

A STATISTICAL APPROACH TO THE INVERSE PROBLEM IN MAGNETOENCEPHALOGRAPHY

BY ZHIGANG YAO^{1,*} AND WILLIAM F. EDDY[†]

*Ecole Polytechnique Fédérale de Lausanne** and *Carnegie Mellon University*[†]

Magnetoencephalography (MEG) is an imaging technique used to measure the magnetic field outside the human head produced by the electrical activity inside the brain. The MEG inverse problem, identifying the location of the electrical sources from the magnetic signal measurements, is ill-posed, that is, there are an infinite number of mathematically correct solutions. Common source localization methods assume the source does not vary with time and do not provide estimates of the variability of the fitted model. Here, we reformulate the MEG inverse problem by considering time-varying locations for the sources and their electrical moments and we model their time evolution using a state space model. Based on our predictive model, we investigate the inverse problem by finding the posterior source distribution given the multiple channels of observations at each time rather than fitting fixed source parameters. Our new model is more realistic than common models and allows us to estimate the variation of the strength, orientation and position. We propose two new Monte Carlo methods based on sequential importance sampling. Unlike the usual MCMC sampling scheme, our new methods work in this situation without needing to tune a high-dimensional transition kernel which has a very high cost. The dimensionality of the unknown parameters is extremely large and the size of the data is even larger. We use Parallel Virtual Machine (PVM) to speed up the computation.

1. Introduction.

1.1. *The basics of magnetoencephalography (MEG).* The anatomy of the brain has been studied intensively for millennia, yet how the brain functions is still not well understood. The neurons in the brain produce macroscopic electrical currents when the brain functions, and those synchronized neuronal currents in the gray matter of the brain induce extremely weak magnetic fields (10–100 femto-Tesla) outside the head. The comparatively recent development of the Superconducting Quantum Interference Device (SQUID) makes it possible to detect those magnetic signals. MEG is an imaging technique using SQUIDS to measure the magnetic signals outside of the head produced by the electrical activity inside the brain [Cohen (1968)]. The primary sources are electric currents within the

Received January 2012; revised January 2014.

¹Supported in part by NSF SES-1061387 and NIH/NIDA R90 DA023420.

Key words and phrases. Ill-posed problem, sequential importance sampling, state space model, parallel computing, source localization.

dendrites of the large pyramidal cells of activated neurons in the human cortex, generally formulated as a mathematical point current dipole. Such focal brain activation can be observed in epilepsy, or it can be induced by a stimulus in neurophysiological or neuropsychological experiments. Due to its noninvasiveness (it is a completely passive measurement method) and its impressive temporal resolution (better than 1 millisecond, compared to 1 second for functional magnetic resonance imaging, or to 1 minute for positron emission tomography), and due to the fact that the signal it measures is a direct consequence of neural activity, MEG is a near optimal tool for studying brain activity, such as assisting surgeons in localizing a pathology, assisting researchers in determining brain function, neurofeedback and others. The skull and the tissue surrounding the brain affect the magnetic fields measured by MEG much less than the electrical impulses measured by electroencephalography (EEG). This means that MEG has higher localization accuracy than the EEG and it allows for a more reliable localization of brain function [Hämäläinen et al. (1993), Okada, Lähteenmäki and Xu (1999)]. MEG has recently been used in the evaluation of epilepsy, where it reveals the exact location of the abnormalities, which may then allow physicians to find the cause of the seizures [Barkley and Baumgartner (2003)]. MEG is also reference free, so that the localization of sources with a given precision is easier for MEG than it is for EEG [Kristeva-Feige et al. (1997)]. The computation associated with estimating the electric source from the magnetic measurement is a challenging problem that needs to be solved to allow high temporal and spatial resolution imaging of the dynamic activity of the human brain.

1.2. Forward and inverse MEG problem. The MEG signals derive from the *primary* current (the net effect of ionic currents flowing in the dendrites of neurons) and the *volume* current (i.e., the additive ohmic current set up in the surrounding medium to complete the electrical circuit). If the electrical source is known and the head model [Kybic et al. (2006)] is specified (e.g., a sphere with homogeneous conductivity), then the “forward problem” is to compute the electric field \mathbf{E} and the magnetic field \mathbf{B} from the source current \mathbf{J} . The calculation uses Maxwell’s equations [see, e.g., Griffiths (1999)],

$$\begin{aligned}\nabla \cdot \mathbf{E} &= \rho / \varepsilon_0, \\ \nabla \times \mathbf{E} &= -\partial \mathbf{B} / \partial t, \\ \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{B} &= \mu_0 (\mathbf{J} + \varepsilon_0 \partial \mathbf{E} / \partial t),\end{aligned}$$

where ε_0 and μ_0 are the permittivity and permeability of a vacuum, respectively, and ρ is the charge density. The total current \mathbf{J} consists of the primary current \mathbf{J}^P plus the volume current \mathbf{J}^V . The source activity in the brain corresponds to the primary current. Under reasonable assumptions [see Hämäläinen et al. (1993)], the

volume current \mathbf{J}^V is not included in the analysis because of its diffuse nature. The terms $\partial \mathbf{B}/\partial t$ and $\partial \mathbf{E}/\partial t$ in Maxwell's equations can be ignored by assuming that the magnetic field varies relatively slowly in time. Rather than working with continuous electric current, the most frequently used computational model assumes that the electric current can be thought of as an electric dipole; this model is called an equivalent current dipole (ECD); see, for example, [Hämäläinen et al. \(1993\)](#). From the perspective of an ECD, a dipole has location, orientation and magnitude; the magnetic field generated by this dipole can explain the MEG measurements. In addition, there is a version of an ECD model assuming multiple dipoles; from Maxwell's equations it is easy to see that this model is simply the sum of the models for each ECD. Such an ECD models a large number of dipoles located at fixed places over the cortical surface. In neuroscience, it is believed that typical MEG data should be explained by only a few dipoles (less than 10), and different criteria or algorithms are used to minimize the number of dipoles in various ECD models; we discuss some of these models in Section 1.3. We assume that \mathbf{E} is generated by \mathbf{J}^P , which in turn comes from the sum of N localized current dipoles at locations \mathbf{r}_n ,

$$\mathbf{J}_n^P(\mathbf{r}) = Q_n \delta(\mathbf{r} - \mathbf{r}_n), \quad n = 1, \dots, N,$$

where $\delta(\cdot)$ is the Dirac delta function. The Q_n is a charged dipole at the point \mathbf{r}_n in the brain volume Ω . Using the quasi-static approximation to Maxwell's equations (i.e., ignoring the partial derivatives with respect to time) given in [Sarvas \(1984\)](#), the magnetic field \mathbf{B} at location \mathbf{r} of a current dipole at \mathbf{r}_n can be calculated by the Biot–Savart equation,

$$\mathbf{B}_n(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_{\Omega} \frac{\mathbf{J}_n^P(\mathbf{r}_n) \times (\mathbf{r} - \mathbf{r}_n)}{|\mathbf{r} - \mathbf{r}_n|^3} d\mathbf{r}_n.$$

In the case of multiple current dipoles, the induced magnetic fields simply add up.

The “inverse problem” comes from the forward model; we want to estimate the dipole parameters from the observed magnetic signal. The difficulty is that there is not a unique solution; there are infinitely many different sources within the skull that produce the same observed data. The goal is to find a meaningful solution among the many mathematically correct solutions. There are three key steps to any source localization algorithm in MEG. First, define the solution space and the parameter space of the electric source. Second, calculate the magnetic field given the information about the head model. Third, according to what criterion the solution must satisfy, perform an extensive search for the solution. Methods of finding the source from the observed MEG signal have been extensively exploited during the past two decades, mostly centered on finding a single estimate of the source. Some of these methods are briefly described in the next subsection. However, finding the distribution of the source in space and (particularly) in time is still a problem requiring investigation.

1.3. *Existing source localization methods.* Most methods for localizing electrical sources in MEG assume that the electrical sources in the brain do not include a temporal component. The data are used to estimate the source parameters at each time point; there is no relation to the estimates for the previous time. This is not the same as assuming the quasi-static approximation to Maxwell's equations. Therefore, those existing methods are restricted to fixed dipole assumptions and are also not able to provide estimates of the variability of source activity. The minimum norm estimate (MNE) [Hämäläinen and Ilmoniemi (1994)] is a regularization method based on the L_2 norm. The L_1 norm regularization yields the minimum current estimates (MCE) [Uutela, Hämäläinen and Somersalo (1999)]. The LORETA approach [Mattout et al. (2006)] is a special case of weighted MNE. The Multiple Signal Classification method (MUSIC) [Mosher, Lewis and Leahy (1992)] searches for a single-dipole model through a three-dimensional head volume and computes projections onto an estimated signal subspace. The source locations are then found as the 3-D locations where the source model gives the best projections onto the subspace. The beamformer methods [Van Veen, Joseph and Hecox (1992)] ignore the ill-posed inverse problem and instead only estimate the current at several fixed locations. Bayesian approaches to the MEG inverse problem try to find the posterior distribution of the dipole parameters [Bertrand et al. (2001), Schmidt, George and Wood (1999)].

The methods mentioned briefly above (MNE, MCE, etc.) have been widely used and produce apparently meaningful solutions; however they have overly restrictive assumptions and lack estimates of the variability of source estimates. By assuming a static localized dipole, these methods are limited in their ability to incorporate problem-specific anatomical or physiological information. It is quite reasonable to consider that the source is time-varying rather than fixed, in which case the noise reduction obtained by averaging over consecutive observations in time is problematic. By utilizing a time-varying source model, we will be able to investigate the distribution of the source at each time point and provide estimates of the variability. Following this idea, the time evolution of the source is modeled by a state space model. Our goal is to find the posterior distribution of the source parameters. Our reformulation of the inverse problem leads to a predictive model of the dipole. It turns out that the posterior source distribution from our predictive model can be interpreted as a statistical solution to the MEG inverse problem.

1.4. *Outline of this paper.* In Section 2 we develop a time-varying source model for the MEG inverse problem. Rather than attempting to “solve” the inverse problem we try to develop estimates of the dipole parameters using a spatio-temporal model. In Section 3 the difficulty of using Markov chain Monte Carlo (MCMC) methods for generating samples from the time-varying model is explained. Then, we introduce the standard Sequential Importance Sampling (SIS) technique. Next, two further Monte Carlo methods are described: (1) the regular SIS method with rejection, and (2) the improved SIS method with resampling.

A simulation study is described in Section 4. We describe our use of the Parallel Virtual Machine (PVM) software to speed up the computations in Section 4.2. We believe this is the first application of parallel computational methods to the problem. In Section 5 a real data application is presented. Section 6 contains a short discussion and conclusion.

2. A probabilistic time-varying source model. Assume that the magnetic field data is measured from the k th sensor at time t as

$$Y_{k,t} = \mathbf{B}_k(\mathbf{J}_t^P) + \mathbf{U}_{k,t}, \quad 1 \leq t \leq T, 1 \leq k \leq L,$$

where $\mathbf{U}_{k,t} \sim N(0, \sigma_1^2)$ denotes the observation noise that is assumed, for simplicity, to be Gaussian, additive and homogeneous for all the sensors, and uncorrelated between every pair of sensors. The assumption of normality is preferred due to the fact that Gaussian sensor noise is present at the MEG sensors themselves, and sensor noise is typically substantially smaller than signals from spontaneous brain activity [Hämäläinen et al. (1993)]. Although correlated sensor noise is more realistic than “homogeneous” sensors, it complicates the problem. Background noise and biological noise can also drown out the brain activity of interest, but these are all very difficult to incorporate. Besides some variation coming from solving the inverse problem, most of the variation of the source localization in MEG is due to the propagation of errors through Maxwell’s equations when solving the forward problem. In order to control variation, we only work with a simple sensor structure. Therefore, we write

$$\mathbf{Y}_t = \mathbf{B}(\mathbf{J}_t^P) + \mathbf{U}_t, \quad 1 \leq t \leq T,$$

where $\mathbf{Y}_t = (Y_{1,t}, \dots, Y_{L,t})^T$, $\mathbf{B}(\mathbf{J}_t^P) = (\mathbf{B}_1(\mathbf{J}_t^P), \dots, \mathbf{B}_L(\mathbf{J}_t^P))^T$ and $\mathbf{U}_t = (\mathbf{U}_{1,t}, \dots, \mathbf{U}_{L,t})^T$. Here, $\mathbf{U}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_1)$, where $\mathbf{\Sigma}_1 = \text{diag}[\sigma_1^2, \dots, \sigma_1^2]$.

The $\mathbf{B}_k(\mathbf{J}_t^P)$, a function of the dipole with parameter vector \mathbf{J}_t^P , is the physical approximation of the Biot–Savart law in Section 1.2. We consider a current dipole located within a horizontally layered conductor [Hämäläinen et al. (1993)]. The noiseless magnetic field at the k th sensor, \mathbf{B}_k , is computed from the source $\mathbf{J}_t^P = (\mathbf{p}_t, \mathbf{q}_t)$ at time t . The vector $\mathbf{p}_t = (p_{1t}, p_{2t}, p_{3t})$ contains the location parameters of the source and the vector $\mathbf{q}_t = (q_{1t}, q_{2t}, q_{3t})$ contains the moments and strength. Thus,

$$(1) \quad \mathbf{B}_k(\mathbf{J}_t^P) = \frac{\mu_0}{4\pi} \frac{\mathbf{q}_t \times (\mathbf{r}_k - \mathbf{p}_t) \cdot \mathbf{e}}{|\mathbf{r}_k - \mathbf{p}_t|^3}.$$

Here, \mathbf{r}_k is the location of the k th sensor. Because the magnetometers measure only the z direction of the magnetic field \mathbf{B} , $\mathbf{e} = (0, 0, 1)$, a unit vector, is used to find the z component of \mathbf{B} . Conventionally, z is perpendicular to the surface of the skull.

To specify the *prior*, the time evolution of the current dipole \mathbf{J}_t^P is modeled by a state space model. A number of other authors have also proposed stating the MEG

inverse problem as a Bayesian dynamic model; see Somersalo, Voutilainen and Kaipio (2003); Campi et al. (2008, 2011); Sorrentino et al. (2009, 2013); Miao et al. (2013). We could choose any of a large variety of state space models but, for theoretical and computational simplicity, we have chosen a six-dimensional first-order autoregression:

$$\mathbf{J}_t^P = \mathbf{m}_{\text{com}} + \boldsymbol{\rho}(\mathbf{J}_{t-1}^P - \mathbf{m}_{\text{com}}) + \mathbf{V}_t, \quad 1 \leq t \leq T,$$

where $\mathbf{V}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_2)$ denotes the state evolution noise. We note that three of the parameters give the spatial location, so this is implicitly a space–time model. Previous work using a space–time model [Ou, Hämäläinen and Golland (2009)] used a novel mixed $L_1 L_2$ -norm estimate for the dipole parameters based on a linear regression model. Jun et al. (2005) have also used MCMC methods for sampling from the posterior of a spatiotemporal Bayesian dynamic model. To reduce the number of parameters, and hence the amount of variation in our estimates, we assume the dipole parameters are uncorrelated. That is, we assume that $\boldsymbol{\Sigma}_2 = \text{diag}[\sigma_{11}^2, \sigma_{22}^2, \dots, \sigma_{66}^2]$ is a known 6 by 6 diagonal matrix and σ_{ii}^2 is the variance of the i th source parameter. The parameter vector \mathbf{m}_{com} is a constant associated with the source \mathbf{J}_t^P for $1 \leq t \leq T$. The initial state is chosen as $\mathbf{J}_0^P \sim \mathcal{N}(\mathbf{m}_{\text{ini}}, \boldsymbol{\Sigma}_2)$, where \mathbf{m}_{ini} is also a constant parameter vector. Both \mathbf{m}_{ini} and \mathbf{m}_{com} are specified in advance. The known diagonal matrix $\boldsymbol{\rho} = \text{diag}[\rho_1, \rho_2, \dots, \rho_6]$ is 6 by 6. Its main diagonal represents the autoregressive coefficients. Hence, at any time t , \mathbf{J}_t^P or $(\mathbf{p}_t, \mathbf{q}_t)$ contains the parameters of interest and $\mathbf{Y}_t = (Y_{1,t}, \dots, Y_{L,t})$ is the (very noisy) data collected from all sensors. Both $\{\mathbf{J}_t^P\}_{t=0}^T$ and $\{Y_{k,t}\}_{t=1}^T$ are assumed to have the following Markov properties:

(i) The \mathbf{J}^P is a first order Markov process. The distribution of each state \mathbf{J}_t^P only depends on its own previous state \mathbf{J}_{t-1}^P ,

$$(2) \quad p(\mathbf{J}_t^P | \mathbf{J}_0^P, \mathbf{J}_1^P, \dots, \mathbf{J}_{t-1}^P) = p(\mathbf{J}_t^P | \mathbf{J}_{t-1}^P)$$

(we are using p as a generic symbol for a probability distribution; the two p 's in this equation are not the same function).

(ii) The process $Y_{k,t}$ (for any $1 \leq k \leq L$) is also a Markov process with respect to the history of \mathbf{J}^P . The density of $Y_{k,t}$ conditioned on $\{\mathbf{J}_t^P\}_0^t$ satisfies

$$f(Y_{k,t} | \mathbf{J}_0^P, \mathbf{J}_1^P, \dots, \mathbf{J}_t^P) = f(Y_{k,t} | \mathbf{J}_t^P)$$

(again f is a generic symbol, in this case, for a joint density function).

(iii) When conditioned on its own history, the unknown \mathbf{J}_t^P does not depend on past measurements. The distribution of \mathbf{J}_t^P based on $\mathbf{Y}^k = (Y_{k,1}, \dots, Y_{k,t-1})$ and \mathbf{J}_{t-1}^P is

$$(3) \quad g(\mathbf{J}_t^P | \mathbf{J}_{t-1}^P, \mathbf{Y}^k) = p(\mathbf{J}_t^P | \mathbf{J}_{t-1}^P), \quad t > 0$$

[the right-hand side of equation (3) in (iii) is the same as the right-hand side of equation (2) in (i)]. The transition kernel, $p(\mathbf{J}_t^P | \mathbf{J}_{t-1}^P)$, is defined here as a first

order Markov process in the state space model above. For a more complex model it could be a higher order Markov process. The choice of more realistic models for this process [e.g., in the situation where the magnetic signal is a response to a stimulus, the source variance might change much more rapidly immediately after the stimulus than before it; the joint density $f(Y_{k,t}|\mathbf{J}_t^P)$ for any $1 \leq k \leq L$ may also vary in time since not all the measurements can be carried out simultaneously] is not the aim of this paper.

Of interest at any time t is the posterior distribution of $\mathcal{J}_t^P = (\mathbf{J}_0^P, \dots, \mathbf{J}_t^P)$. Let $\mathcal{Y}_{\text{obs}}^t = (\mathbf{Y}_1, \dots, \mathbf{Y}_L) = (Y_{1,1}, \dots, Y_{1,t}, \dots, Y_{L,1}, \dots, Y_{L,t})$ be the magnetic measurements, accordingly. By taking all the previous *prior* information and the three assumptions [(i), (ii), (iii)] above into account, our problem can be stated as finding the target distribution, $p(\mathcal{J}_t^P|\mathcal{Y}_{\text{obs}}^t)$, given $\mathcal{Y}_{\text{obs}}^t$. By Bayes' theorem, we have

$$(4) \quad p(\mathcal{J}_t^P|\mathcal{Y}_{\text{obs}}^t) \propto f(\mathcal{Y}_{\text{obs}}^t|\mathcal{J}_t^P)p(\mathcal{J}_t^P) \\ = \left[\prod_{s=1}^t \prod_{k=1}^L f(Y_{k,s}|\mathbf{J}_s^P) \right] \left[\prod_{s=1}^t p(\mathbf{J}_s^P|\mathbf{J}_{s-1}^P) \right] p(\mathbf{J}_0^P).$$

This framework is based on a one-source model ($N = 1$). It can be easily extended to a multiple-source model because the fields generated by distinct sources simply add up. Because it is a high-dimensional distribution ($1 \leq t \leq T$, T is very large) and inherently complicated, sampling from the posterior is difficult. We have chosen to use MCMC methods but they are also complex and are very hard to implement. As we will show later, obtaining $p(\mathcal{J}_t^P|\mathcal{Y}_{\text{obs}}^t)$ can be achieved dynamically by computing the $p(\mathbf{J}_u^P|\mathcal{Y}_{\text{obs}}^u)$ at each time point $1 \leq u \leq t$. These calculations have to be repeated for each $t \leq T$.

3. Solving the MEG inverse problem.

3.1. The difficulty of solving the time-varying model. A problem with MCMC methods (e.g., Metropolis–Hastings) for getting joint posterior samples from $p(\mathcal{J}_t^P|\mathcal{Y}_{\text{obs}}^t)$ occurs when there are a large number of states because it is difficult to find a joint transition kernel which could be used in an MCMC sampler. However, the goal of getting $p(\mathcal{J}_t^P|\mathcal{Y}_{\text{obs}}^t)$ can be achieved by sampling from the distribution $p(\mathbf{J}_s^P|\mathcal{Y}_{\text{obs}}^s)$ for each state s ($1 \leq s \leq t$) separately and the entire outcome could be regarded as the sample from the joint distribution. Gibbs sampling can be used for this restricted goal, but because of the nonlinearity of the model [equation (1)], it is not easy to sample from $p(\mathbf{J}_t^P|\mathbf{J}_{s \neq t}^P, \mathcal{Y}_{\text{obs}}^t)$. One way to alleviate the difficulty is to insert some kind of Metropolis sampler into a Gibbs sampling scheme for each conditional distribution. When we insert a random-walk Metropolis algorithm, where the move depends only on its own state, into the Gibbs sampler, we call it a *random-walk MCMC within Gibbs* sampler, and when we insert a hybrid Metropolis algorithm, where the move may depend on other states, into the Gibbs sampler, we call it a *hybrid MCMC within Gibbs* sampler.

The key to *random-walk MCMC within Gibbs* is to propose a candidate $\mathbf{J}_t^{\text{P}*} \sim \mathcal{N}(\mathbf{J}_t^{\text{P}}, \boldsymbol{\Sigma}_3)$ for each t ($1 \leq t \leq T$), where $\boldsymbol{\Sigma}_3 = \text{diag}[\tau_1^2, \tau_2^2, \dots, \tau_6^2]$ is a 6 by 6 diagonal matrix, and accept $\mathbf{J}_t^{\text{P}*}$ if the acceptance ratio

$$\alpha_t = \frac{\prod_{k=1}^L f(Y_{k,t} | \mathbf{J}_t^{\text{P}*}) p(\mathbf{J}_t^{\text{P}*} | \mathbf{J}_{t-1}^{\text{P}}) p(\mathbf{J}_{t+1}^{\text{P}} | \mathbf{J}_t^{\text{P}*})}{\prod_{k=1}^L f(Y_{k,t} | \mathbf{J}_t^{\text{P}}) p(\mathbf{J}_t^{\text{P}} | \mathbf{J}_{t-1}^{\text{P}}) p(\mathbf{J}_{t+1}^{\text{P}} | \mathbf{J}_t^{\text{P}})} \geq \mathcal{U}(0, 1),$$

where $\mathcal{U}(0, 1)$ is the uniform distribution. The problem is that $\mathcal{N}(\mathbf{J}_t^{\text{P}}, \boldsymbol{\Sigma}_3)$ is not a good proposal for $\mathbf{J}_t^{\text{P}*}$ (i.e., we almost always reject the proposal) and this can not be solved by extensively tuning $\boldsymbol{\Sigma}_3 = \text{diag}[\tau_1^2, \tau_2^2, \dots, \tau_6^2]$ in most practical cases if the dimension of the states is very high. A local linear approximation might be considered, such as performing a Taylor expansion on the joint density function $f(Y_{k,t} | \mathbf{J}_t^{\text{P}})$ and truncating high order terms. The resultant can then be incorporated into the proposal distribution. However, such linearization is not easy due to the highly complex function $f(Y_{k,t} | \mathbf{J}_t^{\text{P}})$; moreover, the extra work of a Taylor expansion might be unnecessary if we only need an efficient sampling scheme in high dimensions.

The *hybrid MCMC within Gibbs* improves upon the random-walk MCMC within Gibbs when the target distribution is not able to be captured by a simple random walk. In Gelman, Roberts and Gilks (1996), a full conditional prior (hybrid MCMC) was proposed. Similar work can also be found in Carter and Kohn (1994), where a single move blocking strategy was developed but bad convergence behavior was discovered. Gamerman (1998) suggested a reparameterization of the model to a prior independent system of disturbances and built a proposal by a weighted least squares algorithm, however, the reparameterization resulted in quadratic computational time. Knorr-Held (1999) suggested an autoregressive prior where the “conditional prior” is drawn independently of the current state but, in general, depends on other states. Here, our hybrid MCMC within Gibbs is built on a single move proposal, that is, $\mathbf{J}_t^{\text{P}*}$ is proposed from the distribution of $p(\mathbf{J}_t^{\text{P}} | \mathbf{J}_{s \neq t}^{\text{P}})$ which can be further reduced to $p(\mathbf{J}_t^{\text{P}} | \mathbf{J}_{t-1}^{\text{P}}, \mathbf{J}_{t+1}^{\text{P}})$ due to the Markov property. One way to update \mathbf{J}_t^{P} is to use a proposal

$$\mathbf{J}_t^{\text{P}*} \sim \mathcal{N}(\boldsymbol{\rho}(\mathbf{J}_{t-1}^{\text{P}} - \mathbf{J}_{t+1}^{\text{P}}) + (\mathbf{I} - \boldsymbol{\rho}\boldsymbol{\rho}')(\mathbf{I} + \boldsymbol{\rho}\boldsymbol{\rho}')^{-1} \mathbf{m}_{\text{com}}, \boldsymbol{\Sigma}_2(\mathbf{I} + \boldsymbol{\rho}\boldsymbol{\rho}')^{-1}).$$

The acceptance ratio then reduces to

$$\alpha_t = \frac{\prod_{k=1}^L f(Y_{k,t} | \mathbf{J}_t^{\text{P}*})}{\prod_{k=1}^L f(Y_{k,t} | \mathbf{J}_t^{\text{P}})}.$$

The performance of a single move could be extended to a block move by sampling a block of states at the same time based on other states. Similarly, the $\mathbf{J}_r^{\text{P}*}, \dots, \mathbf{J}_s^{\text{P}*}$ come from the conditional proposal

$$p(\mathbf{J}_r^{\text{P}}, \dots, \mathbf{J}_s^{\text{P}} | \mathbf{J}_{1,\dots,T}^{\text{P}} / (\mathbf{J}_r^{\text{P}}, \dots, \mathbf{J}_s^{\text{P}})),$$

where $r < s$ and $\mathbf{J}_{1,\dots,T}^P/(\mathbf{J}_r^P, \dots, \mathbf{J}_s^P)$ means a collection of $\mathbf{J}_1^P, \dots, \mathbf{J}_{r-1}^P, \mathbf{J}_{s+1}^P, \dots, \mathbf{J}_T^P$. Thus, the acceptance ratio becomes

$$\alpha_t = \frac{\prod_{k=1}^L \prod_{t=r}^s f(Y_{k,t} | \mathbf{J}_t^{P*})}{\prod_{k=1}^L \prod_{t=r}^s f(Y_{k,t} | \mathbf{J}_t^P)}.$$

Although the block move provides a considerable improvement in the situation where a single move has poor mixing behavior, [Carter and Kohn \(1994\)](#) observed bad mixing and convergence behavior in the blocking strategy.

Recently developed adaptive samplers [[Andrieu and Thoms \(2008\)](#), [Roberts and Rosenthal \(2009\)](#)] might help find the transition kernel within a Gibbs sampler, but these methods are computationally inefficient in high dimension. In addition, although parallel tempering [[Srinivasan \(2002\)](#)] seems reasonable, finding the temperature is not straightforward and significantly increases the computational cost. Again, the MEG data set is extremely large; in particular, we collect hundreds of channels of data at each time and we collect data for hundreds of thousands of time points. It is quite difficult to implement these methods since even a simple model has an extremely large number of states. The computational burden is even more substantial in the multiple-dipole case.

3.2. Sequential importance sampling (SIS). Sequential importance sampling (SIS) [[Liu and Chen \(1998\)](#)] is advocated as a more practical tool for a dynamic system. As we mentioned briefly in Section 2, computing $p(\mathbf{J}_u^P | \mathcal{Y}_{\text{obs}}^u)$ sequentially in u for $1 \leq u \leq t$ can lead to $p(\mathcal{J}_t^P | \mathcal{Y}_{\text{obs}}^t)$. Consider $\pi_t(\mathbf{J}_t^P) = p(\mathbf{J}_t^P | \mathcal{Y}_{\text{obs}}^t)$; calculating $p(\mathcal{J}_t^P | \mathcal{Y}_{\text{obs}}^t)$ or, equivalently, $\pi_t(\mathcal{J}_t^P)$ can be achieved by performing the following two processes in sequential order:

$$(5) \quad \pi_t(\mathbf{J}_t^P) = \frac{f(\mathbf{Y}_t | \mathbf{J}_t^P) \pi_{t-1}(\mathbf{J}_t^P)}{\pi_{t-1}(\mathbf{Y}_t)},$$

$$(6) \quad \pi_t(\mathbf{J}_{t+1}^P) = \int p(\mathbf{J}_{t+1}^P | \mathbf{J}_t^P) \pi_t(\mathbf{J}_t^P) d\mathbf{J}_t^P,$$

where $f(\mathbf{Y}_t | \mathbf{J}_t^P) = \prod_{k=1}^L f(Y_{k,t} | \mathbf{J}_t^P)$ and \mathbf{Y}_t is defined in Section 2. The denominator $\pi_{t-1}(\mathbf{Y}_t)$ in equation (5) is a constant, $\int f(\mathbf{Y}_t | \mathbf{J}_t^P) \pi_{t-1}(\mathbf{J}_t^P) d\mathbf{J}_t^P$. Equation (5) computes the posterior density $\pi_t(\mathbf{J}_t^P)$ and equation (6) is the well-known Chapman–Kolmogorov equation, which allows us to compute the next prior density based on $p(\mathbf{J}_{t+1}^P | \mathbf{J}_t^P)$ [the initial $p(\mathbf{J}_0^P)$ is known]. For each t , most of the MCMC samples are either obtained from sampling the joint $\pi_t(\mathcal{J}_t^P)$ or some other distribution $g_t(\mathcal{J}_t^P)$ and applying an acceptance criterion. However, the random draws of $\pi_t(\mathcal{J}_t^P)$ are never used again when the system proceeds from π_t to π_{t+1} [[Carlin, Polson and Stoffer \(1992\)](#)]. In high dimensions, the posterior samples for each state will have larger variation between iterations and, hence, both convergence and computation problems arise. In contrast, the SIS is able to reuse the

current samples and help create the samples for the next iteration; that improves the computational efficiency and reduces the variation between iterations. For non-linear problems [e.g., nonlinearity of equation (1)] or non-Gaussian densities, SIS requires the use of numerical approximation techniques where the key idea is to represent an approximation to the target posterior distribution by a set of samples and their associated weights.

In practice, suppose a stream $S_t = \{(\mathcal{J}_t^P)^{(j)}, j = 1, \dots, m\}$ (m by t) is a set of random samples properly weighted by the set of weights $\{w_t^{(j)}, j = 1, \dots, m\}$ (m by 1) with respect to $\pi_t(\mathcal{J}_t^P)$ [this can be viewed as approximate posterior samples of $\mathcal{J}_t^P = (\mathbf{J}_1^P, \dots, \mathbf{J}_t^P)$]. Define $g_{t+1}(\mathbf{J}_{t+1}^P | (\mathcal{J}_t^P)^{(j)})$ as a trial function for \mathbf{J}_{t+1}^P ; the recursive SIS procedure produces a new stream S_{t+1} by drawing a new sample \mathbf{J}_{t+1}^P and updating its associated weight. This is summarized as follows:

Algorithm 1: SIS

- (i) Sample a new $(\mathbf{J}_{t+1}^P)^{(j)}$ from the trial distribution $g_{t+1}(\mathbf{J}_{t+1}^P | (\mathcal{J}_t^P)^{(j)})$ and form $(\mathcal{J}_{t+1}^P)^{(j)} = ((\mathcal{J}_t^P)^{(j)}, (\mathbf{J}_{t+1}^P)^{(j)})$.
 - (ii) Compute the incremental weight $u_{t+1}^{(j)} = \frac{\pi_{t+1}((\mathcal{J}_{t+1}^P)^{(j)})}{\pi_t((\mathcal{J}_t^P)^{(j)})g_{t+1}(\mathbf{J}_{t+1}^P | (\mathcal{J}_t^P)^{(j)})}$ and update the weight $w_{t+1}^{(j)} = u_{t+1}^{(j)} w_t^{(j)}$.
 - (ii*) Sample a new stream S'_{t+1} from the stream S_{t+1} based on the updated weights $w_{t+1}^{(j)}$.
 - (iii) Assign equal weights to the samples in S'_{t+1} .
-

It has been proved that the new samples and weights $((\mathcal{J}_{t+1}^P)^{(j)}, w_{t+1}^{(j)})$ are properly weighted samples from π_{t+1} [Liu and Chen (1998)]. As time t increases, a resampling scheme is inserted between adjacent times or one can just resample after the last time. This step is summarized in steps (ii*) and (iii). Shephard and Pitt (1997) showed that resampling [step (ii*)] is only necessary when the weights are very skewed; resampling reduces m and thus reduces the computational burden. A schedule for the resampling scheme in SIS is proposed in Gordon, Salmond and Smith (1993), Kitagawa (1996) and Liu (1996). The choice of trial distribution $g_{t+1}(\mathbf{J}_{t+1}^P | (\mathcal{J}_t^P)^{(j)})$ is crucial in SIS. Choosing $g_{t+1}(\mathbf{J}_{t+1}^P | (\mathcal{J}_t^P)^{(j)}) = \pi_t(\mathbf{J}_{t+1}^P | (\mathbf{J}_t^P)^{(j)})$ is much easier to implement, although it might bring greater variation [see Berzuini et al. (1997)]. This procedure ends up getting $g_{t+1}(\mathbf{J}_{t+1}^P | (\mathcal{J}_t^P)^{(j)}) = p(\mathbf{J}_{t+1}^P | (\mathbf{J}_t^P)^{(j)})$ and incremental weights $f(\mathbf{Y}_{t+1} | (\mathbf{J}_{t+1}^P)^{(j)}) = \prod_{k=1}^L f(Y_{k,t+1} | (\mathbf{J}_{t+1}^P)^{(j)})$. There exist in the literature several kinds of local Monte Carlo methods which could be embedded into SIS to get the weights or even approximate weights no matter what g_{t+1} function we choose. This strategy provides the opportunity to find relatively good weights that could be used in SIS so we can limit our attention to the choice of trial function when we apply SIS. The SIS procedure (Algorithm 1) was initially used in the analysis

of state-space models and is similar to sequential Monte Carlo (SMC) which has recently been applied as an alternative to MCMC for standard Bayesian inference problems [Neal (2001), Del Moral, Doucet and Jasra (2006), Fearnhead (2008)].

3.3. Regular SIS method with rejection. This algorithm [Liu and Chen (1998)] inserts the standard rejection method as a local Monte Carlo scheme into the SIS procedure. At step t , the rejection method is constructed based on sampling the joint distribution of (J, \mathbf{J}_{t+1}^P) . To do this, we draw $J = j$ with probability proportional to $w_t^{(j)}$. Given $J = j$, sample $(\mathbf{J}_{t+1}^P)^{(j)}$ from $p(\mathbf{J}_{t+1}^P | (\mathbf{J}_t^P)^{(j)})$. Next, compute the constant $c_{t+1} = \sup_j \prod_{k=1}^L f(Y_{k,t+1} | (\mathbf{J}_{t+1}^P)^{(j)})$. Then, accept $(j, (\mathbf{J}_{t+1}^P)^{(j)})$ with probability $\prod_{k=1}^L f(Y_{k,t+1} | (\mathbf{J}_{t+1}^P)^{(j)}) / c_{t+1}$. Based on the local samples from the rejection method, the estimates of the associated weights of the samples for each state are computed by the following procedure:

- (i) Estimate the weight $w_{t+1}^{(j)}$ by \hat{f}_j = frequency of $\{J = j\}$ in the sample.
- (ii) Update the sample $(\mathcal{J}_{t+1}^P)^{(j)} = ((\mathbf{J}_t^P)^{(j)}, (\mathbf{J}_{t+1}^P)^*)$ if $\hat{f}_j \neq 0$, where $(\mathbf{J}_{t+1}^P)^*$ is any value of \mathbf{J}_{t+1}^P if the associated $\hat{f}_j \neq 0$, or take a random draw from those with $\hat{f}_j \neq 0$ if the associated $\hat{f}_j = 0$.

Resample m' out of m rows from \mathcal{J}_{t+1}^P without replacement based on the weights $\{w_{t+1}^{(j)}, j = 1, \dots, m\}$. In order to improve the efficiency of SIS, the resampling scheme is used when the SIS arrives at the last time step rather than resampling after every step.

3.4. Improved SIS method with resampling. The disadvantage of the regular SIS with rejection method is that it requires computing the constant c_{t+1} within the embedded rejection method and re-estimation of the weights for the SIS procedure from the samples $\{J_{(l)}, (\mathbf{J}_{t+1}^P)^{(l)}\}_{l=1}^{m'}$. Both of these computations could be quite inefficient in the state space model with high dimension. However, an improvement could be made when the local importance resampling takes place so that the samples are not collected by the accept/reject ratio, but instead by assigning a weight to each sample. Specifically, calculating the constant c_{t+1} or estimating the weights by counting \hat{f}_j is no longer necessary; instead we simply assign to the samples $(\mathbf{J}_{t+1}^P)^{(j)}$ the weights $w_{t+1}^{(j)} = \prod_{k=1}^L f(Y_{k,t+1} | (\mathbf{J}_{t+1}^P)^{(j)})$. It has been proved [Liu and Chen (1998)] that the samples from the local importance resampling method would automatically achieve the resampling effect. Thus, we could just keep those weights from any of the local Monte Carlo methods and iterate the SIS.

4. Simulation study.

4.1. MEG data generation. In a typical MEG experiment, time is measured in milliseconds (the sampling rate is 1 kHz). However, for better understanding, from

now on, we will use timesteps rather than milliseconds. We ran two simulated cases to verify that the methods work. First, we present some preliminary results for the single dipole case with a few parameters and low dimension in time. Second, an extension to the single dipole case with six parameters and high dimension in time is given. We used 40 radially oriented magnetometers in one case, and 100 radially oriented magnetometers in the other. The dipole was restricted to move inside the brain. In order to focus on the source parameters, we fixed several parameters (source noise parameters, measurement noise parameters, etc.) in the model.

Simulated case 1. Before running our algorithms for a long time, we tested a simplified case where the simulation was run for only 15 timesteps with only one of the six parameters allowed to vary. In this very simple example, the dipole only moves in the z dimension in the brain and both the strength and moments of the dipole remain constant. The parameters of the simulated dipole are summarized in Table 1. The regular SIS method with rejection and the improved SIS method with resampling were tested. The random-walk MCMC within Gibbs and the hybrid MCMC within Gibbs were also run for comparison. We randomly generated 25 data sets and tested them under each scenario. Figure 1 shows the trace plots (only 5 overlay plots are shown) for 4 selected timesteps from all the methods. We observe that both the random-walk MCMC within Gibbs and hybrid MCMC within Gibbs do not provide a stable estimate for each timepoint and their samples are highly correlated. Both of our methods produce much nicer samples which oscillate around the true values. To have a quantitative comparison, we conducted a detailed convergence diagnosis for each approach: the sample autocorrelation function of the chain at a selected timestep was computed for each approach (see Figure 2); the effective sample size of the averaged chain from each approach was calculated; Gelman–Rubin’s method was used for evaluating convergence; similar methods such as Geweke, Heidelberger–Welch and Raftery–Lewis were also

TABLE 1
Dipole simulation: the location parameters of the dipole are expressed in terms of Cartesian coordinates $[x \text{ (cm)}, y \text{ (cm)}, z \text{ (cm)}]$, m_1 and m_2 are the dipole moment parameters. $s \text{ (mA)}$ is the strength parameter of a dipole. Only the z component of the dipole is allowed to vary. The other five components are held fixed by setting the diagonal components of the covariance matrix to zero

$\mathbf{m}_{\text{int}} = (x, y, z, m_1, m_2, s)$	$(1, 1, 5, 3, 3, 3)$
$\mathbf{m}_{\text{com}} = (x, y, z, m_1, m_2, s)$	$(0, 0, 0, 0, 0, 0)$
$\boldsymbol{\rho} = \text{diag}[\rho_1, \rho_2, \dots, \rho_6]$	$\text{diag}[1, 1, 0.9, 1, 1, 1]$
$\boldsymbol{\Sigma}_1 = \text{diag}[\sigma_1^2, \sigma_1^2, \dots, \sigma_1^2]$	$\text{diag}[0.0625, 0.0625, \dots, 0.0625]$
$\boldsymbol{\Sigma}_2 = \text{diag}[\sigma_{11}^2, \sigma_{22}^2, \dots, \sigma_{66}^2]$	$\text{diag}[0, 0, 0.0225, 0, 0, 0]$
Number of timesteps	15

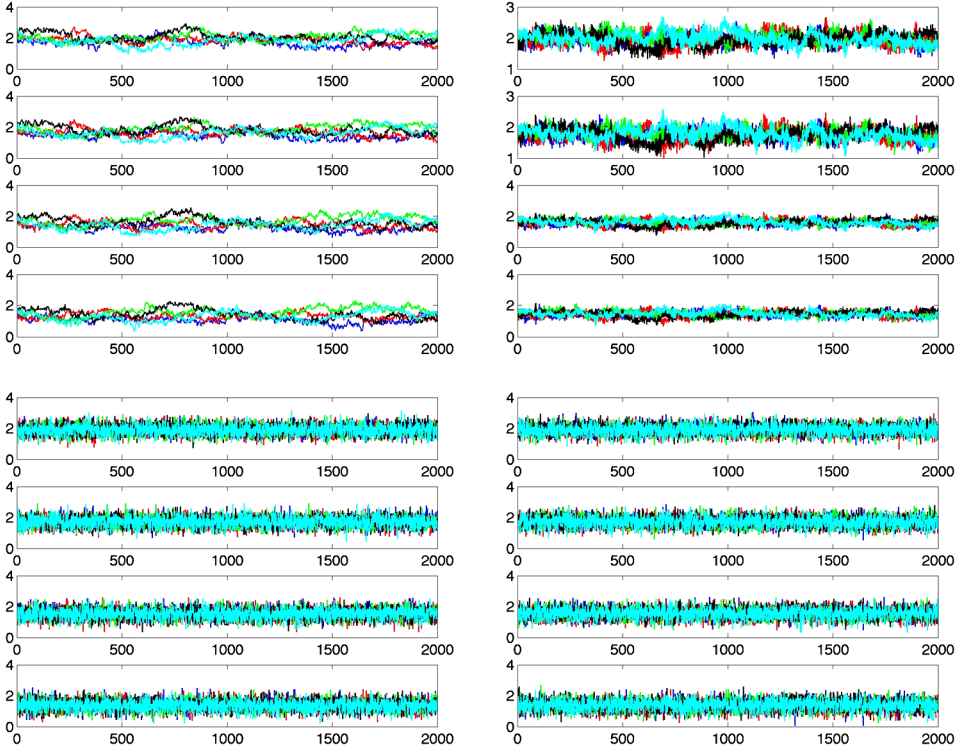


FIG. 1. A simple test case where only one source parameter z is allowed to vary. Top left: trace plots of location parameter z at four selected timesteps (9th, 10th, 11th and 12th) by the random-walk MCMC within Gibbs. Similar plots are also shown for the hybrid MCMC within Gibbs (top right), regular SIS method with rejection (bottom left) and improved SIS method with resampling (bottom right).




FIG. 2. Autocorrelation plots for the four methods. Top left: random-walk MCMC within Gibbs; top right: hybrid MCMC within Gibbs; bottom left: SIS method with rejection; bottom right: SIS with resampling.

TABLE 2

Convergence diagnosis by effective sample size (ESS), Gelman–Rubin (GR) [Gelman and Rubin (1992)]: value near 1 suggests convergence, Geweke (GE) [Geweke (1992)]: z-score for stationary test, Heidelberg–Welch (HW) [Heidelberg and Welch (1983)]: p value for stationary test, Raftery–Lewis (RL) [Raftery and Lewis (1992)]: large value suggests strong autocorrelation

Method	ESS	GR	GE	HW	RL
Random-walk MCMC within Gibbs	26.7	1.29	−14.968	0.05	18.2
Hybrid MCMC within Gibbs	553.5	1.06	−4.985	0.116	1.6
Regular SIS with rejection	1946.9	1.013	−0.64	0.25	1.0
Improved SIS with resampling	2000.0	1.005	−0.28	0.81	1.0

employed for diagnosis (see Table 2). Both Figures 1, 2 and Table 2 show strong evidence that our approaches outperform the MCMC methods.

Simulated case 2. In addition to case 1, a case of multiple-source parameters (three location parameters and three moment and strength parameters) was performed. In this simulation, the source was modeled as a moving dipole following a multivariate autoregressive time series. The dipole moves in the three coordinate directions x , y and z , and both strength and moments of the dipole change as well. The total length of simulation is 100 timesteps (we will run 2000 timesteps for data in Section 4.3). To control the movement of the simulated dipole (to not move outside of the brain when the number of timepoints are large), we restricted the range of each parameter for the dipole. In order to do this, we set boundary values for each parameter (i.e., maximum and minimum). The autoregressive model for \mathbf{J}_t^P in Section 2 occurred only at certain timepoints when specified in advanced. In other words, the dipole had two types of moves: one is a move based on the autoregressive model, and the other is a random-walk move. The dipole moved according to the autoregressive model at certain specified timesteps whereas the random walk was applied to the dipole at the rest of the timepoints. We had similar restrictions on the other parameters of the dipole. The parameters setup is given in Table 3. The plots (histogram) for each dipole location parameter and pairwise plots for the location parameters are shown in Figure 3. These side by side histograms show the distribution of each location parameter at six selected timepoints. Similar plots for the other three moment and strength parameters are also shown in Figure 4. We can see that the distributions (non-Gaussian) of each parameter of the source are varying at each timestep as we expected.

4.2. *Parallel virtual machine (PVM) for high dimension in time.* In practice, the MEG data set we have from an experiment is very large (e.g., hundreds of thousands of timesteps). The same algorithms (Sections 3.3 and 3.4) need to be run for a much longer time. To be exact, if we run for 5000 timesteps with

TABLE 3

Dipole simulation: the location parameters of the dipole are expressed in terms of Cartesian coordinates $[x \text{ (cm)}, y \text{ (cm)}, z \text{ (cm)}]$, m_1 and m_2 are the dipole moment parameters. $s \text{ (mA)}$ is the strength parameter of a dipole. The diagonal elements of Σ_1 and Σ_2 are 0.0625 fT^2 and 0.01 cm^2 , respectively

Initial timepoint

$$\begin{aligned} \mathbf{m}_{\text{int}} &= (x, y, z, m_1, m_2, s) & (6, 7, 8, 3, 5, 5) \\ \mathbf{m}_{\text{com}} &= (x, y, z, m_1, m_2, s) & (0, 0, 0, 0, 0, 0) \\ \rho &= \text{diag}[\rho_1, \rho_2, \dots, \rho_6] & \text{diag}[0.65, 0.7, 0.75, 0.8, 0.85, 0.9] \end{aligned}$$

Random-walk move

$$\begin{aligned} (x, y, z, m_1, m_2, s) & \quad \text{Based on previous value} \\ \text{Number of timesteps} & \quad 10 \end{aligned}$$

Autoregressive move

$$\begin{aligned} (x, y, z, m_1, m_2, s) & \quad \text{Based on previous value} \\ \mathbf{m}_{\text{com}} &= (x, y, z, m_1, m_2, s) & (0, 0, 0, 0, 0, 0) \\ \rho &= \text{diag}[\rho_1, \rho_2, \dots, \rho_6] & \text{diag}[0.65, 0.7, 0.75, 0.8, 0.85, 0.9] \end{aligned}$$

Random-walk move

$$\begin{aligned} (x, y, z, m_1, m_2, s) & \quad \text{Based on previous value} \\ \text{Number of timesteps} & \quad 10 \end{aligned}$$

...

...

Repeat until 100th timepoint

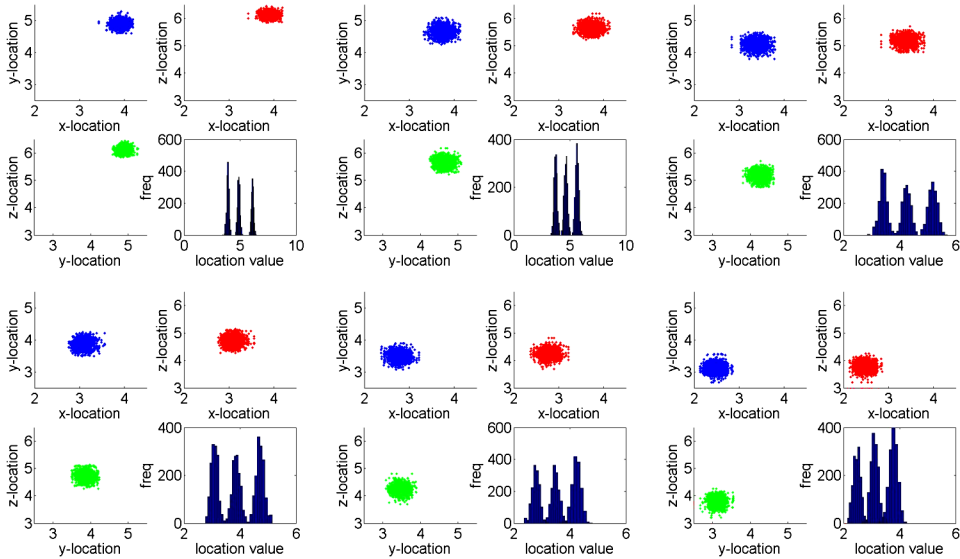


FIG. 3. Source location parameters at six timesteps (a total of six 2×2 subplots). Top left: in this 2×2 subplot, there are three pairwise plots of the source location parameters (x and y , x and z , y and z) at 1st timestep and one side by side histogram plot for the source location parameters (x , y and z) at 1st timestep. The rest of the five subplots give the same information for different timesteps: 20th timestep (top middle), 40th timestep (top right), 60th timestep (bottom left), 80th timestep (bottom middle) and 100th timestep (bottom right).




FIG. 4. Source moment and strength parameters at six timesteps (a total of six 2×2 subplots). Top left: in this 2×2 subplot, there are three pairwise plots of the source moment and strength parameters (m_1 and m_2 , m_1 and s and m_2 and s) at 1st timestep and one side by side histogram plot for the moment and strength parameters (m_1 , m_2 and s) at 1st timestep. The rest of the five subplots give the same information for different timesteps: 20th timestep (top middle), 40th timestep (top right), 60th timestep (bottom left), 80th timestep (bottom middle) and 100th timestep (bottom right).

1500 replications (sample paths) for each \mathbf{J}_t^P , we are supposed to get a stream of $S_{5000} = \{(\mathcal{J}_{5000}^P)^{(j)}, j = 1, \dots, 1500\}$ (S_t is defined in Section 3.2). Because of the sequential character of our algorithms, sample paths $[(\mathbf{J}_t^P)^{(j)}, j = 1, \dots, m]$ for each time are computed in a sequential fashion and the weights updated at each time. Therefore, it is very inefficient to get the sample paths for a longer time.

Note that we always need the sample path from the previous time (\mathbf{J}_{t-1}^P) when we work on the current time (\mathbf{J}_t^P) and they are not independent, therefore, it is not possible to improve the speed in the direction of time [e.g., $(\mathbf{J}_t^P)^{(j)}$ sequentially depends on $(\mathbf{J}_{t-1}^P)^{(j)}$]. However, the sample paths are independent within each timestep; this is to say, at time t , $(\mathbf{J}_t^P)^{(j)}$ is independent $(\mathbf{J}_t^P)^{(j')}$, so they can be computed in a separate fashion. In other words, it is always possible for us to compute several sample paths (several chunks) for the same timestep (at time t) simultaneously. This simultaneous computation for sample paths up to the final timestep (5000) could be achieved by parallel computing where each parallel thread would contain a sequential calculation for all the time t ($1 \leq t \leq 5000$) with fewer samples, so that our sequential problem can be solved in parallel. The Parallel Virtual Machine (PVM) [Geist et al. (1994)], a parallel computing paradigm, is used to speed up the computation. It is designed to allow a network of heterogeneous machines to be used as a single distributed parallel processor, so that a large

scale computing problem can be solved more cost effectively. The PVM structure we use is a Master–Worker model where there are several worker programs performing tasks in parallel and a master program collecting the outcomes from each worker. Each task is to separately compute a subset of the sample paths for all timesteps. The resampling scheme is included in the worker program and there is no parallelism in time. To be exact, if there are three worker programs in the Master–Worker model to generate a stream $S_T = \{(\mathcal{J}_T^P)^{(j)}, j = 1, \dots, m\}$, the way of running PVM is as follows:

Algorithm 2: PVM schedule

- (i) Initialize each worker program and let each worker run for a substream $S'_T = \{(\mathcal{J}_T^P)^{(j)}, j = 1, \dots, \frac{m}{3}\}$.
 - (ii) Stack each S'_T and get a complete S_T .
-

The size of S'_T can be adjusted according to the size of S_T and the number of worker programs that are in use. The speed is mainly influenced by hardware and software components of network and I/O systems. It also depends on the number of worker programs, for example, adding too many parallel workers does not enhance the speed when most of the time is spent on master–worker communication. In practice, deciding on the number of workers requires experience and it varies for different machines. Because the magnetic fields generated by independent dipoles add up, there is no additional complexity (other than increased computation) brought by multiple dipoles.

Since our PVM program involves randomness and a resampling scheme, several issues still need to be resolved. First, if our algorithms were implemented in a single program without parallelism, all samples generated before resampling from this program should be simply related to the random number generator. However, when there are several workers, each of them doing the same thing as a single program but in parallel, the unique randomness within each worker will eventually come up with different but similar samples before resampling. To be exact, in order to have the two programs generate the same results, in the PVM structure we need to explicitly and precisely choose different workers according to a pre-defined random sequence. This random sequence can be obtained from a single program without parallelism. Unfortunately, this needs a lot of work in programming and would surely slow down the computation. Second, in a single program without parallelism, we would only have one resampling procedure. The samples would be generated from the resampling procedure. However, there would be one resampling procedure within each of our worker programs in PVM. The samples would be generated from each of these workers and should eventually be pooled together. In principle, the weights from each worker should be pooled first and then we would perform the resampling procedure. The reason is that each worker might generate different weights so that the normalizing constants might be different. If the resampling happens only one time (at the end of all timesteps), a reasonable

way to solve this problem is that we can do the resampling scheme in the master program after normalizing all the weights when pooled. If there were several re-sampling schemes before the end, we could still return to the master program when necessary. Again, this needs extensive programming and, again, it would surely slow down the computations. In our current program, sums of weights within each worker were almost the same (normalizing constants were almost the same), so we retained the resampling procedure in each worker program. Because the random number generation will not produce the same numbers in a parallel program as in a sequential program without extensive interprocess communication and because re-sampling within each parallel worker program will produce different results than would resampling in the master, we do not expect the identical samples in the parallel version of our sequential program. We do expect the distribution of the samples from the parallel program to be indistinguishable from the distribution of the samples from the sequential version.

4.3. *Numerical results for running PVM for MEG model.* The PVM was first run on a single Linux workstation (Intel Pentium 4 CPU 3.80 GHz, Memory 2 GB) for different configurations. The data size was 2000 MEG timesteps with 1500 sample paths for each timestep. We split the computation into a number of tasks: 1 (without PVM), 3, 5, 10 and 15 workers, respectively, and run for 100, 500, 1000, 1500 and 2000 timesteps. The user CPU time (total number of CPU-seconds for master and worker programs) is used to measure the time spent by each PVM run. The real time elapsed (minutes) is also shown in parentheses beside the user CPU time. The result is shown in Table 4.

We can see that the user CPU time increases roughly linearly in the number of timesteps from 0.008 second to 0.146 second on average. The linear relationship of user CPU time on experiment time is almost the same for each of these PVM configurations as we expected. This can be clearly observed from Figure 5: in Figure 5(a), these lines (user CPU time/Task) are nearly equally distant and stay

TABLE 4
PVM application on a single workstation. Five different PVM configurations were run. The number of workers in PVM is denoted “number of tasks.” The number of sample paths within each worker is denoted “load per task.” Each PVM run eventually generates 1500 sample paths

Number of tasks	Load per task	Time 1 (100)	Time 2 (500)	Time 3 (1000)	Time 4 (1500)	Time 5 (2000)
1	1500	0.008 (1.00)	0.032 (5.12)	0.064 (10.35)	0.120 (16.47)	0.136 (22.08)
3	500	0.008 (0.24)	0.032 (2.05)	0.060 (4.12)	0.096 (6.23)	0.148 (8.40)
5	300	0.008 (0.17)	0.036 (1.27)	0.064 (3.17)	0.104 (4.25)	0.148 (6.47)
10	150	0.008 (0.11)	0.040 (0.59)	0.072 (1.59)	0.096 (3.00)	0.136 (4.51)
15	100	0.008 (0.10)	0.036 (0.50)	0.064 (1.43)	0.124 (2.33)	0.164 (3.24)

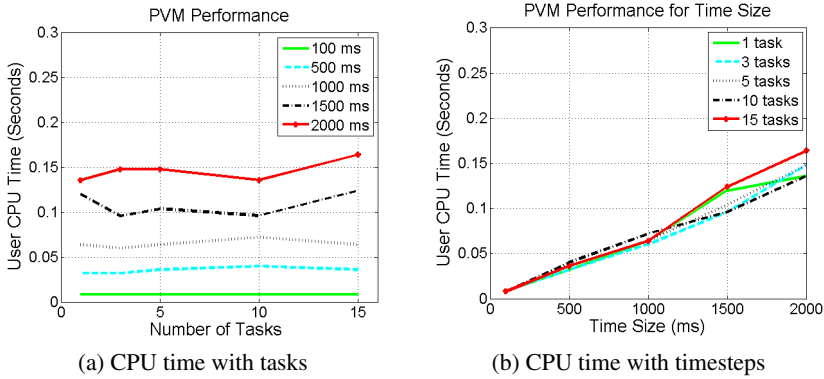


FIG. 5. PVM Performance: user CPU time (seconds) for number of tasks and different time run. (a) Each line (with a specific timestep) is a plot of user CPU time for different number of tasks. (b) Each line (with a specific number of tasks) is a plot of user CPU time for different timesteps.

roughly constant for different tasks within the same steps time run; in Figure 5(b), the slope of each line (user CPU time/Timesteps) is almost the same. Note that there is a significant difference in real time elapsed for different PVM configurations. This should not be considered a contradiction with user CPU time because real time elapsed is mostly affected by other programs and it includes time spent in memory, I/O and other resources.

The performance can still be improved when extra machines are included. Table 5 lists the PVM performance of 1–4 machines with 1500 timesteps. First, since user CPU time is the sum of the CPU time for master and worker programs, it is expected that the user CPU time for each of these PVM runs is roughly 0.120 second. Second, the real elapsed time of each PVM run is cut to 50%–70% if one machine is added. It then goes down to 40%–50% when three computers are employed. The real time elapsed decreases to 10%–30% when four computers are added. These performances are based on our public cluster with heterogeneous CPU speed and cache size. The theoretical reduction in execution time of PVM is not necessarily

TABLE 5

PVM application on multiple workstations. This table shows the user CPU time (seconds) for each PVM run and real time elapsed (minutes) in parentheses using one, two, three and four machines. The length of each PVM run was 1500 timesteps

Number of tasks	Load per task	Time 1 (1500)	Time 2 (1500)	Time 3 (1500)	Time 4 (1500)
3	500	0.084 (7.56)	0.128 (5.46)	0.108 (3.39)	0.096 (2.34)
5	300	0.100 (4.55)	0.084 (3.10)	0.124 (2.23)	0.108 (1.29)
10	150	0.100 (3.19)	0.096 (1.51)	0.104 (1.39)	0.116 (1.03)
15	100	0.124 (2.34)	0.112 (1.31)	0.104 (1.00)	0.112 (0.44)

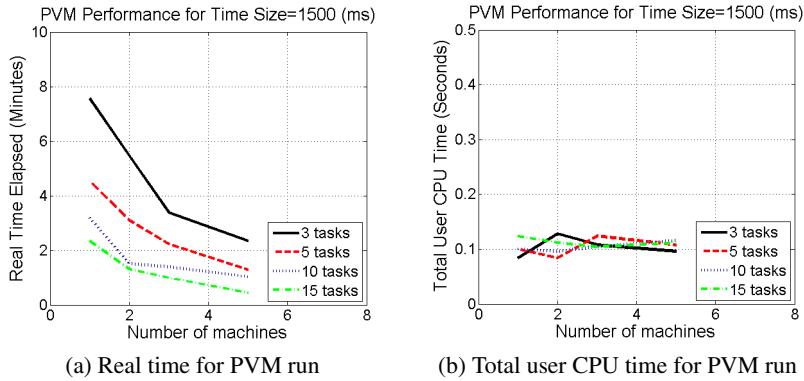


FIG. 6. PVM Performance: real time elapsed (minutes) and user CPU time (seconds) graph for number of machines for 1500 timesteps PVM run. (a) Each line (with a specific number of tasks) is a plot of real time elapsed for a different number of machines. (b) Each line (with a specific number of tasks) is a plot of total user CPU time of master and worker programs for different number of machines.

achieved. Finally, to get better time execution by PVM, we suggest to adjust the number of CPUs and the number of tasks, and to use relatively similar machines. To summarize, Figure 6 is a graphic illustration of both real time elapsed and user CPU time for our PVM run.

5. A real data application. Data was collected by a 306-channel system (Elekta-Neuromag) at the Center for Advanced Brain Magnetic Source Imaging (CABMSI) at UPMC Presbyterian Hospital in Pittsburgh in an experiment related to Brain-controlled interfaces (BCI). A BCI expresses motor commands via neural signals directly from the brain. The experiment involves two parts (see Figure 7): in the first part the subjects were asked to imagine performing the “center-out” task using the wrist (imagined movement task) and in the second part the subjects controlled a 2-D cursor using the wrist to perform the center-out task following a visual target (overt movement task).

The data consists of one trial recording 37,000 milliseconds at 102 MEG sensors (magnetometers). We used this data for testing our model along with our PVM scheme. Instead of analyzing the whole trial of data, we only analyzed about 400 milliseconds (dashed box in Figure 8) after movement onset (12,000 milliseconds–12,400 milliseconds in the original data) from all the channels. To simplify our calculation for the real data, we were only estimating the location of the source (x, y, z) . The moment and strength parameters (m_1, m_2, s) were not of our interest (not varying too much by assumption). The choice of prior for real data is an open question; we used almost the same prior as we did in Section 4 for simplification. We set the mean \mathbf{m}_{ini} of the initial state \mathbf{J}_0^p as $(-4, -4, 11)$ motivated by the minimum norm estimate [Hämäläinen and Ilmoniemi (1994)], which is $(-2, -2, 10)$. We further assumed a unit moment and strength for the dipole

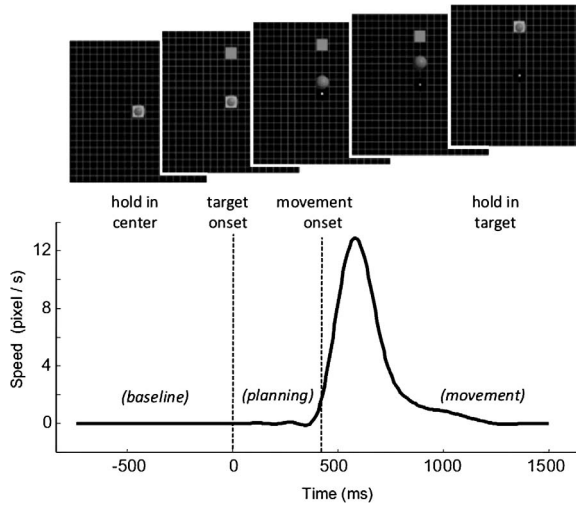


FIG. 7. The subject controls the 2-D cursor position using wrist movements. The cursor needs to go to the center and stay there for a hold period until the peripheral target appears. Then the cursor moves from the center out to the target and stays there for another hold period to complete the trial successfully. The target changes color when hit by the cursor and disappears when the holding period has finished. The bottom trace shows the speed profile of the cursor from a representative trial, and the dotted lines delimit the pre-movement/planning period. Figure and explanation were obtained from Wang *et al.* (2010).

in the data. The starting values for the initial state (x, y, z, m_1, m_2, s) were set to $(-4.06, -3.77, 13.13, 1.11, 0.98, 1.12)$. The empirical density plots of the dipole location parameter (x, y, z) at two selected timesteps are shown in Figure 9. Using the density plots of the location parameters, we were able to find the dipole distribution at different timesteps. Figure 10 shows several snapshots of dipole distribution at six timesteps, that is, the data cloud in each plot tells where the dipole might be located at a given timestep. A full movie of the dipole distribution for 100 milliseconds can be found at http://smat.epfl.ch/~zyao/3dplot_animation.gif.




FIG. 8. The MEG signal of a typical trial at a magnetometer. The horizontal axis is time (ms) and the vertical axis is the magnitude of the signal (fT).




FIG. 9. Empirical density for source location parameter (x, y, z) . Left: density plot for 1st millisecond; right: density plot for 101st millisecond.

Different initial values might have different performance due to the complexity of the problem and the real data, thus, a more realistic prior needs to be investigated in our future work. We ran PVM for 1500 milliseconds (12,000 milliseconds–13,500 milliseconds in the original data) with the same PVM configuration as our simulation; the time spent was very close to that from our previous simulation results.

A typical MEG analysis would report the estimated movement of the dipole at each time step. We can do the same. Additionally, because we have samples from a probability distribution we can provide estimates of the variation of the estimated movement and other source parameters. Other methods cannot provide appropriate estimates of variation. In clinical applications, estimates of variation may provide neurosurgeons a much better basis for their decisions. We still need to overcome the computational burden to make our method feasible for a complete application.

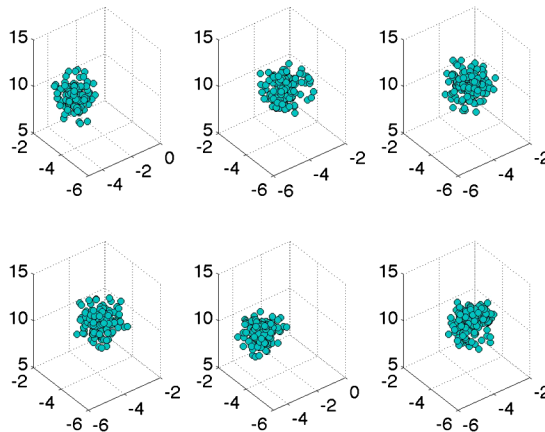


FIG. 10. Time variation of the dipole distribution. Upper row: dipole distribution at 1st, 21st and 41st millisecond (left to right); lower row: dipole distribution at 61st, 81st and 101st millisecond.

6. Conclusion and discussion. We have introduced a general state-space formulation for the time dependency of the parameters of a current dipole which generates the MEG signals. In this paper we have only considered the simplest sort of model, a first order autoregression for the time dependency. However, the framework allows for much more complex models. And, because three of the parameters are spatial coordinates, the model automatically incorporates space–time dependencies. The time dependency in the model has greatly expanded the parameter space. That fact together with the nonlinearity of the model means that more typical MCMC methods converge extremely slowly. The benefit of sequential methods is that we do not attempt to estimate the entire target distribution at once but rather attempt to estimate samples for each time point sequentially. Because of the expanded parameter space, the need for parallel computational methods is obvious. Our initial attempt utilized PVM and provided the expected reduction in running time.

The results so far are mainly based on a one-source model where we assumed there was only one dipole in the data. The extension from one source to multiple sources is natural and only the computational complexity increases. Our algorithms will still work in this multiple-source model case. However, to determine the number of sources in the MEG data is still an open question. In general, there are three ways of finding the number of sources for the advanced model. The first one, which is relatively easy, is to use a predefined number of sources for the data. The second one is to estimate the number of sources from the data in advance [Waldorp et al. (2005), Bai and He (2006), Yao and Eddy (2012)]. The third one is to model the number of the sources using a prior distribution [Bertrand et al. (2001)].

We fixed several parameters when we compared our algorithms with other MCMC methods. In fact, those parameters could be estimated along with the source distribution. The natural way of implementing this is to include the estimation of those parameters and the source distribution in the iterations until all of them become stable. Furthermore, the skewness of weights that arises in sequential importance sampling could be a trade-off between the efficiency of the algorithm and the quality of the source distribution. Naturally, we have observed some skewness in the weights; we do not have enough experience to evaluate whether this skewness should be considered excessive or unusual. Residual sampling [Liu and Chen (1998)] or stratified sampling [Kitagawa (1996)] could replace regular weight sampling and might address excessive skewness.

To summarize, due to its nonuniqueness, finding a good estimate of the MEG source is a challenging problem which is still open. We have proposed a predictive model for finding a distribution for the MEG source and we have applied our methods on both simulated data and real data. In practice, the MEG data sets from different experimental settings are much more complicated. Our methods can be used as a reference with other source localization methods. Driven by the desire of looking at the brain activity in real time, we plan to implement a computing

environment to study the brain activity under the real MEG temporal resolution (1/1000 sec). The computational challenge arises due to the extremely large dimensionality of the problem (high resolution); there is no common computing architecture that could help. We are exploring the use of more advanced forms of parallelism such as CUDA (Compute Unified Device Architecture) and OPENCL to further reduce the running time in the future.

Acknowledgments. We thank Rob Kass and his collaborators for allowing us to use their BCI data to test our methods. Leon Gleser gave helpful comments on the paper. We thank the referee, the Area Editor and the Editor-in-Chief for their helpful comments. We are especially grateful to the Associate Editor for his/her extremely helpful comments on versions of the manuscript and for his patience and persistence. We would also like to thank Dr. Alberto Sorrentino and Prof. Michele Piana of the Dipartimento di Matematica, Università di Genova for providing several references.

REFERENCES

- ANDRIEU, C. and THOMS, J. (2008). A tutorial on adaptive MCMC. *Stat. Comput.* **18** 343–373. [MR2461882](#)
- BAI, X. and HE, B. (2006). Estimation of number of independent brain electric sources from the scalp EEGs. *IEEE Trans. Biomed. Eng.* **53** 1883–1892.
- BARKLEY, G. L. and BAUMGARTNER, C. (2003). MEG and EEG in epilepsy. *J. Clin. Neurophysiol.* **20** 163–178.
- BERTRAND, C., OHMI, M., SUZUKI, R. and KADO, H. (2001). A probabilistic solution to the MEG inverse problem via MCMC methods: The reversible jump and parallel tempering algorithms. *IEEE Trans. Biomed. Eng.* **48** 533–542.
- BERZUINI, C., BEST, N. G., GILKS, W. R. and LARIZZA, C. (1997). Dynamic conditional independence models and Markov chain Monte Carlo methods. *J. Amer. Statist. Assoc.* **92** 1403–1412. [MR1615251](#)
- CAMPI, C., PASCARELLA, A., SORRENTINO, A. and PIANA, M. (2008). A Rao–Blackwellized particle filter for magnetoencephalography. *Inverse Problems* **24** 25023–25037.
- CAMPI, C., PASCARELLA, A., SORRENTINO, A. and PIANA, M. (2011). Highly automated dipole estimation. *Computational Intelligence and Neuroscience* **2011** Article ID 982185, 11 pp.
- CARLIN, B. P., POLSON, N. G. and STOFFER, D. S. (1992). A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *J. Amer. Statist. Assoc.* **87** 493–500.
- CARTER, C. K. and KOHN, R. (1994). On Gibbs sampling for state space models. *Biometrika* **81** 541–553. [MR1311096](#)
- COHEN, D. (1968). Magnetoencephalography: Evidence of magnetic fields produced by alpha-rhythm currents. *Science* **161** 784–786.
- DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68** 411–436. [MR2278333](#)
- FEARNHEAD, P. (2008). Computational methods for complex stochastic systems: A review of some alternatives to MCMC. *Stat. Comput.* **18** 151–171. [MR2390816](#)
- GAMERMAN, D. (1998). Markov chain Monte Carlo for dynamic generalised linear models. *Biometrika* **85** 215–227. [MR1627273](#)
- GEIST, A., BEGUELIN, A., DONGARRA, J., JIANG, W., MANCHEK, R. and SUNDERAM, V. S. (1994). *PVM: Parallel Virtual Machine: A Users' Guide and Tutorial for Network Parallel Computing* (Scientific and Engineering Computation). MIT Press, Cambridge, MA.

- GELMAN, A., ROBERTS, G. O. and GILKS, W. R. (1996). Efficient Metropolis jumping rules. In *Bayesian Statistics* 599–607. Oxford Univ. Press, New York. [MR1425429](#)
- GELMAN, A. and RUBIN, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statist. Sci.* **7** 457–511.
- GEWEKE, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In *Bayesian Statistics* 169–193. Oxford Univ. Press, New York. [MR1380276](#)
- GORDON, N. J., SALMOND, D. J. and SMITH, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)* **140** 107–113.
- GRIFFITHS, D. J. (1999). *Introduction to Electrodynamics*. Prentice Hall, New York.
- HÄMÄLÄINEN, M. S. and ILMONIEMI, R. J. (1994). Interpreting magnetic fields of the brain: Minimum norm estimates. *Med. Biol. Eng. Comput.* **32** 35–42.
- HÄMÄLÄINEN, M. S., HARI, R., ILMONIEMI, R. J., KNUUTILA, J. and LOUNASMAA, O. V. (1993). Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of signal processing in the human brain. *Rev. Modern Phys.* **65** 413–497.
- HEIDELBERGER, P. and WELCH, P. D. (1983). Simulation run length control in the presence of an initial transient. *Oper. Res.* **31** 1109–1144.
- JUN, S. C., GEORGE, J. S., PARÉ-BLAGOEV, J., PLIS, S. M., RANKEN, D. M., SCHMIDT, D. M. and WOOD, C. C. (2005). Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data. *NeuroImage* **28** 84–98.
- KITAGAWA, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Statist.* **5** 1–25. [MR1380850](#)
- KNORR-HELD, L. (1999). Conditional prior proposals in dynamic models. *Scand. J. Stat.* **26** 129–144.
- KRISTEVA-FEIGE, R., ROSSI, S., FEIGE, B., MERGNER, TH., LUCKING, C. H. and ROSSINI, P. M. (1997). The Bereitschaftspotential paradigm in investigating voluntary movement organization in humans using magnetoencephalography (MEG). *Brain Res. Protoc.* **1** 13–22.
- KYBIC, J., CLERC, M., FAUGERAS, O., KERIVEN, R. and PAPADOPOULOU, T. (2006). Generalized head models for MEG/EEG: Boundary element method beyond nested volumes. *Phys. Med. Biol.* **51** 1333–1346.
- LIU, J. S. (1996). Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statist. Comput.* **6** 113–119.
- LIU, J. S. and CHEN, R. (1998). Sequential Monte Carlo methods for dynamic systems. *J. Amer. Statist. Assoc.* **93** 1032–1044. [MR1649198](#)
- MATTOUT, J., PHILLIPS, C., PENNY, W. D., RUGG, M. D. and FRISTON, K. J. (2006). MEG source localization under multiple constraints: An extended Bayesian framework. *NeuroImage* **30** 753–767.
- MIAO, L., MICHAEL, S., KOVVALI, N., CHAKRABARTI, C. and PAPANDREOU-SUPPAPPOLA, A. (2013). Multi-source neural activity estimation and sensor scheduling: Algorithms and hardware implementation. *Journal of Signal Processing Systems* **70** 145–162.
- MOSHER, J. C., LEWIS, P. S. and LEAHY, R. M. (1992). Multiple dipole modeling and localization from spatio-temporal MEG data. *IEEE Trans. Biomed. Eng.* **39** 541–557.
- NEAL, R. M. (2001). Annealed importance sampling. *Stat. Comput.* **11** 125–139. [MR1837132](#)
- OKADA, Y., LÄHTENMÄKI, A. and XU, C. (1999). Comparison of MEG and EEG on the basis of somatic evoked responses elicited by stimulation of the snout in the juvenile swine. *Clin. Neurophysiol.* **110** 214–229.
- OU, W., HÄMÄLÄINEN, M. S. and GOLLAND, P. (2009). A distributed spatio-temporal EEG/MEG inverse solver. *NeuroImage* **44** 932–946.
- RAFTERY, A. E. and LEWIS, S. M. (1992). One long run with diagnostics: Implementation strategies for Markov chain Monte Carlo. *Statist. Sci.* **7** 493–497.

- ROBERTS, G. O. and ROSENTHAL, J. S. (2009). Examples of adaptive MCMC. *J. Comput. Graph. Statist.* **18** 349–367. [MR2749836](#)
- SARVAS, J. (1984). Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys. Med. Biol.* **32** 11–22.
- SCHMIDT, D. M., GEORGE, J. S. and WOOD, C. C. (1999). Bayesian inference applied to the electromagnet inverse problem. *Hum. Brain Mapp.* **7** 195–212.
- SHEPARD, N. and PITTT, M. K. (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* **84** 653–667. [MR1603940](#)
- SOMERSALO, E., VOUTILAINEN, A. and KAIPIO, J. P. (2003). Non-stationary magnetoencephalography by Bayesian filtering of dipole models. *Inverse Problems* **19** 1047–1063. [MR2024688](#)
- SORRENTINO, A., PARKKONEN, L., PASCARELLA, A., CAMPI, C. and PIANA, M. (2009). Dynamical MEG source modeling with multi-target Bayesian filtering. *Hum. Brain Mapp.* **30** 1911–1921.
- SORRENTINO, A., JOHANSEN, A. M., ASTON, J. A. D., NICHOLS, T. E. and KENDALL, W. S. (2013). Dynamic filtering of static dipoles in magnetoencephalography. *Ann. Appl. Stat.* **7** 955–988. [MR3113497](#)
- SRINIVASAN, R. (2002). *Importance Sampling: Applications in Communications and Detection*. Springer, Berlin. [MR1949250](#)
- UUTELA, K., HÄMÄLÄINEN, M. S. and SOMERSALO, E. (1999). Visualization of magnetoencephalographic data using minimum current estimates. *NeuroImage* **10** 173–180.
- VAN VEEN, B., JOSEPH, J. and HECOX, K. (1992). Localization of intra-cerebral sources of electrical activity via linearly constrained minimum variance spatial filtering. In *Proc. IEEE 6th SP Workshop on Statistical Signal and Array Processing* 526–529. Victoria, BC.
- WALDORP, L. J., HUIZENGA, H. M., NEHORAI, A., GRASMAN, R. P. P. P. and MOLENAAR, P. C. M. (2005). Model selection in spatio-temporal electromagnetic source analysis. *IEEE Trans. Biomed. Eng.* **52** 414–420.
- WANG, W., SUDRE, G. P., XU, Y., KASS, R. E., COLLINGER, J. L., DEGENHART, A. D., BAGIC, A. I. and WEBER, D. J. (2010). Decoding and cortical source localization for intended movement direction with MEG. *J. Neurophysiol.* **104** 2451–2461.
- YAO, Z. and EDDY, W. F. (2012). Statistical approaches to estimating the number of signal sources in magnetoencephalography. Unpublished manuscript.

SECTION DE MATHÉMATIQUES
 ECOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE
 EPFL STATION 8, 1015 LAUSANNE
 SWITZERLAND
 E-MAIL: zhigang.yao@epfl.ch

DEPARTMENT OF STATISTICS
 CARNEGIE MELLON UNIVERSITY
 PITTSBURGH, PENNSYLVANIA 15213
 USA
 E-MAIL: bill@stat.cmu.edu