Progress In Electromagnetics Research B, Vol. 74, 93–107, 2017

Inversion of an Inductive Loss Convolution Integral for Conductivity Imaging

Joe R. Feldkamp*

Abstract—Electrical conductivity imaging in the human body is usually pursued by either electrical impedance tomography or magnetic induction tomography (MIT). In the latter case, multiple coils are almost always used, while nonlinear reconstruction is preferred. Recent work has shown that single-coil, scanning MIT is feasible through an analytical 3D convolution integral that relates measured coil loss to an *arbitrary* conductivity distribution. Because this relationship is *linear*, image reconstruction may proceed by any number of linear methods. Here, a direct method is developed that combines several strategies that are particularly well suited for inverting the convolution integral. These include use of a diagonal regularization matrix that leverages kernel behavior; transformation of the minimization problem to standard form, avoiding the need for generalized singular value decomposition (SVD); centering the quadratic penalty norm on the uniform solution that best explains loss data; use of KKT multipliers to enforce non-negativity and manage the rather small active set; and, assignment of the global regularization parameter via the discrepancy principle. The entire process is efficient, requiring only one SVD, and provides ample controls to promote proper localization of structural features. Two virtual phantoms were created to test the algorithm on systems comprised of $\sim 11,000$ degrees of freedom.

1. INTRODUCTION

Electrical conductivity is a property that varies considerably throughout the human body [1]. Thus, an ability to image its spatial variability could be useful from a diagnostic standpoint. Various disease states may cause the conductivity to differ from that exhibited in normal tissue [2, 3]; or, the appearance of growths and other obstructions [4] could alter the way that induced electrical currents behave so that the obstruction becomes visible. Two methods for producing these images in medical applications have appeared over the past few decades — one of them electrical impedance tomography [5], and the other, magnetic induction tomography [6] (MIT). The latter approach offers the considerable advantage of being non-contact as well as noninvasive. Nearly all MIT research has been based upon a multiple coil system placed in the vicinity of the conductive target. One coil functions as a source that interacts with the target, while any other functions as a detector, responding both to the primary field as well as the perturbation of the field due to the target's presence. This multi-coil approach is inherently nonlinear, involving a forward problem that must numerically solve Maxwell's equations prior to image reconstruction — and do so for each position of the coil array. Given that the forward problem is nonlinear, it has been stated that nonlinear image reconstruction is a necessity in order to achieve satisfactory results [7].

More recently, an alternate approach has been evaluated that requires the use of only one coil [8, 9], whose interaction with nearby conductive objects is manifested as an inductive loss, resulting from the secondary field generated by induced eddy currents [10]. The coil under consideration consists of closely-spaced parallel layers of concentric circular loops, all connected in series. Since electrical conductivity is

Received 14 February 2017, Accepted 1 April 2017, Scheduled 12 April 2017
* Corresponding author: Joe R. Feldkamp (jrfeldcinci@gmail.com).
The author is with the Kimberly-Clark Corp., Neenah, WI, USA.

quite low in biological targets ($\lessapprox 5\,\mathrm{S/m}$), rather high excitation frequencies ($\sim 10\,\mathrm{MHz}$) are needed in order to produce a measurable loss in such coils ($\sim 1\,\mathrm{ohm}$) and still maintain high signal to noise ratio. To leverage the usefulness of measured coil loss information, an analytical model was developed that expresses a *linear* relationship between inductive loss in coils consisting of circular loops and an *arbitrary* electrical conductivity distribution. The key assumptions behind the model are that permeability and permittivity are spatially uniform; conductivity is small ($\lessapprox 10\,\mathrm{S/m}$); excitation frequency is sufficiently low that quasi-static field conditions hold; and, that spacing between coil loops residing in different planes is sufficiently small that they essentially share a common plane. The model has been validated numerous times experimentally [8, 9, 11], showing that it is indeed quantitative when permittivity is uniform, and that linearity is preserved even when permittivity shows large sudden spatial variation beyond what is present in biological systems. Furthermore, the analytical model is easily written as a 3D convolution integral that does not require calibration, facilitating its use for image reconstruction.

Because there is just one coil, a scan operation is needed that repeatedly relocates the coil in the vicinity of the target such that the EM field of the coil can intercept the target in a large number of non-redundant ways. With inductive loss written as a convolution integral that convolves the electrical conductivity distribution with a sampling function (kernel), a sufficiently detailed scan should contain adequate information to enable computation of the desired electrical conductivity distribution. The advantages of this approach include reduced hardware complexity, since only one coil has to be managed; and, given the fact that the convolution integral shows a linear relationship between loss and the full conductivity distribution, image reconstruction is made easier.

Nevertheless, image reconstruction has remained challenging, mostly because of the rapid decay of the EM field in a coil's vicinity. But this is also true of coils used in any form of MIT and appears as a similar challenge in other forms of tomography, such as diffuse optical tomography where light scattering greatly attenuates the propagation of light into and out of a sample medium [12]. In our previous work with single coils, inversion proceeded by first discretizing the convolution integral on a finite element mesh that spans the space of the conductive target, so that coil loss $\vec{b}$ can be predicted through a simple matrix operation acting on the unknown conductivity vector $\vec{x}$: $\vec{b} \approx \tilde{A}\vec{x}$. Most commonly, the number of measurements is far smaller than the number of unknowns, necessitating the use of regularized least squares analysis. A solution was pursued by Tikhonov regularization that increasingly penalized those solutions having higher gradient norms. The usual "normal" equations were developed and then passed to an open source solver that imposed non-negativity on the solution via an active set approach developed by Haskell and Hanson [13]. An optimal value for the regularization parameter was chosen using an L-curve technique. Not only was the process inefficient, solutions showed excessive smoothing and were prone to misplace structural features in those circumstances where coil loss measurements were restricted to locations close to a target — both issues were found to be aggravated by instrumental noise.

Here, several individually known techniques are combined in a novel way that accomplishes inversion in a manner that is particularly well-suited to the coil-loss convolution integral. An invertible diagonal regularization matrix is chosen that uses the integral's kernel function to guide how individual solution components ought to be weighted in the penalty term; the penalty term itself is quadratic about a uniform conductivity solution that is simply the single conductivity value that best predicts the measured coil loss vector; since the regularization matrix is invertible, the penalized objective function can be put into "standard form", paving the way for methods based upon singular value decomposition (SVD) rather than GSVD — generalized SVD; non-negativity is enforced by a combination of Karush-Kuhn-Tucker (KKT) multipliers and active set methodology so that the number of iterations tends to be small; since the measurement error norm is known from experiment, the discrepancy principle is used to help guide selection of the global regularization parameter — essentially, the regularization parameter is stepped through a partial sequence of the singular values, progressing from largest to smallest, until solution error norm is near to, but not less than the measurement error norm. Only one SVD computation is required to complete the inversion. Solving for the KKT multipliers associated with the active set is efficient since the system is symmetric and much smaller in size than the number of degrees of freedom for the entire problem. Overall, the inversion is efficient, does not over smooth the image and allows sufficient flexibility to aid improved localization of structural features.

2. SVD NNLS APPLIED TO MIT CONVOLUTION INTEGRAL

The discretized convolution integral (next section) leads to a prediction of coil loss measurements from a conductivity vector $\vec{x}$ consisting of n conductivity values on a finite element (FE) mesh. For the application to single-coil MIT, coil loss measurements $\vec{b}$ are obtained at m scan positions in the vicinity of a conductive target. Coil loss prediction is given by the vector $\tilde{A}\vec{x}$, where the matrix $\tilde{A}$ is $m \times n$ and consists of components that are determined according to coil position and orientation. Measurements are expected to be few in number, such that $m < n$, giving an underdetermined system.

Because conductivity is non-negative, we require that $x_i \geq 0$ for each i. Since scanning inevitably produces many similar measurements, the matrix $\tilde{A}$ is expected to be ill-conditioned (large condition number). Thus, we attempt to find the best conductivity by minimizing a penalized 2-norm, subject to non-negativity:

$$\min \frac{1}{2}\left\|\tilde{A}\vec{x}-\vec{b}\right\|_2^2+\frac{1}{2}\tau^2\left\|\tilde{D}\left(\vec{x}-\vec{\beta}\right)\right\|_2^2 \quad s.t. \quad \vec{x}\geq 0 \tag{1}$$

Problem in Equation (1) is made more flexible, without requiring additional computational effort, by permitting each node to be penalized independently of the others. This is accomplished via the invertible, diagonal regularization matrix $\tilde{D}$, with diagonal entries $d_j > 0$. In order for the weights in the diagonal matrix to produce some benefit to the solution process, guidance on selection of the components of $\tilde{D}$ is essential — we return to this issue later. For now, we mention that the model might inherently weight some members of the solution set more so than others, so that $\tilde{D}$ offers an opportunity to compensate. Furthermore, if some solution members tend to contribute very little to a predicted b_i, or should have little impact on the error norm (first term of Equation (1)), then these could be penalized heavily to ensure a given x_i remains well inside an expected interval through appropriate assignment of β_j and d_j.

Commonly, $\vec{\beta} = 0$, but here we take advantage of the fact that the solution will typically lie within some interval $0 \leq x_i \lessapprox 5\,\mathrm{S/m}$ rather than an unbounded domain. For example, electrical conductivity within human tissues [1] is usually less than $\sim 5\,\mathrm{S/m}$, suggesting that $\vec{\beta}$ ought to be set somewhere inside this expected interval — leading to a "shifted penalty", as we characterize it here. Thus, the regularization term not only promotes a well-conditioned problem with a unique solution [14], but leads to a greater penalty when $x_i \neq \beta_i$. If the regularization parameter τ is made large compared to the largest singular values of $\tilde{A}$, then solutions tend to be heavily localized around $\vec{\beta}$ [15]. Thus, $\vec{\beta}$ is set > 0 so that the number of negative x_i values is reduced when constraints are absent, while τ is set to a non-zero value within the range of the singular values associated with $\tilde{A}$. Non-negativity is addressed using an active set method, but is made more efficient as a result of the shifted penalty, which tends to reduce the number of needed constraints, and thus the size of the active set. We proceed with these definitions to transform minimization problem in Equation (1) into standard form:

$$\vec{z} = \tilde{D}\vec{x} \;\; ; \;\; \vec{\zeta} = \tilde{D}\vec{\beta} \tag{2}$$

Minimization problem in Equation (1) becomes:

$$\min \frac{1}{2}\left\|\tilde{A}\tilde{D}^{-1}\vec{z}-\vec{b}\right\|_2^2+\frac{1}{2}\tau^2\left\|\vec{z}-\vec{\zeta}\right\|_2^2 \quad s.t. \quad \vec{z}\geq 0 \tag{3}$$

Except for the shifted penalty, problem in Equation (3) is now in standard form [16], showing how matrix $\tilde{A}$ is actually pre-processed by $\tilde{D}^{-1}$, which is meant to balance the role of each component of the solution, provided there is some meaningful approach to accomplishing that. As long as the computational cost is acceptable, getting to standard form is also feasible using square regularization matrices other than $\tilde{D}$, provided they are not singular [17]. The primary benefit of getting to standard form, initially proposed by Eldén [18], is improved solution efficiency.

SVD† is now applied to the matrix product appearing in Equation (3), as follows:

$$\tilde{A}\tilde{D}^{-1} = \tilde{U}\tilde{S}\tilde{V}^T \tag{4}$$

† LINPACK, 1979

Orthogonal matrices and singular values are modified due to the pre-multiplication of $\tilde{A}$ by the inverse of $\tilde{D}$ with the intention that solution components have a more equitable chance to contribute to the measurement vector in a noisy environment. To continue our development, we make the additional assignments:

$$\vec{z} = \tilde{V}\vec{y};\ \ \vec{b} = \tilde{U}\vec{b}';\ \ \zeta = \tilde{V}\vec{\beta}' \tag{5}$$

Substitution of Equation (5) into Equation (3) leads to a form for the minimization problem especially suitable for our MIT inversion problem since it is readily amenable to straightforward algebraic solution:

$$\min \frac{1}{2}\left\|\tilde{S}\vec{y} - \vec{b}'\right\|_2^2 + \frac{1}{2}\tau^2\left\|\vec{y} - \vec{\beta}'\right\|_2^2 \quad s.t. \quad \vec{x}\ or\ \vec{z} \geq 0 \tag{6}$$

After expanding the convex objective function in Equation (6) and adding non-negativity constraints for some select members of the solution set, $\{x_{l(k)};\ k = 1, 2, \ldots, K\}$, i.e., the active set, we have the Lagrange objective function:

$$L\left(\vec{y}, \vec{\lambda}\right) = \frac{1}{2}\sum_{i=1}^{r}\left(s_i y_i - b'_i\right)^2 + \frac{1}{2}\sum_{i=r+1}^{m}\left(b'_i\right)^2 + \frac{1}{2}\tau^2\sum_{i=1}^{n}\left(y_i - \beta'_i\right)^2 + \sum_{k=1}^{K}\lambda_k x_{l(k)} \tag{7}$$

Here, we assume that $r \leq m < n$, where the matrix $\tilde{A}$ has rank r; when $r = m$ (usually), the second term on the right side is absent. Constraints, enforced via the KKT multipliers $\{\lambda_k\}$, are only applied to those $\{x_{l(k)}\}$ belonging to the active set — note the mapping implied by the subscript that identifies the correct solution component index associated with the k^{th} constraint. To establish an initial active set, $L(\vec{y}, \vec{\lambda})$ is first minimized without constraints while $\vec{\beta}$ is chosen to deliberately reduce the size of the active set. Nodes with negative conductivity values are identified and placed into the active set for the first iteration following initialization. At each new iteration, KKT multipliers are examined, so that a constraint is released when positive, but maintained when negative. Of course, those nodes not in the active set continue to be examined and placed into the active set if their values are found to be negative.

Optimality equations are written for each of the unknowns in Equation (7), which requires that we set the partial derivative of L with respect to each of the y_j to zero. These may be solved for the $\{y_j\}$ in terms of the KKT multipliers:

$$y_j = \begin{cases} \dfrac{s_j b'_j + \tau^2\beta'_j - \sum_{k=1}^{K}\lambda_k V_{l(k)j}/d_{l(k)}}{s_j^2 + \tau^2};\ j \leq r \\ \\ \dfrac{\tau^2\beta'_j - \sum_{k=1}^{K}\lambda_k V_{l(k)j}/d_k}{\tau^2};\ r < j \leq n \end{cases} ;\quad k \rightarrow l\ mapping \tag{8}$$

The complementarity conditions, which directly enforce non-negativity on members belonging to the active set, are simpler:

$$x_{l(q)} = \sum_{j=1}^{n}\frac{V_{l(q)j}}{d_{l(q)}}y_j = 0; \quad q \rightarrow l\ \ mapping\ implied \tag{9}$$

Substitution of Equations (8) into Equation (9) allows us to obtain the KKT multipliers from a relatively small set of linear equations:

$$\sum_{k=1}^{K}P_{qk}\lambda_k = \sum_{j=1}^{n}\frac{V_{l(q)j}}{d_{l(q)}}y_j^0; \quad q = 1, 2, \ldots, K(\sim 1500) \ll n(\sim 11,000) \tag{10}$$

The elements for the symmetric matrix, P_{qk}, are:

$$P_{qk} = \frac{1}{d_{l(q)}d_{l(k)}}\sum_{j=1}^{n}\frac{V_{l(q)j}V_{l(k)j}}{s_j^2 + \tau^2};\ \ s_j = 0\ if\ j > r;\ q, k = 1, 2, 3, \ldots, K \tag{11}$$

The vector components of $\vec{y}^0$ are just the unconstrained solution, obtained by setting all KKT multipliers to zero in Equation (8). Equation (10) is solved using an LDL^T decomposition method, and is kept smaller in size if $\vec{\beta}$ is chosen within the expected solution interval and the regularization parameter τ is kept nearer to the larger singular values in the set $\{s_j\}$. By choosing τ nearer the middle of the set of singular values, the current solution approach behaves similarly to a truncated SVD method [15] since those singular values which are much smaller than τ have little or no effect on the solution (filtering) — note that all singular values are used here and regularization is controlled by choice of τ rather than truncation. Since the active set may change "membership" with each iteration, matrix $\tilde{P}$ has to be recomputed, but only for those components new to the active set.

Subsequent to solving Equation (10), the KKT multipliers are returned to Equation (8) to get the new solution (either $\vec{z}$ or $\vec{x}$). At that point, both the new solution and KKT multipliers are examined in order to establish a new active set for the next iteration. Though not following the same active set method developed by Haskell and Hanson [13], which adds or drops only one constraint per iteration, we expect similar behavior — i.e., convergence in a finite number of steps [[19], p. 114]. In fact, for the MIT application we consider in the next section, convergence is typically obtained after 5 to 7 iterations — determined by the attainment of a stable active set, with KKT conditions satisfied.

This algorithm requires a choice for $\vec{\beta}$, which establishes a location where the penalty term is zero — optionally, individual components of the penalty may be forced to zero at customized locations, without requiring that the full penalty term is zero at a common location. Though many approaches to setting this vector are worth investigating, here we choose a uniform solution to Equation (1) with penalty term dropped. This sets all solution components to an identical value, β_0, which is easily found by minimizing the error residual with respect to β_0:

$$\beta_0 = \frac{\sum_{i=1}^{m} b_i \sum_{j=1}^{n} A_{ij}}{\sum_{i=1}^{m} \left(\sum_{j=1}^{n} A_{ij} \right)^2} \tag{12}$$

For our MIT application, β_0 becomes the conductivity of a target treated as having no structure, but rather characterized by uniform conductivity. Certainly, it is possible to consider modifications to $\vec{\beta}$ as the iteration process continues — for example, using a previous solution, or even an average of two prior solutions, with the idea that the next solution in the sequence is more penalized as it departs further from a prior solution [20, 21]. This is not explored here — we just use Equation (12) as a "zero order", uniform solution.

To complete the algorithm, a value for the global regularization parameter τ is needed. This value is chosen using the so-called discrepancy principle [17], which requires that some reasonable estimate is available for the error associated with the measurement vector $\vec{b}$:

$$\left\| \vec{b} - \vec{b}_{true} \right\| \approx \varepsilon \tag{13}$$

Since this is something that we have measured for our application as the technology has progressed, the error ϵ is available. Thus, the regularization parameter τ is decreased along a sequence tracking the singular values arising from the SVD of Equation (4), usually jumping several singular values at a time, until the following condition is approximately met:

$$\left\| \tilde{A}\vec{x} - \vec{b} \right\| \approx \varepsilon \tag{14}$$

The rationale for this stopping condition is that it would be unacceptable for the model to predict a measurement vector having greater precision than what is available via experiment. In other words, the quality of the measured data does not support further decrease in τ. There is minimal computational burden involved here since once the model matrix is built and the SVD is obtained, there is no need to redo them. Furthermore, starting from a τ value which has been set to the largest and then proceeds to smallest singular values (i.e., toward less filtering) leads to a sequence of systems for the KKT multipliers that start out very small and gradually grow in size toward the stopping condition (13) or (14). When

ϵ becomes very small (much smaller than what is currently available via experiment), the solution step for KKT multipliers may become burdensome.

3. 3D CONVOLUTION FOR SINGLE-COIL, SCANNING MIT

Previously [8, 9], we have shown that the loss δZ in a planar coil consisting of circular concentric loops of wire is related to the electrical conductivity distribution of the target placed in the vicinity of the coil by:

$$\delta Z = \frac{\mu^2\omega^2}{4\pi^2}\sum_{j,k}\sqrt{\rho_j\rho_k}\int d^3x\frac{\sigma(\vec{r})}{\rho}Q_{1/2}\left(\eta_j\right)Q_{1/2}\left(\eta_k\right) \tag{15}$$

Arguments for the circularly symmetric toroid (or ring) function $Q_{1/2}$ lie in the interval $1 < \eta < \infty$ and are related to field position by:

$$\eta_j = \frac{\rho^2+\rho_j^2+z_c^2}{2\rho\rho_j};\quad \eta_k = \frac{\rho^2+\rho_k^2+z_c^2}{2\rho\rho_k} \tag{16}$$

Using any suitable fixed laboratory coordinate system, other symbols in Equations (15) and (16) are defined by:
$\sigma(\vec{r})$: *Electrical conductivity (real part) at field position:* $\vec{r} = x, y, z$,
ρ_k: *Cylindrical radial distance from coil axis to wire loop 'k'*,
ρ: *Cylindrical radial distance from coil axis to field point*,
z_c: *Perpendicular distance from coil plane to field point*,
μ: *Magnetic permeability — considered uniform*,
ω: *Angular frequency.*

Assumptions leading up to Equation (15) have been discussed in earlier work. Equation (15), which we have called a mapping equation since it relates a coil loss measurement at a particular coil position and orientation to an entire electrical conductivity distribution, is actually a 3D convolution. This is made clear by recasting Equation (15) into the frame of the coil center — with the coil's Z-axis perpendicular to the coil plane, while the coil's X-axis and Y-axis lie within the coil plane. Letting the vector $\vec{c}$ locate the coil center relative to the lab frame origin, and the vector $\vec{r}_c$ locate the field point in the coordinate system (CS) of the coil, we have the revised form for Equation (15):

$$\delta Z(\vec{c}) = \int \sigma_l\left(\vec{c}+\tilde{R}\vec{r}_c\right)G\left(\vec{r}_c\right)dx_cdy_cdz_c = \int \sigma_l\left(\vec{r}\right)G\left(\tilde{R}^T\left(\vec{r}-\vec{c}\right)\right)dxdydz \tag{17}$$

A subscript, designating "lab", has been attached to the conductivity as a reminder that the argument still evaluates to a vector in the lab frame, with the orthogonal rotation matrix $\tilde{R}$ mapping the vector $\vec{r}_c$ from coil frame to lab frame. While arriving at Equation (17), we have made use of the fact that the determinant of the Jacobian transformation matrix is unity. In Equation (17), we have also made use of the definition:

$$G\left(\vec{r}_c\right) = \frac{\mu^2\omega^2}{4\rho\pi^2}\sum_{j,k}\sqrt{\rho_j\rho_k}Q_{1/2}\left(\eta_j\right)Q_{1/2}\left(\eta_k\right) \tag{18}$$

The function $G(\vec{r}_c)$ is recognized as the kernel of the convolution integral and is rapidly evaluated through use of a hypergeometric series for the toroidal functions [22] — the two forms in Equation (17) are equivalent, though the second is used here. Note that if the coil does not rotate relative to the lab frame as it is moved around, then the rotation matrix is just the identity matrix.

A variety of ways exist to discretize the convolution integral, and we continue to explore alternatives inasmuch as computation of the model matrix $\tilde{A}$ becomes time consuming for very large meshes. For our purposes here, we just use a finite element representation of the target and evaluate the convolution integral element-wise. This is accomplished by expanding the electrical conductivity into the usual superposition of shape functions. Here, we use a linear basis set on deformed prismatic elements [[23], Fig. 3.10], with 9-point integration over each element.

4. SETTING "WEIGHTS" IN REGULARIZATION MATRIX $\tilde{D}$

As mentioned, the diagonal matrix $\tilde{D}$ provides an opportunity to selectively penalize contributions from individual solution components, and single coil MIT gives us an excellent example of how this can be used to our advantage. Its use here is motivated by successful application of a position-dependent regularization in diffuse optical tomography [12]. Essentially, the kernel function in the convolution integral naturally weights regions of space according to their contribution to the measured coil loss. Thus, plots of the kernel function, either along the coil's Z-axis or radial axis are helpful in guiding how the components of $\tilde{D}$ are chosen. Here, kernel function plots (or sampling function) are given for two different coil designs to illustrate how "sampling" of space is dependent on coil geometry. The first, coil 'R', consists of five circular concentric loops in each of two layers (ten total) — spacing is 0.5 mm, so that all loops are considered to effectively lie in the same plane. The loops have radii of 4, 8, 12, 16 and 20 mm. Coil 'R' inductance was shown previously to be 2.155 μH, and is used in all calculations for this work. A second coil 'Y' consists of 4 layers of loops, all set at a radius of 20 mm, with a spacing of ~ 0.3 mm between layers — again, spacing is considered to be sufficiently small that they are considered to be in the same plane. Coil 'Y' inductance is easily computed to be ~ 1.8 μH (used here only for comparison purposes; a later publication will consider coil Y).

In a plane parallel to either coil, located at a distance of 5 mm from the coil plane, the kernel function is plotted as a function of radial distance from the coil axis in Figure 1(a). Note that coil 'R' gives a broader or smoother sampling of space while coil 'Y' tends to be more focused. However, in either case, since a coil scan primarily moves the coil in the plane of the coil and with minimal rotation, good sampling of space is achieved by collecting coil loss values on an appropriate lattice of points nominally parallel to the target (e.g., Latin Hypercube). It is likely that coil 'Y' would benefit from more data collection given its "tighter" sampling field. If in-plane sampling provides good "coverage", which is reasonable in the absence of any obstructions that would inhibit such scanning, there is likely no particular reason to base weights upon these plots.

The situation is different for out-of-plane sampling during a scan. Obviously, the coil cannot be located inside the target to collect data, so it is imperative to understand how the kernel function varies with distance, normal to the coil plane. This is shown in Figure 1(b).

Regardless of coil type, the decay with Z-distance in Figure 1(b) is nearly exponential, except in the immediate vicinity of the coil plane. Figure 1(b) indicates that greater weight or importance is given to those points nearest to the coil plane, giving them an advantage in a noisy environment, so

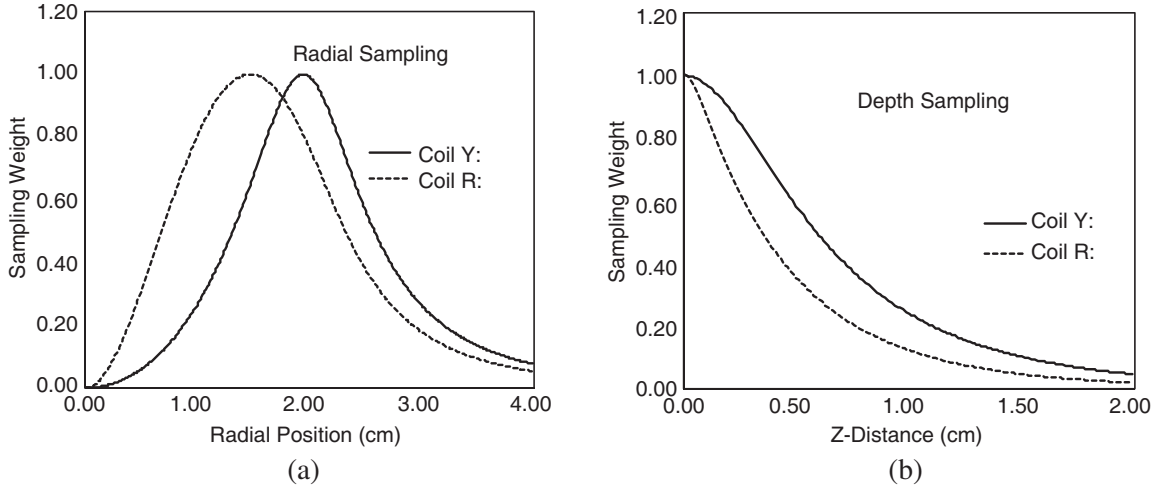

Figure 1. (a) Normalized plots of kernel for two coil types showing spatial weighting radially from the coil axis; plots are obtained in a plane parallel to the coil plane and separated by 5 mm. (b) Normalized plots of the kernel for two coil types showing how spatial weighting varies with distance from the coil plane; plots are obtained at a radial distance of 15 mm from coil axis.

that the solution from any inversion scheme will tend to favor points closest to the coil plane if nothing is done to relieve the bias. A simple procedure to help mitigate the bias applies an exponential decrease of the regularization weights in $\tilde{D}$ as solution components become associated with points farther from the coil plane, or target boundary:

$$d_j = \exp\left(-\eta z_j / z_{\max}\right) \tag{19}$$

Finite element mesh thickness is given as $z_{\max}$, while z_j is distance from target boundary to node location, and η is a decay parameter meant to control the extent of bias mitigation. The other option is to directly use the curves shown in Figure 1(b), though distance-scaled to some extent. Both are tested here.

5. TARGET MIT PHANTOMS

Two virtual phantoms are used for algorithm testing under conditions that adjust regularization and noise. Phantoms consist of buried thick strips of altered conductivity — one with low conductivity strips ($\sim 0.0\,\mathrm{S/m}$) immersed in a higher conductivity matrix ($\sim 1.0\,\mathrm{S/m}$); a second with higher conductivity strips immersed in a lower conductivity matrix. In the first, two thick strips of very low conductivity material, embedded in background material of $1.0\,\mathrm{S/m}$, are located within a $16.5\,\mathrm{cm} \times 16.5\,\mathrm{cm} \times 3.0\,\mathrm{cm}$ slab phantom as shown in Figure 2 — these low-conductivity strips are defined by regions having conductivity beneath $0.1\,\mathrm{S/m}$. The two strips each run parallel to the Y-axis, located at $X = 4.5\,\mathrm{cm}$ and $X = 12.0\,\mathrm{cm}$; and with $Z = 1.0\,\mathrm{cm}$ in either case. Using the definition that strip conductivity is $< 0.1\,\mathrm{S/m}$, then strip thickness is $\sim 1.8\,\mathrm{cm}$, while strip width for each is $\sim 3.8\,\mathrm{cm}$. Separation between strips is $\sim 4.5\,\mathrm{cm}$. Outside of a small region where conductivity is near $0.0\,\mathrm{S/m}$, a squared Gaussian function blends the low conductivity strip with the bulk region at $1.0\,\mathrm{S/m}$. The second phantom is just the inverse of the first, consisting of buried higher conductivity material ($1.0\,\mathrm{S/m}$) embedded in low conductivity matrix. Plotted images use a standard "black body radiation" color scheme and scale, spanning from black (0%) to red (40%) to yellow (80%) and finally white (100%) at full scale — unless noted otherwise.

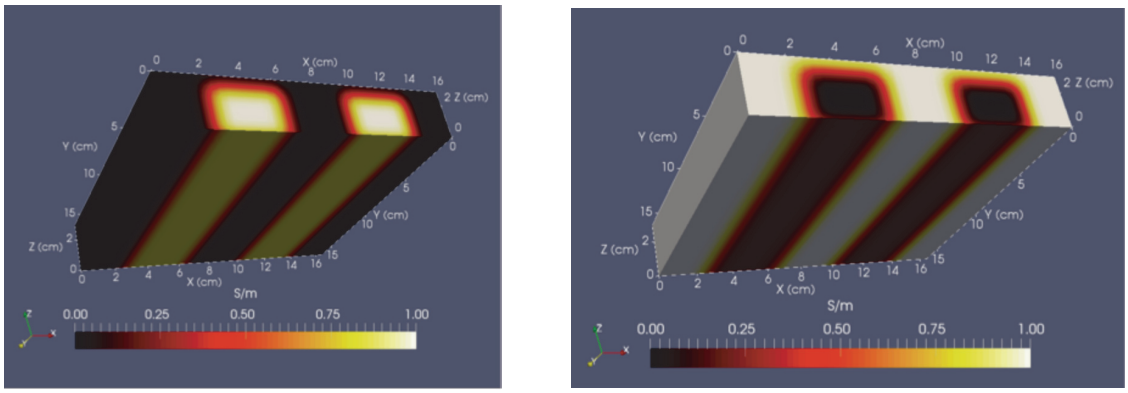

(a) (b)

Figure 2. Virtual phantoms used to test algorithm performance — left (a) phantom consists of buried, high-conductivity strips ($1.0\,\mathrm{S/m}$); right (b) phantom consists of buried, low conductivity strips.

To generate coil loss data associated with a virtual phantom, Equation (15) is used — giving loss in a coil of type 'R' (Section 4). The sample-containing space shown in Figure 2 is subdivided into 10 layers of equal thickness. For numerical integration, the total number of prismatic elements is set to 41,180 and integration on any element follows a 9 point quadrature rule. Integrations are done for each of 486 coil locations above the slab, on a $9 \times 9 \times 6$ lattice. Coil loss is computed for 81 locations on each of the six horizons on a regular grid. Odd layers, starting with the first, all share the same grid, while even numbered horizons share a common grid that is shifted by half a grid interval from the odd numbered horizons — see Figure 3. Sampling horizons are equally spaced from each other over a span of

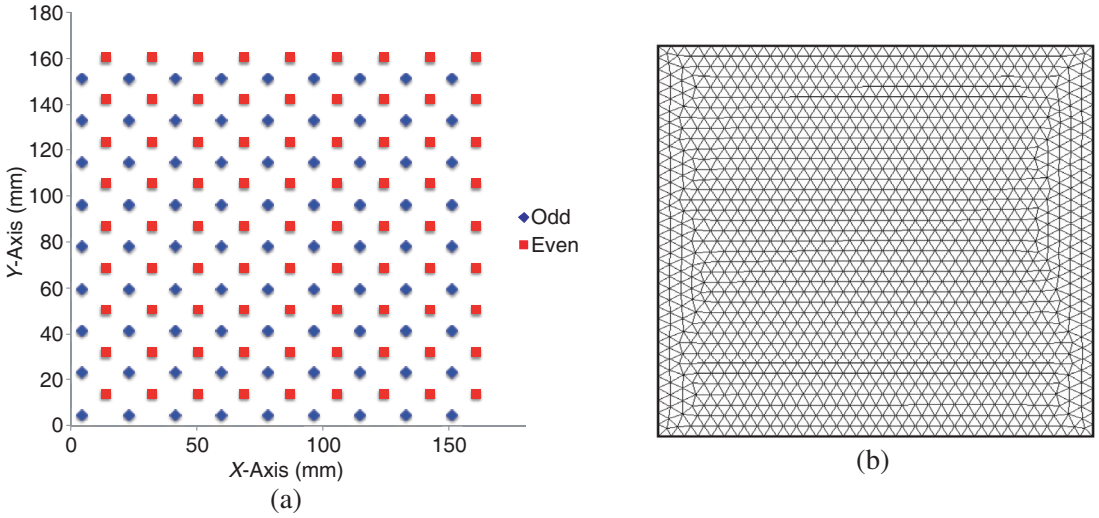

Figure 3. (a) Odd-Even sampling pattern over phantom; (b) 16.5 × 16.5 cm FE mesh — extruded to 3.0 cm thickness.

6.0 mm above the slab. The first sampling horizon keeps the coil parallel to the slab and spaced 2.0 mm above the slab. All other sampling horizons allowed for random coil tilt, with axis tilted as much as 5 degrees from vertical, which alters the height of coil center. Allowing for coil tilt when the coil is clear of the slab better mimics an actual hand scan where coil tilt is unavoidable.

The phantoms shown in Figure 2 provide an effective test of the algorithm's ability to resolve structure over its 3.0 cm depth. Resolving lateral, or XY structure should be less difficult since both strip separation and width are ~ 1 coil diameter. However, as Figure 1(b) suggests, resolving a correct conductivity profile to a depth of 3.0 cm appears more challenging since sampling level at a depth of 1.5 cm has already fallen to 5% of the maximum found at the target boundary, with coil 'R'.

Since all real measurements are contaminated by noise, Gaussian-distributed noise is added to virtual coil loss measurements over each phantom. A base level of noise is added to all "pristine" data to match the noise level of current instrumentation. However, in one additional experiment, Gaussian noise is increased 5× to correspond to older instrumentation. Experimental coil loss is obtained through measurement of the real part of coil admittance Y_{re} while in the vicinity of coil resonance. Admittance change relative to a free space admittance measurement is used to get coil loss [8]:

$$R_{loss} = \omega^2 L^2 \delta Y_{re} \tag{20}$$

Older instrumentation measured admittance with a precision of ±0.12 µS, while newer instrumentation measures admittance with a precision of ±0.025 µS. Further processing of experimental data with Savitzky-Golay denoising yields a precision level of ±0.01 µS in current instrumentation, which corresponds to a loss precision of ±0.000286 ohms — using a frequency of 12.5 MHz and inductance equal to 2.155 µH. This amount of loss uncertainty is used to compute a Gaussian noise component that is added to each of the 486 computed virtual loss values. The coil loss uncertainty can also be used to provide an error bound, and stopping condition as previously discussed, for Equation (13) or (14). Since there are 486 coil loss "measurements", the loss uncertainty for a single solution component is multiplied by $\sqrt{486}$, giving an error norm for Equation (14) of 0.0063 ohms. In one instance, noise is increased 5× to assess its effect on image reconstruction.

6. IMAGE RECONSTRUCTION RESULTS

To build the model matrix from coil position information, the second integral of Equation (17) was discretized into 20,080 elements over eight layers of equal thickness. As with the forward calculation to get virtual data, 9 point integration was used on each element. Image reconstruction requires selection of a particular value for the global regularization parameter τ, chosen here from among the set of

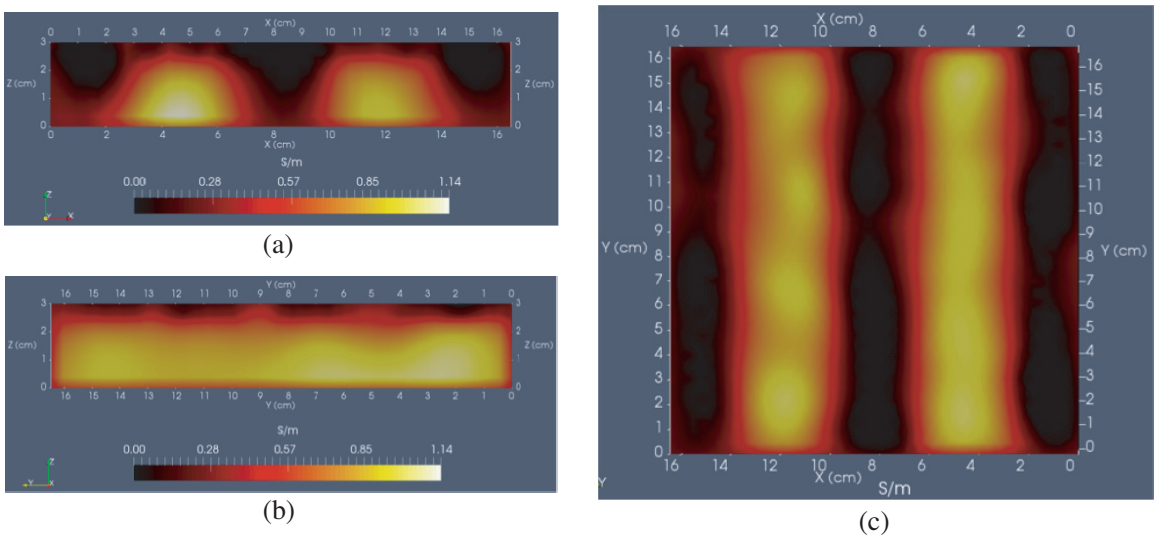

Figure 4. Image reconstructions for phantom having buried high conductivity (1 S/m) strips in a low conductivity (~ 0 S/m) background; slice (a) normal to Y-axis; slice (b) normal to X-axis, $X \sim 4.5$ cm; slice (c) normal to Z-axis; $\beta_0 \approx 0.267$ S/m from (12) to set shift.

singular values, which decrease monotonically from a maximum value at $j = 1$ to a smallest value at $j = 486$ (number of measurements, in this case). For the phantom consisting of two buried strips of 1.0 S/m material, $j = 220$ was selected, which gave an error norm of 0.0047 ohms, close to the target of 0.0063 ohms. The non-dimensional decay parameter discussed in Equation (19) was set to 1.0, as suggested by exponential fits to the 'R' curve in Figure 1(b). Convergence is obtained after 6 iterations of the active set method. Three image slices are shown in Figure 4 — a transverse slice midway along and normal to the Y-axis, a longitudinal slice normal to the X-axis at $X \sim 4.5$ cm, and finally, a Z-normal slice at $Z \sim 1.5$ cm (midway along the Z-axis). To facilitate comparison, these same slice locations are used in all images that follow.

Each slice in Figure 4 indicates that image reconstruction from 486 measurements does a reasonably good job of faithfully representing the original phantom in Figure 2. Conductivity spans from 0.0 to ~ 1.0 S/m, as it should; conductivity stays near 0.0 S/m in the space separating the two strips, from the top of the phantom all the way to the bottom; and, features are arranged in spatially correct position. Of particular importance, variation of structural features with respect to depth track rather well those of the actual phantom, suggesting that the diagonal regularization matrix is performing as intended.

As stated, setting the decay parameter in Equation (19) to 1.0 was based upon Figure 1(b) data. Of interest is the effect of setting it to zero (equal weighting), which is shown in a single slice, normal to the Y-axis (transverse to strips) in Figure 5. The image in Figure 5 is also based upon setting τ equal to singular value $j = 220$, with resulting error norm = 0.0059 ohms (near the stopping condition). Comparing the transverse slices of Figures 4 and 5, we note that when solution components are equally weighted in the penalty term ($\eta \approx 0$, Equation (19)), conductivity in between or to either side of the strips struggles more so to stay near zero over the complete depth of the field compared with the case when decay parameter is set to 1.0. The effect is perhaps more noticeable inside the strips, where in Figure 5, conductivity is considerably higher near the upper portions of the strips. If the decay parameter is increased to $\eta \approx 1.5$, conductivity inside the strips is skewed more toward the bottom, suggesting that 1.0 was near optimal.

The second phantom is perhaps more challenging for image reconstruction, since now, the algorithm has to "see through" the higher conductivity matrix just to reveal the low conductivity strips. Because there is so much matrix material at 1.0 S/m, at least compared to the strips of the first phantom, the variation in signal (coil loss) as the coil completes its scan is less, and the noise we have added would tend to have the effect of corrupting the natural variation in coil loss. For this second phantom, our stopping condition led to a τ value equal to the 210th singular value, similar to before, and a similar

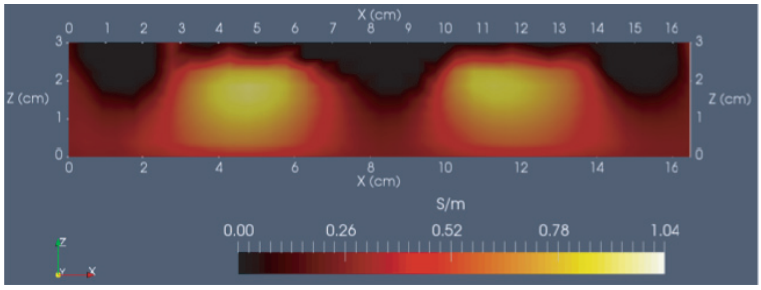

Figure 5. Slice transverse to 1.0 S/m strips, corresponding to transverse image in Figure 4 — except that now, the decay parameter = 0, impairing depth resolution; note that conductivity is skewed upward too much inside strips compared to the corresponding slice in Figure 4 that sets $\eta = 1$.

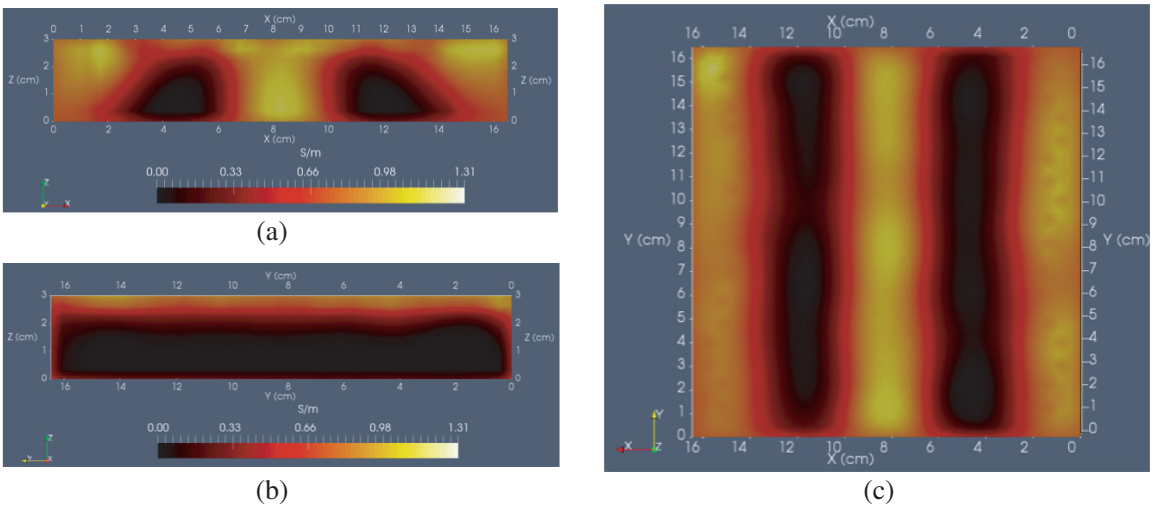

Figure 6. Image reconstructions for phantom having buried low conductivity ($\sim$ 0 S/m) strips in a high conductivity (1.0 S/m) background; slice (a) normal to Y-axis; slice (b) normal to X-axis; slice (c) normal to Z-axis; $\beta_0 \approx 0.73$ S/m from (12) to set shift.

error norm value as well. Figure 6 shows slice locations that correspond to those of Figure 4, but with an expected image inversion of the electrical conductivity distribution. Remarkably, the longitudinal slice along the length of the low-conductivity strip cleanly reveals the thin layer of high conductivity material just above the strip's entire length.

The reconstructed images shown in Figure 6 for buried low conductivity strips again do a reasonable job of faithfully matching the features of the actual phantom shown in Figure 2. Though not apparent in the slices of Figure 6, very few locations inside the full 3D image exhibit conductivity near the maximum of $\sim$ 1.31 S/m shown on the side scale — just a spot near a corner of the upper boundary where sampling is least adequate. Thus, most regions outside the strips are near the 1.0 S/m maximum conductivity, as expected. Perhaps most satisfying is the ability to depth-resolve the slab, especially the electrical conductivity distribution inside the strips and in the space between them. Depth-resolving the regions of the slab nearer the sides is more problematic since the scan includes fewer coil loss values there.

To further understand how noise impacts the image reconstruction process with this algorithm, noise was increased 5$\times$ for the phantom comprised of low conductivity strips buried in 1.0 S/m background, to an admittance precision of $\sim \pm 0.05\,\mu$S. This is 5$\times$ the noise level of newer instrumentation in current testing, and corresponds to a coil loss uncertainty of $\sim \pm 0.00143$ ohms, or a target error norm of $\sim$ 0.0315 ohms (for 486 measurements). Image slices are shown in Figure 7, and may be directly

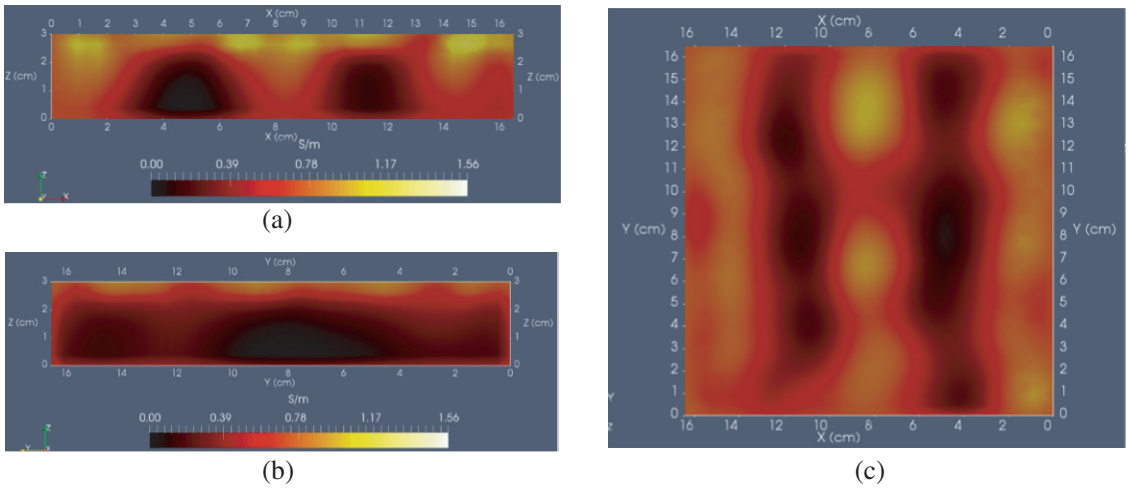

Figure 7. 5× higher noise version of Figure 6. Phantom has buried low conductivity ($\sim 0\,\mathrm{S/m}$) strips in a high conductivity ($1.0\,\mathrm{S/m}$) background; slice (a) normal to Y-axis; slice (b) normal to X-axis; slice (c) normal to Z-axis; $\beta_0 \approx 0.73\,\mathrm{S/m}$ from (12) to set shift.

compared to those in Figure 6 — color transition levels for red and yellow are set to match those in Figure 6 to aid direct comparison (though the difference is barely noticeable).

As expected, this algorithm is not immune to noise — note that geometrical distortions start to develop, which are perhaps most noticeable in the Z-normal slice (image (c)) where strips become wavy. Nevertheless, features continue to be correctly located and readily identifiable relative to the true objects shown in Figure 2. There may be a tendency to want more smoothing. However, we believe it is best to allow the effects of noise to appear in an MIT image reconstruction since even though additional smoothing will help to suppress noise, further smoothing will also remove genuine features or structure from the MIT image.

Finally, we briefly consider a different weighting scheme than the exponential weights proposed in Equation (19). As suggested in Section 4, "Setting Weights", we could directly use the sampling curves plotted in Figure 1(b) — though the Z coordinate needs to be scaled since using the plotted curves "as is" would tend to overemphasize deeper locations. Weights are based on the kernel function given in Equation (18):

$$d_j = G\,(\rho, z_{scale})\,; \quad z_{scale} = \eta z_j / z_{\max}, \tag{21}$$

The scaling parameters have the same significance as before — we just have a different function than exponential. An important difference is that we have to choose a distance from the coil axis for creating weights. Here, the chosen distance is simply the location where G would be maximum along the radial direction — see the plot in Figure 1(a), $\sim 1.5\,\mathrm{cm}$ (coil 'R'). Image reconstruction based upon Equation (21) sets the decay parameter $\eta \sim 0.3$, and with stopping conditions as previously described, singular values having index equal to 220 and 210 were used for buried $1.0\,\mathrm{S/m}$ and $0.0\,\mathrm{S/m}$ strips respectively.

Clearly, the images shown in Figures 8(a)–(b) do not appear much different than the corresponding images given in Figures 4(a) and 6(a). There is a greater tendency for the "kernel-weighted" images to better preserve the square cross sections of the strips — though the exponential weighting parameters could possibly be tweaked to achieve a cross section similar to those in Figures 8(a)–(b).

Figure 8(c) plots the same image as in Figure 8(a), but using the older algorithm described in [8]. Comparison of Figures 8(a) and 8(c) readily shows that the newer algorithm more faithfully preserves the square cross section of the conductive strips, while the older algorithm excessively "rounds" the corners. Furthermore, the image produced by the older algorithm appears distorted. This deficiency of the older algorithm has often been noted in real phantoms as well. However, the most important difference is efficiency, with the new algorithm requiring $\sim 50\times$ less CPU time — and on just a single core.

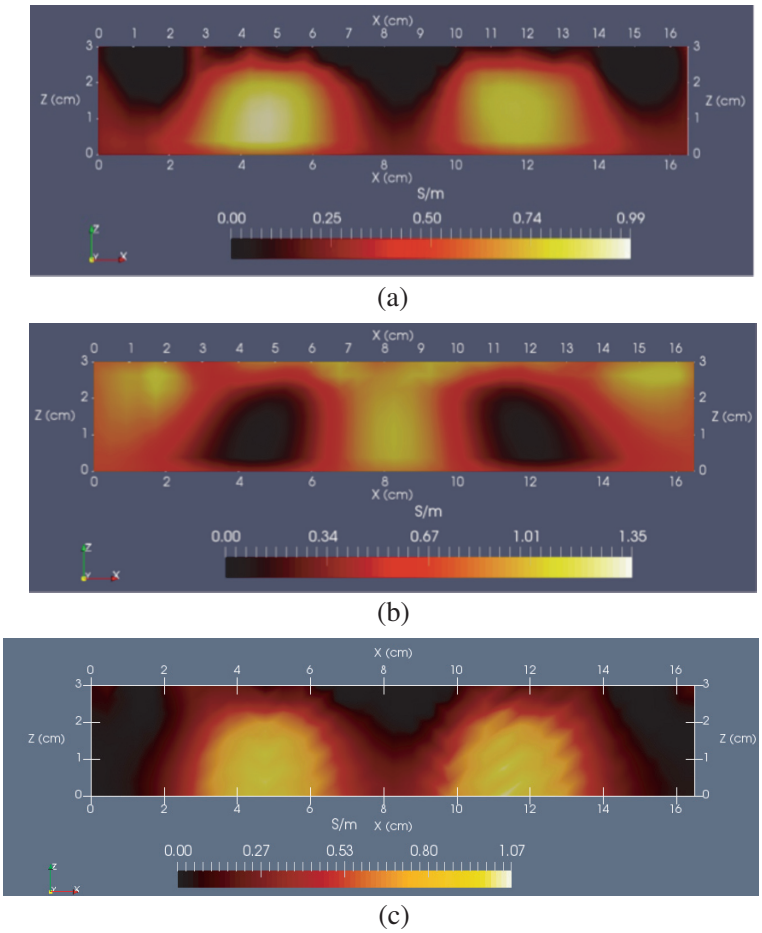

Figure 8. (a)–(b) Weights based upon the kernel function given in (18), though scaling as prescribed by (21) — these slices may be compared with the alternative exponential scaling illustrated by the transverse slices shown in Figures 4 and 6 ((a) images); (c) image reconstruction corresponding to (a), but using the method previously reported by Feldkamp [8].

7. DISCUSSION AND CONCLUSIONS

The regularization or penalty term used here departs from what was used previously in this particular application of single-coil MIT. Before, the regularization matrix consisted of the three components of the conductivity gradient, computed for each element, so that rapidly varying solutions became more heavily penalized as τ increased. Also, an L-curve approach was used to determine a suitable value for the regularization parameter, which worked well but contributed to increased computational overhead. Compared to the images produced by this newer algorithm, the prior approach caused too much smoothing, so that geometrical feature accuracy tended to be less — as noted in Figure 8(c). Furthermore, there were localization issues that could not be corrected, as they are here by appropriate use of the penalty shift and penalty weights applied individually to solution components. Also, the prior formulation was not put into "standard form", as we have done here. As a result, the newer algorithm is much more efficient, running $\sim 50\times$ faster than our older algorithm — even though the new method is run on a single core while the old method ran on 12 cores.

There is no L-curve to prepare as an aid to image selection. The global regularization parameter is set to gradually smaller singular values from the available sequence of singular values until the stopping condition is met. There is minimal overhead in doing this since the model matrix and SVD only need to be computed once. Furthermore, the sequence of solutions starts with choice of a larger singular value,

which very quickly provides a solution and opportunity to check status, and then gradually progresses to the stopping condition. As the stopping condition is approached, the time to acquire a solution becomes longer, as expected due to increasing size of the active set — which is itself a consequence of reduced filtering. The image reconstructions performed in this work took ~ 7 minutes on a single core of a Mac Book Pro (2013 build). Nearly half of that time was used to compute the model matrix $\tilde{A}$.

In our application, the diagonal regularization matrix imposed a higher penalty at those points nearest the boundary of the target, where proximity to the coil was highly likely during the course of a scan. Such an approach is necessary since otherwise, near-boundary nodes would have an outsized advantage during image reconstruction — a direct result of their greater immunity to coil loss measurement noise compared with nodes located at greater depth within a target. The situation is not unlike that encountered in CT X-ray where particular pathways through a medium lead to much greater beam attenuation than others, so that in the absence of compensation, artifacts can develop. Indeed, a diagonal "weighting" matrix has been proposed for CT image reconstruction, resembling what is used here, that applies weights to the residual error norm in accordance with the intensity of signal that exits the target [24]. Our weighting strategy works well provided that points very distant from any one of the locations where the coil visits aren't permitted to exhibit unusually high conductivity values. Though not encountered here, this potential problem could be managed by creating a space or region outside of some "preferred zone" that encompasses the space we really care about, and specifically setting a very high penalty for those outer locations. That would have the desired effect of forcing the solution outside of the "preferred zone" to remain close to some known average conductivity and therefore contribute little if anything to the predicted measurement.

This new algorithm has already been put to work on laboratory phantoms, either biological in origin or derived from agarose and fat. In a recent example reported by Feldkamp and Quirk [25], a cross-cut veal shank was scanned with modified older instrumentation and configured so that IR optical tracking could seamlessly provide coil position and orientation data [26] while acquiring inductive loss. Image reconstruction with the algorithm reported here was shown to faithfully reveal the structure of the veal shank — bone, fat seams and muscle partitions.

As already noted, efficiency is greatly improved in this algorithm compared to its predecessor, though is highly dependent upon maintaining a relatively small active set. Though not an issue as of yet, very large active sets could potentially arise in specimens comprised primarily of fat, which has a conductivity near zero. In such cases, it may be that use of an L1 penalty term could provide a superior path for inversion of the convolution integral considered here. Indeed, the work of Fang et al. [27] provides encouraging results in an MIT setting, suggesting that inversion of the convolution integral might indeed be better accomplished by $L1$ regularization.

REFERENCES

1. Gabriel, C., S. Gabriel, and E. Courthout, "The dielectric properties of biological tissues: I. Literature survey," *Phys. Med. Biol.*, Vol. 41, 2231–2249, 1996.
2. Joines, M. T., Y. Zhang, C. Li, and R. L. Jirtle, "The measured electrical properties of normal and malignant human tissues from 50 to 900 MHz," *Medical Physics*, Vol. 21, No. 4, 547–550, 1994.
3. Haemmerich, D., S. T. Staelin, J. Z. Tsai, S. Tungjitkusolmun, D. M. Mahvi, and J. G. Webster, "In vivo electrical conductivity of hepatic tumors," *Physiol. Meas.*, Vol. 24, 251–260, 2003.
4. Kim, H., A. Merrow, S. Shiraj, B. Wong, P. Horn, and T. Laor, "Analysis of fatty infiltration and inflammation of the pelvic and thigh muscles in boys with Duchenne muscular dystrophy: Grading of the disease involvement on MR imaging and correlation with clinical assessments," *Pediatric Radiology*, Vol. 43, No. 10, 1327–1335, 2013.
5. Borcea, L., "Electrical impedance tomography, *Inverse Problems*, Vol. 18, R99–R136, 2002.
6. Wei, H. Y. and M. Soleimani, "Electromagnetic tomography for medical and industrial applications: challenges and opportunities," *Proc. IEEE*, Vol. 101, 559–564, 2013.
7. Dekdouk, B., C. Ktistis, D. W. Armitage, and A. J. Peyton, "Absolute imaging of low conductivity material distributions using nonlinear reconstruction methods in MIT," *Progress In Electromagnetics Research*, Vol. 155, 1–18, 2016.

8. Feldkamp, J. R., "Single-coil magnetic induction tomographic three-dimensional imaging," *J. Medical Imaging*, Vol. 2, No. 1, 013502, 2015.
9. Feldkamp, J. R. and S. Quirk, "Effects of tissue heterogeneity on single-coil, scanning MIT imaging," *Proc. SPIE 9783, Medical Imaging 2016: Physics of Medical Imaging*, 978359, March 25, 2016.
10. Zaman, A. J. M., S. A. Long, and C. G. Gardner, "The impedance of a single-turn coil near a conducting half space," *J. Nondestructive Eval.*, Vol. 1, No. 3, 183–189, 1980.
11. Feldkamp, J. R. and S. Quirk, "Validation of a convolution integral for conductivity imaging," *Progress In Electromagnetic Research Letters*, Vol. 67, 1–6, 2017.
12. Katamreddy, S. H. and P. K. Yalavarthy, "Model-resolution based regularization improves near infrared diffuse optical tomography," *J. Opt. Soc. Am.*, Vol. 29, No. 5, 649–656, 2012.
13. Haskell, K. H. and R. J. Hanson, "An algorithm for linear least squares problems with equality and nonnegativity constraints," *Math. Program*, Vol. 21, 98–118, 1981.
14. Björck, Ä., "Numerical methods for least squares problems," *Society for Industrial and Applied Mathematics*, Philadelphia, 1996.
15. Mead, J. L. and R. A. Renaut, "Least squares problems with inequality constraints as quadratic constraints," *Linear Algebra and Its Applications*, Vol. 432, 1936–1949, 2010.
16. Hansen, P. C., "Relations between SVD and GSVD of discrete regularization problems in standard and general form," *Linear Algebra and Its Applications*, Vol. 141, 165–176, 1990.
17. Donatelli, M., A. Neuman, and L. Reichel, "Square regularization matrices for large linear discrete ill-posed problems," *Numerical Linear Algebra with Applications*, Vol. 19, 896–913, 2012.
18. Eldén, L., "Algorithms for the regularization of ill-conditioned least squares problems," *BIT*, Vol. 17, 134–145, 1977.
19. Chen, D. and R. J. Plemmons, "Nonnegativity constraints in numerical analysis," *Proceedings: The Birth of Numerical Analysis*, 109–140, Leuven Belgium, 2009.
20. Engl, H. W., "On the choice of the regularization parameter for iterated Tikhonov regularization of ill-posed problems," *J. of Approximation Theory*, Vol. 49, 55–63, 1987.
21. Hochstenbach, M. and L. Reichel, "An iterative method for Tikhonov regularization with a general linear regularization operator," *J. Integral Equ. Appl.*, Vol. 22, 463–480, 2010.
22. Gradshteyn I. S. and Ryzhik, *Table of Integrals, Series and Products*, Corrected and Enlarged Edition, A. Jeffrey, Academic Press, New York, NY, 1980.
23. Lapidus, L. and G. F. Pinder, *Numerical Solution of Partial Differential Equations in Science and Engineering*, Wiley-Interscience, J. Wiley & Sons, NY, 1982.
24. Thibault, J.-B., K. D. Sauer, C. A. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multislice helical CT," *Medical Physics*, Vol. 4, No. 11, 4526–4544, 2007.
25. Feldkamp, J. R. and S. Quirk, "Optically tracked, single-coil, scanning magnetic induction tomography," *Proc. SPIE 10132, Medical Imaging 2017: Physics of Medical Imaging*, 10132P, March 9, 2017.
26. Wiles, A. D., D. G. Thompson, and D. D. Frantz, "Accuracy assessment and interpretation for optical tracking systems," *Medical Imaging Proceedings (SPIE) 5367; Visualization, Image-Guided Procedures, and Display*, 433, 2004.
27. Fang, L., J. F. P. F. Abascal, M. Desco, and M. Soleimani, "Total variation regularization with split Bregman-based method in magnetic induction tomography using experimental data," *IEEE Sensors Journal*, Vol. 17, No. 4, 976–985, 2017.