

Lead Field Basis for FEM Source Localization

*David Weinstein, Leonid Zhukov,
Chris Johnson
Email: dmw@cs.utah.edu,
zhukov@cs.utah.edu, crj@cs.utah.edu*

UUCS-99-014

Center for Scientific Computing and Imaging
Department of Computer Science
University of Utah
Salt Lake City, UT 84112 USA

October 5, 1999

Abstract

In recent years, significant progress has been made in the area of EEG/MEG source imaging. Source imaging on simple spherical models has become increasingly efficient, with consistently reported accuracy of within 5mm. In contrast, source localization on realistic head models remains slow, with sub-centimeter accuracy being the exception rather than the norm. A primary reason for this discrepancy is that most source imaging techniques are based on lead-fields. While the lead-field for simplified geometries can be easily computed analytically, an efficient method for computing realistic domain lead-fields has, until now, remained elusive. In this paper, we propose two efficient methods for computing realistic EEG lead-field bases: the first is element-oriented, and the second is node-oriented. We compare these two bases, discuss how they can be used to apply recent source imaging methods to realistic models, and report timings for constructing the bases.

Introduction

Neural processes - perception, coordination, cognition - are carried out via the propagation of electrical impulses through the brain. These impulses give rise to electromagnetic fields that can be measured extra-cranially by sensitive recording devices. The measurement of electric potentials over time is referred to as electroencephalography, or EEG, and the measurement of magnetic signals as magnetoencephalography, or MEG. The practice through which we infer the inter-cranial sources that give rise to these measurements is termed the *neural source imaging* problem. Neural source imaging is a fundamental problem in neuroscience. Learning precisely which regions of the brain are active at a particular time is a central problem in fields ranging from cognitive science to neuropathology to surgical planning.

The distribution of an electromagnetic field in the head is described by the linear Poisson equation:

$$\nabla \cdot (\sigma \nabla \phi) = \nabla \cdot \mathbf{J}^s, \quad \text{in } \Omega \quad (1)$$

with no-flux Neumann boundary conditions on the scalp:

$$\sigma(\nabla \phi) \cdot \mathbf{n} = 0, \quad \text{on } \Gamma_\Omega \quad (2)$$

where σ is the electrical conductivity tensor, ϕ are the electric potentials, and \mathbf{J}^s are the electric current sources. From the linearity of (1), it follows that the mapping from electric sources within the cranium to scalp recordings on the outside of the scalp can be represented by a linear operator \mathbf{L} . Given a particular configuration of sources, $\mathbf{s} \in \mathbf{S}$, the resultant recordings, $\mathbf{r} \in \mathbf{R}$, and the noise in the system, \mathbf{n} , we represent this relation as $\mathbf{r} = \mathbf{L}\mathbf{s} + \mathbf{n}$. \mathbf{L} is the so-called *lead-field matrix* [1] and contains information about the geometry and conductivity of the model.

The problem of interest in source imaging is this: given a set of recordings, \mathbf{r} , knowing \mathbf{L} , and making certain assumptions about \mathbf{n} , determine \mathbf{s} , the set of sources that gave rise to those recordings. The solution to this problem consists of two steps: building the forward matrix \mathbf{L} , and “inverting” it. Finding the inverse of \mathbf{L} is an ill-posed problem and its solution requires regularization. There exist many different regularization methods, as well as many papers describing their application to EEG/MEG [2, 11, 14]. All of them assume that \mathbf{L} either is known *a priori* or can be easily constructed. While it is true that \mathbf{L} is easy to construct for simple geometries like spheres (or in other cases where an analytic solution to the forward problem exists), building the \mathbf{L} matrix is much more complicated for geometries based on real patient data.

In this paper, we will concentrate on the construction of the \mathbf{L} matrix. Below, we propose two efficient methods for building \mathbf{L} for realistic EEG models with realistic geometries as well as inhomogeneous and anisotropic conductivities.




Figure 1: Depiction of the lead-field mapping from sources (labeled 1-12 inside head) and electrodes (labeled 1-5 outside head) to the matrix \mathbf{L} . Each source corresponds to a column of \mathbf{L} , each electrode corresponds to a row of \mathbf{L} , and each entry of \mathbf{L} corresponds to the potential measured at a particular electrode due to a particular source.

Background

The lead-field matrix, \mathbf{L} , defines a projection from current sources at discrete locations in the cranium to potential measurements at discrete recording sites on the scalp. That is, entry \mathbf{L}_{ij} corresponds to the potential that would be measured at recording site \mathbf{r}_i due specifically to source \mathbf{s}_j . Sources are traditionally defined by three orthogonal dipoles, and source positions are generally located on a regular grid of hexahedral cells covering the domain of interest. A 2D depiction of this traditional lead-field is shown in Fig. 1.

In a number of application papers, researchers have been able to compute the \mathbf{L} basis by exploiting analytic equations for each entry in \mathbf{L} . The equations for the matrix entries, or *kernels*, for each method can be found in Mosher’s review of forward and inverse EEG and MEG methods [10]. For most of them, the \mathbf{L} matrix is constructed one element at a time by evaluating the analytic expression for the potential at each recordings site (the rows), due to a source at each location in the domain (the columns).

While fast analytic expressions exist for the potentials in spherical models, these solutions cannot be applied to realistic head models. Rather, a complete forward numerical solution must be computed in order to determine the electrode potentials due to each source. These realistic model forward solutions can be generated using the boundary element method (BEM) or finite element method (FEM) [5, 6, 8].

Both methods can faithfully represent the complex boundaries and inhomogeneous regions of realistic models. The FEM has the additional advantage that it can capture anisotropic conductivities of the domain. For these reasons, we have chosen to use the FEM for our lead-field calculations.

The main idea behind the FEM is to reduce a continuous problem with infinitely many unknown field values to a finite number of unknowns by discretizing the solution region into elements. The value at any point in the field can then be approximated by interpolation functions within the elements. These interpolation functions are specified in terms of the field values at the corners of the elements, points known as nodes. We note that for linear interpolation functions, the electric field is constant within an element.

Given a geometric model, the FEM proceeds by assembling the matrix equations to build the stiffness matrix \mathbf{A} . This can be done using, for example, a Rayleigh-Ritz or Galerkin method [4]. Boundary conditions are then imposed and source currents are applied. These boundary and source conditions are incorporated within the right hand side of the system (vector \mathbf{s}_n). Details of the FEM method can be found in [4, 5, 8, 18]. Application of the FEM reduces Poisson’s equation to the linear system:

$$\mathbf{A}\phi = \mathbf{s}_n \tag{3}$$

where ϕ are the unknown potentials at the nodes of the volume.

When we use a realistic FEM model (3mm resolution), and the above forward numerical construction, we may require upwards of a minute of CPU time on modern architectures (*e.g.*, an SGI MIPS R10000 processor) to solve a forward simulation. Because of this computational expense, it seems feasible to build a lead-field for only a very sparse grid of sources. Even a 16^3 grid (12mm resolution), for example, with three orthogonal dipole components per cell, would require over a day of continuous computation to build \mathbf{L} .

Methods

Below, we introduce two novel methods for constructing the lead-field matrix \mathbf{L} . Using these methods, we are able to construct \mathbf{L} matrices for a head model with 3mm cell resolution (64^3 grid) in under 10 minutes.




Figure 2: Depiction of the element-oriented lead-field basis. The orthogonal dipole in each element corresponds to a column of \mathbf{L} and, just as in Fig.1, each electrode corresponds to a row of \mathbf{L} , and each entry of \mathbf{L} corresponds to the potential measured at a particular electrode due to a particular source.

Element Basis

As discussed above, the traditional method of constructing the \mathbf{L} matrix is to place three orthogonal sources in each cell of a volume domain, as was shown in Fig. 1. For each dipole source, the voltages at the electrodes are computed. After setting one of the electrodes to ground, we store the potentials as a column of \mathbf{L}_e . For a volume with N elements, this requires computing $(N \times 3)$ forward solutions.

For our new basis, we will still be considering three orthogonal dipole sources per cell, but the shape and distribution of our cells will be different than the traditional configuration. Rather than using a hexahedral, regular grid of source cells, we have chosen to use the unstructured set of tetrahedral cells from our underlying finite element mesh. A 2D depiction of this lead-field is shown in Fig. 2, and a cut-through of our 3D patient-based finite element mesh is shown in Fig. 3.

In constructing the \mathbf{L} matrix, we would like to achieve the maximal possible resolution of sources for our model (one arbitrarily oriented dipole per tetrahedral element). We would also like to take advantage of the fact that when using the finite element method (FEM), we compute the potentials not only on the surfaces (as in the boundary element method), but through the entire volume. We can achieve both goals using the principle of *reciprocity*.

The reciprocity principle was introduced into the biophysical domain by Helmholtz [3] and was first applied to the problem of electrocardiography by Plonsey [15]. It was subsequently adapted to the EEG problem by Rush and Driscoll in 1969 [16], when




Figure 3: Cut-through of our finite element model. The different colored elements correspond to different conductivity regions. Green, blue, yellow, purple and light-blue correspond to scalp, skull, gray matter, white matter and cerebro-spinal fluid, respectively. The model contains approximately 320,000 elements and 60,000 nodes.




Figure 4: Depiction of our reciprocity-based method. A unit current is applied between electrodes ‘3’ and ‘G’. The reciprocity principle states that the voltage difference between ‘3’ and ‘G’ due to a dipole source \mathbf{p} placed in element e^n will be equal to the dot product of \mathbf{p} and the electric field in e^n .

they proved the applicability of reciprocity to anisotropic conductors. The reciprocity principle states that given a dipole (an equivalent source) and a need to know the resulting potential difference between two points A and B, it is sufficient to know the current density at the dipole location resulting from a unit current through points A and B:

$$\frac{(\mathbf{E} \cdot \mathbf{p})}{I} = \phi_{\mathbf{A}} - \phi_{\mathbf{B}} \quad (4)$$

So, rather than placing a source in every element and for each pairing computing a forward solution, we can “invert” this process: we place a source and sink at pairs of electrodes and for each pair compute the resulting electric field in all of the elements. We can then use the reciprocity principle to reconstruct the potential differences at the electrodes for a source placed in any element.

First, we choose one electrode as ground. For each of the other M electrodes, one at a time, we place a unit current source perpendicular to the surface at that electrode and a unit current sink at the ground electrode. The forward solution is then computed, resulting in potentials at each element in the domain. We take the gradient of this potential field, yielding the electric field at each element in the head. Reading off the x, y, and z components of the electric field in each element yields a single row of $\mathbf{L}_{\mathbf{e}}$.

This process is depicted in Fig. 4. The electric field due to a source at electrode ‘3’

and a sink at electrode ‘G’ yields the electric field indicated with the arrows in each element. To compute the potential difference between electrodes ‘3’ and ‘G’ due to a dipole source \mathbf{p} placed in element e^n , we take the dot product of \mathbf{p} and the electric field in element e^n .

This process is repeated for each of the M source electrodes, producing the \mathbf{L}_e matrix:

$$\begin{bmatrix} L_1^{1,x} & L_1^{1,y} & L_1^{1,z} & \dots & L_1^{N,x} & L_1^{N,y} & L_1^{N,z} \\ L_2^{1,x} & L_2^{1,y} & L_2^{1,z} & \dots & L_2^{N,x} & L_2^{N,y} & L_2^{N,z} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ L_M^{1,x} & L_M^{1,y} & L_M^{1,z} & \dots & L_M^{N,x} & L_M^{N,y} & L_M^{N,z} \end{bmatrix} \cdot \begin{bmatrix} s_{1,x} \\ s_{1,y} \\ s_{1,z} \\ \dots \\ s_{N,x} \\ s_{N,y} \\ s_{N,z} \end{bmatrix} = \begin{bmatrix} \phi_1 \\ \phi_2 \\ \dots \\ \phi_M \end{bmatrix} \quad (5)$$

where $\mathbf{L} = \mathbf{E}/I$, N = the number of elements, and $M + 1$ = the number of electrodes. For a system with $M + 1$ recording sites, our reciprocity-based method requires computing only M forward solutions.

Node Basis

The above method for deriving the element-oriented lead-field constructs an \mathbf{L}_e basis that maps dipole components placed at the elements to potentials at the scalp-recording electrodes. One can also think of an alternative formulation based on the divergence of the source current density vector at each node, rather than three orthogonal current dipoles within each element. This lead-field formulation, which we refer to as *node-oriented*, is depicted in Fig. 5. Our node-oriented basis is derived directly from the finite element stiffness matrix, \mathbf{A} , and right-hand side vector, \mathbf{s}_n .

It is straightforward to solve the well-conditioned system (3):

$$\boldsymbol{\phi} = \mathbf{A}^{-1} \mathbf{s}_n \quad (6)$$

to recover the potentials, $\boldsymbol{\phi}$, throughout the volume. For source imaging, however, we are interested not in the potentials everywhere in the volume, but only in the potentials at those few nodes corresponding to scalp electrodes recording sites. We introduce a matrix \mathbf{R} that selects just the electrode potentials from $\boldsymbol{\phi}$. \mathbf{R} is a $[K \times M]$ matrix (number of nodes by one less than the number of recording electrodes). Each row of \mathbf{R} contains a single non-zero entry: the value 1.0 located at the column corresponding to the node index for that electrode.




Figure 5: Depiction of the node-oriented lead-field basis. In contrast to Figs. 1 and 2, where the matrix columns corresponded to orthogonal dipoles, the columns now correspond to nodes. The node-oriented basis has approximately 94% fewer columns than the element-oriented basis, and is better suited to distributed source configurations.

From (6), we now select a subset of ϕ by applying \mathbf{R} :

$$\mathbf{r} = \mathbf{R}\phi = \mathbf{R}\mathbf{A}^{-1}\mathbf{s}_n \quad (7)$$

The above $\mathbf{R}\mathbf{A}^{-1}$ operator is a node-oriented lead-field basis, which we term \mathbf{L}_n .

We can exploit the sparse nature of \mathbf{R} in order to compute $\mathbf{R}\mathbf{A}^{-1}$ in an efficient manner. Since \mathbf{R} contains only M non-zero entries, we need to construct only the corresponding M columns of \mathbf{A}^{-1} . This is accomplished using the identity:

$$\mathbf{A}(\mathbf{A}^{-1})_m = \mathbf{I}_m \quad (8)$$

to solve for each column m of \mathbf{A}^{-1} that corresponds to an electrode node. As with the construction of the \mathbf{L}_e basis above, this technique requires generating M forward solutions.

Lead-field comparison: \mathbf{L}_e vs \mathbf{L}_n

The two lead-fields, element-oriented and node-oriented, and the projections they define,

$$\mathbf{r} = \mathbf{L}_e \mathbf{s}_e \quad (9)$$

$$\mathbf{r} = \mathbf{L}_n \mathbf{s}_n \quad (10)$$

differ in several relevant ways.

The $\mathbf{L_e}$ formulation is based on having a dipole moment of a particular strength and orientation in each element. Because of this, $\mathbf{L_e}$ is most useful for reconstructing discrete dipolar sources. This is an appropriate method for localizing very focal neural activity, such as epileptic seizures or specific motor control tasks.

In contrast, the node-oriented $\mathbf{L_n}$ lead-field is defined with values at the nodes. This means $\mathbf{L_n}$ will work best for recovering less focal, more distributed-type sources. Such a solution should be well-suited to capturing diffuse cognitive events, such as language processing or the performance of complex tasks.

The $\mathbf{L_e}$ basis is $[M \times (N \times 3)]$ (one less than the number of recording electrodes by three times the number of elements). There are often as many as a million elements in a finite element mesh, whereas there are typically only $M = 64$ or 128 recording electrodes. When using this lead-field basis for source imaging (*i.e.*, solving $\mathbf{r} = \mathbf{L_e} \mathbf{s_e}$ for a particular set of electrode recordings), it is clear that the solution will be grossly under-determined.

The node-oriented basis, $\mathbf{L_n}$, is somewhat smaller: $[M \times K]$ (one less than the number of recording sites by the number of nodes). While there are still typically many more nodes (often as many as a hundred thousand) than electrodes, the system $\mathbf{r} = \mathbf{L_n} \mathbf{s_n}$ is less under-determined than the element-based formulation. A typical unstructured finite element mesh contains six times as many elements as nodes, and the $\mathbf{L_e}$ basis requires three sources per element. This means the $\mathbf{L_n}$ basis will be only about $\frac{1}{(6 \times 3)}$ as large as $\mathbf{L_e}$, and therefore will have 94% fewer degrees of freedom and a 94% smaller null-space. These advantages will be significant as we investigate source localization methods based on these lead-fields.

Numerical Simulations

We constructed a realistic finite element head model from a volume MRI scan. The MRI data was segmented at the Brigham and Women’s hospital [17], and a mesh was constructed using Krysl’s variational Delaunay algorithm [7]. The full mesh contained 320,000 elements and 60,000 nodes. From this model, we selected 64 electrode recording sites at appropriate scalp surface nodes.¹

¹We note that if the true electrode locations do not correspond precisely to nodes in the mesh, the bases can still be constructed. The construction of $\mathbf{L_e}$ does not change at all, and the construction of $\mathbf{L_n}$ is only slightly modified. \mathbf{R} can now have three non-zero entries per row (the electrode’s barycentric weightings for the nodes of the scalp triangle containing it), rather than a single non-

We then built \mathbf{L}_e using the reciprocity method described above. For each electrode current source, we solved a forward simulation in order to compute the electric field at each element in our model, and we stored that as a row of \mathbf{L}_e . Each solution required on average 8 seconds of wall-clock time using 8 SGI MIPS R10000 processors, resulting in a total of 9 minutes to compute all of \mathbf{L}_e .

Next, we computed \mathbf{L}_n , the node-oriented basis, by iteratively solving for $\mathbf{R}\mathbf{A}^{-1}$ one row at a time. Each row required approximately 7 seconds of wall-clock time, totaling 8 minutes to construct \mathbf{L}_n .

We then simulated a focal temporal seizure by placing a dipole source in that region of our finite element model and running a forward simulation. We stored the resulting electrode potential as ϕ . To validate our bases, we then recovered the position, orientation, and magnitude of the temporal source using the source localization algorithm described below.

Applications

One of the simplest applications of the \mathbf{L}_e operator is to construct a single dipole *cost-function field*. The cost-function field is constructed as follows: position a dipole in every element one at a time; find the best magnitude and orientation for a source at that location (all locations within the same element are equivalent); compute the two-norm misfit between the forward solution $\hat{\phi}$ due to that dipole, and the “measured” data, ϕ . For each element, computing the best magnitude and orientation for a source requires that we look for an optimal solution vector, \mathbf{s}_e , with only three non-zero entries, corresponding to the (p_x, p_y, p_z) components for a source in that element.

Given that we have *exactly* three non-zeroes in the solution vector, we can form an over-determined system with three unknowns, where the matrix consists of three columns of \mathbf{L}_e . For a dipole p in element n , we form the following system:

$$p_x^n \begin{bmatrix} L_1^{n,x} \\ L_2^{n,x} \\ \vdots \\ L_M^{n,x} \end{bmatrix} + p_y^n \begin{bmatrix} L_1^{n,y} \\ L_2^{n,y} \\ \vdots \\ L_M^{n,y} \end{bmatrix} + p_z^n \begin{bmatrix} L_1^{n,z} \\ L_2^{n,z} \\ \vdots \\ L_M^{n,z} \end{bmatrix} = \begin{bmatrix} \phi_1 \\ \phi_2 \\ \vdots \\ \phi_M \end{bmatrix} \quad (11)$$

The least squares solution to this system will give the three components of the dipole that provides the best fit to the data. We can then compute the two-norm misfit, C , between ϕ and $\hat{\phi}$, and assign this value to the element:

$$C = \|\phi - \hat{\phi}\| = \|\phi - \bar{p}_x \mathbf{L}_e^{n,x} - \bar{p}_y \mathbf{L}_e^{n,y} - \bar{p}_z \mathbf{L}_e^{n,z}\| \quad (12)$$

zero entry. This modification does not affect the efficiency of the construction.




Figure 6: Visualization of the cost-function field. The misfit between the “measured” data (due to the depicted dipole source) and the forward solution due to the optimal source for each element has been color mapped to an axial cutting plane.

Repeating the above steps for all the elements in the model, we construct the cost-function field. Clearly, the minimum of this field corresponds to the optimal dipole position for a single dipole model. We note that while the description above is for the \mathbf{L}_e basis, a similar cost-function can be evaluated for the \mathbf{L}_n basis. For the \mathbf{L}_n basis we solve a 4×4 system, where the four components for an element correspond to contributions from the nodes of that element, rather than from three orthogonal dipole components.

Using the SCIRun problem solving environment [13], we created the \mathbf{L}_e matrix and constructed the cost-function field for an occipital dipole source. A visualization of this field is shown in Fig. 6. The field values have been color-mapped onto a cutting plane positioned axially just below the source. Note the linear increase in the cost-function for elements farther and farther away from the source, as well as the sharp discontinuities in and around the skull elements due to the large conductivity difference between bone and the surrounding tissue.

Though computing the the cost function for a single element requires only a matrix-vector multiplication and a 3×3 system solve, for large models (*e.g.*, with over 100,000 elements), the total time to build the field can reach up to an hour. If our goal is not to compute this field everywhere in the volume, but just to perform a *source localization*, we can be more efficient with our computation. Rather than evaluate the cost-function in all of the elements of the field, we can use a simple search technique such as downhill simplex [12] and evaluate just those locations along a search path. For our model, this reduced the number of cost-function evaluations from 100,000 to 230, resulting in a source localization that can be achieved within a few seconds. However, we note that the downhill simplex search has the short-coming that it may converge on a local, rather than global minimum, and that by definition single dipole source localization will be inadequate for characterizing complex source configurations.

A snapshot of the downhill simplex algorithm running in SCIRun is shown in Fig. 7. The simplex, shown with gray arrows connected by yellow lines, is converging on the true dipole source, indicated by the red arrow.

Advantage of Precomputing $\mathbf{L_e}$

We now compare the time required to localize the source using a downhill simplex search with the standard method and with our new method. For the standard method:

- We search without the benefit of the $\mathbf{L_e}$ matrix; therefore, each evaluation of C requires a forward solution.

For our method:

- We first compute $\mathbf{L_e}$.
- We then search using $\mathbf{L_e}$; each evaluation of C simply requires three matrix-vector multiplications, and the solution of a 3×3 system.

We started the simplex search from a number of different initial positions, from which we were able to localize the source with an average of 230 evaluations. The time required for those 230 evaluations varied significantly between our method, using $\mathbf{L_e}$ and the standard method without $\mathbf{L_e}$.

Using the standard method, we did not precompute the basis $\mathbf{L_e}$, and each evaluation required approximately 9 seconds to solve each of the forward simulations. The total time for the downhill simplex search was 69 minutes.




Figure 7: Visualization of the downhill simplex algorithm converging to a dipole source. The simplex is indicated by the gray vectors joined by yellow lines. The true source is indicated in red. The surface potential map on the scalp is due to the forward solution of one of the simplex vertices, whereas the potentials at the electrodes (shown as small spheres) are the “measured” EEG values (potentials due to the true source).

Using our new method, we first computed the \mathbf{L}_e basis. As described above, this required 9 minutes of CPU time. However, once we had this basis, we were able to iterate through all 230 evaluations of the simplex search in less than 3 seconds. Compared to the 69 minutes for the standard approach, our new method showed a 7x improvement.

Clearly it is worth the 9 minutes of overhead to build \mathbf{L}_e for source localization, since for any search requiring more evaluations than electrodes, this will yield a win. Furthermore, for source imaging methods that require the localization of multiple dipoles [19], these performance improvements will be multiplicative.

Future Work

Having developed an efficient means for computing finite element lead-fields for realistic models, we find many potential applications to consider for these lead-field bases. The most obvious application is to implement the latest source imaging methods [2, 9] for a realistic FEM EEG model and compare the solution accuracy of these methods for various source configurations.

Additionally, we would like to use the lead-field basis as a foundation for developing more physiologically appropriate re-weighting methods. We believe it is possible to better constrain the source imaging solution by incorporating model constraints from the finite element stiffness matrix, \mathbf{A} .

Finally, we would like to evaluate other, more efficient search methods for multi-source localization. As we described, the lead-field basis enables rapid source evaluation. We would now like to incorporate more powerful search methods, such as simulated annealing, in order to make the search for configurations of multiple dipoles computationally tractable.

Acknowledgments

This work was supported in part by the National Science Foundation, the Department of Energy, the National Institutes of Health, and the Utah State Centers of Excellence Program. The authors would like to thank Robert van Uiter for his assistance in generating figures and Yarden Livnat for his assistance in editing.

References

- [1] H.C. Burger and J.B. van Milaan. Heart-vector and leads. Part I. *Br. Heart J.*, 8:157–61, 1946.
- [2] I.F. Gorodnitsky, J.S. George, and B.D. Rao. Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. *Electroencephalography and clinical Neurophysiology*, 95:231–251, 1995.
- [3] H. Helmholtz. Über einige gesetze der vertheilung elektrischer strome in körperlichen leitern, mit anwendung auf die thierischelektrischen versuche. *Annals of Physical Chemistry*, 29:211–233 and 353–377, 1853.
- [4] J. Jin. *The Finite Element Method in Electromagnetics*. John Wiley and Sons, 1993.
- [5] C.R. Johnson. Numerical methods for bioelectric field problems. In J.D. Bronzino, editor, *The Biomedical Engineering Handbook*, pages 162–180. CRC Press, Boca Ratan, 1995.
- [6] C.R. Johnson. Computational and numerical methods for bioelectric field problems. *Critical Reviews in BioMedical Engineering*, 25(1):1–81, 1997.
- [7] P. Krysl and M. Oritz. Variational delaunay approach to the generation of tetrahedral finite element meshes. *International Journal for Numerical Methods in Engineering*, (to appear), 1999.
- [8] C.E. Miller and C.S. Henriquez. Finite element analysis of bioelectric phenomena. *Crit. Rev. in Biomed. Eng.*, 18:181–205, 1990.
- [9] J.C. Mosher and R.M. Leahy. Recursively applied MUSIC: A framework for EEG and MEG source localization. Technical Report LA-UR-96-3829 and USC-SIPI Report No. 304, Los Alamos National Lab, 1996.
- [10] J.C. Mosher, R.M. Leahy, and P.S. Lewis. EEG and MEG: Forward solutions for inverse methods. *Transactions on Biomedical Engineering*, 46, 1999.
- [11] J.C. Mosher, P.S. Lewis, and R.M. Leahy. Multiple dipole modeling and localization from spatio-temporal MEG data. *IEEE Transactions on Biomedical Engineering*, 39(6):541–557, 1992.
- [12] J.A. Nedler and R. Mead. A simplex method for function minimization. *Compt. J.*, 7:308–313, 1965.

- [13] S.G. Parker, D.M. Weinstein, and C.R. Johnson. The SCIRun computational steering software system. In E. Arge, A.M. Bruaset, and H.P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 1–40. Birkhauser Press, 1997.
- [14] R.D. Pasqual-Marqui, C.M. Michel, and D. Lehmann. Low resolution electromagnetic tomography: A new method for localizing electrical activity in the brain. *International Journal of Psychophysiology*, 18:49–65, 1994.
- [15] R. Plonsey. Reciprocity applied to volume conductors and the ecg. *IEEE Transactions of Bio-Medical Electronics*, pages 9–12, 1967.
- [16] S. Rush and D.A. Driscoll. Eeg electrode sensitivity - an application of reciprocity. *IEEE Trans. Biomed. Eng.*, 16:15–22, 1969.
- [17] W.M. Wells, W.E.L. Grimson, R. Kikinis, and F.A. Jolesz. Statistical intensity correction and segmentation of MRI data. In *Visualization in Biomedical Computing*, pages 13–24, 1994.
- [18] Y. Yan, P.L. Nunez, and R.T. Hart. Finite-element model of the human head: scalp potentials due to dipole sources. *Medical & Biological Engineering & Computing*, 29:475–481, 1991.
- [19] L. Zhukov, D. Weinstein, and C. Johnson. Independent component analysis for eeg source localization in realistic head models. *IEEE-EMB*, 2000 (to appear).