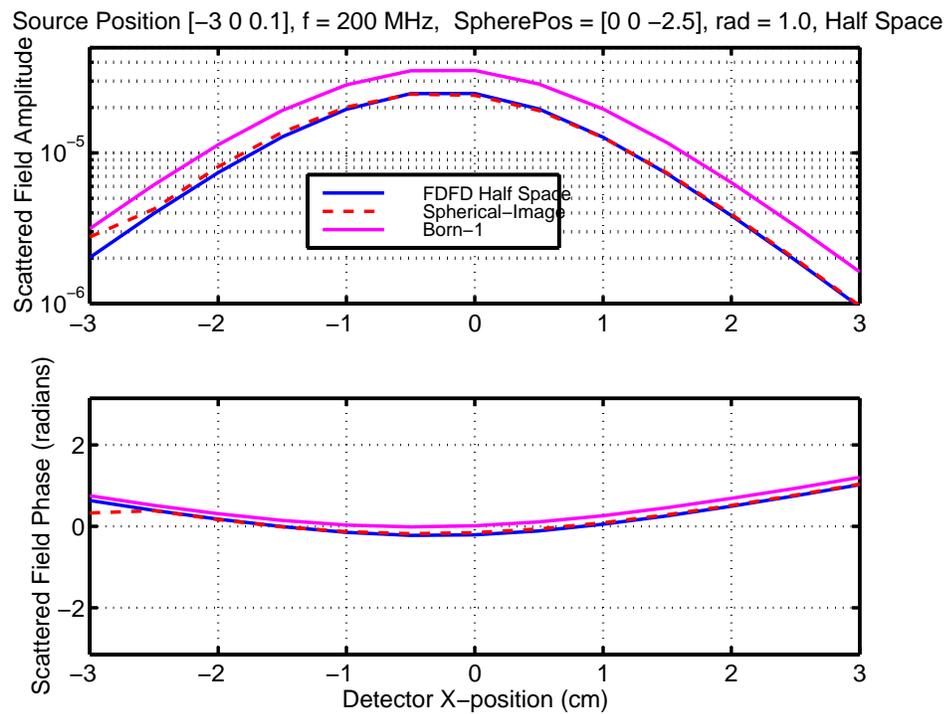


Figure 2.7: The scattered fluence rate amplitude (top) and phase (bottom) relative to the source for the left source position. The simulation scenario consisted of 200 MHz source modulation and a half space medium. The three curves show the scattered fluence rate simulated from three forward solutions, spherical harmonics, FDFD, and first Born.

appears that first Born approximation does capture the spatial variation of the scattered fluence rate well, and provides a good system model for linear model reconstruction, at least in this scenario.

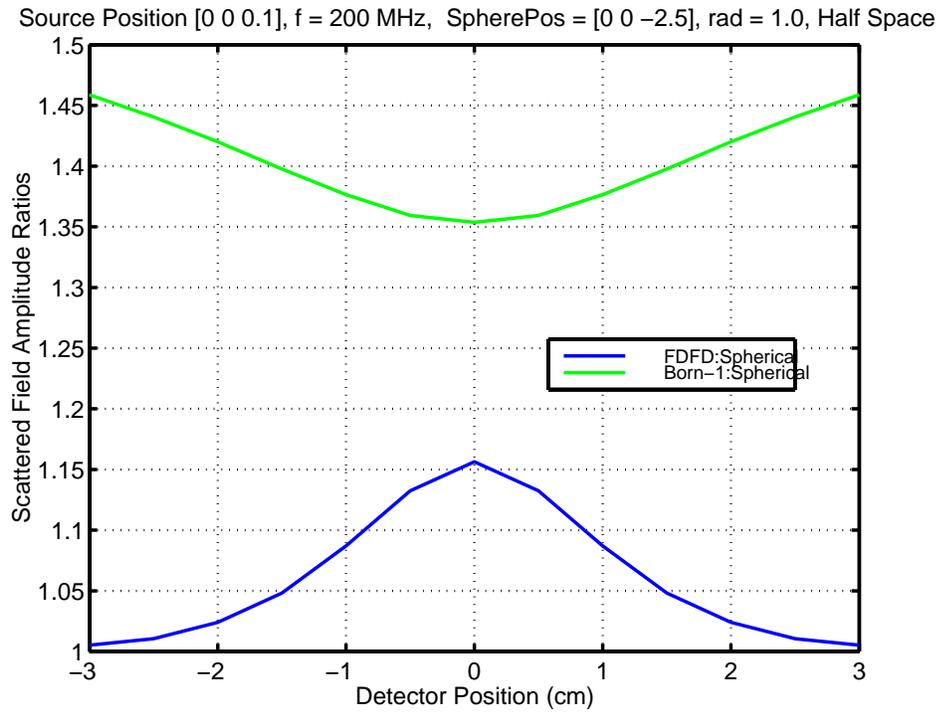


Figure 2.8: Ratio of the amplitude scattered fluence rate relative to the spherical harmonic solution for FDFD and first Born solutions. The simulation scenario is the same as described Figure 2.6.

## 2.3 Noise models

Measured optical DOT data will undoubtedly be corrupted by additive noise. At this point though, it is not yet certain what the characteristics of the noise will be since noise in DOT measurements can come from two general sources. The first is due to thermal noise in the detector electronics. This type of noise is independent of the signal being measured and can be well modeled as a sequence

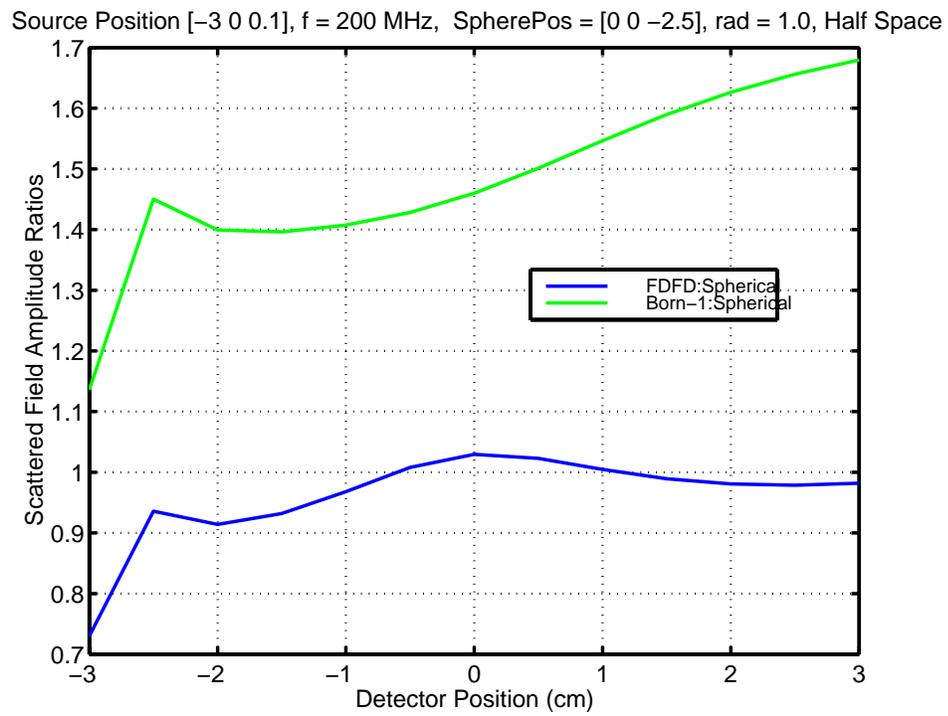


Figure 2.9: Ratio of the amplitude scattered fluence rate relative to the spherical harmonic solution for FDFD and first Born solutions. The simulation scenario is the same as described Figure 2.7.

of independent identically distributed (IID) Gaussian random variables. The second, more complex type of noise, known as shot noise, is due to the quantum fluctuations of the electromagnetic field. Shot noise is typically modeled as Poisson distributed in the rate of photon arrival. We approximate this random process with independent, but not identically distributed, Gaussian random variables with an appropriate mean and standard deviation described below. We base this approximation on the expectation that photon rate will need to be reasonably high to exceed the thermal noise floor of the detectors. In measured DOT signals both of these types of noise are generally present. Whether one or both are significant is dependent upon the detector, the electronics, and the source-detector distance.

### 2.3.1 Detector noise model

To examine the effect of noise on the reconstruction algorithms we will present, we used two different noise models. The simplest noise model we used followed the thermal noise assumption and thus was uniform variance Gaussian noise added to the scattered fluence rate signal. This model allowed us to examine the most basic case without having the noise variance be a function of the measured fluence rate. The signal-to-noise-ratio (SNR) for this model was defined relative to the peak value of the scattered fluence rate and simulated the same noise variance in each channel (source-detector pair measurement). The SNR, in dB, was defined as

$$SNR = 20 \log_{10} \left( \frac{\max(|\Phi^s|)}{\sigma} \right) \quad (2.17)$$

where  $\sigma$  is the standard deviation of the noise in each channel and  $\Phi^s$  is the scattered fluence rate vector from Eq. (2.14). This simulated a detector dominated noise where the dynamic range of all

detectors was constant, set such that they could all measure the maximum signal without saturating.

### 2.3.2 Shot noise model

On the other hand, following the shot noise model, since the variance of the noise in each measurement is a function of the measured fluence rate, we also examined a simple approximation to that case. In this model the variance of the noise for each measurement is proportional to the total fluence rate detected. Mathematically we specified the standard deviation of the noise in channel  $j$ , specified by  $\sigma_j$ , as

$$\sigma_j = |\Phi_j^s + \Phi_j^i| 10^{\frac{-SNR}{20}} \quad (2.18)$$

where  $\Phi^i$  is the incident fluence rate vector.

### 2.3.3 Combined noise model

After many of the simulations in this thesis were completed we became aware of a more accurate model derived from first principles, which we present below. This model takes into account both thermal noise in detector electronics and quantum uncertainty associated with the optical systems. This derivation is taken with permission from a memo written by Prof. Charles DiMarzio.

Consider the complex power of the light emitted by a sinusoidally modulated source represented by the equation

$$P_S(t) = P_{DC} + \beta P_{DC} \exp(j\omega t), \quad (2.19)$$

where  $\beta$  is the modulation depth,  $\beta \in [0,1]$ , and the real part of the equation describes the source

power as a function of time. The power present at a detector is described by

$$P_D(t) = \phi_{DC} \frac{A_D \Omega_D}{4\pi} P_{DC} + \phi_{AC} \frac{A_D \Omega_D}{4\pi} \beta P_{DC} \exp(j\omega t), \quad (2.20)$$

where the real number  $\phi_{DC}$  accounts for the propagation through the diffuse medium at DC, the complex number  $\phi_{AC}$  accounts for the AC diffuse propagation,  $A_D$  is the detector fiber area, and  $\Omega_D$  is the detector acceptance angle.

The current produced in the detector is proportional to the sum of the light power falling on the detector and two noise components, one due to thermal noise in the circuitry of the detector and the other to quantum fluctuations in the expected light power. Thus the total current in the detector can be written

$$I(t) = \frac{\eta e}{\hbar \nu} P_D(t) + I_Q(t) + I_T(t), \quad (2.21)$$

where  $\eta$  is the quantum efficiency of the detector,  $e$  is the charge of an electron,  $\hbar$  is Planck's constant,  $\nu = c/\lambda$  is the photon frequency,  $I_Q(t)$  is the random process representing the quantum fluctuations, and  $I_T(t)$  is the random process representing the thermal noise.  $I_Q$  and  $I_T$  are assumed to be mutually independent. The thermal noise can be modeled as a uniform Gaussian noise process with a standard deviation given by

$$\sigma_T = \sqrt{kTBF_n}, \quad (2.22)$$

where  $k$  is Boltzman's constant,  $T$  is the temperature,  $B$  is the bandwidth of the detector, and  $F_n$  is the noise figure for the detector. Another representation used in optics is the concept of noise-equivalent power (NEP). Using this representation the standard deviation of the thermal noise is given by

$$\sigma_T = \frac{\eta e}{h\nu} (NEP) \sqrt{B}. \quad (2.23)$$

The quantum fluctuations are due to the statistical nature of photon emission. It is assumed that the arrival of the  $n^{\text{th}}$  photon is independent of the previous  $n - 1$  photons. Thus the number of photons arriving in a period  $1/B$  is Poisson distributed. If  $n$  photons are expected to arrive in a period  $1/B$  the standard deviation of the Poisson distribution is  $\sqrt{n}$ . With the expected number of photons varying sinusoidally the actual distribution of the quantum fluctuations will be non-stationary. To simplify dealing with the quantum noise we consider only the DC portion of the fluence rate, resulting in a stationary model for the quantum noise. This is a reasonable model when the detected DC power is much greater than the detected AC power, that is when the modulation depth  $\beta$  is small and/or the ratio of the DC propagation to the AC propagation  $\phi_{DC}/|\phi_{AC}|$  is large. The expected number of photons for the DC fluence rate is given by

$$\bar{n} = \frac{\eta}{h\nu B} \phi_{DC} \frac{A_D \Omega_D}{4\pi} P_{DC}. \quad (2.24)$$

Therefore the standard deviation of the quantum fluctuations is

$$\sigma_Q = \sqrt{\bar{n}} = \sqrt{\frac{\eta}{h\nu B} \phi_{DC} \frac{A_D \Omega_D}{4\pi} P_{DC}}. \quad (2.25)$$

As with the thermal noise, another common formulation in the optics field is in terms of an equivalent optical power that would induce the observed quantum fluctuation. This quantity is called the quantum-noise power  $P_Q$  and given by

$$\sigma_Q = \frac{\eta}{h\nu} P_Q. \quad (2.26)$$

Solving for  $P_Q$  in terms of the detected DC fluence rate gives

$$P_Q = \sqrt{\frac{h\nu B}{\eta} \phi_{DC} \frac{A_D \Omega_D}{4\pi} P_{DC}}. \quad (2.27)$$

Thus with both sources of noise considered, the electrical power of the noise is proportional to the bandwidth. The electrical power of the quantum noise varies proportionally with square root of the detected fluence rate detected while the thermal noise is independent of the detected fluence rate. The other parameters specifying the noise power are fixed physical constants or parameters of the detection equipment.

## Chapter 3

# Linear model reconstruction evaluation

In this chapter we describe a set of experiments to evaluate the performance of two groups of linear model reconstruction techniques under two different reconstruction scenarios. The primary motivations for our exploration of linear model based reconstruction was that they are computationally relatively inexpensive compared to a full non-linear inversion and there is significant interest within the community in this approach. Our goal was to characterize commonly used solutions to ill-posed problems under the DOT imaging scenario and understand the applicability and limitations of the linear model based reconstructions. Additionally, even when a non-linear model is used, a linear model is often employed as part of an iterative inversion process. Therefore, an understanding of the properties and performance of various linear model inversion techniques can also be useful in developing non-linear schemes. The comparison we have performed included two classes of linear model inversion schemes. The more commonly employed algorithms in DOT imaging are known as algebraic techniques. From this class we have examined the Algebraic Reconstruction Technique (ART) and the Simultaneous Iterative Reconstruction Technique (SIRT). The other class of

techniques examined are known as subspace techniques because they operate on subspaces of the forward operator or linear model. From this family of algorithms we have examined the Truncated Singular Value Decomposition (TSVD) and the Truncated Conjugate Gradient (TCG) algorithm. These types of algorithms have been studied much more than the algebraic class with respect to regularization of ill-posed inverse problems [46]. Along with the selection of algorithm type we have also examined how applying regularization techniques affects the reconstruction quality and how to select the regularization parameter.

The reconstruction scenario that is of primary interest to us is full three-dimensional reconstruction where the absorption coefficient is mapped for a large volume under the sources and detectors. Another approach that is often employed is to reconstruct a two-dimensional slab under the source detector region. In this approach the slab has a prespecified depth and thickness and all inhomogeneities are assumed to lie in the slab. We have explored the performance of the linear model reconstruction algorithms under both of these scenarios. Throughout this chapter we label the former F3D (full 3D) and the latter 2DS (2D slice).

### **3.1 Simulation details**

As we described in the previous chapter, we used computer simulations to examine the performance of all of the reconstruction algorithms in this thesis. In the subsections that follow we described the specific details and parameters used with the general models described in Chapter 3.1.

### 3.1.1 Computational geometry

We modeled a semi-infinite slab of tissue with a single spherical absorption inhomogeneity and a reflective imaging geometry. A number of different positions and sizes of the absorption anomaly were examined with similar results. All the results presented here used a 1 cm radius anomaly centered 2.5 cm deep in the tissue, as shown in Figure 2.1.

### 3.1.2 Source and detector parameters

The reflective source-detector geometry we examined simulated a single planar array placed on the surface of the diffuse medium<sup>1</sup>. The detectors formed a four by four array with a 2 cm spacing in both dimensions. The sources were positioned in the interstitial spaces between the detectors, forming a three by three array also with 2 cm spacing. The sources were amplitude modulated at a frequency of 200 MHz. A schematic of the geometric layout of the sources and detectors for the simulations is shown in Figure 2.1. Shown in Figure 3.1 are two examples of a single slice through the volume of the true absorption perturbation. Figure 3.1(a) shows a vertical slice through the center of the anomaly in the X-Z plane. We use this representation for the full three-dimensional reconstructions because it can capture the depth performance of the algorithm in a single image. Figure 3.1(b) shows a horizontal slice through the center of the anomaly in the X-Y plane. We use this representation for two-dimensional slice reconstructions.

---

<sup>1</sup>This geometry models a measurement system currently in use in the Massachusetts General Hospital laboratory.

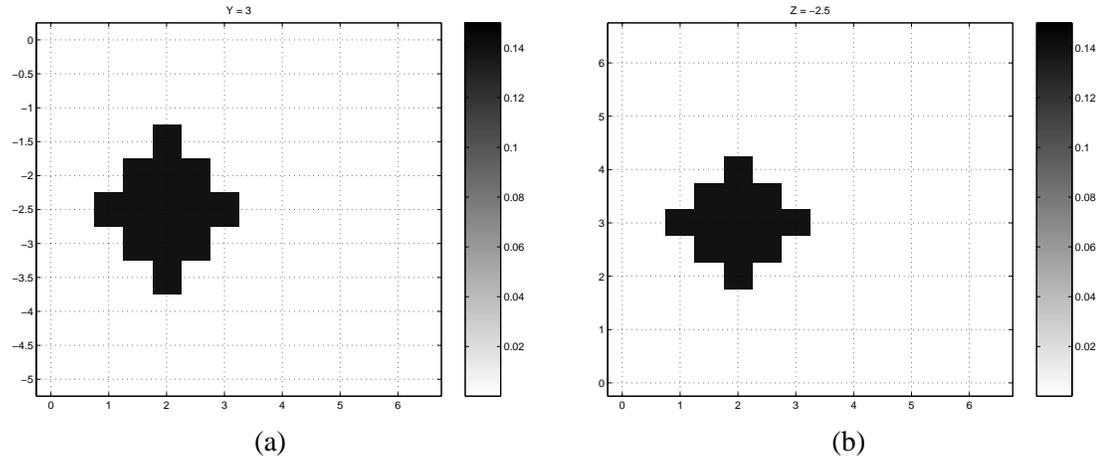


Figure 3.1: (a) A vertical plane of the true absorption perturbation function used in these simulations. This plane is through the center of the absorption anomaly ( $Y=3$  cm). (b) A horizontal plane of the absorption function showing the true absorption perturbation function. This plane is through the center of the absorption anomaly ( $Z=2.5$  cm).

### 3.1.3 Volume discretization

When solving the forward problem we calculated the simulated scattered fluence by discretizing a region just surrounding the absorption inhomogeneity with a fine cubic grid 1 mm on a side, requiring a  $21 \times 21 \times 21$  grid to cover the 1.0 cm radius anomaly. We did not need to discretize outside of this region because of the compact support of the absorption anomaly perturbation.

In contrast, when we solved the inverse problem we assumed that the location of the anomaly within the medium is unknown. Thus the entire volume or slab under the source detector array was discretized uniformly to reconstruct the optical parameter function over the whole region of interest. For the F3D reconstruction, the volume reconstructed, as shown in Figure 2.1, was a cube 7 cm x 7 cm by 5.5 cm deep discretized into voxels. These voxels were 0.5 cm on a side in the  $X$  and  $Y$  dimensions and either 0.5 cm or 2.0 cm deep in the vertical ( $Z$ ) direction. This resulted in 2156 voxels for the 0.5 cm deep voxels or 588 voxels for the 2.0 cm deep voxels.

With the 2DS reconstruction our assumptions were somewhat different. In this reconstruction scenario it is assumed that we have *a priori* knowledge of the width and depth of the absorption anomaly to properly locate the reconstruction slab. The 2DS reconstructions were also implemented using a 0.5 cm grid in the X and Y dimensions. The thickness of the slab was either 0.5 cm or 2.0 cm with the center of the slab aligned with the center of the absorption anomaly. Both of these reconstruction geometries contain 196 voxels. Thus the size of the relevant linear system and the required computations to reconstruct the absorption coefficient map are greatly reduced for the 2DS reconstruction.

In all our reconstructions a finer discretization was used in forward calculations and a coarser one used for the inverse reconstruction. This induces a certain amount of model mismatch between the forward and inverse problems. Additionally, in some of the 2DS reconstructions the thickness of the slice is less than the thickness of the inhomogeneity and thus there is further model mismatch in the inverse reconstructions.

### 3.1.4 Medium optical parameters

The diffuse medium optical parameters for our experiments were selected to simulate human tissue. Specifically, the background scattering coefficient,  $\mu_s$ , was taken as  $100 \text{ cm}^{-1}$ , and the mean cosine of the scattering angle as 0.9, resulting in a reduced scattering coefficient,  $\mu'_s$ , of  $10 \text{ cm}^{-1}$ . The background absorption coefficient was derived from an assumed hemoglobin concentration of  $54.9 \text{ }\mu\text{M}$  and an oxygen saturation of 90%, resulting in an absorption coefficient of  $0.041 \text{ cm}^{-1}$  at a wavelength of 780 nm. The absorption anomaly simulated a hemoglobin concentration of  $0.210 \text{ }\mu\text{M}$  and an oxygen saturation of 60%, resulting in an absorption coefficient of  $0.18 \text{ cm}^{-1}$  or  $0.139$

$\text{cm}^{-1}$  above the background. These concentrations are based on the data in Prah1’s compilation [24].

### 3.1.5 Noise model

We modeled the noise present at each detector as an independent multivariate Gaussian random variable with separate independent real and imaginary components, zero mean, and a standard deviation specified relative to the total fluence at the detector. This is the simplified shot noise model described in Section 2.3.2. It is clear from Eq. (2.18) that the noise does not have a constant variance across measurement pairs, and therefore the noise must be “whitened” to fit the assumptions of the subspace techniques. Whitening was accomplished by weighting the system by the inverse covariance matrix, expressed mathematically as

$$\mathbf{R}^{-1} \tilde{\mathbf{G}} \mathbf{k} = \mathbf{R}^{-1} \tilde{\Phi}^s \quad (3.1)$$

where  $\mathbf{R}^{-1}$  is the inverse of the covariance matrix for the measurements. Because the noise in each channel is assumed to be independent this inverse is easily computed from the variance in each channel

$$\mathbf{R}^{-1} = \text{diag}(\sigma_j^{-1}). \quad (3.2)$$

where  $\sigma_j$  is the noise variance in the  $j^{\text{th}}$  channel.

### 3.1.6 Model mismatch analysis

Even with the different voxel sizes for the forward and inverse models described in Section 3.1.3 there is not a significant mismatch between the two models. From the comparison of forward

models in Section 2.2.2 we know that the first Born forward model overestimates the amplitude of the scattered field. To give an indication of how our results might change when the forward data was generated by a completely different, and more physically accurate, numerical technique, we substituted an FDFD forward model for the first Born forward model in an additional set of experiments. We used a standard finite difference frequency domain (FDFD) forward solution [55] for the same scenario described above to generate the scattered fluence data as described in Section 2.2.2. This data then simply replaced the Born-generated data as input to the noise model and the inversion algorithms. The FDFD simulation used a 0.25 cm grid size with a computational volume of 20 cm on a side for the X and Y domains and 12.5 cm for the Z domain. The X and Y domains were centered on the inhomogeneity. The Z domain was aligned so that the computational boundary was aligned with the extrapolated boundary position as previously described. All of the computational boundaries implemented a zero Dirichlet boundary condition.

### 3.2 Reconstruction algorithms

There are two important characteristics of this inverse problem that reconstruction algorithms must be taken into account. First, the F3D linear system model Eq. (2.16) is typically very underdetermined. Examples we present in this study have 16 detectors and nine sources, for a total of 144 source-detectors pairs. Each source detector pair provides an in-phase and quadrature measurement at 200 MHz, providing a total of 288 measurements. The number of unknowns (the number of voxels) is as large as 2156, and thus we have up to seven times as many unknowns as measurements. Second, this inverse problem in its continuous form is ill-posed. The absorption coefficients do not

vary smoothly with the data, resulting in an ill-conditioned forward operator  $\tilde{\mathbf{G}}$ . This is typical of integral operators with smoothing kernels [46]. Therefore the solution must be regularized if it is not to be dominated by reconstruction noise.

### 3.2.1 Algebraic techniques

As we mentioned previously, the types of algorithms we examined for solving the linear system Eq. (2.16) fall into two classes, algebraic techniques and subspace techniques. Members of the first class solve Eq. (2.16) by projecting an estimate of the solution onto the hyperplanes represented by rows of the linear system. Included in this class are ART and SIRT [25]. ART sequentially projects a solution estimate onto hyperplanes defined by the individual rows of the linear system. This projection becomes the estimate of the solution for the next iteration. This can be expressed mathematically as

$$\hat{\mathbf{k}}_{j+1} = \hat{\mathbf{k}}_j + w \frac{\tilde{\Phi}_i^s - \tilde{\mathbf{G}}_i \hat{\mathbf{k}}_j}{\tilde{\mathbf{G}}_i \tilde{\mathbf{G}}_i^T} \tilde{\mathbf{G}}_i^T,$$

$$j = 0, 1, \dots \quad i = (j \bmod 2m) + 1$$

where  $\hat{\mathbf{k}}_j$  is the  $j^{\text{th}}$  estimate of the absorption anomaly perturbation,  $\tilde{\mathbf{G}}_i$  is the  $i^{\text{th}}$  row of the  $2m \times n$  matrix  $\tilde{\mathbf{G}}$ , and  $\tilde{\Phi}_i^s$  is the  $i^{\text{th}}$  measurement.  $w$  is a relaxation parameter that adjusts the step size of each iteration. SIRT is implemented in a similar manner except that instead of projecting the estimate onto each row in sequence, the component vector that would be projected out is averaged over all rows and then subtracted from the original estimate. Mathematically this is written as

$$\hat{\mathbf{k}}_{j+1} = \hat{\mathbf{k}}_j + w \frac{1}{2m} \sum_{i=1}^{2m} \frac{\tilde{\Phi}_i^s - \tilde{\mathbf{G}}_i \hat{\mathbf{k}}_j}{\tilde{\mathbf{G}}_i \tilde{\mathbf{G}}_i^T} \tilde{\mathbf{G}}_i^T.$$

SIRT has generally been observed to generate smoother reconstructions of the unknown function due to the averaging over a number of projections, at a cost of slower convergence. For underdetermined systems, such as the one we are considering, algebraic techniques will converge to a point on the hyperplane satisfying the linear system that is nearest in 2-norm to the initial guess [25].

If run until convergence, the algebraic techniques produce unreliable oscillations due to the ill-posedness of the problem. To avoid this, regularization is accomplished in both algebraic techniques by limiting the number of iterations. The choice of how many iterations to perform is a difficult topic for these methods; here we avoided the problem by choosing the best regularization (optimal truncation) using prior knowledge of the true solution to compute a performance measure of interest and thus select the optimal truncation parameter.

The problem of selecting an optimal relaxation parameter is still a topic of investigation [58] and outside the scope of this thesis. In order to study this parameter choice in the context of our problem, we examined numerical simulation results over a wide range of SNRs, relaxation parameters, and error measures, as described below. For SIRT, we observed that relaxation parameters over the entire available range did not appreciably alter the best performance metrics achieved. Rather, the rate (number of iterations) at which the best performance measure is achieved varied, so that there is an interaction between relaxation parameter and optimal number of iterations. For ART, at low SNR we observed a small improvement in the best performance when we used a small relaxation parameter ( $w = 0.25$ ), as compared to the best performance achieved with  $w = 1$ . At higher SNR a relaxation parameter closer to one provided the best performance. Again we observed an interaction between relaxation parameter and optimum number of iterations. In no case did we observe any behavior of the algebraic algorithms over this range of relaxation parameters that would change the

basic conclusions we draw below from our comparisons. Thus, since we were not able to find a reliable automatic method to choose the optimal value, and since we are already manually choosing the true optimal truncation parameter for these methods, we simply used a relaxation parameter of one for all of the comparisons reported below.

### 3.2.2 Subspace techniques

The first subspace technique that we examined was the TSVD algorithm. This algorithm is derived from the Singular Value Decomposition (SVD) of the  $2m \times n$  stacked system matrix  $\tilde{\mathbf{G}}$ . The SVD of the system matrix is given by

$$\tilde{\mathbf{G}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad \mathbf{U} \in \mathbb{R}^{2m \times 2m}, \quad \mathbf{\Sigma} \in \mathbb{R}^{2m \times n}, \quad \mathbf{V} \in \mathbb{R}^{n \times n}$$

where  $\mathbf{U}$  and  $\mathbf{V}$  are orthonormal matrices and  $\mathbf{\Sigma}$  is a diagonal matrix with values  $\Sigma_{i,i} = \sigma_i \geq 0$ . The  $\sigma_i$  are known as the singular values of  $\tilde{\mathbf{G}}$  and the decomposition is written such that

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0 \quad \sigma_{r+1}, \sigma_{r+2}, \dots, \sigma_{\min(m,n)} = 0,$$

where  $r$  is the rank of  $\tilde{\mathbf{G}}$ .

Poorly conditioned matrices such as the ones resulting from discretized ill-posed problems have a very wide range of singular values. This ill-posedness is evident in the singular value spectrum of a typical DOT forward matrix, shown in Figure 3.2, which displays a range of 7 orders of magnitude in the singular values. The TSVD algorithm computes the reconstruction by using only the largest  $t$  non-zero singular values and singular vectors to approximately solve  $\tilde{\mathbf{G}}\mathbf{k} = \tilde{\mathbf{\Phi}}^s$ . Mathematically this can be written as

$$\hat{\mathbf{k}} = \mathbf{V}_t \mathbf{\Sigma}_t^{-1} \mathbf{U}_t^T \tilde{\mathbf{\Phi}}^s \quad (3.3)$$

where  $\mathbf{V}_t$  and  $\mathbf{U}_t$  are the first  $t$  columns of  $\mathbf{V}$  and  $\mathbf{U}$  respectively, and  $\Sigma_t^{-1}$  is the inverse of the square diagonal sub-matrix of the largest  $t$  singular values. The value  $t$ , the truncation parameter, controls the amount of regularization in the the inverse. More information on the SVD and TSVD can be found in Golub and Van Loan [44] and Hansen [46]. In this context we note that in Figure 3.2 there is no useful gap in the spectrum (*i.e.* gap in the singular values) that might indicate an obvious choice of a good low-rank model. Truncating at the small jump in the plot of singular values around index 70 did not produce good reconstructions.

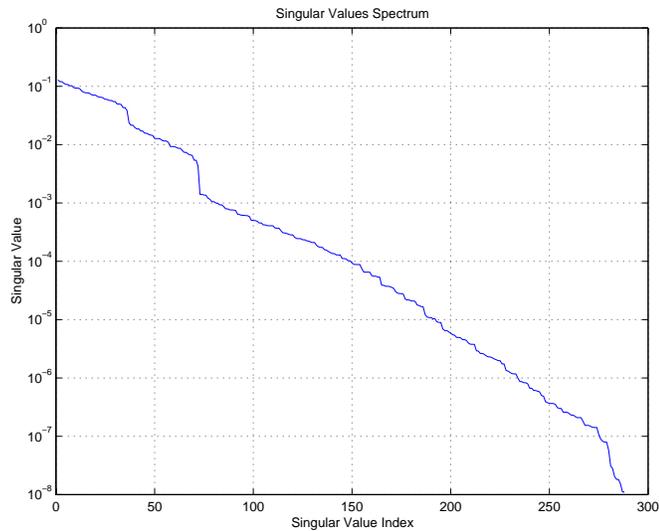


Figure 3.2: Singular value spectrum of a typical forward DOT forward matrix.

The TCG algorithm is derived from the Conjugate Gradient algorithm in a similar manner to the way TSVD follows from the SVD. The Conjugate Gradient algorithm is an iterative technique to solve a symmetric positive definite linear system of equations. As with the previous algorithms, for ill-conditioned systems we can regularize by stopping, or truncating, the iterations before we

reach full convergence [46]. Thus for TCG the number of iterations computed controls the amount of regularization. For most measurement geometries the system matrix  $\tilde{\mathbf{G}}$  is not symmetric positive definite, a requirement for the Conjugate Gradient algorithm to be guaranteed to find a solution. However we can apply the Conjugate Gradient algorithm to the normal equations given by

$$\tilde{\mathbf{G}}^T \tilde{\mathbf{G}} \mathbf{x} = \tilde{\mathbf{G}}^T \tilde{\mathbf{\Phi}}^s.$$

For computational reasons we do this without explicitly computing the normal equations, as given in Table 3.1 (derived from Shewchuk’s [45] implementation).

### 3.2.3 Regularization parameter selection

With any regularization technique, one of the primary issues is the selection of the parameter that controls the trade-off between fidelity to the data and some constraint on the result. There are a large variety of methods available, divided between *a priori* methods which use prior knowledge about the solution, the noise, or both, and *a posteriori* methods which use only the measurements and forward model. For the subspace techniques we use a well-known *a posteriori* method, the L-curve technique [46], which for subspace methods graphs the log of the 2-norm of the residual versus the log of the 2-norm of the estimate while varying the regularization parameter. Thus the regularization parameter itself is represented only parametrically in this graph. An example of an L-curve generated from a TCG reconstruction at a signal-to-noise ratio of 20 dB is shown in Figure 3.3. We generated this graph by plotting the residual norm and reconstruction norm over 300 iterations of the algorithm. The “corner” of the resulting curve is taken as a good choice of regularization parameter because it identifies a point at which there is a balance between increase

$\mathbf{b}_p = \tilde{\mathbf{A}}^T \mathbf{b}; \quad \mathbf{d} = \mathbf{r} = \mathbf{b}_p; \quad \delta = \mathbf{r}^T \cdot \mathbf{r}; \quad \delta_o = \delta; \quad \hat{\mathbf{x}} = \mathbf{0};$ $\text{for } j = 1 : n_{iter}$ $\quad \mathbf{q} = \mathbf{A}^T (\tilde{\mathbf{A}} \mathbf{d})$ $\quad \alpha = \delta / (\mathbf{d}^T \cdot \mathbf{q})$ $\quad \hat{\mathbf{x}} = \hat{\mathbf{x}} + \alpha \mathbf{d}$ $\quad \mathbf{r} = \mathbf{r} - \alpha \mathbf{q}$ $\quad \delta_o = \delta$ $\quad \delta = \mathbf{r}^T \cdot \mathbf{r}$ $\quad \mathbf{d} = \mathbf{r} + \frac{\delta}{\delta_o} \mathbf{d}$ $\text{end}$
---

Table 3.1: Conjugate Gradient algorithm for the normal equations. Note that the normal equations are not computed explicitly. Rather the necessary matrix vector product is first computed and then the transposed matrix is right multiplied by the resulting vector for each instance where the normal equations are needed.

in the residual norm and increase in the solution norm. The diamond drawn on the graph shows the point we manually selected as the L-curve corner, which corresponded to 12 iterations.

For the algebraic techniques *a posteriori* methods such as the L-curve do not work well because standard measures of error such as the residual error or solution norm do not change monotonically as we iterate, so we simply present the best possible result for these methods.

### 3.3 Results and observations

In this section we present results and observations of applying the four reconstruction techniques described in Section 3.2 to the simulation scenario described in Section 3.1. Example reconstructions are shown first to present a qualitative idea of the reconstruction performance of each of the methods. Following this, we report quantitative performance measures for a range of signal-to-noise ratios. Finally, we illustrate the sensitivity of the reconstruction performance to the selection of the

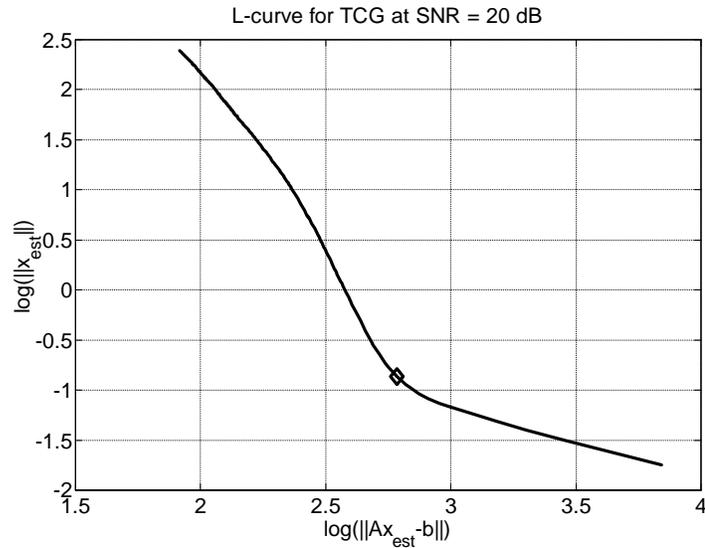


Figure 3.3: L-curve for a TCG reconstruction at a signal-to-noise ratio of 20dB. The diamond at the corner of the graph identifies the point that was selected as the corner of the L-curve.

regularization parameter.

### 3.3.1 Example reconstructions

Figure 3.4 shows a set of images of the reconstructed absorption coefficient in a single vertical (X-Z) plane through the center of the absorption anomaly ( $Y = 3$  cm) at a signal-to-noise ratio of 20 dB. The true absorption anomaly function is shown in Figure 3.1(a). The F3D model was used for the inversion with a 0.5 cm grid in all dimensions. Comparing the true absorption perturbation image with the reconstructions, it is evident that all of the reconstructions show a roughly correct image of the absorption anomaly and that all of the reconstruction techniques underestimate the depth and amplitude of the absorption perturbation function. Note that the scale for the reconstructed images is less than that for the true absorption function. Figure 3.4(a) is the reconstruction using a single

iteration of ART (*i.e.* one iteration cycle through all 288 rows of the system matrix). An absorption anomaly centered approximately at  $X = 1.5$  cm and  $Z = -1.75$  cm with an amplitude of around  $0.04 \text{ cm}^{-1}$  is visible. The reconstructed absorption function does not appear to have the same area as the true absorption function nor does the reconstructed absorption anomaly extend as deep. Figure 3.4(b) shows a reconstruction using the SIRT algorithm with 26 iterations. Qualitatively this reconstruction appears very similar to the ART reconstruction.

Figure 3.4(c) shows a reconstruction using the TSVD algorithm employing 56 singular values. This reconstruction clearly produces a larger amplitude absorption coefficient compared to the algebraic techniques. Additionally, the absorption anomaly center appears at about 2 cm depth, about a quarter centimeter deeper than the algebraic techniques. Thus in both aspects, amplitude and position, it is closer to the true absorption function. Figure 3.4(d) displays the reconstruction for the TCG technique using twelve iterations, which is qualitatively very similar the TSVD reconstruction. In comparing the algebraic and subspace methods we reiterate that for the subspace techniques the number of iterations or singular values was chosen by manually identifying the corner of the L-curve without using knowledge of the true solution while for the algebraic techniques we chose the best solution **after** comparing to the true solution.

Figure 3.5 shows a pair of 2DS reconstructions using the TCG algorithm, again at a SNR of 20 dB. Figure 3.5(a) was generated using a 2.0 cm width reconstruction plane with the number of iterations again chosen by the L-curve method. Figure 3.5(b) was generated using a 0.5 cm width reconstruction plane which *underestimates* the true width of the absorption anomaly. For the 0.5 cm width reconstruction plane the L-curve did not provide useful information on selecting the regularization parameter; this is an indication that the large amount of error in the model will make

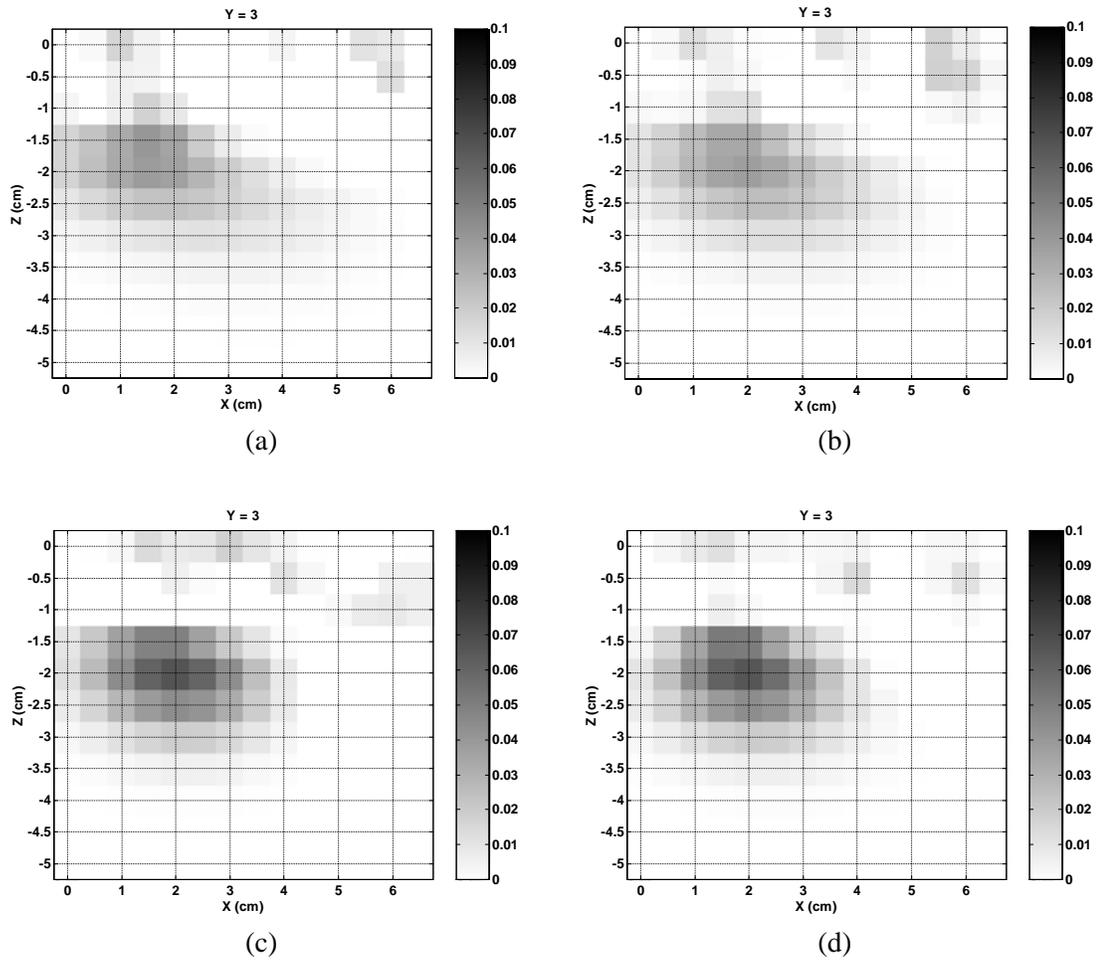


Figure 3.4: F3D reconstruction examples for all four reconstruction techniques at a 20 dB SNR. Each image shows a vertical plane through the center of the absorption anomaly ( $Y = 3$  cm). (a) is the ART algorithm result using one iteration. (b) is the SIRT algorithm result using 26 iterations. (c) is the TSVD algorithm result using 56 singular values. (d) is the TCG algorithm using 12 iterations.

regularization difficult. Therefore for this case we selected the number of iterations that minimized the actual mean square error. For both cases the center of the reconstruction was the true center of the absorption anomaly. We note that the 2DS reconstruction shown is reasonably accurate in the first case but significantly overestimates the absorption anomaly size in the second case.

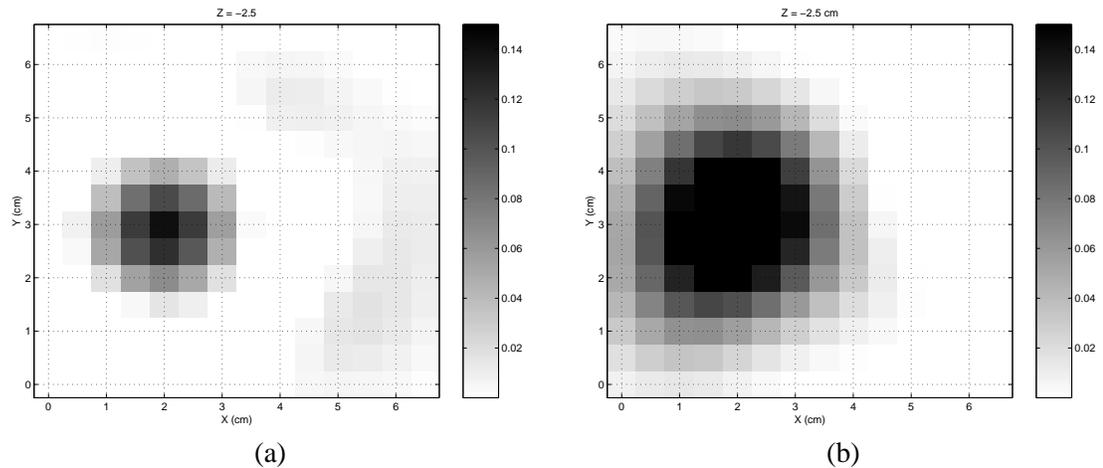


Figure 3.5: 2DS reconstruction with (a) 2.0 cm thick voxels and (b) 0.5 cm thick voxels at a SNR of 20 dB. Both reconstruction were generated using the TCG algorithm. The 2.0 cm thick voxels used 7 iterations. The 0.5 cm voxels used 2 iterations. Unlike the reconstruction shown in Figure 3.4 these reconstructions are shown in a X-Y plane.

### 3.3.2 Performance measures

As noted in the beginning of this chapter, a secondary objective of this work was to compare several error measures. We considered the standard mean square reconstruction error over the entire volume, but also devised measures designed to quantify the position and amplitude of the reconstructed absorption anomaly directly. More precisely, we used the following three error measures:

**Mean squared error: (MSE):** The first performance measure we evaluated was mean square

error over the volume, given by the expression

$$MSE = \text{mean}_{i \in V} ((\mu_a^{true}(i) - \mu_a^{est}(i))^2).$$

**Object centroid error: (OCE)** We evaluated the error in the estimated position of the reconstructed absorption anomaly relative to the true position of the absorption anomaly. Specifically, we set an “object threshold” at 50% of the maximum amplitude in the reconstruction. We considered the largest amplitude voxel as an initial “detected object”, and then iteratively built up a larger “detected object” by including in the object any neighbor of a current “object voxel” whose amplitude was above the threshold. The search finished when there were no more voxels which were neighbors of voxels classified as object voxels whose amplitude was above the threshold. Once the object was detected, we computed its centroid as the weighted average of the position of the object voxels, with the amplitude of these voxels as the weights. The OCE was then calculated by subtracting the object centroid of the true object from the object centroid of the reconstruction.

**Amplitude error: (AE):** The third performance measure we calculated measured the peak amplitude error of the reconstructed absorption coefficient over the known position of the absorption anomaly. We simply calculated the difference between the maximum value of the true absorption anomaly and the maximum value of the reconstructed absorption anomaly over the support of the true absorption anomaly. We note that in general for all the reconstruction methods we considered, the amplitude of all voxels generally increases with decreasing regularization, so for AE errors we always used the regularization parameter which minimized the MSE.

We used these three error measures, MSE, OCE, and AE, to quantitatively compare reconstruction techniques as a function of SNR. SNR was varied and at each SNR we computed reconstruction

estimates and error measures for ten independent realizations of the additive noise. We then averaged the value of each performance measure over the realizations for each reconstruction method. As above, for the algebraic techniques the regularization parameters were selected by finding the minimum MSE reconstruction for the given SNR. With the subspace techniques the regularization parameter was selected using the L-curve whenever the L-curve method proved useful. There was one case (the 0.5 cm width 2DS reconstructions) where the L-curve did not provide useful information even for the subspace techniques, as described below.

Figure 3.6 shows the graphs of the MSE versus SNR for all four F3D reconstruction techniques with (a) showing the 0.5 cm grid discretization and (b) showing the 2.0 cm discretization in the vertical direction. Each line in the figure shows the mean value of the MSE over the ten realizations for a specific reconstruction algorithm. The error bars show the standard deviation over the realizations. With a 0.5 cm grid reconstruction the two subspace techniques outperformed the two algebraic techniques, with TCG slightly outperforming TSVD, over the whole range of signal-to-noise ratios. ART shows a marked degradation in performance below 30 dB SNR. A MSE value for ART at a 10 dB signal-to-noise ratio was computed but was significantly outside the range of the graph. In the 2.0 cm case the subspace methods outperformed the algebraic methods except at high SNR ( $> 30$  dB) where ART provided the same performance as TCG. At lower signal to noise ratios the disparity between the best algebraic method (SIRT) and the worst subspace method was even greater than for the 0.5 cm vertical cell size case. It is also interesting to note that above 40dB signal-to-noise ratio there is no improvement in the MSE with higher SNR for this case.

Figure 3.7 shows the MSE curves for the equivalent 2DS reconstructions. Note that the vertical scale of these graphs is ten times larger than for the previous two graphs. Comparing 2DS and

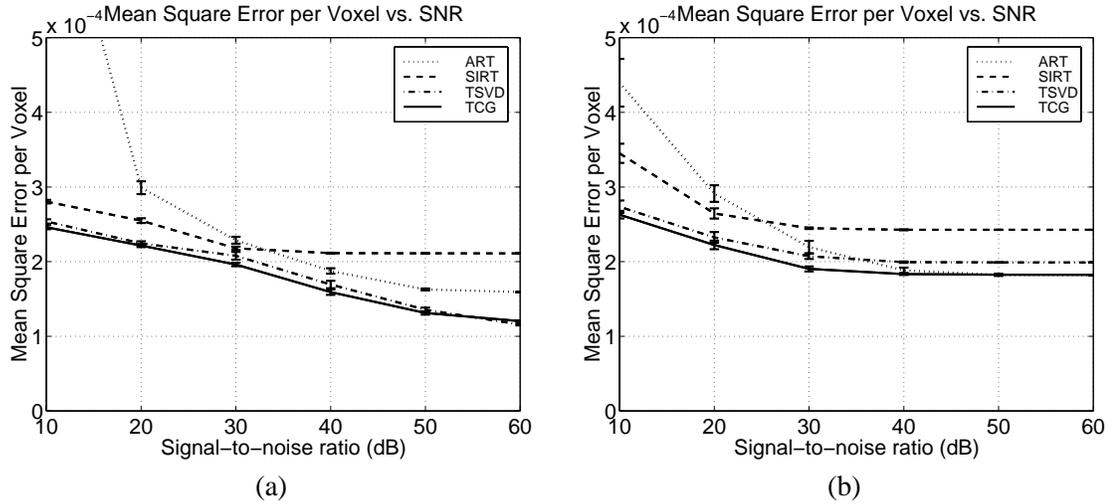


Figure 3.6: MSE for the F3D reconstructions. Panel (a) shows the 0.5 cm vertical cell size reconstruction, (b) shows the 2.0 cm vertical cell size reconstruction. Data points are means over ten realizations and error bars show  $\pm$  one standard deviation over the realizations. See text for details.

F3D results, we see that the MSE of the 2DS reconstructions is much larger than that of the F3D reconstructions. A curious result of the 2DS reconstructions is the lack of change in MSE of the subspace techniques with change in SNR. At present we do not have a good understanding of this phenomenon.

Figure 3.8(a) shows the OCE for the four reconstruction techniques for a F3D with a 0.5 cm vertical step size and Figure 3.8(b) shows the equivalent result for the 2.0 cm 2DS reconstruction. As with the MSE error measure the 0.5 cm vertical step size provided the best performance for F3D reconstructions and a 2.0 cm vertical step size provided the best performance for 2DS reconstructions. For this reason, from here on we present only these two cases. The results are generally similar to those of the MSE curves. Specifically, the two subspace techniques outperform the algebraic techniques over the whole range of SNRs and the OCE for the algebraic techniques degrades significantly below 30 dB SNR. The mean of the planar reconstruction is lower than for the full

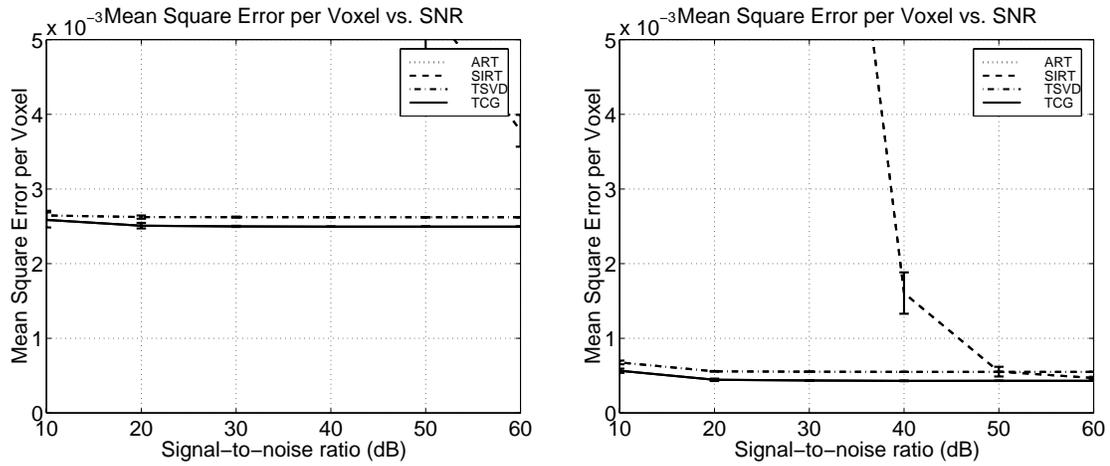


Figure 3.7: MSE for the 2DS reconstructions. Panel (a) shows the 0.5 cm vertical cell size reconstruction, (b) shows the 2.0 cm vertical cell size reconstruction. Data points are means over ten realizations and error bars show  $\pm$  one standard deviation over the realizations. See text for details.

three-dimensional reconstruction. We note, however, that the planar reconstruction used *a priori* information for the depth dimension, which is the largest component of positional error in the F3D reconstruction. Also, the standard deviations of the algebraic 2DS reconstructions were very large.

Figure 3.9(a) shows the AE performance for the F3D 0.5 cm reconstruction while Figure 3.9(b) shows the AE for the 2.0 cm 2DS reconstruction. As with the previous measures, the subspace techniques outperform the algebraic techniques over the range of SNRs evaluated. ART at 10 dB SNR for the F3D reconstruction appears to have a significant increase in performance, but upon examination of the reconstructions we found that this was just due to spurious noise in the region of the absorption anomaly. In other words, despite the low value of the AE measure, there was very little resemblance between the true absorption perturbation function and those generated by ART at a 10 dB signal-to-noise ratio, as indicated by the MSE curves in Fig. 3.6(a).

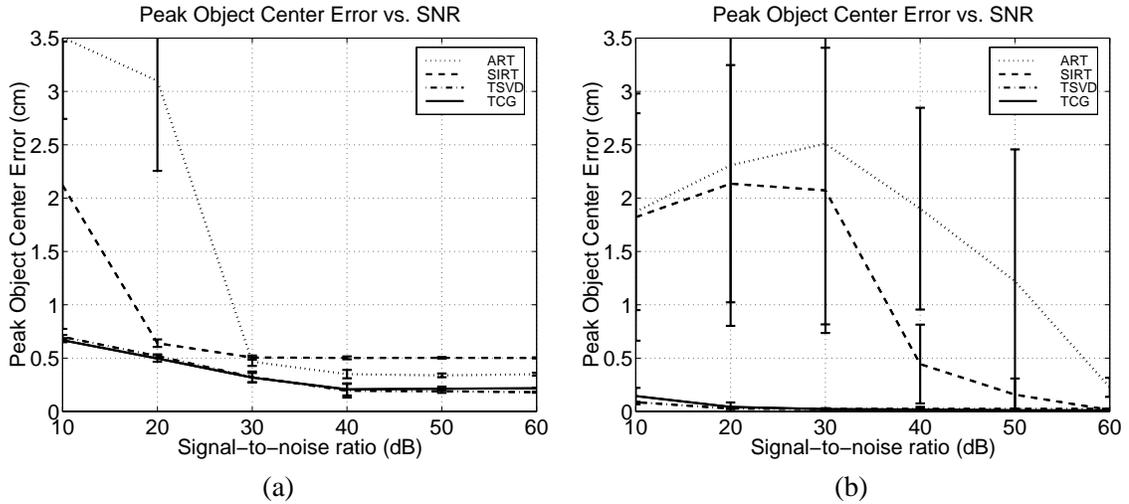


Figure 3.8: Panel (a) shows OCE for a F3D reconstruction with a 0.5 cm vertical step size. Panel (b) shows OCE for 2.0 cm thick planar reconstruction.

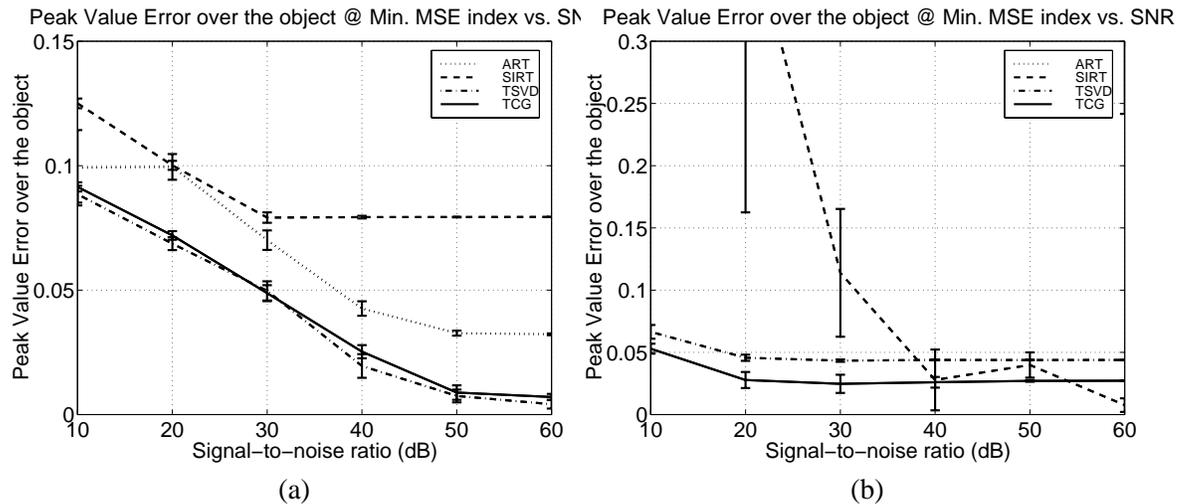


Figure 3.9: Panel (a) shows AE for the F3D 0.5 cm step size reconstruction. Panel (b) shows AE for the 2.0 cm 2DS reconstruction. Note that for the 2DS case the curve for ART is off the scale of the graph.

### 3.3.3 Regularization parameter sensitivity

Although we computed an L-curve for all of the reconstruction algorithms and scenarios, we found that the L-curve only identified a useful truncation parameter for the subspace techniques. Because the algebraic techniques do not monotonically reduce the residual error as the iterations progress, the L-curve graphs for ART and SIRT generally had multiple corners. We did, however, use the L-curve for selecting the regularization parameter for the two subspace techniques for the F3D reconstructions. For these techniques we can study how well the corner of the L-curve identified the “best” regularization parameter. To illustrate this comparison, for the two subspace techniques we graphed the MSE versus the regularization parameter over a range of SNRs for the F3D cases. The results are shown in Figure 3.10. For all cases it can be seen that the L-curve performed well at identifying a regularization parameter that was close to optimum in terms of minimizing the MSE.

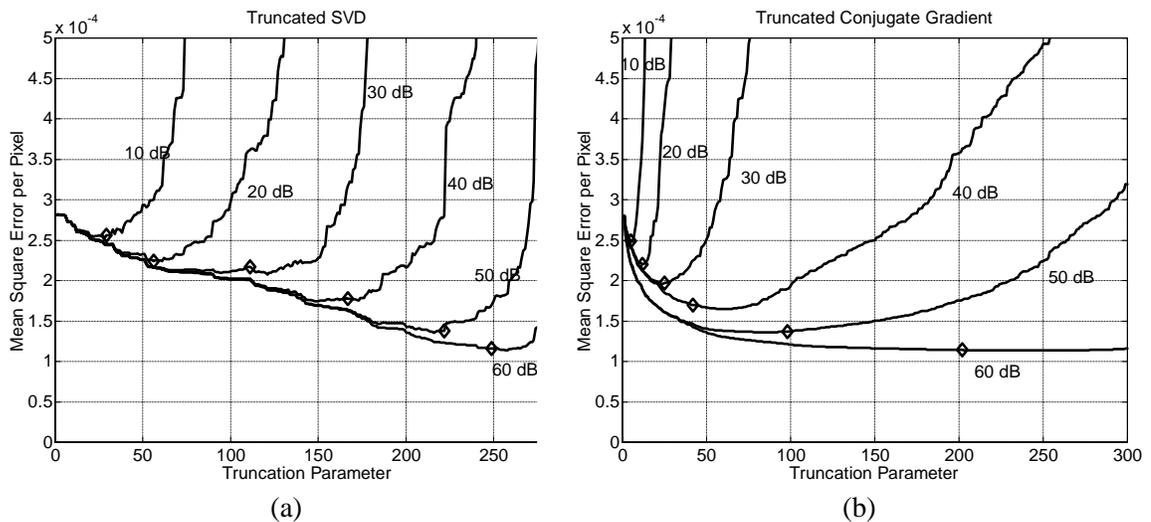


Figure 3.10: Regularization parameter sensitivity for the two subspace reconstruction techniques over a range of SNRs. One MSE curve is shown for each SNR. The diamond on each curve displays the regularization parameter we chose (without knowledge of the MSE curve) by simply manually selecting the corner of the L-curve. Panel (a) gives the results from the TSVD algorithm whereas panel (b) shows results from the TCG algorithm.

As a comparison of the regularization sensitivity for the algebraic techniques we show a plot of the MSE versus regularization parameter in Figure 3.11. From these plots we can see that at low SNRs the algebraic techniques are very sensitive to the regularization parameter and show a dramatic change in MSE with truncation parameter. At high SNRs the opposite behavior is evident.

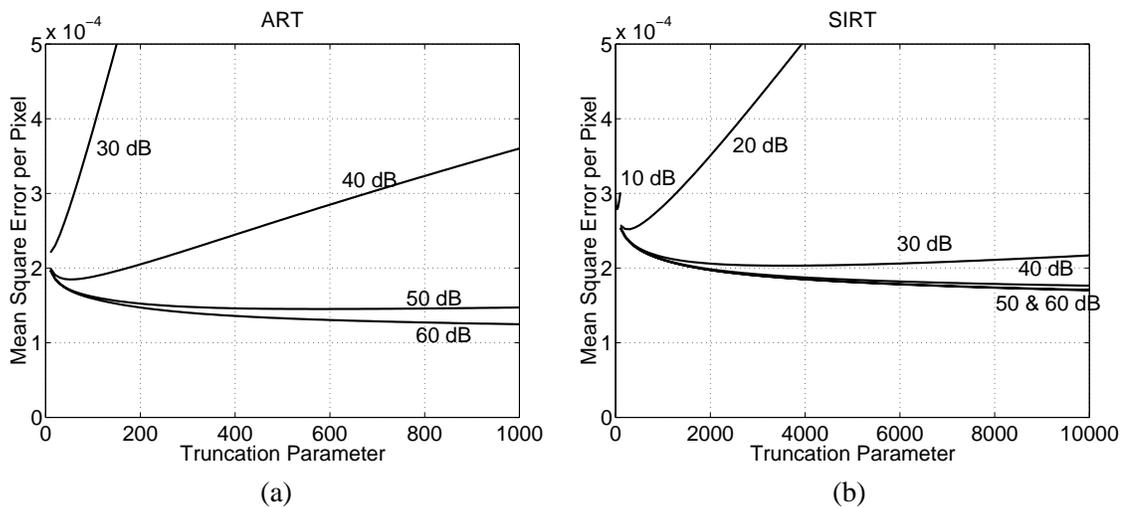


Figure 3.11: Regularization parameter sensitivity for the two algebraic reconstruction techniques over a range of SNRs. One MSE curve is shown for each SNR. Panel (a) is the results from the ART algorithm, the 10 and 20 dB curves are out of the range of the plot. Panel (b) is from the SIRT algorithm.

### 3.4 Linear model discussion

In this chapter we have presented a comparison of four of the most commonly used linear reconstruction techniques for DOT imaging applied to a three-dimensional reflective geometry reconstruction problem, reconstructing both the full three-dimensional volume and a single plane within

the volume. Our results showed that the subspace techniques are superior to the algebraic reconstruction techniques in estimating both the amplitude and the position of an inhomogeneity as well as in the overall fidelity of the absorption function reconstruction. This was true even when we used *a priori* knowledge of the true object function to select the optimal mean squared error regularization parameter for the algebraic techniques while we used only the *a posteriori* L-curve to select the regularization parameter for the subspace techniques.

We have also shown that for the highest level of quantitative accuracy a full three-dimensional reconstruction is necessary. Additionally, mismatch in the width of a two-dimensional slab reconstruction can have a significant adverse effect on the quality of the reconstruction, as is evident in Figures 3.5 and 3.7.

By using the OCE error measure we identified further differences in performance between the subspace and algebraic techniques that were not evident with only a MSE performance measure, particularly at low SNR. This conclusion was supported by direct observation of the reconstructed images at low SNR. We observed that even though the algebraic methods gave fairly good MSE performance at low SNR, the images looked nothing like the true object. The OCE measure identified this for low SNR in the F3D cases. Additionally, the large standard deviation of the OCE measure of the algebraic methods, particularly in the 2DS cases, indicated the poor performance of these techniques. We found that the AE measure did not prove as useful because it was fooled by spurious noise in the region of the object.

Through sensitivity studies we verified that the L-curve is a useful tool in selecting the regularization parameter for the subspace techniques in the three-dimensional reflective geometry reconstruction problem, at least when model mismatch is limited to moderate differences in discretization.

Some preliminary explorations using a FDFD forward model to simulate the noise free measured data show that the basic conclusions presented here hold for a linear reconstruction from a more accurate non-linear forward model. The primary difference we observed was reduction of the performance advantage of the subspace techniques at high SNR. At low and medium SNR we still see a significant performance advantage in the subspace techniques for the MSE and OCE measures.

## Chapter 4

# Two wavelength constrained reconstruction

We saw in the last chapter that there are apparent limitations to the linear model reconstruction approach. This is certainly to be expected given the ill-posed nature of both the two-dimensional and three-dimensional reconstructions and the severity of the underdetermined nature of the three-dimensional reconstruction problem. In this chapter we describe a technique to address this problem through the application of a specific *a priori* constraint on the solution to the system. Solution constraints can come in many forms; some examples include constraints on the size of the solution through norms, and constraints on the smoothness of the solution such as through the norm of the Laplacian. Our approach has been to develop a constraint on the relationship of the absorption coefficient maps at two different optical wavelengths. We were motivated to adopt this approach by the fact that in practice one needs to collect data at at least two wavelengths to compute the total

hemoglobin concentration and oxygen saturation.

## 4.1 Constant absorption coefficient ratio (CACR) constraint

We started first with the assumption that the underlying tissue can be spatially categorized in one of two states, either normal or cancerous. The latter we refer to as an anomalous region. Furthermore, we also assumed that all tissue of each type, be it either normal or anomalous, has the same absorption coefficient, and that its absorption coefficient differs from that of the other type. The physiological rationale is that a cancerous region is expected to have a higher total hemoglobin concentration due to angiogenesis and reduced oxygenation due to a higher metabolic rate of the cancerous tissue [6, 7, 8].

The relationship between the absorption coefficient and the blood volume fraction and blood oxygenation is described by the equation

$$\mu_a(\mathbf{r}) = b_{vf}(\mathbf{r}) \left( R_{Hb0_2}(\mathbf{r}) \mu_a^{Hb0_2} + R_{Hb}(\mathbf{r}) \mu_a^{Hb} \right), \quad (4.1)$$

where  $\mu_a(\mathbf{r})$  is the absorption coefficient as a function of space,  $b_{vf}(\mathbf{r})$  is the spatial blood volume function,  $\mu_a^{Hb0_2}$  and  $\mu_a^{Hb}$  are the molar absorption coefficients for oxygenated and deoxygenated hemoglobin at a specific wavelength, and  $R_{Hb0_2}(\mathbf{r})$  and  $R_{Hb}(\mathbf{r})$  are percentage of oxy- and deoxy-hemoglobin. The functions  $R_{Hb0_2}(\mathbf{r})$  and  $R_{Hb}(\mathbf{r})$  sum to one everywhere, as specified by

$$R_{Hb0_2}(\mathbf{r}) + R_{Hb}(\mathbf{r}) = 1.$$

The structure of Eq. (4.1) also provided us with an important property of the constraint. The spatially varying terms,  $R_{Hb0_2}(\mathbf{r})$  and  $R_{Hb}(\mathbf{r})$ , are independent of optical wavelength, while the

terms that vary with optical wavelength,  $\mu_a^{Hb0_2}$  and  $\mu_a^{Hb}$ , are independent of position. This results in spatial absorption functions,  $\mu_a(r)$ , that have the same boundaries between regions regardless of optical wavelength used to interrogate the medium.

Now consider an implementation of the first Born approximation described in Section 2.2 where the background absorption coefficient is selected to be the same as the absorption coefficient of the normal tissue described above. The equations for the perturbation relationship at the two optical wavelengths,  $\lambda_1$  and  $\lambda_2$  are given by

$$\begin{aligned}\tilde{\mathbf{G}}_{\lambda_1} \mathbf{k}_{\lambda_1} &= \tilde{\Phi}_{\lambda_1}^s & \mathbf{k}_{\lambda_1} &\in \mathbb{R}^n, & \tilde{\Phi}_{\lambda_1}^s &\in \mathbb{R}^{2m} \\ \tilde{\mathbf{G}}_{\lambda_2} \mathbf{k}_{\lambda_2} &= \tilde{\Phi}_{\lambda_2}^s & \mathbf{k}_{\lambda_2} &\in \mathbb{R}^n, & \tilde{\Phi}_{\lambda_2}^s &\in \mathbb{R}^{2m}\end{aligned}\quad (4.2)$$

where  $\tilde{\Phi}_{\lambda_1}^s$  and  $\tilde{\Phi}_{\lambda_2}^s$  are the scattered fluence rate estimates,  $\tilde{\mathbf{G}}_{\lambda_1}$  and  $\tilde{\mathbf{G}}_{\lambda_2}$  are the first Born coefficient matrices, and  $\mathbf{k}_{\lambda_1}$  and  $\mathbf{k}_{\lambda_2}$  are the absorption perturbation functions, and all of the quantities are designated as corresponding to the first and second wavelengths respectively. For  $m$  sensors the measurement vectors are  $2m \times 1$  because we are assuming a sinusoidally modulated source resulting in the real and imaginary part being stacked in real vector as described in Eq. 2.16. Under this implementation the normal tissue should be reconstructed as a zero perturbation and an anomalous region should reconstruct to the difference between the anomalous absorption and the background absorption. This situation holds for both wavelengths, and thus results in perturbation functions at each wavelength that are related by a constant value

$$\mathbf{k}_{\lambda_1} = \alpha \mathbf{k}_{\lambda_2}, \quad (4.3)$$

where  $\alpha$  is the constant that equates the absorption coefficients of the anomalous regions at the two wavelengths. We have called this relationship the Constant Absorption Coefficient Ratio (CACR)

constraint. We also note that this relationship holds only when the background regions are zero in the perturbation formulation such as the Born approximation we have employed.

## 4.2 Applying the CACR constraint

The idea here is to reconstruct absorption perturbation maps that both agree with the measured (or simulated) fluence rate data as well meet the CACR constraint. There are number of methods in which the CACR constraint can be applied to estimating the absorption perturbation functions  $k_{\lambda_1}$  and  $k_{\lambda_2}$ . We examined several approaches to the reconstruction of two-wavelength data, all of which are based on the truncated singular decomposition algorithm described in Section 3.2.2. Two of the approaches we examined applied the CACR constraint to the estimate of the absorption perturbation function. The first of the two constrained approaches we developed applied the CACR constraint after estimating the absorption perturbation function independently at each wavelength. This approach we have called the post-constrained TSVD (PC-TSVD) algorithm. The other constrained approach we developed enforced the constraint within the inversion process itself. This algorithm was derived from the modified TSVD algorithm developed by Hansen [47] and we have called it the joint modified TSVD (JM-TSVD) algorithm. The other two approaches we examined did not apply the constraint. One, called the joint TSVD (J-TSVD), simply considered all of the data from both wavelengths jointly to estimate two separate absorption perturbation functions. The final algorithm used the TSVD independently on the data at each wavelength to estimate the absorption perturbation functions separately. We present the algorithms below followed by a description of the simulation scenario used to analyze their performance. In the last section of this chapter we present

the performance results of each of the algorithms.

#### 4.2.1 Independent TSVD reconstruction

All of the joint two-wavelength reconstruction algorithms are based on the TSVD algorithm. Thus we first present a description of the TSVD algorithm and then develop that into the J-TSVD algorithm, followed by the PC-TSVD algorithm, and finally the JM-TSVD algorithm. As a baseline for evaluating the performance of the two-wavelength reconstruction algorithms we estimated the absorption perturbation functions independently for each wavelength using the TSVD algorithm presented in Section 3.2.2, specifically in Eq. (3.3). Thus the TSVD estimates are

$$\hat{\mathbf{k}}_{\lambda_1} = \mathbf{V}_{\lambda_1,t} \boldsymbol{\Sigma}_{\lambda_1,t}^{-1} \mathbf{U}_{\lambda_1,t}^T \tilde{\boldsymbol{\Phi}}_{\lambda_1}^s$$

$$\hat{\mathbf{k}}_{\lambda_2} = \mathbf{V}_{\lambda_2,t} \boldsymbol{\Sigma}_{\lambda_2,t}^{-1} \mathbf{U}_{\lambda_2,t}^T \tilde{\boldsymbol{\Phi}}_{\lambda_2}^s$$

where  $\hat{\mathbf{k}}_{\lambda_1}$  and  $\hat{\mathbf{k}}_{\lambda_2}$  are the estimates of the absorption perturbation function, and  $\mathbf{V}_{\lambda_1,t} \boldsymbol{\Sigma}_{\lambda_1,t}^{-1} \mathbf{U}_{\lambda_1,t}^T$  and  $\mathbf{V}_{\lambda_2,t} \boldsymbol{\Sigma}_{\lambda_2,t}^{-1} \mathbf{U}_{\lambda_2,t}^T$  are TSVD matrices for  $\tilde{\mathbf{G}}_{\lambda_1}$  and  $\tilde{\mathbf{G}}_{\lambda_2}$  respectively. As in Eq. (3.3),  $t$  specifies the truncation parameter.

#### 4.2.2 Joint truncated singular value decomposition (J-TSVD)

The independent TSVD algorithm above does not take advantage of any relationships between the two wavelengths. Also, the truncation parameters for each system must be selected separately. A straight forward approach to combining the information from each wavelength was to construct a large diagonal system encompassing both of the systems in Eq. (4.2). This allowed us to simultaneously select the best joint subspace from both wavelengths to estimate  $\mathbf{k}_{\lambda_1}$  and  $\mathbf{k}_{\lambda_2}$ . The large

diagonal system is given by

$$\tilde{\mathbf{G}}_{\mathbf{b}} \mathbf{k}_{\mathbf{b}} = \Phi_{\mathbf{b}}^s, \quad \mathbf{k}_{\mathbf{b}} \in \mathbb{R}^{2n}, \quad \Phi_{\mathbf{b}}^s \in \mathbb{R}^{4m}, \quad (4.4)$$

where

$$\mathbf{G}_{\mathbf{b}} = \begin{bmatrix} \tilde{\mathbf{G}}_{\lambda_1} & 0 \\ 0 & \tilde{\mathbf{G}}_{\lambda_2} \end{bmatrix}, \quad \mathbf{k}_{\mathbf{b}} = \begin{bmatrix} \mathbf{k}_{\lambda_1} \\ \mathbf{k}_{\lambda_2} \end{bmatrix}, \quad \Phi_{\mathbf{b}}^s = \begin{bmatrix} \tilde{\Phi}_{\lambda_1}^s \\ \tilde{\Phi}_{\lambda_2}^s \end{bmatrix}. \quad (4.5)$$

We then applied the TSVD algorithm to the large system to get an estimate of  $\mathbf{k}_{\mathbf{b}}$ . This is given by

$$\hat{\mathbf{k}}_{\mathbf{b}} = \mathbf{V}_{b,t} \Sigma_{b,t}^{-1} \mathbf{U}_{b,t}^T \Phi_{\mathbf{b}}^s \quad (4.6)$$

where  $\mathbf{U}_b \Sigma_b \mathbf{V}_b^T$  is the SVD of the large system  $\mathbf{G}_{\mathbf{b}}$ . Because the SVD algorithm is  $\mathcal{O}(n^3)$  we wished not to perform the straight forward computation of the SVD of  $\mathbf{G}_{\mathbf{b}}$ . We avoided this by observing that

$$\begin{bmatrix} \tilde{\mathbf{G}}_{\lambda_1} & 0 \\ 0 & \tilde{\mathbf{G}}_{\lambda_2} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_{\lambda_1} & 0 \\ 0 & \mathbf{U}_{\lambda_2} \end{bmatrix} \begin{bmatrix} \Sigma_{\lambda_1} & 0 \\ 0 & \Sigma_{\lambda_2} \end{bmatrix} \begin{bmatrix} \mathbf{V}_{\lambda_1} & 0 \\ 0 & \mathbf{V}_{\lambda_2} \end{bmatrix}^T. \quad (4.7)$$

Also, since the matrices  $\mathbf{U}_{\lambda_1}$ ,  $\mathbf{U}_{\lambda_2}$ ,  $\mathbf{V}_{\lambda_1}$ , and  $\mathbf{V}_{\lambda_2}$  are orthonormal, the large matrices containing them are also orthonormal. Thus to calculate the SVD of the large system  $\mathbf{U}_b$  we computed the SVD of the two smaller systems  $\tilde{\mathbf{G}}_{\lambda_1}$  and  $\tilde{\mathbf{G}}_{\lambda_2}$ , combined the results as in Eq. (4.7), and then permuted the columns of left and right singular vectors so that the diagonal elements of matrix containing  $\Sigma_{\lambda_1}$  and  $\Sigma_{\lambda_2}$  are decreasing in magnitude. Thus, computing the SVD of the large system had only a slightly higher computational cost than computing both of the SVDs of the small system.

### 4.2.3 Post constrained truncated singular value decomposition (PC-TSVD)

The J-TSVD algorithm above computes a minimum norm reconstruction using the best available information from the large linear system, but without using the CACR constraint. The next approach

in our development of a two-wavelength constrained reconstruction algorithm was to apply the CACR constraint as a post-processing step, which we call PC-TSVD. The PC-TSVD computes intermediate estimates  $\hat{\mathbf{k}}_{\lambda_1}$  and  $\hat{\mathbf{k}}_{\lambda_2}$  using the J-TSVD algorithm. Using the intermediate estimates the constraint was applied in the following manner. First we computed the average of the two intermediate estimates given by

$$\bar{\mathbf{k}} = \frac{1}{2} (\hat{\mathbf{k}}_{\lambda_1} + \hat{\mathbf{k}}_{\lambda_2}). \quad (4.8)$$

Then the CACR constraint relationship between the two absorption perturbation functions was implemented by substituting the constraint equation Eq. (4.3) into Eq. (4.8) which gives

$$\begin{aligned} \hat{\mathbf{k}}_{\lambda_1} &= \frac{2\alpha}{1+\alpha} \bar{\mathbf{k}} \\ \hat{\mathbf{k}}_{\lambda_2} &= \frac{2}{1+\alpha} \bar{\mathbf{k}}, \end{aligned} \quad (4.9)$$

where  $\hat{\mathbf{k}}_{\lambda_1}$  and  $\hat{\mathbf{k}}_{\lambda_2}$  are the PC-TSVD estimates of the absorption perturbation functions for the first and second wavelengths respectively.

#### 4.2.4 Joint modified truncated singular value decomposition (JM-TSVD)

By applying the constraint to the intermediate solutions the in PC-TSVD algorithm we end up with a result that meets the constraint, but to achieve this we have moved our estimates away from matching the approximation of the data chosen by the TSVD algorithm. What we might prefer would be to have the absorption perturbation function estimate agree with the regularized data that has been measured (or simulated) and still attempt to achieve the constraint. One way to do this is to use the range space component of the absorption perturbation function estimate to meet the data and use the null space component to achieve the constraint condition. This makes sense at a heuristic

level since the measurement data  $\Phi_{\mathbf{b}}^s$  results only from range space of the operator  $\mathbf{G}_{\mathbf{b}}$ . We derived this approach from research by Hansen *et al.* [47].

To develop the JM-TSVD algorithm we first expressed the absorption perturbation function estimate as the sum of two orthogonal components

$$\hat{\mathbf{k}}_{\mathbf{b}} = \hat{\mathbf{k}}_{\mathcal{R}} + \hat{\mathbf{k}}_{\mathcal{N}} \quad \hat{\mathbf{k}}_{\mathcal{R}} \in \mathcal{R}(\mathbf{G}_{\mathbf{b}}) \quad \hat{\mathbf{k}}_{\mathcal{N}} \in \mathcal{N}(\mathbf{G}_{\mathbf{b}}) \quad (4.10)$$

where  $\hat{\mathbf{k}}_{\mathcal{R}}$  is the range space component of the absorption perturbation function estimate and  $\hat{\mathbf{k}}_{\mathcal{N}}$  is the null space component. Estimation of the range space component was accomplished using the TSVD algorithm as in the previous techniques. Specifically, this is given by

$$\hat{\mathbf{k}}_{\mathcal{R}} = \mathbf{V}_{b,t} \Sigma_{b,t}^{-1} \mathbf{U}_{b,t}^T \Phi_{\mathbf{b}}^s = \mathbf{G}_{b,t}^{-1} \Phi_{\mathbf{b}}^s \quad (4.11)$$

where, again,  $t$  specifies the truncation parameter and  $\mathbf{G}_{b,t}$  is the reduced rank approximation of  $\mathbf{G}_{\mathbf{b}}$ . Looking at this equation we note that a truncation parameter of  $t$  results in the first  $t$  columns of  $\mathbf{V}_b$  spanning the effective range space of  $\mathbf{G}_{\mathbf{b}}$ . This results in a partition of the right singular vectors, which span the complete solution space, into two sets

$$\mathbf{V}_{\mathbf{b}} = [\mathbf{V}_{b,t} \quad \mathbf{V}_{b,\mathcal{N}}], \quad (4.12)$$

where  $\mathbf{V}_{b,\mathcal{N}}$  is the remaining  $2n - t$  columns of  $\mathbf{V}_b$  and is an orthogonal basis set for the effective null space of  $\mathbf{G}_{\mathbf{b}}$ . Using this orthogonal basis set we then found an appropriate component of the null space that allows the range space portion of the estimate to meet the data and the total solution to minimize the constraint function.

The constraint can be expressed as the minimization

$$\min \|\ [\mathbf{I}_n \quad -\alpha \mathbf{I}_n] \hat{\mathbf{k}}_{\mathbf{b}} \|_2 \quad \text{subject to.} \quad \mathbf{G}_{b,t} \hat{\mathbf{k}}_{\mathbf{b}} = \Phi_{\mathbf{b}}^s. \quad (4.13)$$

Substituting Eq. (4.10) in for  $\hat{\mathbf{k}}_{\mathbf{b}}$  and expressing the null space component in terms of the null space spanning vectors gives

$$\min || [\mathbf{I}_n - \alpha \mathbf{I}_n] \hat{\mathbf{k}}_{\mathcal{R}} + [\mathbf{I}_n - \alpha \mathbf{I}_n] \mathbf{V}_{\mathbf{b},\mathcal{N}} \mathbf{z} ||_2 \quad \text{subject to} \quad \mathbf{G}_{\mathbf{b},t} \hat{\mathbf{k}}_{\mathbf{b}} = \Phi_{\mathbf{b}}^{\mathbf{s}} \quad (4.14)$$

where we have substituted  $\hat{\mathbf{k}}_{\mathcal{N}} = \mathbf{V}_{\mathbf{b},\mathcal{N}} \mathbf{z}$  with  $\mathbf{z}$  being the Fourier coefficients of the null space component. The constraint  $\mathbf{G}_{\mathbf{b},t} \hat{\mathbf{k}}_{\mathbf{b}} = \Phi_{\mathbf{b}}^{\mathbf{s}}$  is enforced implicitly because the range space portion of the estimate is calculated using the TSVD and the minimization only operates on the null space portion of the estimate. Looking at the above minimization we can see that the left hand term in the norm needs to remain constant for the solution to agree with the data, and the right hand term is a matrix vector product with  $2n - t$  unknowns. This has the same form as the standard two-norm minimization problem  $\min || -\mathbf{b} + \mathbf{A}\mathbf{x} ||_2$  where  $\mathbf{b} = -[\mathbf{I}_n - \alpha \mathbf{I}_n] \hat{\mathbf{k}}_{\mathcal{R}}$  and  $\mathbf{A} = [\mathbf{I}_n - \alpha \mathbf{I}_n] \mathbf{V}_{\mathbf{b},\mathcal{N}}$ . In our scenario with  $n$  unknowns at each wavelength and  $m$  measurements resulting in  $[\mathbf{I}_n - \alpha \mathbf{I}_n] \mathbf{V}_{\mathbf{b},\mathcal{N}} \in \mathbb{R}^{n \times 2n-t}$ , the estimations of both the range and the null space components are underdetermined because  $m$  and  $t$  are less than  $n$ . To handle this we have selected the minimum norm solution for both parts of the problem resulting in a minimum norm solution that meets the data and the CACR constraint. The minimum norm solution for the range space portion is given by the TSVD algorithm described above. We computed the minimum norm solution for the null space component by first computing the QR decomposition of the transpose of the matrix above,

$$\mathbf{QR} = ([\mathbf{I}_n - \alpha \mathbf{I}_n] \mathbf{V}_{\mathbf{b},\mathcal{N}})^T. \quad (4.15)$$

Using the QR decomposition the minimum norm solution for  $\mathbf{z}$  is given by [44]

$$\mathbf{z}_{mn} = \mathbf{Q} \left( \mathbf{R}^T \setminus \left( -[\mathbf{I}_n - \alpha \mathbf{I}_n] \hat{\mathbf{k}}_{\mathcal{R}} \right) \right) \quad (4.16)$$

where  $\backslash$  represents solution by back substitution, since  $\mathbf{R}$  is upper triangular.

Like the previously described reconstruction methods, we do not know *a priori* what is the best regularization (truncation) parameter for performing the reconstruction. To select the best truncation parameter we computed the L-curve as we did in the linear model reconstruction technique evaluation described in the previous chapter. This required that we compute both the range and null space components of the solution for a large number of truncation parameters. For the range space component this is efficiently done by computing in additional singular values and right singular vectors. The null space component was more complicated since for each new range space component both the right hand side and the matrix of the null space minimum norm problem changed. To solve the problem efficiently we adopted a recursive update-type approach. We started by computing the solution for the largest truncation parameter,  $tmax$ , we were interested in examining. To simplify the notation we start by identifying the null space minimum norm problem matrix from Eq. (4.15) as

$$\mathbf{N}_{tmax} = ([\mathbf{I}_n \quad -\alpha\mathbf{I}_n] \mathbf{V}_{\mathbf{b},\mathcal{N}})^T \quad (4.17)$$

where  $\mathbf{V}_{\mathbf{b},\mathcal{N}}$  for this truncation parameter is the rightmost  $2n - tmax$  columns of  $\mathbf{V}_{\mathbf{b}}$ . The QR decomposition for the null space minimum norm problem is given by

$$\mathbf{Q}_{tmax}\mathbf{R}_{tmax} = \mathbf{N}_{tmax} \quad (4.18)$$

Reducing the truncation parameter by one removes the right hand column from  $\mathbf{V}_{\mathbf{b},t}$  and places it on the left hand side of  $\mathbf{V}_{\mathbf{b},\mathcal{N}}$ . This results in adding a row to the top of  $\mathbf{N}_{tmax}$  and modifying the

QR relationship by

$$\begin{bmatrix} 1 & 0 \\ 0 & \mathbf{Q}_{tmax} \end{bmatrix} \begin{bmatrix} \mathbf{n}_{tmax-1} \\ \mathbf{R}_{tmax} \end{bmatrix} = \begin{bmatrix} \mathbf{n}_{tmax-1} \\ \mathbf{N}_{tmax} \end{bmatrix}, \quad (4.19)$$

where  $\mathbf{n}_{tmax-1}$  is the new row that results from moving a column of  $\mathbf{V}_{b,t}$  to  $\mathbf{V}_{b,\mathcal{N}}$ . The relationship shown in Eq. (4.19) is close to but not quite a QR decomposition of  $\mathbf{N}_{tmax-1}$ , since the first sub-diagonal of the new  $\mathbf{R}$  matrix is non-zero. This was remedied by Givens rotations applied to zero out the first sub-diagonal. The orthogonal Givens rotation matrices were incorporated into  $\mathbf{Q}$  to generate the new QR decomposition.

To reiterate, the complete two-wavelength reconstruction algorithm is:

- select a range of truncation parameters to examine  $[tmin \ tmax]$
- calculate the SVD of the large system by computing the SVD of the system matrix for each wavelength
- calculate the range space component for the selected set of truncation parameters
- compute the QR decomposition of the null space minimum norm matrix associated with truncation parameter  $tmax$
- solve the null space minimum norm problem using the QR decomposition
- sum the range space and null space component to compute the solution at  $tmax$
- iterate down from  $tmax$  to  $tmin$ 
  - compute the new QR decomposition of the new null space minimum norm matrix using the previous QR decomposition and Givens rotations

- solve the null space minimum norm problem using the new QR decomposition
- sum the range space and null space component to compute the solution at the current truncation parameter

### 4.3 Simulation scenario

To examine the performance of the two-wavelength reconstruction algorithms we implemented a computer simulation similar to that described in the previous chapter for the linear model reconstruction evaluation. The computational volume and absorption anomaly spatial parameters were identical to the previous simulation. Differences in the scenario and additional parameters for this two-wavelength simulation are described below.

The two optical wavelengths that we simulated for this investigation were 780 nm and 830 nm. 780 nm is the same wavelength we used in the linear model reconstruction evaluation, while 830 nm is on the other side of the absorption coefficient crossover point for oxy-hemoglobin and deoxy-hemoglobin (see Figure 1.1). We kept the hemoglobin concentration and oxygen saturation parameters the same, which resulted background absorption coefficients of  $0.041 \text{ cm}^{-1}$  at 780 nm and  $0.054 \text{ cm}^{-1}$  at 830 nm. The absorption anomaly was simulated with absorption coefficients of  $0.18 \text{ cm}^{-1}$  at 780 nm and  $0.20 \text{ cm}^{-1}$  at 830 nm.

While the detector positions in these simulations remained the same as the linear model reconstruction evaluation, two different source geometries were studied. One source geometry had the nine 780 nm and the nine 830 nm sources all collocated in the same locations as the investigation in the previous chapter. A graphical illustration of the source detector geometry for this scenario

is shown in Figure 4.1. The other source geometry was motivated by a hardware implementation being developed in David Boas' laboratory at Massachusetts General Hospital. This source geometry placed all nine 780 nm sources and all nine 830 nm sources in unique locations approximately uniformly interspersed between the detectors. This source layout is illustrated in Figure 4.2.

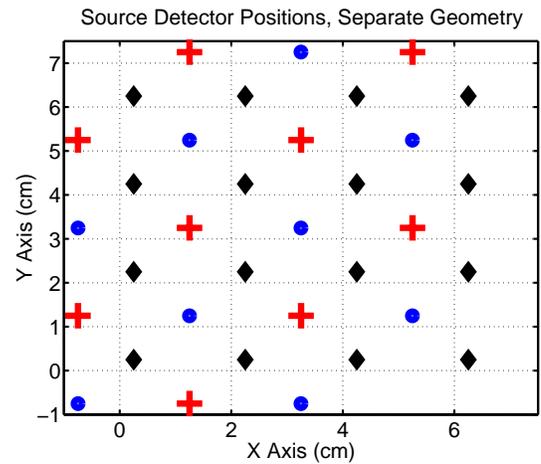
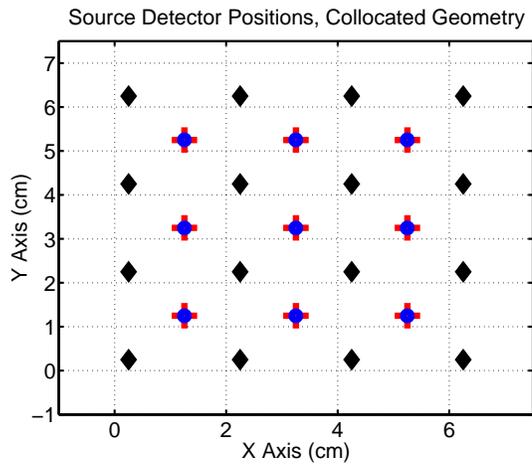


Figure 4.1: Source and detector positions for the collocated source geometry. Red crosses indicate the positions of the 780 nm sources, blue circles indicate the positions of the 830 nm sources which are collocated with the 780 nm sources. Black diamonds illustrate positions of the detectors.

Figure 4.2: Source and detector positions for the separate source geometry. Red crosses indicate the positions of the 780 nm sources, blue circles indicate the positions of the 830 nm sources. Black diamonds illustrate positions of the detectors.

The remaining difference compared to the linear model reconstruction evaluation was the use of two noise models. Both the uniform variance noise model described in Section 2.3.1 and the simplified shot noise model described in Section 3.1.5 were used to study the effects on the reconstruction algorithms.

### 4.3.1 Two wavelength SV spectra

The addition of the sources from the second wavelength and variation in source geometry both lead to modifications of the forward operator. To study the changes in the forward operator and the associated changes in the linear inverse problem we have plotted the singular value spectra for the forward models. Figure 4.3 shows the singular value spectra for the collocated source geometry. Both the singular values from the 780 nm and 830 nm wavelength system matrices are shown. These singular value spectra show very little difference in terms of structure as would be expected because the change from 780 nm to 830 nm just slightly modifies the complex wavenumber for the diffusion equation. By stacking the forward operators from each wavelength and computing the singular value spectrum of this combined matrix, we were able to measure how independent the rows of the forward operators are, and get an idea of how the “posedness” of the two-wavelength problem has changed relative to the single wavelength problem. This singular value spectrum is shown as the solid green line in Figure 4.3. For the large singular values we see that the combined singular value spectrum closely follows that of the two single wavelength spectra. We expected this result also, again because of the small modification to the diffusion equation wavenumber. Only after the singular values of the combined system diminish by several orders of magnitude do we see a significant change in the slope compared to the two single wavelength spectra.

The singular value spectra for the separate source geometry is shown in Figure 4.4. We observe similar features in this set of spectra to those shown in the collocated source geometry singular value spectra. The major difference is the higher value of singular values at which the combined singular value spectrum changes slope compared to the two single wavelength spectra for this geometry.

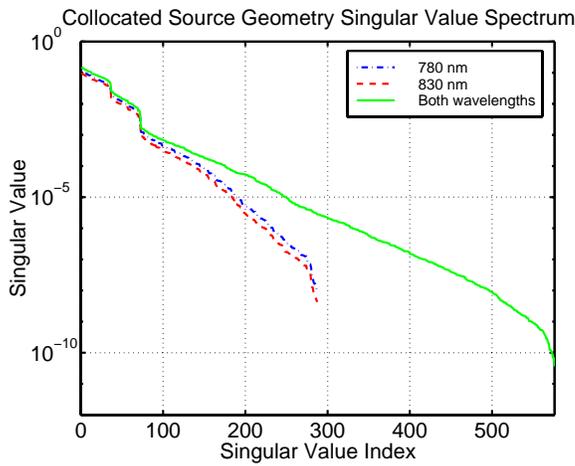


Figure 4.3: Singular value spectra for the collocated source geometry. The blue dash-dotted line shows the singular value spectrum for the 780 nm Born forward operator. The red dashed line shows the singular value spectrum for the 830 nm Born forward operator. The solid green line shows the singular value spectrum for both Born forward operators stacked.

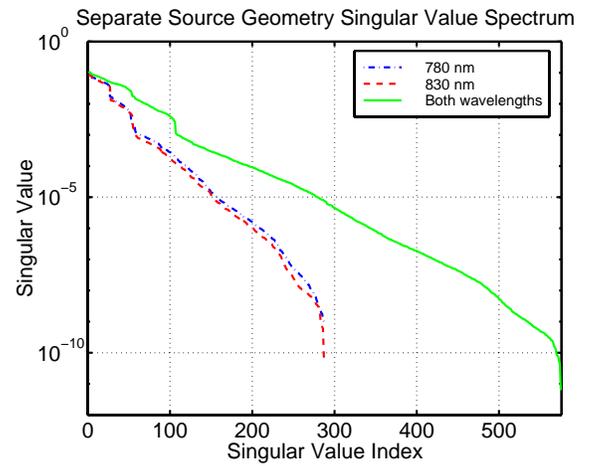


Figure 4.4: Singular value spectra for the separate source geometry. The blue dash-dotted line shows the singular value spectrum for the 780 nm Born forward operator. The red dashed line shows the singular value spectrum for the 830 nm Born forward operator. The solid green line shows the singular value spectrum for both Born forward operators stacked.

## 4.4 Simulation results

We ran simulations of the four different scenarios (two source geometries and two noise models) and then applied each of the four different reconstruction algorithms described in this chapter to create estimates of the absorption perturbation function. We first present some of the qualitative results from two of the simulation scenarios followed by a quantitative comparison using two of the performance measures described in Chapter 3.

### 4.4.1 Qualitative results

To qualitatively compare the reconstruction performance of the four algorithms we have plotted a vertical slice of the reconstructed three-dimensional absorption perturbation function in Figures 4.5 through 4.8 at 780 nm. The slice we show in these figures is at  $Y = 3$  cm, which is through the center of the absorption anomaly, and thus an accurate reconstruction would show a circular object centered  $Z = -2.5$  cm and  $X = 2$  cm. The source geometry used for the simulation was the collocated one. The SNR in all cases was 30 dB and the noise model was a uniform distribution. The truncation parameter selected for each algorithm was the one that minimized the MSE measure.

We observe very similar reconstruction qualities for the single wavelength TSVD and the two wavelength J-TSVD and PC-TSVD algorithms. The peak absorption coefficient perturbation value was approximately  $0.05 \text{ cm}^{-1}$ ,  $1/3$  of the true value. The center of reconstructed anomaly was about 1.5 cm into the medium. The JM-TSVD algorithm gives a slightly higher peak absorption coefficient perturbation and slightly deeper center for the absorption anomaly. In all cases there is noise corruption near the surface of the reconstruction with a small negative region just above the

object.

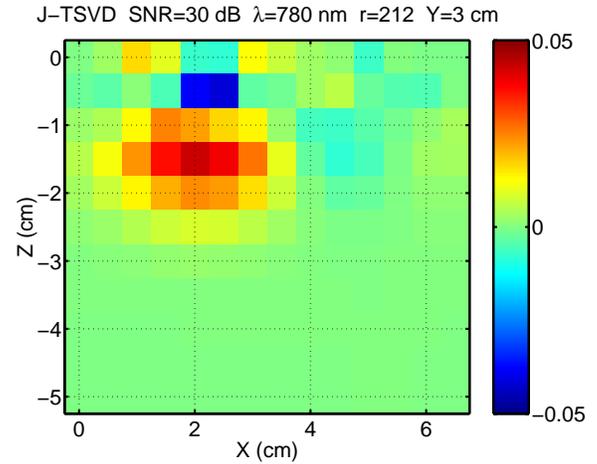
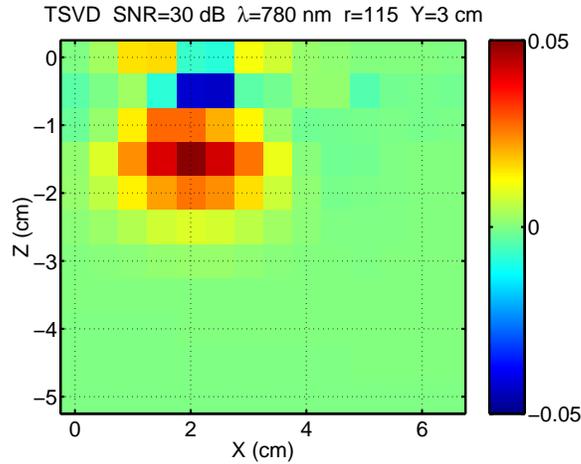


Figure 4.5: Collocated source single wavelength TSVD reconstruction at 30 dB SNR from a uniform noise model.

Figure 4.6: Collocated source two-wavelength J-TSVD reconstruction at 30 dB SNR from a uniform noise model.

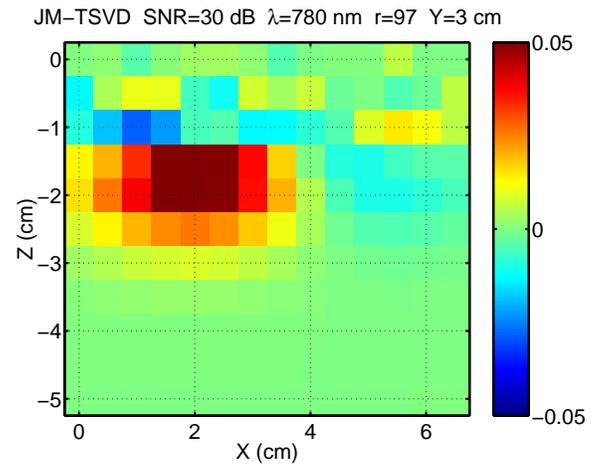
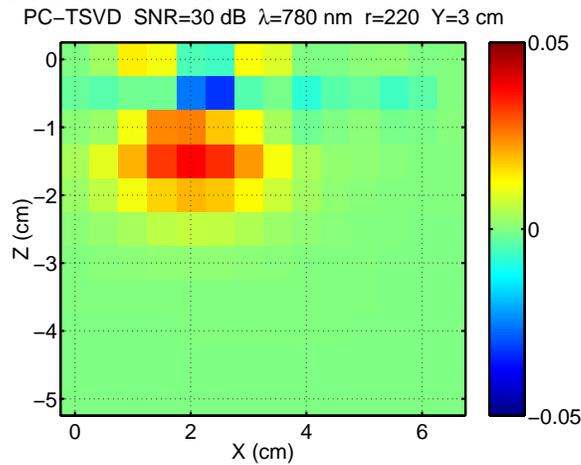


Figure 4.7: Collocated source two-wavelength PC-TSVD reconstruction at 30 dB SNR from a uniform noise model.

Figure 4.8: Collocated source two-wavelength JM-TSVD reconstruction at 30 dB SNR from a uniform noise model.

Figures 4.9 through 4.12 show the same scenario as the previous set of figure except that now the separate source geometry has been modeled. Clearly the reconstructed anomaly amplitude has been reduced in all cases except for the PC-TSVD reconstruction. The range plotted in these figures is

60% of the range for the collocated geometry scenario in the previous figures. With separate source geometry these figures indicate that PC-TSVD algorithm provides a slightly better reconstruction quality. The JM-TSVD algorithm appears to perform much worse than the previous three with an amplitude response that is significantly reduced and spatial structure that is almost indistinguishable from noise.

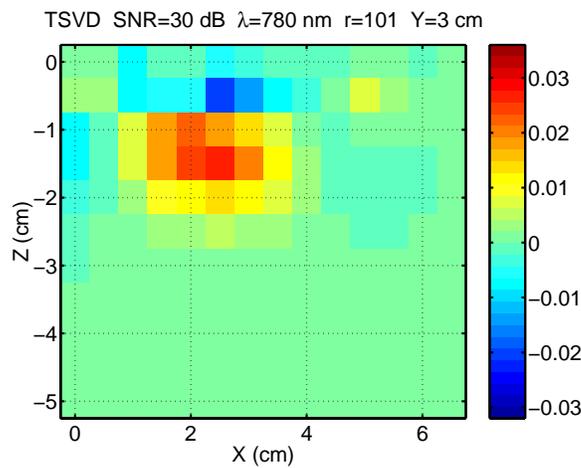


Figure 4.9: Separate source single wavelength TSVD reconstruction at 30 dB SNR from a uniform noise model.

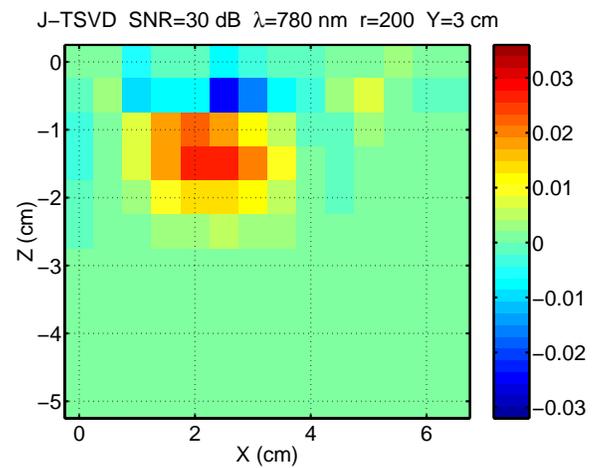


Figure 4.10: Separate source two-wavelength J-TSVD reconstruction at 30 dB SNR from a uniform noise model.

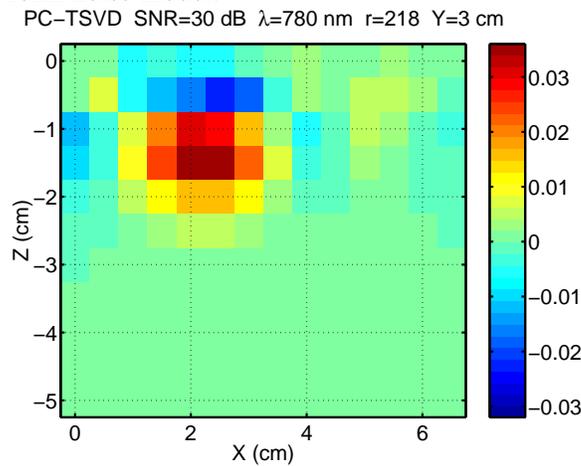


Figure 4.11: Separate source two-wavelength PC-TSVD reconstruction at 30 dB SNR from a uniform noise model.

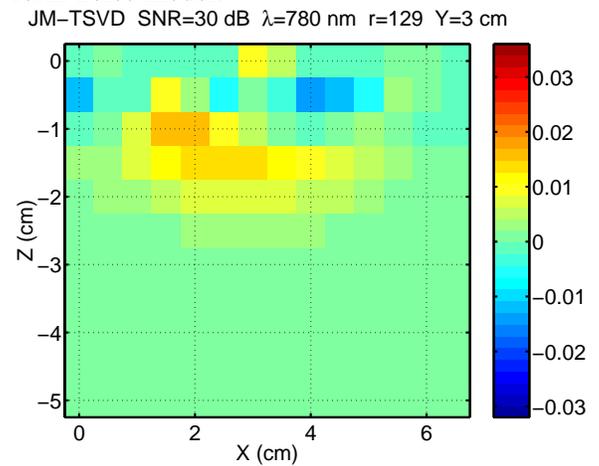


Figure 4.12: Separate source two-wavelength JM-TSVD reconstruction at 30 dB SNR from a uniform noise model.

#### 4.4.2 Quantitative results

To develop a more quantitative comparison of the two wavelength reconstruction algorithms we have examined the MSE and OCE measures while varying the SNR over a 10 dB to 60 dB range as we did in previous chapter. Our initial examination was of the collocated source geometry with the uniform noise model. The MSE measure for this scenario is shown in Figure 4.13. The performance of the JM-TSVD algorithm shows some improvement with respect to the other techniques, as we observed in the qualitative examples. The PC-TSVD algorithm also provides some improvement compared to the two algorithms that do not apply the CACR constraint.

We then changed the noise model to the simplified shot noise model. The MSE performance of the four reconstruction algorithms for this case is shown in Figure 4.14. Under this scenario the JM-TSVD algorithm did not perform as well as any of the other approaches but the PC-TSVD did provide some improvement compared to the non CACR constrained algorithms. Note also that under this noise model MSE at a given SNR is less than the uniform noise model. When we calculated the MSE of the whitened (see Section 3.1.5) problem, the SNR increased by about 27 dB.

Figures 4.15 and 4.16 show the MSE performance for the separate source geometry with the uniform and the simplified shot noise model respectively. The uniform noise model results agree with the qualitative observations (from Figures 4.9 through 4.12) in that the PC-TSVD has a slight advantage over the other algorithms. With the simplified shot noise model the PC-TSVD algorithm also outperforms the other algorithms by a small amount when using the separate source geometry.

The OCE measure was also examined for the same scenarios described in previous section. The

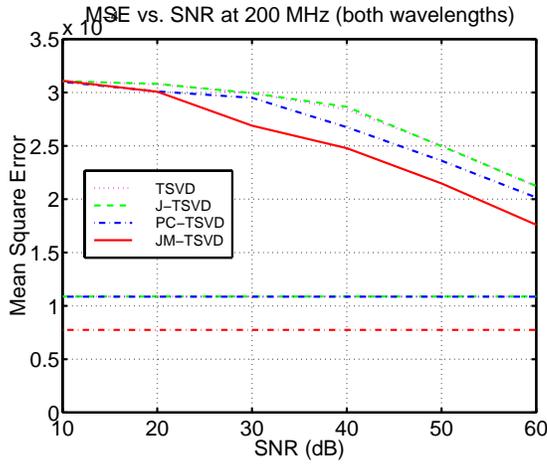


Figure 4.13: Mean square error vs. SNR for the collocated source geometry using the uniform noise model. The curve for the single wavelength TSVD reconstruction lies under the J-TSVD curve for most of the graph. The horizontal lines show the MSE measure when no noise was added to the data, the top horizontal line is same for the TSVD, J-TSVD, and PC-TSVD algorithms.

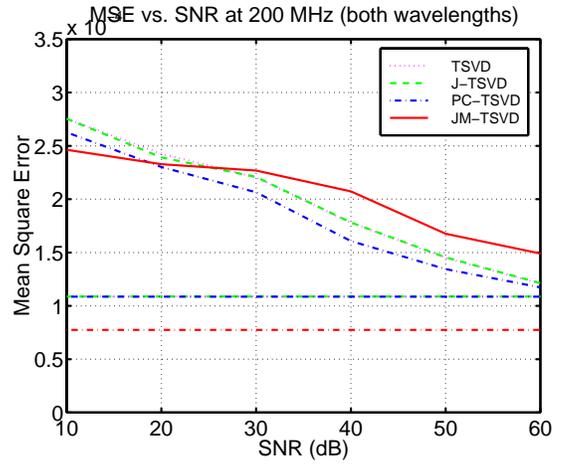


Figure 4.14: Mean square error vs. SNR for the collocated source geometry using the simplified shot noise model. The curve for the single wavelength TSVD reconstruction lies under the J-TSVD curve for most of the graph. The horizontal lines show the MSE measure when no noise was added to the data.

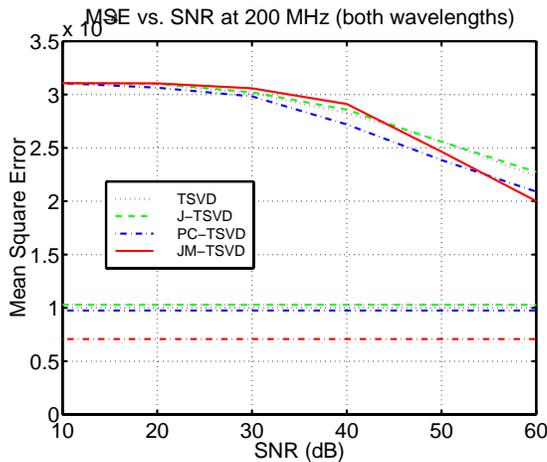


Figure 4.15: Mean square error vs. SNR for the separate source geometry using the uniform noise model. The horizontal lines show the MSE measure when no noise was added to the data.

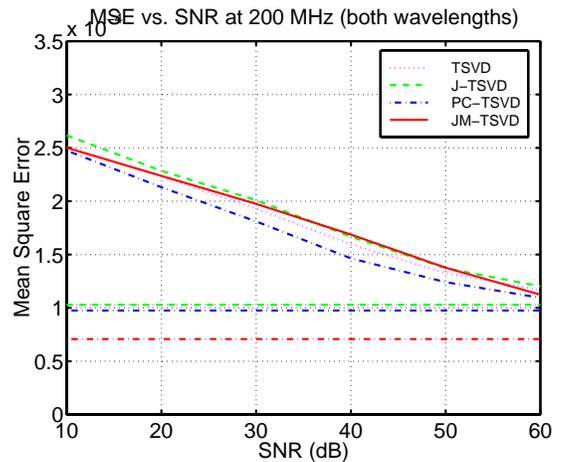


Figure 4.16: Mean square error vs. SNR for the separate source geometry using the simplified shot noise model. The horizontal lines show the MSE measure when no noise was added to the data.

OCE measure results for the collocated source geometry with the uniform noise model are shown in Figure 4.17, as with the MSE measure the JM-TSVD algorithm performs best for this scenario. Observations of the reconstructions at 10 dB SNR showed that the other algorithms did not generate a clearly identifiable anomaly and thus the OCE measure was detecting a noise response.

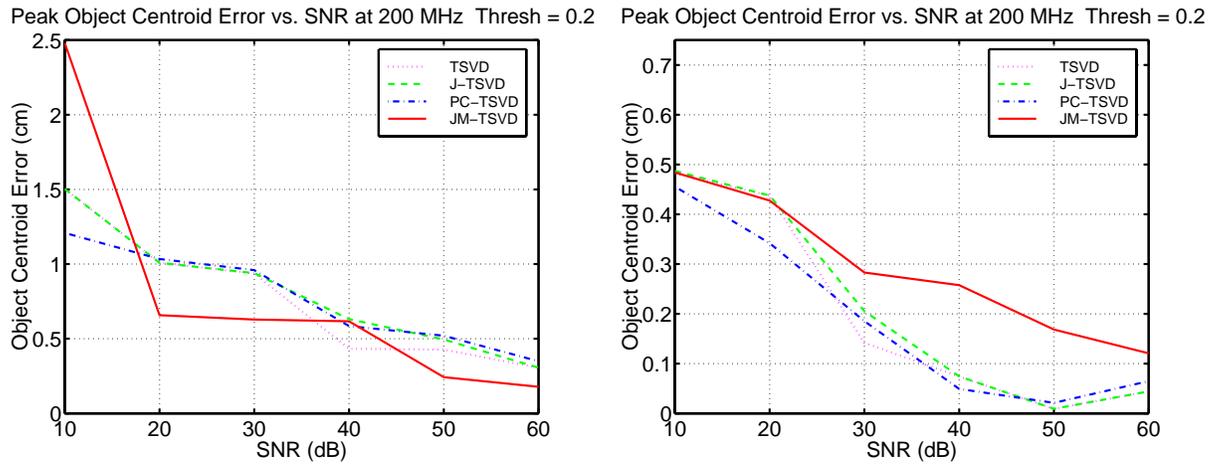


Figure 4.17: OCE vs. SNR for the collocated source geometry using the uniform noise model.

Figure 4.18: Peak object centroid error vs. SNR for the collocated source geometry using the simplified shot noise model.

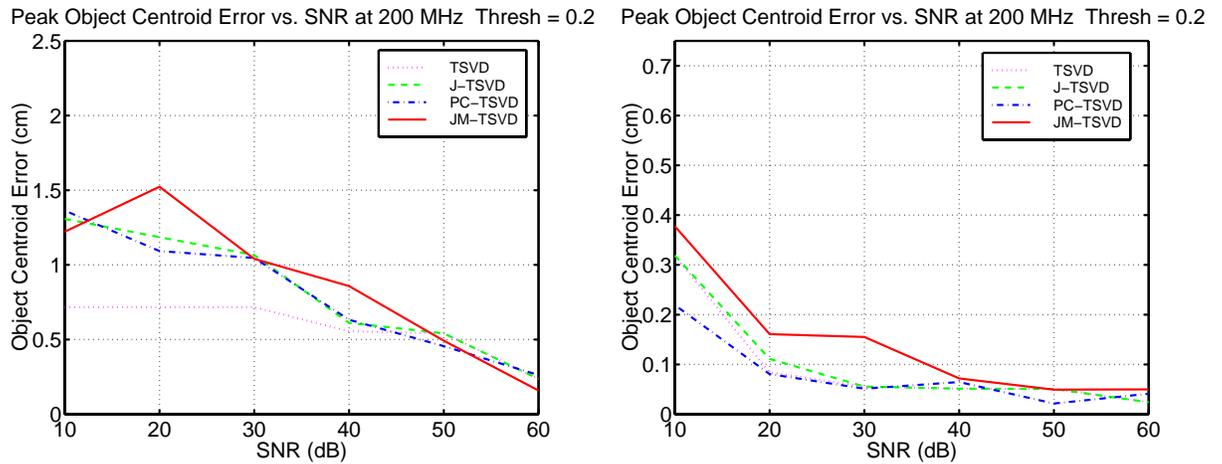


Figure 4.19: Peak object centroid error vs. SNR for the separate source geometry using the uniform noise model.

Figure 4.20: Peak object centroid error vs. SNR for the separate source geometry using the simplified shot noise model.

The OCE measure for the collocated source geometry with the simplified shot noise model is shown in Figure 4.18. This measure shows similar behavior as the MSE measure, in that the PC-TSVD algorithm performs the best over the range of SNRs and the JM-TSVD algorithm performs the worst in this case.

The uniform noise model and simplified shot noise model OCE measures for the separate source model are shown in Figures 4.19 and 4.20 respectively. These also show similar traits as the MSE measure. With the uniform noise model there is very little difference in OCE performance between the algorithms. Under the simplified shot noise model the PC-TSVD algorithm shows a slight advantage over the other algorithms.

## 4.5 Two wavelength discussion

In our simulations we have observed that the CACR constraint can provide some improvement in the reconstruction fidelity of the absorption coefficient perturbation function. But the amount of improvement does seem to be highly dependent upon the actual measurement scenario and environment.

It is also a somewhat surprising result that the separate source geometry does not produce better results. Given the improvement in the singular value spectrum compared to the collocated source geometry we had expected to see reconstruction with better amplitude accuracy and better depth visibility. This does indicate that there is a significant need to better study the relationship between the source detector geometry and the limits of the reconstruction quality.

limits of the reconstruction quality.

## Chapter 5

# Admissible solution DOT

In the previous chapter we demonstrated that the incorporation of a particular *a priori* constraint on the reconstruction improves the fidelity of the reconstruction under certain circumstances. This raises the question as to whether other constraints might improve the reconstruction fidelity under other scenarios, or if other constraints can further improve the reconstruction fidelity compared to the CACR constraint. By incorporating other constraints we might be able to narrow the large range of possible solutions that results from the underdetermined nature of the problem, thus improving the reconstruction fidelity. We might also be able to stabilize further the ill-posed nature of the problem, thus improving the reconstruction reliability. Among the undesirable features of the reconstructions presented in the previous two chapters are the large amount of reconstruction noise near the surface and negative reconstructed values for the absorption coefficient, a physical impossibility. In this chapter, by incorporating total variation and a min/max constraint on the solution we try to suppress some of these undesirable features and improve the fidelity of the reconstructions.

Other constraints examined were the 2-norm of the Laplacian of the reconstruction, which is commonly used in ill-posed reconstruction algorithms, and a normalized version of the total variation constraint.

## 5.1 The admissible solution approach

In our work we apply what we call an admissible solution approach to incorporate additional constraints into the inversion process. With the admissible solution approach we start with any number of *a priori* constraints on the solution. Each of these constraints defines a particular region in the solution space; the intersection of all of these constraint regions defines our admissible solution region. With these constraints in mind we search the solution space, using some appropriate algorithm, until a point in the admissible solution region is found. This point is then considered the solution to the problem. The admissible solution approach differs from a minimization approach in that it does not try to minimize a particular constraint or combination of the constraints. To illustrate this idea of an admissible region we show a Venn diagram in Figure 5.1 of three hypothetical constraints and the resulting admissible region. The horizontally oriented ellipse illustrates the admissible region that might be implied by a weighted 2-norm constraint, the square region illustrates a  $\infty$ -norm constraint, and the diagonally oriented ellipse illustrates a residual 2-norm constraint. The admissible region in this example is the somewhat triangular region where all three constraints intersect.

The acceptable range for each constraint is typically decided upon *a priori* from some knowledge of the problem or observation of the data. For example, we used an estimate of the variance

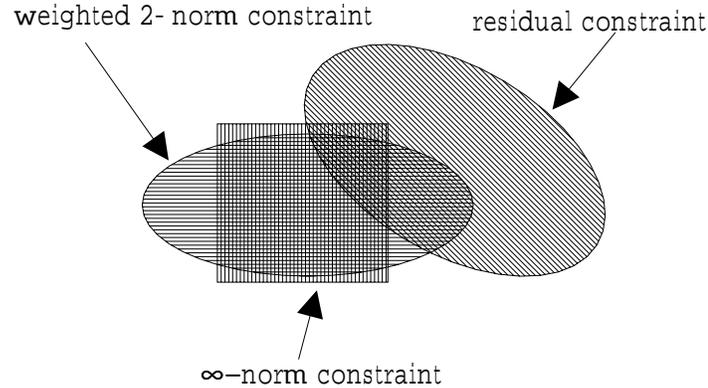


Figure 5.1: Admissible solution approach illustration. Three constraint regions are shown, the square region represents a  $\infty$ -norm constraint, the horizontal ellipse represents a weighted 2-norm constraint, and the rotated ellipse represents and residual 2-norm constraint. The admissible solution region is the intersection of all three constraints.

of the data to set the constraint values for the 2-norm of the residual. This differs from the methods presented in previous chapters where the parameters controlling the reconstruction method are typically chosen through an *a posteriori* method, such as the L-curve.

The admissible solution approach itself does not specify a specific algorithm to find the admissible region. In fact, several different algorithms have applied to solve the admissible solution approach. Projections Onto Convex Sets (POCS) has been applied in the image deblurring context [59, 60]. Other researchers in our laboratory have used the Ellipsoid Algorithm to find an admissible solution to the inverse electrocardiography problem [51]. We have chosen to apply the Ellipsoid Algorithm to the DOT imaging problem because it can be implemented without knowledge of the exact projectors for the constraints.

## 5.2 The Ellipsoid Algorithm

The Ellipsoid Algorithm is a geometrically based optimization algorithm that is applicable when the constraints are convex functions [49, 50]. This is a more relaxed requirement than gradient descent type algorithms where the gradient of the constraint function must exist or complicated modifications need to be employed. The Ellipsoid Algorithm is particularly applicable to our problem since all of the constraint functions we were interested in examining were convex, but not all differentiable (*e.g.* a total variation constraint). The underlying idea behind the Ellipsoid Algorithm is a generalization of the gradient known as the subgradient. Below we describe the concept of a subgradient and then how it used to developed the Ellipsoid Algorithm.

### 5.2.1 Subgradients

The concept of a subgradient is very similar to that of a gradient. In fact, where a function is differentiable the subgradient is exactly the gradient. The subgradient of a convex function  $f(\cdot)$  at  $\mathbf{x}$  is defined as any vector  $\mathbf{g}$  such that

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{g}^T(\mathbf{y} - \mathbf{x}). \quad (5.1)$$

To demonstrate this concept a simple one-dimensional non-differentiable convex function is shown in Figure 5.2. At  $\mathbf{x}_1$  the subgradient, represented by the thin tangent line  $g(x_1)$  is exactly the gradient. At  $\mathbf{x}_2$  the gradient does not exist but the whole family of tangent lines from the gradient of  $f(\mathbf{x}_{2-})$  to the gradient of  $f(\mathbf{x}_{2+})$  satisfy Eq. (5.1) and are thus elements of the subgradient of  $f(\mathbf{x}_2)$ . Two elements of the subgradient  $f(\mathbf{x}_2)$  are shown as the tangent lines  $g_1(x_2)$  and  $g_2(x_2)$ .

We can use the properties of the subgradient, specifically that we are working with a convex

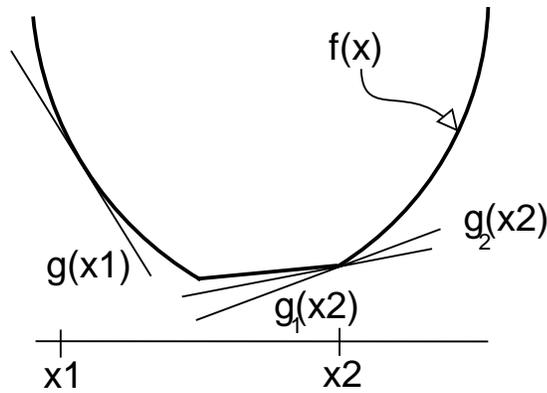


Figure 5.2: One-dimensional subgradient illustration. The thick line  $f(x)$  shows a convex function. The thin lines show the value (or some values) of the subgradient at those specific points by their slopes. At  $x_1$  there exists only one value for the subgradient signified by the tangent line  $g(x_1)$ . At  $x_2$  the subgradient is a continuous range of values from the slope of  $f(x_{2-})$  to the slope of  $f(x_{2+})$ . Two values of this range are illustrated by the slopes of the lines  $g_1(x_2)$  and  $g_2(x_2)$ .

function and the relationship described in the definition Eq. (5.1), to create a search algorithm for the minimum of the convex function. To see this consider the function shown in Figure 5.2. Examining the subgradient at  $\mathbf{x}_1$  tells us that  $f(\mathbf{x} < \mathbf{x}_1)$  is greater than  $f(\mathbf{x}_1)$  and examining the subgradient at  $\mathbf{x}_2$  tells us that  $f(\mathbf{x} > \mathbf{x}_2)$  is greater than  $f(\mathbf{x}_2)$ . Knowing that the minimizer lies somewhere between  $\mathbf{x}_1$  and  $\mathbf{x}_2$  we can then test the subgradient at some point in this region, say half way between  $\mathbf{x}_1$  and  $\mathbf{x}_2$ . This will tell us in which half of  $[\mathbf{x}_1 \ \mathbf{x}_2]$  the minimizer lies. This process can be iterated until we are sufficiently close to the minimum point of  $f(\mathbf{x})$  (or in our case where we are looking for an admissible solution we stop iterating once  $f(\mathbf{x})$  is admissible).

### 5.2.2 Extending the search algorithm to multiple dimensions

Consider the problem where we have an arbitrary number of  $l$  convex constraints on  $\mathbf{x} \in \mathcal{R}^n$  that must be met, written as

$$\begin{aligned} \text{find } \mathbf{x} \text{ such that: } & f_0(\mathbf{x}) < c_0 \\ & f_1(\mathbf{x}) < c_1 \\ & \vdots \\ & f_l(\mathbf{x}) < c_l, \end{aligned} \tag{5.2}$$

where  $f_i(\mathbf{x})$  is a convex function on  $\mathcal{R}^n$ . The algorithm is initialized with an ellipsoid in which some or all of the admissible set is known to exist. An ellipsoid in an  $n$ -dimensional space is defined by the equation

$$\mathcal{E}_o(\mathbf{x}_o, \mathbf{E}_o) = \left\{ \mathbf{z} \mid (\mathbf{z} - \mathbf{x}_o)^T \mathbf{E}_o^{-1} (\mathbf{z} - \mathbf{x}_o) \leq 1 \right\} \tag{5.3}$$

where  $\mathbf{x}_o$  is the centroid of the ellipsoid and  $\mathbf{E}_o$  is a real, symmetric positive definite matrix. The square roots of the eigenvalues of  $\mathbf{E}_o$  are the axis radii and the eigenvectors describe the axis directions. Each of the constraint functions is evaluated at the centroid of the ellipsoid to see if the respective constraints are met. If any of the constraints is not met, one of the unmet constraints is selected and the subgradient of that constraint function is evaluated at the center of the ellipsoid. Eq. (5.1) tells us that a half of the current ellipsoid can be ruled out because the constraint function is greater or equal to the value at the center which does not meet the constraint. Specifically, if we draw a hyperplane through the centroid that is normal to the subgradient, elements of the ellipsoid on the side of the hyperplane to which the subgradient vector points are not admissible. Using this

information we define a new smaller volume ellipsoid that contains the half of the current ellipsoid in which the admissible region must exist. The center of the new ellipsoid is given by [49, 50]

$$\begin{aligned}\tilde{\mathbf{g}}_i &= \frac{1}{\sqrt{\mathbf{g}_i^T \mathbf{E}_i \mathbf{g}_i}} \mathbf{g}_i \\ \mathbf{x}_{i+1} &= \mathbf{x}_i - \mathbf{E}_i \tilde{\mathbf{g}}_i,\end{aligned}\tag{5.4}$$

where  $\mathbf{g}_i$  is the subgradient of constraint at the center of the  $i^{\text{th}}$  ellipsoid. The new ellipsoid matrix is given by [49, 50]

$$\mathbf{E}_{i+1} = \frac{n^2}{n^2 - 1} \left( \mathbf{E}_i - \frac{2}{n+1} \mathbf{E}_i \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^T \mathbf{E}_i \right).\tag{5.5}$$

All of the constraint functions are re-evaluated at the center of the new ellipsoid. If any constraints are not met, the process of selecting an unmet constraint, calculating the subgradient, and generating a new ellipse is re-iterated. Otherwise the center of final ellipsoid meets all constraints and we have a solution to the problem in Eq. (5.2).

The sequence of ellipsoids is decreasing in volume at each step thus we are guaranteed to converge to a solution if an admissible set exists within the initial ellipsoid. At each step the volume of the new ellipsoid is related to the previous ellipsoid by the relationship [49, 50]

$$\text{vol}(E_{i+1}) = \left( \frac{n}{n+1} \right)^{\frac{n+1}{2}} \left( \frac{n}{n-1} \right)^{\frac{n-1}{2}} \text{vol}(E_i),\tag{5.6}$$

where  $\text{vol}(\cdot)$  represents the volume of the ellipsoid defined by  $\mathbf{E}$ .

Also, at each iteration we can test to see if a particular constraint is no longer feasible given the current ellipsoid. If the solution becomes infeasible this implies that there was no intersection of the admissible regions. Specifically if the constraint is in the form of one of Eq. 5.2 then if

$$f_i(\mathbf{x}_i) - \sqrt{\mathbf{g}_i^T \mathbf{E}_i \mathbf{g}_i} > c_i\tag{5.7}$$

the constraint is infeasible.  $\sqrt{\mathbf{g}_i^T \mathbf{E}_i \mathbf{g}_i}$  bounds the maximum change in value for the constraint  $f_i(\mathbf{x})$  at the furthest point in the current ellipsoid [50]. If this quantity is not large enough to reduce the current value of the constraint to less than that of the constraint bound the current ellipsoid does not intersect with the admissible region defined by  $f_i(\mathbf{x}) < c_i$ .

Since the Ellipsoid Algorithm is geometrically based, an instructive way to obtain an idea of how the algorithm performs is through a simple two-dimensional example. We consider two constraints on the results, a residual 2-norm constraint on a linear system relationship and an infinity norm constraint that bounds both elements of the solution to the interval [-1 1]. The linear system relationship is given by

$$\mathbf{G}\mathbf{x} = \mathbf{b} \quad \mathbf{G} = \begin{bmatrix} 1 & 1.5 \\ 1 & 0.5 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 2.25 \\ 1.35 \end{bmatrix}. \quad (5.8)$$

The true solution for this equation is  $\mathbf{x} = [0.9 \ 0.9]^T$ . Now if we know the measured data  $\mathbf{b}$  is perturbed by a vector with a 2-norm less than one we can use these constraints to find an admissible solution. The constraints are expressed mathematically as

$$\|\mathbf{G}\mathbf{x}_i - \mathbf{b}\|_2 < 1 \quad \text{and} \quad \|\mathbf{x}\|_\infty < 1 \quad (5.9)$$

These constraints are illustrated in Figure 5.3. The thick lined green square shows the infinity norm constraint and the magenta ellipse with the dash-dot pattern shows the residual constraint mapped to the solution space.

An initial ellipse with axis radii of 5 for both dimensions (a circle) was centered at (0,0). This is illustrated by the dashed red circle in the first iteration plot in Figure 5.3. We start by examining which constraints are not met by the center of the initial ellipsoid. In our graphical example we can

see that the center of the initial ellipsoid meets the infinity norm constraint but does not meet the residual constraint. The subgradient of the residual constraint is illustrated by the thick red dashed line coming from the center of the initial ellipse in the first iteration plot.

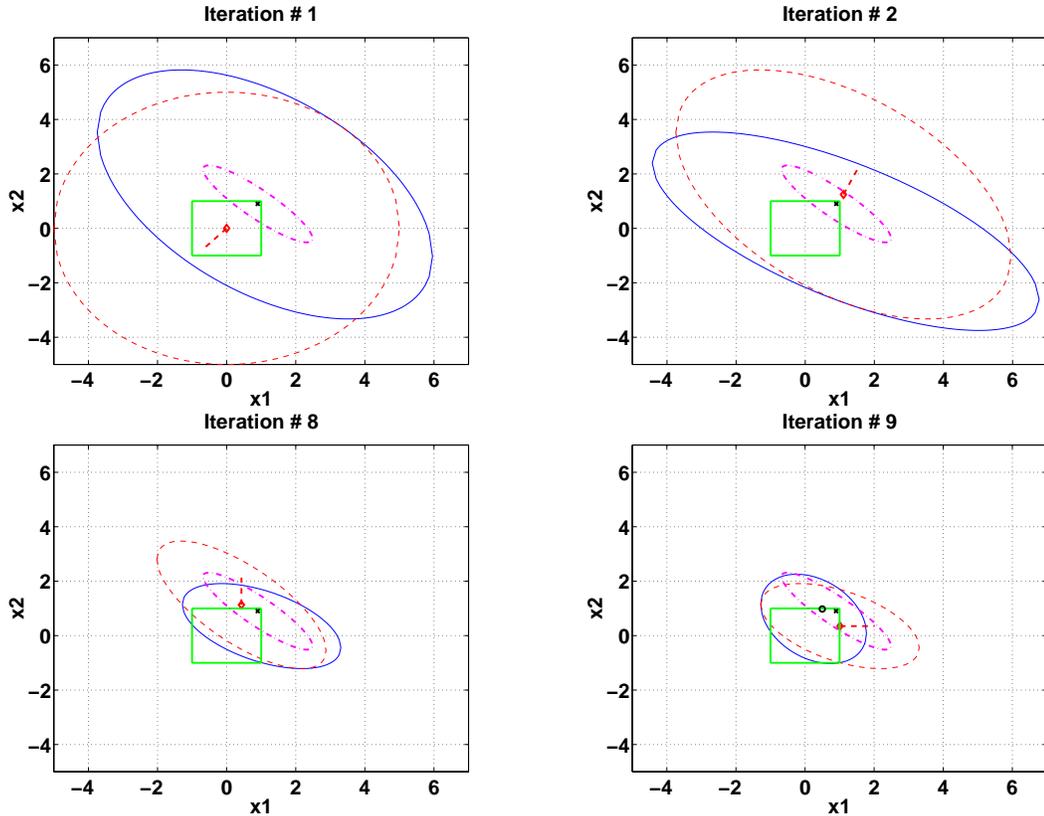


Figure 5.3: Two-dimensional example of the Ellipsoid Algorithm used to solve an admissible solution problem. One constraint is a maximum and minimum constraint of 1 and -1, represented by the green square. The other constraint is a 2-norm constraint on the residual with a value of 1, represented by the thick magenta ellipse with a dash-dot line pattern. The admissible region is the intersection of these two. The first (1), second (2), second to last (8), and last (9) iterations are shown. The dashed red ellipse shows the current ellipsoid region and dashed red line shows subgradient of the selected constraint function at the center of this ellipse. The thin lined blue ellipse shows the next ellipsoid resulting from evaluating the subgradient at the center of dashed red ellipse. The true value of the vector is shown by the small black x, the first admissible solution is shown by the small black circle in the lower right figure.

The new ellipse and centroid are then calculated using Eq. (5.4) and Eq. (5.5). This new ellipse

is shown with the thin solid blue line in first iteration. For the second iteration this now becomes the old ellipse, identified by the dashed red ellipse in the plot labeled Iteration #2. The center of this ellipse, identified by the small red diamond lies within the residual constraint but outside of the infinity-norm constraint. Thus in this iteration the subgradient for the infinity-norm is calculated at the new centroid (shown again as a thick dashed red line) and a new ellipse calculated (thin lined blue ellipse). Skipping ahead to the ninth (and last) iteration illustrated in the bottom right of Figure 5.3 we see that the center of the final ellipse lies in the intersection of both of the constraints and thus we have found a solution.

### 5.2.3 Derivation of subgradients

One topic that we have avoided until this point is the calculation of the subgradients of the constraint functions we were interested in evaluating. In this section we derive the subgradient functions for all of the constraints we have explored. We note that we use the standard notation for the partial derivative  $\partial/\partial\mathbf{x}$  to represent the subgradient, following the notation of Boyd and Barratt [50].

#### Residual 2-norm subgradient

The subgradient 2-norm of the residual is exactly the gradient of the 2-norm of the residual function since the 2-norm of the residual is differentiable. To simplify the numerical calculation of the subgradient for the 2-norm of the residual we actually calculate the subgradient of the square of the 2-norm of the residual. We start first by writing the partial derivative of the squared 2-norm of the residual with respect to the vector  $\mathbf{x}$

$$\frac{\partial}{\partial\mathbf{x}}\|\mathbf{G}\mathbf{x} - \mathbf{b}\|_2^2 = \frac{\partial}{\partial\mathbf{x}}\|\mathbf{G}_1x_1 + \mathbf{G}_2x_2 + \dots + \mathbf{G}_nx_n + -\mathbf{b}\|_2^2 \quad (5.10)$$

where  $\mathbf{G}_i$  is the  $i^{\text{th}}$  column of the system matrix  $\mathbf{G}$  and  $x_n$  is the  $i^{\text{th}}$  element of  $\mathbf{x}$ . Looking at the derivative with respect to  $x_i$  and rewriting the squared 2-norm gives

$$\frac{\partial}{\partial x_i} \|\mathbf{G}\mathbf{x} - \mathbf{b}\|_2^2 = \frac{\partial}{\partial x_i} (\mathbf{G}_1 x_1 + \mathbf{G}_1 x_1 + \dots - \mathbf{b})^T (\mathbf{G}_1 x_1 + \mathbf{G}_1 x_1 + \dots - \mathbf{b}). \quad (5.11)$$

For the  $i^{\text{th}}$  element only terms containing  $x_i$  are non-zero giving

$$\frac{\partial}{\partial x_i} \|\mathbf{G}\mathbf{x} - \mathbf{b}\|_2^2 = \frac{\partial}{\partial x_i} \left\{ 2\mathbf{G}^T_i \mathbf{G}_1 x_1 x_i + \dots + \mathbf{G}^T_i \mathbf{G}_i x_i^2 + \dots + 2\mathbf{G}^T_i \mathbf{G}_n x_n x_i - 2\mathbf{G}^T_i \mathbf{b} x_i \right\}. \quad (5.12)$$

Evaluating the derivative gives

$$\frac{\partial}{\partial x_i} \|\mathbf{G}\mathbf{x} - \mathbf{b}\|_2^2 = 2\mathbf{G}^T_i \mathbf{G}_1 x_1 + \dots + 2\mathbf{G}^T_i \mathbf{G}_i x_i + \dots + 2\mathbf{G}^T_i \mathbf{G}_n x_n - 2\mathbf{G}^T_i \mathbf{b}. \quad (5.13)$$

Thus the  $i^{\text{th}}$  element of Eq. (5.10) is

$$\frac{\partial}{\partial x_i} \|\mathbf{G}\mathbf{x} - \mathbf{b}\|_2^2 = 2\mathbf{G}^T_i (\mathbf{G}\mathbf{x} - \mathbf{b}). \quad (5.14)$$

The complete (sub)gradient vector is then given by

$$\frac{\partial}{\partial \mathbf{x}} \|\mathbf{G}\mathbf{x} - \mathbf{b}\|_2^2 = 2\mathbf{G}^T (\mathbf{G}\mathbf{x} - \mathbf{b}). \quad (5.15)$$

### Laplacian 2-norm subgradient

The 2-norm of the Laplacian is a commonly employed measure in ill-posed problems. The Laplacian for a discrete function represented as a vector is implemented as a linear operator on  $\mathbf{x}$

$$\nabla^2 \mathbf{x} = \mathbf{L}\mathbf{x}, \quad (5.16)$$

where  $\mathbf{L}$  implements the finite difference operator to the requested order, summing all dimensions for each element. Thus, in general, we want to calculate the subgradient of an arbitrary linear

operator on  $\mathbf{x}$

$$\frac{\partial}{\partial \mathbf{x}} \|\mathbf{L}\mathbf{x}\|_2^2. \quad (5.17)$$

This is just a special case of the residual subgradient describe above and is given by

$$\frac{\partial}{\partial \mathbf{x}} \|\mathbf{L}\mathbf{x}\|_2^2 = 2\mathbf{L}^T \mathbf{L}\mathbf{x}. \quad (5.18)$$

### Total variation subgradient

The total variation for a continuous function  $x(\mathbf{z})$  is commonly defined as [61]

$$tv(x(\mathbf{z})) = \int |\nabla x(\mathbf{z})| d\mathbf{z}, \quad (5.19)$$

where we have used  $\mathbf{z}$  as the independent variable so as to not confuse it with the spatial variable with function we are trying to reconstruct  $x(\mathbf{z})$ . Since we are dealing with a discrete function of  $x(\mathbf{z})$  one approximation of the total variation measure is given by

$$tv(\mathbf{x}) = \sum_{j=1}^n |\mathbf{L}_j \mathbf{x}| \quad (5.20)$$

where  $\mathbf{L}_j$  is the gradient approximation around the  $j^{th}$  element of  $\mathbf{x}$ . Combining this into a matrix vector product yields

$$tv(\mathbf{x}) = \|\mathbf{L}\mathbf{x}\|_1. \quad (5.21)$$

In deriving the subgradient of the total variation measure we are actually computing the subgradient of an arbitrarily weighted 1-norm

$$\frac{\partial}{\partial \mathbf{x}} tv(\mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} \|\mathbf{L}\mathbf{x}\|_1. \quad (5.22)$$

Looking at the  $i^{th}$  differential element of Eq. (5.22) we have

$$\frac{\partial}{\partial x_i} \left\{ \sum_{j=1}^m |\mathbf{L}_j \mathbf{x}| \right\} = \frac{\partial}{\partial x_i} \{ |\mathbf{L}_1 \mathbf{x}| + \dots + |\mathbf{L}_m \mathbf{x}| \} \quad (5.23)$$

where  $\mathbf{L}_j$  is the  $j^{th}$  row. Now consider the  $k^{th}$  term of the sum

$$\frac{\partial}{\partial x_i} |\mathbf{L}_k \mathbf{x}| = \begin{cases} \mathbf{L}_{k,i} & \mathbf{L}_k \mathbf{x} > 0 \\ -\mathbf{L}_{k,i} & \mathbf{L}_k \mathbf{x} < 0, \end{cases} \quad (5.24)$$

where  $\mathbf{L}_{k,i}$  is the element at the  $k^{th}$  row and  $i^{th}$  column of  $\mathbf{L}$ . This can also be expressed as

$$\frac{\partial}{\partial x_i} |\mathbf{L}_k \mathbf{x}| = \mathbf{L}_{k,i} \text{sign}(\mathbf{L}_k \mathbf{x}). \quad (5.25)$$

For the complete  $i^{th}$  differential element we now have

$$\frac{\partial}{\partial x_i} |\mathbf{L} \mathbf{x}| = \mathbf{L}_{1,i} \text{sign}(\mathbf{L}_1 \mathbf{x}) + \mathbf{L}_{2,i} \text{sign}(\mathbf{L}_2 \mathbf{x}) + \dots + \mathbf{L}_{m,i} \text{sign}(\mathbf{L}_m \mathbf{x}). \quad (5.26)$$

Thus, the subgradient for the total variation measure is given by

$$\frac{\partial}{\partial \mathbf{x}} tv(\mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} \|\mathbf{L} \mathbf{x}\|_1 = \mathbf{L}^T \text{sign}(\mathbf{L} \mathbf{x}). \quad (5.27)$$

### Normalized total variation subgradient

In examining the performance of the total variation constraint function we have observed, as will be described later, that because some of the reconstructions have a very small amplitude the total variation constraint does not affect the results. Yet in these reconstructions there is still a significant amount of variation relative to the amplitude of the image. To account for this we have developed a normalized total variation measure and derived its associated subgradient. The expression for the normalized total variation is

$$ntv(\mathbf{x}) = \left\| \mathbf{L} \frac{\mathbf{x}}{\|\mathbf{x}\|_2} \right\|_1. \quad (5.28)$$

Taking the partial with respect to  $\mathbf{x}$  gives

$$\frac{\partial}{\partial \mathbf{x}} ntv(\mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} \sum_{j=1}^m \left| \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| \quad (5.29)$$

where  $\mathbf{L}_j$  is the  $j^{\text{th}}$  row of the difference operator describe in the previous section. Considering just the  $i^{\text{th}}$  differential element results in

$$\frac{\partial}{\partial x_i} \sum_{j=1}^m \left| \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| = \frac{\partial}{\partial x_i} \left\{ \left| \mathbf{L}_1 \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| + \left| \mathbf{L}_2 \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| + \dots + \left| \mathbf{L}_m \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| \right\} \quad (5.30)$$

From the  $i^{\text{th}}$  differential element we examine just the  $j^{\text{th}}$  term given by

$$\frac{\partial}{\partial x_i} \left| \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| = \frac{\partial}{\partial x_i} \left\{ \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right\} \text{sign} \left( \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right) \quad (5.31)$$

Applying the partial derivative to this expression gives

$$\frac{\partial}{\partial x_i} \left\{ \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right\} = \frac{-2x_i}{\|\mathbf{x}\|_2^4} \mathbf{L}_j \mathbf{x} + \frac{1}{\|\mathbf{x}\|_2^2} \mathbf{L}_{j,i}. \quad (5.32)$$

Thus the differential for the  $j^{\text{th}}$  term is

$$\frac{\partial}{\partial x_i} \left| \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| = \text{sign}(\mathbf{L}_j \mathbf{x}) \left( \frac{-2x_i}{\|\mathbf{x}\|_2^4} \mathbf{L}_j \mathbf{x} + \frac{1}{\|\mathbf{x}\|_2^2} \mathbf{L}_{j,i} \right). \quad (5.33)$$

Combining all of the terms over the  $j^{\text{th}}$  index gives the complete  $i^{\text{th}}$  differential

$$\frac{\partial}{\partial x_i} \sum_{j=1}^m \left| \mathbf{L}_j \frac{\mathbf{x}}{\|\mathbf{x}\|_2^2} \right| = \frac{-2x_i}{\|\mathbf{x}\|_2^4} \sum_{j=1}^m \mathbf{L}_j \mathbf{x} \text{sign}(\mathbf{L}_j \mathbf{x}) + \frac{1}{\|\mathbf{x}\|_2^2} \sum_{j=1}^m \mathbf{L}_{j,i} \text{sign}(\mathbf{L}_j \mathbf{x}). \quad (5.34)$$

Finally, the complete normalized total variation subgradient vector can be expressed as

$$\frac{\partial}{\partial \mathbf{x}} ntv(\mathbf{x}) = -2 \frac{\|\mathbf{L}\mathbf{x}\|_1}{\|\mathbf{x}\|_2^4} \mathbf{x} + \frac{1}{\|\mathbf{x}\|_2^2} \mathbf{L}^T \text{sign}(\mathbf{L}\mathbf{x}). \quad (5.35)$$

**Min/max subgradient**

The last constraint that we examined was a min/max constraint. The subgradient is found by taking the partial with respect to  $\mathbf{x}$  of the max function

$$\frac{\partial}{\partial \mathbf{x}} \max(\mathbf{x}) \quad (5.36)$$

If  $x_i$  is not the maximum element of the vector  $x_i$  there is no change with respect to  $x_i$ . If  $x_i$  is the largest element of the vector  $x_i$  then the derivative with respect to that element is one. This results in a binary vector identifying which element is the maximum. We express this mathematically as

$$\frac{\partial}{\partial x_i} \max(\mathbf{x}) = \begin{cases} 0 & \max(\mathbf{x}) \neq x_i \\ 1 & \max(\mathbf{x}) = x_i. \end{cases} \quad (5.37)$$

The same applies for the minimum function.

**5.3 Admissible solution simulation scenario**

To analyze the performance of the admissible solution approach we have implemented a computer simulation similar to those described in Chapters 3 and 4. The primary difference in these simulations is that only a single vertical 2-D slab was examined and the anomaly only existed within this slab. Also, unlike the slab experiments in Chapter 3, this slab was positioned vertically so that it extended into the medium. We oriented the slab in this manner to examine the reconstruction fidelity with respect to depth, which the previous simulations have shown to be the most problematic dimension. We chose to examine the vertical 2-1/2 dimensional simulation because of the expected decreased computational complexity of the Ellipsoid Algorithm compared to the linear algebraic approaches of the previous chapters.

A rectangular absorption anomaly was simulated centered at 2.25 cm deep and 0.75 cm to the right of the source-detector array center line. The width of the anomaly was the same as the width of the reconstruction slab, 0.25 cm. An image of the anomaly is shown in Figure 5.4.

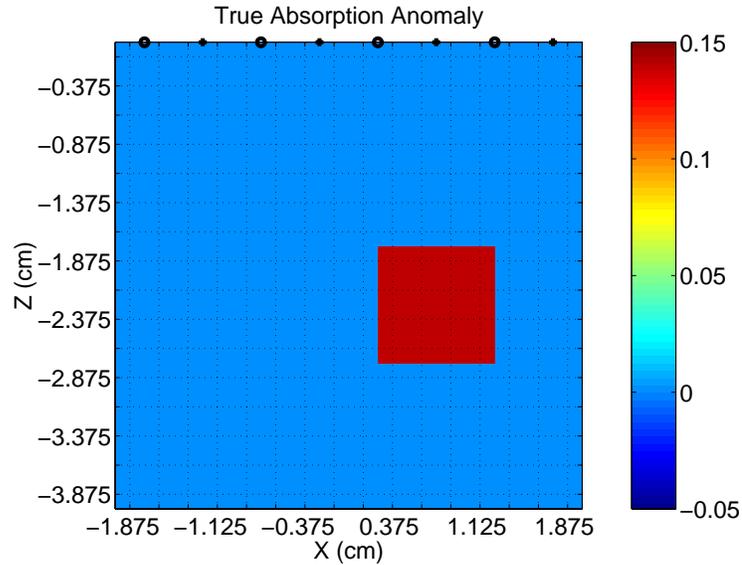


Figure 5.4: True absorption perturbation image. The absorption coefficient of the anomaly is 0.139 above the background

We simulated four sources and four detectors, all collinear along the air-tissue edge of the slab. The forward and inverse models were identical (Born-1) including the discretization of the slab. The discretization was on a 0.25 cm grid with 16 x 16 (256) voxels. The simulated modulation frequency was 200 MHz, resulting in 32 measurements and 256 unknowns. This is a similar measurements-to-unknowns-ratio to that used in our previous 3D experiments. The complete set of simulation parameters are presented in Table 5.1.

To get an idea of how much information is lost to the null space of the forward operator we calculated both the SVD of the forward operator, shown in Figure 5.5, and the minimum norm

<i>Forward Model Information</i>				
Index of refraction:		1.370		
Propagation velocity (cm/s):	2.189781e+10			
Mean cosine:		0.9		
Scattering coef. (cm <sup>-1</sup> ):		100		
Reduced scat. coef. (cm <sup>-1</sup> ):		10		
Absorption coef. (cm <sup>-1</sup> ):		0.041		
Wavelength (nm):		780		
Modulation frequency (MHz):		200		
Number of measurements:		16		
<i>Source optodes</i>				
Type:	uniform			
X positions (cm):	-1.75	-0.75	0.25	1.25
Y positions (cm):	0.00			
Z positions (cm):	0.00			
Normal vector:	0.00	0.00	-1.00	
Amplitude:	1.00			
Numerical aperture:	1.00			
<i>Detector optodes</i>				
Type:	uniform			
X positions (cm):	-1.25	-0.25	0.75	1.75
Y positions (cm):	0.00			
Z positions (cm):	0.00			
Normal vector:	0.00	0.00	-1.00	
Efficiency:	1.00			
Numerical aperture:	1.00			
Medium boundary geometry:	Semi-infinite			
Computational method:	Born	order:	1	
<i>Computational volume</i>				
Type :	uniform			
Min X: -1.88	Max X: +1.88	Step: +0.25	Count 16	
Min Y: +0.00	Max Y: +0.00	Step: +0.25	Count 1	
Min Z: -3.88	Max Z: -0.12	Step: +0.25	Count 16	
Total Voxels:	256			
<i>Inverse Model Information :</i> identical to forward model				
<i>Object Information</i>				
Type:	block			
Position:	0.750	0.000	-2.250	
Dims:	X=1.000	Y=0.500	Z=1.000	
$\mu_a$ :	0.180			
$\mu_s$ :	bkgnd			

Table 5.1: Admissible solution analysis simulation parameters.

reconstruction with no added noise, shown in Figure 5.6. The SVD shows a range of about 6 orders of magnitude, a little less than our three-dimensional simulation. The SVD of this forward operator also shows similar characteristics to the forward operator of the full three-dimensional simulation shown in Figure 3.2. Small breaks in the singular value spectrum and a gradual decay of the singular values were present in both forward models.

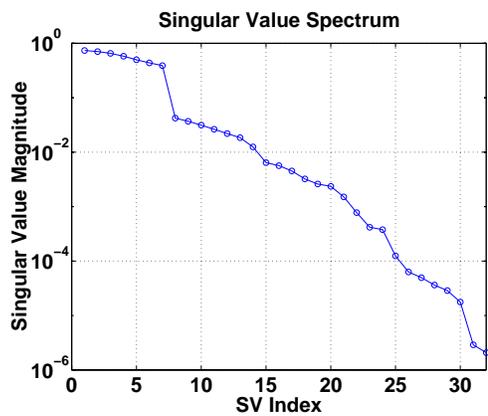


Figure 5.5: Singular value spectrum of the forward operator.

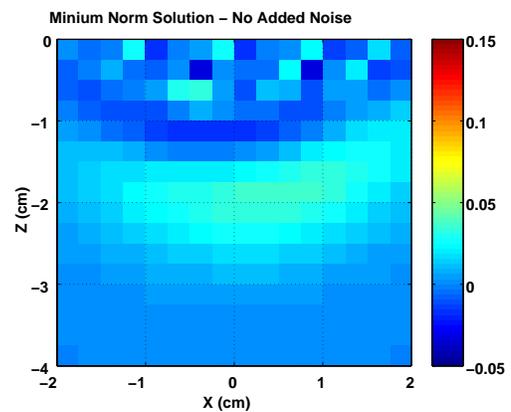


Figure 5.6: Minimum norm solution with no added noise to the measurements.

The minimum norm reconstruction shows the best possible (in the MSE sense) reconstruction using a linear model approach without constraints. With these simulations we are again underestimating the depth and absorption coefficient magnitude and also somewhat incorrectly estimating the X position of the anomaly using an unconstrained linear model reconstruction.

## 5.4 Admissible solution performance

We examined several aspects of the performance of the admissible solution approach applied to DOT. These included an investigation of which of the constraints produced the best reconstruction

fidelity, an examination of reconstruction sensitivity to specific noise vector instances, the effects of variation in absorption anomaly parameters, reconstruction of multiple absorption anomalies, and the effect of background variation on the reconstructions. All of these results are presented in the next several sections.

### 5.4.1 Constraint comparisons

We have described a number of possible constraints on the absorption coefficient reconstruction. In this section we wish to understand which of the constraints actually improves the reconstruction fidelity relative to the linear algebraic approaches, and how these constraints behave. We implemented admissible solution algorithms for all of the constraints described along with combinations of several of the more promising constraints. We applied these algorithms to simulated data that was modified with uniform variance additive Gaussian noise. The standard deviation of this noise was scaled to produce signal-to-noise ratios from 60 dB to 10 dB in 10 dB steps, relative to the peak scattered fluence rate. For each SNR we examined the mean square error of the absorption coefficient reconstruction to analyze the performance of the given constraint(s).

#### **Residual 2-norm constraint**

To examine the basic behavior of the admissible solution approach to the DOT problem we have performed reconstructions with just a residual 2-norm constraint. Using the residual 2-norm constrained admissible solution approach we expected similar results to the TSVD algorithm given that both algorithms are minimizing the 2-norm of the residual, the difference being that TSVD reduces the residual by incorporating smaller and smaller singular values and the ellipsoid follows a complex

path until an admissible solution is found.

We examined a number of values for the stopping criteria on the residual 2-norm constraint. For each SNR we estimated the absorption coefficient using the largest value of the residual constraint. We then used this solution along with the associated ellipsoid to initialize the Ellipsoid Algorithm for the next smaller value of the constraint. This process was iterated until we calculated an admissible solution for each residual constraint value or found that as some point the problem became infeasible as the constraint was reduced. The value of the residual constraint was calculated relative to the expected 2-norm of the noise vector and is expressed by

$$\|\mathbf{G}\hat{\mathbf{k}} - \mathbf{\Phi}^s\|_2 < [2 : -0.1 : 0.1]\sigma\sqrt{m-1}, \quad (5.38)$$

where  $\sigma$  is the standard deviation of the noise and  $m$  is the number of measurements. The term  $\sigma\sqrt{m-1}$  is the expected 2-norm of the noise vector.

Figures 5.7 and 5.8 present a qualitative comparison of the TSVD and residual constraint admissible solution reconstruction. From these figures it is evident that qualitative performance of the residual 2-norm constrained admissible solution approach is very similar to that of the TSVD algorithm.

The mean square error performance versus SNR is shown in Figure 5.9. As would be expected the performance of these two algorithms is similar since they are both trying to minimize the same quantity.

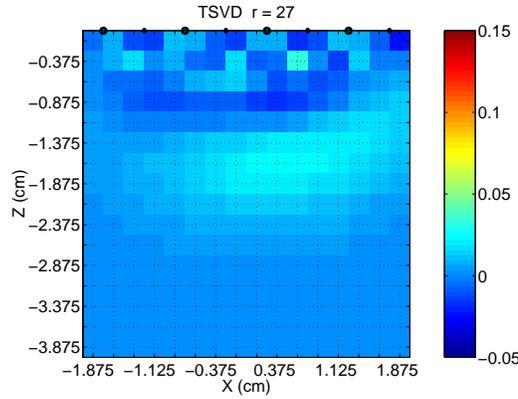


Figure 5.7: TSVD reconstruction for 30 dB SNR with  $r = 27$ . The position of the sources and detectors are shown as circles and crosses at the top of the plot.

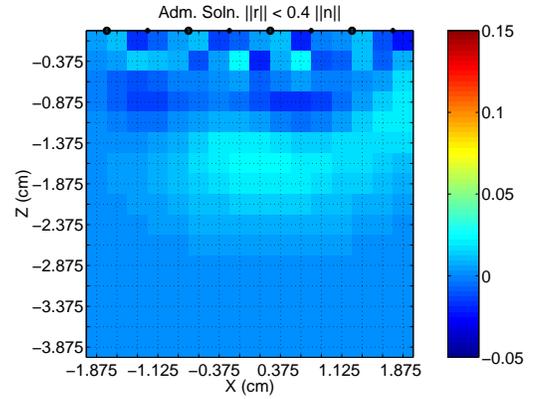


Figure 5.8: Admissible solution reconstruction with the residual constraint  $\|r\| < 0.4\|n\|$  at a SNR of 30 dB.

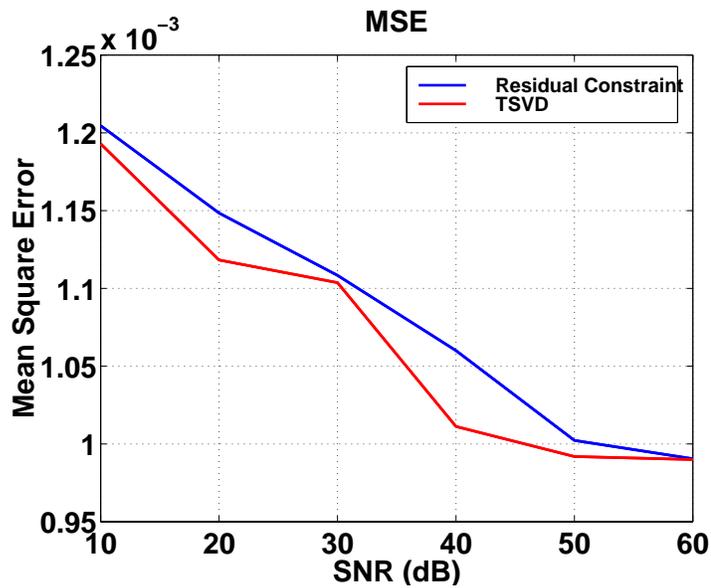


Figure 5.9: Mean Square Error performance for the TSVD and residual constraint admissible solution algorithms.

**Residual and total variation constraints**

The first additional constraint on the solution we explored was a total variation constraint. This was implemented as the sum of the absolute value of the X and Y derivatives. The value of the total variation for the true absorption anomaly perturbation was 2.240. In addition to the range of residual constraints described above, we examined a number of TV constraints. These included TV values of 2.75, 2.5, 2.25, 2.0, and 1.75. An example reconstruction at 30 dB SNR using the residual and total variation admissible solution approach is shown in Figure 5.10. This reconstruction had the lowest MSE of constraint pairs examined. The residual constraint was  $\|r\| < 0.3\|n\|$  and the total variation constraint was  $TV(x) < 1.75$ . The MSE for this reconstruction was 1.0336e-3 while the residual constraint alone and TSVD reconstructions were  $> 1.1e-3$ , as is evident from the MSE curve in the previous section.

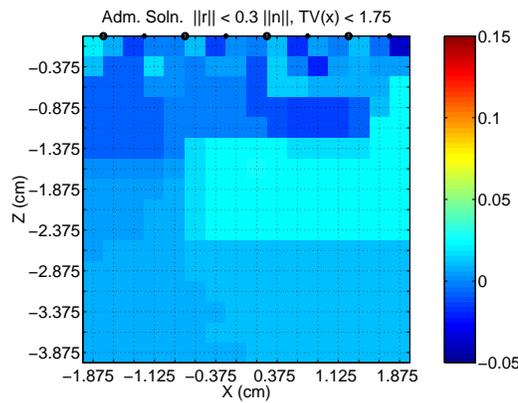


Figure 5.10: Admissible solution reconstruction with the residual and constraint  $\|r\| < 0.3\|n\|, tv(x) < 1.75$ . The MSE for this reconstruction was 1.0336e-3.

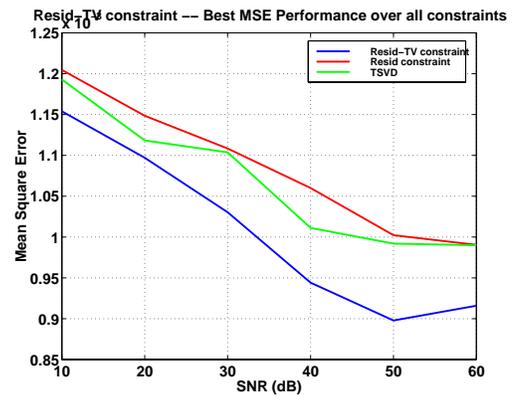


Figure 5.11: MSE versus SNR for the residual-TV constrained admissible solution approach. Also shown is the residual constrained admissible solution results described in the previous section and TSVD for comparison.

The MSE performance of the residual and TV constrained admissible solution algorithm is

shown in Figure 5.11. The MSE values at each SNR are the best MSE performance over the whole range of residual and TV constraint values described above. Although the residual and TV constrained admissible solution algorithm shows significant improvement over most of the SNR range, the lower SNRs (10 and 20 dB) do not qualitatively show a significant improvement over TSVD when examined manually. For SNRs of 30 dB and above the reconstruction improvement is evident both in SNR and the qualitative look of the images.

To explore how the TV constraint comes into play in these reconstructions we have plotted the best MSE for each TV constraint value in Figure 5.12 along with the TSVD performance. Again at 30 dB and above the residual-TV constrained algorithm beats the TSVD approach over most of the values for the TV constraint.

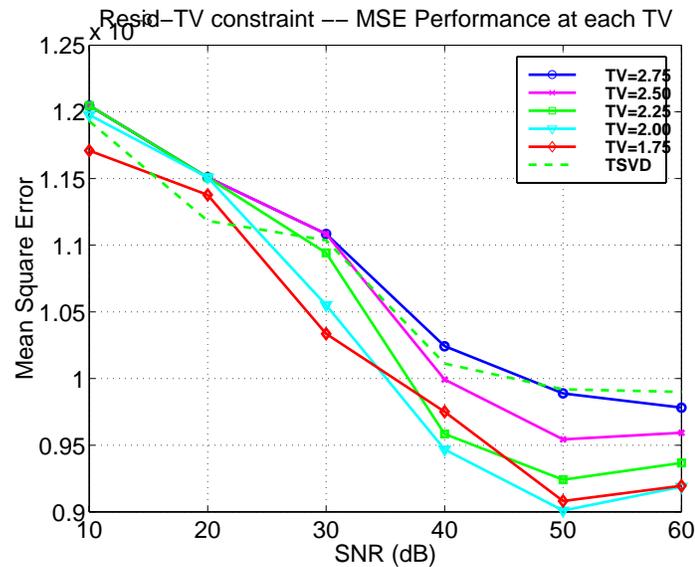


Figure 5.12: MSE versus SNR for the residual-TV constrained admissible solution approach. This graph shows the performance of each value of the TV constraint separately. Also shown is the TSVD MSE for comparison.

**Residual and normalized total variation**

The reconstructions generated from the residual and TV constrained admissible solution approach still have a large amount of “variation” even when the TV constraint is less than the TV measure of the true absorption anomaly function. This can occur because the amplitude of the reconstructed anomaly perturbation function is typically less than the true absorption anomaly function, thus reducing the value of the TV measure. To address this we developed a normalized total variation constraint and employed this in a admissible solution approach. However, simulations with no added noise showed no improvement over the residual and TV constrained admissible solution approach, so we did not pursue this line of work any further.

**Residual and Laplacian constraints**

We also explored the use of a Laplacian constraint on the solution in addition to the residual approach and the residual and TV constraint approach. In all of these cases the 2-norm of the Laplacian of the intermediate and final estimates remained well below the 2-norm of the Laplacian of the true absorption anomaly perturbation function. Therefore we did not pursue this constraint any further.

**Residual, total variation and min/max constraints**

One phenomenon that is evident in all of the previous reconstructions is a large amount “speckle” noise near the boundary. One possibility is that noise near the boundary is accounting for a significant portion of the energy in the measured data and is one of the causes for the low reconstructed absorption coefficient magnitude we see. By applying a min/max constraint, specifically limiting

the lower bound on the solution, we hope to force the solution toward better results. By removing the freedom of the solution to have large negative values near the surface we are also limiting the positive “speckle” noise. This results from the interaction of the minimum constraint and the residual 2-norm constraint. With the minimum constraint limiting the negative values near the surface, the positive reconstruction noise is limited in order for the reconstruction to agree with the data through the residual constraint.

We examined lower bounds on the absorption coefficient perturbation of -.03, -.025, -.02, -.015, -.010, -.005, and 0.0. The upper bound was set at 0.15 although we had not observed any absorption coefficient estimates that high. An example reconstruction at 30 dB SNR with the lower bound constraint set at 0.0, the TV constraint set at 2.5, and the residual constraint set at  $0.8\|n\|$  is shown below in Figure 5.13. These specific values of the constraints produced the best MSE performance for this set of constraints at 30 dB SNR. Shown next to this reconstruction is the true absorption anomaly perturbation function for reference.

Clearly we get a markedly improved reconstruction with the application of the range constraints for this scenario. The reconstructed anomaly value is very near the true value of the absorption anomaly. The location of the anomaly is accurate. The only significant distortion is the “shadow” effect below the anomaly.

The MSE performance of this algorithm with a minimum value constraint of zero is shown in Figure 5.15. This graph shows the best MSE performance at each SNR over all the residual and TV constraint values examined. Consistent with the qualitative example above we see a huge improvement in MSE over most of the range in SNR. Additionally, even at the lower SNRs, our observations were that the qualitative reconstruction of the images were greatly improved.

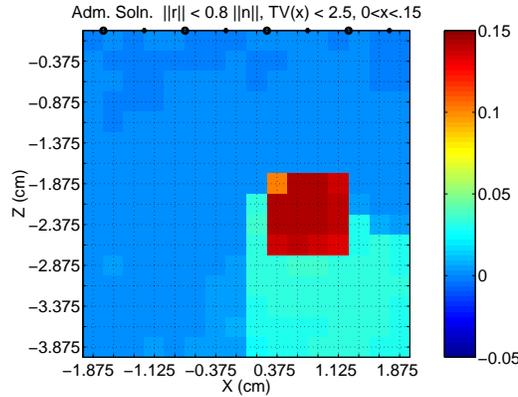


Figure 5.13: Admissible solution reconstruction with the residual and constraint  $\|r\| < 0.8\|n\|$ ,  $tv(x) < 2.5$  and a min/max range constraint of  $0 < x < 0.15$ . The SNR was 30 dB as in the previous example images. This resulted in a MSE of  $1.978e-04$  which is less than 25% of the MSE of the minimum norm solution with no added noise.

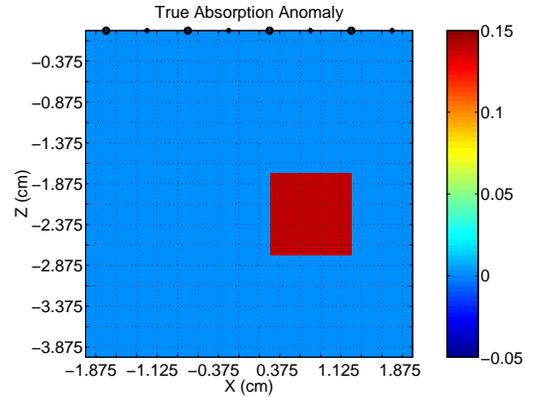


Figure 5.14: True absorption anomaly perturbation. The absorption coefficient of the anomaly is 0.139 above the background

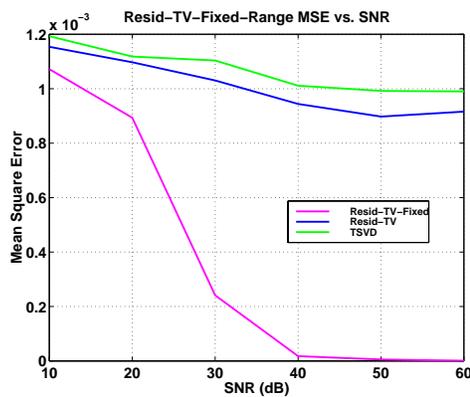


Figure 5.15: MSE versus SNR for the residual-TV-min/max constrained admissible solution approach. Also shown is the residual-TV constrained admissible solution results described in the previous section and TSVD for comparison.

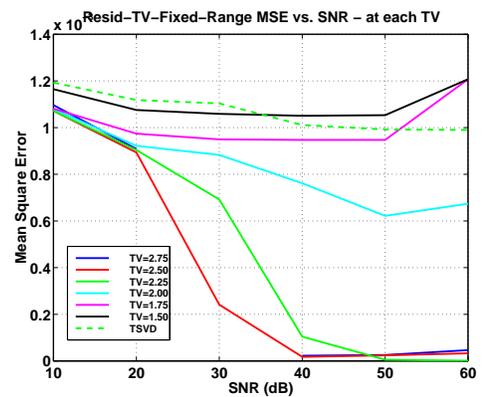


Figure 5.16: MSE versus SNR for the residual-TV-min/max constrained admissible solution approach. This graph shows the MSE performance of each value of the TV constraint separately. Also shown is the TSVD MSE for comparison.

To see how sensitive the MSE performance was to the value of TV constraint we examined each TV constraint value separately. The best MSE performance from each value of the TV constraint versus SNR is shown in Figure 5.16. From these results it is evident that we do not want to over-constrain the TV, for the  $TV < 2.0$  or less performance is much degraded. For constraints of  $TV < 2.25$  or greater (recall that the true  $TV = 2.25$ ) the MSE performance is good over the whole range of SNRs. This is in contrast to the residual-TV reconstructions in which the best MSE performance was produced by a low TV constraint value of 1.75.

Finally, we plotted the minimum MSE at 30dB SNR for versus the value of the minimum constraint. This is shown below in Figure 5.17. This plot appears to indicate that a fairly tight bound on the lower value is needed to gain all of the improvement offered by the minimum value constraint.

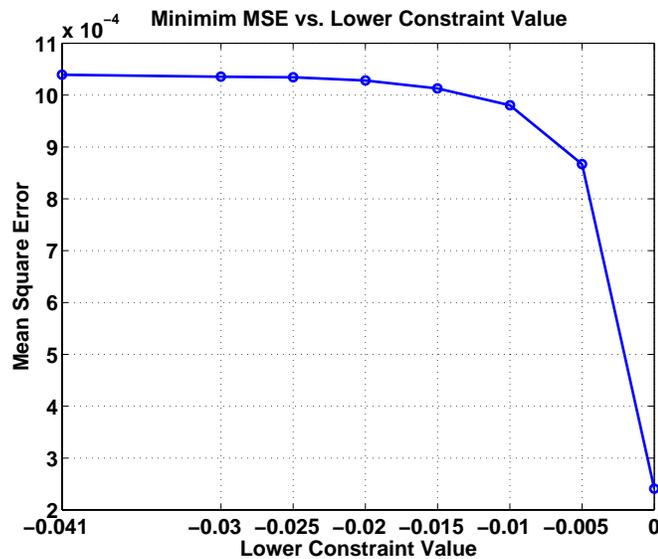


Figure 5.17: Best MSE versus the minimum value constraint for the residual-TV-min/max constraint admissible solution algorithm.

### 5.4.2 Monte Carlo noise test

All the reconstructions presented in the previous section were performed using data with the same added noise realization scaled to the desired SNR. This was done so that we could compare the various constraint sets without concern that a particular noise instance was causing ambiguous results. Computing a number of noise instances for each constraint set examined was avoided due to the large amount of computation necessary. From the previous section it is evident that the admissible solution approach with a residual 2-norm, total variation, and min/max constraint set is the most promising approach of those tested. To be sure that this marked improvement in reconstruction fidelity is not related to a particular noise instance we repeated the analysis of this approach using number of noise instances and averaged the performance results.

The constraint values examined were:

- The total variation constraint was set at 2.5
- The min/max constraint was [0 0.15]
- The residual constraint was varied over  $[2 : -0.1 : 0.1]E\{\|n\|_2\}$

As in the previous analyses, an estimate for a particular residual constraint value is computed by initializing the Ellipsoid Algorithm with the estimate and final ellipse from the next greater residual.

The largest residual constraint value is initialized with a zero vector and unit spheroid.

For each set of constraint values 10 noise realizations were examined. These were the same 10 for all values of SNR since MATLAB was restarted for each value of SNR. The SNR (uniform Gaussian random noise relative to the peak fluence response) was stepped from 10 to 60 dB in 10

dB increments as in the previous experiments. At a given SNR the best MSE performance over all constraint values was found for each noise instance. The average MSE over the noise instances was computed for each SNR, the results are shown Figure 5.18. Also shown by the error bars in Figure 5.18 is the standard deviation of the best MSE over the noise instance at each SNR value.

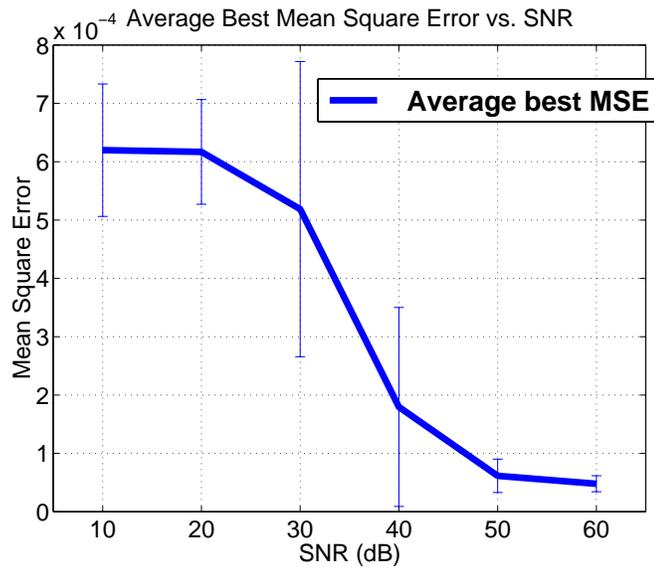


Figure 5.18: Mean square error of the estimate vs. SNR. For each noise instance the best MSE over the set of constraint values was found. The curve displays average of these best values over the ten noise instances at each SNR. The error bars show the standard deviation of the best MSE measure at each SNR.

### 5.4.3 Object parameter tests

To understand how the admissible solution approach handled variations in anomaly parameters we ran a large number of simulations varying the absorption anomaly size, position, and contrast with the background. Additionally we examined the algorithm response to multiple anomalies by placing two anomalies at varying depths.

### Object position simulations

In this experiment the location of a 1 cm (X) by 1 cm (Z) by 0.25 cm (Y) absorption anomaly was moved to nine different positions throughout the reconstruction volume. The anomaly was centered at three different depths of -3, -2, and -1 cm. For each depth, the center of the anomaly was examined at three different horizontal locations: -1, 0, and 1 cm displacements from the center of the source-detector array. At each position uniform noise was added at the typical SNRs of 10 through 60 dB in 10 dB increments. For each SNR and anomaly position a group of reconstructions was computed by successively reducing the the residual constraint. The initial starting point and ellipse was the solution from the next higher residual constraint. The constraint parameters used for these reconstructions were:

- The total variation constraint was set at 2.25
- The min/max range constraint was [0 0.15]
- The residual constraint was varied over  $[2 : -0.1 : 0.1]E\{\|n\|_2\}$

Examples of three of the nine anomaly positions are shown in Figures 5.19 through 5.21. Figure 5.19 shows the deepest and left- most position, Figure 5.20 shows the center position in both the horizontal and depth dimensions, and Figure 5.21 shows the shallowest and right most positions. Again, in each of these figures the positions of the sources is indicated by the black circles at the top of the image and the position of the detectors is indicated by the black crosses.

For each of the object simulation scenarios described previously we calculated the MSE for each reconstruction. From this data we computed the minimum MSE over the varying residual constraint

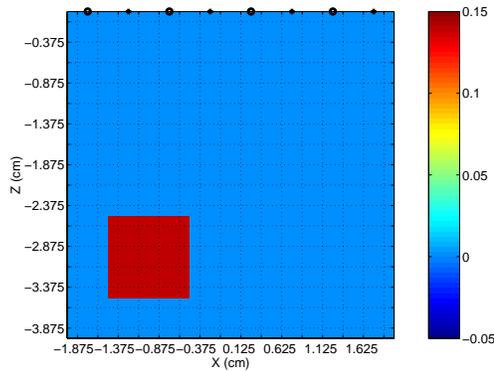


Figure 5.19: The true absorption anomaly at the deepest and left most position. The center of the anomaly in this case was at  $X = -1$  cm and  $Z = -3$  cm.

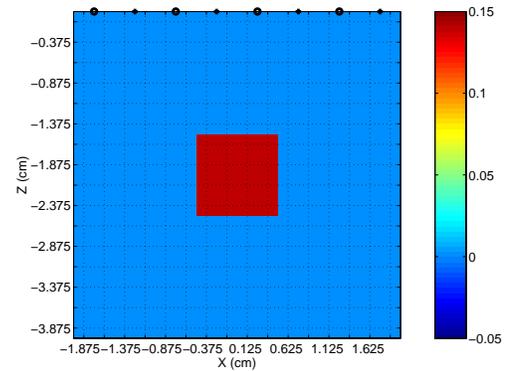


Figure 5.20: The true absorption anomaly at the center position in both the  $X$  and  $Z$  dimensions. The center of the anomaly in this case was at  $X = 0$  cm and  $Z = -2$  cm.

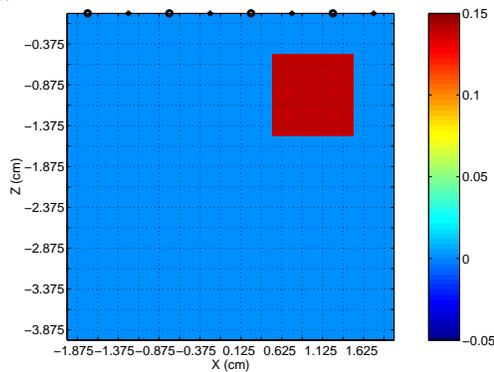


Figure 5.21: The true absorption anomaly at the shallowest and right most position. The center of the anomaly in this case was at  $X = 1$  cm and  $Z = -1$  cm.

parameter for each anomaly position and SNR. The average value of the minimum MSE vs. SNR is shown in Figure 5.22.

To see how the MSE vs. SNR curve varied with depth we averaged over each of the three horizontal anomaly positions separately at each depth. This result is shown in Figure 5.23. The two shallower anomaly depths result in much better performance than the deepest anomaly.

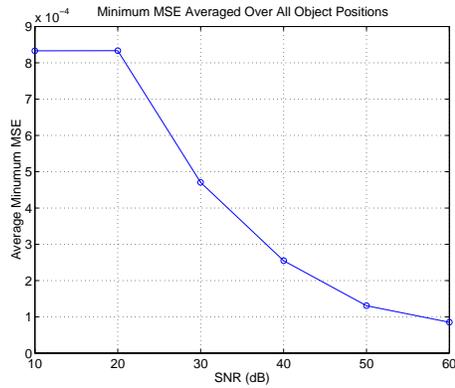


Figure 5.22: The minimum MSE averaged over all nine anomaly positions as a function of SNR.

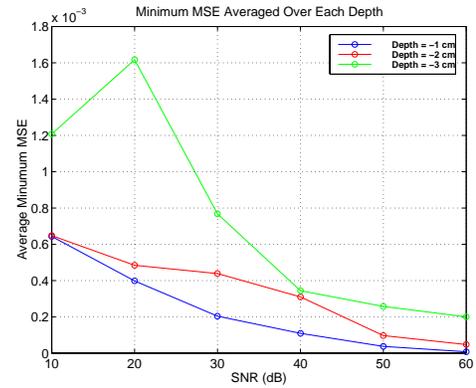


Figure 5.23: The minimum MSE averaged over each of the three anomaly positions separately for each of the object depths examined.

### Anomaly size simulations

In this set of experiments the X and Z dimensions of the anomaly were varied concurrently from 0.25 cm to 1.75 cm in 0.25 cm steps. The Y dimension was held constant at 0.25 cm, which like all of the previous experiments in this chapter was the width of the reconstruction volume. The center of the anomaly was moved slightly relative to previous tests to  $X = 0.6875$  cm and  $Z = -2.0625$  cm to facilitate enlarging the anomaly by one computational volume cell at a time (0.25 cm).

Since the size of the anomaly was being varied in this test the total variation of the true image also varied from 0.5560 for a 0.25 by 0.25 cm anomaly to 3.892 for the 1.75 by 1.75 cm anomaly. Thus a total variation constraint of 4 was used for all of the reconstructions

Examples of the true absorption perturbation function are shown in Figures 5.24 and 5.25 for the smallest and largest anomalies respectively.

As in the anomaly position experiment the minimum MSE for each anomaly size and SNR was computed. An average minimum MSE over all anomaly sizes at each SNR was then computed. The

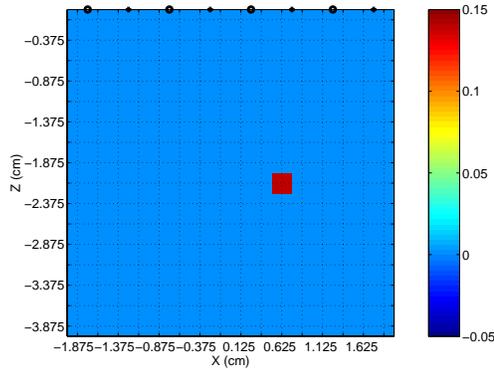


Figure 5.24: The true absorption anomaly image for the smallest size examined, the dimensions are 0.25 cm in X and Z.

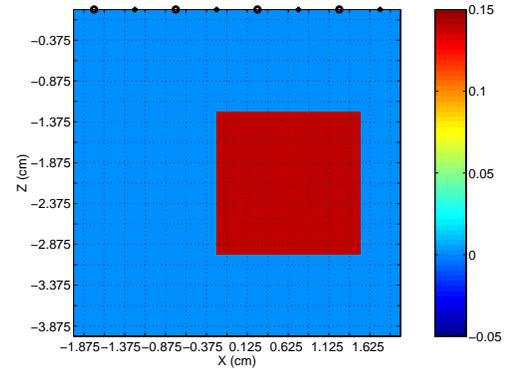


Figure 5.25: The true absorption anomaly image for the largest size examined, the dimensions are 1.75 cm in X and Z.

result is shown in Figure 5.26.

The minimum MSE for each anomaly size individually is shown in Figure 5.27. The behavior of some of the curves in this graph is somewhat surprising. For instance, for the 1.0 cm anomaly size curve, the scenario is very similar to the multiple noise realization study above. But the behavior of the minimum MSE vs. SNR curve is very different. Looking at the raw results it is evident that total variation of the reconstructions is high, typically in the high threes, compared to a TV measure of 2.25 for the true absorption perturbation function. Thus not having a tight constraint on the total variation seems to result in a degradation in MSE performance. Qualitatively though, the reconstructions still look promising. Typically, very deep high valued artifacts appear to be the result of the relaxed total variation constraint. Two example reconstructions of a 1 cm anomaly are shown in Figures 5.28 and 5.29. Figure 5.28 shows the minimum MSE 60 dB SNR reconstruction and Figure 5.29 shows the minimum MSE 20 dB SNR reconstruction.

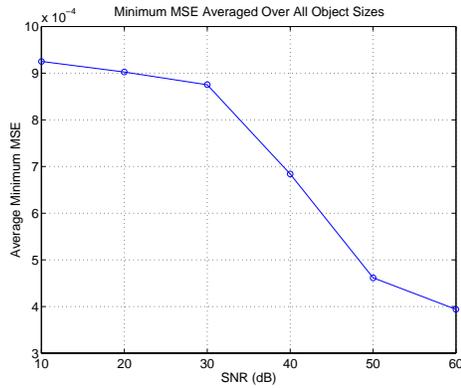


Figure 5.26: The minimum MSE averaged over all seven anomaly size as a function of SNR.

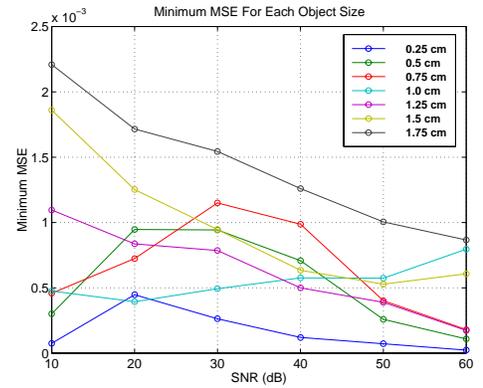


Figure 5.27: The minimum MSE plot as a function of SNR for each anomaly size separately.

### Anomaly contrast simulations

For this set of experiments the absorption coefficient of the anomaly was varied from 2 to 5 times the background absorption coefficient. The geometric parameters of the absorption anomaly were the standard 1 cm (X) x 1 cm (Z) centered 2.25 cm into the medium and 0.75 cm to the right of the center of the source-detector array.

In this experiment we used a FDFD forward model as described in Section 2.2.2 to generate the scattered fluence data. If we had used the first Born forward model as in the previous studies, changing the anomaly contrast would have no effect on the outcome of the study because of the linear nature of the model. The computational volume for the FDFD forward model was 16 cm  $\times$  16 cm in the X and Y dimensions and extended from the extrapolated boundary position to 10 cm deep into the medium. The grid size was 0.25 cm for all dimensions.

As in the anomaly size experiment described in the previous section, the total variation of the true anomaly varies as the contrast ratio changes. Since the “tightness” of the TV constraint appeared to have a significant effect on the MSE performance in the anomaly size experiment we

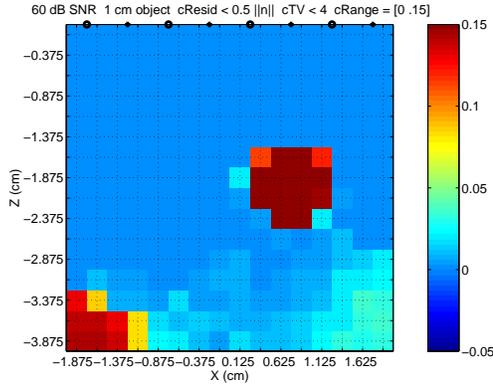


Figure 5.28: Reconstruction of a 1cm (X) by 1cm (Z) absorption anomaly at 60 dB SNR. The constraints used for the reconstruction were a residual constraint of  $0.5||n||$ , a total variation constraint of 4, and a min/max constraint of  $[0 .15]$ . This residual constraint gave the minimum MSE for this size anomaly.

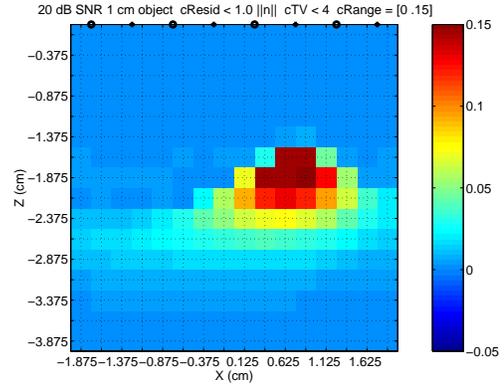


Figure 5.29: Reconstruction of a 1cm (X) by 1cm (Z) absorption anomaly at 20 dB SNR. The constraints used for the reconstruction were a residual constraint of  $1.0||n||$ , a total variation constraint of 4, and a min/max constraint of  $[0 .15]$ . This residual constraint gave the minimum MSE for this size anomaly.

explored four different constraint scenarios in this experiment. The constraint scenarios consisted of all four possible combinations of tight and loose constraints for both the total variation and the min/max constraints. The loose total variation constraint was set at 2.75 which was greater than the largest true image total variation of 2.6240. The loose min/max constraint was set at  $[0 0.18]$ , since the largest  $\mu_a$  perturbation was 0.1640 at a contrast ratio of 5. For the tight constraint scenarios, the true total variation and min/max for each anomaly instance was used.

As in the previous test the minimum MSE for each case was computed. The minimum MSE averaged over all anomaly contrasts is shown below in Figure 5.30. The minimum MSE for each anomaly contrast is shown in Figures 5.31 through 5.34

From these graphs we can observe that the tight TV and tight min/max constraint set along with the tight TV and loose min/max constraint set generally perform well. On the opposite end of the spectrum the loose TV and loose min/max constraint set generally did not perform as well until

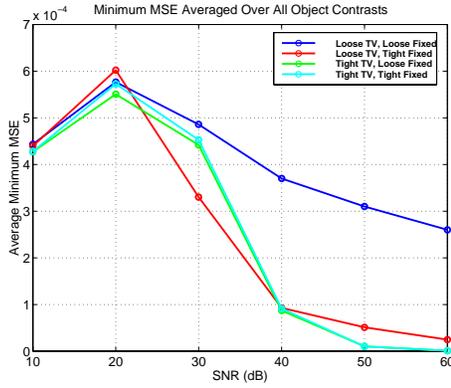


Figure 5.30: The minimum MSE averaged over all four anomaly contrasts as a function of SNR.

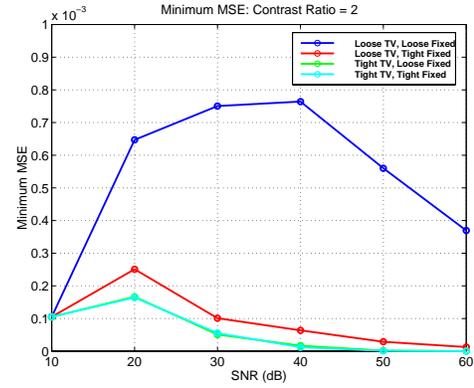


Figure 5.31: The minimum MSE for an anomaly with a contrast ratio of 2:1 in absorption coefficient as a function of SNR.

higher contrast ratios were achieved, at which point the constraints were pretty tight.

### Multiple anomaly simulation

In order to understand how the admissible solution approach with the residual-TV-min/max constraint set performed with multiple absorption anomalies we ran 3 sets of simulations. The right hand side of the reconstruction volume contained the absorption anomaly used in the previous studies at a depth of 2 cm.. On the left hand side we placed an absorption anomaly of the same size (1 cm x 1 cm x 0.25 cm) in one of three positions. The left hand anomaly center depths were 1 cm, 2 cm and 3 cm. The X position of the left hand anomaly was -1 cm for all three instances. The absorption coefficient of both anomalies was  $0.18 \text{ cm}^{-1}$  or  $0.139 \text{ cm}^{-1}$  above the background. Images of all three scenarios are shown in Figures 5.35 through 5.37.

The constraints employed for this experiment were a min/max constraint of  $[0 \ 0.15]$  and TV constraint of 4.5. The total variation of the true image was 4.48. The residual constraint, as in the previous experiments, was reduced from  $2\|n\|$  to  $0.1\|n\|$  in steps of  $0.1\|n\|$ .

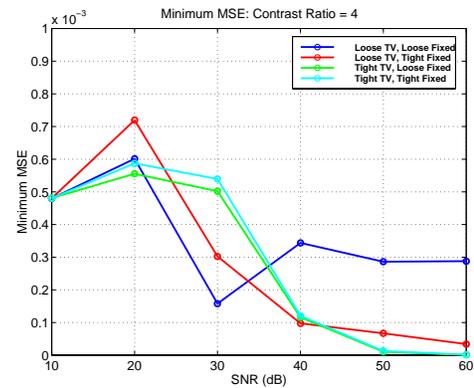
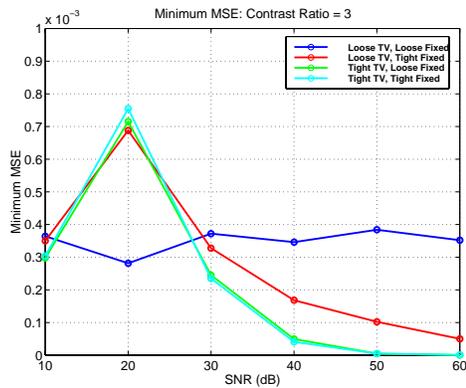


Figure 5.32: The minimum MSE for an anomaly with a contrast ratio of 3:1 in absorption coefficient as a function of SNR.

Figure 5.33: The minimum MSE for an anomaly with a contrast ratio of 4:1 in absorption coefficient as a function of SNR.

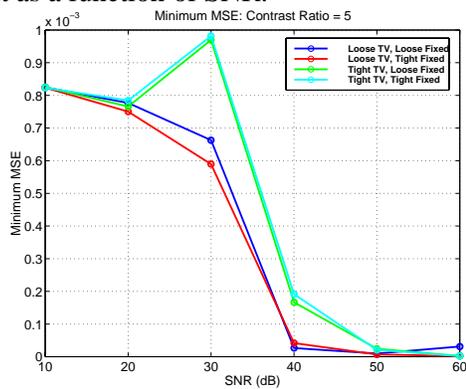


Figure 5.34: The minimum MSE for an anomaly with a contrast ratio of 5:1 in absorption coefficient as a function of SNR.

From this image we can clearly resolve both anomalies, their positions are correct, and their amplitudes are very close to the true amplitude. It appears that more of the scattered field energy is assigned to the shallower anomaly since its amplitude is a little overestimated whereas the amplitude of the deeper anomaly is slightly underestimated.

At 30 dB SNR two anomalies are still more-or-less visible but at a clearly reduced fidelity compared to the 60 dB SNR case. The energy redistribution to the shallower anomaly appears to be

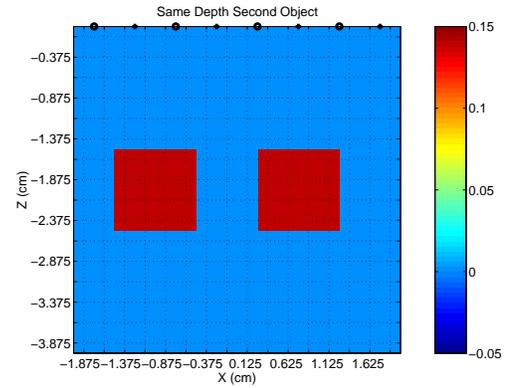
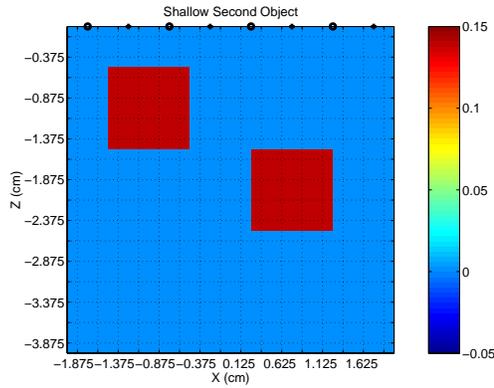


Figure 5.35: The true absorption anomalies with the left hand anomaly at the shallow position of 1 cm.

Figure 5.36: The true absorption anomalies with both at the same depth of 2 cm.

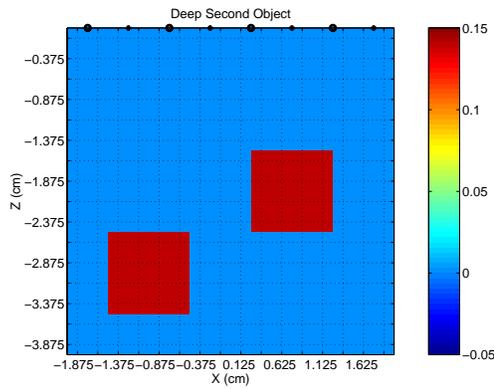


Figure 5.37: The true absorption anomalies with the left hand anomaly at the deep position of 2 cm.

even more evident in this case. The shallow anomaly is significantly overestimated.

When the anomalies are at the same depth the reconstruction fidelity is the best. In the 60 dB SNR case above both anomalies are accurately reconstructed. The anomaly amplitudes are a little high and the bottom set of voxels is a little underestimated.

At 30 dB SNR the 2 anomalies are still clearly evident but a distorted region between the anomalies is present. Also it appears that having the left hand anomaly centered 1 cm from centerline and the right hand anomaly centered at 0.75 cm from centerline attributes more energy to the right hand

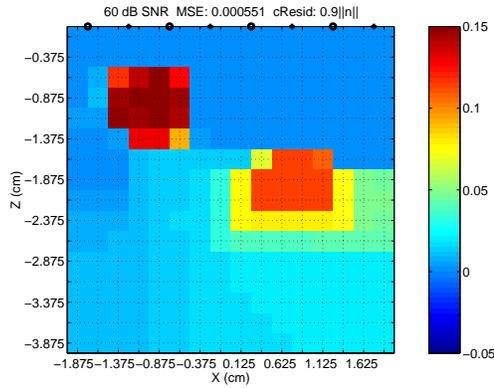


Figure 5.38: The minimum MSE reconstruction for the shallow left hand anomaly scenario at a SNR of 60 dB.

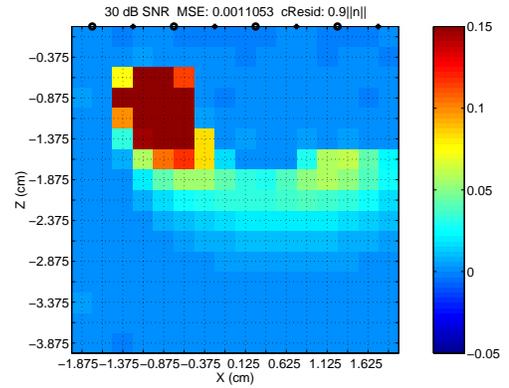


Figure 5.39: The minimum MSE reconstruction for the shallow left hand anomaly scenario at a SNR of 30 dB.

anomaly.

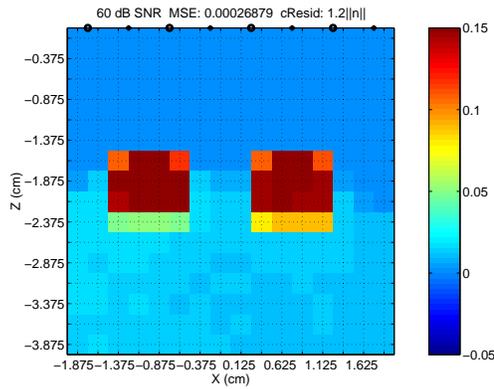


Figure 5.40: The minimum MSE reconstruction for the same depth anomalies at a SNR of 60 dB.

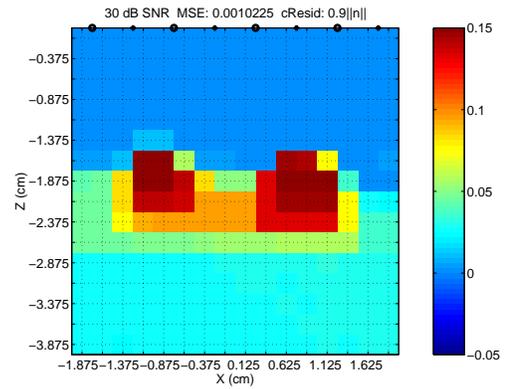


Figure 5.41: The minimum MSE reconstruction for the same depth anomalies at a SNR of 30 dB.

With the deep left hand anomaly there is some evidence of the deep anomaly at 60 dB SNR but the actual physical dimensions and amplitude are difficult to identify.

In the 30 dB SNR deep left anomaly case above the deep anomaly is not detectable. It does appear that the additional scattered field from the second anomaly causes some distortion in the reconstruction of the shallower right hand anomaly.

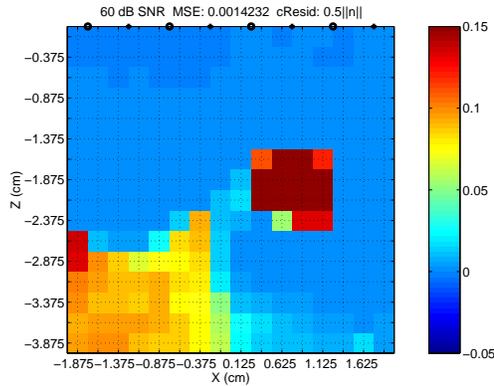


Figure 5.42: The minimum MSE reconstruction for the deep left hand anomaly scenario at a SNR of 60 dB.

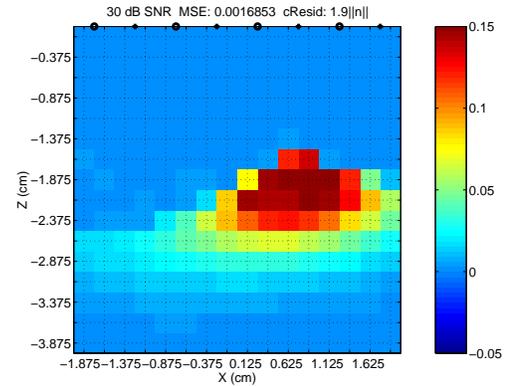


Figure 5.43: The minimum MSE reconstruction for the deep left hand anomaly scenario at a SNR of 30 dB.

The minimum MSE for each of the scenarios as function of SNR is show in Figure 5.44. The performance across SNR is as expected from the raw results on the previous page. The best MSE performance is given by the scenario where the anomalies are the same depth followed by the shallow left hand anomaly scenario.

#### 5.4.4 Variable background test

Since the lower min/max constraint of zero made such a significant difference in the reconstruction performance we wanted to explore what effect a true variation in the background would have on the reconstructions under such a tight min/max constraint. The background variation was simulated by adding a IID Gaussian random vector to the perturbation vector used in the Born-1 forward model. The magnitude of the background variation was varied so that the standard deviation of the Gaussian random vector was either 0.1%, 0.2%, 0.5%, 1.0%, 2.0%, 5.0% or 10% of the background absorption coefficient. The voxel size used in the forward simulation was 0.125 cm in the X and Z dimensions. The parameters chosen to model the background variation were somewhat arbitrary

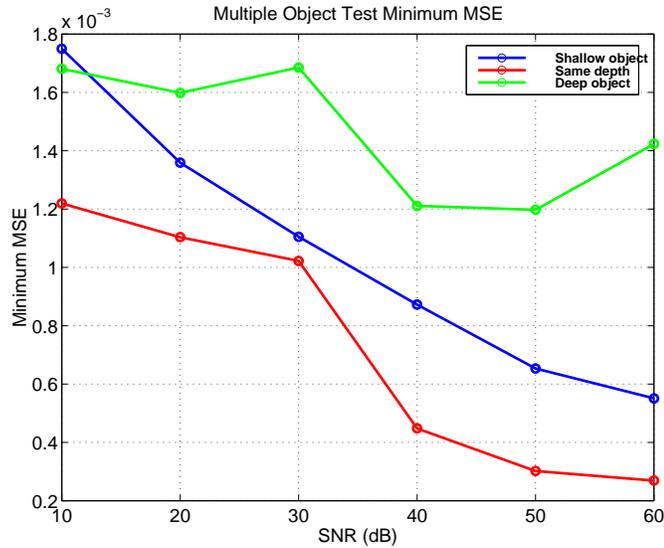


Figure 5.44: The minimum MSE versus SNR for all three multiple anomaly scenarios.

because we were unable to find any information on the expected range and structure of background variations in tissue. We hope that we have covered enough of a range to provide some meaningful results when applied to true tissue measurements.

We considered three ways to handle the expected absorption values below the background absorption coefficient. The first was to basically ignore it. In this case both the forward and inverse models use the same background optical parameters. The minimum value constraint was still left at 0 forcing an incorrect result where the background variation induced a absorption parameter less than the background. This approach we called “unmodified”. The second approach was to set the minimum value constraint to the minimum value of the true absorption perturbation. This approach we called “modified range”. The final approach was to change the background absorption coefficient used in the inverse model to the minimum absorption coefficient in the reconstruction volume.

This ensured that the reconstructed image could be bounded below by zero. This approach we called “modified model”.

The minimum MSE performance for each of the variable background tests are shown in Figures 5.45 through 5.51. Clearly leaving the admissible solution approach unmodified does not work very well. Reducing the lower-bound on the min/max constraint, however, does work well.

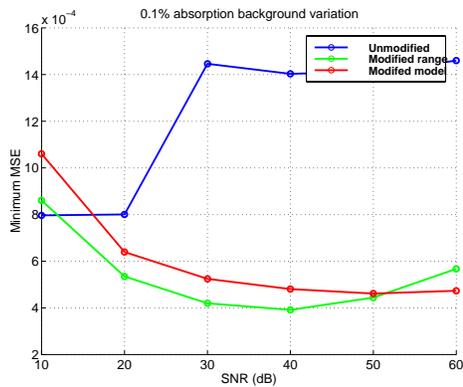


Figure 5.45: Minimum MSE for the three different approaches with 0.1% background variation.

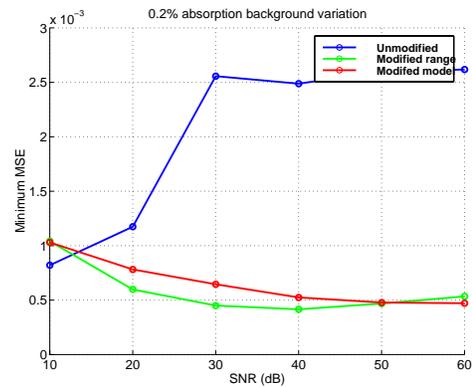


Figure 5.46: Minimum MSE for the three different approaches with 0.2% background variation.

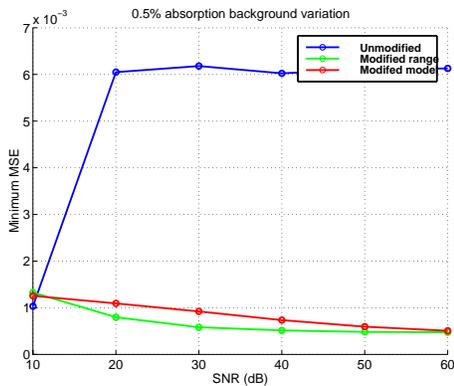


Figure 5.47: Minimum MSE for the three different approaches with 0.5% background variation.

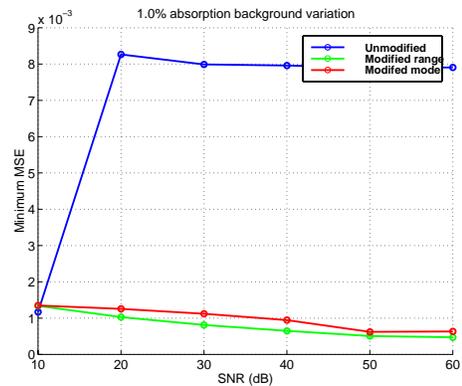


Figure 5.48: Minimum MSE for the three different approaches with 1.0% background variation.

To get an idea of how the background variation and noise interact we plotted minimum MSE contours versus both SNR and background variation for the modified range and modified model

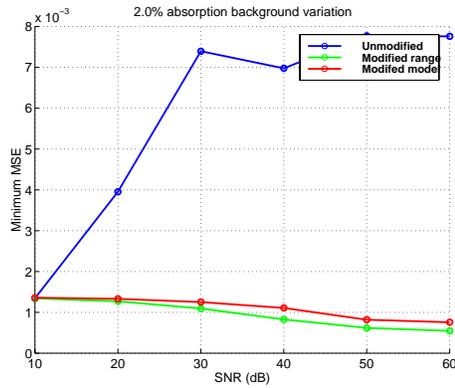


Figure 5.49: Minimum MSE for the three different approaches with 2.0% background variation.

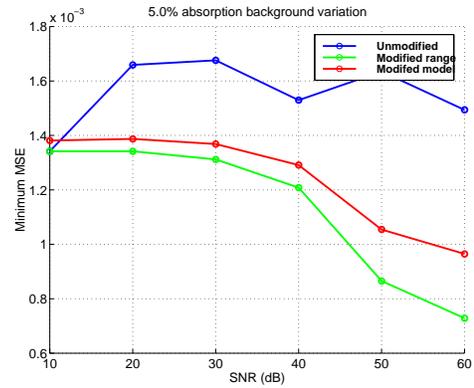


Figure 5.50: Minimum MSE for the three different approaches with 5.0% background variation.

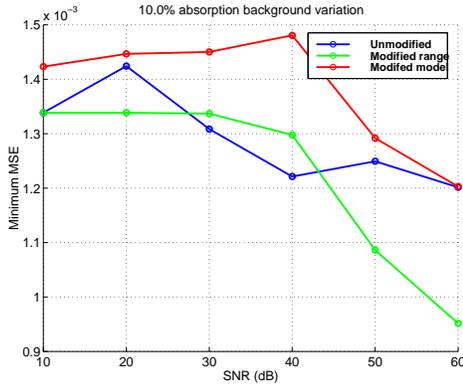


Figure 5.51: Minimum MSE for the three different approaches with 10.0% background variation.

approaches. These are shown in Figures 5.52 and 5.53

Finally, example reconstructions using the three approaches at 30 dB SNR and 1% background variation are shown in Figures 5.54 through 5.56.

## 5.5 Computational considerations

The computational work required to implement the admissible solution approach is significantly higher than the linear algebraic reconstruction technique presented in Chapter 3. Using the FLOPS

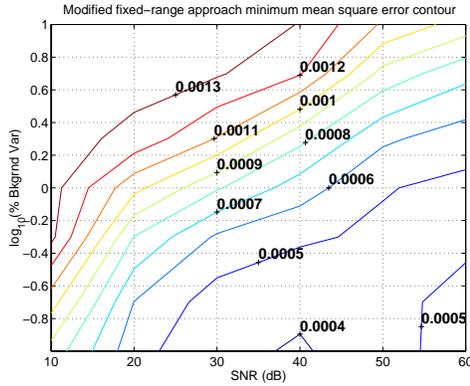


Figure 5.52: Minimum MSE contour for the modified range approach as function of SNR and background variation.

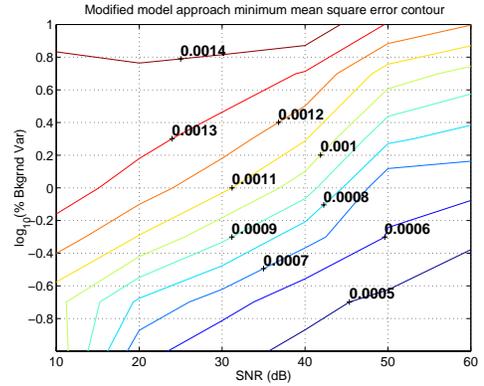


Figure 5.53: Minimum MSE contour for the modified model approach as function of SNR and background variation.

counter within MATLAB we counted between 15E9 FLOPS for the residual constraint alone to 550E9 FLOPS for the residual, total variation, and min/max constrained solution for a single constraint set. In comparison, to compute the “economy” SVD for the forward operator requires 2.5E6 FLOPS.

## 5.6 Admissible solution discussion

It is evident from the results in these simulations that the incorporation of additional constraints did improve the reconstruction fidelity of the absorption coefficient perturbation.

Incorporating the initial additional constraint, the total variation constraint, brought a modest improvement to the reconstructed absorption coefficient perturbation. Qualitatively we observed, a deeper and larger response, associated with the anomaly. Quantitatively we observed a small improvement in the MSE measure over the whole range of SNR. The sensitivity to the value of this constraint was not too great, and as expected if the constraint was too loose the algorithm

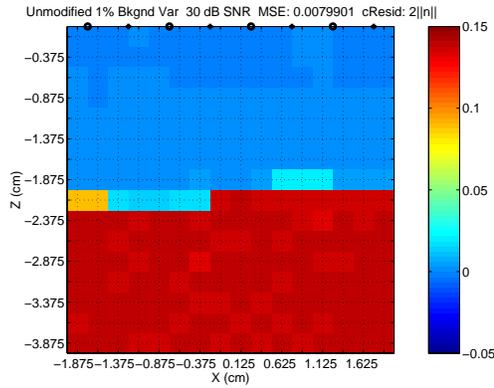


Figure 5.54: Unmodified minimum MSE reconstruction at 30 dB SNR and 1% background variation.

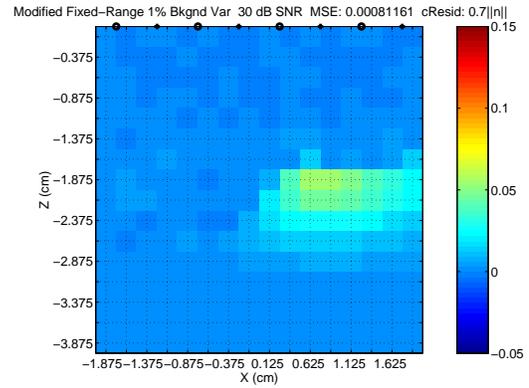


Figure 5.55: Modified range minimum MSE reconstruction at 30 dB SNR and 1% background variation.

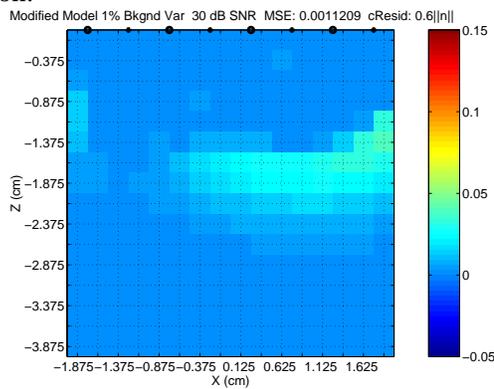


Figure 5.56: Modified model minimum MSE reconstruction at 30 dB SNR and 1% background variation.

performance reverted back to that of the constrained residual approach. The reduction in expected improvement from the addition of the total variation constraint appears to be a result of the small amplitude of the reconstruction. Because of the large effective null space, having the solution meet the data does not require a reconstruction with a similar magnitude as the true absorption anomaly perturbation function. This small amplitude reduces the total variation measure and thus limits the effects of the constraint on the reconstruction. This was the motivation for development of the

normalized total variation.

Unfortunately the observed effect of the normalized total variation measure in the noiseless case was to induce a large constant negative value region deep into the medium. Therefore further investigation of this constraint was not warranted at this time. Future investigation should attempt to understand why this response was induced and whether there is a more appropriate manner in which to normalize the total variation constraint.

Incorporating the min/max constraint is where we really saw a large improvement in the reconstruction fidelity. We hypothesize that this improvement due to the addition of the min/max constraint is caused by the following relationship. First, looking at the residual alone and residual 2-norm and total variation constrained reconstructions in Figures 5.8 and 5.10 we observe a noisy region near the surface that has both negative and positive variations. The minimum value constraint we have applied restricts these negative absorption coefficients near the surface, but for the residual constraint to be met, the positive regions of noise must also disappear. This leaves more measured signal energy that must be accounted for by the absorption anomaly, resulting in an increase in the amplitude of the reconstructed absorption anomaly. A further illustration of the effect of the minimum constraint is the MSE measure's sensitivity to this constraint as shown in Figure 5.17. This graph clearly identifies the minimum value constraint as a primary factor in improving the reconstruction fidelity.

Our multiple noise realization tests, object parameter test, and background variation test all showed that this approach is fairly robust to changes in the reconstruction scenario and environment. As expected, deep absorption anomalies, either alone or with another absorption anomaly present, are more difficult to reconstruct. Also in our investigation of the contrast ratio, our use of the FDFD

forward solution and the good quality reconstructions produced provides further evidence that a linear model approach is feasible.

Finally, the background variation examination provided some interesting results. Before running these experiments we expected that the modified model approach would work the best, since it was still allowed to use a minimum constraint of zero and suppress any negative responses as described above. The results were somewhat surprising in that the modified range approach performed better than the other two approaches even though it did allow some negative responses. It appears that somehow allowing for negative responses when they do actually exist (*i.e.* true variations in the background less than the background) is not as detrimental as when they are not actually there. This is a positive result in the sense that a more relaxed set of constraints may be appropriate instead of using a more perturbed model when dealing with expected background variation.

## Chapter 6

# Conclusions and future work

As we have seen throughout this thesis, accurate reconstruction of the optical absorption coefficient from fluence rate measurements at the body surface is a difficult problem. For the full three dimensional reconstruction the problem is both ill-posed and severely underdetermined. We have demonstrated that when a linear model based reconstruction is employed, the use of subspace based techniques resulted in higher fidelity reconstruction of the absorption coefficient. These techniques also have an added advantage in that an *a posteriori* technique, the L-curve, is available for the selection of the regularization parameter.

By the incorporation of additional *a priori* constraints we have been able to significantly improve the fidelity of the absorption coefficient reconstructions. We have explored two general approaches to incorporating additional *a priori* constraints into the reconstruction process. The two-wavelength approach did not result in as much improvement as we had hoped but did show that additional constraints can improve the reconstruction fidelity. Further investigation into this technique might

indicate how or where it can be best applied. The most useful technique we have found to incorporate additional constraints is through an admissible solution approach implemented with the ellipsoid algorithm. Specifically, through the use of a combined residual 2-norm, total variation, and min/max constraint set, we were able to achieve near perfect reconstruction of the absorption parameter with an approximation to the full three dimensional problem. Additionally, this approach has shown to be fairly robust to changes in the reconstruction scenario such as absorption anomaly parameter variations and variations in the background characteristics.

## 6.1 Future work

There are a number of aspects of this problem that still need to be explored. We have presented a more accurate noise model that was developed after the simulations we have presented were run. Incorporation of this more accurate noise model and the examination of its effects on the reconstruction still needs to be studied. The use of more accurate forward models to simulate data is also planned. Although basic analysis of the forward problem has shown that the first Born approximation provides reasonable results, we have observed in several instances that scenario changes have produced unexpected results. Now that we have a basic foundation for the understanding of linear model reconstruction approaches to DOT we can advance our simulations to employ more accurate forward models.

We also plan on implementing the full three dimensional admissible solution approach. We feel that the full three dimensional admissible solution approach may produce better results than the vertical two dimensional slice approximation we have used in our initial study. With a 3D array

the performance at greater depth will hopefully be improved due to the available larger source-detector distances and the greater number of sampling “banana patterns”. Even though we matched the singular value spectra and underdeterminedness ratio between the three dimensional and two dimensional problems, we cannot be sure that the problems are similar in difficulty until we have performed more studies. This also identifies another area of investigation described in more detail below, namely the effect of the source detector geometry.

Due to the computational requirements of this algorithm, the full three dimensional approach will likely have to be implemented on a parallel architecture. Additionally, we plan on investigating the use of the Deep Cut Ellipsoid Algorithm [50] which can significantly reduce the number of iterations required to find an admissible solution.

Since the admissible solution approach requires the use of *a priori* constraints on the solution, we plan on investigating further the appropriate values of those constraints, and their sensitivity under different measurement scenarios. To expand upon the requirement of having *a priori* constraints we would also like to investigate the possibility of the constraint being modified on the fly and estimated from the data as we iterate, thus developing an adaptive constraint approach. We also plan on investigating the sensitivity of the min/max constraint to background variation when the extent of that variation is unknown.

In regard to the two-wavelength approach, more work needs to be done to understand the effects of the constraint and how to best implement the constraint. We plan on examining a smaller two dimensional problem that will allow us to efficiently examine more scenarios for the CACR constraint.

Finally, our observations in this thesis have been that the reconstruction fidelity is very sensitive

to the geometry of the sources and detectors. We propose to investigate improved source/detector geometries with the application restrictions in mind and to investigate if there is a way to specify an optimal, in some sense, source/detector distribution for a given problem.

## **Appendix A**

### **Table of Parameters**

Since the topics discussed in this thesis cover several different areas of study, we provide a table of parameters and variable symbols we have used. We have attempted to be consistent with this notation throughout the thesis.

Parameter	Description	Page reference
$\alpha$	Absorption coefficient ratio for the CACR constraint	76
$\beta$	Modulation depth	44
$\mu_a$	Optical absorption coefficient	8
$\mu_a^o$	Background optical absorption coefficient	27
$\Delta\mu_a(\mathbf{r})$	Optical absorption coefficient perturbation function	27
$\eta$	Detector quantum efficiency	45
$\lambda$	Photon wavelength	76
$\mu_s$	Optical scattering coefficient	8
$\mu_s'$	Reduced optical scattering coefficient	8
$\mu_t$	Optical transport coefficient	8
$\nu$	Photon frequency	45
$\Omega_D$	Detector fiber acceptance angle	45
$\phi_{AC}$	AC propagation constant	45
$\phi_{DC}$	DC propagation constant	45
$\Phi(\mathbf{r}, t), \Phi(\mathbf{r}, \omega)$	Photon fluence rate	9
$\Phi^i(\mathbf{r})$	Incident photon fluence rate	29
$\Phi^i$	Discrete incident photon fluence rate	44
$\Phi^s(\mathbf{r})$	Scattered photon fluence rate	29
$\Phi^s$	Discrete scattered photon fluence rate	33
$\tilde{\Phi}^s$	Stacked real discrete scattered photon fluence rate	34
$A_D$	Detector fiber area	45
$B$	Detector bandwidth	45
$D$	Diffusion coefficient	10
$e$	Electron charge	45
$f(\hat{\Omega}, \hat{\Omega}')$	Scattering phase function	8
$F_n$	Noise figure	45
$G(\mathbf{r}, \mathbf{r}')$	Green's function	44
$\mathbf{G}$	Forward operator matrix	44
$\tilde{\mathbf{G}}$	Stacked real forward operator matrix	34
$g$	Mean cosine of the phase function	8
$\mathbf{J}(\mathbf{r}, t)$	Photon flux	9
$k$	Diffuse photon density wavenumber	11
$k_o$	Background diffuse photon density wavenumber	29
$\Delta k(\mathbf{r})$	Diffuse photon density wavenumber perturbation function	29
$\hbar$	Planck's constant	45
$\mathbf{k}$	Discrete Diffuse photon density wavenumber perturbation function	33
$L(\mathbf{r}, \hat{\Omega}, t)$	Radiance	8
$NEP$	Noise equivalent power	45
$\mathbf{R}$	Covariance matrix	53
$v$	Electromagnetic propagation velocity	8

Table A.1: Table of Parameters

## **Appendix B**

# **Photon Migration Imaging Toolbox**

### **B.1 PMI Toolbox Introduction**

This document describes the architecture of the Photon Migration Imaging (PMI) toolbox. Specifically the Photon Migration Imaging structure and some of the high level MATLAB functions that operate on this structure. The PMI structure is implemented as a tree of MATLAB structures containing all of the information about an experiment. For most operations, such as simulating measured data, or reconstructing absorption maps, this top level data structure is passed as the argument to a function and received as the result. Additionally, the PMI data structure can be supplied as the argument to the GUI function `iPMI` and all operations can be run from within this user interface.

The PMI data structure is primarily a container for other structures that describe the parameters of the experiment and hold the results. The PMI data structure and its child structures are described in Table 1 through Table 11. The layout of these structures is valid for version 3.11 of the PMI toolbox.

## B.2 High Level Algorithms

A number of high level functions are available that perform common diffuse optical tomography operations. These functions are useful both for implementing automated computer experiments through the use of command sequences or scripts and for conducting interactive simulations. Each of these algorithms takes a data structure containing the parameters needed by the algorithm and supply results in the appropriate fields within a returned data structure. A simple example of creating a Truncated Conjugate Gradient (TCG) image of a 1st order born simulated measurement at 780 nm is shown below.

```
%
% Load in the standard 780 nm simulation parameters from the
% template file.
%
load pmi780v3

%
% Modify some of the parameters as necessary.
%
pmi780.Fwd.ModFreq = 200;
pmi780.Inv.ModFreq = 200;

pmi780.Recon.ReconAlg = 'TCG';
pmi780.Recon.TCGnIter = 10;

%
% Generate the first order Born approximation forward matrix for
% both the forward and inverse problem as well as the measured data.
%
pmi780.Fwd = genFwdMat(pmi780.Fwd);
pmi780.Inv = genFwdMat(pmi780.Inv);
pmi780 = genMeasData(pmi780);
pmi780 = genNoise(pmi780);

%
% Compute the requested reconstruction
%
pmi780 = genRecon(pmi780);

%
% Display the reconstructed absorption function over the
% computational volume
%
```

```
showImage(pmi780)
```

The high level algorithms available include the following categories:

### Forward Problem Algorithms

<code>genFwdMat</code>	Generate a forward matrix and incident field.
<code>genMeasData</code>	Generate simulated measured data
<code>genNoise</code>	Add the requested noise model to the measured data.

### Reconstruction Algorithms

<code>genRecon</code>	Generate the absorption image using the specified reconstruction algorithm.
-----------------------	---

### Visualization Algorithms

<code>showImage</code>	Display the resulting reconstruction.
<code>multImage</code>	Display a number of reconstructions with the ability to scan by a scroll bar.

## B.2.1 The Photon Migration Imaging data structure

The following tables describe the fields present in the Photon Migration Imaging data structure and its children structures. The hierarchical relationship between these structures is shown graphically in Figure 1.

Items shown in *italics* are only included when the `Debug` field is greater than zero (`true`). Most are included when `Debug` is one or greater, some less important information is included only when `Debug` is greater than one. See the algorithm that generates the specific field to identify exactly when a debug dependent field is generated.

Table B.1: The Photon Migration Imaging data structure. This is the top-level data structure that most of the high-level operations take as an argument and return. It contains all of the information and results from a particular experiment.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Version	scalar	The version number for this data structure.
Debug	scalar	A flag specifying the level of debug messages to display and the amount of debug info to carry along in the data structure.
Fwd	structure	A Model structure. This contains containing all of the information to calculate the forward problem.
Object	cell array	A cell array of Object structures. This contains all of the information on This contains all of the information regarding absorption objects and the resulting perturbation inhomogeneities present in the computational volume. Each element in the cell array contains an Object data structure and specifies an individual object in the medium.
Noise	structure	A Noise structure. This specifies the noise models to apply and contains the weighting vector for "whitening" the system. Also available if debug is enabled are the noise vectors and individual standard deviations.
PhiTotal	matrix	The total measured fluence rate without noise. Each column represents the fluence rate measured at a particular wavelength. Each column is stacked running through all of the detector for the first source and then incrementing the source
PhiTotalN	matrix	The total measured fluence perturbed by the noise specified in the Noise structure.
PhiTotalNw	matrix	The total measured fluence perturbed by the noise specified in the Noise structure and weighted such that the variance of each measurement is one.
Recon	structure	A Reconstruction structure. This specifies the reconstruction algorithm and it's parameters along with holding the estimate(s) of the absorption perturbation function(s).
Visualize	structure	A Visualization structure. This specifies the parameters for the selected visualization technique.

Table B.2: The Model data structure, this is used for both the forward and inverse problems.

<b>Field</b>	<b>Type</b>	<b>Description</b>
idxRefr	vector	The index of refraction in the medium. For multiple wavelengths this should be a vector specifying the index at each wavelength.
v	vector	The velocity of electromagnetic propagation in the medium. For multiple wavelengths this should be a vector specifying the propagation velocity at each wavelength.
g	vector	The mean cosine of the scattering angle. For multiple wavelengths this should be a vector specifying the mean cosine of the scattering angle at each wavelength.
Cphore	vector	The vector of chromophore structures, each element specifies a different chromophore identified in the structure.
Mu_s	vector	The scattering coefficient for the medium. For multiple wavelengths this should be a vector specifying the scattering coefficient at each wavelength.
Mu_sp	vector	The reduced scattering coefficient for the medium. For multiple wavelengths this should be a vector specifying the reduced scattering coefficient at each wavelength.
Mu_a	vector	The absorption coefficient for the medium. For multiple wavelengths this should be a vector specifying the absorption coefficient at each wavelength.
Src	structure	An Optode structure specifying the position(s) and parameters such as numerical aperture and amplitude of the source(s).
Det	structure	An Optode structure specifying the position(s) and parameters such as numerical aperture and amplitude of the detector(s).
Lambda	vector	The wavelengths of the sources used.
ModFreq	vector	The modulation frequency(s) (in MHz) of the sources. Multiple modulation frequencies generate multiple sets of measurements that are stacked in the order of frequency specified in this vector.

Table B.2 the Model data structure continued.

<b>Field</b>	<b>Type</b>	<b>Description</b>
MeasList	matrix	This matrix defines the order in which measurements are present in the measurement data vector. It is an $n \times 4$ matrix where each row identifies a specific measurement and each column indexes the source, detector, frequency and wavelength respectively.
Boundary	structure	A Boundary structure defining the boundary conditions and locations to use.
Method	structure	This is a Method structure. It defines the method used to calculate forward solution or create a forward operator for the inverse solution as well as any parameters that are required by the method specified.
CompVol	structure	This is a Computational Volume structure. It specifies the values in three dimensions on which the forward problem or inverse problem will be evaluated (if appropriate).
A	array	The forward operator (if appropriate).
Aw	array	The weighted forward operator (if appropriate). Column weighting is used to "whiten" noise present in the measurement.
delMu_a	matrix	The absorption function perturbation. This is the value of the absorption function minus the background absorption for each voxel. The 3 dimensional array is converted into a vector by stacking, the stacking order is y x z. Each column contains the absorption function perturbation for a different wavelength.
delMu_sp	matrix	The scattering function perturbation. This is the value of the scattering function minus the background absorption for each voxel. This field has the same structure as delMu_a.
PhiInc	matrix	The computed incident field(s).
PhiScat	matrix	The computed scattered field(s)
PhiScatw	matrix	The weighted computed scattered field(s), the weighting is performed to whiten the noise present in the measure data. This field will most likely only be present in the inverse model (Inv) data structure.
IPR	matrix	The incident-to-perturbation ratio for each source-detector pair at each frequency and wavelength.

Table B.3: The Chromophore data structure. This structure describes the concentration and resultant absorption coefficient of the specified chromophore.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Name	string	The name of the chromophore.
Conc	scalar	The concentration of the chromophore in the specified region.
Mu <sub>a</sub>	vector	The absorption coefficient of the chromophore at each wavelength.

Table B.4: The Optode data structure. This is used to specify the positions, orientation, and optical parameters of the sources and detectors.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Type	string	This specifies the type of position data. This can either be 'uniform' where X, Y, and Z, specify a uniform array of positions or 'list' where Pos is a n x 3 matrix specify a list of positions
X	vector	The X positions of the optode array.
Y	vector	The Y positions of the optode array.
Z	vector	The Z positions of the optode array.
Pos	matrix	A n x 3 matrix that specifies a list of arbitrary points in 3-space.
Normal	matrix	A n x 3 matrix that specifies the outward normal for each source or detector.
Amplitude	vector	The amplitude of each source or relative efficiency of each detector. Efficiency should be specified between 0 and 1.
NA	vector	The numerical aperture of each source or detector.

Table B.5: The Boundary data structure. This describes the type of boundary model to use and any parameters that are needed with the type specified.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Geometry	string	A description of the boundary geometry. Valid string descriptors are: 'Infinite' - An infinite medium. 'Semi-infinite' - A semi-infinite medium, the air-tissue interface is assumed to be at $z = 0$ . 'Slab' - A slab medium. One boundary is assumed to be at $z = 0$ , the other is at $Z = \text{Thickness}$ .
Thickness	scalar	The thickness of a slab medium.

Table B.6: The Method data structure. This describes the method used to compute either the forward problem or forward operator for the inverse problem.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Type	string	The forward model method used to either generate the data or create a forward matrix. The currently known values are: 'Born', 'Rytov', 'Spherical', 'MatlabVariable', 'DFD', 'FEM'.
Order	scalar	A value dependent field. If Type is 'Born', or 'Spherical' then this field specifies the order of the approximation.
MatlabVarName	string	A value dependent field. If Type is 'Matlab Var' then this field specifies the name of the Matlab variable to read in the measured data from. The MATLAB variable should be an $n \times 2$ matrix where the columns contain the total fluence and measurement variance respectively.

Table B.7: The Computational Volume data structure. This is used to specify the points at which the reconstruction will be performed and if appropriate where the forward solution will be calculated.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Type	string	This specifies the type of voxelation data. This can either be 'uniform' where X, Y, and Z, specify a uniform array of voxels or 'list' where Pos is a n x 3 matrix specify a list of positions
X	vector	The centers of the voxels in the X direction.
XStep	vector	The step size in the X direction for a 'uniform' voxelation
Y	vector	The centers of the voxels in the Y direction.
YStep	vector	The step size in the Y direction for a 'uniform' voxelation
Z	vector	The centers of the voxels in the Z direction.
ZStep	vector	The step size in the Z direction for a 'uniform' voxelation
Pos	matrix	A n x 3 matrix that specifies the arbitrary points in the volume used for voxelation.

Table B.8: Object data structure. This data structure specifies the optical heterogeneities present in the medium.

<b>Field</b>	<b>Type</b>	<b>Description</b>
Type	string	The type of heterogeneity described in this structure. This can be either a geometric object such as 'sphere' or 'block', or a description based on MRI data or some other source. The latter has yet to be clearly defined.
Pos	vector	The position of the center of the object.
Radius	scalar	If Type specifies a sphere then this field provides the radius.
Dims	vector	If Type specifies a block then this field provides the dimensions of the block. The format of the row vector is [lx ly lz] where lx, ly and lz are the x, y and z dimensions of the block, respectively.
Conc	vector	A vector of concentration data structures. This data along with spectra functions can be used to automatically fill in the Mu_a and Mu_s fields.
Mu_a	vector	The absorption coefficient(s) of the region specified. A value should be specified for each wavelength in forward model data structure.
Mu_s	scalar	The scattering coefficient of the region specified. If this field is not present the scattering coefficient is assumed to match the background.

Table B.9: The Noise data structure. This specifies the noise models to be used along with the weighting coefficients for whitening the system.

<b>Field</b>	<b>Type</b>	<b>Description</b>
SrcSNRflag	scalar	This flag specifies source relative noise. The noise standard deviation is relative to the total fluence measured at the detector for each source detector pair.
SrcSNR	scalar	The signal-to-noise-ratio in dB for the source relative noise. This is only employed if SrcSNRflag is true.
SrcNoiseSTD	matrix	The standard deviation of the noise for each source detector pair. This has the same structure as the total fluence matrix.
SrcNoise	matrix	The noise vectors that were added to the total fluence. This has the same structure as the total fluence matrix.
DetSNRflag	scalar	This flag specifies detector noise. The noise standard deviation is uniform across source-detector pairs and is relative to the maximum fluence detected for all source-detector pairs.
DetSNR	scalar	The signal-to-noise-ratio in dB for the detector noise. This is only employed if DetSNRflag is true.
DetNoiseSTD	vector	The standard deviation of the detector noise for each wavelength.
DetNoise	matrix	The noise vectors that were added to the total fluence. This has the same structure as the total fluence matrix.
ScatSNRflag	scalar	This flag specifies a noise model with a uniform variance across source-detector pairs relative to the peak scattered fluence. This is a useful model for exploring basic reconstruction performance under AWGN.
ScatSNR	scalar	The signal-to-noise-ratio in dB for the scattered field noise. This is only employed if ScatSNRflag is true.
ScatNoiseSTD	vector	The standard deviation of the scattered field noise for each wavelength.
ScatNoise	matrix	The scattered field relative noise vectors that were added to the total field. This has the same structure as the total fluence matrix.

Table B.9 The Noise data structure continued.

<b><i>Field</i></b>	<b><i>Type</i></b>	<b><i>Description</i></b>
ShotSNRflag	scalar	This flag specifies a shot noise model is to be used.
ShotSNR	vector	The signal-to-noise-ratio in dB for the shot noise.
<i>ShotNoiseSTD</i>	matrix	The standard deviation of the shot noise for each source detector pair. This has the same structure as the total fluence matrix.
<i>ShotNoise</i>	matrix	The shot noise vectors that were added to the total fluence. This has the same structure as the total fluence matrix.
ElectronicSNRflag	scalar	This flag specifies electronics noise. This still need work.
ElectronicSNR	scalar	
<i>ElectronicNoiseSTD</i>	vector	
<i>ElectronicNoise</i>	matrix	
<i>TotalVar</i>	matrix	The total variance of the noise added to the each measurement (source-detector pairing). This has the same structure as the total fluence matrix.
w	matrix	The weight vector for each wavelength. This has the same structure as the total fluence matrix.

Table B.10: The Reconstruction data structure. This specifies the reconstruction technique and any parameters required.

<b>Field</b>	<b>Type</b>	<b>Description</b>
ReconAlg	string	The reconstruction algorithm to use. Algorithms currently identified by genRecon are: 'Back Projection', 'ART', 'SIRT', 'Min. Norm', 'TSVD', 'MTSVD', 'TCG'.
xEst	matrix	The estimate(s) of the absorption function. Each column contains the reconstruction at a different wavelength.
<i>xError</i>	matrix	The absolute error function(s): $xEst - delMu_a$
<i>xResid</i>	matrix	The residual function(s).
ARTnIter	scalar	The number of iterations to compute for ART.
SIRTnIter	scalar	The number of iterations to compute for SIRT.
TSVDnSV	scalar	The number of singular values to use for TSVD.
flgNeedEconSVD	scalar	A flag specifying that the weighted forward system has changed since the last calculation of the economy SVD.
Uecon, Secon, Vecon	arrays	The economy singular value decomposition.
MTSVDnSV	scalar	The number of singular values to use for MTSVD.
MTSVDlambda	scalar	The weighting parameter for the null space in the MTSVD algorithm.
flgNeedFullSVD	scalar	A flag specifying that the weighted forward system has changed since the last calculation of the full SVD.
U, S, V	arrays	The full singular value decomposition.
xTSVD	matrix	The truncated SVD intermediate solution for use in the MTSVD algorithm.
xNS	matrix	The null space component of the reconstruction for use in the the MTSVD algorithm.
TCGnIter	scalar	The number of iterations to compute for TCG.

Table B.11: The Visualization data structure. This specifies how to visualize the results

<b>Field</b>	<b>Type</b>	<b>Description</b>
VisPlane	string	The plane in which to perform the visualization: 'X', 'Y', or 'Z'.
Type	string	The type of visualization: 'image' or 'contour'.
PlaneIndices	vector	The indices of the plane(s) to visualize. This also specifies the order in which the planes are laid out in the figure.
LayoutVector	vector	A two element vector specifying the number of axes [rows columns] to display.
Cmap	string	The colormap to use in rendering. 'bgyor' for a blue-green-yellow-orange-red colormap. 'grey' for white to black colormap.
Crange	string	The color ranging method to use. 'fixed' displays all axes over the same color range. 'auto' adjusts the range for each axes to the range of the data present in that plane.
CrangeVal	vector	Specifies the range for the color map. If this is zero and Crange is 'fixed' then the rendering function automatically sets the color range to the min and max of the whole volume. If Crange is 'auto' this is ignored.
nCLines	scalar	The number of lines to render if the 'contour' type is selected.

# Bibliography

- [1] Brian C. Wilson, Eva M. Sevick, Michael S. Patterson, and Britton Chance. Time-dependent optical spectroscopy and imaging for biomedical applications. *Proceedings of the IEEE*, 80(6):918–930, June 1992.
- [2] Gerhard Mitic, Jochen Kolzer, Johann Otto, Erich Plies, Gerald Solkner, and Wolfgang Zinth. Time-gated transillumination of biological tissues and tissuelike phantoms. *Applied Optics*, 33(28):6699–6710, 1 October 1994.
- [3] Maria Angela Franceschini, Vlad Toronov, Mattia Filiaci, Enrico Gratton, and Sergio Fantini. On-line optical imaging of the human brain with 160-ms temporal resolution. *Optics Express*, 6(3):49–57, 31 January 2000.
- [4] Britton Chance, Endla Anday, Shoko Nioka, Shuoming Zhou, Long Hong, Katherine Worden, C. Li, Murray T. Ovetsky, D. Pidikiti, and R. Thomas. A novel method for fast imaging of brain function, non-invasively, with light. *Optics Express*, 2(10):411–423, 11 May 1998.
- [5] Yu Chen, Shuoming Zhou, Shoko Nioka, Britton Chance, Maria Delivoria-Papadopoulos, and Endla Anday. Novel portable system for neonatal brain imaging. In Britton Chance, Robert R. Alfano, and Bruce J. Tromberg, editors, *Optical Tomography and Spectroscopy of Tissue III*, volume 3597, pages 262–269. Proc. SPIE, January 1999.
- [6] Peter W. Vaupel, Stanley Frinak, and Hiam I. Bicher. Heterogenous Oxygen Partial Pressaure and pH Distribution in C3H Mouse Mammary Adenocarcinoma. *Cancer Research*, 41:2008–2013, May 1981.
- [7] F. Kallinowski, K. H. Schlenger, S. Runkel, M. Kloes, M. Stohrer, P. Okunieff, and P. Vaupel. Blood Flow, Metabolism, Cellular Microenvironment, and Growth Rate of Human Tumor Xenografts. *Cancer Research*, 49:3759–3764, 15 July 1989.
- [8] David M. Brizel, Sean P. Scully, John M. Harrelson, Lester J. Layfield, Joseph M. Bean, Leonard R. Prosnitz, and Mark W. Dewhirst. Tumor Oxygenation Predicts for the Likelihood of Distant Metastases in Human Soft Tissue Sarcoma. *Cancer Research*, 56:941–943, 1 March 1996.
- [9] K. M. Case and P. F. Zweifel. *Linear Transport Theory*. Addison-Wesley, Reading, MA, 1967.

- [10] A. Ishimaru. *Wave propagation and scattering in random media*. IEEE Press, New York, NY, 1997.
- [11] R.C. Haskell, L. O. Svaasand, T.-T. Tsay, T.-C. Feng, M.S. McAdams, and B.J. Tromberg. Boundary conditions for the diffusion equation in radiative transfer. *J. Opt. Soc. Am. A*, 11(10):2727–2741, October 1994.
- [12] D. A. Boas. *Diffuse Photon Probes of Structural and Dynamical Properties of Turbid Media: Theory and Biomedical Applications*. PhD thesis, University of Pennsylvania, 1996.
- [13] S. R. Arridge. Optical tomography in medical imaging. *Inverse Problems*, 15(2):R41–R93, April 1999.
- [14] T. Durduran, A. G. Yodh, B. Chance, and D. A. Boas. Does the photon-diffusion coefficient depend on absorption. *J. Opt. Soc. Am. A*, 14(12):3358–3365, December 1997.
- [15] D. J. Durian. The diffusion coefficient depends on absorption. *Optics Letters*, 23(19):1502–1504, 1 October 1998.
- [16] Klose A. D. and Hielscher A. H. A transport-theory based reconstruction algorithm for optical tomography. In B. Chance, R. R. Alfano, and B. Tromberg, editors, *Optical Tomography and Spectroscopy of Tissue III*, volume 3597, pages 26–35. SPIE, January 1999.
- [17] M. A. O’Leary, D. A. Boas, B. Chance, and A. G. Yodh. Refraction of diffuse photon density waves. *Physical Review Letters*, 69(18):2658–2661, 2 November 1992.
- [18] J. B. Fishkin and E. Gratton. Propagation of photon-density waves in strongly scattering media containing an absorbing semi-infinite plane bounded by a straight edge. *J. Opt. Soc. Am. A*, 10(1):127–140, January 1993.
- [19] D. A. Boas, M. A. O’Leary, B. Chance, and A.G. Yodh. Scattering of diffuse photon density waves by spherical inhomogeneities within turbid media: Analytical solutions and applications. *Proc. Natl. Acad. Sci. USA*, 91:4887–4891, May 1994.
- [20] Scott A. Walker, David A. Boas, and Enrico Gratton. Photon density waves scattered from cylindrical inhomogeneities: theory and experiments. *Applied Optics*, 37(10):1935–1944, 1 April 1998.
- [21] Brian W. Pogue and Michael S. Patterson. Frequency-domain optical absorption spectroscopy of finite tissue volumes using diffusion theory. *Phys. Med. Biol.*, 39:1157–1180, 1994.
- [22] Wai fung Cheong, Scott A. Prahl, and Ashley J. Welch. A review of the optical properties of biological tissues. *IEEE Journal of Quantum Electronics*, 26(12):2166–2185, December 1990.
- [23] Frederic Bevilacqua, Dominique Piguet, Pierre Marquet, Jeffrey D. Gross, Bruce J. Tromberg, and Christian Depeursinge. *In vivo* local determination of tissue optical properties: applications to human brain. *Applied Optics*, 38(22):4939–4950, 1 August 1999.

- [24] Scott A. Prahl. Optical absorption of hemoglobin. Technical report, Oregon Medical Laser Center, <http://omlc.ogi.edu/spectra/hemoglobin/index.html>, 15 December 1999.
- [25] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, New York, NY, 1988.
- [26] M. A. O’Leary, D. A. Boas, B. Chance, and A. G. Yodh. Experimental images of heterogenous turbid media by frequency-domain diffusing-photon tomography. *Optics Letters*, 20(5):426–428, 1 March 1995.
- [27] M. A. O’Leary. *Imaging with diffuse photon density waves*. PhD thesis, University of Pennsylvania, 1996.
- [28] A. Knüttel, J.M. Schmitt, and J. R. Knutson. Spatial localization of absorbing bodies by interfering diffusing photon-density waves. *Applied Optics*, 32(4):381–389, 1 February 1993.
- [29] H.L. Graber, R.L. Barbour, and J. Chang. Algebraic reconstruction of images of a diffusive medium containing strong absorbers: comparative study of different illuminations schemes and the effect of restricted view angle. In B. Chance and R.R. Alfano, editors, *Optical Tomography: Photon Migration and Spectroscopy of Tissue and Model Media: Theory, Human Studies, and Instrumentation*, volume 2389 of *Proc. SPIE*, pages 431–447, 1995.
- [30] H. L. Graber, J. Chang, and R.L. Barbour. Imaging of multiple targets in dense scattering media. In R. Barbour, M. Carvlin, and M. Fiddy, editors, *Experimental and Numerical Methods for Solving Ill-Posed Inverse Problems: Medical and Nonmedical Applications*, volume 2570, pages 219–234. *Proc. SPIE*, 1995.
- [31] Jenghwa Chang, Raphael Aronson, Harry L. Graber, and Randall L. Barbour. Imaging diffusive media using time-independent and time-harmonic sources: dependence of image quality on imaging algorithms, target volume, weight matrix and view angles. In B. Chance and R.R. Alfano, editors, *Optical Tomography: Photon Migration and Spectroscopy of Tissue and Model Media: Theory, Human Studies, and Instrumentation*, volume 2389 of *Proc. SPIE*, pages 448–464, 1995.
- [32] Jenghwa Chang, Harry L. Graber, and Randall L. Barbour. Image reconstruction of dense scattering media from CW sources using constrained CGD and matrix rescaling technique. In B. Chance and R.R. Alfano, editors, *Optical Tomography: Photon Migration and Spectroscopy of Tissue and Model Media: Theory, Human Studies, and Instrumentation*, volume 2389 of *Proc. SPIE*, pages 682–691, 1995.
- [33] J. Chang, H. L. Graber, R. L. Barbour, and Raphael Aronson. Recovery of optical cross-section perturbations in dense scattering media by transport-theory-based imaging operators and steady-state simulated data. *Applied Optics*, 35(20):3963–3975, 10 July 1996.
- [34] Wenwu Zhu, Yao Wang, Jengwa Chang, Harry L. Graber, and Randall L. Barbour. A total least squares approach for the solution of the perturbation equations. In B. Chance and R.R. Alfano,

- editors, *Optical Tomography: Photon Migration and Spectroscopy of Tissue and Model Media: Theory, Human Studies, and Instrumentation*, volume 2389 of *Proc. SPIE*, pages 420–430, 1995.
- [35] Wenwu Zhu, Yao Wang, Yuqi Yao, Jengwa Chang, Harry L. Graber, and Randall L. Barbour. Iterative total least-squares image reconstruction algorithm for optical tomography by the conjugate gradient method. *J. Opt. Soc. Am. A*, 14(4):799–807, April 1997.
- [36] H. Jiang, K.D. Paulsen, U. L. Osterberg, B. W. Pogue, and M. S. Patterson. Optical image reconstruction using frequency domain data: simulations and experiments. *J. Opt. Soc. Am. A*, 13:253–266, October 1996.
- [37] Yuqi Yao, Yao Wang, Yaling Pei, Wenwu Zhu, and Randall L. Barbour. Simultaneous reconstruction of absorption and scattering distributions in turbid media using a Born iterative method. In R. Barbour, M. Carvlin, and M. Fiddy, editors, *Experimental and Numerical Methods for Solving Ill-Posed Inverse Problems: Medical and Nonmedical Applications*, volume 2570, pages 96–107. *Proc. SPIE*, 1995.
- [38] Yuqi Yao, Yao Wang, Yaling Pei, Wenwu Zhu, Jenhun Hu, and Randall L. Barbour. Frequency domain optical tomography in human tissue. In R. Barbour, M. Carvlin, and M. Fiddy, editors, *Experimental and Numerical Methods for Solving Ill-Posed Inverse Problems: Medical and Nonmedical Applications*, volume 2570, pages 254–266. *Proc. SPIE*, 1995.
- [39] Wenwu Zhu, Yao Wang, Yining Deng, and Randall L. Barbour. Multiresolution regularized least squares image reconstruction based on wavelet in optical tomography. In R. Barbour, M. Carvlin, and M. Fiddy, editors, *Experimental and Numerical Methods for Solving Ill-Posed Inverse Problems: Medical and Nonmedical Applications*, volume 2570, pages 186–197. *Proc. SPIE*, 1995.
- [40] Yuqi Yao, Yao Wang, Yaling Pei, Wenwu Zhu, and Randall L. Barbour. Frequency-domain optical imaging of absorption and scattering distributions by a Born iterative method. *J. Opt. Soc. Am. A*, 14(1):325–342, January 1997.
- [41] Margaret J. Eppstein, David E. Dougherty, Tamara L. Troy, and Eva M. Sevick-Muraca. Biomedical optical tomography using dynamic parameterization and Bayesian conditioning on photon migration measurements. *Applied Optics*, 38(10):2138–2150, 1 April 1999.
- [42] Misha E. Kilmer, Eric L. Miller, David A. Boas, Dana H. Brooks, Charles A. DiMarzio, and Richard J. Gaudette. Direct object localization and characterization from diffuse photon density wave data. In B. Chance, R. R. Alfano, and B. Tromberg, editors, *Optical Tomography and Spectroscopy of Tissue III*, volume 3597, pages 45–54. *Proc. SPIE*, 1999.
- [43] V. Kolehmainen, S. R. Arridge, W. R. B. Lionheart, M. Vauhkonen, and J. P. Kaipio. Recovery of region boundaries of piecewise constant coefficients of an elliptic PDE from boundary data. *Inverse Problems*, 15:1375–1391, 1999.

- [44] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, MD, 2nd edition, 1989.
- [45] J. R. Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. Technical report, <http://www.cs.cmu.edu/~jrs/jrspapers.html#cg>, July 1994.
- [46] P.C. Hansen. *Rank-Deficient and Discrete Ill-Posed Problems*. SIAM Press, Philadelphia, PA, 1998.
- [47] P.C. Hansen, T. Sekii, and H. Shibahashi. The modified truncated svd method for regularization in general form. *SIAM J. Sci. Stat. Compt.*, 13:1142–1150, 1992.
- [48] L. G. Khachiyan. A polynomial algorithm in linear programming. *Doklady Akademiia Nauk SSSR*, 244:1093–196, 1979. translated in *Soviet Mathematics Doklady* 20, 191-194 1979.
- [49] Robert G. Bland, Donald Goldfarb, and Michael J. Todd. The Ellipsoid Method: A Survey. *Operations Research*, 29(6):1039–1091, November 1981.
- [50] Stephen P. Boyd and Craig H. Barratt. *Linear Controller Design: Limits of Performance*. Prentice Hall, 1991.
- [51] Ghandi F. Ahmad, Dana H. Brooks, and Robert S. MacLeod. An Admissible Solution Approach to Inverse Electrocardiography. *Annals of Biomedical Engineering*, 26:278–292, 1998.
- [52] Richard J. Gaudette, David A. Boas, Dana H. Brooks, Charles A. DiMarzio, Misha E. Kilmer, and Eric L. Miller. A comparison of linear reconstruction techniques for 3D DPDW imaging of absorption coefficient. In B. Chance, R. R. Alfano, and B. Tromberg, editors, *Optical Tomography and Spectroscopy of Tissue III*, volume 3597, pages 55–66. Proc. SPIE, 1999.
- [53] Richard J. Gaudette, Dana H. Brooks, Charles A. DiMarzio, Misha E. Kilmer, Eric L. Miller, Tom Gaudette, and David A. Boas. A comparison study of linear reconstruction techniques for diffuse optical tomographic imaging of absorption coefficient. *Phys. Med. Bio.*, 45(4):1051–1070, April 2000.
- [54] Robert McOwen. *Partial Differential Equations*. Prentice Hall, 1996.
- [55] Matthew N. O. Sadiku. *Numerical Techniques in Electromagnetics*. CRC Press, Boca Raton, FL, USA, 1992.
- [56] David Boas, Xingde Li, Maureen O’Leary, Britton Chance, Arjun Yodh, Martin Ostermeyer, Steve Jacques, Goro Nishimura, and Scott Walker. Photon migration imaging (pmi). Technical report, [http://www.nmr.mgh.harvard.edu/DOT/DOT\\_Lab/PMI/pmi.html](http://www.nmr.mgh.harvard.edu/DOT/DOT_Lab/PMI/pmi.html), 1998.
- [57] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, 2nd edition, 1992.
- [58] A. van der Sluis and H. A. van der Vorst. Sirt- and cg-type methods for the iterative solution of sparse linear least squares problems. *Linear Algebra and It’s Applications*, 130:257–302, 1990.

- [59] D.C. Youla. *Image Recovery: Theory and Application*, chapter 2, pages 29–77. Academic Press, 1987.
- [60] Mehmet K. Ozkan, A. Murat Tekalp, and M. Ibrahim Sezan. POCS-Based Restoration of Space-Varying Blurred Images. *IEEE Transactions on Image Processing*, 3(4):450–454, July 1994.
- [61] Tony F. Chan and Chiu-Kwong Wong. Total Variation BLind Deconvolution. *IEEE transactions on Image Processing*, 7(3):370–375, March 1998.