A comparison study of linear reconstruction techniques for diffuse optical tomographic imaging of absorption coefficient

Richard J Gaudette†, Dana H Brooks†, Charles A DiMarzio‡,

Misha E Kilmer§, Eric L Miller†, Thomas Gaudette‖ and David A Boas‖

† CenSSIS, CDSP Center, 442 DA, Department of Electrical and Computer Engineering,
Northeastern University, Boston, MA 02115, USA
‡ CenSSIS, CER, 235 Forsyth Building, Department of Electrical and Computer Engineering,
Northeastern University, Boston, MA 02115, USA
§ Department of Mathematics, Tufts University, Boston, MA, USA
‖ CIMIT and Mass General Hospital NMR Center, Harvard Medical School, Boston, MA, USA

Abstract. We compare, through simulations, the performance of four linear algorithms for diffuse optical tomographic reconstruction of the three-dimensional distribution of absorption coefficient within a highly scattering medium using the diffuse photon density wave approximation. The simulation geometry consisted of a coplanar array of sources and detectors at the boundary of a half-space medium. The forward solution matrix is both underdetermined, because we estimate many more absorption coefficient voxels than we have measurements, and ill-conditioned, due to the ill-posedness of the inverse problem. We compare two algebraic techniques, ART and SIRT, and two subspace techniques, the truncated SVD and CG algorithms. We compare three-dimensional reconstructions with two-dimensional reconstructions which assume all inhomogeneities are confined to a known horizontal slab, and we consider two 'object-based' error metrics in addition to mean square reconstruction error. We include a comparison using simulated data generated using a different FDFD method with the same inversion algorithms to indicate how our conclusions are affected in a somewhat more realistic scenario. Our results show that the subspace techniques are superior to the algebraic techniques in localization of inhomogeneities and estimation of their amplitude, that two-dimensional reconstructions are sensitive to underestimation of the object depth, and that an error measure based on a location parameter can be a useful complement to mean squared error.

1. Introduction

Over the past 10 years there has been considerable research into the use of near-infrared light to image inside the human body, a technique known variously as diffuse optical tomography (DOT), photon migration imaging (PMI) and diffuse photon density wave (DPDW) imaging (see Arridge (1999) for a recent review). One of the primary goals of this research is to image the distribution of the optical absorption coefficient, which at near-infrared wavelengths (700–900 nm) is primarily influenced by haemoglobin in its various forms. Thus a mapping of the density of haemoglobin can be inferred from an image of the absorption coefficient. Recently considerable work has been done developing both linear and nonlinear physical models and back-propagation techniques for constructing these images. Nonlinear techniques such as those developed by Arridge *et al* (1992) and Jiang *et al* (1996) are attractive because they minimize the number of assumptions regarding both the medium and the physics, but they are computationally very expensive. Back-propagation, which has been explored by

groups such as Colak *et al* (1997) and Matson *et al* (1997), is computationally economical but sensitive to noise and does not deal well with multiple absorbing objects. Linear perturbation techniques attempt to specify the imaging problem as a perturbation to a known or estimated background medium and thus pose the relationship between the absorption coefficient and the measured data as a linear system of equations. Linear methods include both the Born and Rytov approximations, and have been explored by O'Leary *et al* (1995) and Chang *et al* (1995), among others.

The linear inverse problem is ill-posed due to absorption and scattering in the medium; the mathematical consequence is that the matrix representing the forward problem will be ill-conditioned. Thus some type of regularization or stabilization technique is required to obtain physically realistic results. Once a linear forward model is assumed, the inverse problem can be approached in a linear algebraic context. A number of linear algebraic techniques have been applied to linear DOT imaging. However, there have been few, if any, systematic comparisons of these methods as applied to DOT. In this paper we present such a comparison. Specifically, we examined the performance of two representatives each of two classes of linear reconstruction techniques, algebraic and subspace methods, in a simple, but clinically relevant reflective imaging geometry. The algebraic techniques we examined were the algebraic reconstruction technique (ART) and the simultaneous iterative reconstruction technique (SIRT) (Kak and Slaney 1988), and the subspace techniques were the truncated singular value decomposition (TSVD) and truncated conjugate gradient (TCG) algorithm (Hansen 1998). The algebraic methods have been widely used for DOT as well as many other tomographic problems, are simple to understand and program, and can be relatively computationally efficient for certain classes of problems. The subspace techniques are often regarded as being more accurate in many applications of linear algebra. The main purpose of this investigation was to quantify the performance of these two classes of reconstruction methods in the DOT framework. In this context we address some related issues such as the choice of, and sensitivity to, the required regularization parameter which controls the amount of regularization applied in a given method. We point out that although the discussion in this paper is framed entirely in terms of a linear forward model, many nonlinear algorithms depend on a succession of linear problems which may need to be regularized, so we believe that the results obtained here may well be more widely applicable. To give an indication of how the methods perform when a different, nonlinear approach is used to generate the data, we include a brief comparison of linear reconstructions obtained from data simulated by a frequency domain finite difference (FDFD) forward model.

In many DOT imaging scenarios, three-dimensional volumetric reconstruction of a region of tissue under study is computationally challenging. Additionally, such reconstructions often require more parameters to be estimated than the number of measurements available, so that the resulting problem is mathematically underdetermined. Many investigators have dealt with this problem by assuming that all absorption inhomogeneities lay in a known planar slice in the volume (Cheng and Boas 1998, Chang *et al* 1995, Jiang *et al* 1996). By assuming that the background values are known or can be estimated, the number of unknowns is dramatically reduced and, geometrically, the inverse problem is reduced to the estimation of absorption coefficients in a two-dimensional slice. The effect of this assumption on the accuracy of the inverse solution has frequently not been carefully tested, especially in the case where the assumption is itself not accurate. A second purpose of the current work, then, is to quantify the performance of full three-dimensional reconstructions in comparison with two-dimensional planar reconstructions.

In a simulation scenario where the goal is to find and quantify a region of absorption inhomogeneity in an otherwise homogeneous three-dimensional volume, it is not clear that simple error measures such as the mean squared error over the volume between a reconstruction and

the true distribution is the best measure to use; in applications, for instance, one may be more interested in estimating the location or peak amplitude of the inhomogeneity. As a secondary objective of this work, then, we compare three different error measures on our reconstructions to determine whether, or when, simple mean square error measures are adequate.

In summary, then, this work has three goals:

- to compare algebraic and subspace reconstructions in the same imaging scenario;
- to compare full three-dimensional volumetric reconstructions to reconstructions which assume an inhomogeneity is confined to a known two-dimensional slice; and
- to compare mean squared error with more object-oriented *ad hoc* error measures.

The paper is structured as follows. In section 2 we describe the simulations and propagation model we used for the study. In section 3 we present the details of the four reconstruction algorithms we compared. In section 4 we show both qualitative and quantitative comparisons of the reconstruction techniques under the simulation scenario described, and in section 5 we describe our observations and conclusions.

2. Methods

All of the results presented in this study were calculated from computer simulations of detector measurements. Our study simulated a half-space region of diffuse media with a single spherical inhomogeneity embedded in the region. Using computer simulations we were easily able to exercise precise control over parameters of interest such as signal-to-noise ratio and object position.

2.1. Discrete DPDW model

A model of light propagation in a highly scattering medium is necessary both to compute the simulated fluence at the detectors and to map the fluence values back to the spatial absorption function. The solution to either problem provides a forward model; however, in studying inverse solutions the forward model used for the first computation (i.e. the forward computation of simulated fluence) may not be the same as that used for the second computation (i.e. the inverse solution). In this subsection we describe briefly the mathematical approach we used to formulate our forward models (for more detail see Kak and Slaney (1988) and O'Leary (1996)); in subsequent subsections we describe the distinct forward models we computed based on this general mathematical model.

One useful and commonly employed model for the photon fluence in a highly scattering medium is the Helmholtz frequency domain diffusion equation (O'Leary *et al* 1995, Ishimaru 1997, Fishkin and Gratton 1993)

$$\left(\nabla^2 + \frac{j\omega - v\mu_a(r)}{D}\right)\phi(r) = \frac{-v}{D}S(r) \tag{1}$$

where $\phi(r)$ is the photon fluence at position r, $\mu_a(r)$ is the spatially varying absorption coefficient, v is the electromagnetic propagation velocity in the medium, ω is the frequency in rad s^{-1} and $S(r)$ is the source function. D is the diffusion coefficient, given by

$$D = \frac{v}{3\mu_s'}$$

where μ_s' is the reduced scattering coefficient. Note that this equation only considers spatial variations in the absorption coefficient and is functionally different if D is spatially varying.

Using a perturbation approach we first rewrite (1) as

$$\left(\nabla^2 + \frac{\mathrm{j}\omega - v\mu_a^0}{D} - \frac{v\Delta\mu_a(r)}{D}\right)(\phi_i(r) + \phi_s(r)) = \frac{-v}{D}S(r) \tag{2}$$

where we have written the absorption function as a sum of the background absorption, μ_a^0, and a spatially varying perturbation $\Delta\mu_a(r)$, and the fluence as the sum of the incident field, $\phi_i(r)$, due to the source acting on the background medium and a scattered fluence, $\phi_s(r)$, due to the inhomogeneities.

Subtracting the homogeneous medium equation (described by (1) with $\mu_a(r) = \mu_a^0$) from (2) we are left with the equation for the scattered fluence

$$[\nabla^2 + k_0^2]\phi_s(r) = -\Delta k^2(r)(\phi_i(r) + \phi_s(r)) \tag{3}$$

where $k_0^2 = (\mathrm{j}\omega - v\mu_a^0)/D$ and $\Delta k^2(r) = (v/D)\Delta\mu_a(r)$. Using a Green's function approach and the assumption that $\phi_i(r) \gg \phi_s(r)$, we can approximate the scattered field by the integral equation

$$\phi_s(r) \approx \int_V G(r, r')\phi_i(r')\Delta k^2(r')\,\mathrm{d}r' \tag{4}$$

providing a linear relationship between the scattered fluence and the absorption perturbation. The Green's function for the half-space medium is

$$G(r, r') = \frac{-1}{4\pi|r - r'|}\mathrm{e}^{\mathrm{j}k_0|r-r'|} + \frac{1}{4\pi|r - r_i'|}\mathrm{e}^{\mathrm{j}k_0|r-r_i'|} \tag{5}$$

where r_i' is the image point of r'. Reflection across the boundary is discussed below.

This linearization, which is based on ignoring the contribution of the scattered field on the right side of (3), is known as the first Born approximation. Physically it 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 field from one part of an inhomogeneity on the field incident on another part.†

Thus, for each source we calculate the incident field everywhere in the domain using the Green's function and then calculate the scattered field present at each detector using (4). Since (4) gives a linear expression relating $\Delta k^2(r)$, the spatially varying perturbation of the background medium, to the measured scattered field, $\phi_s(r)$, we can discretize the problem into a system of linear equations of the form

$$\mathbf{Ax} = \mathbf{b} \tag{6}$$

where each row of $\mathbf{A}$ corresponds to a different source–detector pairing and the columns of $\mathbf{A}$ index small volumes (voxels) within the region of interest V. The elements of $\mathbf{A}$ are integrals of the Green's function multiplied by the incident field for the specified source detector pair over the corresponding voxel. The vector $\mathbf{x}$ is a discretization of $\Delta k^2(r)$ for each voxel and $\mathbf{b}$ is a vector whose elements are the detector scattered field measurements for each source–detector pair. To ensure that the reconstructed $\Delta k^2(r)$ is real we can separate the real and imaginary parts of $\mathbf{A} = \mathbf{A_r} + \mathrm{j}\mathbf{A_i}$ and $\mathbf{b} = \mathbf{b_r} + \mathrm{j}\mathbf{b_i}$ and stack them as separate rows in (6). This is expressed as

$$\tilde{\mathbf{A}}\mathbf{x} = \begin{bmatrix} \mathbf{A_r} \\ \mathbf{A_i} \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{b_r} \\ \mathbf{b_i} \end{bmatrix}.$$

We evaluated the integrals in (4) by using a pulse basis for the functions $G(r, r')$ and $\phi_i(r)$.

† We also note that to use this approach for inverse calculations one needs to be able to measure the scattered field alone. In simulations one simply computes only the scattered field; in practice, of course, this is usually impossible. Hence, in practice some approximate homogeneous field is estimated and subtracted from measurements.

We modelled our sources as point sources one mean free path length into the medium (Haskell *et al* 1994). The air–tissue boundary for the incident field and in (4) is accounted for using an extrapolated boundary condition (Haskell *et al* 1994). This extrapolated boundary condition specifies a zero Dirichlet boundary at a distance out of the medium specified by

$$z_{\text{bnd}} = \frac{2}{3\mu'_s} \left(\frac{1 + R_{\text{eff}}}{1 - R_{\text{eff}}} \right)$$

where R_{eff} is the effective reflection coefficient from inside the medium. For a value of R_{eff} we used 0.4664, which we computed by linearly interpolating for an index of refraction of 1.37 (commonly used to represent tissue) from table 2 in Haskell *et al* (1994). We accounted for the boundary condition by reflecting both the modelled sources and the equivalent source term $\phi_i(r)\Delta k^2(r)$ across the boundary. This resulted in a Green's function for the half-space given by (5).

In this study we used the approach just described both to generate simulated measured data and to model the forward system upon which the inversion was based. However, as described below, the specific forward model was different in the two cases because the spatial discretization of the model was different.

2.2. Computational geometry

We modelled 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 centred 2.5 cm deep in the tissue, as shown in figure 1.
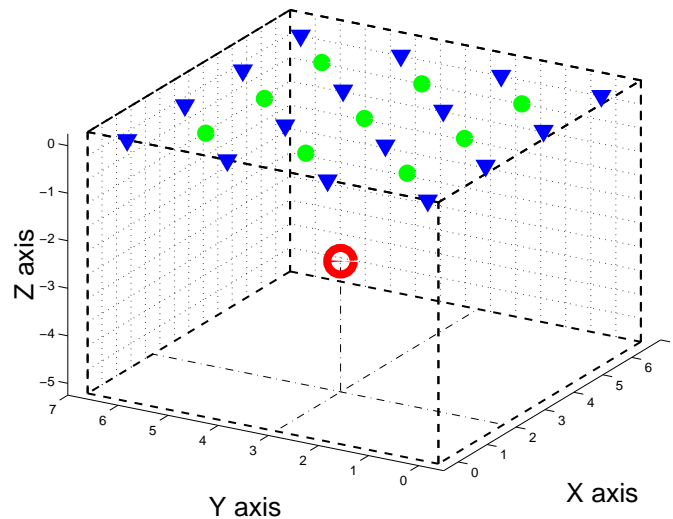
Figure 1. A schematic of the geometric layout of the simulations in this study. The sources and detectors lie along the tissue–air boundary ($z = 0$ cm). The entire volume was 7 cm × 7 cm × 5.5 cm deep. The absorption anomaly is 2.5 cm into the medium at $x = 2$ cm and $y = 3$ cm. The grid size used to discretize the volume for some of the inverse computations is illustrated by the grid shown in the figure.
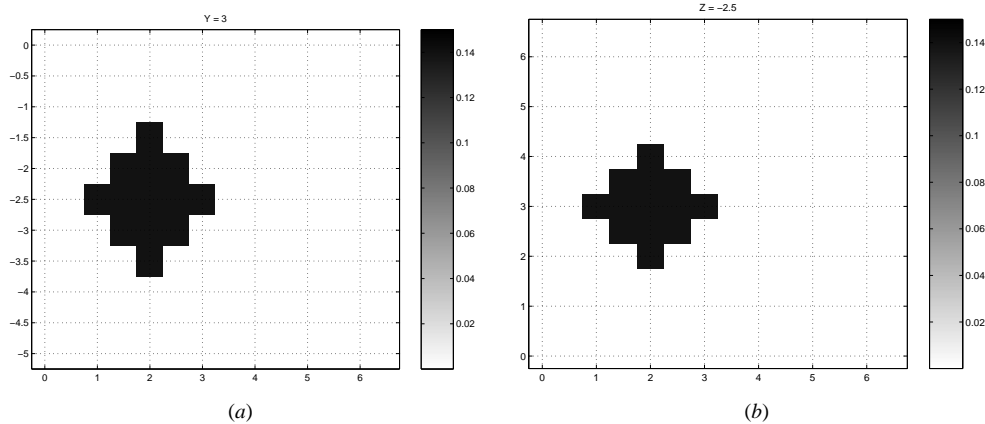
Figure 2. (*a*) A vertical plane of the absorption function showing the true object image. This plane is through the centre of the object ($Y = 3$ cm). (*b*) A horizontal plane of the absorption function showing the true object image. This plane is through the centre of the object ($Z = 2.5$ cm).

2.2.1. Source and detector parameters. The reflective source–detector geometry we examined simulated a single planar array placed on the surface of the diffuse medium.† 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 1. Shown in figure 2 are two examples of a single slice through the volume of the true absorption perturbation. Figure 2(*a*) shows a vertical slice through the centre 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 2(*b*) shows a horizontal slice though the centre of the anomaly in the X–Y plane. We use this representation for two-dimensional slice reconstructions.

2.2.2. 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 this region because of the compact support of the object.

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 under the source detector array was discretized uniformly for the forward model used in the inverse reconstructions. Specifically, the volume modelled, as shown in figure 1, was a cube 7 cm $\times$ 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. In addition to the full three-dimensional discretization as described in section 1, we also examined reconstructions using only a single horizontal plane of voxels centred on the anomaly. In this method the source terms, the unknowns in the inverse solution, were assumed to be zero outside the chosen horizontal slice during the inverse computation. Thus the size of the relevant linear system and the required computations are greatly reduced.

† This geometry models a measurement system currently being constructed in one of our laboratories.

In all our reconstructions a finer discretization was used in forward calculations and a coarser one used for the inverse reconstruction. Additionally, in some of the single-slice 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. In the sequel, when we compare reconstruction accuracy using the full three-dimensional forward model and the single slice forward model for inverse solutions, we will refer to the former case as F3D (full 3D) and the latter as 2DS (2D slice).

2.3. Medium optical parameters

The diffuse medium optical parameters for our experiments were selected to simulate human tissue. Specifically, the background scattering coefficient, μ_s, was taken as 100 cm^{-1} and the mean cosine of the scattering angle as 0.9, resulting in a reduced scattering coefficient, μ_s', of 10 cm^{-1}. The background absorption coefficient was derived from an assumed haemoglobin concentration of 0.0920 mM and an oxygen saturation of 90%, resulting in an absorption coefficient of 0.041 cm^{-1} at a wavelength of 780 nm. The absorption anomaly simulated a haemoglobin concentration of 0.368 mM and an oxygen saturation of 60%, resulting in an absorption coefficient of 0.18 cm^{-1} or 0.139 cm^{-1} above the background.

2.4. Noise model

We modelled 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 relative to the total fluence at the detector. Specifically the noise standard deviation used in these results was calculated using the formula

$$\gamma(i) = \phi(i) 10^{-\text{SNR}/20} \tag{7}$$

where $\gamma(i)$ is the noise standard deviation for the ith source–detector pairing, $\phi(i)$ is the total fluence (incident + scattered) computed at the detector and SNR is the signal-to-noise ratio in dB. This noise model is based on the assumption that shot noise from the sources is the dominant source of measurement noise. It is clear from (7) that 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

$$\mathbf{DAx} = \mathbf{Db}, \quad \mathbf{D} = \text{diag}(\gamma(i)^{-1}).$$

2.5. Finite difference frequency domain forward model

A full analysis of a more complete forward model is beyond the scope of this paper. However, as mentioned earlier, we felt it was of some interest to give an indication of how our results might change when the forward data were generated by a completely different, and more physically accurate, numerical technique, but the inverse solutions still used the Born forward model. To this end we used a standard finite difference frequency domain (FDFD) forward solution (Saduiku 1992) for the same scenario described above to generate the simulated measurement data. These 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 centred on the inhomogeneity. The Z domain was aligned so that the computational boundary was aligned with the extrapolated

boundary position previously described. All of the computational boundaries implemented a zero Dirichlet boundary condition.

3. Reconstruction algorithms

There are two important characteristics of this inverse problem that reconstruction algorithms must take into account. First, the 3D linear system model (6) is typically very underdetermined. Examples we present in this study have 16 detectors and 9 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 $\mathbf{A}$. This is typical of integral operators with smoothing kernels (Hansen 1998). Therefore the solution must be regularized if it is not to be dominated by reconstruction noise.

3.1. Algebraic techniques

The types of algorithms we examined for solving the linear system (6) fall into two classes, algebraic techniques and subspace techniques. Members of the first class solve (6) 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 (Kak and Slaney 1988). 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

$$\hat{\mathbf{x}}_{j+1} = \hat{\mathbf{x}}_j + w \frac{b_i - \mathbf{a}_i \hat{\mathbf{x}}_j}{\mathbf{a}_i \mathbf{a}_i^{\mathbf{T}}} \mathbf{a}_i^{\mathbf{T}} \qquad j = 0, 1, \ldots \quad i = (j \bmod 2m) + 1$$

where $\hat{\mathbf{x}}_j$ is the jth estimate of the object function, $\mathbf{a}_i$ is the ith row of the $2m \times n$ matrix $\tilde{\mathbf{A}}$, and b_i is the ith 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 turn, 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

$$\hat{\mathbf{x}}_{j+1} = \hat{\mathbf{x}}_j + w \frac{1}{2m} \sum_{i=1}^{2m} \frac{b_i - \mathbf{a}_i \hat{\mathbf{x}}_j}{\mathbf{a}_i \mathbf{a}_i^{\mathbf{T}}} \mathbf{a}_i^{\mathbf{T}}.$$

SIRT has generally been observed to generate smoother reconstructions of the object 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 to the initial guess (Kak and Slaney 1988). Regularization is accomplished by limiting the number of iterations in both algebraic techniques. 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 and well beyond the scope of this paper (van der Sluis and van der Vorst 1990). In order to study this parameter choice in the context of this paper, 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 do not appreciably alter the best performance metrics achieved. Rather, the rate (number of iterations) at which the best performance measure is achieved varies, 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 with the best performance achieved with $w = 1$). At higher SNR a relaxation parameter closer to 1 provided the best performance. Again we observed an interaction between relaxation parameter and optimum number of iterations. In no case did we observe any behaviour of the algebraic algorithms over this range of relaxation parameters that would change the basic conclusions we draw 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 1 for all of the comparisons reported below.

3.2. 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 $\tilde{\mathbf{A}}$. The SVD of the system matrix is given by

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

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

$$\sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant \sigma_r \qquad \sigma_{r+1}, \sigma_{r+2}, \ldots, \sigma_{\min(m,n)} = 0$$

where r is the rank of $\tilde{\mathbf{A}}$. 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, which displays a range of seven 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 $\tilde{\mathbf{A}}\mathbf{x} = \mathbf{b}$. Mathematically this can be written as

$$\hat{\mathbf{x}} = \mathbf{V}_t \mathbf{\Sigma}_t^{-1} \mathbf{U}_t^{\mathbf{T}} \mathbf{b}$$

where $\mathbf{V}_t$ and $\mathbf{U}_t$ are the first t columns of $\mathbf{V}$ and $\mathbf{U}$ respectively, and $\mathbf{\Sigma}_t^{-1}$ is the inverse of the square diagonal submatrix of the largest t singular values. The value t, the truncation parameter, controls the amount of regularization in the inverse. More information on the SVD and TSVD can be found in Golub and Van Loan (1989) and Hansen (1998). In this context we note that in figure 3 there is no useful gap in the spectrum 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.

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. For ill-conditioned systems we can regularize by stopping, or truncating, the iterations before we reach full convergence (Hansen 1998). Thus for TCG the number of iterations computed controls the amount of regularization. For most measurement geometries the system matrix $\tilde{\mathbf{A}}$ 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 technique to the normal equations given by

$$\tilde{\mathbf{A}}^{\mathbf{T}}\tilde{\mathbf{A}}\mathbf{x} = \tilde{\mathbf{A}}^{\mathbf{T}}\mathbf{b}.$$
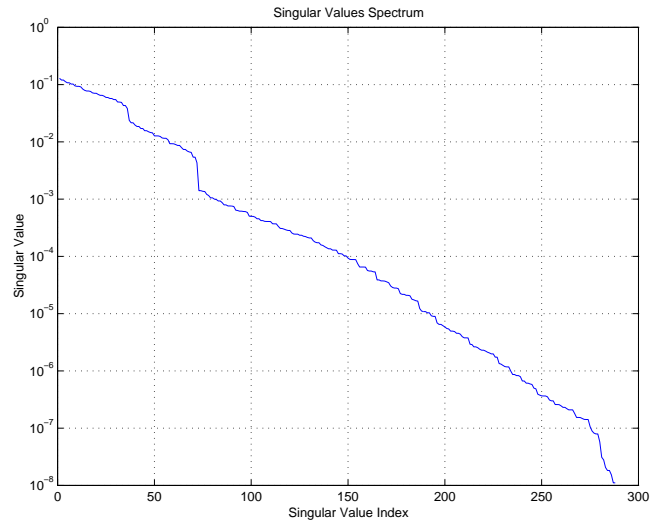
Figure 3. Singular value spectrum of a typical forward DOT forward matrix.

Table 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.

$\mathbf{b_p} = \tilde{\mathbf{A}}^\mathbf{T}\mathbf{b}; \quad \mathbf{d} = \mathbf{r} = \mathbf{b_p}; \quad \delta = \mathbf{r}^\mathbf{T} \cdot \mathbf{r}; \quad \delta_0 = \delta; \quad \hat{\mathbf{x}} = \mathbf{0};$

for $j = 1 : n_{\text{iter}}$

$\qquad \mathbf{q} = \tilde{\mathbf{A}}^\mathbf{T}(\tilde{\mathbf{A}}\mathbf{d})$

$\qquad \alpha = \delta/(\mathbf{d}^\mathbf{T} \cdot \mathbf{q})$

$\qquad \hat{\mathbf{x}} = \hat{\mathbf{x}} + \alpha\mathbf{d}$

$\qquad \mathbf{r} = \mathbf{r} - \alpha\mathbf{q}$

$\qquad \delta_0 = \delta$

$\qquad \delta = \mathbf{r}^\mathbf{T} \cdot \mathbf{r}$

$\qquad \mathbf{d} = \mathbf{r} + \frac{\delta}{\delta_0}\mathbf{d}$

end

For computational reasons we do this without explicitly computing the normal equations, as given in table 1 (derived from Shewchuck's (1994) implementation).

3.3. 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 is 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 (Hansen 1998), 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 4. We generated this graph by plotting the
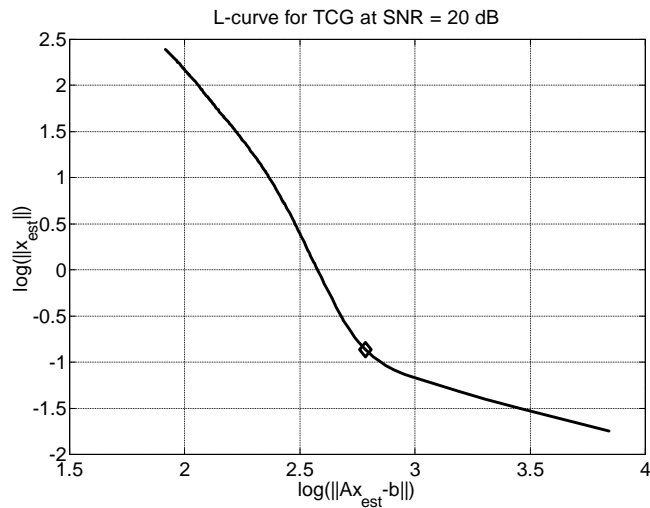
Figure 4. L-curve for a TCG reconstruction at a signal-to-noise ratio of 20 dB. The diamond at the corner of the graph identifies the point that was selected as the corner of the L-curve.

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 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.

As mentioned earlier, 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.

4. Results and observations

In this section we present results and observations of applying the four reconstruction techniques described in section 3 to the simulation scenario described in section 2. 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 regularization parameter.
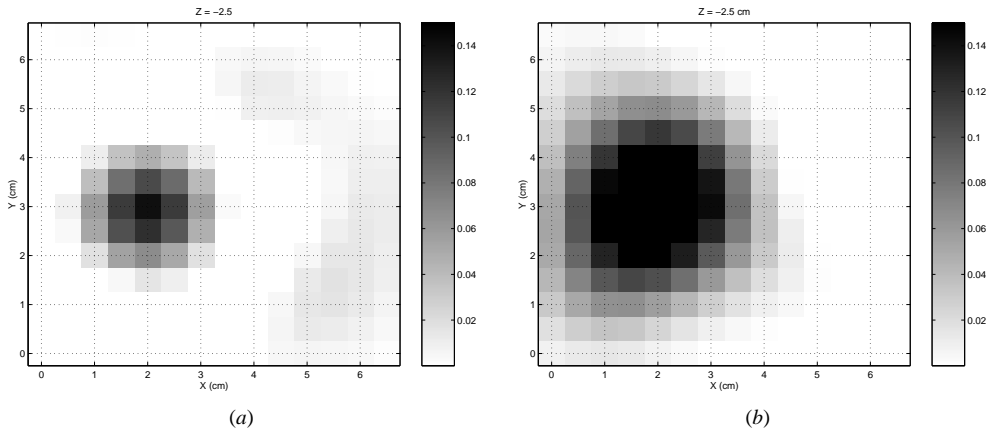
4.1. Example reconstructions

Figure 5 shows a set of images of the reconstructed absorption coefficient in a single vertical $(X-Z)$ plane through the centre of the object ($Y = 3$ cm) at a signal-to-noise ratio of 20 dB. The true object is shown in figure 2(a). The F3D model was used for the inversion with a 0.5 cm grid in all dimensions. Comparing the true absorption image with the reconstructions, it is evident that all of the reconstructions show a roughly correct image of the object and that all of the reconstruction techniques underestimate the depth and amplitude of the object. Note that the scale for the reconstructed images is less than that for the true absorption function. Figure 5(a) is the reconstruction using a single iteration of ART (i.e. one iteration cycle through all 288 rows
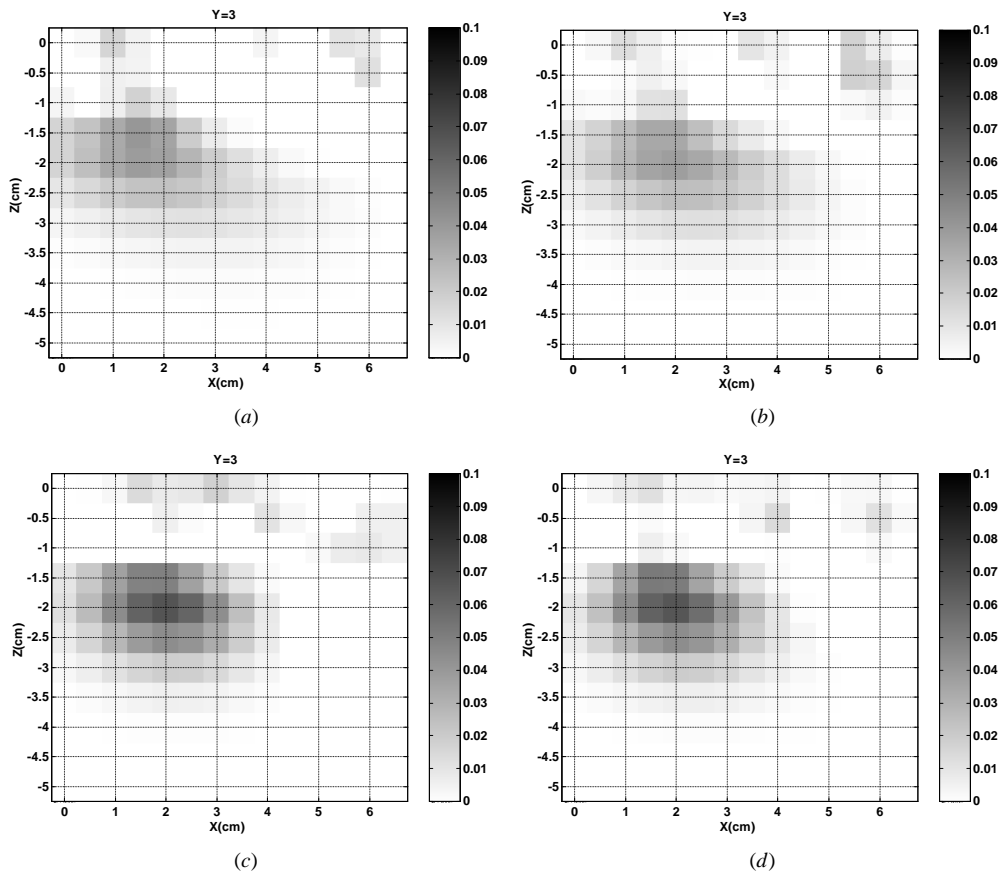
Figure 5. Reconstruction examples for all four reconstruction techniques at a 20 dB SNR. Each image shows a vertical plane through the centre of the absorption anomaly ($Y = 3$ cm): (*a*) is the ART algorithm result using one iteration, (*b*) is the SIRT algorithm result using 26 iterations, (*c*) is the TSVD algorithm result using 56 singular values and (*d*) is the TCG algorithm using 12 iterations.

of the system matrix). An object centred approximately at $X = 1.5$ cm and $Z = -1.75$ cm with an amplitude of around 0.04 cm^{-1} is visible. The object does not appear to have the same area as the true absorption function nor does the reconstruction extend as deep. Figure 5(*b*) shows a reconstruction using the SIRT algorithm with 26 iterations. Qualitatively this reconstruction appears very similar to the ART reconstruction.

Figure 5(*c*) shows a reconstruction using the TSVD algorithm employing 56 singular values. This reconstruction clearly produces a larger-amplitude absorption coefficient than the algebraic techniques. Additionally, the object centre appears at about 2 cm depth, about 0.25 cm deeper than the algebraic techniques. Thus in both aspects, amplitude and position, it is closer to the true absorption function. Figure 5(*d*) displays the reconstruction for the TCG technique using 12 iterations, which is qualitatively very similar to 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 6. 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 5 these reconstructions are shown in a X-Y plane.

Figure 6 shows a pair of 2DS reconstructions using the TCG algorithm, again at an SNR of 20 dB. Figure 6(*a*) was generated using a 2.0 cm width reconstruction plane with the number of iterations again chosen by the L-curve method. Figure 6(*b*) was generated using a 0.5 cm width reconstruction plane which *underestimates* the true width of the object. 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 regularization difficult. Therefore for this case we selected the number of iterations that minimized the actual mean square error. For both cases the centre of the reconstruction was the true centre of the absorption anomaly. We note that the 2DS reconstruction shown is reasonably accurate in the first case but significantly overestimates the object size in the second case.

4.2. Performance measures

As noted in section 1, 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 object directly. More precisely, we used the following three error measures:

(a) *Mean squared error (MSE).* The first performance measure we evaluated was mean square error, given by the expression

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

(b) *Object centroid error (OCE).* We evaluated the error in the estimated position of the reconstructed object 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 neighbour of a current 'object voxel' whose amplitude was above the threshold. The search finished when there were no more voxels which were neighbours of voxels classified as object voxels whose amplitude was above the threshold. Once the object was detected, we computed its centroid
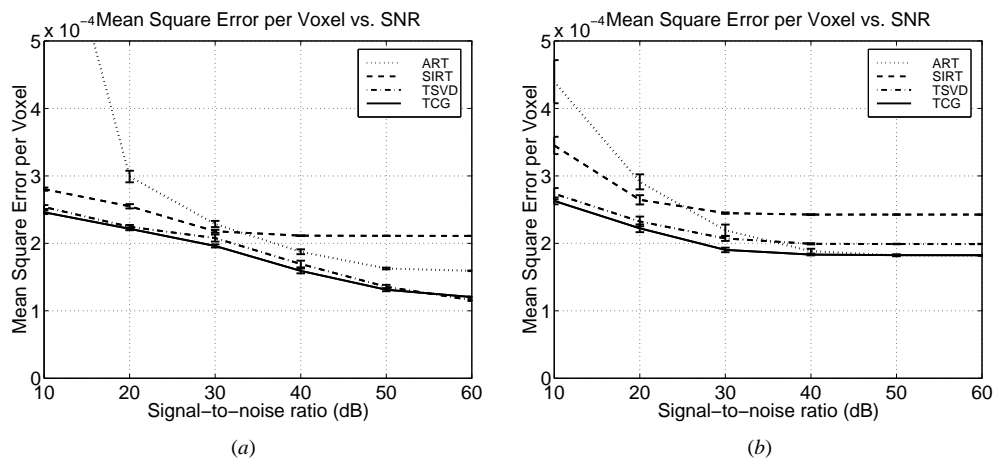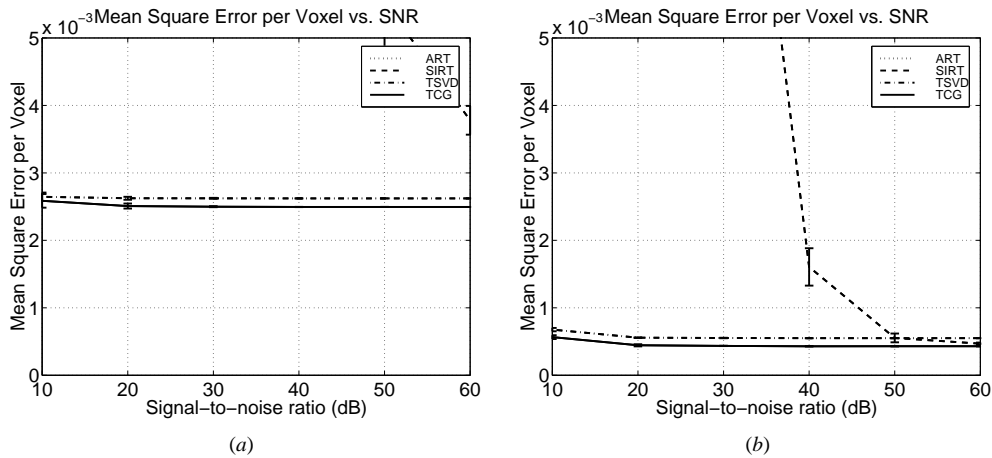
Figure 7. MSE for the F3D reconstructions. (*a*) The 0.5 cm vertical cell size reconstruction. (*b*) The 2.0 cm vertical cell size reconstruction. Data points are means over 10 realizations and error bars show ±1 SD over the realizations. See text for details.

as the weighted average of the position of the object voxels, with the amplitude of these voxels as the weights.

(c) *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 object. We simply calculated the difference between the maximum value of the true object and the maximum value of the reconstructed object over the support of the true object. 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 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 for 10 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 7 shows the graphs of the MSE versus SNR for all four F3D reconstruction techniques with 7(*a*) showing the 0.5 cm grid discretization and 7(*b*) showing the 2.0 cm discretization in the vertical direction. Each curve in the figure shows the mean value of the MSE over the 10 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

Figure 8. MSE for the 2DS reconstructions. (*a*) The 0.5 cm vertical cell size reconstruction. (*b*) The 2.0 cm vertical cell size reconstruction. Data points are means over 10 realizations and error bars show ±1 SD over the realizations. See text for details.
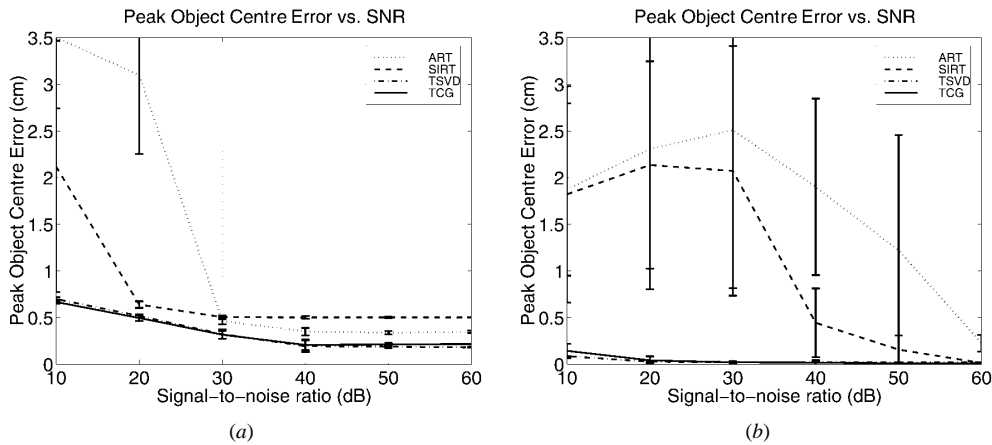
**Figure 9.** (*a*) OCE for a F3D reconstruction with a 0.5 cm vertical step size. (*b*) OCE for 2.0 cm thick planar reconstruction.

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 40 dB signal-to-noise ratio there is no improvement in the MSE with higher SNR for this case.

Figure 8 shows the MSE curves for the equivalent 2DS reconstructions. Note that the vertical scale of these graphs is 10 times larger than for the previous two graphs. Comparing 2DS and 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 9(*a*) shows the OCE for the four reconstruction techniques for an F3D with a 0.5 cm vertical step size and figure 9(*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
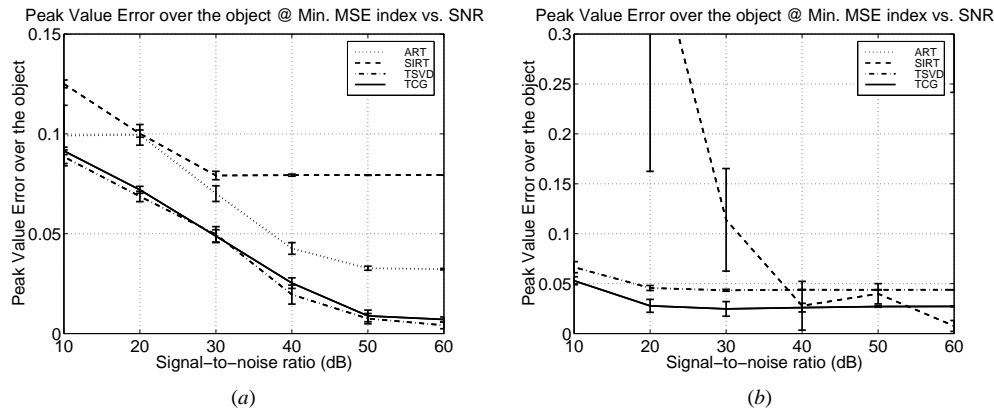
Figure 10. (*a*) AE for the F3D 0.5 cm step size reconstruction. (*b*) AE for the 2.0 cm 2DS reconstruction. Note that for the 2DS case the curve for ART is off the scale of the graph.

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 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 10(*a*) shows the AE performance for the F3D 0.5 cm reconstruction while figure 10(*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 object. In other words, despite the low value of the AE measure, there was very little resemblance between the true object absorption function and those generated by ART at a 10 dB signal-to-noise ratio, as indicated by the MSE curves in figure 7(*a*).
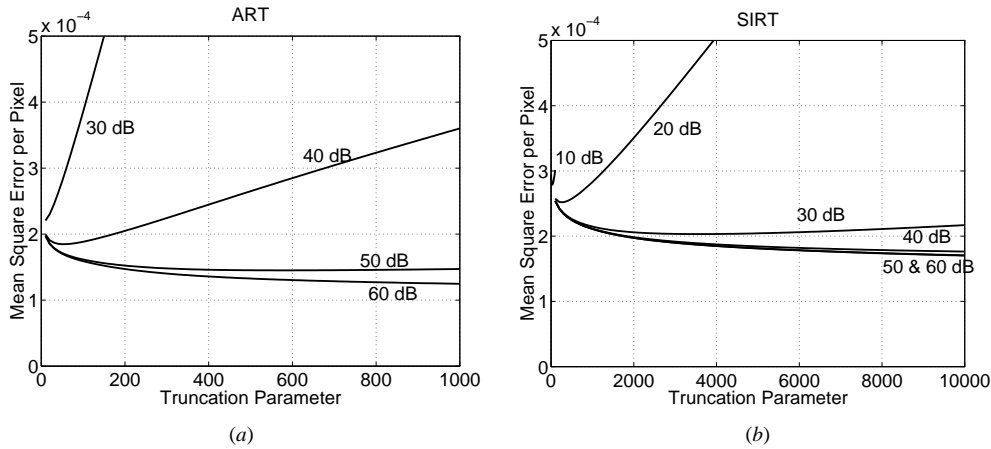
4.3. Regularization parameter sensitivity

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 11. 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.
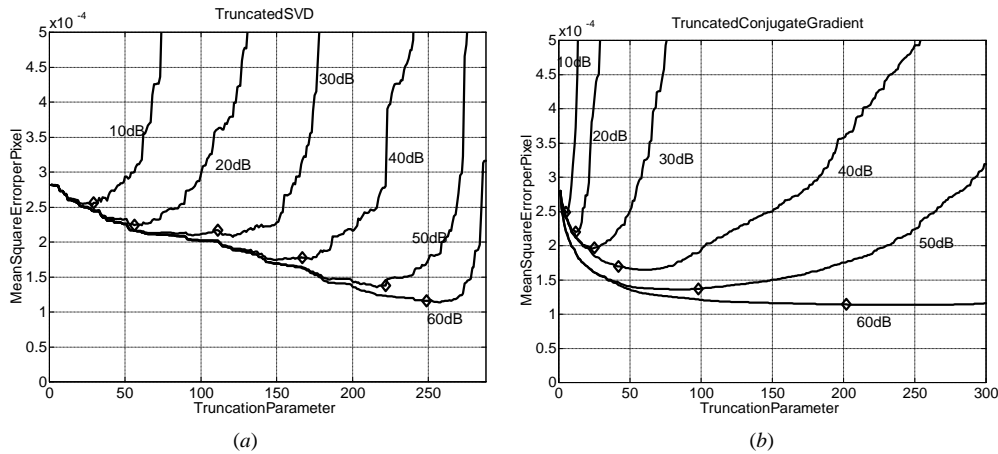
Figure 11. Regularization parameter sensitivity for the two subspace reconstruction techniques over a range of SNRs. One MSE curve is shown for each SNR. The diamond on each curve displays the regularization parameter we chose (without knowledge of the MSE curve) by simply manually selecting the corner of the L-curve. (*a*) The results from the TSVD algorithm. (*b*) The results from the TCG algorithm.


Figure 12. Regularization parameter sensitivity for the two algebraic reconstruction techniques over a range of SNRs. One MSE curve is shown for each SNR. (*a*) The results from the ART algorithm, the 10 and 20 dB curves are out of the range of the plot. (*b*) Results from the SIRT algorithm.

As a comparison of the regularization sensitivity for the algebraic techniques we show a plot of the MSE versus regularization parameter in figure 12. 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 behaviour is evident.

4.4. FDFD forward model results

As described above, we also generated data with a FDFD forward model to feed into our comparison of inversion algorithms. We performed an F3D 0.5 cm reconstruction using all
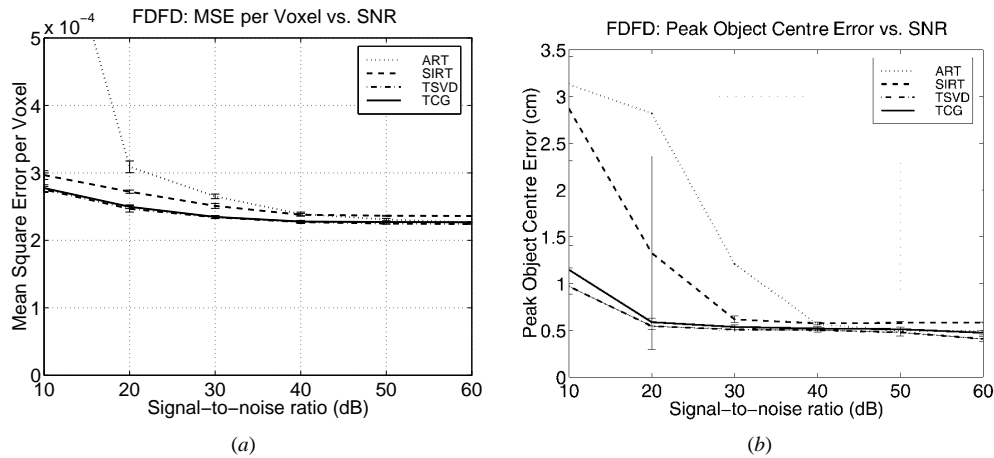
Figure 13. F3D, 0.5 cm vertical cell size performance with simulated data generated using a FDFD forward model. (*a*) The MSE for the four reconstruction algorithms versus SNR. (*b*) OCE error for these cases. The standard deviation bars for ART in the OCE plot are removed because they were large enough to obscure the other algorithms.

four algorithms over the range of SNRs previously described. Results are shown for MSE and OCE error in figure 13. At low SNRs (below 40 dB) we observe performance similar to that displayed when a first Born forward model was used to generate the simulated forward data. At higher SNRs the performance of all of the algorithms levels off at a higher error than with the Born generated data, with the subspace algorithms slightly outperforming the algebraic algorithms.

5. Conclusions and future work

In this paper we have presented a comparison of four of the most commonly used linear reconstruction techniques for diffuse photon density wave 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 6 and 8.

By using the OCE error measure we identified further differences in performance between the subspace and algebraic techniques that were not evident with only an 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 an 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 nonlinear 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.

In future work we plan to repeat this comparison, still using a Born forward model for the inverse, with measured experimental data, so that we can explore the effects of increased model mismatch on the ability of each method to produce quality reconstructions. Another area of work we will investigate is the development of a scheme with a spatially varying regularization parameter. Observations of our results indicate that there is a significant decrease in the level of reconstruction noise with depth. In particular there tends to be significant reconstruction noise near the surface. Thus a spatially varying regularization parameter might be a way to take this observation into account to improve the accuracy and reliability of the reconstruction.

Acknowledgments

References

Arridge S R 1999 Optical tomography in medical imaging *Inverse Problems* **15** R41–R93
Arridge S R, Schweijerand M and Delpy D T 1992 Iterative reconstruction of near infrared absorption images *Proc. SPIE* **1767** 372–83
Chang J, Graber H L and Barbour R L 1995 Image reconstruction of dense scattering media from CW sources using constrained CGD and a matrix rescaling technique *Proc. SPIE* **2389** 682–91
Cheng X and Boas D A 1998 Diffuse optical reflection tomography using continuous wave illumination *Opt. Express* **3**(3) August
Colak S B, Papaioannou D G, t'Hooft G W, van der Mark M B, Schomberg H, Paasschens J C J, Melissen J B M and van Asten N A A J 1997 Tomographic image reconstruction from optical projections in light-diffusing media *Appl. Opt.* **36** 180–213

Fishkin J B and Gratton E 1993 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** 127–40

Golub G H and Van Loan C F 1989 *Matrix Computations* 2nd edn (Baltimore, MD: The Johns Hopkins University Press)

Hansen P C 1998 *Rank-Deficient and Discrete Ill-Posed Problems* (Philadelphia, PA: SIAM Press)

Haskell R C, Svaasand L O, Tsay T T, Feng T C, McAdams M S and Tromberg B J 1994 Boundary conditions for the diffusion equation in radiative transfer *J. Opt. Soc. Am.* A **11** 2727–41

Ishimaru A 1997 *Wave Propagation and Scattering in Random Media* (New York: IEEE Press)

Jiang H, Paulsen K D, Osterberg U L, Pogue B W and Patterson M S 1996 Optical image reconstruction using frequency domain data: simulations and experiments *J. Opt. Soc. Am.* A **13** 253–66

Kak A C and Slaney M 1988 *Principles of Computerized Tomographic Imaging* (New York: IEEE Press)

Matson C L, Clark N, McMackin L and Fender J S 1997 Three-dimensional tumor localization in thick tissue with the use of diffuse photon-density waves *Appl. Opt.* **36** 214–20

O'Leary M A 1996 Imaging with diffuse photon density waves *PhD Thesis* University of Pennsylvania

O'Leary M A, Boas D A, Chance B and Yodh A G 1995 Experimental images of heterogenous turbid media by frequency-domain diffusing-photon tomography *Opt. Lett.* **20** 426–8

Saduiku M N O 1992 *Numerical Techniques in Electromagnetics* (Boca Raton, FL: CRC Press) ch 3

Shewchuck J R 1994 An introduction to the conjugate gradient method without the agonizing pain *Technical Report* http://www.cs.cmu.edu/~jrs/jrspapers.html#cg

van der Sluis A and van der Vorst H A 1990 SIRT- and CG-type methods for the iterative solution of sparse linear least-squares problems *Linear Algebr. Appl.* **130** 257–302