ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 23 (2004) S289–S299

Mapping human brain function with MEG and EEG: methods and validation

F. Darvas, D. Pantazis, E. Kucukaltun-Yildirim, and R.M. Leahy*

Department of Electrical Engineering, Signal and Image Processing Institute, University of Southern California, Los Angeles, CA 90089 2564, United States

Available online 25 September 2004

We survey the field of magnetoencephalography (MEG) and electroencephalography (EEG) source estimation. These modalities offer the potential for functional brain mapping with temporal resolution in the millisecond range. However, the limited number of spatial measurements and the ill-posedness of the inverse problem present significant limits to our ability to produce accurate spatial maps from these data without imposing major restrictions on the form of the inverse solution. Here we describe approaches to solving the forward problem of computing the mapping from putative inverse solutions into the data space. We then describe the inverse problem in terms of low dimensional solutions, based on the equivalent current dipole (ECD), and high dimensional solutions, in which images of neural activation are constrained to the cerebral cortex. We also address the issue of objective assessment of the relative performance of inverse procedures by the free-response receiver operating characteristic (FROC) curve. We conclude with a discussion of methods for assessing statistical significance of experimental results through use of the bootstrap for determining confidence regions in dipole-fitting methods, and random field (RF) and permutation methods for detecting significant activation in cortically constrained imaging studies.

Keywords: MEG; EEG; Inverse methods; Dipole fitting; Imaging methods; ROC analysis; Statistical parametric tests

Introduction

Magnetoencephalography (MEG) and electroencephalography (EEG) provide a unique window on the human brain. Both modalities measure the electromagnetic signals produced by electrical activity in the brain. It is widely believed that the primary source of these signals is current flow in the apical dendrites of pyramidal cells in the cerebral cortex. Coherent activation of a large number of pyramidal cells small area of cortex can be modeled as an equivalent current dipole (ECD), which, because of the columnar organization of cortex, is oriented normally to its surface (Okada et al., 1997). The current dipole is therefore the basic element used to represent neural activation in EEG- and MEG-based inverse methods, and these dipoles are often constrained to lie within cortical gray matter.

EEG data are measurements of potential differences on the scalp resulting from ohmic currents induced by electrical brain activity. Instrumentation for EEG consists of a set of scalp electrodes coupled to high-impedance amplifiers and a digital data acquisition system. Because EEG signals are produced by ohmic current flow in the head, they are highly sensitive to the conductivity of the brain, skull, and extracranial tissue. Consequently, solving an inverse problem to localize regions of neural activation requires accurate knowledge of these properties. In contrast, MEG measures the magnetic field outside the head induced by current flow within the brain. In this case, while the signals are affected by ohmic currents, the major contributor to the signal is the field induced directly by neural current generators, sometimes called the primary currents. It is these current generators, rather than the secondary or ohmic currents, that are of interest in solving the inverse problem because they are localized to regions in which the brain is activated during a particular evoked response study. Thus, the underlying electromagnetic principles of MEG and EEG make the latter less sensitive to the conductivity of the head and hence better able to provide accurate localization of neural activation from noninvasive measurements (Hamalainen et al., 1993). MEG instrumentation typically requires the use of superconducting SQUID-based magnetometers housed in a magnetically shielded room. Recently, MEG and EEG have come to be viewed as complementary rather than competing modalities, and most MEG protocols routinely include simultaneous acquisition of multichannel EEG data.

MEG and EEG differ fundamentally from fMRI and PET, the other major functional brain imaging modalities, both in the physiological processes that are measured and in the properties of the inverse problem that are solved to produce functional images. Neuronal activation can be detected indirectly with both fMRI and PET through imaging of localized contrast changes that result from the hemodynamic effects of this activation. Aside from the potential problem that the neuronal activation and hemodynamic changes

* Corresponding author. Department of Electrical Engineering, Signal and Image Processing Institute, University of Southern California, 3740 McClintock Avenue, Los Angeles, CA 90089 2564. Fax: +1 213 740 4651.
E-mail address: leahy@sipi.usc.edu (R.M. Leahy).
Available online on ScienceDirect (www.sciencedirect.com.)

1053-8119/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2004.07.014

may not always be co-located, these methods are limited in temporal resolution by hemodynamic time constants on the order of 1 s. In contrast, MEG and EEG signals are the direct extracranial manifestations of neuronal activation and can record changes at the millisecond scale at which they occur. Consequently, it may be possible with MEG or EEG not only to detect networks of neuronal assemblies, but also to determine causality within these networks. The other major difference between MEG or EEG and fMRI or PET is in the form of the inverse problems. While the latter are mildly ill-posed, both fMRI and PET involve reconstruction of images from a large number of independent samples so that stable images can always be computed at the expense of some mild loss in resolution. In contrast, the MEG or EEG inverse problem involves estimation of brain activation from at most a few hundred spatial samples so that highly restrictive assumptions are required to compute a stable inverse. This presents significant challenges both in the selection of a model that reflects these assumptions and also in the interpretation of the resulting inverse solutions.

In this article, we survey the issues underlying the problem of model selection in EEG and MEG inverse problems and describe the major classes of inverse procedures. We also describe how the associated forward mapping is solved and the impact of the assumed head model on the accuracy of these methods. We then describe objective methods for assessing and comparing the relative performance of different inverse procedures. Finally, we discuss the issue of assessing the statistical significance of different inverse solutions computed from experimental data.

Forward models

To estimate the neural sources of scalp potentials and neuromagnetic fields, we must first be able to solve the associated forward problem, that is, we need a forward model that maps a source of known location, strength, and orientation to an array of EEG or MEG sensors. The frequencies of interest are relatively low (typically <100 Hz) so the forward model is governed by the quasi-static versions of Maxwell's equations (Hamalainen et al., 1993). Because the permeability of biological tissue is approximately that of free space, the electrical conductivity and geometry of the head are the factors that determine the forward model.

The current density in the head associated with neural activation is the sum of the primary current $\boldsymbol{j}_p$, which is assumed to be the physical correlate of neuronal activity, and the volume or return currents

$$\boldsymbol{j}_v = -\sigma(\boldsymbol{r})\nabla U(\boldsymbol{r}) \tag{1}$$

that result from passive ohmic conductance; here $\sigma(\boldsymbol{r})$ represents the electrical conductivity of the head and $U(\boldsymbol{r})$ the electric potential. It is the primary currents that are of interest when solving EEG and MEG inverse problems because they represent neuronal activation; however, the effects of volume currents must still be considered when solving the forward problem because they contribute to scalp potentials and neuromagnetic fields.

The elemental source model in MEG or EEG is the equivalent current dipole (ECD), which represents an idealized point source. Due to the linearity of Maxwell's equations, the forward model for arbitrary source configurations can be written as a linear superposition of forward models for these point sources. The measurement m for either an EEG or MEG sensor at location $\boldsymbol{r}_d$, due to an ECD with orientation and strength $\boldsymbol{q}$ at location $\boldsymbol{r}_s$, can be written as:

$$\boldsymbol{m}(\boldsymbol{r}_d) = \boldsymbol{G}(\boldsymbol{r}_d, \boldsymbol{r}_s, \sigma(\boldsymbol{r}))\boldsymbol{q} \tag{2}$$

and the measurement for an arbitrary distribution of sources $\boldsymbol{j}_p(\boldsymbol{r}_s)$ inside a volume V as:

$$\boldsymbol{m}(\boldsymbol{r}_d) = \int_V \boldsymbol{G}(\boldsymbol{r}_d, \boldsymbol{r}_s, \sigma(\boldsymbol{r}))\boldsymbol{j}_p(\boldsymbol{r}_s)d^3\boldsymbol{r}_s \tag{3}$$

Note that here the effects of the return currents $\boldsymbol{j}_v$ have been factored into $\boldsymbol{G}$ because they are an implicit function of the primary currents. The specific form of the forward model $\boldsymbol{G}(\boldsymbol{r}_d, \boldsymbol{r}_s, \sigma(\boldsymbol{r}))$ depends on the assumed shape and conductivity of the head.

Head conductivity

The conductivity profile of the head, $\sigma(\boldsymbol{r})$, is typically determined by segmenting an anatomical MR image into its various components, e.g., skin, skull, CSF, and brain tissue. Recent studies use up to 10 different tissue types, adding soft bone, fat, eye, cerebellum, muscle, and a separation of white and gray matter in their models (Haueisen et al., 2002). A good estimate of skull conductivity and shape (Huiskamp et al., 1999) is of particular importance because the large difference in conductivity of skull and soft tissue has a large impact on scalp potentials. Segmentation of the anatomical image is typically performed either manually or using automated software such as BrainSuite (Shattuck and Leahy, 2002), as illustrated in Fig. 1, or BrainVoyager (Goebel, 1997). A conductivity value is then assigned to each tissue type. These are typically standard values that have been measured in vitro using excised tissue (Geddes and Baker, 1967). Alternatively, electrical impedance tomography (EIT) can be used to compute in vivo estimates of conductivities. EIT can make use of an existing EEG measurement setup, injecting current through one electrode and measuring potential differences with the remaining electrodes. Determination of the conductivity $\sigma(\boldsymbol{r})$ from the known current source and measured surface potentials poses an ill-posed nonlinear inverse problem (Goncalves et al., 2000; Saulnier et al., 2001) so that it is difficult to obtain accurate estimates, and

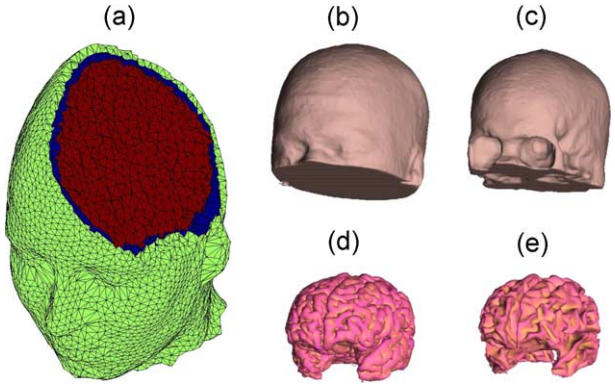

Fig. 1. (a) Sample FEM tetrahedral mesh with tissue classification: green = skin, blue = bone, red = CSF or brain. Tessellated surfaces extracted automatically from a T1-weighted MR volume using BrainSuite: (b) scalp, (c) skull, (d) gray matter, and (e) white matter.

consequently these methods are rarely used in practice. In most forward models, conductivities are assumed isotropic, but recently anisotropic conductivity, caused by fiber bundles in white matter for example, has been estimated using diffusion tensor MR imaging (Tuch et al., 1999). This information can be directly incorporated into the forward computation using the FEM methods discussed below.

With a segmentation of the head volume and a set of conductivities for each tissue type, a solution of Laplace's equation,

$$\nabla(\sigma(\boldsymbol{r})\nabla U(\boldsymbol{r})) = \nabla \boldsymbol{j}_p(\boldsymbol{r}) \tag{4}$$

can be computed. Eq. (4) with the boundary condition that no currents leave the head volume and that the potential U is fixed at some reference point gives a full mathematical description of the EEG forward problem. Once the potential $U(\boldsymbol{r})$ is known, one can also compute the forward model for MEG. Closed form analytic solutions exist for MEG and EEG if we assume the head consists of a set of nested concentric spheres, each with homogeneous and isotropic conductivity. For realistic human head geometries numerical solutions using Boundary Element Methods (BEM), Finite Element Methods (FEM), and Finite Difference Methods (FDM) are required (Fuchs et al., 2001; Johnson, 1997).

Spherical head models

The MEG forward solution for the spherical case is independent of the specific conductivities of the layers of the sphere. All that is required to compute the solution is the center of the sphere and the location and orientation of the sensors (Sarvas, 1987). The sphere center is typically computed from anatomical landmarks: nasion, left and right preauricular points. If scalp surface data are also available, the single sphere model in MEG can be expanded by locally fitting a sphere to each sensor (Huang et al., 1999). Phantom studies have shown that this local spherical approximation of the head can perform with similar accuracy to a BEM solution using the realistic surface geometry of the phantom (Leahy et al., 1998). Further refinement of the analytic spherical forward models to approximate realistic geometries can be achieved by using spherical harmonic expansions to more closely match the head shape and extending the analytical solutions to these nonspherical surfaces (Nolte, 2003).

Closed form solutions also exist for EEG for the spherical head geometry, although in this case the conductivities of each layer must be known (Geselowitz, 1967). Computation cost for this model is higher than for MEG because multilayer EEG models require summation of an infinite series. Computationally efficient approximations have been developed to overcome this problem (Zhang, 1995). While spherical models provide good approximations for MEG forward solutions, this is typically not the case for EEG. Spherical models can give a reasonable good approximation in superior regions of the brain, where the head is roughly spherical in shape, but in general, realistically shaped models are necessary to achieve good localization accuracy throughout the brain (Buchner et al., 1996).

Numerical solution of the forward problem: BEM and FEM

A BEM can be used to compute a numerical solution for Eq. (4) under the assumption that $\sigma(\boldsymbol{r})$ is piecewise homogeneous. In this case, the differential equation for the potentials on the surfaces can be transformed to an integral equation over the interfaces of the different conductivity compartments plus the analytical infinite homogeneous space solution for the primary current (Mosher et al., 1999). By discretizing these boundaries, the integrals are turned into sums, and the solution of the problem is found by solving a set of linear equations. The surfaces are discretized by triangle meshes and the complexity of the linear system is defined by the total number of nodes in all meshes. One disadvantage of the BEM is that the system of linear equations is dense so that a model with a set of surfaces with a total of 10,000 nodes would require about 1 GB of memory and take on the order of several minutes to solve on a modern PC. Repeated recomputation of the forward solution is prohibitive if BEM is used directly as part of an iterative inverse method. However, precomputation and the use of interpolation methods can reduce the cost of using BEM to that of spherical models (Ermer et al., 2001). An important consideration in the use of BEM methods is that numerical accuracy depends on the size of the tessellation elements relative to the distance of the source to the nearest boundary, so that finer tessellation's are required when sources are close to boundaries (Fuchs et al., 2001).

The finite element model can overcome many of the problems of the boundary element method by discretizing Eq. (4) over the entire head volume (Awada et al., 1997; van den Broek et al., 1996). The nodes of the grid are typically connected by tetrahedrons for an irregular grid, as illustrated in Fig. 1, or by cubes in the case of a regular grid, and the unknown potential $U(\boldsymbol{r})$ is interpolated over these nodes. The conductivity can therefore be defined locally for each element and the model allows the incorporation of an anisotropic conductivity tensor instead of scalar values for $\sigma(\boldsymbol{r})$. The discretization for FEM also leads to a linear system of equations. Because the system only involves interactions between neighboring nodes, the matrix for the FEM is sparse and it can be efficiently stored and inverted. Unlike the BEM, the FEM in its standard formulation does not contain the analytical infinite homogeneous space solution for the current dipole. However, the FEM can be modified to include the analytical potential, thus allowing point like sources (Awada et al., 1997; Schimpf et al., 2002). If cubic elements are used, the FEM becomes similar to the finite difference method. The disadvantage of FDM, and of using cubic elements in FEM, is that all cubes must be of the same size and therefore local mesh refinement is not possible. An advantage of the regular elements is a well-conditioned system matrix, which leads to faster convergence. The use of tetrahedral elements allows for elements that can vary in size and can therefore be used to model the geometry very accurately. Tetrahedral meshes can also be locally refined to increase the accuracy of the model.

Choice of head model

In principal, an FEM model incorporating spatially varying anisotropic conductivity will provide the most accurate forward model. For an isotropic model with piecewise constant conductivity, the BEM offers similar accuracy. However, it is important to note that accuracy of these methods is dependent both on knowledge of the true conductivities and on the numerical details of the FEM and BEM implementation, in particular the resolution of the mesh on which the solution is computed. No method currently exists to produce accurate in vivo high resolution images

of conductivity, and in practice piecewise constant approximate conductivities are typically used. Furthermore, because of computational costs, BEM and FEM meshes are often of relatively low resolution. Under these circumstances, and particularly for MEG, the spherical model may often be an adequate alternative that avoids the potential for numerical instability associated with BEM and FEM. For example, in a phantom study based on dipoles in a human skull (Leahy et al., 1998), we found that there was little difference in localization error (<1 mm averaged over 32 dipoles) between a locally fitted sphere model and BEM for both MEG and EEG data. When BEM and FEM methods are used, it is important that their accuracy be assessed for the specific application in which they are employed.

Inverse methods

With a forward model $\boldsymbol{G}(\boldsymbol{r}_d, \boldsymbol{r}_s, \sigma(\boldsymbol{r}))$ and a set of measurements $\boldsymbol{m}(t)$ at time t, the inverse problem is to find the current density distribution $\boldsymbol{j}_p(\boldsymbol{r}_s, t)$ that generated the measurements according to Eq. (3). With a discrete set of c detector locations $(\boldsymbol{r}_d^1, \boldsymbol{r}_d^2,\ldots,\boldsymbol{r}_d^c)$ and $\boldsymbol{m}(t) = (m(\boldsymbol{r}_d^1, t), m(\boldsymbol{r}_d^2, t),\ldots,m(\boldsymbol{r}_d^c, t))$, Eq. (3) can be written as

$$\boldsymbol{m}(t) = \int_V \boldsymbol{G}(\boldsymbol{r}_s)\boldsymbol{j}_p(\boldsymbol{r}_s, t)d^3\boldsymbol{r}_s \tag{5}$$

where $\boldsymbol{G}(\boldsymbol{r}_s) \in R^{c\times 3}$ maps an ECD at $\boldsymbol{r}_s$ onto the set of c sensors. There are two key approaches to finding $\boldsymbol{j}_p$ and thus solving the inverse problem. One can assume that $\boldsymbol{j}_p$ consists of a small set of q ECDs, which can be written as a sum of delta functions:

$$\boldsymbol{j}_p(\boldsymbol{r}_s, t) = \sum_{i=1}^{q} \delta\left(\boldsymbol{r}_s - \boldsymbol{r}_q^i\right)\boldsymbol{q}_i(t)$$

The integral in Eq. (5) becomes a sum

$$\boldsymbol{m}(t) = \sum_{i=1}^{q} \boldsymbol{G}\left(\boldsymbol{r}_q^i\right)\boldsymbol{q}_i(t).$$

If one stacks the $\boldsymbol{q}_i(t)$ into a vector $\boldsymbol{q}(t) = (\boldsymbol{q}_1(t), \boldsymbol{q}_2(t),\ldots,\boldsymbol{q}_q(t))^T \in R^{3q}$ and the forward models $\boldsymbol{G}(\boldsymbol{r}_q^i)$ in a matrix $\mathbf{G} = (\boldsymbol{G}(\boldsymbol{r}_q^1), \boldsymbol{G}(\boldsymbol{r}_q^2),\ldots,\boldsymbol{G}(\boldsymbol{r}_q^q)) \in R^{c\times 3q}$, we can write,

$$\boldsymbol{m}(t) = \boldsymbol{G}\left(\boldsymbol{r}_q^1, \boldsymbol{r}_q^2,\ldots,\boldsymbol{r}_q^q\right)\boldsymbol{q}(t) \tag{6}$$

If the orientations of each of the ECDs in $\boldsymbol{q}_i(t)$ are constrained to be invariant over time, $\boldsymbol{q}_i(t)$ can be written as $\boldsymbol{q}_i s_i(t)$, where $s_i(t)$ encodes the strength and time course of the source and $\boldsymbol{q}_i$ the normalized orientation. The data and source time series are sampled at discrete intervals, that is, $s_i(t) \rightarrow \boldsymbol{s}_i^T = (s_i(t_1), s_i(t_2),\ldots,s_i(t_T))^T$ and $\boldsymbol{M} = (\boldsymbol{m}(t_1),\boldsymbol{m}(t_2),\ldots,\boldsymbol{m}(t_T))$, with T the number of samples. The constrained orientation can be combined with the forward matrix to form the matrix $\mathbf{A} = [\boldsymbol{G}(\boldsymbol{r}_q^1)\boldsymbol{q}_1, \boldsymbol{G}(\boldsymbol{r}_q^2)\boldsymbol{q}_2,\ldots,\boldsymbol{G}(\boldsymbol{r}_q^q)\boldsymbol{q}_q] \in R^{c\times q}$ and the time series s_i can be stacked as $\boldsymbol{S}^T = (s_1, s_2,\ldots,s_q)^T$, leading to the discrete, constrained version of Eq. (5)

$$\boldsymbol{M} = \boldsymbol{A}\left(\boldsymbol{r}_1, \boldsymbol{q}_1, \boldsymbol{r}_2, \boldsymbol{q}_2,\ldots, \boldsymbol{r}_q, \boldsymbol{q}_q\right)\boldsymbol{S}^T = \boldsymbol{A}\boldsymbol{S}^T \tag{7}$$

for dipolar sources (Baillet et al., 2001; Mosher and Leahy, 1998). If the orientations are unconstrained, we can generalize Eq. (6) in the obvious way to $\boldsymbol{M} = \boldsymbol{G}\boldsymbol{Q}^T$. The inverse problem is then to estimate the number, location, and time series of each of the ECDs.

The second approach to the inverse problem is to allow a large number of sources that represent a continuous distribution of neural current generators. Because these generators are believed to lie in the cerebral cortex, the sources are usually constrained to lie on the cortical surface, with orientation normal to the surface. The inverse problem then reduces to the solution of a linear inverse problem in the set of unknown time series for each elemental source on a tessellated representation of the cortical surface (Fig. 2). In this case, the forward problem can also be formulated as in Eq. (7), but because the source locations are fixed (one per surface element), the inverse problem involves estimation of $\boldsymbol{S}^T$ only.

Dipole fitting and scanning methods

Dipole fitting methods use the formulation in Eq. (7) and solve for the parameters $(\{\boldsymbol{r}, \boldsymbol{q}\}, \boldsymbol{S}^T)$ (Mosher and Leahy, 1998; Scherg, 1990). The problem is often solved using least squares, to find $(\{\boldsymbol{r}, \boldsymbol{q}\}, \boldsymbol{S}^T)$ such that

$$\|\boldsymbol{M} - \boldsymbol{A}(\{\boldsymbol{r}, \boldsymbol{q}\})\boldsymbol{S}^T\|_2^2 \rightarrow \min \tag{8}$$

The number of sources is unknown and the cost function is nonconvex in the source locations, which results in the existence of multiple local minima. Consequently, minimization of Eq. (8) is a challenging problem to which a large number of search strategies have been applied. The model depends nonlinearly on the source locations and linearly on the orientation and time series. The problem can therefore be solved in two steps: the linear parameters can be expressed as an implicit function of the nonlinear parameters, so that we can solve first for the nonlinear parameters (Mosher et al., 1992). Direct approaches to solving for location parameters include use of gradient descent or the Nelder–Meade simplex algorithm (Huang et al., 1998). These methods do not deal directly with the problem of multiple local minima. Methods to cope with this issue include the use of multistart, genetic algorithms and simulated annealing (Huang et al., 1998; Khosla et al., 1997; Uutela et al., 1998).

Another approach to dealing with the existence of local minima is to use a scanning method. These methods use a discrete grid to search for optimal dipole positions throughout the source volume. Source locations are then determined as those for which a metric computed at that location exceeds a given threshold. While these approaches do not lead to true least squares solutions, they can be used to initialize a local least squares search. It is important to note that these scanning methods are explicitly based on essentially the same model as least squares dipole fitting, that is, that the data can be explained by a few equivalent low-rank or dipolar sources. This fact is sometimes overlooked because the outputs of the scanning methods are often viewed as an "image" as illustrated in Fig. 3.

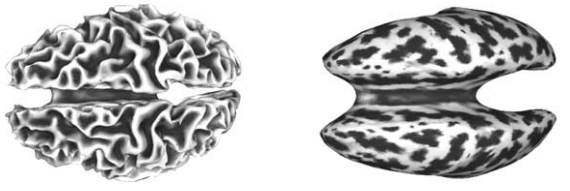

Fig. 2. Example of a cortical surface and its smooth representation for easy visualization of cortically constrained activation maps.

The most common scanning methods are the linearly constrained minimum variance (LCMV) beamformer and multiple signal classification (MUSIC) and their variants. The LCMV beamformer was developed in the array signal processing community as a means for combining the outputs of multiple antennae to produce a single output with the characteristics of a single larger antenna. By analogy, some researchers view the LCMV beamformer in MEG and EEG as a "virtual depth electrode", because at each time sample it forms a linear combination of the external field measurements to monitor a single point in the brain while minimizing contributions to the beamformer output from all other sources. By scanning the location of the depth electrode throughout the brain, we can monitor all possible source locations in the brain. The limitation of this approach is that the process of minimization of sensitivity to other sources is data adaptive: sources that are correlated with the true signal from the scanning location can cause cancellation of the signal of interest. Thus, the beamformers will work best when there are a limited number of independent components in the data that are not strongly correlated with each other.

The LCMV beaformer is formulated mathematically as minimizing the output power of a linear filter under the constraint that its gain is unity at $\boldsymbol{r}_s$, the location of interest, i.e., $\min_{\mathbf{W}}\ tr(\boldsymbol{W}^T\boldsymbol{C}_m\boldsymbol{W})$ subject to $\boldsymbol{W}^T\boldsymbol{G}(\boldsymbol{r}_s)=\boldsymbol{I}$ and $\boldsymbol{C}_m=E(\boldsymbol{m}(t)\boldsymbol{m}(t)^T)$; where *tr* denotes the trace of a matrix. The solution for the matrix $\boldsymbol{W}$ of filter weights is given by (van Veen and Buckley, 1988)

$$\boldsymbol{W}^T(\boldsymbol{r}_s)=\left[\boldsymbol{G}(\boldsymbol{r}_s)^T\boldsymbol{C}_m^{-1}\boldsymbol{G}(\boldsymbol{r}_s)\right]^{-1}\boldsymbol{G}(\boldsymbol{r}_s)^T\boldsymbol{C}_m^{-1} \tag{9}$$

By recomputing this solution for each point $\boldsymbol{r}_s$ on the brain surface or volume and computing the output as $\boldsymbol{S}=\boldsymbol{W}^T\boldsymbol{M}$, we obtain a time series at each scan location. To obtain a scalar output, we compute the average signal power over time. For low noise data, $\boldsymbol{C}_m^{-1}$ may be ill conditioned and needs to be regularized, that is, $\boldsymbol{C}_m^{-1}$ is replaced by $(\boldsymbol{C}_m+\lambda\boldsymbol{I})^{-1}$ (Robinson and Vrba, 1999).

The concept of MUSIC and recursively applied (RAP) MUSIC (Mosher and Leahy, 1998, 1999; Mosher et al., 1992) is based on finding those dipolar topographies, (columns of G corresponding to a particular dipole) that project into an estimated signal subspace. The p-dimensional signal subspace is defined as corresponding to that part of the column space of $\boldsymbol{M}$ whose corresponding singular values $\lambda_1,\lambda_2,\ldots,\lambda_p$ lie above a noise floor. Writing the singular value decomposition (SVD) of the data matrix as $\boldsymbol{M}=\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}^T$, the first p left singular vectors of $\boldsymbol{M}$ define the signal subspace $\boldsymbol{U}_p$, i.e., $\boldsymbol{U}_p$ is formed from the first p columns of $\boldsymbol{U}$. The cosine of the smallest principal angle:

$$\mathrm{subcorr}\left(\boldsymbol{U}_p,\boldsymbol{G}(\boldsymbol{r}_s)\right)_1=\max_{\boldsymbol{q}}\sqrt{\frac{\boldsymbol{q}^T\boldsymbol{G}(\boldsymbol{r}_s)^T\boldsymbol{U}_p^T\boldsymbol{U}_p\boldsymbol{G}(\boldsymbol{r}_s)\boldsymbol{q}}{\boldsymbol{q}^T\boldsymbol{G}(\boldsymbol{r}_s)^T\boldsymbol{G}(\boldsymbol{r}_s)\boldsymbol{q}}} \tag{10}$$

is used as the scanning metric. Source locations are found as those for which $\mathrm{subcorr}(\boldsymbol{U}_p,\boldsymbol{G}(\boldsymbol{r}_s))_1$ is sufficiently close to unity given the SNR in the data. In RAP-MUSIC, the dipolar topographies with maximum subspace correlation are determined recursively. In each recursion, the signal subspace is projected away from the topographies already found, thus avoiding the search for local maxima of the correlation metric. The algorithm terminates if no source with a correlation above the user-specified threshold can be found. As with the LCMV beamformer, the method assumes discrete sources. Provided the sources are not perfectly synchronous, MUSIC does not suffer from the signal cancellation problems of LMCV, but the dependence of the method on correlation measures can make it less robust to modeling errors.

Imaging methods

Imaging methods solve the linear inverse problem in Eq. (7) for the current density $\boldsymbol{S}^T$, usually constrained to a tessellated representation of the cerebral cortex. Typically, the number p ($\approx 10^4$) of sources in the discrete grid is much larger than the number of detectors c ($\approx 10^2$), so that the problem is highly underdetermined. Also, the singular value spectrum of $\mathbf{A}$ decays approximately exponentially, making the problem ill-posed. This ill-posedness is a principal property of the bioelectromagnetic inverse problem: the current distribution inside a conductor cannot be uniquely determined by knowledge of the fields and potentials outside (von Helmholtz, 1853).

To compute a unique solution we must introduce some form of regularization or prior on the solution. While not all inverse methods are explicitly formulated in this way, it is useful to view imaging methods in a common Bayesian perspective.

Given the conditional distribution $p(\boldsymbol{M}_0|\boldsymbol{S})$ of the noiseless measurements $\boldsymbol{M}_0$ with respect to the sources $\boldsymbol{S}$, $p(\boldsymbol{M}|\boldsymbol{M}_0)$ the probability of observing the measured data $\boldsymbol{M}$ given the noiseless data $\boldsymbol{M}_0$, and $p(\boldsymbol{S})$ the prior source distribution, the posterior probability for the sources is given by

$$p(\boldsymbol{S}|\boldsymbol{M})=c_0p(\boldsymbol{S})\int p(\boldsymbol{M}|\boldsymbol{M}_0)p(\boldsymbol{M}_0|\boldsymbol{S})d\boldsymbol{M}_0 \tag{11}$$

where c_0 is a normalization constant (Tarantola, 1987). The forward model is specified in $p(\boldsymbol{M}_0|\boldsymbol{S})$ and will be simply $\delta(\boldsymbol{M}_0-\boldsymbol{A}\boldsymbol{S}^T)$ in the case where there is no uncertainty in the forward calculation. Uncertainty can be incorporated in this model if the conductivities are unknown by treating them as random variables with a known statistical distribution. The measurement noise model is described by $p(\boldsymbol{M}|\boldsymbol{M}_0)$ and all prior information about $\boldsymbol{S}$ is encoded in $p(\boldsymbol{S})$. With the known posterior distribution, many statistical properties of $\boldsymbol{S}$ could be computed using Markov Chain Monte Carlo (MCMC) methods. In practice, most inverse methods are restricted to finding the $\boldsymbol{S}$ that maximizes the posterior density. For the case where there is no forward model uncertainty, measurement noise is a white Gaussian process, and the current density is assumed zero mean Gaussian with covariance $\boldsymbol{C}_s$, the maximum a posteriori (MAP) estimate is the minimizer of

$$\|\boldsymbol{M}-\boldsymbol{A}\boldsymbol{S}^T\|_2^2+\mathrm{tr}\left(\boldsymbol{S}\boldsymbol{C}_s^{-1}\boldsymbol{S}^T\right) \tag{12}$$

For the case $\boldsymbol{C}_s=\lambda^{-1}\boldsymbol{I}$, this is the Tikhonov regularized version of the inverse problem (Tikhonov and Arsenin, 1977) with solution:

$$\boldsymbol{S}^T=\left(\boldsymbol{A}^T\boldsymbol{A}+\lambda\boldsymbol{I}\right)^{-1}\boldsymbol{A}^T\boldsymbol{M} \tag{13}$$

Different choices of the prior distribution yield other commonly used linear methods such as column weighted minimum norm, which is designed to reduce the preference for superficial cortical sources. In this case, $\boldsymbol{C}_s=\lambda^{-1}\boldsymbol{W}$ with $(\boldsymbol{W})_{ii}=\|\boldsymbol{a}_i\|^2$, where $\boldsymbol{a}_i$ is the ith column of $\boldsymbol{A}$. In low-resolution electromagnetic tomography (LORETA) (Pascual-Marqui et al., 1994), which uses Laplacian

weighting to regularize the solution, $\boldsymbol{C}_s$ is defined by $\boldsymbol{C}_s^{-1} = \lambda \boldsymbol{K}\boldsymbol{K}^{\mathrm{T}}$ and

$$(\boldsymbol{K})_{ij} = \begin{cases} 1 & i = j \\ -\frac{1}{n} & j \in \mathcal{N}(i) \\ 0 & \text{else} \end{cases}$$

with $\mathcal{N}(i)$ the set of nearest neighbors of the source location i on the discrete grid and n the cardinal number of $\mathcal{N}(i)$.

The minimum norm approach can be extended to include nonlinear procedures by using the L_p norm, $1 \leq p \leq 2$ (Scales and Gersztenkorn, 1988). Among the L_p norm solutions, the L_1 norm has particularly interesting properties as it produces sparse solutions and can be formulated as a linear programming problem. The properties of linear programs guarantee that there exists an optimal solution for which the number of nonzero values does not exceed the number of measurements. This sparseness is an attractive property for studies where a relatively small fraction of cortex is expected to be activated. The idea of sparse cortical activation can also explicitly be built into the prior. Markov random fields (RFs) provide a rich framework for building prior distributions on spatial processes that can incorporate sparseness, focal activation, and statistical characterization of the number and size of activated regions (Baillet and Garnero, 1997; Phillips et al., 1997). These high dimensional nonlinear methods are far more computationally demanding than the linear minimum norm approaches and often include multiple hyperparameters on the prior distributions. However, they are better able to reflect the expected characteristics of brain activation than linear minimum norm images that typically exhibit very low resolution characteristics as illustrated in Fig. 3.

Choice of inverse method

The choice of inverse method critically determines the manner in which users interpret their data. Dipole methods can yield quite precise and accurate results in the case of highly focal activation, for example, in somatosensory stimulation (Buchner et al., 1995) or in analysis of epileptic brain activity (Barkley and Baumgartner, 2003). However, in application to cognitive experiments, where the number of active regions in the brain cannot be predicted and large areas of the brain may be involved in the response, dipole models can perform poorly and imaging methods might be more suitable. Linear imaging methods produce solutions that often show activity over large portions of the brain surface. This is due to the low resolution that results from mapping $c \approx 10^2$ detectors onto $p \approx 10^4$ nodes on the cortical surface. Consequently, the imaging approach lacks the precision of a multiple dipole fit, but avoids the potentially overrestrictive nature of the dipole model. Imaging solutions can be made more precise by using a more restrictive prior. However, increased precision reflects increased accuracy only so far as the prior reflects the true complexity of brain activation.

With this broad array of inverse procedures, selection of the appropriate method is a difficult task. One suggestion we can offer is that dipole methods be used in cases where a few stationary focal sources are expected, and that imaging methods be used to study more complex activation. However, further studies, both in simulation and with experimental data, are required to better understand the properties and relative merits of these approaches. Such evaluation can use object task-based evaluation criteria, estimation of confidence intervals, and detection of significant activation, as described in the following section. It is also important to note that the ability of MEG and EEG to resolve brain activity is ultimately limited by the underlying physics, which tells us that we cannot uniquely determine the current field within the brain from external measurements. Therefore, all inverse methods will provide different approximations of this unknown quantity, with the approximations reflecting the underlying assumptions implicit or explicit in each method.

Validation and statistical analysis

The MEG and EEG literature encompasses a wide variety of reconstruction and localization methodologies. It is important to evaluate the relative performance of these methods under different experimental settings such as the number, location,

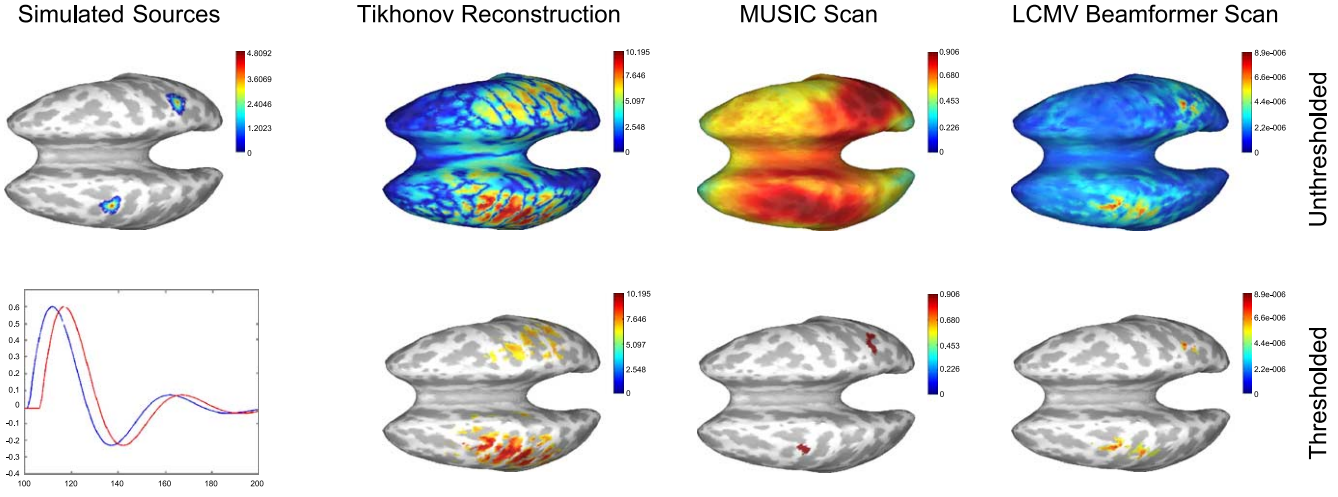

Fig. 3. Localization maps for two simulated sources in the left motor area and the right occipital lobe; their time series are shown in the lower left panel. The Tikhonov regularized density map is given at time $t = 112$, whereas MUSIC and LCMV beamformer methods use all post-stimulus time points to obtain the scalar output of principal angle and signal power at each scan location on the cortex, respectively. The second row shows subjectively thresholded activation maps to demonstrate localization ability of these methods.

and time series of neural sources. Furthermore, several methods require fine tuning of parameters, for example, the Tikhonov regularization parameter in current density reconstruction, or the subspace correlation threshold for the MUSIC algorithm. Adjusting these parameters affects the sensitivity and specificity of each method. Here we illustrate the use of the receiver operating characteristic (ROC) curve to study the trade-off between sensitivity and specificity and to compare different inverse methods.

In addition to evaluating the relative performance of different methods, it is also important to establish some degree of confidence in the results of real data analysis. Dipole scanning methods often produce unstable solutions and the reproducibility of the reconstructed dipoles is not guaranteed. Consequently, it is important to quantify the accuracy of the dipole locations and time series. Conversely, imaging methods are hugely underdetermined, resulting in low-resolution localization maps; interpretation is further confounded by the presence of additive noise exhibiting a highly nonuniform spatial correlation. In this case, we need a mechanism to decide which features in the data are indicative of true activation vs. those that are noise artifacts. While parametric models can be assumed as the basis for developing statistical tests, nonparametric resampling methods such as the bootstrap and permutation methods are well suited to the analysis of MEG and EEG data. Event-related studies typically involve multiple repetitions that can be treated as a set of identically distributed and independent realizations from which we can resample. Here we describe the application of resampling methods and parametric random field methods to the analysis of MEG and EEG inverse solutions.

Task-based evaluation of inverse procedures

The ROC curve was developed during World War II for quantifying performance in target detection using radar and sonar. ROC analysis was first applied for the assessment of medical imagery in the 1970s (Metz et al., 1973). Several variants of ROC analysis have been proposed, but all produce a plot of the variation in true positive fraction (TPF) vs. false positive fraction (FPF) as a function of some decision threshold; they differ in their definitions of these quantities. Because ROC methods require knowledge of the ground truth, they are typically restricted to simulation studies or cases where truth can be established using an independent technique, such as invasive measurements in the case of MEG and EEG; however, recently there have been attempts to apply ROC methodologies in cases where the ground truth is not known (Nandy and Cordes, 2003).

Standard ROC analysis (Metz et al., 1973) is used in diagnostic studies where we are interested in discriminating between the presence and absence of a lesion or target and the location of the lesion is not of concern (Swets and Pickett, 1982). In location-response ROC (LROC) studies the observer must also specify the location of the target in each image (Starr et al., 1975). LROC is closer than ROC to a realistic medical image interpretation task; however, the observer is allowed to identify only one target per image, even if many targets are present. Bunch et al. (1978) proposed the more realistic free-response ROC (FROC) method, which allows detection of multiple targets per image. Localization within an acceptable distance of a lesion is accepted as a true event and all other detections are classified as false. The FROC curves generated from the detection data show the probability that a true lesion is detected vs. the expected value of the number of false-positive detections per image.

In neuroimaging studies, we often expect to see several simultaneously activated brain areas. FROC is therefore an appropriate framework for assessing the specificity and sensitivity of MEG or EEG inverse methods. Here we use a computer observer model, where cortical maps of neural activation are reconstructed from simulated data and thresholded to determine active and inactive cortical regions. By varying the threshold and comparing results to the ground truth, we produce FROC curves for each inverse method.

Before giving an example of this procedure, we first give some definitions. A voxel is the basic element in the cortical map at which we test for activation. An actual positive is a voxel in which a signal is truly present, or is within a predefined acceptable distance of an activated voxel. An actual blob is a contiguous set of actual positives. An image may contain many actual blobs and the number of actual voxels included in each blob changes according to the size of the activated region. We apply a global threshold to the cortical map and voxels exceeding this threshold are defined as active; a group of adjacent active voxels forms an active blob. Negatives are defined analogously. To apply FROC in MEG or EEG studies, we assume a true positive occurs if there is at least one active voxel that intersects with an actual blob, and the true positive fraction is defined as the ratio of true positives to the total number of actual blobs. Conversely, an active voxel that is an actual negative is treated as a false positive and the points on the abscissa of the curve for each threshold represent the ratio of false positives to actual negatives per image. Because the number of true positives is assured not to decrease as the average number of false positives increases, the resulting FROC curve will be nondecreasing.

To illustrate the FROC methodology, we simulated 300 MEG data sets for a CTF Systems Inc. Omega 151 system, each consisting of an array of data representing the measured magnetic field at each sensor as a function of time. The first 100 of the data sets consisted of measurements with no activation on the cortical surface and white noise only at the sensors; the remaining 200 data sets represented two temporally uncorrelated activated patches randomly positioned on the cortical surface. We used three methods to produce detection maps: Tikhonov regularized minimum norm reconstruction (Eq. (13)), MUSIC (Eq. (10)), and LCMV beamformer (Eq. (9)). Before thresholding these maps, we first normalized the minimum norm and LCMV maps to produce a statistic that gives uniform spatial specificity in the absence of activation. This normalization involves division by the expected output in the absence of activation. For the minimum norm images, we compute the index $d(\boldsymbol{r}_s)$ as

$$d(\boldsymbol{r}_s) = E_t\left\{\tilde{\boldsymbol{s}}^2_{\boldsymbol{r}_s}(t)\right\}, \quad \boldsymbol{S}^T = \boldsymbol{\Sigma}\left(\boldsymbol{A}^T\boldsymbol{A} + \lambda\boldsymbol{I}\right)^{-1}\boldsymbol{A}^T\boldsymbol{M} \tag{14}$$

where $\tilde{\boldsymbol{s}}_{\boldsymbol{r}_s}(t)$ represents the reconstructed time series computed at location $\boldsymbol{r}_s$ at time t, and $\boldsymbol{\Sigma}$ is a diagonal matrix equal to the inverse of the noise sensitivity: $\boldsymbol{\Sigma} = (\mathrm{diag}\{\boldsymbol{H}\boldsymbol{C}_n\boldsymbol{H}^T\})^{-1/2}$ with $\boldsymbol{C}_n$ the estimated sensor noise covariance (Dale et al., 2000) and $\boldsymbol{H} = (\boldsymbol{A}^T\boldsymbol{A} + \lambda\boldsymbol{I})^{-1}\boldsymbol{A}^T$. Similarly, for the LCMV beamformer, we compute a "neural activity index" (van Veen et al., 1997) as:

$$d(\boldsymbol{r}_s) = \frac{\mathrm{tr}\left\{\boldsymbol{W}^T(\boldsymbol{r}_s)\boldsymbol{C}_m\boldsymbol{W}(\boldsymbol{r}_s)\right\}}{\mathrm{tr}\left\{\boldsymbol{W}^T(\boldsymbol{r}_s)\boldsymbol{C}_n\boldsymbol{W}(\boldsymbol{r}_s)\right\}} = \frac{\mathrm{tr}\left\{\left[\boldsymbol{G}^T(\boldsymbol{r}_s)\boldsymbol{C}_m^{-1}\boldsymbol{G}(\boldsymbol{r}_s)\right]^{-1}\right\}}{\mathrm{tr}\left\{\left[\boldsymbol{G}^T(\boldsymbol{r}_s)\boldsymbol{C}_n^{-1}\boldsymbol{G}(\boldsymbol{r}_s)\right]^{-1}\right\}} \tag{15}$$

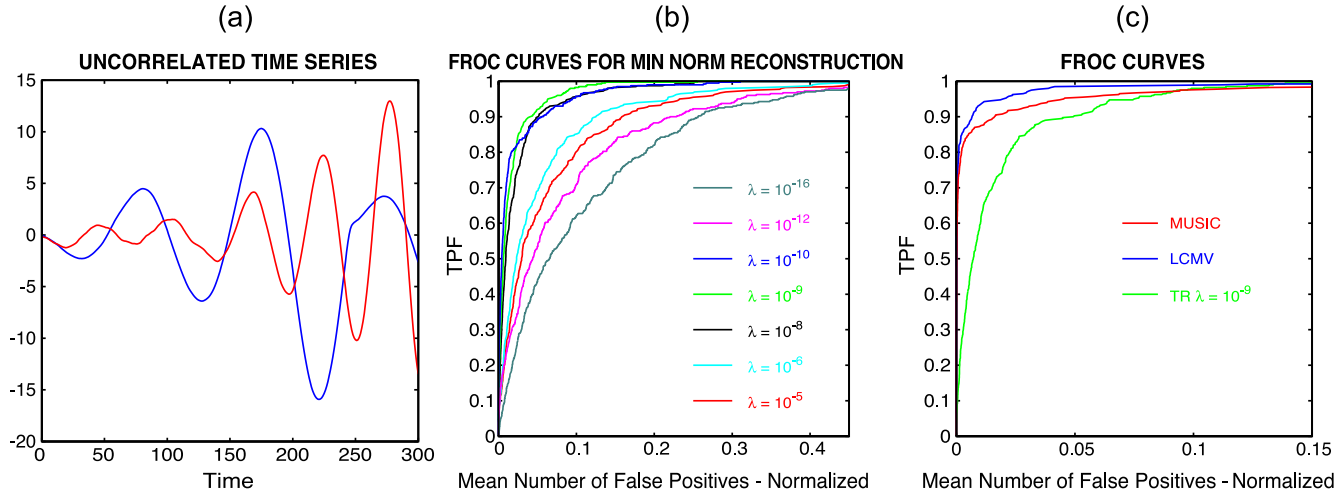

Fig. 4. (a) Time series of the two temporally uncorrelated activated patches randomly positioned on the cortical surface, (b) FROC curves for different regularization parameters (λ) for noise-normalized Tikhonov-regularized minimum norm reconstruction; (c) FROC curves for MUSIC, LCMV Beamformer, and the minimum norm approach (with $\lambda = 10^{-9}$). FROC curves in b and c are plotted for small values of false positives per image to illustrate differences between the methods.

where $\boldsymbol{C}_n$ is the noise-only covariance. In the case of MUSIC, the method precludes this form of normalization, and we use the index:

$$d(\boldsymbol{r}_s) = \text{subcorr}\left(\boldsymbol{U}_p, \boldsymbol{G}(\boldsymbol{r}_s)\right)_1 \tag{16}$$

Fig. 4a shows the time courses used in our simulations. Using the FROC procedure, we determined the optimal Tikhonov regularization parameter λ for minimum norm reconstruction, $\lambda = 10^{-9}$, as that which maximized the area under the FROC curve (Fig. 4b). Finally, Fig. 4c compares FROC curves for the three different methods and shows that, for this task, the LCMV beamformer has a larger area under the FROC curve than MUSIC or minimum norm. At matched sensitivity, LCMV has higher specificity than MUSIC and minimum norm. More generally, we have illustrated how the ROC methodology can be used to objectively compare the performance of different inverse methods.

The bootstrap and confidence regions for dipole localization

We turn now to the assessment of uncertainty in sources estimated from experimental data in which the ground truth is not known. A number of different approaches have been investigated for assessing dipole localization accuracy, including Cramer Rao lower bounds, perturbation analysis, and Monte Carlo simulation; most make assumptions about the underlying distribution of the data, as reviewed in Braun et al. (1997) and Darvas et al. (in press). Here we restrict our attention to the nonparametric bootstrap (Efron and Tibshirani, 1986), a method for assessing accuracy of an estimator by sampling with replacement from a set of independent trials. By repeating this process and estimating parameters from each bootstrap resample, we can learn the approximate distribution of the estimator. The principle underlying the bootstrap is that although the distribution of the data is unknown, it can be approximated by the empirical distribution of a set of independent trials. To achieve an acceptable signal-to-noise ratio (SNR) in an event-related MEG or EEG study, a large number (typically $\approx$ 100) of repeated trials or epochs are recorded. The average response is then computed from these trials through stimulus-locked averaging. The raw trials can be viewed as independent realizations of the brain's response to a particular event, and hence the bootstrap methodology can be directly applied by sampling with replacement from the epochs, and averaging the resulting samples (DiNocera and Ferlazzo, 2000). The advantage of this bootstrap approach is that no specific assumptions are made regarding the distribution of the noise, the dipole time series, or the number of dipoles.

As an illustration, we have applied the bootstrap to MUSIC-based source localization from MEG data from somatosensory stimulation of left and right fingers. Clustering the resulting source locations and computing the principal eigenvectors of the sample covariance of each cluster, we can determine a confidence ellipsoid around the locations of each estimated source. The clusters for four digits of each hand show the standard somatotopic mapping in primary somatosensory cortex (SI) and the degree of uncertainty for each digit as illustrated in Fig. 5. The potential advantage of this approach over methods based on analysis of residual error or prestimulus data is that the bootstrap is able to include all sources of variability in the data, including statistical variability of the brain response itself, and hence should lead to more accurate quantification of localization accuracy uncertainty or confidence interval estimates.

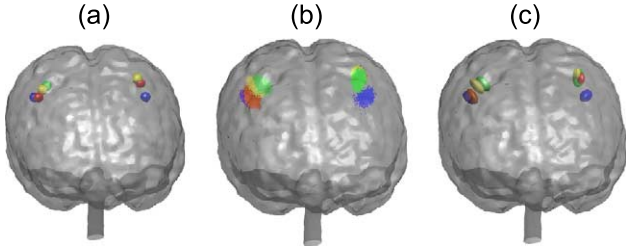

Fig. 5. Bootstrap results from an MEG somatosensory study involving delivery of electrical pulses to the fingers of both hands of a healthy subject; (a) locations from original data of the strongest dipole source for each of the four digits demonstrating somatotopic mapping in the sensory cortex; (b) scatter distribution of bootstrapped dipoles, color coded for each digit; (c) confidence ellipsoids constructed with principal axes along the eigenvectors of the cluster covariances; axes of length equal to two times the standard deviation along each principal axis represent an 87% confidence region for a Gaussian distribution. (From Darvas et al. (in press).)

Detection of regions of significant activation in cortical maps

In contrast to dipole localization, where the number of spatial parameters is far fewer than the number of detectors, cortically constrained maps typically contain far more surface elements than detectors. This overparameterization leads to high spatial correlation in the maps and presents difficulties in determining a suitable threshold for detecting statistically significant activation. Applying a simple Bonferroni correction and testing at each surface element will lead to a very conservative threshold. Conversely, an unacceptably high false-positive rate may arise if no correction is made for multiple hypothesis testing. The standard approach to this problem is to control the familywise error rate (FWER), that is, the chance of any false positives under the null hypothesis (type 1 error). Parametric random field methods and nonparametric permutation methods address this problem by estimating family-wise-corrected thresholds. Here we review these approaches as they might be applied to the detection maps in Eqs. (14–16).

Parametric statistical methods are well established in fMRI and PET neuroimaging studies, where image maps are analyzed using massively univariate hypothesis tests against the null hypothesis of no activation. The theory of random fields (RFs) is used to approximate the upper tail of the maximum distribution of the image statistics using the expected value of the Euler characteristic of the thresholded image (Worsley et al., 1996). Adapting the theory to MEG or EEG inverse problems is complicated by their highly nonuniform spatial correlation structure. A thorough treatment of such fields is given by Worsley et al. (1999).

Barnes and Hillbrand (2003) present an application of RF theory to MEG data. Their method is specifically tailored to LCMV beamforming solutions, and the statistic images are formed as maps of source power change between an active and a passive state in a specific time–frequency interval. In Pantazis et al. (in press), we apply RF results to the noise normalized minimum-norm reconstruction maps defined in Eq. (14). Expressions for the RESEL count and the Euler characteristic density are based on the prestimulus cortical maps and are used to derive an analytical solution for the maximum distribution of the image statistics. A transformation similar to Worsley et al. (1999) is used to compensate for the heterogeneity of the RFs. The application of statistical parametric mapping to low resolution electromagnetic tomography (LORETA) image maps computed from event-related potentials is treated by Park et al. (2002). They construct volumetric current density maps, apply three-dimensional Gaussian smoothing, and analyze the data using the traditional Gaussian RF theory.

The permutation method (Nichols and Holmes, 2001) is a nonparametric approach to controlling the FWER. It is restricted to data that satisfy an exchangeability criterion, that is, the requirement that under a null hypothesis H_0, the distribution of the statistic of interest remains unaltered when the labels of the data are permuted. In the case of MEG or EEG data, we can permute pre- and post-stimulus data under the H_0 that there is no event-related activation. With permutation tests, we have the flexibility of using any statistic without the need to know its distribution. We then calculate this statistic for all permutation samples and the resulting empirical distribution is used to define familywise-corrected thresholds.

Blair and Karnisky (1994) were the first to use permutation tests in an EEG study. Permutation samples were created by exchanging frequent and rare channel waveforms from a standard oddball paradigm, and permutation statistics obtained by subtracting the time courses between these two experimental conditions and averaging to form a single time course. The empirical distribution of the maximum statistic over time was used to extract a global threshold and temporally localize significant activity. Permutation tests can be extended to provide combined spatial and temporal localization of brain activation. Pantazis et al. (in press) used a permutation scheme where maps of cortical activation are permuted. The goal was to localize those surface elements in space and time that exhibit significant experimental effects. By randomly permuting pre- and post-stimulus data from the collection of individual epochs in an event-related study, we learn the maximum distribution of the statistic and use this to control the FWER. This approach implicitly accounts for both spatial and temporal correlation in the cortical maps. An alternative permutation scheme proposed by Singh et al. (2003) detects event-related synchronization or desynchronization components in a MEG study involving visual stimulation and an LCMV beamformer applied to data decomposed in multiple frequency bands.

We close by illustrating RF and permutation methods applied to a simulation of sources in the right and left somatosensory areas. Shown in Fig. 6 are examples of significant activation maps as derived in Pantazis et al. (in press). By comparing Figs. 6a and b, we see that spatiotemporal thresholds are more stringent than spatial thresholds, because the false positives are controlled over all time slices. Further, spatial smoothing increases SNR at the

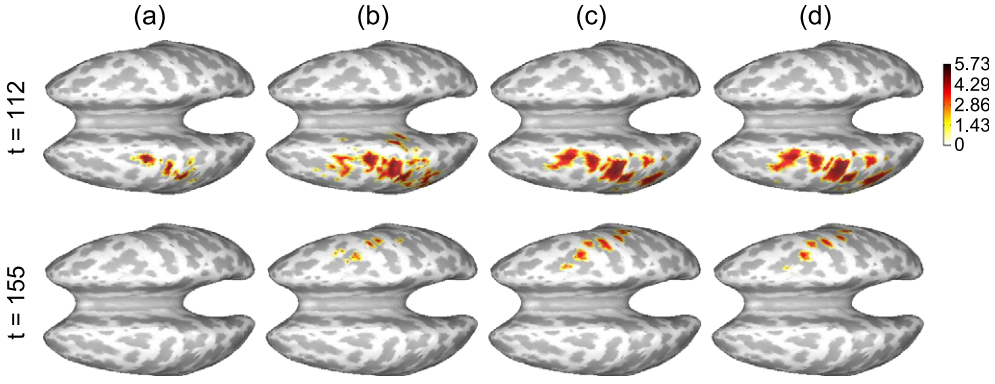

Fig. 6. Examples of significant activation maps for permutation and random field methods for two simulated sources on the right and left somatosensory area; (a) permutation method controlling spatiotemporal FWER and using unsmoothed current densities, (b) permutation method controlling spatial FWER and using unsmoothed current densities, (c) permutation method controlling spatial FWER and using smoothed current densities, (d) RF method controlling spatial FWER and using smoothed current densities. Smoothing is optional for permutation tests, but necessary for RFs to avoid conservative thresholds.

expense of lower resolution (Figs. 6b and c). In simulation studies, permutation and random field methods perform similarly with smoothed cortical activation maps (Figs. 6c and d).

We should comment here that permutation and RF tests do not address the limited resolution of MEG reconstruction methods. The MEG inverse problem is ill-posed and cortical activation maps are of low resolution and tend to mislocalize source activation. If the inverse method identifies experimental variation in some region, permutation and random field tests will identify these regions regardless of the presence of an actual source at those locations. The activation maps in Fig. 6 show the signal leakage onto neighboring gyri that is typical in MEG imaging studies.

Concluding remarks

We have provided a brief survey of the issues in solving the forward and inverse problem in EEG and MEG. An open source Matlab toolbox that implements many of the methods discussed can be found at http://neuroimage.usc.edu/brainstorm/ (Baillet et al., 2004). We have also discussed methods for comparing different inverse procedures and for assessing statistical significance of inverse results from experimental data. In this short article, there are many important aspects and recent developments in this field that we are not able to address. We believe that of particular importance among these are the following: (i) the use of independent components analysis as a means of separating physiological and other noise processes from brain activation, and possibly for also separating distinct components of neural activation; (ii) the use of time–frequency analysis and coherence analysis applied to single trial data to better elucidate the mechanisms underlying communication between neuronal assemblies; and (iii) the development of methods for simultaneous acquisition and analysis of EEG and MEG data with other modalities, principally fMRI, which have the potential to combine the high temporal resolution of electrophysiological data with the higher spatial resolution of hemodynamic effects measured with fMRI (Dale et al., 2000). The analysis of MEG and EEG data remains a rich research area, which together with the other brain imaging modalities, will lead to an increased understanding of the cortical specialization and neuronal networks that underlie human brain function.

Acknowledgments

This work was supported by NIBIB under Grant R01 EB002010. The authors thank Sabine Meunier of the Physiology and Physiopathology of Human motricity Lab, La Salpetriere hospital, INSERM, Paris, for providing the somatosensory MEG data.

References

Awada, K., Jackson, D., Williams, J., Wilton, D., Baumann, S., Papanicolaiu, A., 1997. Computational aspects of finite element modeling in eeg source localization. IEEE Trans. Biomed. Eng. 44, 736–752.

Baillet, S., Garnero, L., 1997. A Bayesian approach to introducing anatomofunctional priors in the EEG/MEG inverse problem. IEEE Trans. Biomed. Eng. 44, 374–385.

Baillet, S., Mosher, J., Leahy, R., 2001. Electromagnetic brain mapping. IEEE Signal Process. Mag. 18 (6), 14–30.

Baillet, S., Mosher, J.C., Leahy, R.M., 2004. Electromagnetic brain imaging using Brainstorm. Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), Arlington, VA, April 2004.

Barkley, G., Baumgartner, C., 2003. MEG and EEG in epilepsy. J. Clin. Neurophysiol. 20 (3), 163–178.

Barnes, G., Hillbrand, A., 2003. Statistical flattening of MEG beamformer images. Hum. Brain Mapp. 18, 1–12.

Blair, R., Karnisky, W., 1994. Distribution-free statistical analyses of surface and volumetric maps. In: Thatcher, R.W., Hallett, M., Zeffiro, T., Jony, E.R., Huerta, M. (Eds.), San Diego: Academic Press, Functional Neuroimaging: Technical Foundations, pp. 19–28.

Braun, C., Kaiser, S., Kinces, W., Elbert, T., 1997. Confidence interval of single dipole locations based on eeg data. Brain Topogr. 10 (1), 31–39.

Buchner, H., Kauert, C., Rademacher, I., 1995. Short term changes of finger representation at the somatosensory cortex in humans. Neurosci. Lett. 198, 57–59.

Buchner, H., Waberski, T., Fuchs, M., Wischmann, H., Wagner, M., Drenckhahn, R., 1996. Comparison of realistically shaped boundary element and spherical head models in source localization of early somatosensory evoked potentials. Brain Topogr. 8, 137–143.

Bunch, P., Hamilton, J., Sanderson, G., Simmons, A., 1978. A free response approach to the measurement and characterization of radiographic-observer performance. J. Appl. Photogr. Eng. 4, 166–172.

Dale, A., Liu, A., Fischl, B., Buckner, R., Belliveau, J., Lewine, J., Halgren, E., 2000. Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. Neuron 26, 55–67.

Darvas, F., Rautianinen, M., Pantazis, D., Baillet, S., Benali, H., Mosher, J., Garnero, L., Leahy, R., 2004. Investigations of dipole localization accuracy in MEG using the Bootstrap. NeuroImage (in press).

DiNocera, F., Ferlazzo, F., 2000. Resampling approach to statistical inference: bootstrapping from event-related potentials data. Behav. Res. Methods Instrum. Comput. 32 (1), 111–119.

Efron, B., Tibshirani, R.J., 1986. Bootstrap methods for standard errors, confidence intervals and other measures of statistical accuracy. Stat. Sci. 1, 54–77.

Ermer, J., Mosher, J., Baillet, S., Leahy, R., 2001. Rapidly recomputable EEG forward models for realistic headshapes. Phys. Med. Biol. 46 (4), 1265–1281.

Fuchs, M., Wagner, M., Kastner, J., 2001. Boundary element method volume conductor models for EEG source reconstruction. Clin. Neurophysiol. 112 (8), 1400–1407.

Geddes, L., Baker, L., 1967. The specific resistance of biological materials—A compendium of data for the biomedical engineer and physiologist. Med. Biol. Eng. 5, 271–293.

Geselowitz, D., 1967. On bioelectric potentials in an inhomogeneous volume conductor. Biophys. J. 7, 1–11.

Goebel, R., 1997. Brainvoyager: Ein programm zur analyse und visualisierung von magnetresonanztomographiedaten. In: Plesser, T., Wittenburg, P. (Eds.), Forschung und Wissenschaftliches Rechnen.

Goncalves, S., deMunck, J., Heethar, R., LopesDaSilva, F., vanDijk, B., 2000. The application of electrical impedance tomography to reduce systematic errors in the eeg inverse problem—A simulation study. Physiol. Meas. 21, 379–393.

Hamalainen, M., Hari, R., Ilmoniemi, R.J., Knuutila, J., Lounasmaa, O.V., 1993. Magnetoencephalography—Theory, instrumentation, and applications to noninvasive studies of the working human brain. Rev. Mod. Phys. 65 (2), 413–498.

Haueisen, J., Tuch, D., Ramon, C., Schimpf, P., Weeden, V., George, J., Belliveau, J., 2002. The influence of brain tissue anisotropy on human EEG and MEG. NeuroImage 15 (1), 159–166.

Huang, M., Aine, C., Supek, S., Best, E., Ranken, D., Flynn, F., 1998. Multi-start downhill simplex method for spatio-temporal source localization in magnetoencephalography. Electroencephalogr. Clin. Neurophysiol. 108 (1), 32–44.

Huang, M., Mosher, J., Leahy, R., 1999. A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG. Phys. Med. Biol. 44, 423–440.

Huiskamp, G., Vreijenstijn, M., Dijk, R., Wieneke, G., van Huffelen, A., 1999. The need for correct realistic geometry in the inverse eeg problem. IEEE Trans. Biomed. Eng. 46, 1281–1287.

Johnson, C., 1997. Computational and numerical methods for bioelectric field problems. Crit. Rev. Biomed. Eng. 25, 1–103.

Khosla, D., Singh, M., Don, M., 1997. Spatio-temporal EEG source localization using simulated annealing. IEEE Trans. Biomed. Eng. 44, 716–723.

Leahy, R., Mosher, J., Spencer, M., Huang, M., Lewine, J., 1998. A study of dipole localization accuracy for MEG and EEG using a human skull phantom: electroencephalography. Clin. Neurophysiol. 107 (2), 159–173.

Metz, C., Goodenough, D., Rossmann, K., 1973. Evaluation of receiver operating characteristic curve data in terms of information theory, with applications in radiography. Radiology 109 (2), 297–303.

Mosher, J.C., Leahy, R.M., 1998. Recursive MUSIC: a framework for EEG and MEG source localization. IEEE Trans. Biomed. Eng. 45 (11), 1342–1355.

Mosher, J.C., Leahy, R.M., 1999. Source localization using recursively applied and projected (RAP) music. IEEE Trans. Signal Process. 47 (2), 332–340.

Mosher, J., Leahy, R., Lewis, P., 1992. Multiple dipole modeling and localization from spatiotemporal meg data. IEEE Trans. Biomed. Eng. 39 (6), 541–557.

Mosher, J., Leahy, R., Lewis, P., 1999. EEG and MEG: forward solutions for inverse methods. IEEE Trans Biomed. Eng. 46 (3), 245–259.

Nandy, R., Cordes, D., 2003. Novel ROC-type method for testing the efficiency of multivariate statistical methods in fMRI. Magn. Reson. Med. 49 (6), 1152–1162.

Nichols, T., Holmes, A., 2001. Nonparametric permutation tests for functional neuroimaging: a primer with examples. Hum. Brain Mapp. 15, 1–25.

Nolte, G., 2003. The magnetic lead field theorem in the quasi-static approximation and its use for magnetoencephalography forward calculation in realistic volume conductors. Phys. Med. Biol. 48 (22), 3637–3652.

Okada, Y., Wu, J., Kyuhou, S., 1997. Genesis of MEG signals in a mammalian CNS structure. Electroencephalogr. Clin. Neurophysiol. 103, 474–485.

Pantazis, D., Nichols, T., Baillet, S., Leahy, R., 2004. A comparison of random field theory and permutation methods for the statistical analysis of MEG data. NeuroImage (in press).

Park, H., Kwon, J., Youn, T., Pae, J., Kim, J., Kim, M., Ha, K., 2002. Statistical parametric mapping of LORETA using high density EEG and individual MRI: application to mismatch negativities in schizophrenia. Hum. Brain Mapp. 17, 168–178.

Pascual-Marqui, R., Michel, C., Lehman, D., 1994. Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. Int. J. Psychophysiol. 18, 49–65.

Phillips, J., Leahy, R., Mosher, J., 1997. MEG-based imaging of focal neuronal current sources. IEEE Trans. Med. Imag. 16 (3), 338–348.

Robinson, S., Vrba, J., 1999. Functional neuroimaging by synthetic aperture magnetometry (SAM). In: Yoshimoto, T., Kotani, M., Kuriki, S., Karibe, H., Nakasato, N. (Eds.), Recent Advances in Biomagnetism. Tohoku Univ. Press, Sendai, Japan, pp. 302–305.

Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. Phys. Med. Biol. 32, 11–22.

Saulnier, G., Blue, R., Newell, J., Isaacson, D., Edic, P., 2001. Electrical impedance tomography. IEEE Signal Process. Mag. 18 (6), 31–43.

Scales, J., Gersztenkorn, A., 1988. Fast lp solution of a large, sparse linear systems: application to seismic travel time tomography. J. Comput. Phys. 75, 314–333.

Scherg, M., 1990. Fundamentals of dipole source potential analysis. Adv. Audiol. 6, 40–69.

Schimpf, P., Ramon, C., Haueisen, J., 2002. Dipole models for the EEG and MEG. IEEE Trans. Biomed. Eng. 49 (5), 409–418.

Shattuck, D., Leahy, R., 2002. Brainsuite: an automated cortical surface identification tool. Med. Image Anal. 6, 129–142.

Singh, K., Barnes, G., Hillebrand, A., 2003. Group imaging of task-related changes in cortical synchronization using nonparametric permutation testing. NeuroImage 19, 1589–1601.

Starr, S., Metz, C., Lusted, L., Goodenough, D., 1975. Visual detection and localization of radiographic images. Radiology 116, 533–538.

Swets, J., Pickett, R., 1982. Evaluation of Diagnostic Systems: Methods from Signal Detection Theory. Academic Press, New York.

Tarantola, A., 1987. Inverse Problem Theory. Elsevier, New York.

Tikhonov, A., Arsenin, V., 1977. Solutions to Ill-Posed Problems. Wiley, New York.

Tuch, D., Weeden, V., Dale, A., George, J., Belliveau, J., 1999. Conductivity mapping of biological tissue using diffusion MRI. Ann. N. Y. Acad. Sci. 888, 314–316.

Uutela, K., Haemaelaeinen, M., Samelin, R., 1998. Global optimization in the localization of neuromagnetic sources. IEEE Trans. Biomed. Eng. 45 (6), 716–723.

van den Broek, S., Zhou, H., Peters, M., 1996. Computation of neuromagnetic fields using finite-element method and Biot–Savart law. Med. Biol. Eng. 34, 21–26.

van Veen, B., Buckley, K., 1988. Beamforming: a versatile approach to spatial filtering. IEEE ASSP Mag. 5, 4–23.

van Veen, B., van Drongelen, W., Yuchtman, M., Suzuki, A., 1997. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. IEEE Trans. Biomed. Eng. 44-9, 867–880.

von Helmholtz, H., 1853. Ueber einige gesetze der vertheilung elektrischer stroeme in koerperlichen leitern, mit anwenung auf die thierisch-elektrischen versuche. Ann. Phys. Chem., 89, 211–233, 353–377.

Worsley, K., Marrett, S., Neelin, P., Vandal, A., Friston, K., Evans, A., 1996. A unified statistical approach for determining significant signals in images of cerebral activation. Hum. Brain Mapp. 4, 58–73.

Worsley, K., Andermann, M., Koulis, T., MacDonald, D., Evans, A., 1999. Detecting changes in non-isotropic images. Hum. Brain Mapp. 8, 98–101.

Zhang, Z., 1995. A fast method to compute surface potentials generated by dipoles within multilayer anisotropic spheres. Phys. Med. Biol. 40, 335–349.