Source Localization in Magnetoencephalography using an Iterative Weighted Minimum Norm Algorithm

Irina F. Gorodnitsky and Bhaskar D. Rao
Dept. of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92037-0411

John George
Biophysics Group
Los Alamos National Laboratory
Los Alamos, NM 87545

Abstract

Imaging of brain activity based on magnetoencephalography (MEG) requires high resolution estimates that closely approximate the spatial distribution of the underlying currents. We examine the physics of the MEG problem to motivate the development of a new algorithm that meets its unique requirements. The technique is a nonparametric, iterative, weighted norm minimization procedure with posteriori constraints. We develop the algorithm and determine the necessary requirements for convergence. Issues of initialization and bias equalization for MEG reconstruction, and techniques for analysis of noisy data are discussed.

1 Introduction

Magnetoencephalography measures the weak magnetic fields outside the head produced by electric currents in the brain. The electric currents are due to the simultaneous firings of large populations of neurons, often concentrated within localized discrete aggregates that extend tens to hundreds of square millimeters [1]. Solving the neuromagnetic inverse problem, the identification and characterization of the generators of the measured fields, may reveal much about the function of the brain. The complexity of the problem however requires development of special estimation techniques.

We begin our approach by examining the physical and mathematical nature of the problem to develop an algorithm to accomodate its unique requirements. The inverse problem is intrinsically ill-posed, as a myriad of current configurations within the head can produce the same external magnetic field. The underdetermination is exacerbated because source reconstruction at any given instant is based on only a few measurements contaminated by noise. Furthermore, the reconstructed current configuration must be consistent with physiological evidence which suggests that the energy of the signal is often located in a few small discrete centers. In signal processing terms, the problem is nearfield, signal-to-noise ratio is low and three-dimensional field determination is required.

The nature of the problem suggests a number of requirements for neuromagnetic current estimation. Imaging of small discrete sources requires a high resolution technique. *Accurate three-dimensional localization of neuromagnetic sources requires that the spatial extent of the sources be closely reconstructed.* The source depth estimate is particularly sensitive to the spatial extent reconstruction. None of the currently used methods addresses this issue adequately.

Two principal source models are presently employed in MEG with variations closely tied to these two. Several methods employ the multiple dipole model assumption where each active neuronal aggregate is modeled as an equivalent point current dipole [2]. The resulting low order model simplifies the estimation process. However, the inaccuracy of such models can lead to substantial errors in source localization when the size of the sources is significant relative to the distance at which they are recorded [3]. Alternatively, the nonparametric minimum 2-norm model [4] yields high dimensional ($O(10^3)$), poorly constrained solutions, limited to the outer surface of the reconstruction volume. This type of solution is at odds with expectations of non-surficial localized sources. More recently, we and others have investigated weighted minimum norm models.

In ill-posed problems, extra assumptions about the source signal are necessary to obtain a unique estimate. However, in general the available *a priori* information in neuromagnetic imaging is not sufficient to constrain the solution to the required degree of high accuracy. Here we present a new nonparametric method in MEG. It is an iterative weighted norm minimization[1] technique constrained *a priori* to a class of models representing highly localized but arbitrary shaped sources, and on *posteriori* information. These kind of assumptions are consistent with expectations based on physi-

[1] All norms in this paper refer to the 2-norm

1058-6393/92 $03.00 © 1992 IEEE

ological knowledge. Provided proper initialization, the algorithm accurately recovers the extent and locations of the sources. An iterative algorithm for MEG with *posteriori* constraints was first suggested in [5], but the results were presented in 2-D reconstructions only. The results indicated that the method had limited capacity to spatially concentrate sources. A similar algorithm was also suggested in the spectral estimation context [6].

The nonparametric techniques have a number of general advantages. The often complicated step of parametric source modeling is avoided, and a wealth of supplementary information can be incorporated relatively easily into the estimate. The algorithm finds a solution based on an instantaneous snapshot of data and can be used with nonstationary processes such as highly dynamic brain activity. In fact, the approach requires no assumptions about the nature of the signals, for example that they are uncorrelated or are special statistical processes.

2 Discrete Forward Model

The neuromagnetic forward problem is naturally continuous and can be described by the *Biot-Savart* law. Here we employ a discrete model approximation to the law to simplify the presentation. Advantages of the continuous model in MEG are undermined by noise in the data. Different methods of discretization lead to approximations of differing accuracy. Because the accuracy of the discretization is irrelevant for the purpose of this paper, we choose a simple method of spatial sampling. The algorithm could be used equally well with any type of discretized model.

In discrete form, a set of measurements of the magnetic field $b(\vec{r_i})$ at locations $\vec{r_i}$ around the head are related to the discretized neural electric current x by a linear transformation, with each equation corresponding to a single measurement

$$b(\vec{r_i}) = Ax, \tag{1}$$

where A is $m \times n$, $m < n$ matrix of rank m. Each three elements of x and the three corresponding columns of A describe the three components (x, y and z) of a single point current vector. For the sensor location $\vec{r_i}$, the three elements of the A matrix corresponding to the single point current vector located at $\vec{r_j}$ are defined as

$$[a_{ij_x}\ a_{ij_y}\ a_{ij_z}] = k\frac{(\vec{r_i} - \vec{r_j}') \times s(\vec{r_i})}{|\vec{r_i} - \vec{r_j}'|^3}, \tag{2}$$

where $s(\vec{r_i})$ denotes the orientation of the *ith* sensor and k is a constant in our model.

3 Inverse Algorithm

3.1 Background

In discrete form the problem is becomes one of finding a solution to an underdetermined system of linear equations

$$Ax = b. \tag{3}$$

The infinite set of solutions to (3) is a linear variety

$$x = x_{mn} + v, \tag{4}$$

where v is any vector in the null space of A, $x_{mn} = A^+b$ is the minimum 2-norm (Euclidian norm) solution, and A^+ denotes the Moore-Penrose inverse. Our objective as discussed above is to find the true solution from the set (4) of possible ones. An element of the linear variety (4) can be reached through a weighted minimum norm solution as [7]

$$x = W(AW)^+b = W^2A^T(AW^2A^T)^{-1}b \tag{5}$$

where W is a symmetric weight matrix. This is equivalent to solving a constrained norm minimization problem.

$$Minimize\ \|W^{-1}x\|, \quad subject\ to\ Ax = b. \tag{6}$$

The weighted norm minimization procedure provides the basis for our algorithm.

3.2 Algorithm

The constraint information in W is assumed to be *a priori*. We maintain that in general, the *a priori* information available to build the weight matrix in MEG is not sufficient to provide the required degree of constraint. Instead, we propose using *posteriori* constraints. The motivation for the algorithm comes from the nature of the localized sources often encountered in the MEG inverse problem. The reconstructed current has to be sparse, with few active elements and the rest being zeros. A successful algorithm should find a low dimensional solution to the problem. We define low dimensional solutions to be basic solutions which are estimates with the number of non-zero elements not exceeding the number of measurements m. In the algorithm, a diagonal weight matrix W_k is constructed from elements of the solution x_k of the previous iterative step. The kth step of the algorithm then solves (6) by minimizing the norm $\|W^{-1}x_{k+1}\| = \sum_{i=1}^{n}\left(\frac{x^i_{k+1}}{x^i_k}\right)^2$. The effect, as explained below, is a gradual increase of a few elements of x at the expense of others until only a minimum number of elements remain, and the algorithm converges to a low dimensional estimate. The algorithm can be written as follows [8,9,10]

$$W_k = diag(x_k) \tag{7}$$

$$q_{k+1} = (AW_k)^+ b \tag{8}$$

$$x_{k+1} = W_k q_{k+1} \tag{9}$$

We postpone discussion of the initialization of the iterative algorithm, along with other key issues effecting convergence, until after the properties of the algorithm are described.

3.3 Properties of the Algorithm

Property 1. Convergence to a localized solution: Except for a few special cases, the algorithm converges to a low dimensional solution[2]. In the rare cases of special geometry, the algorithm can converge to a solution with a slightly larger number of nonzero elements. A setup of sensor configuration that is not perfectly symmetric avoids this problem entirely. The convergence is determined by the minimization at each step of the norm

$$||W_k^{-1} x_{k+1}||^2 = \sum_{i=1}^{n} \left(\frac{x_{k+1}^i}{x_k^i} \right)^2 .$$

If the element x_k^i is small in magnitude, it will tend to decrease the magnitude of x_{k+1}^i in the next iteration, unless the corresponding column in A is necessary in fitting b, when taken in conjunction with the current linear combination of other columns of A. This implies that once a favorable weighting is obtained, the selected elements continue to be favored at the expense of others until convergence to the solution with only these few elements is completed.

Property 2. Stationary Fixed Points: The algorithm is a discrete nonlinear process (a map) whose geometry consists of stable fixed points (s-f-ps) which are the basic p-dimensional ($p \leq m$) solutions, and unstable fixed points (u-f-ps) located in l-dimensional ($l > m$) hyperplanes. The higher dimensional solutions in the cases of special symmetry are saddle points. The s-f-ps are the solutions to the reduced square systems of equations with the left hand sides obtained from the original system by selecting all sets of linearly independent columns from the A matrix. The algorithm converges to one of the s-f-ps unless the initialization falls exactly on an u-f-p.

Property 3. Robustness: The method is robust with respect to the choice of W. The convergence to some fixed point with $p < m$ non-zero elements occurs for any choice of W that weighs the corresponding p columns of A favorably compared to the remaining $n-p$ columns regardless of the precise values of the weights. It implies that the algorithm is insensitive to changes in the weights that preserve the relative differences, and that the weight function does not need to be a solution to the system (1).

[2] The dimension of the exact estimate never drops below the dimension of x_0 but asymptotically converges to a low dimensional vector. In finite precision the solution becomes low dimensional.

3.4 Initialization and Bias Equalization

A particular localized solution is not unique. The solution to which the algorithm converges depends on two key factors: the initialization and the size of the basins of attraction around each localized solution. The issues are fairly complex. Here we briefly present methods that have proven reliable for many of the cases likely to be encountered in MEG. A complete discussion will appear in [11].

The s-f-ps divide the solution space into regions of attraction. The final solution is determined by the region into which the initialization x_0 falls. The closer the x_0 is to the true solution, the greater is the chance of it being in the right region of attraction. The objective of the initialization is then to capture as many true features of the signal as possible. Solutions to the given problem, for example the minimum norm solution, a weighted minimum norm solution that utilizes *a priori* information about the process, or a solution based on a highly localized or even multiple dipole model with a small residual added to zero valued elements are all potential candidates for initialization. Alternatively, the initialization does not need to be a solution to the given system, allowing great flexibility, including properly guessed functions. In analyses of a time sequence of measurements, a preceding snapshot within a sequence may be used to initialize the reconstruction of the next. We have found that initialization with the minimum norm estimate, scaled to remove the bias as described below, works well for many cases.

A special property of MEG and similar physical problems allows one to judge the acceptability of a found solution. *In MEG, in general, there is only one highly localized solution with dimension less than m, and all other localized solutions are of the highest possible dimension m.* If the solution obtained is a thicket of m disconnected single point sources, it is often an indication that the true signal was not found. If this occurs, different initializations may be tried until the highly localized solution is found. A well constructed initial estimate that fails, sometimes may be perturbed in some logical direction to yield a localized solution.

The outcome of the algorithm is also determined by the shape of the regions of attraction of localized solutions. An initialization that is in close proximity to some s-f-p can still be in the attracting region of another far removed s-f-p. To converge to the true solution with a small attracting region, the initialization has to start very close to it. This is highly undesirable, and so we want to equalize the areas of the regions as much as possible. The size of the regions is set dynamically based on the particular b and A. While A can be scaled to at least partly counter the bias contained in it, we generally do not have control over b. Thus the best we can do is to equalize general regional size differences due to the bias in A. Such strategies do not preclude the existence of singular cases with poor convergence. This issue is still being actively investigated.

In neuromagnetic inverse procedures based on norm minimization, the bias contained in A is toward solutions closest to the sensors, because magnetic field strength decreases in proportion to the square of the distance from the source. Countering of even the known bias is not easy, and it is not clear if exact cancellation can be achieved at all. The intrinsic bias works similarly to the way external weights bias the solution toward some regions of the source space in the weighted minimum norm algorithm. Some columns of A can have large effective weights due to the large magnitudes of their elements relative to other columns. This can produce disproportionally large values in the corresponding solution elements. We adopt an obvious approach to scaling A, using counter weights by incorporating them into the existing weights of our iterative algorithm. Since the bias in A affects most iterations, the scaling is used at each step rather than just for initialization.

The goal of the bias removal is to eliminate any natural size difference between the columns of A. Because the size of vectors is not uniquely defined there is no simple recipe for compensation. The total size of the column vectors as measured by the 1- or 2-norm, as well as disproportionately larger or smaller individual column elements or a greater range of element magnitudes within a column compared with other columns, can all contribute to the bias. All these factors must be adjusted simultaneously and are dependent on a particular A. Even within a single application like MEG, the matrix A is not universal but is dependent on the sensor geometry. We have had good success with first normalizing the columns of A by dividing each element in a column by the 2-norm of that column and fine tuning the weights further. For the geometry we use, which is described below, the normalization tends to slightly bias the deeper sources. Hence we scale the normalization weights by a factor proportional to the distance to a source element from the center of the sphere.

It is important to note that the A matrix in MEG is frequently ill-conditioned, which can lead to poor results. Numerically robust computational schemes are required. If numerical instability persists, the methods for stabilizing the solution in the presence of noise that we discuss in Section 5 may be used in noiseless simulations as well.

4 Simulations

A single dipole case and a combination of a single dipole and two extended sources were used to demonstrate the performance of the algorithm. The geometry of the 37-sensor BTI Squid Magnetometer centered on the z-axis above the head was used in the simulations. The 3-D reconstruction examples are shown in Figs. 1 and 2. The 3-D results were projected onto a 2-D surface, using an orthographic projection. The values of the distributed estimates in Figs. 1(b) and 2(b) were encoded into five intensity levels with values below a threshold not shown. For clarity a maximum intensity projection was used; only the highest of the three-dimensional values projecting onto the same point on the plane was plotted. The minimum norm estimate was used to initialize the single dipole solution. A weighted minimum norm estimate with the scaled A matrix as described above, was used for initialization in the multiple sources example. The algorithm correctly reconstructed all sources and converged after only 5 and 4 iterations respectively in the two examples. In the case of significantly extended sources, the algorithm sometimes produces an extra z-component in voxels immediately under a true source location. Simple matching of all three components of the current eliminates this artifact.

5 Reconstruction with Noisy Data

No analysis of an ill-posed inverse problem is complete without examining the stability of the solution to noise in the measurements. Because of limited space we only briefly state the results here. In general, neuromagnetic source reconstructions that provide an exact fit to the data are extremely sensitive to noise in the measurements. Even though our technique provides low dimensional solutions, the sensitivity remains. Techniques to stabilize the solution are required. Here we discuss the two existing methods for handling noisy data: regularization and singular value decomposition (SVD) truncation. A third method - the spectral approach is equivalent in principle to the SVD truncation method but has worse numerical properties. Because our algorithm works to reduce the dimension of the solution, the SVD truncation method facilitates convergence. Regularization, on the other hand, increases the size of the smallest eigenvalues of A, which does not allow the solution to drop below m dimensions. We successfully employed the SVD truncation method in conjunction with our algorithm (7-9) to correctly recover sources from noisy simulated data. In our implementation, the number of the truncated singular values was increased successively with each iteration. Ultimately more robust methods utilizing both techniques may be developed.

References

[1] S.J. Williamson. "Neuromagnetic Localization of Sensory and Cognitive Functions". *Third Swiss EEG-EP Mapping Meeting*, May 1990.

[2] D.B. Geselowitz and W.T. Miller III. "Extracorporeal magnetic fields generated by internal bioelectric sources". *IEEE Trans. Magn.*, MAG-9:392–398, 1989.

[3] A.Z. Snyder. "Dipole source localization in the study of EP generators: a critique". *EEG. Cl. Neur.*, 80:321–325, 1991.

[4] C.W. Crowley, R.E. Greenblatt, and I. Khalil. "Minimum norm estimation of current distributions in realistic geometries". In S.J. Williamson et. al., editor, *Advances in Biomagnetism*. New York:Plenum, 1989.

[5] A. A. Ioannides, J. P. R. Bolton, and C. J. S. Clarke. "Continuous Probabilistic Solutions to the Biomagnetic Inverse problem". *Inverse Problems*, pages 523–542, 1990.

[6] H. Lee, D. P. Sullivan, and T. H. Huang. "Improvement of discrete band-limited signal extrapolation by iterative subspace modification". In *Proc. ICASSP*, volume 3, pages 1569–1572, Dallas, Texas, April 1987.

[7] L.C. Potter and K.S. Arun. "Energy Concentration in Band-Limited Extrapolation". *IEEE Trans. ASSP*, ASSP-37(7):1027–1041, July 1989.

[8] I.F. Gorodnitsky, J.S. George, and P.S. Lewis. "Weighted Linear Estimator Proceedures for Neuromagnetic Source Reconstruction". *Am. Soc. of Neuroimaging Annual Meeting*, Feb. 1992.

[9] I.F. Gorodnitsky and B.D. Rao. "A new iterative weighted norm minimization algorithm and its applications". In *Sixth SP Workshop on Statistical Signal and Array Processing*, Oct. 1992.

[10] I.F. Gorodnitsky, J.S. George, and P.S. Lewis. "A weighted iterative algorithm for neuromagmetic imaging". *IEEE in Med. and Biol.* 14^{th} *Intr. Conf. Sattellite Symposium on Neuroscience and Technology*, Nov. 1992.

[11] I.F. Gorodnitsky at al. "Neuromagnetic source imaging using iterative weighted minimum norm algorithm". *manuscript in preparation.*

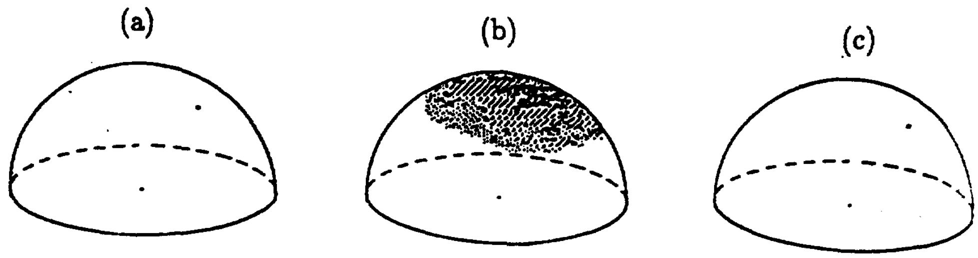

Figure 1: A 2-D projection of the 3-D reconstruction of a single dipole: (a) the true source position. (b) the minimum norm estimate, (c) the estimate from the new algorithm after 5 iterations.

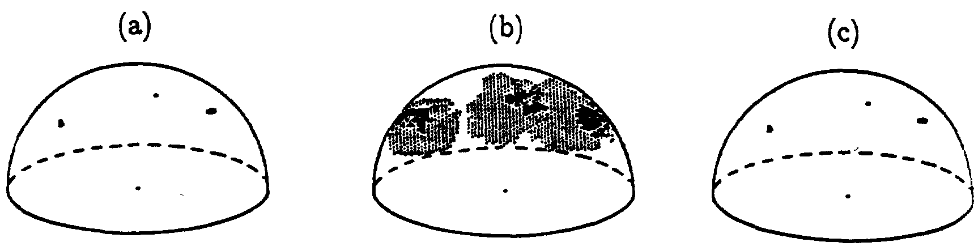

Figure 2: A 2-D projection of the 3-D reconstruction of a single dipole and two extended sources: (a) the true source position, (b) the minimum norm estimate, (c) the estimate from the new algorithm after 5 iterations.