

1-1-2003

A General Framework for Nonlinear Multigrid Inversion

Seungseok Oh

Adam B. Milstein

Charles A. Bouman

Kevin J. Webb

Follow this and additional works at: <http://docs.lib.purdue.edu/ecetr>

Oh, Seungseok ; Milstein, Adam B. ; Bouman, Charles A. ; and Webb, Kevin J. , "A General Framework for Nonlinear Multigrid Inversion" (2003). *ECE Technical Reports*. Paper 135.
<http://docs.lib.purdue.edu/ecetr/135>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

A General Framework for Nonlinear Multigrid Inversion

Seungseok Oh, Adam B. Milstein, Charles A. Bouman,
and Kevin J. Webb

School of Electrical and Computer Engineering
1285 Electrical Engineering Building
Purdue University
West Lafayette, Indiana 47907-1285

This work was supported by the National Science Foundation under contract CCR-0073357.

CONTENTS

I	Introduction	1
II	Multigrid Inversion Framework	4
II-A	Inverse Problems	4
II-B	Fixed-Grid Inversion	5
II-C	Multigrid Inversion Algorithm	6
II-D	Convergence of Multigrid Inversion	10
II-E	Stabilizing Functionals	13
III	Application to Optical Diffusion Tomography	15
IV	Numerical Results	20
IV-A	Evaluation of Required Forward Model Resolution	20
IV-B	Multigrid Performance Evaluation	23
V	Conclusions	28
	References	29
	Appendix I: Proof of Multigrid Monotone Convergence	33
	Appendix II: Computational Complexity	35

LIST OF TABLES

I	Distortion-to-noise (DNR) ratio for various forward model resolutions. Coarse discretization increased forward model error, and source/detector pairs on the same face had much higher DNR.	22
II	The normalization parameter σ that yields the best reconstruction and the resulting RMS image error between the reconstructions and the decimation of the true phantom.	22
III	Complexity comparison for each algorithm. Theoretical complex multiplications are estimated with (53), and experimental relative complexity is the ratio of user time required for one iteration.	24

LIST OF FIGURES

1	The role of adjustment term $r^{(q+1)}x^{(q+1)}$. (a) When the gradients of the fine scale and coarse scale cost functionals are different at the initial value, the updated value may increase the fine grid cost functional's value. (b) When the gradients of the two functionals are matched, a properly chosen coarse scale functional can guarantee that the coarse scale update reduces the fine scale cost.	9
2	Pseudo-code specification of a two-grid inversion algorithm. The notation $c^{(q+1)}(x^{(q+1)}; y^{(q+1)}, r^{(q+1)})$ is used to make the cost functional's dependency on $y^{(q+1)}$ and $r^{(q+1)}$ explicit.	10
3	Pseudo-code specification of (a) the main routine for multigrid inversion and (b) the subroutine for the Multigrid-V inversion. The Multigrid-V algorithm is similar to the 2-grid algorithm, but recursively calls itself to perform the coarse grid update.	11
4	Pseudo-code specification of fixed grid and multigrid inversion methods for ODT problem showing (a) main routine for ODT problems, (b) fixed-grid update, and (c) Multigrid-V inversion.	19
5	(a) Source and (b) detector pattern on each face of the cube geometry. Two data set scenarios were considered: one containing all source/detector pairs, and a second containing only source/detector pairs on different faces.	21
6	A cross-section through (a) the inhomogeneous phantom, and the best reconstructions obtained using source detector pairs on different faces with (b) $65 \times 65 \times 65$ resolution, (c) $33 \times 33 \times 33$ resolution, (d) $17 \times 17 \times 17$ resolution, and (e) all source detector pairs with $65 \times 65 \times 65$ resolution. The $65 \times 65 \times 65$ reconstruction with different face source/detector pairs produced substantially better quality reconstruction. Reconstructions using all source/detector pairs failed even at $65 \times 65 \times 65$ reconstruction.	23
7	Convergence of (a) cost function and (b) RMS image error when reconstructions were initialized with average values of true phantom. All multigrid algorithms converges about 13 times faster than the fixed-grid algorithm.	25

- 8 Cross-sections of reconstructions on the $z = 0$ plane using (a) 4 level multigrid with 19.35 iterations, (b) 3 level multigrid with 19.95 iterations, (c) 2 level multigrid with 18.24 iterations, and (d) 270 fixed grid iterations. All the multigrid reconstructions have better image quality than the fixed grid reconstruction. 26
- 9 Convergence of (a) cost function and (b) RMS image error with a poor initial guess. For higher level multigrid algorithms, the convergence was faster. In particular, the four level multigrid algorithm converged almost as fast as when the reconstruction was initialized with the true phantom's average value. 27

Abstract

A variety of new imaging modalities, such as optical diffusion tomography, require the inversion of a forward problem that is modeled by the solution to a 3-D partial differential equation. For these applications, image reconstruction is particularly difficult because the forward problem is both nonlinear and computationally expensive to evaluate.

In this paper, we propose a general framework for nonlinear multigrid inversion that is applicable to a wide variety of inverse problems. The multigrid inversion algorithm results from the application of recursive multigrid techniques to the solution of optimization problems arising from inverse problems. The method works by dynamically adjusting the cost functionals at different scales so that they are consistent with, and ultimately reduce, the finest scale cost functional. In this way, the multigrid inversion algorithm efficiently computes the solution to the desired fine scale inversion problem. Importantly, the new algorithm can greatly reduce computation because both the forward and inverse problems are more coarsely discretized at lower resolutions. An application of our method to optical diffusion tomography shows the potential for very large computational savings. Numerical data also indicates robust convergence with a range of initialization conditions for this non-convex optimization problem.

Keywords

multigrid algorithms, inverse problems, optical diffusion tomography, multiresolution

I. INTRODUCTION

A large class of image processing problems, such as deblurring, high-resolution rendering, image recovery, image segmentation, motion analysis, and tomography, require the solution of inverse problems. Often, the numerical solution of these inverse problems can be computationally demanding, particularly when the problem must be formulated in three dimensions.

Recently, some new imaging modalities, such as optical diffusion tomography (ODT) [1], [2], [3], [4] and electrical impedance tomography (EIT) [5], have received much attention. For example, ODT holds great potential as a safe, non-invasive medical diagnostic modality with chemical specificity [6]. However, the inverse problems associated with these new modalities present a number of difficult challenges. First, the forward models are described by the solution of a partial differential equation (PDE) which is computationally demanding to solve. Second, the unknown image is formed by the coefficients of the PDE, so the forward model is highly nonlinear, even when the PDE is itself linear. Finally, these problems typically are inherently 3-D due to the 3-D propagation of energy in the scattering media being modeled. Since many phenomena in nature are mathematically described by PDEs, numerous other inverse problems have similar computational difficulties, including microwave tomography [7], thermal wave tomography [8], and inverse scattering [9].

To solve inverse problems, most algorithms, such as conjugate gradient (CG), steepest descent (SD), and iterative coordinate descent (ICD) [10] work by performing all computations using a fixed discretization grid. While tremendous progress has been made in reducing the computational complexity of these fixed grid methods, computational cost is still of great concern. Perhaps more importantly, fixed grid optimization methods are essentially performing a local search of the cost function, and are therefore more susceptible to being trapped in local minima that can result in poorer quality reconstructions.

Multiresolution techniques have been widely investigated to reduce computation for inverse problems. Even simple multiresolution approaches, such as initializing fine resolution iterations with coarse solutions [11], [12], [13], [14], [15], have been shown to be effective in many imaging problems. Wavelets have been studied for Bayesian tomography [16], [17], [18], [19], [20], and both wavelet and multiresolution models have been applied in Bayesian formulations

of emission tomography [21], [22], [23], [24] and thermal wave tomography [25]. For ODT, a two resolution wavelet decomposition was used to speed inversion of a problem linearized with a Born approximation [26].

Multigrid methods are a special class of multiresolution algorithms which work by recursively operating on the data at different resolutions [27], [28], [29], [30]. Multigrid algorithms originally attracted interest for numerical analysis to facilitate the computation of PDE solvers by effectively removing smooth error components which are not damped in some fixed-grid relaxation schemes. This advantage of the multigrid methods has been used to expedite convergence in various image processing problems, for example, lightness computation [31], shape-from-shading [31], optical flow estimation [31], [32], [33], [34], adaptive smoothing of signals [35], multispectral MRI image analysis [36], image matching [37], image restoration [38], and anisotropic diffusion [39].

More recently, multigrid algorithms have been used to solve image reconstruction problems. Bouman and Sauer showed that nonlinear multigrid algorithms could be applied to inversion of Bayesian tomography problems [40]. This work used nonlinear multigrid techniques to compute maximum *a posteriori* (MAP) reconstructions with non-Gaussian prior distributions and a non-negativity constraint. McCormick and Wade [41] applied multigrid methods to a linearized EIT problem, and Borcea [42] used a nonlinear multigrid approach to EIT based on a direct nonlinear formulation analogous to the full approximation scheme (FAS) in nonlinear multigrid PDE solvers. Johnson *et al.* [43] applied an algebraic multigrid algorithm to inverse bioelectric field problems formulated with the finite-element method. In [44], [45], Ye, *et al.* formulated the multigrid approach directly in an optimization framework, and used the method to solve ODT problems. In related work, Nash and Lewis formulated multigrid algorithms for the solution of a broad class of optimization problems [46], [47]. Importantly, both the approaches of Ye and Nash are based on the matching of cost functional derivatives at different scales.

In this paper we propose a method we call multigrid inversion. Multigrid inversion is a general approach for applying nonlinear multigrid optimization to the solution of inverse problems. A key innovation in our approach is that the resolution of both the forward and inverse

models are varied. This makes our method particularly well suited to the solution of inverse problems with PDE forward models for a number of reasons:

- The computation can be dramatically reduced by using coarser grids to solve the forward model PDE. In previous approaches, the forward model PDE was solved only at the finest grid. This means that coarse grid updates were either computationally costly, or a linearization approximation was made for the coarse grid forward model [41], [44], [45].
- The coarse grid forward model can be modeled by a correctly discretized PDE, preserving the nonlinear characteristics of the forward model.
- A wide variety of optimization methods can be used for solving the inverse problem at each grid. Hence, common methods such as pre-conditioned conjugate gradient and/or adjoint differentiation [48], [49] can be employed at each grid resolution.

While the multigrid inversion method is motivated by the solution of inverse problems such as ODT and EIT, it is generally applicable to any inverse problem in which the forward model can be naturally represented at differing grid resolutions.

The multigrid inversion method is formulated in an optimization framework by defining a sequence of optimization functionals at decreasing resolutions. In order for the method to have well behaved convergence to the correct fine grid solution, it is essential that the cost functionals at different scales be consistent. To achieve this, we propose a recursive method for adapting the coarse grid functionals which guarantees that the fine grid solution is also a fixed point of the multigrid algorithm. In addition, we show that under certain conditions, the nonlinear multigrid inverse algorithm is guaranteed to produce monotone convergence of the fine grid cost functional. We present experimental results for the ODT application which show that the multigrid inversion algorithm can provide dramatic reductions in computation when the inversion problem is solved at the resolution necessary to achieve a high quality reconstruction.

This paper is organized as follows. Section II introduces the general concept of the multigrid inversion algorithm, and Section II-D discusses its convergence. In Section III, we illustrate the application of the multigrid inversion method to the ODT problem, and its numerical results are provided in Section IV. Finally, Section V makes concluding remarks.

II. MULTIGRID INVERSION FRAMEWORK

In this section, we overview regularized inverse methods and then formulate the general multigrid inversion approach.

A. Inverse Problems

Let Y be a random vector of (real or complex) measurements, and let x be a finite dimensional vector representing the unknown quantity, in our case an image, to be reconstructed. For any inverse problem, there is a forward model $f(x)$ given by

$$E[Y|x] = f(x) \quad (1)$$

which represents the computed means of the measurements given the image x . For many inverse problems, such as ODT, the forward model $f(x)$ is given by the solution of a PDE where x determines the coefficients of the discretized PDE. We will assume that the measurements Y are conditionally Gaussian given x , so that

$$\log p(y|x) = -\frac{1}{2\alpha} \|y - f(x)\|_{\Lambda}^2 - \frac{P}{2} \log(2\pi\alpha|\Lambda|^{-1}), \quad (2)$$

where Λ is a positive definite weight matrix, P is the dimensionality of the measurement, α is a parameter proportional to the noise variance, and $\|w\|_{\Lambda}^2 = w^H \Lambda w$. Note that the measurement noise covariance matrix is equal to $\alpha\Lambda^{-1}$. When the data values are real valued, P is equal to the length of the vector Y , but when the measurements are complex, then P is equal to twice the dimension of Y .

Our objective is to invert the forward model of (1) and thereby estimate x from a particular measurement vector y . There are a variety of methods from performing this estimation including maximum a posteriori (MAP) estimation, penalized maximum likelihood, and regularized inverse. All of these methods work by computing the value of x which minimizes a cost functional of the form

$$\frac{1}{2\alpha} \|y - f(x)\|_{\Lambda}^2 + \frac{P}{2} \log(2\pi\alpha|\Lambda|^{-1}) + S(x), \quad (3)$$

where $S(x)$ is a stabilizing functional used to regularize the inverse. Note that in the MAP approach, $S(x) = -\log p(x)$, where $p(x)$ is the prior distribution assumed for x . We will

estimate both the noise variance parameter α and x by jointly maximizing over both quantities [50]. Minimization of (3) with respect to α yields the condition $\hat{\alpha} = \frac{1}{P} \|y - f(x)\|_{\Lambda}^2$. Substitution of $\hat{\alpha}$ into (3) and dropping constants yields the cost functional to be optimized as

$$c(x) = \frac{P}{2} \log \|y - f(x)\|_{\Lambda}^2 + S(x), \quad (4)$$

where we will generally assume $c(x)$ be a continuously differentiable function of x .

We have found that joint optimization over α and x has a number of important advantages. First, in many applications the absolute magnitude of the measurement noise is not known in advance, while the relative noise magnitude may be known. In such a scenario, it is useful to simultaneously estimate the value of α along with the value of x [51], [44], [45]. More importantly, we have found that the logarithm in the expression of (4) makes optimization less susceptible to being trapped in local minima [52]. In any case, the multigrid methods we describe are equally applicable to the case when α is fixed. In this case, the cost functional is given by $c(x) = \frac{1}{2\alpha} \|y - f(x)\|_{\Lambda}^2 + S(x)$, instead of (4).

B. Fixed-Grid Inversion

Once the cost functional of (4) is formulated, the inverse is computed by solving the associated optimization problem

$$\hat{x} = \arg \min_x \left\{ \frac{P}{2} \log \|y - f(x)\|_{\Lambda}^2 + S(x) \right\}. \quad (5)$$

Most optimization algorithms, such as CG, SD, and ICD, work by iteratively minimizing the cost functional. We express a single iteration of such a fixed grid optimizer as

$$x_{update} \leftarrow \text{Fixed_Grid_Update}(x_{init}, c(\cdot)), \quad (6)$$

where $c(\cdot)$ is the cost functional being minimized, x_{init} is the initial value of x , and x_{update} is the updated value.¹ We will generally assume that the fixed grid algorithm reduces the cost functional with each iteration, unless the initial value of x is at a local minimum of the

¹We use the \leftarrow symbol to denote assignment of a value to a variable, thereby eliminating the need for time indexing in update equations.

cost functional. Therefore, we say that an update algorithm is **monotone** if $c(x_{update}) \leq c(x_{init})$, with strict inequality when $\nabla c(x_{init}) \neq 0$ or $x_{update} \neq x_{init}$. Repeated application of a monotone fixed grid optimizer will produce a sequence of estimates with monotonically decreasing cost. Thus, we may approximately solve (5) through iterative application of (6).

In many inverse problems, such as ODT, the forward model computation requires the solution of a 3-D PDE which must be discretized for numerical solution on a computer. Although a fine discretization grid is desirable because it reduces modeling error and increases the resolution of the final image, these improvements are obtained at the expense of a dramatic increase in computational cost. For a 3-D problem, the computational cost typically increases by a factor of 8 each time the resolution is doubled. Solving problems at fine resolution also tends to slow convergence. For example, many fixed grid algorithms such as ICD² effectively eliminate error at high spatial frequencies, but low frequency errors are damped slowly [27], [10].

C. Multigrid Inversion Algorithm

In this section, we derive the basic multigrid inversion algorithm for solving the optimization of (5). Let $x^{(0)}$ denote the finest grid image, and let $x^{(q)}$ be a coarse resolution representation of $x^{(0)}$ with a grid sampling period of 2^q times the finest grid sampling period. To obtain a coarser resolution image $x^{(q+1)}$ from a finer resolution image $x^{(q)}$, we use the relation $x^{(q+1)} = I_{(q)}^{(q+1)} x^{(q)}$, where $I_{(q)}^{(q+1)}$ is a linear decimation matrix. We use $I_{(q+1)}^{(q)}$ to denote the corresponding linear interpolation matrix.

We first define a coarse grid cost functional, $\tilde{c}^{(q)}(x^{(q)})$, with a form analogous to that of (4), but with quantities indexed by the scale q , as

$$\tilde{c}^{(q)}(x^{(q)}) = \frac{P}{2} \log \|y^{(q)} - f^{(q)}(x^{(q)})\|_{\Lambda}^2 + S^{(q)}(x^{(q)}). \quad (7)$$

Notice that the forward model $f^{(q)}(\cdot)$ and the stabilizing functional $S^{(q)}(\cdot)$ are both evaluated at scale q . This is important because evaluation of the forward model at low resolution substantially reduces computation due to the reduced number of variables. The specific form of $f^{(q)}(\cdot)$ generally results from the physical problem being solved with an appropriate grid

²ICD is generally referred to as Gauss-Seidel in the PDE literature.

spacing. In Section III, we will give a typical example for ODT where $f^{(q)}(\cdot)$ is computed by discretizing the 3-D PDE using a grid spacing proportional to 2^q . The quantity $y^{(q)}$ in (7) denotes an adjusted measurement vector at scale q . The stabilizing functional at each scale is fixed and chosen to best approximate the fine scale functional. We give an example of such a stabilizing functional later in Section II-E.

In the remainder of this section, we explain how the cost functionals at each scale can be matched to produce a consistent solution. To do this, we define an adjusted cost functional

$$\begin{aligned} c^{(q)}(x^{(q)}) &= \tilde{c}^{(q)}(x^{(q)}) - r^{(q)}x^{(q)} \\ &= \frac{P}{2} \log \|y^{(q)} - f^{(q)}(x^{(q)})\|_{\Lambda}^2 + S^{(q)}(x^{(q)}) - r^{(q)}x^{(q)}, \end{aligned} \quad (8)$$

where $r^{(q)}$ is a row vector used to adjust the functional's gradient. At the finest scale, all quantities take on their fine scale values and $r^{(q)} = 0$, so that $\tilde{c}^{(0)}(x^{(0)}) = c^{(0)}(x^{(0)}) = c(x)$. Our objective is then to derive recursive expressions for the quantities $y^{(q)}$ and $r^{(q)}$ that match the cost functionals at fine and coarse scales.

Let $x^{(q)}$ be the current solution at grid q . We would like to improve this solution by first performing an iteration of fixed grid optimization at the coarser grid $q + 1$, and then using this result to correct the finer grid solution. This coarse grid update is

$$\tilde{x}^{(q+1)} \leftarrow \text{Fixed_Grid_Update}(I_{(q)}^{(q+1)}x^{(q)}, c^{(q+1)}(\cdot)), \quad (9)$$

where $I_{(q)}^{(q+1)}x^{(q)}$ is the initial condition formed by decimating $x^{(q)}$, and $\tilde{x}^{(q+1)}$ is the updated value. We may now use this result to update the finer grid solution. We do this by interpolating the *change* in the coarser scale solution by

$$\tilde{x}^{(q)} \leftarrow x^{(q)} + I_{(q+1)}^{(q)}(\tilde{x}^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)}). \quad (10)$$

Ideally, the new solutions $\tilde{x}^{(q)}$ should be at least as good as the old solution $x^{(q)}$. Specifically, we would like $c^{(q)}(\tilde{x}^{(q)}) \leq c^{(q)}(x^{(q)})$ when the fixed grid algorithm is monotone. However, this may not be the case if the cost functionals are not consistent. In fact, for a naively chosen set of cost functionals, the coarse scale correction could easily move the solution away from the optimum.

This problem of inconsistent cost functionals is eliminated if the fine and coarse scale cost functionals are equal within an additive constant.³ This means we would like

$$c^{(q+1)}(\tilde{x}^{(q+1)}) \simeq c^{(q)}(x^{(q)} + I_{(q+1)}^{(q)}(\tilde{x}^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)})) + \text{constant} \quad (11)$$

to hold for all values of $\tilde{x}^{(q+1)}$. Our objective is then to choose a coarse scale cost functional which matches the fine cost functional as described in (11). We do this by the proper selection of $y^{(q+1)}$ and $r^{(q+1)}$. First, we enforce the condition that the initial error between the forward model and measurements be the same at the coarse and fine scales, giving

$$y^{(q+1)} - f^{(q+1)}(I_{(q)}^{(q+1)}x^{(q)}) = y^{(q)} - f^{(q)}(x^{(q)}) . \quad (12)$$

This yields the update for $y^{(q+1)}$

$$y^{(q+1)} \leftarrow y^{(q)} - [f^{(q)}(x^{(q)}) - f^{(q+1)}(I_{(q)}^{(q+1)}x^{(q)})] . \quad (13)$$

Intuitively, the term in the bracket compensates for the forward model mismatch between resolutions.

Next, we use the condition introduced in [44], [45], [46], [47] to enforce the condition that the gradients of the coarse and fine cost functionals be equal at the current values of $x^{(q)}$ and $x^{(q+1)} = I_{(q)}^{(q+1)}x^{(q)}$. More precisely, we enforce the condition that

$$\nabla c^{(q+1)}(x^{(q+1)}) \Big|_{x^{(q+1)}=I_{(q)}^{(q+1)}x^{(q)}} = \nabla c^{(q)}(x^{(q)}) I_{(q+1)}^{(q)} . \quad (14)$$

This condition is essential to assure that the optimum solution is a fixed point of the multigrid inversion algorithm [45], and is illustrated graphically in Fig. 1. In Section II-D, we will also show how this condition can be used along with other assumptions to ensure monotone convergence of the multigrid inversion algorithm.

The equality of (14) can be enforced at the current value $x^{(q)}$ by choosing

$$r^{(q+1)} \leftarrow \nabla \tilde{c}^{(q+1)}(x^{(q+1)}) \Big|_{x^{(q+1)}=I_{(q)}^{(q+1)}x^{(q)}} - \left(\nabla \tilde{c}^{(q)}(x^{(q)}) - r^{(q)} \right) I_{(q+1)}^{(q)} , \quad (15)$$

³A constant offset has no effect on the value of x which minimizes the cost functional.




Fig. 1. The role of adjustment term $r^{(q+1)}x^{(q+1)}$. (a) When the gradients of the fine scale and coarse scale cost functionals are different at the initial value, the updated value may increase the fine grid cost functional's value. (b) When the gradients of the two functionals are matched, a properly chosen coarse scale functional can guarantee that the coarse scale update reduces the fine scale cost.

where $\tilde{c}^{(q)}(\cdot)$ is the unadjusted cost functional defined in (7). By evaluating the gradients and using the update relation of (13), we obtain

$$r^{(q+1)} \leftarrow g^{(q+1)} - (g^{(q)} - r^{(q)}) I_{(q+1)}^{(q)}, \quad (16)$$

where $g^{(q)}$ and $g^{(q+1)}$ are the gradients of the unadjusted cost functional at the fine and coarse scales, respectively, given by

$$g^{(q)} = -\frac{P}{\|y^{(q)} - f^{(q)}(x^{(q)})\|_{\Lambda}^2} \text{Re} \left\{ (y^{(q)} - f^{(q)}(x^{(q)}))^H \Lambda A^{(q)} \right\} + \nabla S^{(q)}(x^{(q)}) \quad (17)$$

$$g^{(q+1)} = -\frac{P}{\|y^{(q)} - f^{(q)}(x^{(q)})\|_{\Lambda}^2} \text{Re} \left\{ (y^{(q)} - f^{(q)}(x^{(q)}))^H \Lambda A^{(q+1)} \right\} + \nabla S^{(q+1)}(I_{(q)}^{(q+1)}x^{(q)}), \quad (18)$$

where H is the conjugate transpose (Hermitian) operator, and $A^{(q)}$ denotes the gradient of the forward model or Fréchet derivative given by

$$A^{(q)} = \nabla f^{(q)}(x^{(q)}) \quad (19)$$

$$A^{(q+1)} = \nabla f^{(q+1)}(x^{(q+1)}) \Big|_{x^{(q+1)}=I_{(q)}^{(q+1)}x^{(q)}} \cdot \quad (20)$$

```

 $x^{(q)} \leftarrow \text{Twogrid\_Update}(q, x^{(q)}, y^{(q)}, r^{(q)}) \{$ 
  Repeat  $\nu_1^{(q)}$  times
     $x^{(q)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q)}, c^{(q)}(\cdot; y^{(q)}, r^{(q)}))$  //Fine grid update
   $x^{(q+1)} \leftarrow I_{(q)}^{(q+1)} x^{(q)}$  //Decimation
  Compute  $y^{(q+1)}$  using (13)
  Compute  $r^{(q+1)}$  using (16)
  Repeat  $\nu_1^{(q+1)}$  times
     $x^{(q+1)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q+1)}, c^{(q+1)}(\cdot; y^{(q+1)}, r^{(q+1)}))$  //Coarse grid update
   $x^{(q)} \leftarrow x^{(q)} + I_{(q+1)}^{(q)}(x^{(q+1)} - I_{(q)}^{(q+1)} x^{(q)})$  //Coarse grid correction
  Repeat  $\nu_2^{(q)}$  times
     $x^{(q)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q)}, c^{(q)}(\cdot; y^{(q)}, r^{(q)}))$  //Fine grid update
  Return  $x^{(q)}$  //Return result
 $\}$ 

```

Fig. 2. Pseudo-code specification of a two-grid inversion algorithm. The notation $c^{(q+1)}(x^{(q+1)}; y^{(q+1)}, r^{(q+1)})$ is used to make the cost functional's dependency on $y^{(q+1)}$ and $r^{(q+1)}$ explicit.

As a summary of this section, Fig. 2 shows pseudocode for implementing the two-grid algorithm. In this figure, we use the notation $c^{(q+1)}(x^{(q+1)}; y^{(q+1)}, r^{(q+1)})$ to make the dependency on $y^{(q+1)}$ and $r^{(q+1)}$ explicit. Notice that $\nu_1^{(q)}$ fixed grid iterations are done before the coarse grid correction, and that $\nu_2^{(q)}$ iterations are done afterwards. The convergence speed of the algorithm can be tuned through the choice of $\nu_1^{(q)}$ and $\nu_2^{(q)}$ at each scale.

The Multigrid-V algorithm [27] is obtained by simply replacing the fixed grid update at resolution $q + 1$ of the two-grid algorithm with a recursive subroutine call, as shown in the pseudocode in Fig. 3(b). We can then solve (5) through iterative application of the Multigrid-V algorithm, as shown in Fig. 3(a). The Multigrid-V algorithm then moves from fine to coarse to fine resolutions with each iteration.

D. Convergence of Multigrid Inversion

Multigrid inversion can be viewed as a method to simplify a potentially expensive optimization by temporarily replacing the original cost functional by a lower resolution one. In fact, there is a large class of optimization methods which depend on the use of so-called surrogate functionals, or functional substitution methods to speed or simplify optimization. A classic example of a surrogate functional is the Q-function used in the EM algorithm [53], [54]. More

```

main() {
  Initialize  $x^{(0)}$  with a background estimate
   $r^{(0)} \leftarrow 0$ 
   $y^{(0)} \leftarrow y$ 
  Choose number of fixed grid iterations  $\nu_1^{(0)}, \dots, \nu_1^{(Q-1)}$  and  $\nu_2^{(0)}, \dots, \nu_2^{(Q-1)}$ 
  Repeat until converged:
     $x^{(0)} \leftarrow \text{MultigridV}(q, x^{(0)}, c^{(0)}(\cdot; y^{(0)}, r^{(0)}))$ 
}

```

(a)

```

 $x^{(q)} \leftarrow \text{MultigridV}(q, x^{(q)}, y^{(q)}, r^{(q)})$  {
  Repeat  $\nu_1^{(q)}$  times
     $x^{(q)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q)}, c^{(q)}(\cdot; y^{(q)}, r^{(q)}))$  //Fine grid update
  If  $q = Q - 1$ , return  $x^{(q)}$  //If coarsest scale, return result
   $x^{(q+1)} \leftarrow I_{(q)}^{(q+1)} x^{(q)}$  //Decimation
  Compute  $y^{(q+1)}$  using (13)
  Compute  $r^{(q+1)}$  using (15)
   $x^{(q+1)} \leftarrow \text{MultigridV}(q + 1, x^{(q+1)}, y^{(q+1)}, r^{(q+1)})$  //Coarse grid update
   $x^{(q)} \leftarrow x^{(q)} + I_{(q+1)}^{(q)} (x^{(q+1)} - I_{(q)}^{(q+1)} x^{(q)})$  //Coarse grid correction
  Repeat  $\nu_2^{(q)}$  times
     $x^{(q)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q)}, c^{(q)}(\cdot; y^{(q)}, r^{(q)}))$  //Fine grid update
  Return  $x^{(q)}$  //Return result
}

```

(b)

Fig. 3. Pseudo-code specification of (a) the main routine for multigrid inversion and (b) the subroutine for the Multigrid-V inversion. The Multigrid-V algorithm is similar to the 2-grid algorithm, but recursively calls itself to perform the coarse grid update.

recently, De Pierro discovered that this same basic method could be applied to tomography problems in a manner that allowed parallel updates of pixels in the computation of penalized ML reconstructions [55], [56]. De Pierro's method has since been exploited to both prove convergence and allow parallel updates for ICD methods in tomography [57], [58].

However, the application of surrogate functionals to multigrid inversion is unique in that the substituting functional is at a coarser scale and therefore has an argument of lower dimension. As with traditional approaches, the surrogate functional should be designed to guarantee monotone convergence of the original cost functional. In the case of the multigrid algorithm, a sequence of optimization functionals at varying resolutions should be designed so that the

entire multigrid update decreases the finest resolution cost function.

Figure 1 graphically illustrates the use of surrogate functionals in multigrid inversion. Figure 1(a) shows the case in which the gradients of the fine scale and coarse scale (i.e. surrogate) functions are different at the initial value. In this case, the surrogate function can not upper bound the value of the fine scale functional, and the updated value may actually increase the fine grid cost functionals value. Figure 1(b) illustrates the case in which the gradients of the two functionals are matched. In this case, a properly chosen coarse scale functional can upper bound the fine scale functional, and the coarse scale update is guaranteed to reduce the fine scale cost.

The concepts illustrated in Fig. 1 can be formalized into conditions that guarantee the monotone convergence of the multigrid algorithms. The following theorem, proved in Appendix I, gives a set of sufficient conditions for monotone convergence of the multigrid inversion algorithm.

Theorem: (Multigrid Monotone Convergence)

For $0 \leq q < Q - 1$, define the functional $\xi^{(q+1)} : \mathbb{R}^{N^{(q+1)}} \rightarrow \mathbb{R}$

$$\xi^{(q+1)}(x^{(q+1)}) = \tilde{c}^{(q+1)}(x^{(q+1)}) - \tilde{c}^{(q)}(x^{(q)} + I_{(q+1)}^{(q)}(x^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)})), \quad (21)$$

where $N^{(q+1)}$ is the number of voxels in $x^{(q+1)}$, \mathbb{R} is the set of real numbers, and the functions $\tilde{c}^{(q)}(\cdot)$ and $\tilde{c}^{(q+1)}(\cdot)$ are continuously differentiable. Assume that the following conditions are satisfied:

- 1) The fixed grid update is monotone for $0 \leq q < Q$.
- 2) $\xi^{(q)}(\cdot)$ is convex on $\mathbb{R}^{N^{(q)}}$ for $0 < q < Q$.
- 3) The adjustment vector $r^{(q+1)}$ is given by (15) for $0 \leq q < Q$.
- 4) $\nu_1^{(q)} + \nu_2^{(q)} \geq 1$ for $0 \leq q < Q$.

Then, the multigrid algorithm of Fig. 3 is monotone for $c^{(0)}(\cdot)$.

The conditions 1, 3, and 4 of the Theorem are easily satisfied for most problems. However, the difficulty lies in satisfying condition 2, convexity of $\xi^{(q)}(\cdot)$ for $q > 0$. If the eigenvalues of the Hessian of $\xi^{(q)}(\cdot)$ are lower-bounded, the convexity condition can be satisfied by adding a convex term, such as $\gamma \|x^{(q)}\|^2$, to $\tilde{c}^{(q)}(\cdot)$ for $q > 0$, where γ is a sufficiently large con-

stant. However, addition of such a term tends to slow convergence by making the coarse scale corrections too conservative.

When the forward model is given by a PDE, it can be difficult or impossible to verify or guarantee the convexity condition of 2. Nonetheless, the theorem still gives insight into the convergence behavior of the algorithm; and in Section IV we will show that empirically, for the difficult problem of ODT, the convergence of the multigrid algorithm is monotone in all cases, even without the addition of any convex terms.

E. Stabilizing Functionals

The coarse scale stabilizing functionals, $S^{(q)}(x^{(q)})$, may be derived through appropriate scaling of $S(x)$. A general class of stabilizing functional has the form

$$S(x) = \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} \rho \left(\frac{|x_i - x_j|}{\sigma} \right), \quad (22)$$

where the set \mathcal{N} consists of all pairs of adjacent grid points, b_{i-j} represents the weighting assigned to the pair $\{i, j\}$, σ is a parameter that controls the overall weighting, and $\rho(\cdot)$ is a symmetric function that penalizes the differences in adjacent pixel values. Such a stabilizing functional results from the selection of a prior density $p(x)$ corresponding to a Markov random field (MRF) [59]. A wide variety of functionals $\rho(\cdot)$ have been suggested for this purpose [60], [61], [62]. Generally, these methods attempt to select these functionals so that large differences in pixel value are not excessively penalized, thereby allowing the accurate formation of sharp edge discontinuities.

The stabilizing functional at scale q must be selected so that

$$S^{(q)}(x^{(q)}) \cong S(x). \quad (23)$$

This can be done by using a form similar to (22) and applying scaling factors to result in

$$S^{(q)}(x^{(q)}) = 2^{qd} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} \rho \left(\frac{|x_i^{(q)} - x_j^{(q)}|}{2^q \sigma} \right), \quad (24)$$

where d is the dimension of the problem. Here we assume that $x_i - x_j \cong (x_i^{(q)} - x_j^{(q)})/2^q$, and we use the constant 2^{qd} to compensate for the reduction in the number of terms as the sampling grid is coarsened.

In our experiments, we use the generalized Gaussian Markov random field (GGMRF) image prior model [62], [52], [45], [13], [14] given by

$$p(x) = \frac{1}{\sigma^N z(p)} \exp \left[-\frac{1}{p\sigma^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i - x_j|^p \right], \quad (25)$$

where σ is a normalization parameter, $1 \leq p \leq 2$ controls the degree of edge smoothness, and $z(p)$ is a partition function. For the GGMRF prior, the stabilizing functional is given by

$$S(x) = \frac{1}{p\sigma^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i - x_j|^p, \quad (26)$$

and the corresponding coarse scale stabilizing functionals are derived using (24) to be

$$S^{(q)}(x^{(q)}) = \frac{1}{p(\sigma^{(q)})^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i^{(q)} - x_j^{(q)}|^p, \quad (27)$$

where $\sigma^{(q)}$ is given by

$$\sigma^{(q)} = 2^{(q)(1-\frac{d}{p})} \cdot \sigma^{(0)}. \quad (28)$$

III. APPLICATION TO OPTICAL DIFFUSION TOMOGRAPHY

Optical diffusion tomography is a method for determining spatial maps of optical absorption and scattering properties from measurements of light intensity transmitted through a highly scattering medium. In frequency domain ODT, the measured modulation envelope of the optical flux density is used to reconstruct the absorption coefficient and diffusion coefficient at each discretized grid point. However, for simplicity, we will only consider reconstruction of the absorption coefficient.

The complex amplitude $\phi_k(r)$ of the modulation envelope due to a point source at position a_k and angular frequency ω satisfies the frequency domain diffusion equation

$$\nabla \cdot [D(r)\nabla\phi_k(r)] + [-\mu_a(r) - j\omega/c]\phi_k(r) = -\delta(r - a_k), \quad (29)$$

where r is position, c is the speed of light in the medium, $\mu_a(r)$ is the absorption coefficient, and $D(r)$ is the diffusion coefficient. The 3-D domain is discretized into N grid points, denoted by r_1, r_2, \dots, r_N . The unknown image is then represented by an N dimensional column vector $x = [\mu_a(r_1), \mu_a(r_2), \dots, \mu_a(r_N)]^T$ containing the absorption coefficients at each discrete grid point, where T is the transpose operator. We will use the notation $\phi_k(r; x)$ in place of $\phi_k(r)$, in order to emphasize the dependence of the solution on the unknown image x . Then the measurement of a detector at location b_m resulting from a source at location a_k can be modeled by the complex value $\phi_k(b_m; x)$. The complete forward model function is then given by ⁴

$$f(x) = [\phi_1(b_1; x), \phi_1(b_2; x), \dots, \phi_1(b_M; x), \phi_2(b_1; x), \dots, \phi_K(b_M; x)]^T. \quad (30)$$

Note that $f(x)$ is a highly nonlinear function because it is given by the solution to a PDE using coefficients x . The measurement vector is also organized similarly as

$y = [y_{11}, y_{12}, \dots, y_{1m}, y_{21}, \dots, y_{KM}]^T$, where y_{km} is the measurement with the source at a_k and the detector at b_m .

⁴For simplicity of notation, we assume that all source-detector pairs are used. However, in our experimental simulations we use only a subset of all possible measurements. In fact, practical limitations can often limit the available measurements to a subset so that $P \neq 2KM$.

Our objective is to estimate the unknown image x from the measurements y . In a Bayesian framework, the MAP estimate of x is given by

$$\hat{x}_{MAP} = \arg \max_{x \geq 0} \{ \log p(y|x) + \log p(x) \}, \quad (31)$$

where $p(y|x)$ is the data likelihood and $p(x)$ is the prior model for image x , which is assumed to be strictly positive in value. We use an independent Gaussian shot noise model [52] with the form given in (2), where the weight matrix Λ is given by

$$\Lambda = \text{diag}\left(\frac{1}{|y_{11}|}, \dots, \frac{1}{|y_{1M}|}, \frac{1}{|y_{21}|}, \dots, \frac{1}{|y_{KM}|}\right). \quad (32)$$

For the prior model, we use the GGMRF density of (25) for $p(x)$. Using the formulation of Section II-A, the ODT imaging problem is reduced to the optimization

$$(\hat{x}_{MAP}, \hat{\alpha}) = \arg \max_{x \geq 0} \max_{\alpha} \left\{ -\frac{1}{2\alpha} \|y - f(x)\|_{\Lambda}^2 - \frac{P}{2} \log \alpha - \frac{1}{p\sigma^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i - x_j|^p \right\}, \quad (33)$$

where constant terms are neglected. Minimizing (33) with respect to α reduces the cost functional to

$$c(x) = \frac{P}{2} \log \|y - f(x)\|_{\Lambda}^2 + \frac{1}{p\sigma^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i - x_j|^p. \quad (34)$$

This cost functional has the same form as (4) with the stabilizing functional given by (26).

The gradient terms of the stabilizing functional used in (17) and (18) are given by

$$\nabla S(x) = \frac{1}{\sigma^p} \sum_{j \in \mathcal{N}_n} b_{n-j} |x_n - x_j|^{p-1} \text{sgn}(x_n - x_j). \quad (35)$$

We use multigrid inversion to solve the required optimization problem with coarse grid cost functionals of the form

$$c^{(q)}(x^{(q)}) = \frac{P}{2} \log \|y^{(q)} - f^{(q)}(x^{(q)})\|_{\Lambda}^2 + \frac{1}{p(\sigma^{(q)})^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i^{(q)} - x_j^{(q)}|^p - r^{(q)} x^{(q)}, \quad (36)$$

where $\sigma^{(q)}$ is given by (28) with $d = 3$.

At each scale q , we must also select a fixed grid optimization algorithm. For simplicity, we minimize (36) by alternatively minimizing with respect to α and x using the update formulas

$$\alpha \leftarrow \frac{1}{P} \|y - f(x)\|_{\Lambda}^2 \quad (37)$$

$$x \leftarrow \approx \arg \min_{x \geq 0} \left\{ \frac{1}{2\alpha} \|y - f(x)\|_{\Lambda}^2 + \frac{1}{p\sigma^p} \sum_{\{i,j\} \in \mathcal{N}} b_{i-j} |x_i - x_j|^p - rx \right\}, \quad (38)$$

where all expressions are interpreted as their corresponding scale q quantities. The fixed scale optimization (38) is performed using ICD optimization, as described in [52]. ICD requires the evaluation of the Fréchet derivative matrix of (19). For the ODT problem, it can be shown that the Fréchet derivative is given by

$$\begin{aligned} A_{(k-1)M+m, n} &= \frac{\partial [f(x)]_{(k-1)M+m}}{\partial x_n} \\ &= \frac{\partial \phi_k(b_m; x)}{\partial x_n} \\ &= -G(a_k, r_n; x)G(b_m, r_n; x)V, \end{aligned} \quad (39)$$

where V is the voxel volume, $G(r_s, r_o; x)$ is the diffusion equation Green's function for the problem domain computed using the image x , with r_s as the source location and r_o as the observation point, and domain discretization errors are ignored [63], [14]. Since the ODT problem is inherently 3-D, the Fréchet derivative matrix is usually very large. Fortunately, the separable structure of the Fréchet derivative can be use to substantially reduce memory requirements by storing the two quantities

$$\phi = [G(a_1, r_1; x), \dots, G(a_1, r_N; x), G(a_2, r_1; x), \dots, G(a_K, r_N; x)] \quad (40)$$

$$\psi = [G(b_1, r_1; x), \dots, G(b_1, r_N; x), G(b_2, r_1; x), \dots, G(b_M, r_N; x)] \quad (41)$$

and computing A on the fly.

The ICD algorithm is initialized by setting a state vector \hat{y} equal to the forward model output for the current value of x , giving

$$\hat{y} \leftarrow f(x). \quad (42)$$

Each ICD iteration is then computed by visiting each voxel n once using a random order, and updating each pixel value x_n and the state \hat{y} using the following expressions

$$x_{old,n} \leftarrow x_n \tag{43}$$

$$x_n \leftarrow \arg \min_{u \geq 0} \left\{ \frac{1}{2\alpha} \|y - \hat{y} - A_{*n}(u - x_n)\|_{\Lambda}^2 + \frac{1}{p\sigma^p} \sum_{j \in \mathcal{N}_n} b_{n-j} |u - x_j|^p - r_n u \right\} \tag{44}$$

$$\hat{y} \leftarrow \hat{y} + A_{*n}(x_n - x_{old,n}) , \tag{45}$$

where A_{*n} is the n^{th} column of the matrix A . Note that the state \hat{y} keeps a running estimate of the forward model output by (45), so that subsequent state updates can be computed efficiently.

Figure 4 shows a detailed pseudo-code specification for the fixed grid and multigrid algorithms for the ODT application. In particular, it explicitly shows the computation of the quantities $\phi^{(q)}$ and $\psi^{(q)}$ used in the computation of the Fréchet derivative.

```

main() {
  Initialize  $x^{(0)}$  with a background estimate
  For  $q = 1, 2, \dots, Q - 1$ ,  $x^{(q)} \leftarrow I_{(q-1)}^{(q)} x^{(q-1)}$ 
  For  $q = 0, 1, \dots, Q - 1$ ,  $r^{(q)} \leftarrow 0$  and  $y^{(q)} \leftarrow y$ 
  Repeat until converged: {
    Compute  $\phi^{(0)}, \psi^{(0)}$  and  $\hat{y} \leftarrow f^{(0)}(x^{(0)})$ 
    If Multigrid Inversion :
      Choose  $\nu_1^{(0)}, \dots, \nu_1^{(Q-1)}$  and  $\nu_2^{(0)}, \dots, \nu_2^{(Q-1)}$ 
       $x^{(0)} \leftarrow \text{MultigridV}(0, x^{(0)}, y^{(0)}, r^{(0)}, \phi^{(0)}, \psi^{(0)}, \hat{y})$ 
    If Fixed Grid Inversion :
       $x^{(0)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(0)}, y^{(0)}, r^{(0)}, \phi^{(0)}, \psi^{(0)}, \hat{y})$ 
  }
}

```

(a)

```

 $x \leftarrow \text{Fixed\_Grid\_Update}(x, y, r, \phi, \psi, \hat{y})$  {
  Compute  $\alpha \leftarrow \frac{1}{P} \|y - \hat{y}\|_{\Lambda}^2$ 
  For  $n = 0, \dots, N - 1$  (in random order), {
    Compute column vector  $A_{*n}$  with (39)
    Update  $x_n$ , as described by Ye, et al.[52]:
       $x_{old,n} \leftarrow x_n$ 
      
$$x_n \leftarrow \arg \min_{u \geq 0} \left\{ \frac{1}{2\alpha} \|y - \hat{y} - A_{*n}(u - x_n)\|_{\Lambda}^2 + \frac{1}{p\sigma^p} \sum_{j \in \mathcal{N}_n} b_{n-j} |u - x_j|^p - r_n u \right\}$$

       $\hat{y} \leftarrow \hat{y} + A_{*n}(x_n - x_{old,n})$ 
  }
}

```

(b)

```

 $x^{(q)} \leftarrow \text{MultigridV}(q, x^{(q)}, y^{(q)}, r^{(q)}, \phi^{(q)}, \psi^{(q)}, \hat{y})$  {
  For  $\nu = 1, \dots, \nu_1^{(q)}$ 
     $x^{(q)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q)}, y^{(q)}, r^{(q)}, \phi^{(q)}, \psi^{(q)}, \hat{y})$  //Fine grid update
  If  $q - 1$  is coarsest desired grid, return  $x^{(q)}$  //Return result
   $x^{(q+1)} \leftarrow I_{(q)}^{(q+1)} x^{(q)}$  //Decimation
  Compute  $\phi^{(q+1)}, \psi^{(q+1)}$  and  $\hat{y} \leftarrow f^{(q+1)}(x^{(q+1)})$ 
  Compute  $y^{(q+1)}$  using (13)
  Compute  $r^{(q+1)}$  using (16)
   $x^{(q+1)} \leftarrow \text{MultigridV}(q + 1, x^{(q+1)}, y^{(q+1)}, r^{(q+1)}, \phi^{(q+1)}, \psi^{(q+1)}, \hat{y})$  //Coarse grid update
   $x^{(q)} \leftarrow x^{(q)} + I_{(q+1)}^{(q)}(x^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)})$  //Coarse grid correction
  For  $\nu = 1, \dots, \nu_2^{(q)}$ 
     $x^{(q)} \leftarrow \text{Fixed\_Grid\_Update}(x^{(q)}, y^{(q)}, r^{(q)}, \phi^{(q)}, \psi^{(q)}, \hat{y})$  //Fine grid update
  Return  $x^{(q)}$  //Return result
}

```

(c)

Fig. 4. Pseudo-code specification of fixed grid and multigrid inversion methods for ODT problem showing (a) main routine for ODT problems, (b) fixed-grid update, and (c) Multigrid-V inversion.

IV. NUMERICAL RESULTS

This section contains the results of numerical experiments using simulated data sets. In all cases, our simulated physical measurements were generated using a $257 \times 257 \times 257$ grid discretization of the domain and the MUDPACK [64] PDE solver. We used the highest practical resolution for the forward model simulation, so as to achieve the best possible accuracy of the simulated measurements. Since the sources and detectors are not located exactly on the grid points, a three-dimensional linear interpolation of the nearest grid points was also used.

Our experiments used two tissue phantoms, which we refer to as the homogeneous and non-homogeneous phantoms. Both phantoms had dimensions of $10 \times 10 \times 10$ cm, and with each face containing eight sources and nine detectors with modulation frequencies of 100 MHz, as shown in Fig. 5. Some experiments used all source/detector pairs, while others only used source/detector pairs on different faces of the cube. A zero-flux boundary condition on the outer boundary was imposed to approximate the physical boundary condition [52], [63], [14].

The homogeneous phantom had the constant values $\mu_a = 0.02 \text{ cm}^{-1}$ and $D = 0.03 \text{ cm}$. For the inhomogeneous phantom of Fig. 6(a), the μ_a background was linearly varied from 0.01 cm^{-1} to 0.04 cm^{-1} in the horizontal direction, except for the outermost region of width 1.25 cm, which was homogeneous with $\mu_a = 0.025 \text{ cm}^{-1}$. Two spherical μ_a inhomogeneities were centered along the z -axis of the phantom ($z = 0$) with values of $\mu_a = 0.1 \text{ cm}^{-1}$ (left-top) and $\mu_a = 0.12 \text{ cm}^{-1}$ (right-bottom). The diffusion coefficient D was homogeneous with $D = 0.03 \text{ cm}$. For both phantoms, the reconstruction was performed for all voxels except the eight, four, and two outermost layers of grid points for $65 \times 65 \times 65$, $33 \times 33 \times 33$, and $17 \times 17 \times 17$ reconstruction resolutions, respectively. These border regions were fixed to their true values in order to avoid singularities near the sources and detectors. These regions have also been excluded from all cross-section figures and the evaluation of root-mean-square (RMS) reconstruction error.

A. Evaluation of Required Forward Model Resolution

The objective of this section is to experimentally determine the forward model resolution required to produce a high quality reconstruction. To do this, we first evaluated the accuracy of




Fig. 5. (a) Source and (b) detector pattern on each face of the cube geometry. Two data set scenarios were considered: one containing all source/detector pairs, and a second containing only source/detector pairs on different faces.

the forward model as a function of resolution using the homogeneous phantom. The forward model PDE was first solved as resolutions corresponding to $129 \times 129 \times 129$, $65 \times 65 \times 65$, $33 \times 33 \times 33$, and $17 \times 17 \times 17$ grid points. We then computed the distortion-to-noise ratio (DNR) for two scenarios. The first scenario included all source/detector pairs, and the second only included source/detector pairs on different faces. This was done because the close proximity of source/detector pairs on the same face can result in susceptibility to discretization errors in the forward model. The DNR was computed as

$$DNR = \frac{2}{P} \sum_{i=1}^{P/2} \frac{|y_{(257)i} - y_{(N)i}|^2}{|y_{(257)i}|}, \quad (46)$$

where i is the index of source-detector pairs, $y_{(N)i}$ is the i -th simulated complex measurement with resolution N , and $P/2$ is the number of complex measurements. Notice that $|y_{(257)i}|$ is proportional to the assumed noise variance defined in (2) and (32).

Table I lists the DNR as a function of resolution for the two scenarios. Notice that for all resolutions the DNR is uniformly higher when source/detector pairs on the same face are included. As expected, the DNR also monotonically decreases as the resolution of the forward model is increased.

Next, we examined the reconstruction quality as a function of resolution using the inhomogeneous phantom. Figure 6 shows a $z = 0$ cross-section of the original phantom and the corresponding reconstructions for a variety of resolutions and data set scenarios.⁵ Each recon-

⁵These reconstructions were all produced using the multigrid algorithm with the mean phantom value as the initial condition

Forward Model Resolution	Distortion-to-noise ratio	
	All measurements	Source/detector pairs on different faces
$17 \times 17 \times 17$	6.74×10^{-4}	9.96×10^{-7}
$33 \times 33 \times 33$	9.66×10^{-5}	2.85×10^{-8}
$65 \times 65 \times 65$	2.44×10^{-6}	3.35×10^{-9}
$129 \times 129 \times 129$	1.74×10^{-6}	1.04×10^{-10}

TABLE I

Distortion-to-noise (DNR) ratio for various forward model resolutions. Coarse discretization increased forward model error, and source/detector pairs on the same face had much higher DNR.

Resolution/Data Set	σ	RMS image error
$65 \times 65 \times 65/\text{diff. faces}$	0.018	0.0069
$33 \times 33 \times 33/\text{diff. faces}$	0.008	0.0079
$17 \times 17 \times 17/\text{diff. faces}$	0.004	0.0093
$65 \times 65 \times 65/\text{all}$	0.03	0.0099

TABLE II

The normalization parameter σ that yields the best reconstruction and the resulting RMS image error between the reconstructions and the decimation of the true phantom.

struction used $p = 1.2$, but the value of σ was varied in the range of 0.002 to 0.12, in order to minimize the RMS image error between the reconstructions and the decimation of the true phantom. The parameters and the resulting RMS errors are summarized in Table II.

Figure 6 is consistent with the DNR measurement. The $65 \times 65 \times 65$ reconstruction from source/detector pairs on different faces has the best quality. Reconstructions at lower resolutions degrade rapidly, with very poor quality at $17 \times 17 \times 17$ resolution. Perhaps it is surprising that even the $65 \times 65 \times 65$ resolution reconstruction fails when all source/detector pairs are because in each case this method best minimized the relevant cost functional.




Fig. 6. A cross-section through (a) the inhomogeneous phantom, and the best reconstructions obtained using source detector pairs on different faces with (b) $65 \times 65 \times 65$ resolution, (c) $33 \times 33 \times 33$ resolution, (d) $17 \times 17 \times 17$ resolution, and (e) all source detector pairs with $65 \times 65 \times 65$ resolution. The $65 \times 65 \times 65$ reconstruction with different face source/detector pairs produced substantially better quality reconstruction. Reconstructions using all source/detector pairs failed even at $65 \times 65 \times 65$ reconstruction.

used. This result emphasizes the importance of using sufficiently high resolution, particularly when source/detector pairs are closely spaced.

B. Multigrid Performance Evaluation

Based on the results of Section IV-A, all comparisons of fixed-grid and multigrid inversion algorithms were performed for the $65 \times 65 \times 65$ resolution using only source/detector pairs on different faces. Our simulations compared fixed-grid inversion with multigrid inversion using 2, 3, and 4 levels of resolution. In order to make fair comparisons of computational speed, we scale the number of iterations for all methods into units of single fixed grid iterations at the finest scale. To do this, we use the approximate theoretical number of multiplies and

Algorithm		Parameters				Theoretical		Experimental
		$\nu^{(0)}$	$\nu^{(1)}$	$\nu^{(2)}$	$\nu^{(3)}$	Multiplications ($\times 10^6$)	Relative Complexity	Relative Complexity
Fixed-grid		1	.	.	.	4,020	1	1
Multigrid-V	2 levels	1	20	.	.	14,069	4.31	4.96
	3 levels	1	10	40	.	11,557	3.79	4.56
	4 levels	1	8	40	60	11,023	3.67	4.60

TABLE III

Complexity comparison for each algorithm. Theoretical complex multiplications are estimated with (53), and experimental relative complexity is the ratio of user time required for one iteration.

the corresponding relative complexity shown in Table III. However, we note that Table III indicates that the theoretical complexity of the multigrid iterations was somewhat lower than the experimentally measured complexity. See Appendix II for details of this conversion. All reconstructions were done using the inhomogeneous phantom and a prior model with $p = 1.2$ and $\sigma = 0.018 \text{ cm}^{-1}$. Gaussian shot noise was added to the data, so that the average signal-to-noise ratio for sources and detectors on opposite faces was 35 dB.

For the first experiment, all algorithms were initialized with the average values of the true phantom, which were $\mu_a = 0.026 \text{ cm}^{-1}$ and $D = 0.03 \text{ cm}$.⁶ Figure 7 shows that the multigrid algorithms converged much faster than the fixed grid algorithm, both in the sense of cost and RMS error. The multigrid algorithms converged in only 20 iterations, while the fixed algorithm required 270 iterations. Even after 200 iterations, the fixed grid algorithm still changed very little in the convergence plots.

Figure 8 shows reconstructions produced by the four algorithms. The reconstructed image quality for all three multigrid algorithms is nearly identical, but the reconstructed quality is significantly worse for the fixed grid algorithm. In fact, the multigrid algorithms converged to

⁶In practice, this is not possible since the average value is not known, but it was done because it favors the fixed-grid algorithm.




Fig. 7. Convergence of (a) cost function and (b) RMS image error when reconstructions were initialized with average values of true phantom. All multigrid algorithms converges about 13 times faster than the fixed-grid algorithm.

slightly lower values of the cost functional (-39833 to -39763) than the fixed-grid algorithm (-39392), and the RMS image error for the multigrid reconstructions ranged from 0.0069 to 0.007, while the fixed algorithm converged to the higher RMS error of 0.0081.

To investigate the sensitivity of convergence with respect to initialization, we performed reconstructions with a poor initial estimate. The initial image was homogeneous, with a value of 1.75 times the true phantom's average value. The plots in Fig. 9 show that the three and four level multigrid algorithms converged rapidly. In particular, the four level multigrid algorithm converges almost as rapidly as it did when initialized with the true phantom's average value. The fixed grid algorithm changed very little from the initial estimate even after 300 iterations, and the two grid algorithm progressed slowly. These results suggest that higher level multigrid algorithms are necessary to overcome the effects of a poor initial estimate.






Fig. 8. Cross-sections of reconstructions on the $z = 0$ plane using (a) 4 level multigrid with 19.35 iterations, (b) 3 level multigrid with 19.95 iterations, (c) 2 level multigrid with 18.24 iterations, and (d) 270 fixed grid iterations. All the multigrid reconstructions have better image quality than the fixed grid reconstruction.



(a)



(b)

Fig. 9. Convergence of (a) cost function and (b) RMS image error with a poor initial guess. For higher level multigrid algorithms, the convergence was faster. In particular, the four level multigrid algorithm converged almost as fast as when the reconstruction was initialized with the true phantom's average value.

V. CONCLUSIONS

We have proposed a nonlinear multigrid inversion algorithm which works by simultaneously varying the resolution of both the forward model and inverse computation. Multigrid inversion is formulated in a general framework and is applicable to a wide variety of inverse problems, but it is particularly well suited for the inversion of nonlinear forward problems such as those modeled by the solution of PDEs.

We performed experimental simulations for the application of multigrid inversion to optical diffusion tomography. These simulations indicate that multigrid inversion can dramatically reduce computation, particularly if the reconstruction resolution is high, and the initial condition is inaccurate. Perhaps more importantly, multigrid inversion showed robust convergence under a variety of conditions and while solving an optimization problem that is subject to local minima. Our experiments also indicated the importance of adequate resolution in the forward model.

REFERENCES

- [1] D. A. Boas, D. H. Brooks, E. L. Miller, C. A. DiMarzio, M. Kilmer, R. J. Gaudette, and Q. Zhang, "Imaging the body with diffuse optical tomography," *IEEE Signal Proc. Magazine*, vol. 18, no. 6, pp. 57–75, Nov. 2001.
- [2] S. R. Arridge, "Optical tomography in medical imaging," *Inverse Problems*, vol. 15, pp. R41–R93, 1999.
- [3] J. C. Hebden, S. R. Arridge, and D. T. Delpy, "Optical imaging in medicine: I. experimental techniques," *Phys. Med. Biol.*, vol. 42, pp. 825–840, 1997.
- [4] S. R. Arridge and J. C. Hebden, "Optical imaging in medicine: II. Modelling and reconstruction," *Phys. Med. Biol.*, vol. 42, pp. 825–840, 1997.
- [5] G. J. Saulnier, R. S. Blue, J. C. Newell, D. Isaacson, and P. M. Edic, "Electrical impedance tomography," *IEEE Signal Proc. Magazine*, vol. 18, no. 6, pp. 31–43, Nov. 2001.
- [6] J. S. Reynolds, A. Przada, S. Yeung, and K. J. Webb, "Optical diffusion imaging: a comparative numerical and experimental study," *Appl. Optics*, vol. 35, no. 19, pp. 3671–3679, July 1996.
- [7] E. C. Fear, S. C. Hagness, P. M. Meaney, M. Okoniewski, and M. A. Stuchly, "Enhancing breast tumor detection with near field imaging," *IEEE Microwave Magazine*, pp. 48–56, March 2002.
- [8] R. L. Thomas, "Reflections of a thermal wave imager: tow decades of research in photoacoustics and photothermal phenomena," *Analytical Sciences*, vol. 17, pp. s1–s4, April 2001.
- [9] R. Pike and P. Sabatier, *Scattering: Scattering and Inverse Scattering in Pure and Applied Science*, Academic Press, San Diego, 2002.
- [10] K. Sauer and C. A. Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. on Signal Processing*, vol. 41, no. 2, pp. 534–548, February 1993.
- [11] M. V. Ranganath, A. P. Dhawan, and N. Mullani, "A multigrid expectation maximization reconstruction algorithm for positron emission tomography," *IEEE Trans. on Medical Imaging*, vol. 7, no. 4, pp. 273–278, Dec. 1988.
- [12] T. Pan and A. E. Yagle, "Numerical study of multigrid implementations of some iterative image reconstruction algorithms," *IEEE Trans. on Medical Imaging*, vol. 10, no. 4, pp. 572–588, Dec. 1991.
- [13] A. B. Milstein, S. Oh, J. S. Reynolds, K. J. Webb, C. A. Bouman, and R. P. Millane, "Three-dimensional Bayesian optical diffusion tomography using experimental data," *Opt. Letters*, vol. 27, pp. 95–97, Jan. 2002.
- [14] S. Oh, A. B. Milstein, R. P. Millane, C. A. Bouman, and K. J. Webb, "Source-detector calibration in three-dimensional bayesian optical diffusion tomography," *J. Optical Society America A*, vol. 19, no. 10, pp. 1983–1993, Oct. 2002.
- [15] A. B. Milstein, S. Oh, K. J. Webb, C. A. Bouman, Q. Zhang, D. A. Boas, and R. P. Millane, "Fluorescence optical diffusion tomography," *Applied Optics*, submitted.
- [16] B. Sahiner and A. Yagle, "Image reconstruction from projections under wavelet constraints," *IEEE Trans. on Signal Processing*, vol. 41, no. 12, pp. 3579–3584, 1993.
- [17] M. Bhatia, W. C. Karl, and A. S. Willsky, "Wavelet-based method for multiscale tomographic reconstruction," *IEEE Trans. on Medical Imaging*, vol. 15, no. 1, pp. 92–101, 1996.
- [18] M. Bhatia, W. C. Karl, and A. S. Willsky, "Tomographic reconstruction and estimation based on multiscale natural -pixel bases," *IEEE Trans. on Image Processing*, vol. 6, no. 3, pp. 463–478, March 1997.
- [19] N. Lee, *Wavelet-vaguelette decompositions and homogenous equations*, Ph.D. dissertation, Purdue University, West Lafayette, IN, 1998.
- [20] A. Delaney and Y. Bresler, "Multiresolution tomographic reconstruction using wavelets," *IEEE Trans. on Image Processing*, vol. 4, no. 6, pp. 799–813, June 1995.

- [21] Z. Wu, G. T. Herman, and J. A. Browne, "Edge preserving reconstruction using adaptive smoothing in wavelet domain," in *Proc. of IEEE Nucl. Sci. Symp. and Med. Imaging Conf.*, San Francisco, California, October 31 - November 6 1993, vol. 3, pp. 1917–1921.
- [22] S. S. Saquib, C. A. Bouman, and K. Sauer, "A non-homogeneous MRF model for multiresolution Bayesian estimation," in *Proc. of IEEE Int'l Conf. on Image Proc.*, Lausanne Switzerland, September 16-19 1996, vol. 2, pp. 445–448.
- [23] R. Nowak and E. D. Kolaczyk, "A multiscale MAP estimation method for Poisson inverse problems," in *Proceedings of the 32nd Asilomar Conference on Signals, Systems & Computers*, Pacific Grove, CA, November 1-4 1998, vol. 2, pp. 1682–1686.
- [24] T. Frese, C. A. Bouman, and K. Sauer, "Adaptive wavelet graph model for Bayesian tomographic reconstruction," *IEEE Trans. on Image Processing*, vol. 11, no. 7, pp. 756–770, July 2002.
- [25] E. L. Miller, L. Nicolaidis, and A. Mandelis, "Nonlinear inverse scattering methods for thermal wave slice tomography: A wavelet domain approach," *J. Optical Society America A*, vol. 15, no. 6, pp. 1545–1556, June 1998.
- [26] W. Zhu, Y. Wang, Y. Deng, Y. Yao, and R. Barbour, "A wavelet-based multiresolution regularization least squares reconstruction approach for optical tomography," *IEEE Trans. on Medical Imaging*, vol. 16, no. 2, pp. 210–217, April 1997.
- [27] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A Multigrid Tutorial, 2nd Ed.*, Society for Industrial and Applied Mathematics, Philadelphia, 2000.
- [28] S. McCormick, Ed., *Multigrid Methods*, Society for Industrial and Applied Mathematics, Philadelphia, 1987.
- [29] W. Hackbusch, *Multigrid Methods and Applications*, Springer Series in Computational Mathematics. Springer-Verlag, Berlin, 1985.
- [30] P. Wesseling, *An Introduction to Multigrid Methods*, John Wiley & Sons, Chichester, 1992.
- [31] D. Terzopoulos, "Image analysis using multigrid relaxation methods," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 2, pp. 129–139, March 1986.
- [32] E. Enkelmann, "Investigations of multigrid algorithms for the estimation of optical flow fields in image sequences," *Comput. Vision Graphics and Image Process.*, vol. 43, pp. 150–177, 1988.
- [33] E. Memin and P. Perez, "Dense estimation and object-based segmentation of the optical flow with robust techniques," *IEEE Trans. on Image Processing*, vol. 7, no. 5, pp. 703–719, May 1998.
- [34] P. Hellier, C. Barillot, E. Mémin, and P. Pérez, "Hierarchical estimation of a dense deformation field for 3-d robust registration," *IEEE Trans. on Medical Imaging*, vol. 20, no. 5, pp. 388–402, May 2001.
- [35] P. Saint-Marc, J. Chen, and G. Medioni, "Adaptive smoothing: a general tool for early vision," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 13, no. 6, pp. 514–529, June 1991.
- [36] D. L. Pham and J. L. Prince, "Adaptive fuzzy segmentation of magnetic resonance images," *IEEE Trans. on Medical Imaging*, vol. 18, no. 9, pp. 737–752, Sep. 1999.
- [37] S. Henn and K. Witsch, "A multigrid approach for minimizing a nonlinear functional for digital image matching," *Computing*, vol. 64, pp. 339–348, 2000.
- [38] K. Zhou and C. K. Rushforth, "Image restoration using multigrid methods," *Appl. Optics*, vol. 30, no. 20, pp. 2906–2912, July 1991.
- [39] S. T. Acton, "Multigrid anisotropic diffusion," *IEEE Trans. on Image Processing*, vol. 7, no. 3, pp. 280–291, March 1998.
- [40] C. A. Bouman and K. Sauer, "Nonlinear multigrid methods of optimization in Bayesian tomographic image reconstruc-

- tion,” in *Proc. of SPIE Conf. on Neural and Stochastic Methods in Image and Signal Processing*, San Diego, CA, July 19-24 1992, vol. 1766, pp. 296–306.
- [41] S. F. McCormick and J. G. Wade, “Multigrid solution of a linearized, regularized least-squares problem in electrical impedance tomography,” *Inverse Problems*, vol. 9, pp. 697–713, 1993.
- [42] L. Borcea, “Nonlinear multigrid for imaging electrical conductivity and permittivity at low frequency,” *Inverse Problems*, vol. 17, pp. 329–359, April 2001.
- [43] C. R. Johnson, M. Mohr, U. Ruede, A. Samsonov, and K. Zyp, “Multilevel methods for inverse bioelectric field problems,” in *Lecture Notes in Computational Science and Engineering - Multiscale and Multiresolution Methods: Theory and Applications*, Eds. T.J. Barth, T.F. Chan, R. Haimes., Springer-Verlag Publishing, Heidelberg, Oct 2001, vol. 20.
- [44] J. C. Ye, C. A. Bouman, R. P. Millane, and K. J. Webb, “Nonlinear multigrid optimization for Bayesian diffusion tomography,” in *Proc. of IEEE Int’l Conf. on Image Proc.*, Kobe, Japan, October 25-28 1999.
- [45] J. C. Ye, C. A. Bouman, K. J. Webb, and R. P. Millane, “Nonlinear multigrid algorithms for Bayesian optical diffusion tomography,” *IEEE Trans. on Image Processing*, vol. 10, no. 6, pp. 909–922, June 2001.
- [46] S. G. Nash, “A multigrid approach to discretized optimization problems,” *J. of Optimization methods and software*, vol. 14, pp. 99–116, 2000.
- [47] R. M. Lewis and S. G. Nash, “A multigrid approach to the optimization of systems governed by differential equations,” in *8-th AIAA/USAF/ISSMO Symp. Multidisciplinary Analysis and Optimization*, Long Beach, CA, 2000.
- [48] S. S. Saquib, K. M. Hanson, and G. S. Cunningham, “Model-based image reconstruction from time-resolved diffusion data,” in *Proc. of SPIE Conf. on Medical Imaging: Image Processing*, Newport Beach, CA, February 25-28 1997, vol. 3034, pp. 369–380.
- [49] A. H. Hielscher, A. D. Klose, and K. M. Hanson, “Gradient-based iterative image reconstruction scheme for time-resolved optical tomography,” *IEEE Trans. on Medical Imaging*, vol. 18, no. 3, March 1999.
- [50] A. Mohammad-Djafari, “Joint estimation of parameters and hyperparameters in a Bayesian approach of solving inverse problems,” in *Proc. of IEEE Int’l Conf. on Image Proc.*, Lausanne, Switzerland, September 16-19 1996, vol. II, pp. 473–476.
- [51] A. Mohammad-Djafari, “On the estimation of hyperparameters in Bayesian approach of solving inverse problems,” in *Proc. of IEEE Int’l Conf. on Acoust., Speech and Sig. Proc.*, Minneapolis, Minnesota, April 27-30 1993, pp. 495–498.
- [52] J. C. Ye, K. J. Webb, C. A. Bouman, and R. P. Millane, “Optical diffusion tomography using iterative coordinate descent optimization in a Bayesian framework,” *J. Optical Society America A*, vol. 16, no. 10, pp. 2400–2412, October 1999.
- [53] L. E. Baum and T. Petrie, “Statistical inference for probabilistic functions of finite state Markov chains,” *Ann. Math. Statistics*, vol. 37, pp. 1554–1563, 1966.
- [54] L. Baum, T. Petrie, G. Soules, and N. Weiss, “A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains,” *Ann. Math. Statistics*, vol. 41, no. 1, pp. 164–171, 1970.
- [55] A.R. De Pierro, “A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography,” *IEEE Trans. on Medical Imaging*, vol. 14, no. 1, pp. 132–137, March 1995.
- [56] J. Browne and A. R. De Pierro, “A row-action alternative to the EM algorithm for maximizing likelihoods in emission tomography,” *IEEE Trans. on Medical Imaging*, vol. 15, no. 5, pp. 687–699, October 1996.
- [57] J.A. Fessler, E.P. Ficaro, N.H. Clinthorne, and K. Lange, “Grouped-coordinate ascent algorithms for penalized-

- likelihood transmission image reconstruction,” *IEEE Trans. on Medical Imaging*, vol. 16, no. 2, pp. 166–175, April 1997.
- [58] J. Zheng, S. S. Saquib, K. Sauer, and C. A. Bouman, “Parallelizable Bayesian tomography algorithms with rapid, guaranteed convergence,” *IEEE Trans. on Image Processing*, vol. 9, no. 10, pp. 1745–1759, Oct. 2000.
- [59] J. Besag, “Spatial interaction and the statistical analysis of lattice systems,” *Journal of the Royal Statistical Society B*, vol. 36, no. 2, pp. 192–236, 1974.
- [60] T. Hebert and R. Leahy, “A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors,” *IEEE Trans. on Medical Imaging*, vol. 8, no. 2, pp. 194–202, June 1989.
- [61] D. Geman and G. Reynolds, “Constrained restoration and the recovery of discontinuities,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 14, no. 3, pp. 367–383, March 1992.
- [62] C. A. Bouman and K. Sauer, “A generalized Gaussian image model for edge-preserving MAP estimation,” *IEEE Trans. on Image Processing*, vol. 2, no. 3, pp. 296–310, July 1993.
- [63] J. C. Ye, K. J. Webb, R. P. Millane, and T. J. Downar, “Modified distorted Born iterative method with an approximate Fréchet derivative for optical diffusion tomography,” *J. Optical Society America A*, vol. 16, no. 7, pp. 1814–1826, July 1999.
- [64] J. C. Adams, “MUDPACK: Multigrid portable FORTRAN software for the efficient solution of linear elliptic partial differential equations,” *Appl. Math. Comput.*, vol. 34, pp. 113–146, 1989.

APPENDIX I

PROOF OF MULTIGRID MONOTONE CONVERGENCE

We begin with two lemmas which give sufficient conditions to guarantee monotone decrease of the finer grid cost functional in the two-grid algorithm. All lemmas assume that the functions $\tilde{c}^{(q)}(\cdot)$ and $\tilde{c}^{(q+1)}(\cdot)$ are continuously differentiable.

Lemma 1: Assume that the following conditions are satisfied for a resolution $q \geq 0$:

- 1) The fixed grid update is monotone at resolutions q and $q + 1$.
- 2) A functional $\eta^{(q+1)} : \mathbb{R}^{N^{(q+1)}} \rightarrow \mathbb{R}$, defined by

$$\eta^{(q+1)}(x^{(q+1)}) = c^{(q+1)}(x^{(q+1)}) - c^{(q)}(x^{(q)} + I_{(q+1)}^{(q)}(x^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)})), \quad (47)$$

has a global minimum at $x^{(q+1)} = I_{(q)}^{(q+1)}x^{(q)}$, where $x^{(q)}$ is the value resulting after the initial $\nu_1^{(q)}$ fixed grid iterations.

- 3) $\nu_1^{(q)} + \nu_2^{(q)} \geq 1$.

Then, the two-grid inversion algorithm of Fig. 2 is monotone for the functional $c^{(q)}(\cdot)$.

Proof of Lemma 1:

By the definition of monotonicity, the updated value $\tilde{x}^{(q+1)}$ of (9) satisfies

$$c^{(q+1)}(\tilde{x}^{(q+1)}) \leq c^{(q+1)}(I_{(q)}^{(q+1)}x^{(q)}). \quad (48)$$

Applying the definition of $\eta^{(q+1)}(\cdot)$ and the second condition, we have

$$\begin{aligned} \eta^{(q+1)}(\tilde{x}^{(q+1)}) &\geq \eta^{(q+1)}(I_{(q)}^{(q+1)}x^{(q)}) \\ c^{(q+1)}(\tilde{x}^{(q+1)}) - c^{(q)}(x^{(q)} + I_{(q+1)}^{(q)}(\tilde{x}^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)})) &\geq c^{(q+1)}(I_{(q)}^{(q+1)}x^{(q)}) - c^{(q)}(x^{(q)}) \end{aligned} \quad (49)$$

From the inequalities (48) and (49), it follows that

$$c^{(q)}(x^{(q)} + I_{(q+1)}^{(q)}(\tilde{x}^{(q+1)} - I_{(q)}^{(q+1)}x^{(q)})) \leq c^{(q)}(x^{(q)}). \quad (50)$$

This inequality means that the coarse grid update and its subsequent coarse grid correction decreases the cost functional $c^{(q)}(\cdot)$. With the first condition, this guarantees the inequality in the definition of monotone convergence for $c^{(q)}(\cdot)$. Furthermore, by the first and fourth conditions, if $\nabla c^{(q)}(x^{(q)}) \neq 0$, the update at resolution q either before or after the coarse grid

update strictly decreases $c^{(q)}(\cdot)$. Therefore, the two-grid algorithm is monotone under these assumptions.

Lemma 2. (Two-Grid Monotone Convergence)

Assume that the following conditions are satisfied for a resolution $q \geq 0$:

- 1) The fixed grid update is monotone at resolutions q and $q + 1$.
- 2) $\xi^{(q+1)}(\cdot)$ is convex on $\mathbb{R}^{N^{(q+1)}}$.
- 3) The adjustment vector $r^{(q+1)}$ is given by (15).
- 4) $\nu_1^{(q)} + \nu_2^{(q)} \geq 1$.

Then, the two-grid inversion algorithm of Fig. 2 is monotone for the functional $c^{(q)}(\cdot)$.

Proof of Lemma 2:

It is enough to show that the third and second conditions of this lemma imply the second condition in Lemma 1. By condition three, we know that

$$\eta^{(q+1)}(x^{(q+1)}) = \xi^{(q+1)}(x^{(q+1)}) + vx^{(q+1)} + \text{constant} \quad (51)$$

for some row vector v of length $N^{(q+1)}$. In fact, we know that equation (15) selects the vector v so that the gradients of the coarse and fine scale cost functionals are matched, and therefore

$$\nabla \eta^{(q+1)}(x^{(q+1)}) \Big|_{x^{(q+1)} = I_{(q)}^{(q+1)} x^{(q)}} = 0. \quad (52)$$

By (51) we also know that $\eta^{(q+1)}(\cdot)$ is a continuously differentiable convex function. Therefore, we know that $\eta^{(q+1)}(\cdot)$ must take on its global minimum value at $x^{(q+1)} = I_{(q)}^{(q+1)} x^{(q)}$.

Proof of Multigrid Monotone Convergence Theorem:

Our proof is by induction. Consider the case when $q = Q - 2$, then we have the two-grid case, and the proof is trivial by Lemma 2.

Now consider $q < Q - 2$. By induction, we assume that the Multigrid-V algorithm applied at resolution $q + 1$ is monotone for the function $c^{(q+1)}(\cdot)$. This then meets condition 1 of Lemma 2, since the multigrid algorithm serves as the coarse grid optimizer in a two-grid algorithm. Therefore, the Multigrid-V algorithm applied at resolution q is monotone for the function $c^{(q)}(\cdot)$, and the induction is complete.

APPENDIX II

COMPUTATIONAL COMPLEXITY

In this appendix, we compare the computational cost of the proposed multigrid inversion algorithm for ODT problems with that of the fixed-grid ICD algorithm [52]. We use the number of complex multiplications required for one iteration of the V-cycle algorithm as a measure of computational complexity.

First, let us consider the computation required for one iteration of Fixed_Grid_Update(). Here, we use the analysis from [52]. Assuming F iterations are used for the linear PDE solver, the computation of Green's functions of (40) and (41) needs $5(K + M)FN_0$ multiplications, where N_0 is the number of grid points in the PDE domain. Then, we need PN and $\frac{5}{2}PN$ multiplications to compute (39) and (44), respectively, where N is the updated image size.⁷ Thus, the total computational cost for one iteration of the ICD fixed-grid update is $5(K + M)FN_0 + \frac{7}{2}PN$ multiplications.

Now, let us estimate the computation required for one iteration of MultigridV() which operates at resolutions $0, \dots, Q - 1$. For simplicity, we neglect the computational cost required for decimation and interpolation of images and the correction vector. In other words, we assume that the main computational cost at resolution q consists of the fixed-grid update on $x^{(q)}$ and the computation of $r^{(q)}$. To update $x^{(q)}$, one iteration of MultigridV() involves $\nu^{(q)} = \nu_1^{(q)} + \nu_2^{(q)}$ iterations of Fixed_Grid_Update(), which requires $[5(K + M)FN_0^{(q)} + \frac{7}{2}PN^{(q)}]\nu^{(q)}$ multiplications. Since $N_0^{(q)} = 8^{-q}N_0$ and $N^{(q)} = 8^{-q}N$ in 3-D problems, this is equal to $8^{-q} \times [5(K + M)FN_0 + \frac{7}{2}PN]\nu^{(q)}$ multiplications.

The correction vector $r^{(q)}$ is computed only once when the inversion proceeds from resolution q to $q + 1$. Since $g^{(q+1)}$ is computed in the optimization for the update of $x^{(q+1)}$, the only additional computation for $r^{(q)}$ is computation of $g^{(q)}$ given by (17). To compute $g^{(q)}$, we first compute the Green's functions of (40) and (41) and then use them to compute Fréchet derivatives by (39), which requires $5(K + M)FN_0^{(q)}$ and $PN^{(q)}$ multiplications, respectively [52]. Then, $PN^{(q)}$ multiplications are required to evaluate the expression in the brace of (17). The resulting total complexity for computation of $r^{(q)}$ is $8^{-q} \times [5(K + M)FN_0 + 2PN]$ mul-

⁷In Section IV, we do not update the outermost region to avoid singularity problems, so N and N_0 are different in this case.

tuplications.

Thus, for resolutions $q = 0, \dots, Q - 2$, the total complexity of the Multigrid-V algorithm is $8^{-q} \times \left[\{5(K + M)FN_0 + \frac{7}{2}PN\} \nu^{(q)} + \{5(K + M)FN_0 + 2PN\} \right]$ multiplications. At the coarsest resolution $q = Q - 1$, we do not need $r^{(Q-1)}$, so the complexity is $8^{-(Q-1)} \times \{5(K + M)FN_0 + \frac{7}{2}PN\} \nu^{(Q-1)}$ multiplications. Therefore, the total complexity for one Multigrid-V is

$$\sum_{q=0}^{Q-2} \left[8^{-q} \times \left\{ \left\{ 5(K + M)FN_0 + \frac{7}{2}PN \right\} \nu^{(q)} + \{5(K + M)FN_0 + 2PN\} \right\} \right] + 8^{-(Q-1)} \times \left\{ 5(K + M)FN_0 + \frac{7}{2}PN \right\} \nu^{(Q-1)}, \quad (53)$$

where K is the number of sources, M is the number of detectors, P is the number of measurements, N_0 is the PDE image size, N is the updated image size, F is the number of iterations required for the linear forward solver, and $\nu^{(q)}$ is the number of iterations of fixed grid update at resolution q .

Table III lists the estimated number of complex multiplications required for each iteration of the fixed-grid and Multigrid-V algorithms for typical values of parameters which we use in the simulations of Section IV-B. The values of the parameters are $K = 48$, $M = 54$, $P = 2160$, $N_0 = 65 \times 65 \times 65$, $N = 49 \times 49 \times 49$, and $F = 16$. We also provide the experimental computation time. One fixed-grid iteration took 55.5 minutes of user time on a Pentium-III 697 MHz Linux machine, and the complexity per iteration is $4.56 \sim 4.96$ times larger for the multigrid algorithm. However, one multigrid iteration involves many coarser grid iterations, and the simulation results show that the number of iterations required for the multigrid algorithms to converge is substantially less than is required using the fixed grid algorithm.