

Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm

Irina F. Gorodnitsky^{a,*}, John S. George^b, Bhaskar D. Rao^a

^a *Electrical Engineering Department, University of California at San Diego, La Jolla, CA 92093-0407, USA*

^b *Biophysics Group, Los Alamos National Laboratory, Los Alamos, NM, USA*

Accepted for publication: 11 May 1995

Abstract

The paper describes a new algorithm for tomographic source reconstruction in neural electromagnetic inverse problems. Termed FOCUSS (FOCaL Underdetermined System Solution), this algorithm combines the desired features of the two major approaches to electromagnetic inverse procedures. Like multiple current dipole modeling methods, FOCUSS produces high resolution solutions appropriate for the highly localized sources often encountered in electromagnetic imaging. Like linear estimation methods, FOCUSS allows current sources to assume arbitrary shapes and it preserves the generality and ease of application characteristic of this group of methods. It stands apart from standard signal processing techniques because, as an initialization-dependent algorithm, it accommodates the non-unique set of feasible solutions that arise from the neuroelectric source constraints. FOCUSS is based on recursive, weighted norm minimization. The consequence of the repeated weighting procedure is, in effect, to concentrate the solution in the minimal active regions that are essential for accurately reproducing the measurements. The FOCUSS algorithm is introduced and its properties are illustrated in the context of a number of simulations, first using exact measurements in 2- and 3-D problems, and then in the presence of noise and modeling errors. The results suggest that FOCUSS is a powerful algorithm with considerable utility for tomographic current estimation.

Keywords: Recursive weighted minimum norm algorithm; Neuromagnetic source localization; EEG/MEG

1. Introduction

Electroencephalography and magnetoencephalography (EEG and MEG) are powerful techniques for the noninvasive characterization of brain function. In contrast to the hemodynamic or metabolic correlates of neural activity measured by functional MRI (fMRI) or positron emission tomography (PET), scalp electric potential distributions and magnetic fields are direct physical consequences of intracranial electric currents associated with neuronal firing. EEG and MEG provide temporal resolution on the millisecond timescale characteristic of neural population activity and can potentially provide spatial localization accuracy of a few millimeters for focal current sources, under favorable conditions. Estimation of the location and

distribution of current sources within the brain from electromagnetic recordings at or near the head surface requires the solution of an inverse problem. This problem is ill-posed by the nature of the physics (Helmholtz, 1853), meaning the solutions are not completely determined by the measured data and therefore are not unique. Field superposition permits generation of identical external electromagnetic fields for different source configurations, so that even complete knowledge of these fields does not determine internal sources uniquely. In addition, only a modest number of mathematically independent measurements are available, increasing the underdetermined nature of the problem. To solve an underdetermined problem, modeling assumptions about the solution must be made, yet accuracy of the solution is very dependent on errors and approximations in these assumptions. Models for neuromagnetic sources are obtained from physiological evidence which suggests that the underlying activity is often limited in spatial extent (Williamson, 1990), but otherwise

* Corresponding author. Tel.: (619) 534-6194; Fax: (619) 534-2486; E-mail: igorodni@ece.ucsd.edu.

is distributed over arbitrarily shaped areas. Such sources are called localized energy sources in signal processing.

Two general approaches have been applied to the electromagnetic inverse problem, the typically termed equivalent current dipole modeling methods (e.g., Geselowitz and Miller, 1973; Henderson et al., 1975; Sidman et al., 1978; Darcey, 1979; Romani and Leoni, 1984; Wood and McCarthy, 1984; Scherg and Von Cramon, 1986; Romani and Rossini, 1988; Achim et al., 1988; Scherg, 1989; Mosher et al., 1992) and linear estimation or tomographic reconstruction methods (e.g., Hämäläinen and Ilmoniemi, 1984, 1994; Singh et al., 1984; Crowley et al., 1989; Kullman et al., 1989; Roth et al., 1989; Clarke et al., 1989; Ioannides et al., 1990; Smith et al., 1990; George et al., 1991; Wang et al., 1992). The standard methods used in the two approaches have underlying limitations for MEG application. Multiple dipole modeling methods suffer from their reliance on simplified models which may not adequately describe sources of significant extent (Achim et al., 1991; Snyder, 1991). In addition, a number of method-specific problems exist (see Discussion). A major difficulty associated with the implementation of these methods is that they require explicit assumptions to be made about the model order, i.e., number of dipoles to be fit. The standard linear estimation techniques suffer from lack of resolution needed to image localized sources, including problems with source depth resolution.

In this paper we explore a novel tomographic technique for EEG/MEG source reconstruction which produces localized solutions but retains the generality and ease of application of tomographic procedures. As a tomographic reconstruction method the FOCUSS (FOCal Underdetermined System Solution) algorithm (Gorodnitsky et al., 1992) uses a forward model that assigns a current to each element within a predetermined reconstruction region. This formulation preserves the linear nature of the physical problem, by specifying the non-linear current location parameters so that the unknown current values can be related linearly to the measurements. The approach avoids the need for explicit assumptions about the model order or the shape and extent of sources. Models of arbitrary complexity, such as a multishell boundary integral or a 3-D finite element model of the head can be easily utilized in such procedures. However, the resulting large scale linear inverse problem is highly underdetermined (i.e., the number of measurements \ll the number of parameters), and the number of possible solutions for any given set of measurements is infinite. The localized energy constraint does not necessarily define a unique solution (Gorodnitsky, 1995). FOCUSS brings two features into the MEG inverse problem. As a high resolution non-parametric technique FOCUSS provides solutions that represent compact but otherwise arbitrarily shaped areas of activation in the head. By finding the true extent of localized sources, FOCUSS is able to resolve the depth of the sources correctly. As an initialization dependent algorithm, FOCUSS is also able to

deal with the non-uniquely defined set of localized energy solutions (Gorodnitsky, 1995).

The algorithm is a recursive linear estimation procedure, based on a weighted pseudo-inverse solution. The weights at each step are derived from the solution of the previous iterative step. Starting from some initial estimate, the algorithm converges to a source distribution in which the number of parameters required to describe source currents does not exceed the number of measurements (Gorodnitsky and Rao, 1993b). The initialization determines to which of these localized solutions the algorithm converges. A priori constraints can be incorporated into the algorithm through simple modifications.

Recursive constraints in MEG were first used by Ioannides et al. (1990) with a continuous source model. His algorithm was demonstrated on planar reconstructions only and appeared to have limited resolution of highly localized sources such as dipoles. The concept of recursive constraints had also been suggested, but in a limited context and without analysis, in earlier general signal processing literature (Lee et al., 1987; Cabrera and Parks, 1991). Dale and Sereno (1993) proposed a hybrid method that uses a solution from a dipole estimation procedure to determine weights for a weighted pseudo-inverse solution. This method can be seen as a first step of a recursive weighted estimation procedure initialized using a weighted MUSIC solution.

We exclusively discuss the MEG source localization problem as a vehicle for presentation of the FOCUSS algorithm in this paper. The algorithm, however, can be used equally well for EEG reconstructions and other applications that fit the FOCUSS model assumptions (Gorodnitsky and Rao, 1992, 1993a; Gorodnitsky, 1995).

2. Forward neuromagnetic model

The forward model defining the relationship between continuous head current flow and the external magnetic field is described by the Biot-Savart law (Sarvas, 1987) (Appendix A). The head in the MEG forward model is most commonly approximated by a homogeneous conducting sphere and to concentrate primarily on the inverse problem, we employ this simplified model here. In this model, the radial component of the magnetic field can be shown to reflect only the “primary” currents corresponding to intracellular ionic currents in regions of neural activation, the currents we are interested in. “Secondary” or “volume” currents, which are generated in the head by the primary current flow, make no contribution to the radial measurements. For this reason, only the radial (with respect to the head) component of the magnetic field is typically measured. The spherical model is particularly appropriate for cases where magnetic field data are measured over the regions of the head that can be approximated locally by a section of a sphere, such as the back of

the head (Cuffin, 1990). In general, the simplifying modeling assumptions are not essential and forward calculations of arbitrary complexity can be utilized as we discussed in the Introduction.

We also utilize a discrete formulation where a quasi-static approximation is used to model the continuous current. This is justified because biological currents are sufficiently slowly varying and because the medium is a good conductor (Plonsey and Heppner, 1967). An accurate discrete representation of the current is assured by sufficiently fine sampling (i.e., the discrete volumetric grid) relative to the smoothness of the continuous signal. The resolution of the data is limited, in any case, by noise and other model errors, which negate the accuracy that could be derived from exact continuous models. The discrete formulation significantly simplifies the computations and possible extensions to the algorithm. For example, incorporation of prior constraints or paring down of the reconstruction space can be easily accommodated. It also offers clarity in presenting the algorithm, making the ideas more transparent and (we hope) more accessible.

To discretize, the reconstruction volume is divided into voxels (discrete VOLUME ELEMENTS). The value of the putative continuous current is approximated in each voxel by a point dipole. Dipole strengths are unknown vector quantities that we want to estimate. They are represented in the Cartesian coordinates (x, y, z). An alternative model can also be used to represent dipoles, which employs only two orthogonal vector components tangential to the radius of the sphere and does not account for radially oriented currents in the head, which cannot be observed. Such a model eliminates from the formulation the intrinsic ill-posedness associated with the magnetic imaging problem, and in turn improves the robustness of solutions.

The measurement at each sensor is produced by a linear combination of all dipole components, some of which may be zero, indicating no current at their locations. The reconstructions are done from single time slice measurements. In MEG, these correspond to the electric activity at that instant only. Stacking up all the unknown dipole amplitude components into a vector I we can write the linear relationship between the putative dipole currents and the external magnetic field, which we express in the matrix equation

$$\begin{matrix} \begin{bmatrix} B_1 \\ \vdots \\ B_m \end{bmatrix} \\ \text{vector } B \end{matrix} = \begin{matrix} [G_{1,x}, G_{1,y}, G_{1,z}, G_{2,x}, \dots, G_{n/3,z}] \\ \text{matrix } G \end{matrix} \cdot \begin{matrix} \begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_n \end{bmatrix} \\ \text{vector } I \end{matrix}, \quad (1)$$

where B is a $m \times 1$ vector containing m radial magnetic field measurements at the m sensors. G is the $m \times n$ basis matrix whose known elements are explicitly specified in the model (the formula is given in Appendix B). The number of columns is n , the number of current compo-

nents, which (in this study) is 3 times the number of voxels in the reconstruction volume. Thus, each column $G_{i,j}$, shown above, corresponds to the j th component of a dipole located in the i th voxel and each row corresponds to one sensor. Two indices for the columns of G are used to help clarify how the three components of a single voxel are entered into the forward model. Each element of G , called the spatial weighting, specifies the contribution to a given magnetic measurement from one dipole component of unit strength. Since intracranial magnetic permeability is assumed constant here, the elements of G are pure geometric factors. Since G is constant and can be pre-calculated for a particular head and sensor geometry, it is feasible to input very complex and detailed forward models incorporating realistic geometry and major compartments within the head volume conductance, as discussed earlier. The current in each voxel is represented by groups of 3 sequential elements of vector I , which is the $n \times 1$ vector, described above. The vector I is the unknown we want to solve for. We show the elements of I numbered from 1 to n sequentially to simplify future discussions, but each entry of I corresponds to one column of G and their respective sets of 3 elements and columns correspond to one voxel.

3. Methods

3.1. The FOCUSS algorithm

We expect that most of the components in the solution for localized sources will be zero, indicating no current at the corresponding locations. Another way of looking at the solutions is through the columns of the matrix G , which are in one-to-one correspondence with the elements of vector I . For localized sources, we want to select only a small number of columns of G and, hence, the corresponding elements of vector I that can produce the magnetic measurements. This differs from the objectives of other reconstruction solutions, such as the minimum norm estimate (Appendix C) used for comparison in this paper, which produces many non-zero voxel elements representing current dispersed over a large area of the brain.

The FOCUSS algorithm finds localized solutions by starting with a distributed estimate that can be readily computed, such as the minimum norm. Then, in a series of repetitive steps, FOCUSS recursively enhances the values of some of the initial solution elements, while decreasing the rest of the elements until they become zero. In the end, only a small number of winning elements remain non-zero, yielding the desired type of localized energy solution. Computations of this type are also termed competitive networks (Grossberg, 1988).

The numerical strategy for computing recursive solutions in the FOCUSS algorithm is the weighted minimum norm. It is a pseudo-inverse based solution, like the mini-

imum norm, but it includes an additional constraint, or weight, on the solution. It is computed as

$$I = W(GW)^+ B = WW^T G^T (GWW^T G^T)^{-1} B. \quad (2)$$

W is an $n \times n$ matrix that provides a constraint on or a weighting of the solution, which in turn causes the enhancement of some of the elements in I . In the MEG FOCUSS algorithm we use only diagonal W matrices, but other extensions of the algorithm can incorporate non-diagonal W . Each step of the FOCUSS algorithm varies only in the weight matrix W . To enhance some of the already prominent elements, W is constructed directly from the amplitudes of the elements of I of the preceding steps.

Such W matrices are dimensionless scaling factors, which can be varied to produce any number of recursive schemes. Most simply, W can be constructed by taking its diagonal elements to be the previous iterative step solution, as follows:

$$W_k = \begin{bmatrix} I_{1_{k-1}} & & 0 \\ & \ddots & \\ 0 & & I_{n_{k-1}} \end{bmatrix}, \quad (3)$$

where $I_{i_{k-1}}$ represents the i th element of the vector I at the $k-1$ iteration, and k is the index of the iteration step.

Another version of the scheme can be derived by taking the weight W_k to be the compound product of all the preceding solutions. It can be computed by taking the product of W_{k-1} with the matrix of Eq. 3, that is the diagonal current matrix from the preceding step:

$$W_k = W_{k-1} \cdot \begin{bmatrix} I_{1_{k-1}} & & 0 \\ & \ddots & \\ 0 & & I_{n_{k-1}} \end{bmatrix}. \quad (4)$$

We call this scheme the compound version of the algorithm. The two versions work similarly, but from our experience, the compound version is faster and more robust, i.e., it is less likely to produce localized energy solutions very far from the initializations. We use this version almost exclusively in the presentation here.

The FOCUSS algorithm can be summarized in the following steps:

(1) Construct a W matrix from the corresponding elements of the previous step solution(s), for example Eq. 3 or Eq. 4. Other desired constraints can also be included in W .

(2) Compute the weighted minimum norm solution, Eq. 2, using W from step 1.

The two steps are repeated until the solution I no longer changes, indicating that the algorithm has converged.

Analysis and use of the algorithm

Each diagonal element of W corresponds to one element of the current. In the weighted minimum norm

computation, some large elements of W in conjunction with the data make the corresponding elements in I larger, and vice versa for small elements. By taking W to be the previous step solution, we reinforce the positive (negative) changes in the elements of I with each step. The process continues until most elements are reduced to zero. We emphasize that the algorithm does not simply magnify the few elements that are the largest at the initialization, unless these elements can form a valid solution. Often the non-zero elements in the final solution are not the ones that were the largest at initialization.

A convergence analysis of FOCUSS is presented in (Gorodnitsky and Rao, 1993b; Gorodnitsky, 1995), where the algorithm is proved to always converge to a localized solution with no more than m non-zero elements, which is the number of measurements, with at least a quadratic rate of convergence. More precisely, FOCUSS can be viewed as a novel optimization procedure part way between classical optimization routines which optimize known cost criteria and neural networks (Gorodnitsky, 1995). FOCUSS has several stable points which for a given set of measurements constitute all possible localized solutions to Eq. 1 with no more than m non-zero elements. Initialized in the basin of attraction of a particular stable point, the algorithm always leads to the localized solution represented by this stable point. Since in our model each 3 elements of I describe a single current element, the number of non-zero voxels never exceeds one-third of the number of measurements, unless some components of some active voxels are zero. We also have shown that, in general, localized solutions are non-unique (Gorodnitsky, 1995). FOCUSS is different from the standard signal processing techniques in that it can resolve this non-uniqueness because it is a local optimizer only that relies on a robust initialization to find the right solution. Thus the initialization and the shape of the basins of attraction play crucial roles in the performance of the algorithm, and they are discussed next.

It should be noted that, as is typical of non-linear dynamical processes and neural networks, the FOCUSS algorithm may get stuck in an intermediate non-localized solution, with a probability of occurrence of almost zero. If this occurs, a small perturbation in the initialization of the algorithm resolves the problem.

Initialization

From above, in the absence of a priori information about the sources, the unbiased estimate of the true solution should be used for the initialization to assure that the algorithm starts in the correct basin of attraction. Such an estimate is achieved through what we call an unbiased minimum norm estimate found as explained below and in much greater detail, including its motivation (Gorodnitsky, 1995). Initialization using the standard minimum norm solution was also evaluated and found to be adequate in the case of some simple sources, but it was not a reliable choice in general. Other potential initialization choices for

model noise present in the data. This results in unstable solutions, where even very low levels of noise and modeling errors give rise to gross deviation of the solutions from the true sources. This explains, incidentally, why exhaustive search methods which typically use exact fits to the data may fail in these kinds of problems.

The potential for instability of a solution is determined by the forward model, the matrix G . In our model, G depends only on the geometric variables: the number of sensors and their geometric configuration, the coarseness of the voxel grid, the distance from the sources to the sensors, and the geometry of the volumetric reconstruction space. A moderately well conditioned MEG model G could be achieved by choosing optimal geometric and grid spacing parameters. In the general case, however, the G matrix in MEG is ill conditioned. This, combined with high levels of error in the data and in the model G , makes regularization of MEG solutions imperative.

Stable reconstructions are achieved through the use of techniques called *regularization*. Since the measurements and the model contain noise, the true solution is no longer the minimum of the error function and cannot be determined from the measurements. Regularization techniques achieve a compromise between a close fit to the data and stability of the solution. Regularization does not intentionally filter noise from the data, but most regularization techniques remove the high frequency component from the solution, which is largely generated by noise. This effectively filters a portion of noise. A comprehensive strategy for dealing with noisy data also includes data filtering and an optimal selection of geometric parameters, such as sensor positioning. These are general topics not specific to our method and are not discussed in this paper.

Sources of error in MEG

The errors in MEG vary from simple additive noise to more complex non-linear components. The following is the list of main sources of error in MEG.

1. Noise and resolution limitations of the sensors.
2. Background noise, i.e., signals other than the ones of interest, such as unrelated body and brain activity, and environmental interference.
3. Discretization of the continuous forward model, which leads to errors in the lead field representation; this is also called error due to off-grid sources.
4. Differences between the true physics and the assumed model of head geometry and sensor physics.
5. Errors in assumed locations of sensors and the coordinate frame of the head.

The first two types of noise are simply added to the magnetic field generated by the neuronal sources of interest. The resulting errors are linearly dependent on the noise. Typically, this noise affects mostly the high frequency component of the solution and a large part of this contribution is removed from the solution in a process of regularization.

The next two sources of error cause a discrepancy between the measurements predicted by our forward model and the ones that would be produced by an actual physical system. Included in this group are “off-grid errors,” which refer to errors due to the approximation by discrete stationary charges at the grid nodes of a true continuous current that falls between the grid points. By making the discretization grid sufficiently fine, “off-grid errors” can be made small relative to errors from other sources so that the accuracy of the solution is not compromised. This strategy, however, increases the computational burden and the underdetermination of the formulation. Better discretization methods can also reduce this error. Both sources of error may generate a low frequency component in a solution, which has a greater impact on accuracy than the high frequency component from the first two sources of error. The last error listed, in the position parameters, is introduced into solutions through non-linear terms in the forward model and its effect can be significant relative to the size of other errors, as was demonstrated through Monte Carlo simulations in Medvick et al. (1989).

All of the above errors are additive, rather than multiplicative, i.e., they represent a deviation of the predictions of our forward model from the actual measurements. Thus, in the simulations, we simply model all errors as correlated or uncorrelated noise added to the measurements.

Regularization

The availability of different regularization techniques, combined with the need to repeatedly apply regularization in the multiple step algorithm, presents a multitude of options in the implementation of regularization procedures. Our goal here is to validate the FOCUSS algorithm, so we limit ourselves to basic regularization issues, demonstrating that existing tools can be applied to produce stable FOCUSS solutions. For this purpose, we use two standard regularization techniques and indicate their strength and limitations in their applications with FOCUSS.

Regularizations can be considered as the minimization of a cost function which contains a stabilization factor. Practical implementations of regularization techniques, however, typically amount to modifications of the forward model, i.e., G in our case, and we find it useful to discuss the topic from this point of view. The concept of singular value decomposition (SVD) (Golub and Reinsch, 1970) (Appendix C), which appears in many aspects of data analysis, is closely tied to regularization. In SVD a matrix, in our case G or GW , representing m sensors measuring output from n dipolar source components ($m < n$), is factored into a product of an orthogonal matrix, a diagonal matrix and another orthogonal matrix. The 3 matrices together describe m linearly independent features in G (or GW). The diagonal elements of the middle matrix are referred to as *singular values*. These values, m in number, can be thought to represent the magnitude of the contribution of each independent feature of a source to the mea-

sured data. The existence of the relatively small singular values results in unstable solutions. The utility of SVD in regularization lies in the fact that only a small number of independent components, typically much less than the number of instantaneous measurements, are needed to represent highly localized MEG sources to a high degree of accuracy (see Fig. 9c). (We assume the number of measurements used is reasonable given the complexity of underlining activity.) Usually the components associated with the small singular values are not important in describing the solution and can be altered to achieve stability.

Two basic regularization techniques exist. Tikhonov regularization (Foster, 1961; Tikhonov, 1963) is often used in medical imaging. Spectral truncation and truncated SVD (TSVD) (Lawson and Hanson, 1974) constitute the second approach. The two techniques are closely related, but TSVD is numerically more stable than spectral truncation and will be used here.

Both approaches work to modify the basis matrix to eliminate the presence of very small singular values. Their use requires that an optimal regularization parameter be identified. In Tikhonov regularization, the small regularization parameter is added to all singular values of GW , thus disproportionately increasing the small ones. In TSVD, the pseudo-inverse of GW is derived from its SVD as shown in Appendix C, with the exception that small singular values below some chosen threshold are replaced by zeros. The two methods have been shown to be comparable for the computation of the minimum norm solution, when the regularization parameters are chosen optimally (Hansen, 1990).

A popular method for choosing the regularization parameter is to search through a range of values. The stopping criteria is determined by the trade-off curve between the solution norm that we minimize and the residual between actual data and the one predicted by the solution. The curve closely resembles an L-shape. The optimal parameter is located in the “corner” of this curve, where the solution norm that we minimize levels off while the residual norm remains small (immediately before it starts to grow significantly). Parameter selection by most other methods can be shown to reduce to the same key idea of locating the corner of this or a closely related curve. Recently, Sano presented an alternative algorithm for finding optimal parameter values (Sano, 1993). More recently, we have investigated the use of alternative methods for regularization (Gorodnitsky and Rao, 1994b) and the use of additional criteria for selecting regularization parameters based on the error due purely to the process of regularization itself (Gorodnitsky and Rao, 1994a). Our results indicate that popular regularization procedures should work well in the majority of applications.

The FOCUSS algorithm operates by concentrating information in a few vectors of the GW matrix. It follows that the TSVD method, which reduces the rank of the GW matrix by making the corresponding small singular values

zero, facilitates the convergence of the FOCUSS algorithm. The elimination of the selected dimensions amounts to faster reduction of the corresponding elements in the solution. We also found the TSVD method easier to apply than the Tikhonov regularization.

In simulations, we successfully determined the regularization parameter for both methods using the trade-off curve as described above. In the case of TSVD, we also tried an approach where, for iterations subsequent to the first one, we dropped the small singular values, selected by visual inspection. In general, this strategy is not well justified. In the case of FOCUSS, however, the level below which the singular values are not significant is determined in the first iteration. After that, the algorithm successively increases a few of already dominant singular values and reduces the rest. We assumed that the singular values that became much smaller were insignificant and truncated them. The method worked consistently well in our simulations. The advantage of the approach is that it is computationally less expensive than a search through a range of parameter values. The success of this technique confirms the robustness of the TSVD method, as a range of truncation parameter choices obtained from different criteria, all led to good approximate solutions. A theoretical possibility exists for excessive truncation using the TSVD method, but this was not observed in our simulations. This can lead to an overconcentration of highly distributed sources or a loss of very weak sources.

Tikhonov regularization, in contrast to TSVD, preserves the original dimension of GW and, thus might impede convergence of FOCUSS to highly localized sources, although this also was not observed in our simulations. In any case, the convergence to highly localized solutions is not aided by this regularization method, since the small singular values that FOCUSS works to reduce may be increased through regularization. The method should be more appropriate for less localized sources. We also found the solutions produced by Tikhonov regularization to be more sensitive to the choice of the regularization parameters, than the solutions produced by the TSVD method. Given these considerations, we prefer the use of TSVD with the FOCUSS algorithm. The limitations of the two approaches, however, are somewhat different and the use of both methods with FOCUSS should be investigated further.

4. Results

4.1. Noise-free simulations

2-D reconstruction

To demonstrate the inherent properties of the algorithm with respect to MEG application, we start by examining the performance of FOCUSS in the absence of noise. Any

errors in the solutions are then directly attributable to the algorithm. For simpler illustration, the algorithm is first demonstrated on 2-D examples. Selected 3-D examples are shown at the end. The unbiased minimum norm solutions were used for initialization in all cases unless otherwise indicated.

The 2-D reconstruction grid and sensor geometry are shown in Fig. 1. An example of the simulated magnetic field (SMF) measurements used in the example of Fig. 3 is shown in Fig. 1a. The mesh plot presentation format used in Figs. 2 and 7 and 8 is shown in Fig. 1b. The reconstruction space is divided into area elements called pixels (Picture ELements). The 2-D plane can be thought of as a slice through the head and sensor array, so that the measurement plane is an extension of the reconstruction plane. In this geometry the sensors form a 1-D array. The use of 2-D sensor arrays with a 2-D reconstruction area can lead to unrealistically favorable performance by reconstruction algorithms. All currents are taken to be normal to the reconstruction plane, since any tangential components result in zero field at sensors in the same plane. The key physical properties of the 3-D problem, which are the ill-posed, underdetermined nature of the problem and the difference in distance between sources and sensors, are preserved in our 2-D simulation configuration.

Single dipole reconstruction. Fig. 2 illustrates the reconstruction of a single dipole randomly placed in the grid.

This particular example was computed using the non-compound version of the FOCUSS algorithm. We show the results of each of the iterations performed by the FOCUSS algorithm, starting with the minimum norm solution, which was used for the initialization. With this choice, the algorithm required 6 iterations to converge. Initialized with the unbiased minimum norm solution, the algorithm required only 3 iterations to converge. The compound version of the algorithm with unbiased minimum norm initialization converged in only 2 iterations in this example.

Systematic performance studies. Figs. 3–6 show the results from systematic studies of the performance of the FOCUSS algorithm. Fig. 3 displays reconstructions of a single dipole at different depth levels. The FOCUSS algorithm was found to reconstruct the sources exactly, at any depth level, with a maximum of 5 iterations. Lateral shifting of the sources did not affect the general results. The minimum norm and the unbiased minimum norm solutions are shown for comparison. The tendency of the minimum norm solution to produce diffuse and surficial source distributions is apparent in these plots. Note that the unbiased minimum norm produced source distributions that were no longer surficial, since to a large extent the procedure compensates for the distance dependence of the fields. The diffuse nature of the reconstruction, however, was characteristic of minimum norm based reconstructions, regardless of distance compensation.

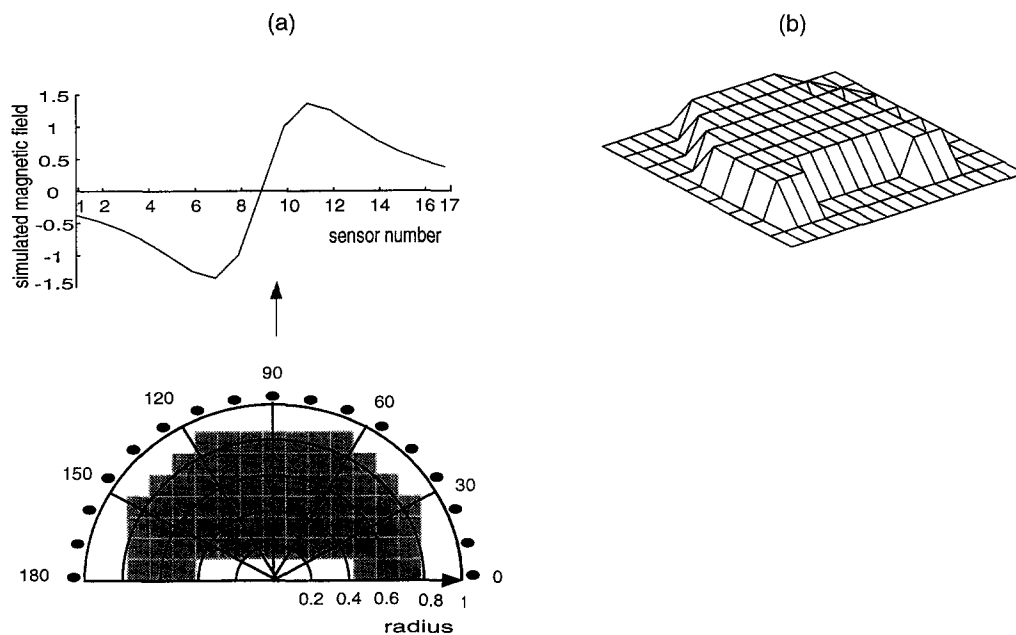


Fig. 1. 2-D geometry used in simulations. The planar geometry consisted of 17 sensors and a 72-point reconstruction grid. (a) Geometry and presentation format used in Figs. 3–6 and a plot of simulated magnetic field (SMF) distribution across the 1-D array. SMF magnitude (ordinate), for an assumed dipole source at the center of the grid, is plotted as a function of sensor numbers (abscissa). Sensors are indicated by small (filled) circles. All source currents are assumed to be normal to the plane of reconstruction. For reference in subsequent figures, sensors are numbered from 1 to 17, starting at 180° and moving clockwise. Pixels are numbered top to bottom in each column, starting in the leftmost column and moving to the right. The head radius is normalized to 1 and all radii are expressed as fractions of the radius of the head. (b) Mesh plot presentation format for the planar geometry used in Figs. 2, 7 and 8. The raised surface indicates the reconstruction region.

Fig. 4 illustrates a computational experiment similar to that in Fig. 3, except that a composite source made up of 5 equal intensity point sources was used instead of a single dipole. The source was again placed at different depths and the resulting field distribution was used to reconstruct current sources. Algorithm performance in these trials was comparable to that observed in the single dipole trials (i.e., the locations were identified exactly). Both of the minimum norm based solutions were poor and, as evident from the figure, neither one could be used reliably to determine the position or extent of the source. The minimum norm solution placed most of the current in the 3 center pixels at

the surface for all 3 depth levels of the source. The unbiased minimum norm solutions were extremely diffuse. Only the pixel with maximum power shifted from the surface to the bottom of the reconstruction grid corresponding to the placement of the source at the two boundaries.

Figs. 5 and 6 depict two systematic studies of FOCUSS reconstructions for a single dipole and for two dipoles. The same intensity scale was used for all the subimages within each figure. In the study illustrated in Fig. 5, a single dipole was placed successively at every point on the grid. The intensity of each pixel in the grayscale image repre-

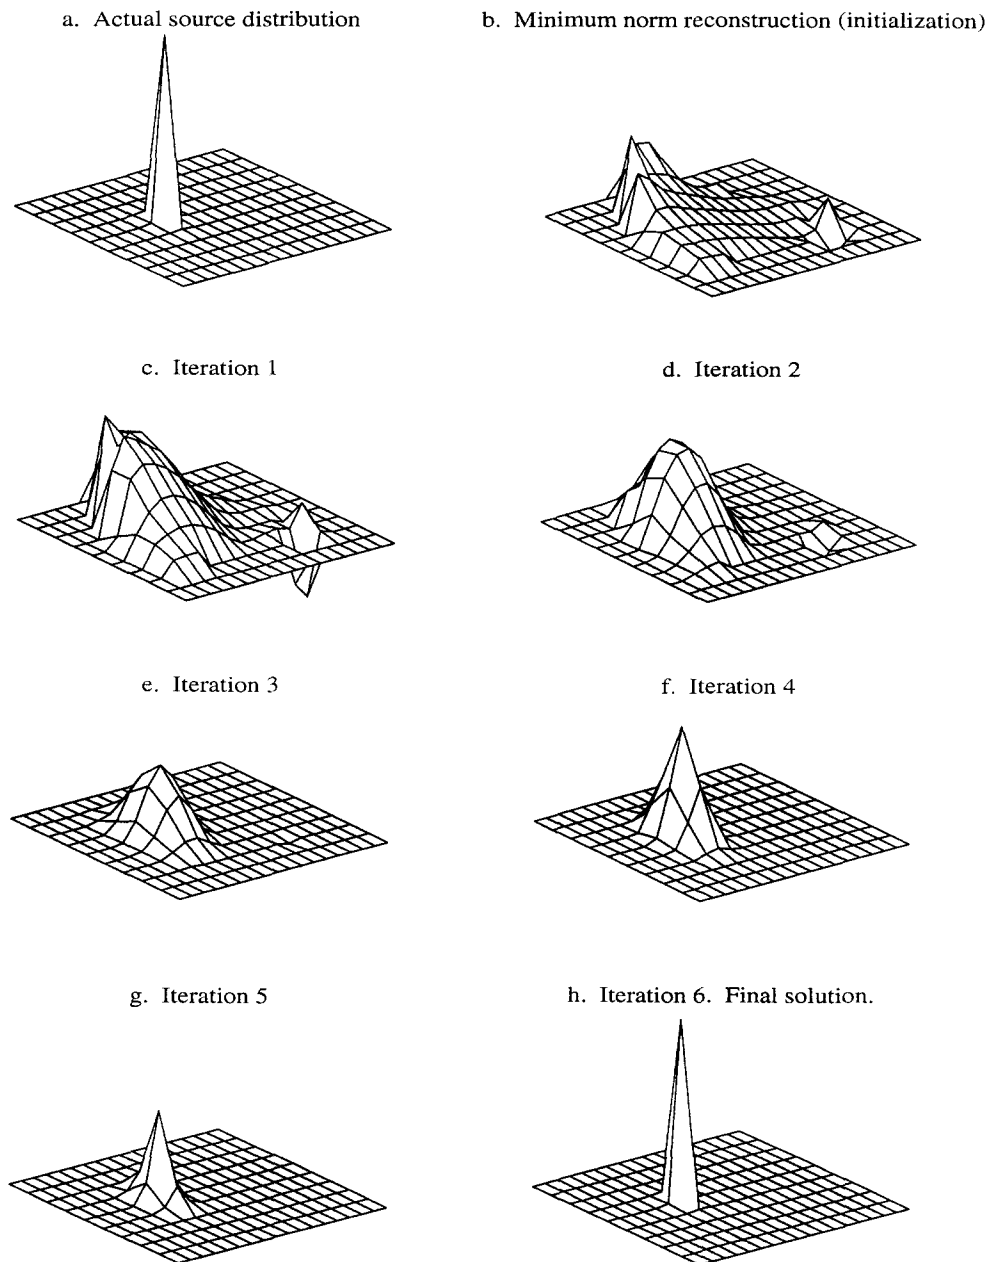


Fig. 2. Reconstruction of a single current dipole source by the FOCUSS algorithm. The simulated source distribution and the initial minimum norm estimate are illustrated in the first row. Source estimates produced by subsequent FOCUSS iterations are illustrated in subsequent panels.

sents the error in the reconstruction when the moving source was positioned in this pixel. Various error measures are possible. We chose a measure that penalizes errors in location and dispersion in the reconstructions. The error metric used for the example in Fig. 5 was computed as the sum of the distances from each element in the reconstruction to the true source, weighted by that element's intensity:

$$error = \sum_i a_i d_i \quad (6)$$

where a_i is the intensity of the i th reconstruction element and d_i is the Euclidian distance from that element to the true source.

Fig. 6 illustrates a similar study using two dipoles. The first dipole was fixed on the grid, in the same position as

the mid-depth source in Fig. 3. The second dipole was moved successively to each point of the grid. To calculate the error the true sources were first subtracted from the reconstructions. Then weighted distances from each reconstruction point to the two true sources were computed similarly to the previous example:

$$error = \sum_i a_i (s_{i1} d_{i1} + s_{i2} d_{i2}) \quad (7)$$

where a_i is the intensity as above, and d_{i1} and d_{i2} are the Euclidian distances from each element to the corresponding true source. An extra scaling s_i was used to reduce the penalty due to the distance between a source and a part of the solution near the other source. The part of the solution close to a true source should contribute only a small error in the reconstruction, but such a solution could be unduly

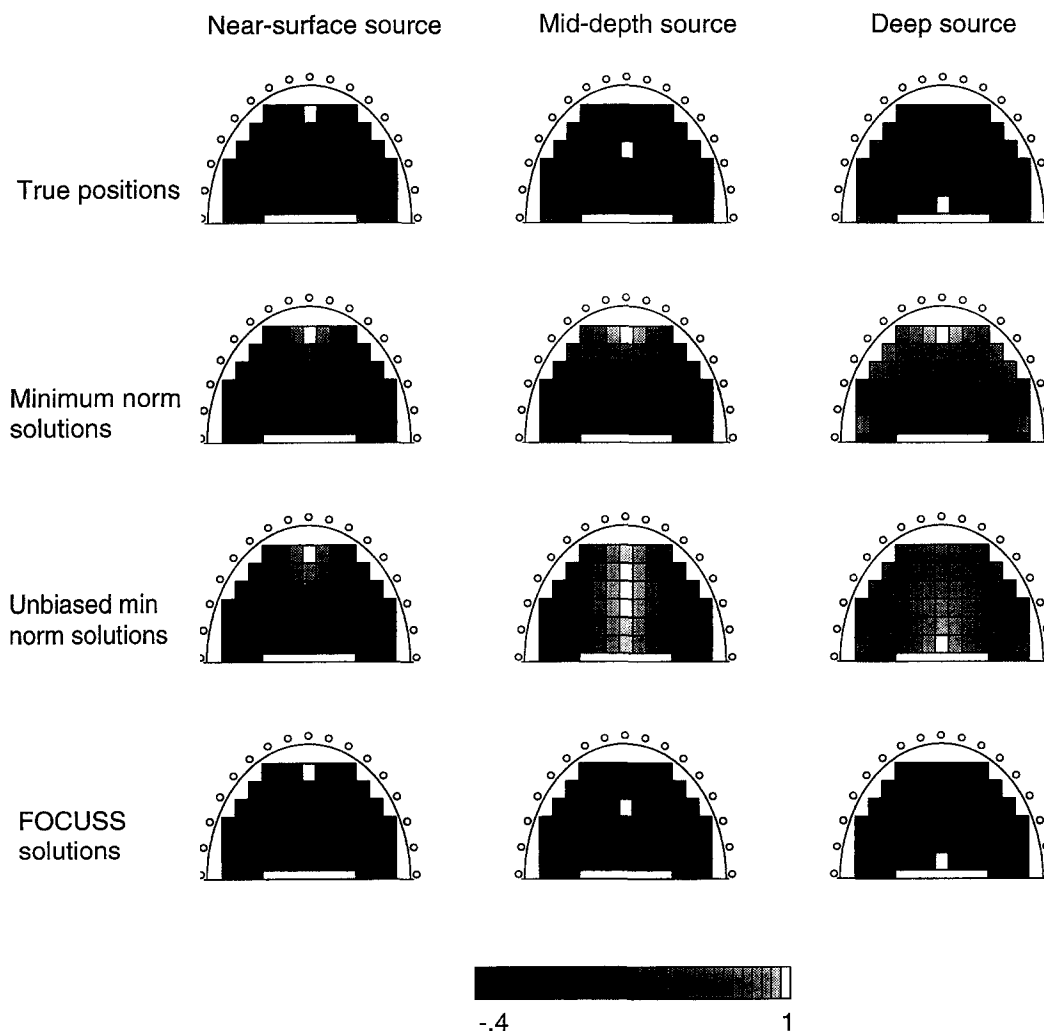


Fig. 3. Resolution achieved with different reconstruction techniques for a single dipole source varying in depth. Each column illustrates a different depth source (near-surface, mid-depth and deep) and each row illustrates a different reconstruction method (minimum norm, unbiased minimum norm, and FOCUSS). Simulated current distributions are illustrated in the top row. The compound version of FOCUSS with bias compensation was used. Each subimage was normalized by dividing each element by the maximum value in the image.

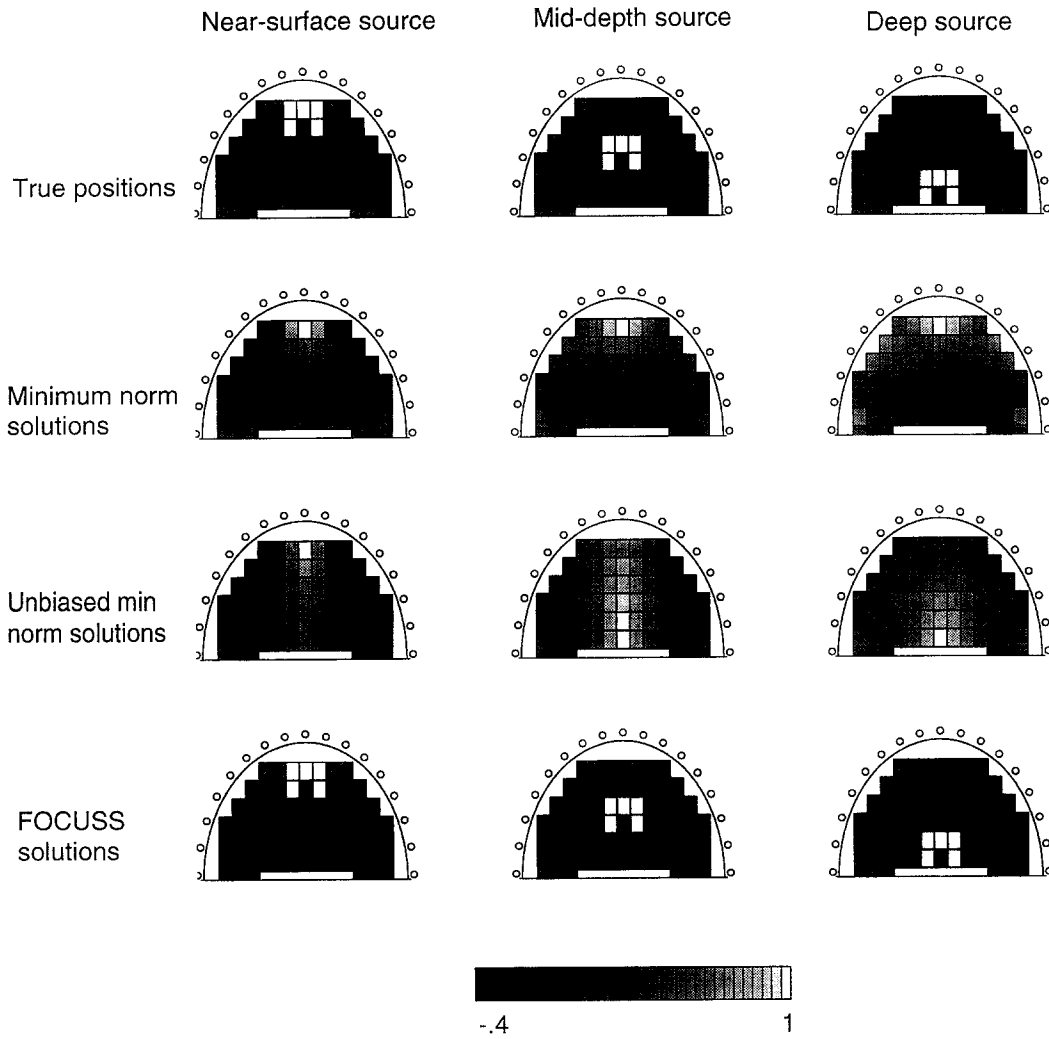


Fig. 4. Resolution achieved with different reconstruction techniques for an extended source at different depths. Each column illustrates a different depth source (near-surface, mid-depth and deep) and each row illustrates a different reconstruction method (minimum norm, unbiased minimum norm, and FOCUSS). Simulated current distributions are illustrated in the top row. The compound version of FOCUSS with bias compensation was used. Subimages are normalized as in Fig. 3.

penalized by the potentially large distance to the other source. The s_i factors in our examples were computed as:

$$s_{i1} = \begin{cases} 1.5d_{i2}, & \text{for } d_{i2} < 0.34 \text{ cm} \\ 1, & \text{otherwise} \end{cases} \quad (8)$$

and similarly for s_{i2} using d_{i1} .

The exact factors used in the error metric are not important for the type of analysis presented here. The choice does not influence the basic results, which are that FOCUSS reconstructions in these simulations were exact while the minimum norm and unbiased minimum norm solutions contained various degrees of error. We also note,




Fig. 5. Summary results of individual trials for reconstruction of a single dipole source. The intensity at each pixel represents the error in the reconstruction when a single dipole is placed at that pixel. The error metric is calculated as the sum of the Euclidean distance between each element in the reconstruction space and the true source pixel, weighted by that element intensity. White corresponds to zero error. Black is set to the maximum error among the 3 methods.




Fig. 6. Summary results of individual trials for reconstruction of 2-dipole configurations. One dipole is fixed at the center of the grid and the second dipole is placed at each pixel in the grid as in the previous example. The intensity at each pixel represents the error in reconstruction when the second dipole is placed at that pixel. The error metric (defined in the text) is a slightly modified metric described in Fig. 5. The greyscale is defined as in Fig. 5.

however, that some error metrics may not be as suitable for the type of solutions we are comparing. For example the correlations between the columns of G indicate the proximity of the sources and can be used to calculate the “distances” between the true and the estimated solutions. We found this metric to be a somewhat inconsistent measure for comparison of reconstructions that may be distributed everywhere on the grid. Alternatively, the summed squared error between simulated and reconstructed source distributions might serve as a reasonable error metric if the simulated sources and solutions were rather distributed. However such a metric would unduly penalize minor location errors in point sources.

Reconstruction of complex source distributions. Fig. 7 shows the reconstruction of a dipole and two extended sources. The selected problem was difficult, because the

two weaker sources on the right side are separated by only a single pixel. Fig. 8 shows another reconstruction of multiple sources which are close to each other and have varying current intensities within each source. We display only the final results from the FOCUSS algorithm along with the minimum norm and the unbiased minimum norm solutions for comparison. In both cases the FOCUSS algorithm reconstructed the true sources.

Image of the basis matrices. To understand how the different minimum norm based solutions are obtained it is useful to examine the basis matrices used by the inverse procedures. Fig. 9 shows the visualization of the G and GW matrices used in the minimum norm, unbiased minimum norm, and the FOCUSS 2-D reconstructions. The elements of the matrices are encoded by intensity values, from black for large negative field amplitudes to white for

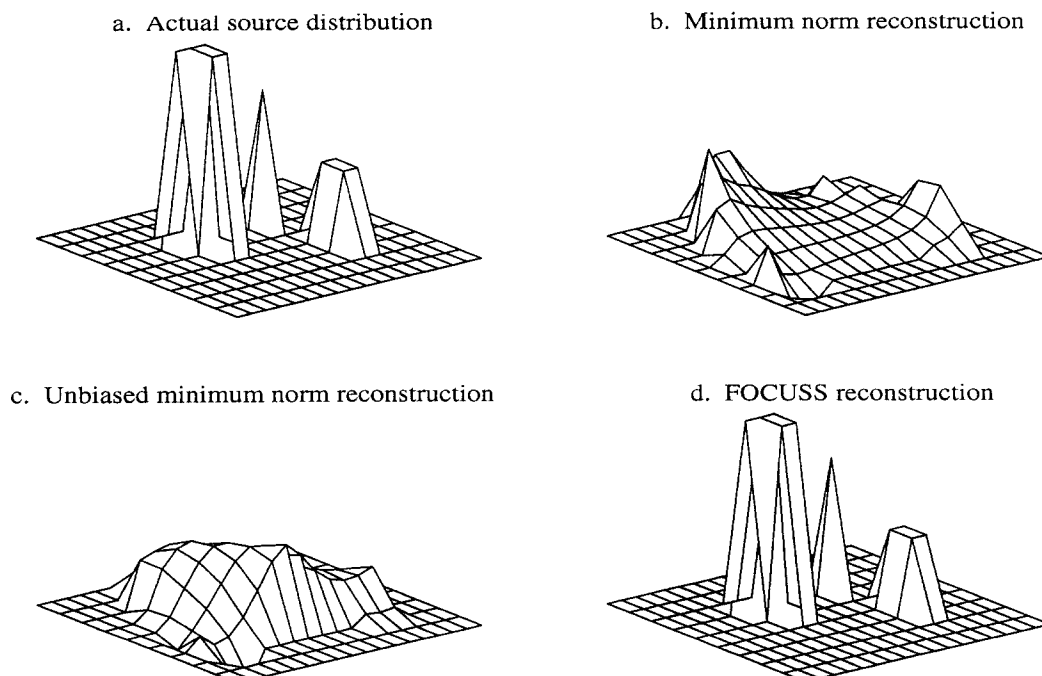


Fig. 7. Reconstruction of a complex source configuration including a dipole and two extended sources. Geometry is as illustrated in Fig. 1b. The compound version of FOCUSS with bias compensation was used.

large positive values. Each element of a matrix represents the value of the magnetic field at a sensor due to a unit strength dipole at a grid point. Each column matrix corresponds to one grid point of the reconstruction space, and each row represents one sensor. Thus, for example, the (3,5) element of the G matrix is the field at the third sensor due to a unit strength dipole at the fifth grid point. The higher the intensity of an element of G , (either positive or negative) the greater the burden placed on that grid point to explain the measurements of the corresponding sensor.

The intensity coded G matrix in Fig. 9a demonstrates that the grid points nearest to a sensor have significantly higher element values and thus higher contrast within a column. This simply reflects the effect of the magnetic field fall off with distance. In minimum norm solutions these elements are employed to account for most of the power in measurements at that sensor, thus producing solutions with most power near the surface of the head.

Fig. 9b illustrates the weighted base matrix used in the unbiased minimum norm solutions. In the unbiased estimates we used the bias compensation procedures described above, to equalize the burden placed on each element of the basis matrix to explain the measurements. Note the increased intensity and contrast associated with deeper voxels.

Fig. 9c illustrates the GW matrix from the last iteration of the FOCUSS algorithm for the multiple source reconstruction example in Fig. 7. When all but a few elements

of W were diminished, the corresponding columns of matrix GW were also diminished, i.e., their elements became zero. A regular minimum norm reconstruction with GW as the basis matrix would produce a localized solution, whose only non-zero elements are ones that retained representation in the weighted basis matrix. In effect, the FOCUSS algorithm concentrated all the energy of the basis matrix in the few appropriate columns and in this way produced a localized solution.

3-D reconstruction

The two examples shown are reconstructions of a single dipole and a configuration of 3 composite sources. The sensor geometries, system axes, and source position coordinate conventions used in the 3-D reconstruction examples are illustrated in Fig. 10. The reconstruction volume consists of the top half of a sphere with the restriction that sources do not lie within a small radius of the center. A $16 \times 16 \times 8$ reconstruction grid was used. Measurements were simulated using the 37-sensor ‘‘Magnes’’ neuromagnetometer geometry (Biomagnetic Technologies, Inc., San Diego, CA), centered on the vertical axis directly above the sphere. Thirty-seven measurements were used to reconstruct a single dipole. For the 3-source configuration, the number of measurements was increased to 73 by first computing 37 measurements for the original magnetometer position and then rotating it around the vertical axis and computing an additional 36 measurements at the new locations. Both examples use sources with large tangential

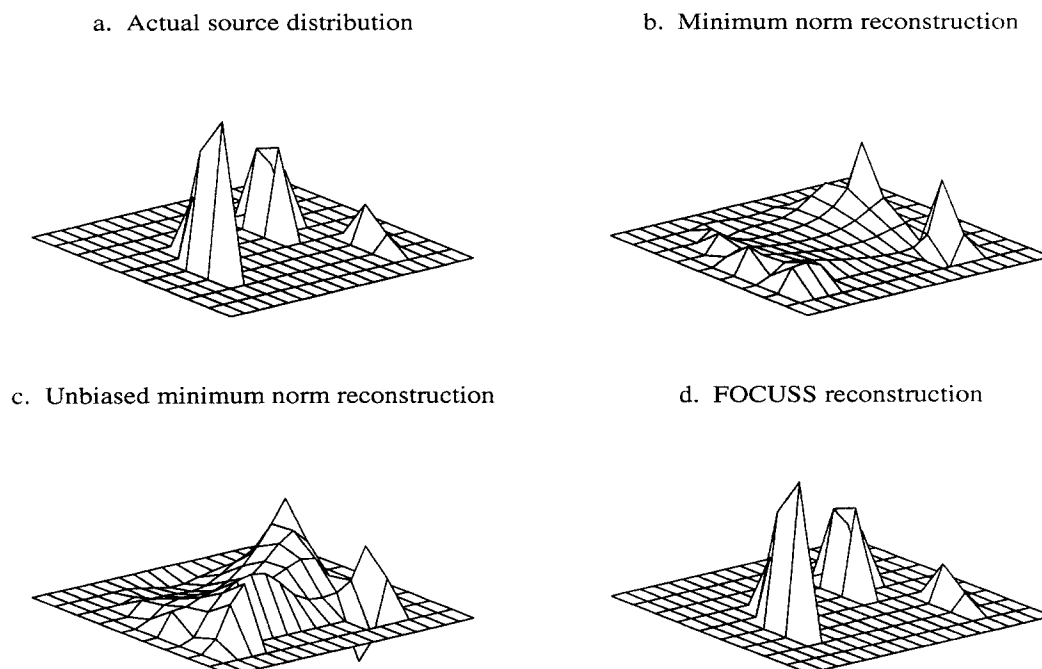


Fig. 8. Reconstruction of a complex source configuration with varying current intensities within each source. The compound version of FOCUSS with bias compensation was used.

components in each of the 3 axes. The polar coordinates of the source positions in each example are listed in Table 1. The coordinates for the center of each source are given in the case of the extended sources. To visualize the 3-D results, source currents were orthographically projected onto a plane. For clarity, a maximum intensity projection was used, that is, only the highest of the values projecting onto the same point on the plane was plotted.

3-D examples. The first example is the reconstruction of a single dipole that was placed in an arbitrary location inside the half sphere shown in Fig. 11. The algorithm resolved the arbitrarily placed dipole exactly, after only 5 iterations. The minimum norm solution was used for initialization and is shown for comparison.

Fig. 12 shows the reconstruction of a combination of a single point dipole and two extended sources, simulated with 4 and 5 static dipoles respectively. The orientations of the dipoles within each extended source were nearly parallel, but the intensities of the dipoles within each source varied from 0.2 to 0.4. The total intensity of both the 4 and 5 point extended sources was 1.5, compared to an intensity of 1.0 for the single dipole source. The FOCUSS algorithm reconstructed the exact shapes of the extended as well as the single dipole sources, converging after 5 iterations. When the 3 axial components per dipole model was used with significantly extended sources, as was done here, we found that the algorithm sometimes assigned small radially oriented currents in voxels adjacent to the true sources.

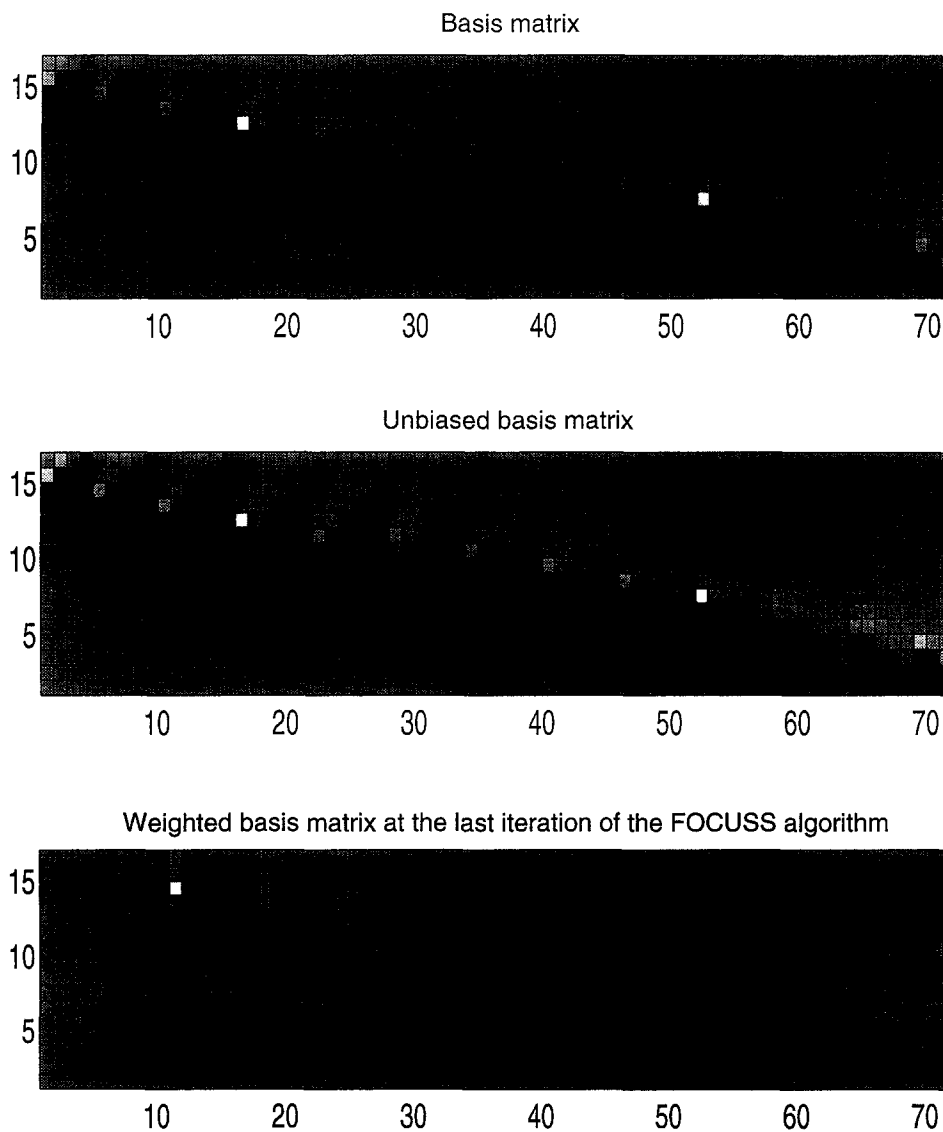


Fig. 9. Intensity coded gain matrices. Each row corresponds to a sensor. Each column corresponds to a pixel in the reconstruction space and represents the field intensities at the sensor locations that are producible by a dipole of unit strength in the given pixel location. The column values are analogous to the ones shown in Fig. 1c. Numbering conventions are as described in Fig. 1a. (a) Basis matrix G for the 2-D slice geometry illustrated in Fig. 1a. (b) Weighted basis matrix computed as matrix G with bias compensation weighting applied. (c) Weighted basis matrix GW from the last iteration of the FOCUSS algorithm for the example illustrated in Fig. 7.




Fig. 10. Reconstruction geometry, system axes, and source position coordinate conventions utilized for 3-D simulations and source reconstructions. The assumed sensor array consisted of 37 first order gradiometers distributed in a geometry approximating that of a commercially available instrument (BTi, San Diego). The reconstruction volume consisted of a $16 \times 16 \times 8$ grid with voxels adjacent to the center of the sphere being omitted.

This sort of phenomenon reflects the inherent ambiguity of the neural electromagnetic inverse problem. Since the extraneous currents are radial, they can be easily detected. This should not be a significant problem in any case, since solutions are expected to be contaminated by even larger errors from noise when empirical data are used. This problem can also be avoided by employing the two axial components per dipole model.

The speed at which the algorithm converges depends on its initialization, but is generally high. From our experience, the compound version of the algorithm converged in

Table 1
Polar coordinates of the true sources shown in Figs. 11 and 12

	Sources	r^a	θ ($^\circ$)	ϕ ($^\circ$)
Fig. 11	Source 1	0.76	33	128
Fig. 12	Source 1	0.78	64.7	86.7
	Source 2	0.73	27	185
	Source 3	0.84	41.5	243.3

^a Radial distance is expressed as a fraction of the radius of the head.

no more than 10 iterations, and often half as many. The algorithm is numerically robust, assuming that the computational steps are implemented optimally. We have not experienced any numerical instability problems with the algorithm when used in standard double precision on SUN workstations. The regularization necessary for processing empirical data provides further computational stability for the algorithm.

4.2. Reconstructions with noisy data

In the next series of studies we illustrate the performance of the FOCUSS algorithm in the presence of noise. Two stabilization methods, Tikhonov regularization and TSVD, discussed above, were used with the FOCUSS algorithm, applying the chosen method at each iteration. The optimal regularization parameters were selected as described above. Additionally we used parts of generated noise sequences to simulate the noise that can be recorded prior to administration of stimuli. This was then used to “prewhiten” the data before the algorithms were applied.

Reconstructions from noisy data will always contain some error. To avoid having errors in source amplitudes




Fig. 11. A 2-D projection of the 3-D reconstruction of a single dipole. The compound version of FOCUSS with bias compensation was used. Figures illustrated are maximum intensity projections (analogous to X-ray projections), where only the maximum value projected into the same point is shown. (a) The simulated source dipole location. (b) The minimum norm estimate. (c) The FOCUSS algorithm estimate after 5 iterations.




Fig. 12. A 2-D projection of the 3-D reconstruction of a single dipole and two extended sources. Computation and visualization conventions are as in Fig. 11. (a) The simulated source distribution. (b) The unbiased minimum norm estimate. (c) The FOCUSS algorithm estimate after 5 iterations.

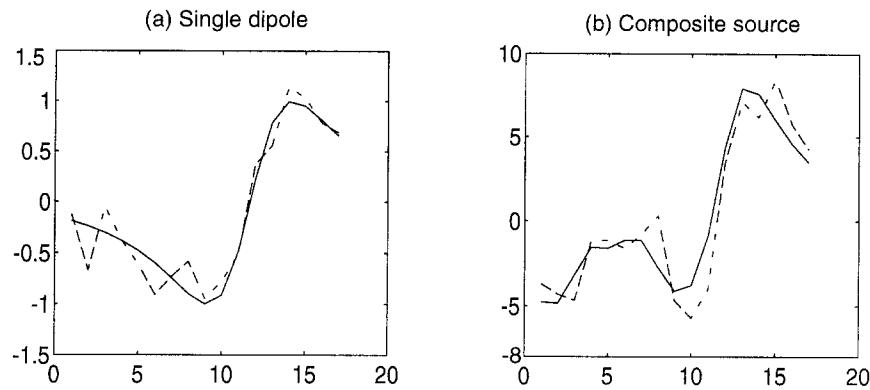


Fig. 13. Noise-free and noise corrupted simulated magnetic field measurements (15% noise level) as a function of position along the sensors for the 2-D geometry. (a) Single dipole example from Fig. 2. (b) Composite source example from Fig. 7.

obscure the localizations, we use contour plots of the absolute values to display source positions only. Planar reconstructions with the same 2-D geometry as above are used for clearer presentation, but the techniques work similarly well in 3-D.

For a given source configuration, the measurements were computed by calculating the exact magnetic field at the sensors, using Eq. 1, and adding random noise to it. Noise levels with standard deviation values ranging from 3% to 15% of peak signal intensity were used to test the algorithm. Both, correlated and uncorrelated noise were tested, but correlated noise did not appear to have a significant effect on the regularized solutions. The FOCUSS performance was also consistent for all noise levels,

and we show only the results of simulations using uncorrelated noise with a standard deviation of 15% of the peak signal value or equivalently, a noise variance of 22.5%. This level of noise is significant from the computational point of view. The examples here are intended only to demonstrate that stable solutions are achievable with FOCUSS and do not explore the ultimate resolution capabilities of the algorithm. Such a study would require the development of an optimal regularization procedure for FOCUSS as well as the use of more realistic separation between sources.

The single dipole (Fig. 2) and the example with a single dipole in combination with two distributed sources (Fig. 7), from above, are used for source configurations. Fig. 13

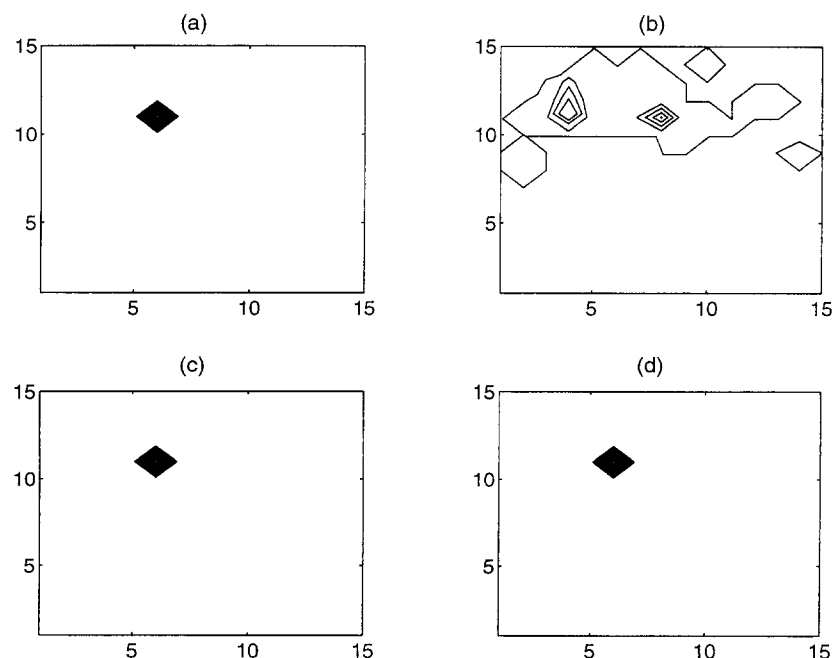


Fig. 14. Solutions with and without regularization for the single dipole example of Fig. 1 with 15% noise level. (a) Simulated dipole source. (b) Reconstruction without regularization. (c) Solution using Tikhonov regularization. (d) Truncated singular value decomposition (TSVD) solution using 1 singular value.

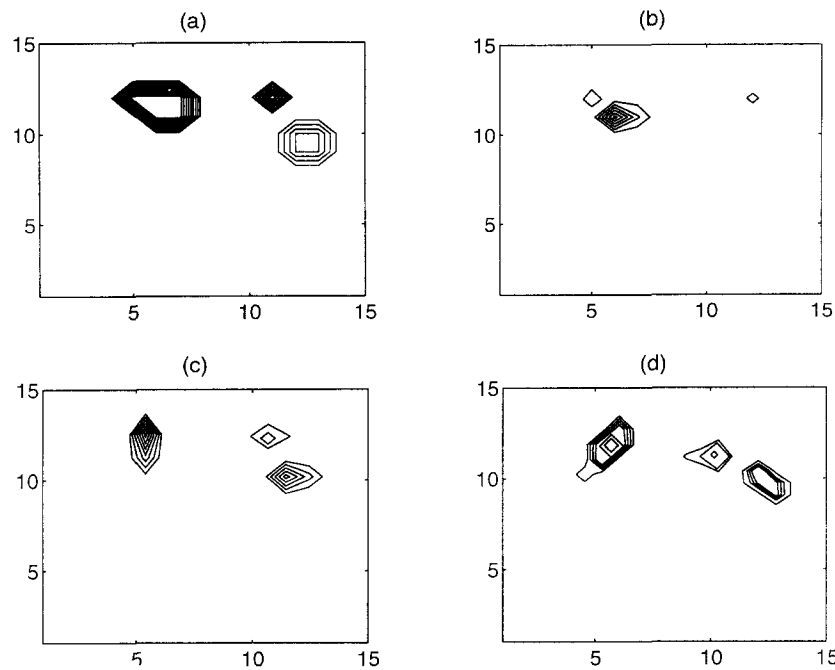


Fig. 15. Solutions with and without regularization for the composite source example of Fig. 7 with 15% noise level. (a) True sources. (b) Reconstruction without regularization. (c) Solution using Tikhonov regularization. (d) TSVD solution.

shows the noise-free and the noisy magnetic field measurements for both examples. The second example is quite challenging, as noted above, because the two weaker sources are separated by only one pixel. This separation is further obscured with the addition of noise. We chose a more challenging case from the randomly generated noise samples. The noise in the selected case amplifies the field due to the strongest source, and significantly distorts the rest of the field, greatly obscuring the waveform generated by the two close sources on the right.

Single dipole reconstruction. Fig. 14 shows the reconstruction of a single dipole. The true source configuration (a), the solution with no regularization (b), and the regularized solutions using Tikhonov regularization (c) and TSVD methods (d) are shown. The non-compound version of the algorithm was used for this and the next example. The unregularized solution is different from the true source, as predicted. The TSVD solution shown was computed using only one singular value of the GW matrix in the final iteration. The number of the singular values kept in the final iterations was determined by visually inspecting the plots of these values. The plots clearly show the growth of one dominant singular value, which we used, truncating the rest. If a solution with 3 or 4 singular values was used instead, it yielded the same correct source reconstruction, but the resulting source was slightly spread out, i.e., a fraction of the source was placed over a second voxel. This demonstrates the relative insensitivity of the TSVD method to the exact choice of the truncation threshold.

A complex source reconstruction. The second example is shown in Fig. 15. The parts a, b, c and d of the figure

correspond to those of Fig. 14. In both regularized solutions the original single dipole gets spread over two neighboring voxels, so it looks a little extended. The error, however, is small for both methods, since the current assigned to the extra voxel is minimal compared to the size of the dipole source. Some error in the reconstruction of the sources is unavoidable, and given the level of noise and the difficulty of this case, we consider the regularized reconstructions successful.

5. Discussion

Investigation of the functional organization of the brain using neural electromagnetic techniques has been severely limited by the deficiencies of inverse methods used. The FOCUSS algorithm offers a powerful alternative to the standard methods used to solve the EEG/MEG inverse problem. The algorithm finds solutions that represent spatially localized generators of the measured fields, which are thought to be typical of certain types of neural activity. FOCUSS is also able to handle non-uniquely defined localized energy sources. The algorithm does not require the use of prior constraints and instead can constrain the solution based solely on the measurements. If desired, however, the algorithm easily accommodates incorporation of additional constraints, such as anatomical constraints. The algorithm is capable of very fine spatial resolution, limited only by the uncertainties in the data and model.

The standard strategies applied in EEG/MEG inverse modeling have exhibited a number of problems, which are addressed by the FOCUSS algorithm. The equivalent cur-

rent dipole modeling methods assume that surface electrical or magnetic measurements can be adequately modeled by a small number of equivalent current dipoles (ECDs). Parameters describing the location, orientation and strength of each component vector current are found using non-linear minimization techniques. The model becomes well determined (i.e., the number of parameters < number of measurements), when the number of sources is assumed to be small. A major difficulty associated with the implementation of multiple dipole modeling methods is that they require explicit assumptions to be made about the model order, i.e., the number of dipoles to be fit, which strongly influences solutions (Supek and Aine, 1993). Although the simplified ECD model is valid for highly localized sources, in the case of sources of significant extent, such model assumptions can introduce sizable localization errors (e.g., Achim et al., 1991; Snyder, 1991). More specific limitations associated with individual methods in this category are as follows. High resolution estimation methods, such as the MUSIC algorithm attempt to fit dipoles at fixed locations to some underlying activity over some time span. Thus, any changes in the location of activity, such as the appearance of an additional source or the extinction of an existing source during this time span can result in significant errors. Likewise, sources that are highly correlated temporally present a significant problem. MUSIC also assumes that both signal and noise can be modeled as stochastic processes, requirements that are not necessarily satisfied by MEG data. Another problem observed with MUSIC is that, at least in some implementations (using real data), it can produce completely different MEG solutions using model vectors with different scale factors. In non-linear least-squares based search methods, the error surface, though low dimensional, may have many local extrema, so that the solution obtained can be extremely sensitive to the choice of initial parameter estimates. Dependence on the initialization is necessary when localized energy solutions are non-unique, as discussed earlier. However, in order to use initialization to resolve non-uniqueness, the minima of the error surface must correspond to the valid localized energy solutions only. This, in general, is not the case for non-linear least-squares methods. These implementation complexities often make the non-linear least-squares solutions critically dependent on the diligence and insight of the analyst.

Alternatively, the tomographic reconstruction techniques do not require specification of the model order, and can conceptually accommodate sources of arbitrary shape or extent, but they suffer from an underdetermined formulation. Standard tomographic reconstructions typically have poor resolution, and often display problems with source depth estimation.

The most common tomographic reconstruction used in neural electromagnetic imaging is the minimum norm solution (Appendix C) (e.g., Hämäläinen and Ilmoniemi, 1984, 1994; Crowley et al., 1989; Kullman et al., 1989; George

et al., 1991; Wang et al., 1992). Unconstrained minimum norm reconstructions produce solutions that are typically diffuse and superficial (i.e., currents are maximal at sites closest to the sensors), stemming from the strong dependence of field strength on the distance between the source and sensor, and the underdetermination of the linear problem. The uniqueness of the minimum norm solution has been considered to be a desirable property that justifies the use of this solution in MEG. The minimum norm solution, however, is unique only in a mathematical sense; i.e., there is only one solution that both minimizes the norm of the current and provides the exact fit to the data, but it is not the only possible exact fit to the data. A unique solution for a given set of assumptions does not guarantee an accurate solution, if those assumptions do not fit the characteristics of the true sources. The dispersed and superficial solutions that are characteristic of the minimum norm are counter to the expectations of localized sources for certain classes of neural activity, such as those produced by focal sensory stimulation.

Recursively weighted norm minimization methods, such as FOCUSS, provide alternatives to the standard techniques. The performance of these novel techniques depends on a number of issues and implementation options, many of which we have investigated and discuss here.

The FOCUSS algorithm finds solutions in which the number of non-zero elements does not exceed the number of measurements. A number of such solutions may exist for a given set of measurements (Gorodnitsky, 1995). An important feature of FOCUSS is that the particular solution the algorithm finds is initialization dependent. As with all initialization dependent localization methods, the initialization has to be within some proximity to the true sources, in order to find the correct solution. In the case of FOCUSS, this means that at least a small number of non-zero elements in the true solution must be among the larger elements in the initialization. The precise relationship, however, depends on the intrinsic bias (or the portion thereof remaining after bias compensation) and on the spatial extent of the underlying sources. The greater the extent of the true sources, the more accurate the initialization has to be. For example, in our experience, a single dipole or a small extended source is always found exactly with FOCUSS, whether initialized with the regular or the unbiased minimum norm solution. As the number of non-zero elements in the true sources, as described by the forward model, approaches the number of measurements, the available initialization may not lead exactly to the true sources. However, the errors that may result in the FOCUSS-generated distributed source reconstructions are smaller than the errors that can occur in reconstructing more compact sources. Decent initializations will typically lead to reconstructions that are acceptably close to the actual sources.

Robust implementation of the recursive weighting strategy along with bias compensation and optimal initializa-

tion, are the factors that are vital for successful utilization of the algorithm. This paper gives two examples of the implementation of recursive constraints, of which the compound version is considerably more reliable. We also describe some successful methods for bias compensation and for robust initialization.

The numerical stability of the solutions is affected by inaccuracies in the models and data, requiring some form of regularization. Many variations in regularization procedures are possible and optimal approaches need to be established. The two standard regularization methods, Tikhonov regularization and TSVD method may favor slightly different source distributions and their use in combination could be advantageous.

The discretization of the reconstruction volume affects the model accuracy and numerical stability of the solutions. The quasi-static discrete model simply samples continuous functions at the rate determined by the discretization grid. The grid has to be sufficiently fine compared to the smoothness of the current to provide good approximation. However, too fine a grid exacerbates the problem of underdetermination, and increases the instability of reconstructions in the presence of noise. More accurate discretization methods such as quadrature rule approximations of volume integrals may offer greater accuracy with a coarser reconstruction grid, without compromising numerical stability.

Reconstructions using a hierarchical grid may hold practical utility for the FOCUSS algorithm as well. In such procedures, the grid is progressively refined in regions of activity and pared down in the regions with no activity. Since the number of non-zero current elements in the FOCUSS solutions is limited by the number of measurements, such a strategy improves the chances that there will be enough non-zero solution points to accurately describe the active regions.

To increase the robustness of the inverse procedures and, particularly, to improve the resolution of standard tomographic methods, the inherent ambiguity of the MEG inverse problem can be minimized by combining MEG data with other functional imaging data (EEG and fMRI) and with structural information on cortical anatomy (from MRI) (e.g., Gevins et al., 1990; Pantev et al., 1990; George et al., 1991; Wang et al., 1992; Dale and Sereno, 1993; Mosher et al., 1993). The approach, however, carries an additional cost, and some of the underlying problems with the standard reconstruction techniques, such as reconstruction bias, remain (e.g., George et al., 1991). The use of anatomical constraints and additional functional imaging modalities with FOCUSS have not yet been explored. We anticipate that these strategies would improve the performance of the algorithm.

We note that specification of the reconstruction space in any inverse method must be approached with care. Selection of a limited reconstruction space that is poorly justified and does not represent all of the true sources can lead

to completely incorrect solutions, rather than partially correct ones. Solutions that are largely concentrated at the borders of the reconstruction space are a good indicator of this kind of problem.

Although neuromagnetic activity is believed to be often localized, distributed patterns of activity are also known to exist, for example in epileptic seizures. Given the severely underdetermined nature of the neuromagnetic imaging problem, an inverse method can only provide solutions that fit its specific objectives, e.g. reconstructing localized vs. extended sources. For example, when the underlying electric activity becomes relatively extensive the accuracy of FOCUSS solutions diminishes and other methods such as the unbiased minimum norm solution may provide more reliable estimates. Presently, methods for estimation of the size of the underlying sources are not established. Although a trade off between methods with different model objectives for different extents of the source activity can be attempted, we expect that in the future, successful methods in MEG will rely on combining a number of algorithms in a single procedure. The potential combination of methods may include an intelligent synthesis of similar reconstructions from several methods and an evaluation of consistency of estimates from different methods with closely related objectives. The differences in objectives determine the sensitivity of methods to the errors in their assumptions and lack of consistency among a group of methods with a similar objective could indicate a mismatch between the objective and the true sources. Given its accuracy, flexibility in the implementation of additional constraints, and the distinctive ability to handle potentially non-unique localized energy solutions, we believe the FOCUSS algorithm has considerable utility for the class of neural electromagnetic inverse problems in which localized sources are expected. As such, it can be a valuable tool either on its own or as a part of a larger imaging strategy that combines a number of techniques.

Acknowledgements

The authors want to thank P.S. Lewis, C.J. Aine, M. Sereno, B. Helton, C.C. Wood, H.A. Schlitt and 3 unknown referees for helpful discussions and/or comments on the manuscript. Their contributions are respectfully acknowledged and appreciated.

The work of I.G. and B.R. has been supported by a grant from NSF, No. MIP-9220550. J.G. has been supported in part by the National Eye Institute Grant No. EY08610 and the Department of Energy Grant No. KP0601000.

Appendix A

The Biot-Savart law gives the relationship between the current density J in some volume conductor V to the

induced magnetic field B :

$$B(\vec{r}) = k \int J(\vec{r}') \times \frac{\vec{r} - \vec{r}'}{|\vec{r} - \vec{r}'|^3} d\vec{r}'$$

where $k = \mu_0/4\pi$ is magnetic permeability coefficient which is constant when a homogeneous conductor is assumed, \vec{r} is the position parameter of current J , \vec{r} is the vector of positions where the measurements are taken and \times indicates a cross-product operation.

Appendix B

Each 3 columns of the matrix G contain spatial weights for the x, y and z components of one source element in I . Each row of G corresponds to one sensor. Thus g_{ij_x} represents the i th entry of the j_x column G_{j_x} of the matrix G . Each 3 elements $g_{ij_x}, g_{ij_y}, g_{ij_z}$ correspond to the 3 orthogonal dipolar components of a j th voxel and an i th sensor and are computed by:

$$[g_{ij_x}, g_{ij_y}, g_{ij_z}] = k \frac{(\vec{r}_i - \vec{r}_j) \times s(\vec{r}_i)}{|\vec{r}_i - \vec{r}_j|^3},$$

where \vec{r}_j is now the location of the j th point source of I and $s(\vec{r}_i)$ denotes the orientation of the i th sensor.

Appendix C

The minimum norm solution to the system $B = GI$ is defined as the solution which has the minimum Euclidian or the 2-norm. The Euclidian norm is defined as $\|I\|_2 = \sqrt{I^T I}$. In biomagnetic imaging, this solution gives the measured field generating current of the minimum mean square strength summed (integrated) over an entire space. Denoting this solution by I_{mn} , it is found as:

$$I_{mn} = G^+ B$$

where G^+ denotes the Moore-Penrose or pseudo-inverse. The standard definition for the pseudo-inverse of an under-determined matrix is $G^+ = G^T(GG^T)^{-1}$, where $(\cdot)^{-1}$ denotes the regular inverse of a matrix. The computation of G^+ using this definition is not advantageous for numerical reasons. A better way to evaluate G^+ is to compute the singular value decomposition of G , i.e., $G = USV^T$, where U and V^T are two orthogonal matrices, with $(\cdot)^T$ indicating a transpose of a matrix and S is a diagonal matrix containing the singular values (s_i) of G . The pseudo-inverse of G is then found by transposing and reversing the order of the

U and V^T matrices and replacing the singular values in S by their reciprocal ($1/s_i$) to obtain

$$G^+ = VS^{-1}U^T.$$

References

- Achim, A., F. Richer and J. Saint-Hilaire (1988) Methods for separating temporally overlapping sources of neuroelectric data. *Brain Topogr.*, 1: 22–28.
- Achim, A., F. Richer and J. Saint-Hilaire (1991) Methodological considerations for the evaluation of spatio-temporal source models. *Electroenceph. clin. Neurophysiol.*, 79: 227–240.
- Cabrera, S.D. and T.W. Parks (1991) Extrapolation and spectral estimation with iterative weighted norm modification. *IEEE Trans. Sign. Proc.*, 39: 842–851.
- Clarke, C.J.S., A.A. Ioannides and J.P.R. Bolton (1989) Localized and distributed source solutions for the biomagnetic inverse problem I. In: S. Williamson, M. Hoke, G. Stroink and M. Kotani (Eds.), *Advances in Biomagnetism*. Plenum, New York, pp. 587–590.
- Crowley, C.W., R.E. Greenblatt and I. Khalil (1989) Minimum norm estimation of current distributions in realistic geometries. In: S.J. Williamson, M. Hoke, G. Stroink and M. Kotani (Eds.), *Advances in Biomagnetism*. Plenum, New York.
- Cuffin, B.N. (1990) Effects of head shape on EEGs and MEGs. *IEEE Trans. Biomed. Eng.*, 37: 44–52.
- Dale, A.M. and M.I. Sereno (1993) Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. *J. Cogn. Neurosc.*, 5: 162–176.
- Darcey, T.M. (1979) *Methods for the Localization of Electrical Sources in the Human Brain and Applications to the Visual System*. PhD Thesis. California Institute of Technology, University Microfilms Order No. 7920351.
- Foster, M. (1961) An application of the Wiener-Kolmogorov smoothing theory to matrix inversion. *J. SIAM*, 9: 387–392.
- George, J.S., P.S. Lewis, D.M. Ranken, L. Kaplan and C.C. Wood (1991) Anatomical constraints for neuromagnetic source models. *SPIE Med. Imaging V: Image Phys.*, 1443: 37–51.
- George, J.S., P.S. Lewis, H.A. Schlitt, L. Kaplan, I. Gorodnitsky and C.C. Wood (1993) Strategies for source space limitation in tomographic inverse procedures. In: *Proc. 9th Int. Conf. on Biomagnetism*, Vienna.
- Geselowitz, D.B. and W.T. Miller (1973) Extracorporeal magnetic fields generated by internal bioelectric sources. *IEEE Trans. Magn.*, MAG-9: 392–398.
- Gevins, A., P. Brickett, B. Costales, J. Le and B. Reutter (1990) Beyond topographic mapping: towards functional-anatomical imaging with 124-channel EEGs and 3-D MRIs. *Brain Topogr.*, 3: 53–64.
- Golub, G.H. and C. Reinsch (1970) Singular value decomposition and least squares solutions. *Num. Math.*, 14: 403–420.
- Gorodnitsky, I.F. (1995) *A Novel Class of Recursively Constrained Algorithms for Localized Energy Solutions: Theory and Application to Magnetoencephalography and Signal Processing*. PhD Thesis. University of California, San Diego, La Jolla, CA.
- Gorodnitsky, I.F. and B.D. Rao (1992) A new iterative weighted norm minimization algorithm and its applications. In: *Proc. 6th SP Workshop on Statistical Signal and Array Processing*, pp. 412–415.
- Gorodnitsky, I.F. and B.D. Rao (1993a) A recursive weighted minimum norm algorithm: analysis and applications. In: *Proc. Int. Conf. on Acoustic, Speech and Sign. Processing, III*, pp. 456–459.

- Gorodnitsky, I.F. and B.D. Rao (1993b) Convergence analysis of a class of adaptive weighted norm extrapolation algorithms. In: Proc. 27th Asilomar Conf. on Signals, Systems and Computers, Vol. 1, pp. 339–343.
- Gorodnitsky, I.F. and B.D. Rao (1994a) Analysis of regularization error in Tikhonov regularization and truncated singular value decomposition methods. In: Proc. 28th Asilomar Conf. on Signals, Systems and Computers, Vol. 1.
- Gorodnitsky, I.F. and B.D. Rao (1994b) Truncated total least squares regularization algorithm for underdetermined problems. In: Proc. 7th SP Workshop on Statistical Signal and Array Processing, pp. 19–22.
- Gorodnitsky, I.F., J.S. George and P.S. Lewis (1992) Weighted linear estimator procedures for neuromagnetic source reconstruction. *J. Neuroimaging*, 1: 54.
- Grossberg, S. (1988) *Neural Networks and Natural Intelligence*. MIT Press, Cambridge, MA.
- Hämäläinen, M.S. and R.J. Ilmoniemi (1984) Interpreting measured magnetic fields of the brain: estimates of current distributions. Tech. Rep. TKK-F-A559, Helsinki University of Technology, Espoo.
- Hämäläinen, M.S. and R.J. Ilmoniemi (1994) Interpreting magnetic fields of the brain – minimum norm estimates. *Med. Biol. Eng. Comput.*, 32: 35–42.
- Hansen, P.C. (1990) Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numerical rank. *SIAM J. Sci. Statist. Comput.*, 11: 503–518.
- Helmholtz, H. (1853) Über einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern, mit Anwendung auf die thierisch-elektrischen Versuche. *Pogg. Ann. Phys. Chem.*, 89: 211–233, 353–377.
- Henderson, C.J., S.R. Butler and A. Glass (1975) The localization of equivalent dipoles of EEG sources by the application of electric field theory. *Electroenceph. clin. Neurophysiol.*, 39: 117–130.
- Ioannides, A.A., J.P.R. Bolton and C.J.S. Clarke (1990) Continuous probabilistic solutions to the biomagnetic inverse problem. *Inverse Probl.*, 6: 523–542.
- Kullman, W.H., K.D. Jandt, K. Rehm, H.A. Schlitt, W.J. Dallas and W.E. Smith (1989) A linear estimation approach to biomagnetic imaging. In: S.J. Williamson, M. Hoke, G. Stroink and M. Kotani (Eds.), *Advances in Biomagnetism*. Plenum, New York, pp. 571–574.
- Lawson, C. and R. Hanson (1974) *Solving Least-Squares Problems*. Prentice-Hall, Englewood Cliffs, NJ.
- Lee, H., D.P. Sullivan and T.H. Huang (1987) Improvement of discrete band-limited signal extrapolation by iterative subspace modification. In: Proc. ICASSP, Dallas, TX, pp. 1569–1572.
- Medvick, P.A., P.S. Lewis, C.J. Aine and E.R. Flynn (1989) Monte Carlo analysis of localization errors in magnetoencephalography. In: S.J. Williamson, M. Hoke, G. Stroink and M. Kotani (Eds.), *Advances in Biomagnetism*. Plenum Press, New York.
- Mosher, J., P. Lewis and R. Leahy (1992) Multiple dipole modeling and localization from spatial temporal MEG data. *IEEE Trans. Biomed. Eng.*, 39: 541–557.
- Mosher, J.C., M.E. Spencer, R.M. Leahy and P.S. Lewis (1993) Error bounds for EEG and MEG dipole source location. *Electroenceph. clin. Neurophysiol.*, 86: 303–321.
- Pantev, C., M. Hoke, K. Lehnertz, B. Lütkenhöner, G. Fahrenndorf and U. Stöber (1990) Identification of sources of brain neuronal activity with high spatiotemporal resolution through combination of neuromagnetic source localization (NMSL) and magnetic resonance imaging (MRI). *Electroenceph. clin. Neurophysiol.*, 75: 173–184.
- Plonsey, R. and D.B. Heppner (1967) Considerations of quasi-stationarity in electrophysiological systems. *Bull. Math. Biophys.*, 29: 657–664.
- Romani, G. and R. Leoni (1984) Localization of cerebral sources by neuromagnetic measurements. In: G. Weinberg et al. (Eds.), *Biomagnetism: Applications and Theory*. Pergamon Press, New York, pp. 205–220.
- Romani, G.L. and P. Rossini (1988) Neuromagnetic functional localization. *Brain Topogr.*, 1: 5–21.
- Roth, B.J., N.G. Sepulveda and J.P. Wikswo, Jr. (1989) Using a magnetometer to image a two-dimensional current distribution. *J. Appl. Physiol.*, 65: 361–372.
- Sano, A. (1993) Optimally regularized inverse of singular value decomposition and application to signal extrapolation. *Sign. Proc.*, 30: 163–176.
- Sarvas, J. (1987) Basic mathematical and electromagnetic concepts of the biomagnetism inverse problem. *Phys. Med. Biol.*, 32: 11–22.
- Scherg, M. (1989) Fundamentals of dipole source potential analysis. In: M. Hoke et al. (Eds.), *Auditory Evoked Magnetic Fields and Potentials: Advances in Audiology*. Karger, Basel.
- Scherg, M. and D. Von Cramon (1986) Evoked dipole source potentials of the human auditory cortex. *Electroenceph. clin. Neurophysiol.*, 65: 344–360.
- Sidman, R.D., V. Giambalvo, T. Allison and P. Bergey (1978) A method for localization of sources of human cerebral potentials evoked by sensor stimuli. *Sens. Proc.*, 2: 116–129.
- Singh, M., D. Doria, V. Henderson, G. Huth and J. Beatty (1984) Reconstruction of images from neuromagnetic fields. *IEEE Nucl. Sci., NS-31*: 585–589.
- Smith, W.E., W.J. Dallas, W.H. Kullman and H.A. Schlitt (1990) Linear estimation theory applied to the reconstruction of a 3-D vector current distribution. *Appl. Opt.*, 29: 658–667.
- Snyder, A.Z. (1991) Dipole source localization in the study of EP generators: a critique. *Electroenceph. clin. Neurophysiol.*, 80: 321–325.
- Spencer, M.E., R.M. Leahy, J.C. Mosher and P.S. Lewis (1992) Adaptive filters for monitoring localized brain activity from surface potential time series. In: Proc. 26th Asilomar Conf. on Signals, Systems and Computers, Vol. 1, pp. 152–156.
- Supek, S. and C.J. Aine (1993) Model order determination and limits of source resolution for multisource neuromagnetic data: simulation studies. *IEEE Trans. Biomed. Eng.*, 40: 529–540.
- Tikhonov, A.N. (1963) Regularization of incorrectly posed problems. *Soviet Math.*, 4: 1624–1627.
- Van Veen, B., J. Joseph and K. Hecox (1992) Localization of intracerebral sources of electrical activity via linearly constrained minimum variance spatial filtering. In: Proc. 26th Asilomar Conf. on Signals, Systems and Computers, Vol. 1, pp. 157–161.
- Wang, J.Z., S.J. Williamson and L. Kaufman (1992) Magnetic source images determined by a leadfield analysis: the unique minimum norm least squares estimation. *IEEE Trans. Biomed. Eng.*, 39: 665–675.
- Williamson, S.J. (1990) Neuromagnetic localization of sensory and cognitive functions. In: *Third Swiss EEG/EP Mapping Meeting*, May 1990.
- Wood, C. and G. McCarthy (1984) Principal component analysis of event-related potentials: simulation studies demonstrate misallocation of variance across components. *Electroenceph. clin. Neurophysiol.*, 59: 249–260.