Bayesian Inference Applied to the Electromagnetic Inverse Problem[†]

David M. Schmidt,* John S. George, and C.C. Wood

Biophysics Group, MS-D454, Los Alamos National Laboratory, Los Alamos, New Mexico

Abstract: We present a new approach to the electromagnetic inverse problem that explicitly addresses the ambiguity associated with its ill-posed character. Rather than calculating a single "best" solution according to some criterion, our approach produces a large number of likely solutions that both fit the data and any prior information that is used. Whereas the range of the different likely results is representative of the ambiguity in the inverse problem even with prior information present, features that are common across a large number of the different solutions can be identified and are associated with a high degree of probability. This approach is implemented and quantified within the formalism of Bayesian inference, which combines prior information with that of measurement in a common framework using a single measure. To demonstrate this approach, a general neural activation model is constructed that includes a variable number of extended regions of activation and can incorporate a great deal of prior information on neural current such as information on location, orientation, strength, and spatial smoothness. Taken together, this activation model and the Bayesian inferential approach yield estimates of the probability distributions for the number, location, and extent of active regions. Both simulated MEG data and data from a visual evoked response experiment are used to demonstrate the capabilities of this approach. *Hum. Brain Mapping 7:195–212, 1999*   **Published 1999 Wiley-Liss, Inc.**[†]

Key Words: MRI; MEG; EEG; brain mapping; inverse problem; Monte Carlo; biomagnetic
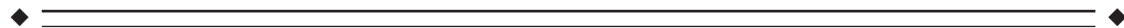
INTRODUCTION

Under suitable conditions of spatial and temporal synchronization, neuronal currents are accompanied by electric potentials and magnetic fields that are sufficiently large to be recorded noninvasively from the surface of the head. These are known as the electroencephalogram (EEG) and magnetoencephalogram (MEG), respectively. In contrast to PET and fMRI, which measure cerebral vascular changes secondary to changes in neuronal activity, EEG and MEG are direct physical consequences of neuronal currents and are capable of resolving temporal patterns of neural activity in the millisecond range [Hämäläinen et al., 1993; Aine, 1995; Toga and Mazziotta, 1996]. Unlike PET and fMRI, however, the problem of estimating the current distribution in the brain from surface EEG and MEG measurements (the so-called electromagnetic inverse problem) is mathematically ill-posed, i.e., it has no unique solution in the most

general, unconstrained case [von Helmholtz, 1853; Nunez, 1981].

Existing approaches to the electromagnetic inverse problem fall into two broad categories: (1) "few-parameter models" (i.e., those in which M ≪ N, where M is the number of parameters to be estimated in the model and N is the number of recording sites), and (2) "many-parameter models" (i.e., those in which M ≥ N). A well-known example of the "few parameter" approach is the single- or multiple-dipole model [e.g., Kavanaugh et al., 1978; Scherg and von Cramon, 1986; Mosher et al., 1992], in which the current is assumed to be represented by a few point-dipoles, the "order" of the model is estimated using Chi-square or related statistical techniques, and the best-fitting values of the dipole parameters (locations, orientations, and magnitudes) are estimated using nonlinear numerical minimization techniques. A well-known example of the "many-parameter" approach is the "minimum-norm linear inverse" [e.g., Hämäläinen and Ilmoniemi, 1984; Dale and Sereno, 1993; Hämäläinen et al., 1993], in which the problem is underdetermined (because M ≥ N) and a strictly mathematical criterion is used to select among the many solutions that fit the data equally well. In the case of the minimum-norm approach, the mathematical criterion is the solution that minimizes the sum of squared current strengths.

Here, we introduce a new probabilistic approach to the electromagnetic inverse problem, based on Bayesian inference [e.g., Bernardo and Smith, 1994; Gelman et al., 1995]. Unlike other approaches to this problem, including other recent applications of Bayesian methods [Baillet and Garnero, 1997; Phillips et al., 1997], our approach does not result in a single "best" solution to the problem. Rather, we estimate a probability distribution of solutions upon which all subsequent inferences are based. This distribution provides a means of identifying and estimating the likelihood of features of current sources from surface measurements that explicitly emphasize the multiple solutions that can account for any set of surface EEG/MEG measurements.

In addition to emphasizing the inherent probabilistic character of the electromagnetic inverse problem, Bayesian methods provide a formal, quantitative means of incorporating additional relevant information, independent of the EEG/MEG measurements themselves, into the resulting probability distribution of inverse solutions. Such information might include constraints derived from anatomy on the likely location and/or orientation of current [Wang et al., 1992; George et al., 1995; Baillet and Garnero, 1997; Dale, 1997], maximum current strength, spatial and/or temporal smoothness of current, etc. The Bayesian approach also provides a way to marginalize over nuisance variables that cannot be determined or resolved from the data. A related approach, concerning Bayesian confidence regions, was presented in Hämäläinen et al. [1993].

We begin with an overview of the general techniques of Bayesian inference. Then, we show how these techniques may be applied to the EEG/MEG inverse problem and demonstrate their use in examples from both simulated MEG data and MEG data from a visual evoked response experiment.

BAYESIAN INFERENCE

Bayesian inference (BI) is a general procedure for constructing a (posterior) probability distribution for quantities of interest from the measurements given (prior) probability distributions for all of the uncertain parameters—both those that relate the quantities of interest to the measurements and the quantities of interest themselves. The method is conceptually simple, using basic laws of probability, making its application even to complicated problems relatively straightforward. The posterior probability distribution is often too complicated to be calculated analytically, but usually can be adequately sampled using modern computer techniques, even in problems with many parameters. The method is outlined here; more detailed presentations can be found elsewhere [e.g., Gelman et al., 1995].

The starting point for BI is Bayes' rule of probability:

$$P(\theta, y) = P(\theta|y)P(y), \tag{1}$$

where $P(\theta, y)$ is the joint probability distribution for the quantities, θ and y, $P(\theta|y)$ is the conditional probability distribution of θ given y, and $P(y)$ is the marginalized probability distribution of y; $P(y) = \Sigma_\theta P(\theta, y)$ (or $P(y) = \int P(\theta, y)d\theta$ for continuous θ). If θ represents parameters about which we wish to learn and y represents data bearing upon θ, then the probability of θ given y can be constructed from Bayes' rule as:

$$P(\theta|y) = \frac{P(\theta, y)}{P(y)} = \frac{P(y|\theta)P(\theta)}{P(y)}. \tag{2}$$

Here, $P(\theta)$ is the prior probability distribution of θ, which represents one's knowledge of θ prior to the measurement. This is modified by the data through the likelihood function, $P(y|\theta)$, to produce the posterior probability distribution, $P(\theta|y)$. Since $P(y)$ is independent of θ it can be considered a normalizing constant and can be omitted from the unnormalized posterior density:

$$P(\theta|y) \propto P(y|\theta)P(\theta). \tag{3}$$

As summarized in Gelman et al. [1995], "These simple expressions encapsulate the technical core of Bayesian inference: the primary task of any specific application is to develop the model $P(\theta, y)$ and perform the necessary computations to summarize $P(\theta|y)$ in appropriate ways."

Bayesian inference applied to the EEG/MEG inverse problem

Activity model

The methods of BI applied to the EEG/MEG inverse problem are demonstrated within the context of a model for the regions of activation, which is intended to be applicable in evoked response experiments. There is both theoretical and experimental evidence that EEG and MEG recorded outside the head arise primarily from neocortex, in particular from apical dendrites of pyramidal cells [e.g., Allison et al., 1986; Dale and Sereno, 1993; Hämäläinen et al., 1993]. We, therefore, construct a model that assumes a variable number of variable size cortical regions of stimulus-correlated activity in which current may be present. Specifically, an active region is assumed to consist of those locations that are identified as being part of cortex and are located within a sphere of some radius r centered on some location w, also in cortex. There can be any number n of these active regions up to some maximum n_{max}, and the radius can have any value up to some maximum r_{max}. The goal is to determine the posterior probability values for the set of activity parameters $\boldsymbol{\alpha} = \{n, w, r\}$, which govern the number, location, and extent of active regions.

Probability model for activity parameters

The first step in BI is to construct a probability model that relates the activity parameters to the measurements. Let the N measurements at one instant in time be denoted by $\mathbf{b} = \{b_1, \ldots, b_N\}$. The conditional probability of the activity parameters given the observed data, $P(\alpha|b)$, can be expressed using Bayes' rule of probability as

$$P(\alpha|b) \propto P(b|\alpha)P(\alpha) \tag{4}$$

where $P(\alpha)$ is the prior probability for the activity parameters and $P(b|\alpha)$ is the probability of the data given a particular set of values for the activity parameters. The prior probability for the activity parameters will be set by the experimenter using physiological information about the particular experiment being analyzed. Because the data do not depend on the activity parameters directly, but rather on a given current distribution $\mathbf{j}$, the function, $P(b|\alpha)$, cannot be specified until first expanding it to include the dependence of the measurements on the current. This may be accomplished by marginalizing out the current in the joint probability of the data and current such that

$$\begin{aligned} P(b|\alpha) &= \int P(b, \mathbf{j}|\alpha) \, \oslash \mathbf{j} \\ &= \int P(b|\mathbf{j}, \alpha)P(\mathbf{j}|\alpha) \, \oslash \mathbf{j} \end{aligned} \tag{5}$$

where the integral is a functional integral over all current distributions. The function, $P(b|\mathbf{j}, \alpha)$, is the likelihood function of the data; there is no explicit dependence upon α since $\mathbf{j}$ is all that is needed completely to specify $P(b|\mathbf{j}, \alpha)$. In particular, since most evoked response experiments are repeated many times and averaged, it is assumed that the likelihood function is Gaussian such that

$$P(b|\mathbf{j}, \alpha) \propto \exp\left\{-\frac{1}{2}\sum_{k,l=1}^{N} (b_k - \langle a_k, \mathbf{j}\rangle)C_{kl}^{-1}(b_l - \langle a_l, \mathbf{j}\rangle)\right\}. \tag{6}$$

Here, C is the covariance matrix of the noise or background in the measurements and $\mathbf{a}$ are the forward fields or measurement kernel such that if there were no noise or background the measurements would be related to the current by the inner product:

$$b_k = \langle a_k, \mathbf{j}\rangle = \int a_k(x) \cdot \mathbf{j}(x)d^3x. \tag{7}$$

We find it convenient to use an equivalent representation of Eq. 6, which has the noise covariance absorbed into a new set of effective measurements and forward

fields, $\tilde{b}$ and $\tilde{a}$, such that Eq. 6 becomes

$$P(b|j, \alpha) \propto \exp\left\{-\frac{1}{2}\sum_{k=1}^{N} (\tilde{b}_k - \langle\tilde{a}_k, j\rangle)^2\right\}. \quad (8)$$

The function $P(j|\alpha)$, which gives the probability of any current given a particular set of activity parameters, needs to be constructed. Clearly, the current should be zero outside of active regions. Furthermore, we would like to be able to incorporate prior information about the limits of current strength, spatial variability, and orientation of the current within active regions. For example, high-resolution current source density estimates suggest that net cortical currents are oriented predominately perpendicular to the cortical surface [Mitzdorf, 1985]. These forms of prior information may be incorporated in a manner that simplifies computing the marginalization over j in $P(b, j|\alpha)$, by using a Gaussian distribution such that

$$P(j|\alpha) \propto |V_\alpha|^{-1/2} \exp\left\{-\frac{1}{2}\langle j, V_\alpha^{-1}j\rangle\right\} \quad (9)$$

where V_α^{-1} is the inverse of the covariance operator (matrix) of the current. The diagonal elements, or the variances, serve to limit the current strength, and the off-diagonal elements, which are related to the correlation coefficients, can serve to restrict the smoothness and orientation of the current distribution. The variance at locations that are not part of any active region for a given α is set to zero. The experimenter needs to set the values of the covariance matrix, based on knowledge of the experiment to be analyzed, using prior information about the strength, orientation, and spatial variability of current within active regions.

The full probability model for the activity parameters is

$$P(\alpha|b) \propto P(\alpha)|V_\alpha|^{-1/2}$$

$$\int \exp\left[-\frac{1}{2}\left\{\sum_{k=1}^{N} (\tilde{b}_k - \langle\tilde{a}_k, j\rangle)^2 + \langle j, V_\alpha^{-1}j\rangle\right\}\right]\oslash j \quad (10)$$

and $P(\alpha)$ is set by the experimenter. The integral over the current may be constructed using the eigenvalues, $\{\lambda_\theta(\alpha)\}$, and normalized eigenvectors, $\{\psi_\theta(\alpha)\}$ ($\theta = 1, \ldots, N$), of the matrix $G_{k,l}(\alpha) = \langle\tilde{a}_k, V_\alpha\tilde{a}_l\rangle$; all of which may be calculated using standard numerical techniques. Using these eigenvalues and eigenvectors the formula for the

posterior probability distribution becomes

$$P(\alpha|b) \propto P(\alpha)\exp\left[-\frac{1}{2}\left\{\sum_{k,\theta,l} \tilde{b}_k \frac{\psi_{k,\theta}(\alpha)\psi_{l,\theta}(\alpha)}{1 + \lambda_\theta(\alpha)} \tilde{b}_l\right.\right.$$

$$\left.\left.+ \sum_\theta \ln(1 + \lambda_\theta(\alpha))\right\}\right]. \quad (11)$$

This formula is well behaved and is not overly sensitive to very small eigenvalues. Moreover, it is relatively simple to compute because it depends only on the N by N matrix, $G_{k,l}(\alpha)$.

Sampling the posterior

The next step in BI is to use the posterior probability distribution in order to answer questions related to the activity parameters in terms of probability. Examples of such questions include: what is the probability that there were m regions of activity? What are the locations for these active regions at a 95% probability level? In cases where the number of different possible sets of activation parameters is small, one can evaluate the complete posterior distribution. Generally, however, the number of different possible sets of activation parameters is large. In such cases, the method of Markov Chain Monte Carlo (MCMC) can be used to generate a sample of sets of activity parameters that are distributed according to the posterior distribution. This is known as sampling the posterior, the techniques for which are described in detail elsewhere [e.g., Gelman et al., 1995].

Examples

Although the methods described above apply to models for both EEG and MEG data, in the remainder of this report, we use MEG data to illustrate the properties of the approach. Both simulated and empirical MEG data for a Neuromag-122 whole-head system were used [Ahonen et al., 1993]. The physical setup of the actual MEG experiment was used to determine the location of the subject's head relative to the sensors in the simulated data examples. In addition, an anatomical MRI data set acquired from the subject in the MEG experiment was used to determine the location of cortex (actually gray matter) using MRIVIEW (Fig. 1), a software tool developed in our laboratory [Ranken and George, 1993]. About 50,000 voxels were tagged and the normal directions for each of these voxels was
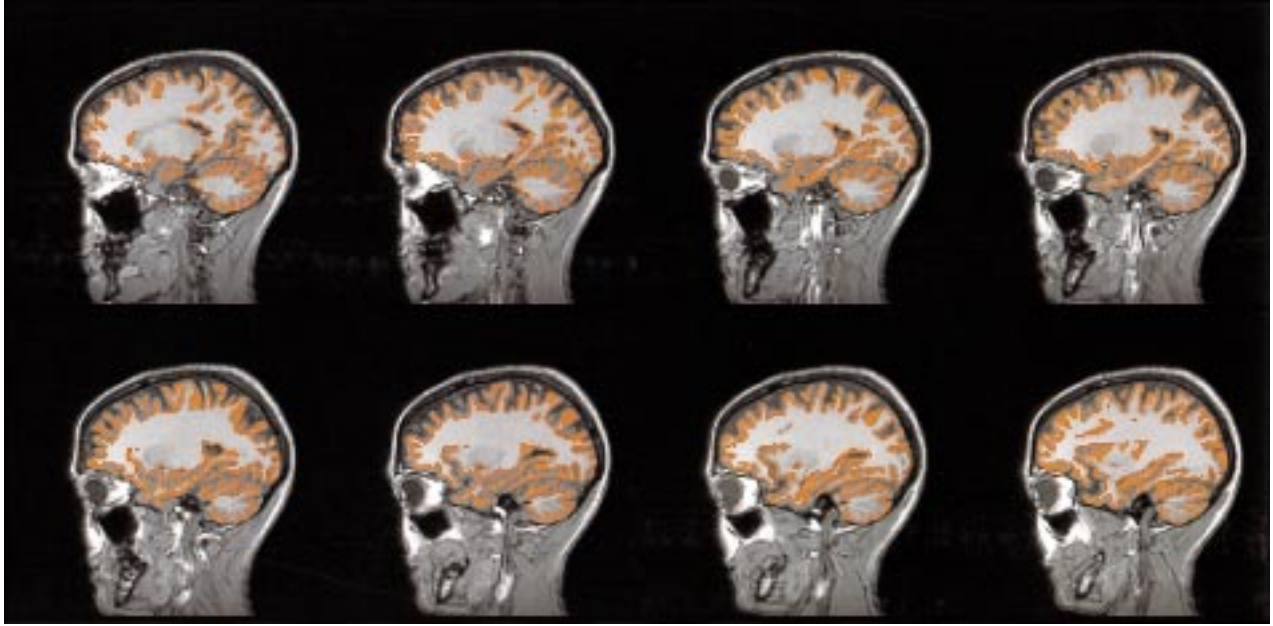
Figure 1.
Gray matter regions are tagged (in red) from anatomical MRI data. These tagged voxels constitute the anatomical model used to implement the cortical location and orientation prior information.

then determined by examining the curvature of the local tagged region.

A spherically symmetric conductivity model was used to calculate the expected measurements given a current source both for the simulated data and in the likelihood calculations [Sarvas, 1987]. The same prior assumptions were used with the simulated data sets, with only minor changes for the real data example. Specifically, the prior probability function $P(\alpha)$ was uniform so that each set of activation parameters had the same prior probability. The number of active regions was allowed to range from 0 to 8, and the radius of any region of activity was allowed to range from 0 to 10 mm.

The covariance matrix was factored such that

$$^{\beta\gamma}V_{\alpha}(i, j) = \sigma_{\alpha}(i)\sigma_{\alpha}(j)\rho(i - j)^{\beta\gamma}\Omega(i, j) \qquad (12)$$

where β and γ are orientation indices, i and j are location indices, $\sigma_{\alpha}(i)$ is the standard deviation at location i, $\rho(i - j)$ is the spatial correlation function, and $^{\beta\gamma}\Omega(i, j)$ is the orientation covariance. The correlation function was chosen to be a Gaussian with zero mean and 7 mm standard deviation, which imposes spatial smoothness on scales of ~7 mm or less. Because of this prior information concerning spatial correlation, the continuous current distributions and integrals

of the previous section may be well approximated by discrete distributions and sums over the volume elements (voxels) that were tagged from the anatomical MRI data. For example, in evaluating the posterior probability value using Eq. 11, the matrix G is calculated in the following examples by approximating the continuous integral with a sum over tagged voxels. This is a good approximation because the covariance operator has a correlation length of 7 mm, which is larger than the voxel dimensions of 2 mm on a side.

To complete the specification of the covariance operator, a value of 2 nAm was used for $\sigma_{\alpha}(i)$ at all locations in active regions and 0 nAm elsewhere. The orientation covariance was chosen such that there was no correlation between the orientations at different locations and the orientation distribution at any given location was symmetric with respect to the direction normal to the cortical surface at that location and had a mean equal to the cortical norm direction and a standard deviation of 30°. Unlike other recent implementations of cortical constraints in distributed inverse solutions [e.g., Dale and Sereno, 1993; Baillet and Garnero, 1997], this procedure results in a distribution of orientations around the perpendicular, not a fixed normal orientation.

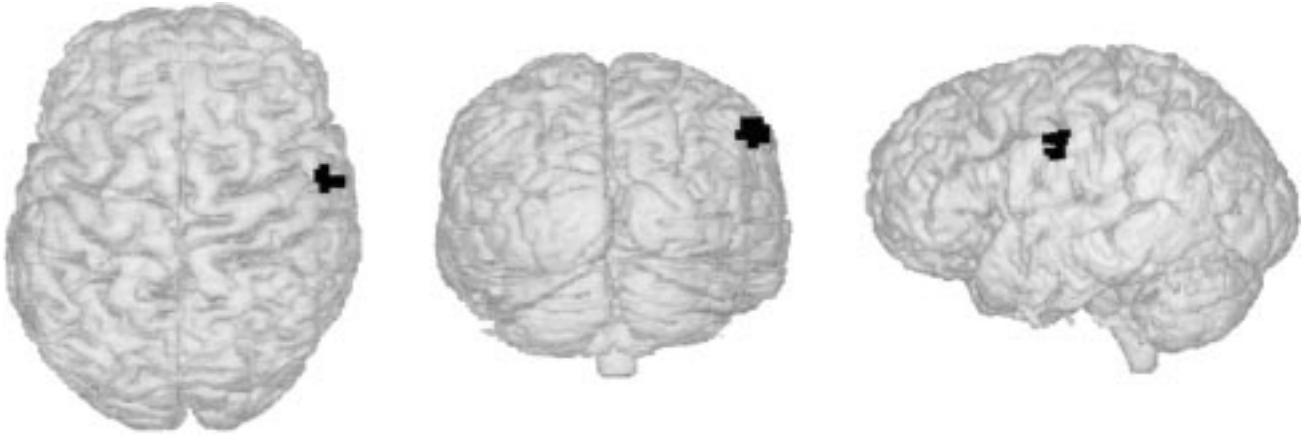Finally, the same noise was added to all simulated data sets, which was Normal with a standard

Figure 2.
Maximum intensity projection of the location and extent (in black) of the active region used to generate simulated MEG data for Example 1.

deviation of 10 fT. The values used here in the prior probability distribution are meant to be an example of what one might choose for a MEG analysis and should be chosen for each particular MEG data set.

Example 1

The location and extent of the active region used to generate the simulated MEG data are shown in Figure 2. The bounding radius of the active region was 5 mm and the current dipole strength at each voxel was 2 nAm oriented in the cortical normal direction. A plot of the simulated data and some views of its field pattern are shown in Figure 3. Ten-thousand samples were drawn from the posterior distribution using a MCMC algorithm. The program ran overnight, not having been optimized, on a current personal computer. It took about 600 samples to progress from the starting point, which had a low probability to one that had a high probability and was therefore representative of the posterior distribution. Only the final 9,000 samples, a few of which are shown in Figure 4, were used in making probabilistic inferences as discussed below. All of the samples shown in Figure 4 are among the 95% most probable and therefore fit both the data and the prior expectations quite well. Any of these could have produced the given MEG data; yet there are clearly vast differences among the samples. The number of active regions ranges from 1 to 5, the sizes of the regions vary greatly, and the

locations of the active regions vary nearly across the entire tagged region of the brain (when considering all 9,000 samples). This variability is a representation of the degree of the ambiguity of the inverse problem for these MEG data, even with the prior information present.

Despite the degree of variability among the samples in Figure 4, a property common to all is apparent; namely, an active region in the dorsal, lateral region of the right hemisphere. A feature such as this, common to all or most of the samples, is associated with a high degree of probability. This probability can be quantified because the MCMC samples are distributed according to the posterior probability distribution. The smallest set of voxels that contains the center of the active region in the dorsal, lateral region in 95% of the samples was identified and is shown in Figure 5. This region, which contains a center of activity with a probability of 95%, in fact encompasses the region of activity that was used to produce the simulated data set (Fig. 2). Although it is good to see this agreement, it is not sufficient to justify this or any MEG inverse method based solely on whether it produces results consistent with the true active regions because any of the sets of active regions shown in Figure 4 also could have been used to generate the same MEG data. Any robust and highly probable result or inference, therefore, should be consistent with the wide range of possible sets of active regions, as is the result in Figure 5 by construction. This is a very important feature of BI, which is necessarily missing from any other analysis method that only considers just one possible result,
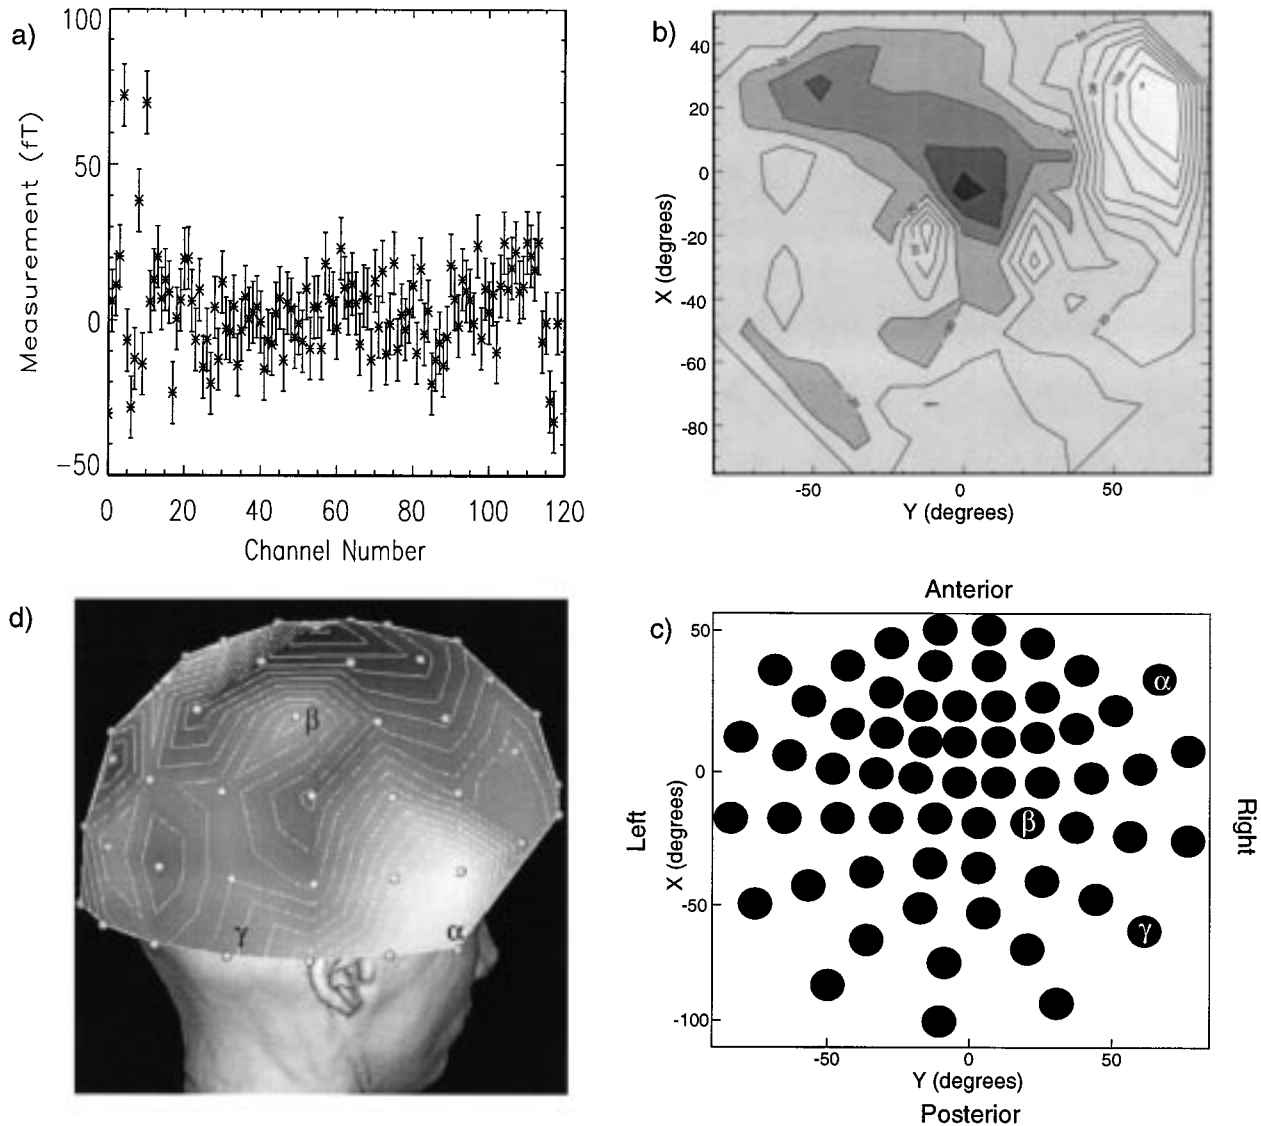
Figure 3.

The simulated data used in Example 1 as a function of channel number. (**a**): To aid in visualizing the field pattern associated with this data panels (**b**) and (**d**) display iso-amplitude contours of a field on the sensor helmet surface that is consistent with the data shown in (a). Positive fields (flux emerging from the head) are displayed in light shades and negative fields (reentering flux) are shown in dark shades. (b) is a polar projection, viewing the helmet surface from above, for which the relative position of the sensors and orientation of the subject's head are shown in (**c**). The sensors labeled α, β and γ in (c) correspond to those with the same labels in (d).

even if it happens to be the most likely result within a given model.

In addition to the information about the locations of probable regions of activity, the Bayesian approach combined with this activity model also provides probabilistic information about the number and size of active regions. The posterior distribution for the number of active regions was constructed by histograming the number of regions across the MCMC samples. This histogram is shown in Figure 6a. One active region is the most probable; however, two active regions are quite likely as well. Although the location of one active region was identified above, the location of a second could not be well localized because it occurred in a wide range of locations across the MCMC samples. Assuming there was only one active
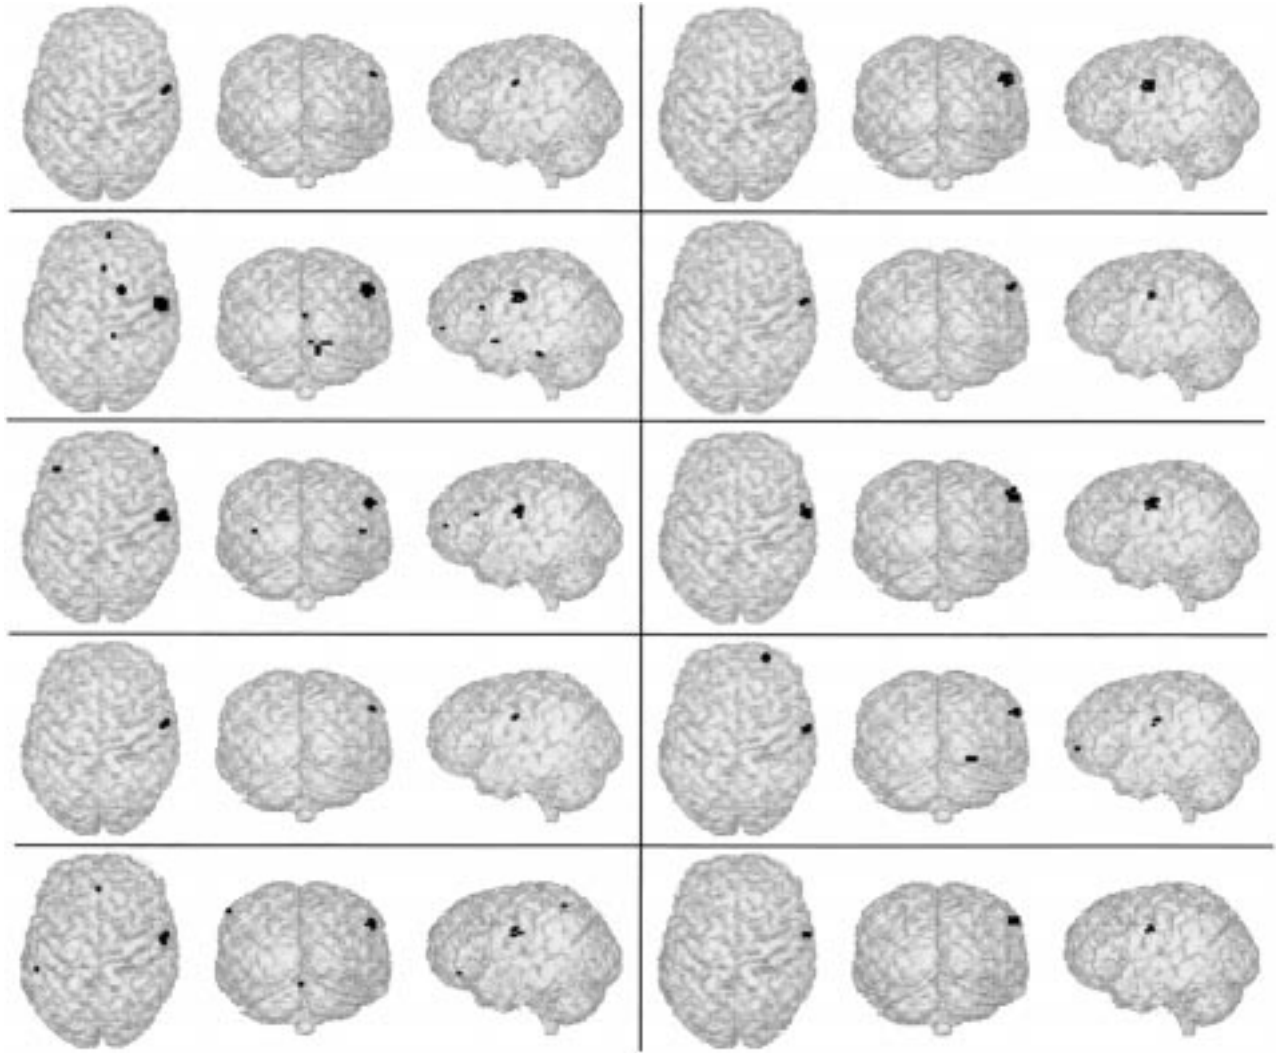
Figure 4.
A few of the 9,000 samples drawn from the posterior probability distribution of Example 1. Each panel shows three views of the maximum intensity projection of all of the active regions from a single sample. All of these samples could have produced the same MEG data set.

region present, we can obtain information about its size by histograming the size of active regions in the MCMC samples that had only one region present. This histogram is shown in Figure 6b and represents the posterior probability for the bounding radius of the active region, assuming that there was only one region active. Regions <2 mm and >8 mm in radius are very unlikely, whereas regions that are ~5 mm in radius are likely. The size of the region used to produce the simulated data was 5 mm. We believe that much of the

information on size derives from prior information about location, orientation, and strength of neural current.

Other inferences could be drawn using the MCMC samples in a similar manner. For example, one could construct the probability for the size of the active region, assuming there was one centered within the 95% probability region shown in Figure 5, rather than assuming there was only one active region present throughout the entire head as was done above.
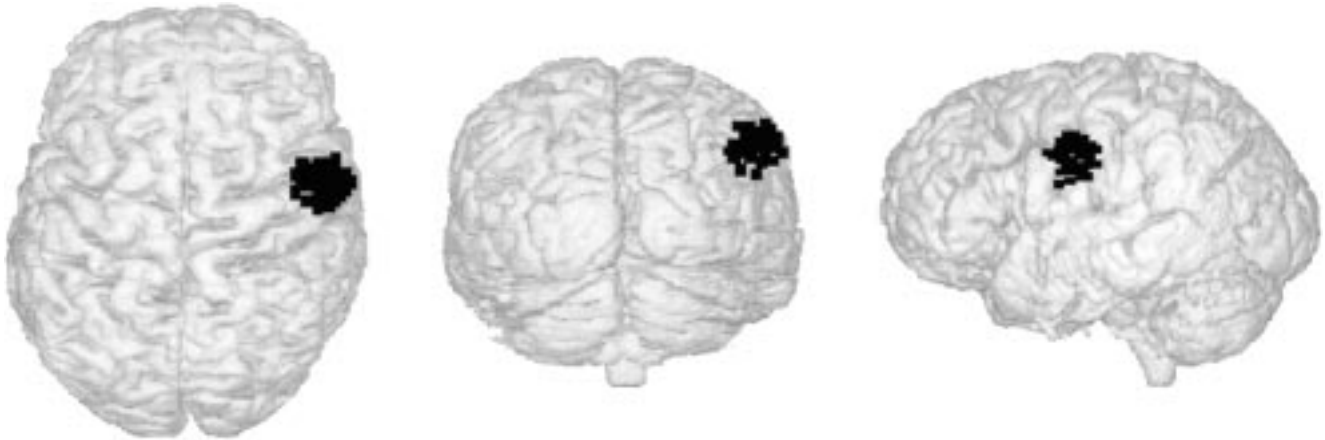
Figure 5.
Maximum intensity projections of the location and extent of a region containing a center of activity at a 95% probability level in Example 1.

Example 2

A second simulated data set was generated using the three active regions of different sizes shown in Figure 7. The most anterior region is centered at the same location as the region in the first example except for this case it has a bounding radius of 8 mm. A current dipole strength of 2 nAm oriented normal to cortex was used at each voxel within this bounding sphere. The nearby more posterior region had a bounding radius of 5 mm and a current dipole strength of 2.5 nAm was used. The most posterior region had a 3 mm bounding radius and a current dipole strength of 1.5 nAm. The same noise and prior assumptions where used here as for the first example. Figure 8 shows a plot of the resulting simulated data and the field pattern.

Ten-thousand samples were drawn from the posterior, of which the final 8,000 were used to make probabilistic inferences. Just as in the first example, we expect there to be many different locations where activity may be found in these samples. Since we are interested in those locations that contain activity in most of the samples, it is useful to make a histogram of the locations of the centers of active regions across the 8,000 samples. This histogram is shown in Figure 9. It is relatively simple to determine those regions that contain centers of activity at a 95% probability level from this histogram by centering each region on the local peaks in the histogram and expanding the radius of each region until a 95% level is reached. This was done for the three peaks present in Figure 9 and is shown in Figure 10. These regions are consistent with

the locations of the active regions used to generate the simulated data, but what is more important is that these regions are consistent with at least 95% of the likely sets of active regions that could have also generated this data. This is true even when allowing a variable number of active regions of variable extent. Furthermore, these regions are not necessarily the only regions that could have been active. As shown in Figure 11, there is significant probability that more than three regions may have been active. Shown in Figure 10 are the locations of those active regions that occurred consistently in well-localized areas across the MCMC samples. Other possible active regions were not so well localized.

In order to learn about the extent or size of each of the active regions localized in Figure 10, a histogram of the radius of the active regions present in each of the areas shown in Figure 10 across the samples was made. This represents the posterior probability for the size of active regions, assuming there was an active region in each of these areas. These plots are shown in Figure 12. Recall that the radii of the actual regions used to generated the data were 8 mm, 5 mm, and 3 mm for the regions in anterior-to-posterior order. The agreement between actual radii and posterior probabilities is especially remarkable given the variation in the current strengths of the regions used to generate the data. Such information on extent can be very useful, is not present in most other current methods for analyzing MEG data, and is affirmation of the likely utility of anatomical and physiological prior information.
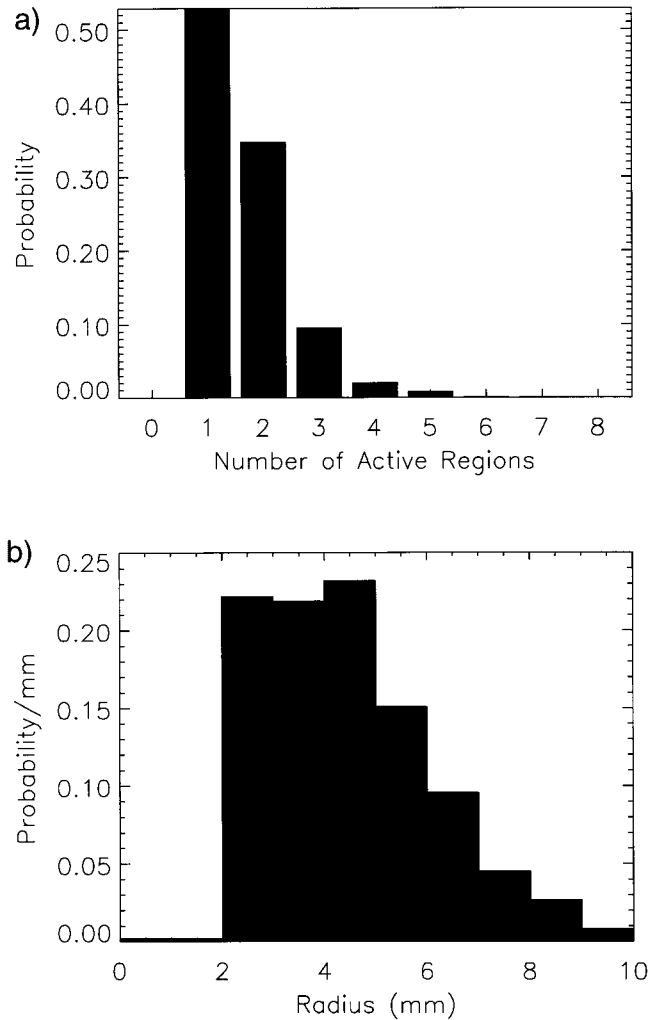
Figure 6.
The posterior probability for (**a**) the number of active regions present in Example 1 and (**b**) the radius of the sphere bounding activity in Example 1, assuming there was only one active region present.

Example 3

The final example, based on MEG data from a visual evoked response experiment [Aine et al., 1997a], illustrates the feasibility and the value of the approach with actual data. In order to examine the sensitivity of the Bayesian approach to detect known features of human visual cortex organization, we compared Bayesian analyses of MEG responses to visual stimuli in the left and right visual fields. Based on the crossed anatomical projections of the visual fields to the brain and on previous lesion, MEG, and fMRI studies in humans

[e.g., Horton and Hoyt, 1991; Sereno et al., 1995; Aine et al., 1996], initial cortical activation for stimuli in the left and right visual fields should occur near the calcarine fissure in the occipital region of the contralateral hemispheres.

The visual stimuli were black-white circular sinusoidal patterns, 1.0° in diameter, presented near the horizontal meridian at 6.2° in the left and right visual fields. The stimulus duration was 250 ms and the average interstimulus interval was 1.0 s. One-hundred epochs (from 100 ms before each stimulus to 400 ms after each stimulus) were averaged; bad channels were identified and removed before data analysis. The variance of the noise was estimated by calculating the variance of the prestimulus epoch. The same model and the same prior information used in Examples 1 and 2 were used for the Bayesian analyses in this example, except that the standard deviation of the current strength was assumed to be 8 nAm instead of 2 nAm. This value is consistent with the maximum current strength measurements in Okada et al. [1998].

The model was applied separately to the data for each visual field stimulus at 10 ms intervals from 110 ms to 160 ms poststimulus. Ten-thousand samples of the posterior probability were generated for each latency. The results to be presented here are from 110 ms and 150 ms following stimulus onset; latencies that should include robust activation of the calcarine region [Aine et al., 1996]. Figure 13 presents the field distributions for these data.

The top of Figure 14 presents maximum intensity projections of the probability of activity for each voxel in the anatomical model for the left and right visual field stimuli at two different latencies following stimulus onset (110 and 150 ms, respectively). This probability distribution was constructed by calculating the fraction of MCMC samples in which each voxel had activity and is a marginalization of the full posterior probability distribution onto the space of anatomical voxels. The bottom of Figure 14 presents the posterior probability marginalized onto the number of active regions for each latency and visual field combination.

For the left visual field stimulus, maximal probability of activation at 110 ms was located in the right (contralateral) hemisphere, centered upon the calcarine region. This pattern was reversed for the right visual field stimulus at 110 ms, consistent with the predictions from anatomy, and from the lesion, fMRI, and previous MEG studies cited above. In order to
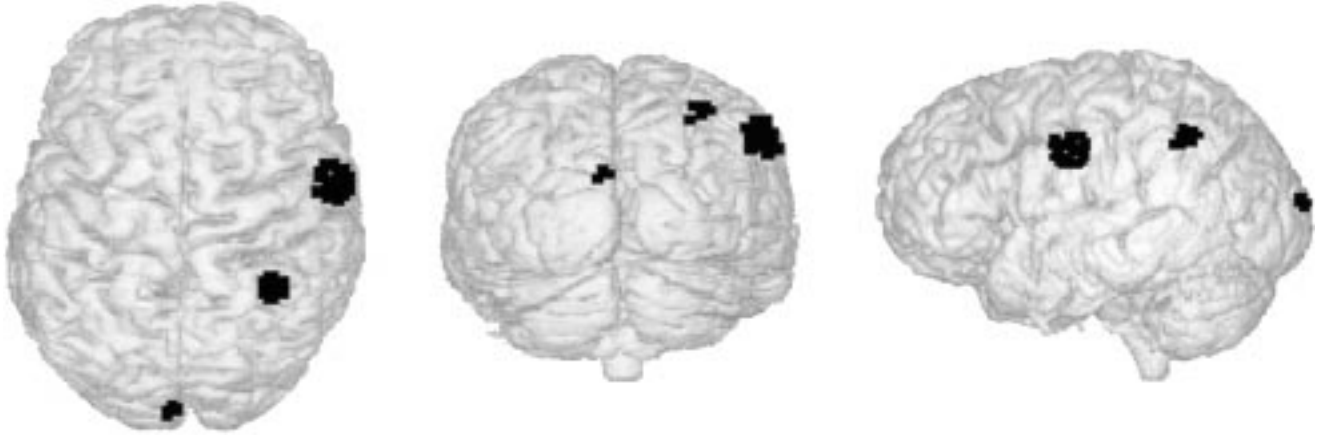
Figure 7.
Maximum intensity projections of the location and extent of the active regions used to generate the simulated MEG data for Example 2.

show this more clearly, regions that contained activity at a probability level of 95% were identified and are shown in greater detail in Figure 15, which depicts relative probability of activation within these regions on a color scale in three orthogonal slices through the calcarine region and a three-dimensional rendering of the occipital region.

For both the left and right visual field stimuli, the most probable number of active regions at 110 ms latency was two, suggesting that active regions in addition to the most probable ones in the calcarine regions of each hemisphere were needed to account for the data and prior information. However, these regions were inconsistently located over the Monte Carlo
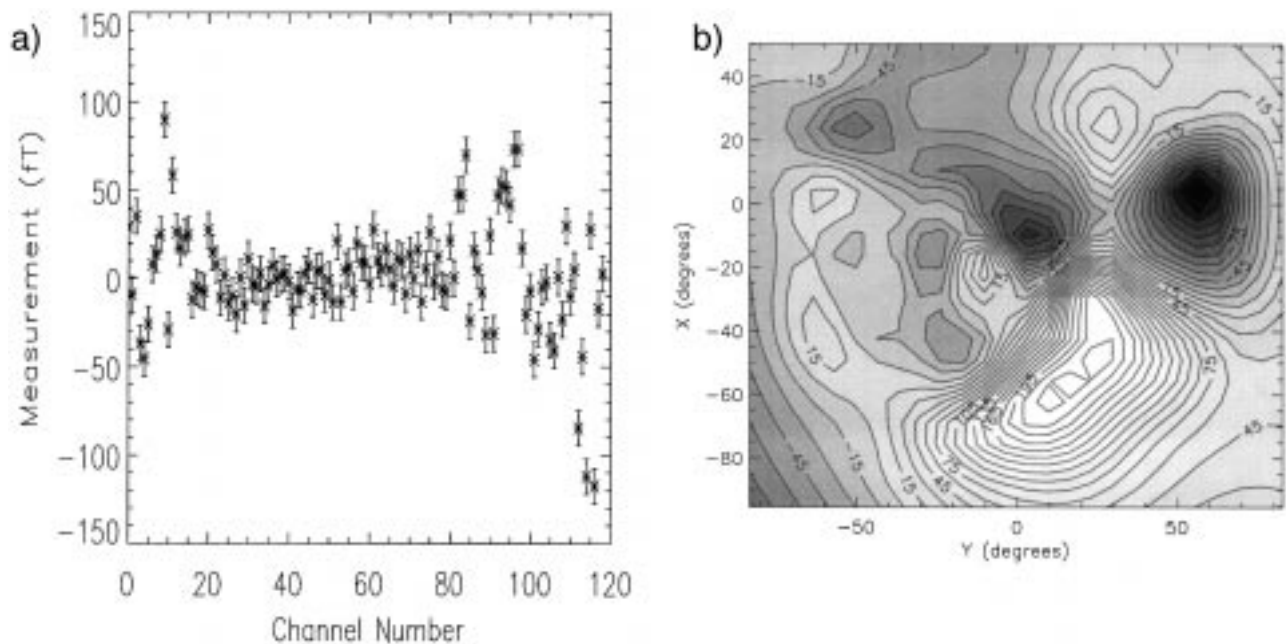

Figure 8.
The simulated data of Example 2, (**a**) as a function of channel number and (**b**) as a field pattern in a top projection view that is described in Figure 3.
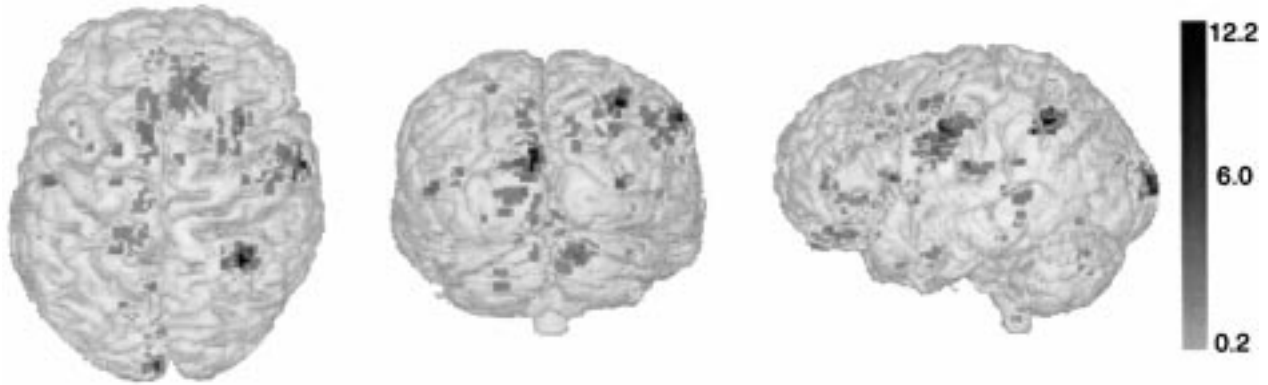
Figure 9.
Maximum intensity projections of the histogram of centers of active regions across the MCMC samples in Example 2, shown on top of surface renderings of cortex. The color bar shows the mapping between shade of grey and histogram amplitude as a percentage of the number of samples used.

samples, as indicated by the relatively widespread regions of low probability in addition to the focus of high probability in Figure 14.

At 150 ms, the most probable number of active regions increased to five for the left visual field stimulus and to three for the right visual field stimulus. Regions of highest probability in each case were located in parieto-occipital and temporo-occipital regions of the hemisphere contralateral to the visual field stimulated. These results are consistent in general terms with MEG and fMRI evidence of multiple regions of extra-striate activity [Aine et al., 1997a,b; Shah et al., 1998], although much additional work is

needed to obtain a definitive comparison of Bayesian inference, multiple-dipole, and fMRI estimates of activity in such experiments.

Two additional features of the results in Example 3 should be noted. First, although maximal probability of activation at the 110 ms latency was indeed located in the opposite hemisphere, there exists sizable probability for activity in the ipsilateral hemisphere near the midline. The extent of the 95% probability regions shown in Figure 15 is indicative of both the extent of estimated activation and the degree of error or uncertainty in that estimate even allowing for the possibility of different numbers of active regions of variable
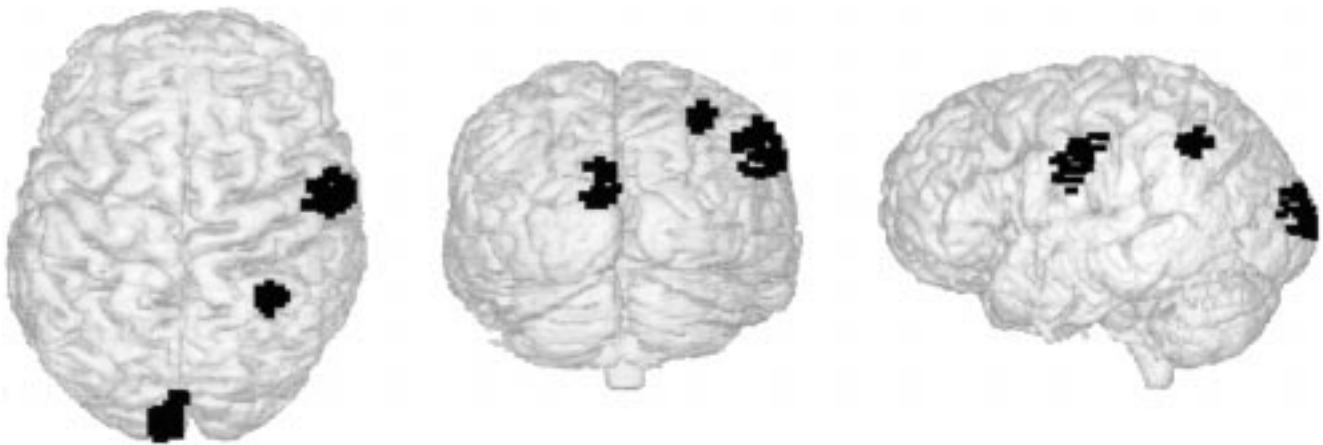

Figure 10.
Maximum intensity projections of the location and extent of the three regions that were found to contain centers of activity at a probability level of at least 95% in Example 2.
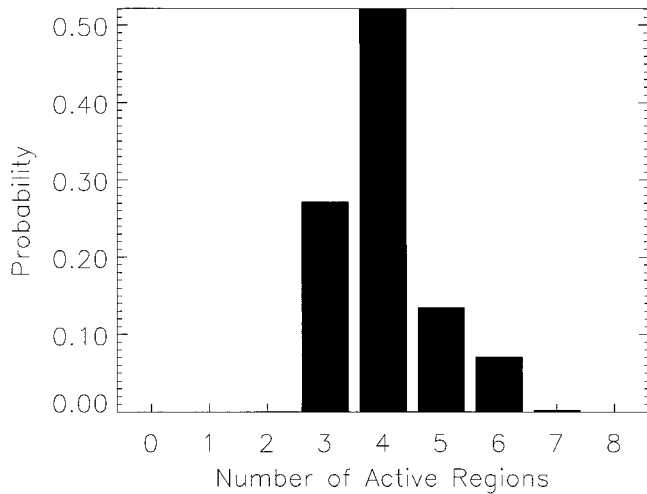
Figure 11.
The posterior probability for the number of active regions in Example 2.

extent. Second, although not shown in detail here, analyses at other latencies suggest a progressively increasing number of probable regions of activation, in both the ipsilateral and contralateral hemispheres, over the latency region from 110–160 ms following stimulus onset. It will be of considerable interest to explore the time dependence of the Bayesian inference analyses in relation to evidence for multiple, functionally organized areas of striate and extra-striate visual cortex and to examine the value of temporal prior information (not included in the current activation model) in the form of, for example, temporal covariance constraints.

DISCUSSION

We have demonstrated a method for analyzing EEG/MEG data that directly addresses the ill-posed character of the electromagnetic inverse problem by allowing probabilistic inferences to be drawn about regions of activation from a large number of possible solutions that both fit the data and the prior expectations made explicit by the Bayesian approach. In addition, we have introduced a model for the current distributions corresponding to neural activity that produce MEG (and EEG) data that are not overly restrictive, allow extended regions of activity, and can easily incorporate prior information such as anatomical constraints from MRI [Dale and Sereno, 1993; George et al., 1995].

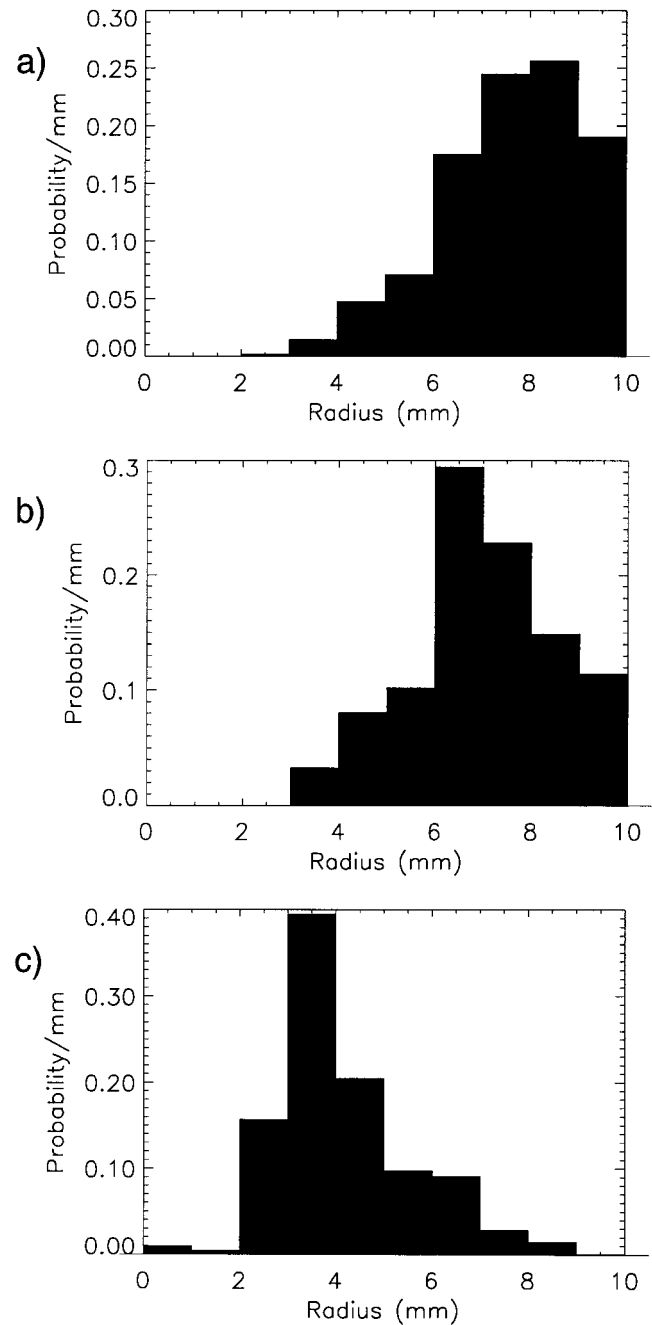Other investigators have applied the Bayesian formalism to various models for EEG or MEG inverse

Figure 12.
The posterior probability distributions for the size of the three active regions whose centers are shown in Figure 10 in anterior to posterior order. The true sizes of the regions used to generate the simulated data are: (**a**) 8 mm, (**b**) 5 mm, and (**c**) 3 mm.

problems [e.g., Phillips et al., 1997; Baillet and Garnero, 1997]. Any Bayesian approach requires: (1) a model that relates the EEG/MEG measurements to underlying neuronal currents, and (2) an implementation of
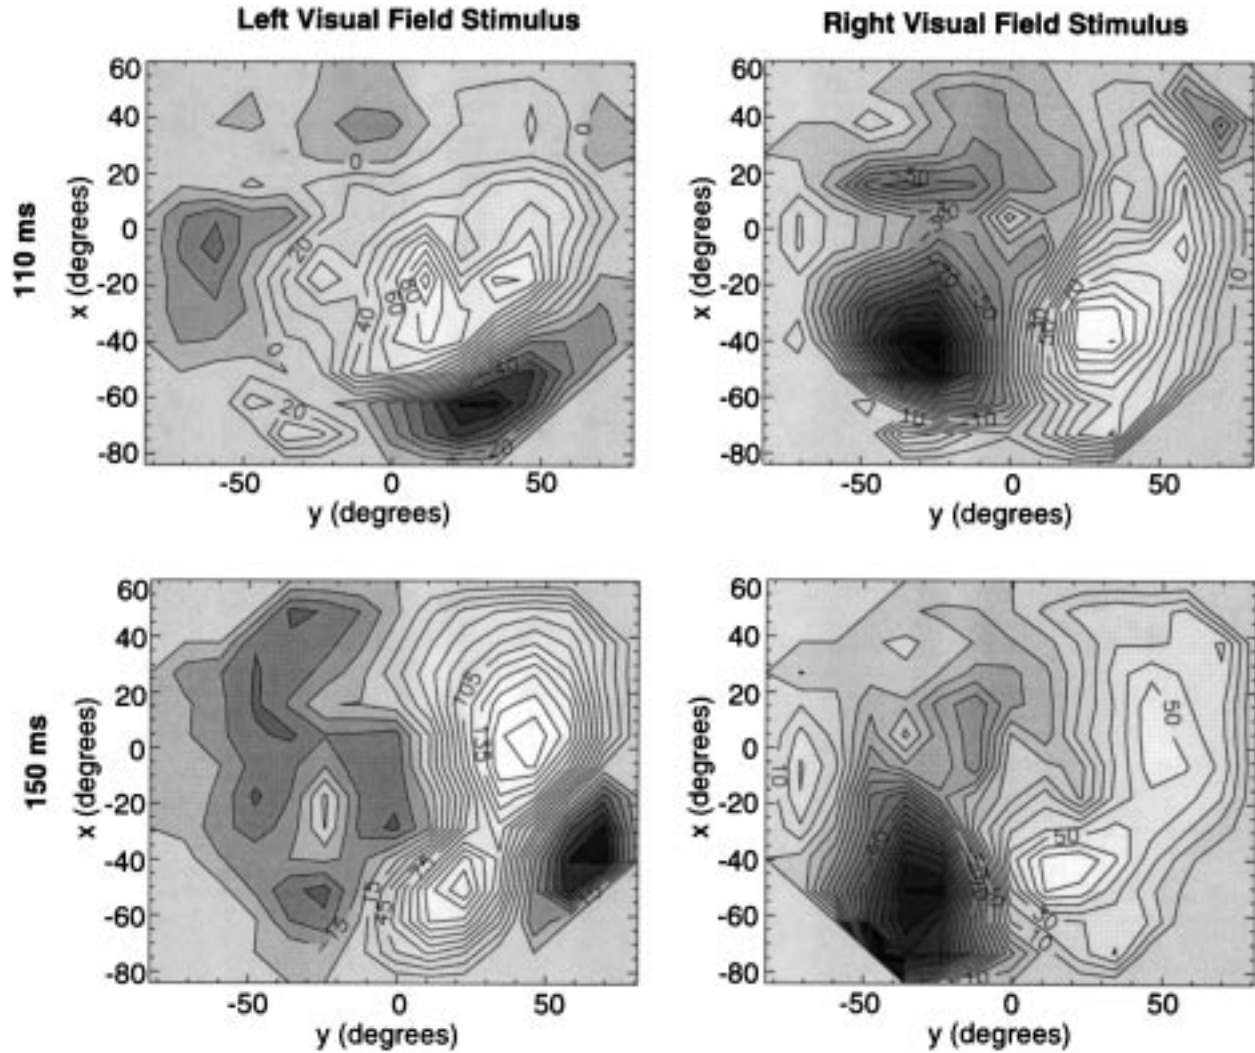
Left Visual Field Stimulus   **Right Visual Field Stimulus**

Figure 13.
The field patterns for the real MEG data in Example 3 in a top projection view that is described in Figure 3.

that model within the Bayesian formalism, including the nature and parameterization of independent prior information. The approach presented here differs from those of Phillips et al. [1997] and Baillet and Garnero [1997] both in the form of the activity model employed and in the manner in which the Bayesian formalism is exploited. Much additional work is needed to determine which combination of source model and prior information is most useful for the electromagnetic inverse problem. However, we believe it is clear even at this early stage that the strategy of estimating the probability distributions for model parameters (e.g., the number, location, and extent of active regions in the present model) is a richer, more

robust, and more realistic basis for inference than estimating a single "best-fitting" solution within a given model. As we have demonstrated, considering only one solution, even if it is the most likely, is not necessarily representative of the range of possible solutions that both fit the data and the prior information. Only by considering this full range of possible solutions can one construct robust, reliable inferences from the data.

Bayesian approaches to data analysis in general are often criticized for the lack of objectivity associated with prior information [e.g., Efron, 1986]. Those criticisms apply as well to Bayesian applications of the electromagnetic inverse problem, and it is essential
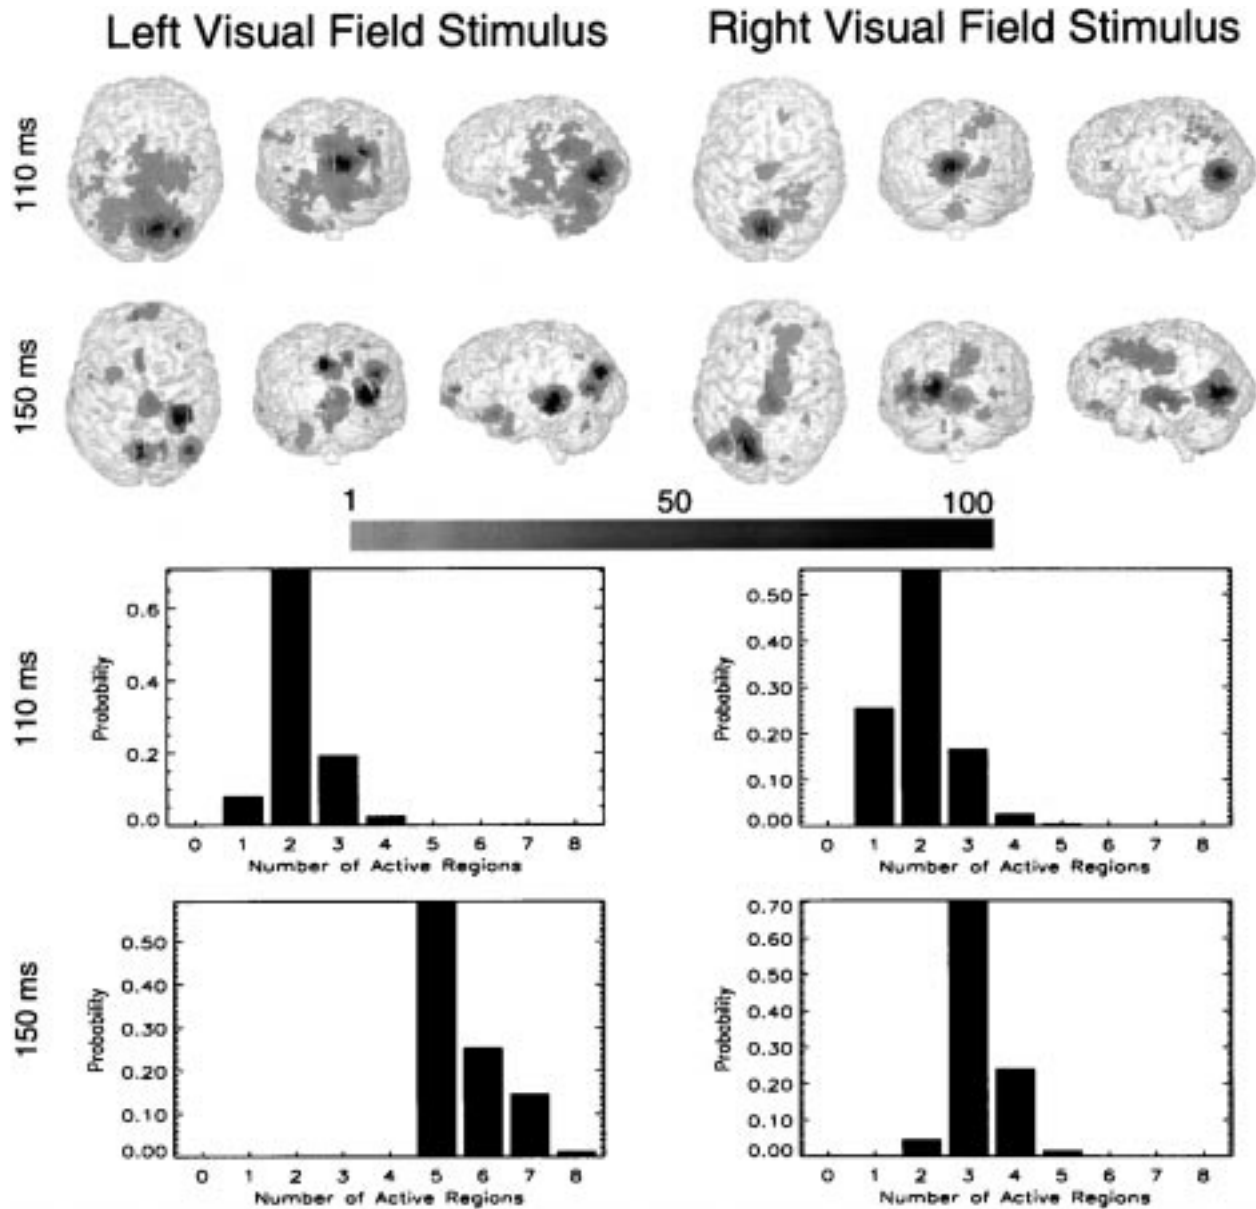
Figure 14.

Posterior probability distributions for the real MEG data of Example 3, marginalized onto anatomical location and onto number of active regions. The probability of activation as a function of location distributions are shown as maximum intensity projections over surface renderings of anatomy in the top half of the figure. The color bar shows the mapping between shades of grey and probability. The distributions for the number of active regions are shown in the bottom half of the figure. These results show evidence for activation contralateral to the stimulus at both 110 ms and 150 ms latencies.

to attempt to justify both the choice of the source model and the nature and specific values of the prior information as thoroughly and rigorously as possible. However, it is important to note that any attempt to solve the electromagnetic inverse problem forces the investigator to make analogous assumptions, even though they are rarely explicit. For example, widely used inverse approaches such as dipole models, minimum norm, FOCUSS [Gorodnitsky et al., 1995], or LORETA [Pascual-Marqui et al., 1994] all require restrictive assumptions regarding the nature and form of the allowable current distributions. A
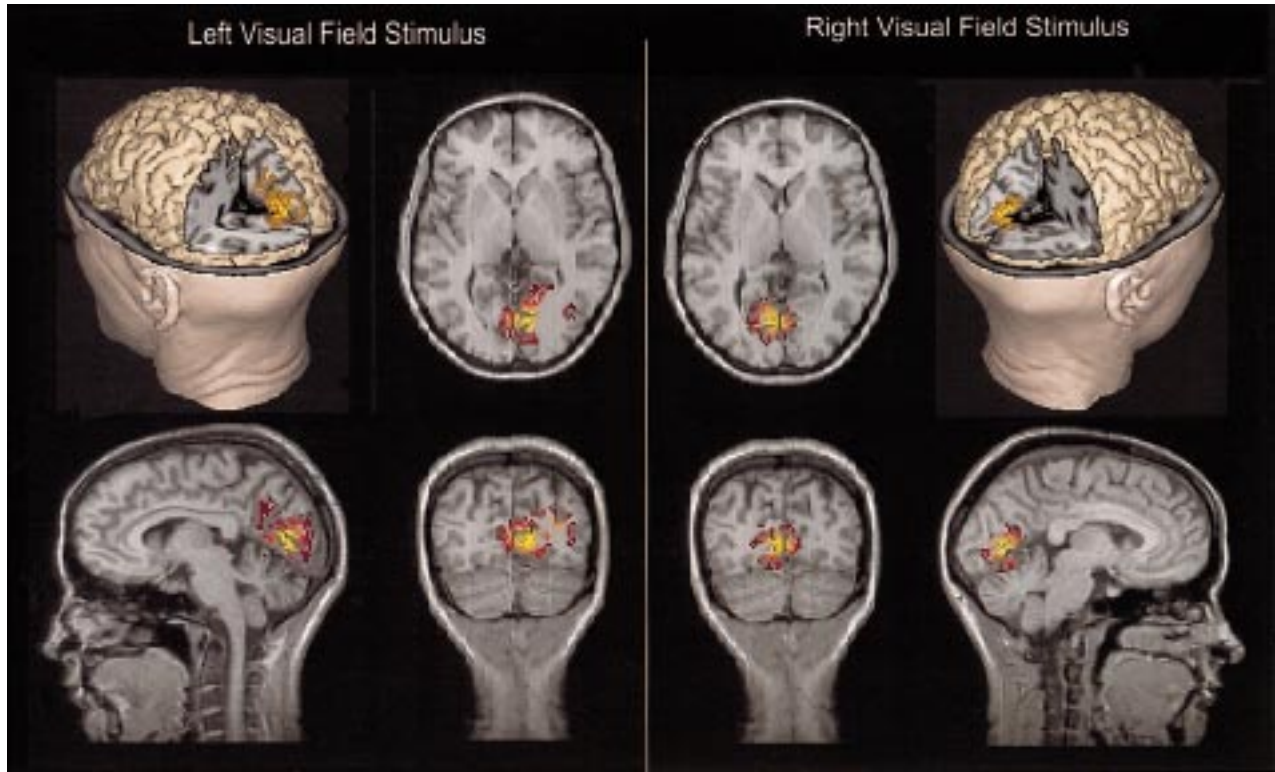
Figure 15.

Four views of a region that was found to contain activity at a 95% probability level in Example 3, for both a left and a right visual field stimulus, at 110 ms latency. The two-dimensional views show the regions (in color) within the anatomical MRI data (grey-scale). Shades of color represent relative probability within the regions on a temperature-like scale; bright yellow represents the highest probability. The horizontal and coronal views are from the top and from the back of the subject, respectively; the sagital views are from the left for the left visual field stimulus panel and from the right for the right visual field stimulus panel. The three-dimensional views are useful for showing the location of the regions relative to other brain structures. These results indicate that the probability of activity is maximal in the calcarine region of the hemisphere contralateral to the visual field stimulated.

Bayesian approach: (1) generalizes this strategy by weighting the possible current distributions on a probabilistic continuum instead of restricting the possibilities to those that are allowed, and (2) requires that the assumptions and prior information be made explicit and their associated prior probability distributions be justified explicitly. This formal, explicit treatment of prior information in Bayesian approaches is, therefore, a useful general feature for applications to inverse problems.

Finally, we emphasize that the activity model and examples described here are meant to illustrate the techniques and capabilities of BI in EEG/MEG and that other activity models, conductivity models, or sets of parameters of interest might be more appropriate for different experimental conditions. Our major objective is to present and illustrate the value of the Bayesian inferential approach, not to argue for the universal applicability of the particular activity model and prior information employed. Nevertheless, we believe the activity model described here is useful for many functional imaging applications and can be readily extended in a number of ways. These include incorporating temporal prior information in the form of temporal covariance constraints [e.g., Dale and Sereno, 1993], or explicit temporal models for evoked response studies. In addition, the Bayesian approach provides a natural means for incorporating information from other functional imaging modalities such as PET or fMRI [George et al., 1995; Belliveau, 1997; Dale, 1997]. The latter can be readily achieved with the Bayesian framework and with this activity model by assigning prior probabilities to possible locations of active regions based on results from the other modality

or modalities. Such a Bayesian formulation of multimodality integration would yield an inherently probabilistic result in which the quantity estimated would be the probability of activation as a function of both space and time.

ACKNOWLEDGMENTS

REFERENCES

Ahonen AI, Hämäläinen MS, Kajola MJ, Knuutila JET, Laine PP, Lounasmaa OV, Parkkonen LT, Simola JT, Tesche CD. 1993. 122-channel SQUID instrument for investigating the magnetic signals from the human brain. Physica Scripta T49A:198–205.

Aine CJ. 1995. A conceptual overview and critique of functional neuroimaging techniques in humans: I. MRI/fMRI and PET. Critical Reviews Neurobiol 9:229–309.

Aine C, Supek S, George J, Ranken D, Lewine J, Sanders J, Best E, Tiee W, Flynn E, Wood C. 1996. Retinotopic organization of human visual cortex: departures from the classical model. Cerebral Cortex 6:354–361.

Aine C, Chen HW, Ranken D, Mosher J, Best E, George J, Lewine J, Paulson K. 1997a. An examination of chromatic/achromatic stimuli presented to central/peripheral visual fields: an MEG study. Neuroimage 5:153.

Aine C, Schlitt H, Shah J, Ranken D, Krause B, George J, Müller-Gärtner HW. 1997b. MEG/fMRI differentiates between ventral/dorsal processing streams. Soc Neurosci 23:1312.

Allison T, Wood CC, McCarthy G. 1986. The central nervous system. In: Donchin E, Porges S, Coles M, editors. Psychophysiology: systems, processes and applications: a handbook. New York: Guilford, p 5–25.

Baillet S, Garnero L. 1997. A bayesian approach to introducing anatomo-functional priors in the EEG/MEG inverse problem. IEEE Trans Biomed Eng 44:374–385.

Belliveau J. 1997. Dynamic human brain mapping using combined fMRI, EEG and MEG. Presentation to the symposium on approaches to cognitive neuroscience by means of functional brain imaging, Caen, France.

Bernardo JM, Smith AFM. 1994. Bayesian theory. New York: John Wiley & Sons.

Dale AM, Sereno MI. 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. J Cognitive Neurosci 5:162–176.

Dale AM. 1997. Strategies and limitations in integrating brain imaging and electromagnetic recording. Soc Neurosci Abstr, 23:1.

Efron B. 1986. Why isn't everyone a Bayesian? Am Statistician 40:1–5; commentary 5–11.

Gelman A, Carlin JB, Stern HS, Rubin DB. 1995. Bayesian data analysis. London: Chapman & Hall.

George JS, Aine CJ, Mosher JC, Schmidt DM, Ranken DM, Schlitt HA, Wood CC, Lewine JD, Sanders JA, Belliveau JW. 1995. Mapping function in the human brain with MEG, anatomical MRI, and functional MRI. J Clin Neurophys 12:406–431.

Gorodnitsky IF, George JS, Rao BD. 1995. Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. Electroenceph Clin Neurophysiol 95:231–251.

Hämäläinen M, Ilmoniemi R. 1984. Interpreting measured magnetic fields of the brain: estimates of current distributions. Technical Report TKK-F-A559, Helsinki University of Technology, Finland.

Hämäläinen M, Hari R, Ilmoniemi RJ, Knuutila J, Lounasmaa OV. 1993. Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain. Rev Mod Phys 65:413–497.

Horton JC, Hoyt WF. 1991. The representation of the visual field in human striate cortex. Proc Roy Soc London 132:348–361.

Kavanaugh RN, Darcey TM, Lehmanh D, Fender DH. 1978. Evaluation of methods for three-dimensional localization of electrical sources in the human brain. IEEE Trans Biomed Eng 25:421–429.

Mitzdorf U. 1985. Current source-density method and application in cat cerebral-cortex: investigation of evoked-potentials and EEG phenomena. Physiological Rev 65:37–100.

Mosher JC, Lewis PS, Leahy RM. 1992. Multiple dipole modeling and localization from spatio-temporal MEG data. IEEE Trans Biomed Eng 39:541–557.

Nunez P. 1981. Electrical fields of the brain: neurophysics of the EEG. Oxford: Oxford University Press.

Okada YC, Papuashvili N, Xu C. 1998. Maximum current dipole moment density as an important physiological constraint in MEG inverse solutions. Proceedings of the Tenth International Conference on Biomagnetism.

Pascual-Marqui RD, Michel CM, Lehmann D. 1994. Low resolution electromagnetic tomography: a new method for localizing electrical activity of the brain. Int J Psychophysiol 18:49–65.

Phillips JW, Leahy RM, Mosher JC. 1997. MEG-based imaging of focal neuronal current sources. IEEE Trans Med Imag 16:338–348.

Ranken DM, George JS. 1993. MRIVIEW: an interactive computational tool for investigation of brain structure and function. Visualization '93, p. 324–331, IEEE Computer Society.

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

Scherg M, von Cramon D. 1986. Evoked dipole source potentials of the human auditory cortex. Electroenceph Clin Neurophysiol 65:344–360.

Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RBH. 1995. Borders of multiple visual areas in human revealed by functional magnetic resonance imaging. Science 268:889–893.

Shah NJ, Aine C, Schlitt H, Krause BJ, Ranken D, George J, Müller-Gärtner HW. 1998. Activation of the ventral stream in humans: an fMRI study. In: Gulyas B, Müller-Gärtner HW, editors. Positron emission tomography: a critical assessment of recent trends. Dordrecht: Kluwer.

Toga AW, Mazziotta JC. 1996. Brain mapping: the methods. New York: Academic Press.

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

Wang JZ, Williamson SJ, Kaufman L. 1992. Magnetic source images determined by a lead-field analysis: the unique minimum-norm. IEEE Trans Biomed Eng 39:665–675.