CONDITIONALLY GAUSSIAN HYPERMODELS FOR CEREBRAL SOURCE LOCALIZATION

DANIELA CALVETTI, HARRI HAKULA, SAMPSA PURSIAINEN,
AND ERKKI SOMERSALO

ABSTRACT. Bayesian modeling and analysis of the MEG and EEG modalities provide a flexible framework for introducing prior information complementary to the measured data. This prior information is often qualitative in nature, making the translation of the available information into a computational model a challenging task. We propose a generalized gamma family of hyperpriors which allows the impressed currents to be focal and we advocate a fast and efficient iterative algorithm, the Iterative Alternating Sequential (IAS) algorithm for computing maximum a posteriori (MAP) estimates. Furthermore, we show that for particular choices of the scalar parameters specifying the hyperprior, the algorithm effectively approximates popular regularization strategies such as the Minimum Current Estimate and the Minimum Support Estimate. The connection between priorconditioning and adaptive regularization methods is also pointed out. The posterior densities are explored by means of a Markov Chain Monte Carlo (MCMC) strategy suitable for this family of hypermodels. The computed experiments suggest that the known preference of regularization methods for superficial sources over deep sources is a property of the MAP estimators only, and that estimation of the posterior mean in the hierarchical model is better adapted for localizing deep sources.

1. INTRODUCTION

The human brain contains approximately 10^{11} exitable neurons whose resting state is characterized by a cross-membrane voltage difference. Electromagnetic signals propagate as perturbations of this voltage difference, or action potentials, along the axons and are transferred across the synaptic gaps by neurotransmitters, creating a post-synaptic potential in the receiving neurons. The post-synaptic potential may be relatively stable over a period of milliseconds and, being well localized, it can be modelled mathematically as current dipole. The neurons are organized in bundles, and when thousands of neighboring neurons are simultaneously in the post-synaptic excitation state, the net effect of the post-synaptic potentials gives rise to a localized current approximately parallel to the neuron bundle. This elementary impressed source current drives an Ohmic volume current in the brain tissue, and the net electromagnetic field can be registered on or outside the skull. Imaging of the neuronal activity based on the registered electric voltage (EEG) or on the magnetic field (MEG) has become a standard research tool in clinical and cognitive studies. Furthermore, when coupled to functional imaging methods,

MEG and EEG have a great potential to gather pertinent information about the coupling between neuronal activity and cerebral hemodynamics.

The advantage of electromagnetic brain imaging modalities is their good temporal resolution, about a millisecond, while their spatial resolution is limited by various factors, including weak signal-to-noise ratio, ambiguities in the source estimation due to the non-uniqueness of the inverse source problem [25] and uncertainties in the model. For example, while EEG suffers from lack of knowledge of the electric conductivity distribution and anisotropy of the head, MEG is known to be less affected by that [9, 34]. For both modalities, the anisotropy of the white matter may result in a strong bias of the volume currents, an effect that should be taken into account in a detailed model.

The properties of the solutions to electromagnetic inverse source problems depend on the *a priori* information implemented in the algorithm. One piece of information that may certainly improve the spatial resolution, if properly incorporated in the algorithm, is the focal nature of the impressed currents. In fact, since the electric and magnetic fields outside the skull depend linearly on the impressed currents, but the component of the source belonging to the significantly nontrivial null space of the forward map has no effect on the data, the determination of a physiologically meaningful solution must be based on complementary information. The implementation of feasible selection criteria has led to different solutions. For example, setting the null space component to zero gives the *minimum norm estimate* (MNE) [12], while minimization of the current, or ℓ^1 norm of the source results in the *minimum current estimate* (MCE) [31]. Solutions such as the Low Resolution Electromagnetic Tomography (LORETA) [22] are based on the assumption of smoothness of the source. In [13], the localization is pursued by local multipole expansions of the fields. Low-dimensional parametric models, for example those using only a few current dipoles [8, 19, 11], lead to localized solutions but limit the number of possible source configurations. Prior anatomical information such as the location of the sulci has been shown to improve the performance of the localization [2].

Bayesian methods are widely used in EEG/MEG to implement pertinent prior information such as anatomic constraints or functional information based on other imaging modalities [2, 14, 27, 30]. An additional level of flexibility to Bayesian modelling is provided by hierarchical models that allow uncertainties in the prior model itself [1, 21, 26]. In particular, the model hierarchy is a powerful tool for including prior information that is qualitative rather than quantitative [7].

The connection between classical regularization methods and Bayesian Maximum A Posteriori (MAP) estimates is well known. In particular, the Gaussian models lead to a MAP estimate that coincides with the standard Tikhonov regularized solution with a quadratic penalty, while non-Gaussian models are needed for non-quadratic penalties that are often more useful but computationally more challenging. In this work, we construct a *conditionally* Gaussian hierarchical parametric model that has the computational advantages of Gaussian prior models but leads to a rich class of MAP estimators that have the desirable qualitative properties of numerous commonly used non-Gaussian models. Using the conditional normality, we construct a fast, efficient and simple MAP estimation algorithm, and show that with proper choices of the few model parameters, the algorithm can be interpreted as a fixed point iteration for solving the minimum norm estimate, the

minimum current estimate and, more generally, the minimum ℓ^p estimate, while with a different choice, the algorithm approximates the *minimum support estimate* (MSE) [20]. Hence, our approach puts these methods in a unified framework.

The estimation algorithms require the solution of large linear systems, the system size depending on the discretization of the model. When a realistic three-dimensional model is used as in the present paper, iterative linear systems solvers are indispensable. It is common practice to improve the performance of the iterative solvers by preconditioners [24]. Recently, the authors introduced the concept of *priorconditioner* [6], a preconditioner that is based on the prior model rather than on linear algebraic properties of the system matrix. It is shown in this article that the priorconditioning based on the conditionally Gaussian hierarchical models yields an effective implementation of adaptive regularization similar to the FOCUSS (FOCal Underdetermined System Solver) algorithm [10].

A well-known feature of many regularization based source localization algorithms, including the minimum norm and minimum current estimates, is their tendency to favor surface sources, leading sometimes to a gross misplacement of deep sources. To suppress this bias, different weighting methods have been proposed to favor deep sources over the shallow ones, see, e.g., [16] for a recent account. From the statistical point of view, however, a depth dependent prior covariance which compensates for the preference of the likelihood towards superficial sources appears rather dubious, since there is no reason to believe a priori that a deep source should have higher variance than a superficial one. In this article, we confirm with numerical experiments that surface biasing may be a property of the MAP estimate, while Markov Chain Monte Carlo (MCMC) sampling based posterior mean estimates may localize better deep sources without requiring weight compensation. This result reinforces the concept, well known to the Bayesian statistics community, that the mode of the posterior distribution may do a poor job at representing the distribution, demonstrating that the Bayesian formulation of the inverse problem is much more than "yet another way to regularize an ill-posed problem", as is sometimes incorrectly stated.

Preliminary testing of the inversion algorithms and the MCMC sampling is done using a simple geometry with constant conductivity, modelling the distributed currents by a field of current dipoles. Subsequently, numerical experiments with a realistic three dimensional head model, with different electric conductivities in the scalp, skull, cerebrospinal fluid and the brain tissue are also presented. The volume current calculations are carried out with a finite element algorithm developed for this purpose. The distributed currents are represented by Raviart-Thomas elements that are a reasonable substitute for singular dipoles in the FEM context.

2. Forward model, EEG and MEG

We start by introducing the notations to be used in the sequel and review some basic facts concerning the forward model, which is based on the standard approach using the quasi-static approximation of Maxwell's equations [25]. We denote by $D \subset \mathbb{R}^3$ the head with boundary surface S and scalar conductivity distribution $\sigma > 0$. If we let $\mathbf{J}$ denote the impressed current density in D, the total current density, consisting of the impressed current and the Ohmic volume current, is $\mathbf{J}_{\mathrm{tot}} = \mathbf{J} + \sigma \mathbf{E}$, where $\mathbf{E}$ is the electric field induced by the current. Under the quasi-static approximation, we may assume that $\mathbf{E}$ is conservative, $\mathbf{E} = -\nabla u$. Neglecting the

electric displacement current in the Maxwell-Ampère equation, the total current density is divergence free, leading to the Poisson equation for the electric potential,

$$(1) \qquad \nabla \cdot (\sigma \nabla u) = \nabla \cdot \mathbf{J}, \quad \left. \frac{\partial u}{\partial n} \right|_S = 0,$$

where the Neumann boundary condition follows from the assumption that the conductivity vanishes outside D. The magnetic field $\mathbf{B}$ outside the head induced by the total current is, according to the Biot-Savart law, obtained as

$$(2) \qquad \mathbf{B}(x) = \frac{\mu_0}{4\pi} \int_D \mathbf{J}_{\text{tot}} \times \frac{x - y}{|x - y|^3} dy, \quad x \in \mathbb{R}^3 \setminus D.$$

The computation of the magnetic field therefore consists of two steps: the numerical solution of the Poisson equation (1) to find u and thus the total electric current density $\mathbf{J}_{\text{tot}} = \mathbf{J} - \sigma \nabla u$, and the computation of the integral yielding the magnetic field $\mathbf{B}$ (2). Each of these steps poses computational challenges.

The solution of the boundary value problem (1) can be formally expressed in terms of the Neumann Green's function $\mathcal{G}_\mathrm{N}$ of the diffusion operator $\nabla \cdot \sigma \nabla$,

$$(3) \qquad u(x) = \int_D \mathcal{G}_\mathrm{N}(x, y) \nabla \cdot \mathbf{J}(y) dy.$$

We assume that the electric potential is measured at locations x_ℓ, $1 \le \ell \le L$ on the surface S of the head. The approximation of the impressed current by a finite linear combination $\mathbf{J} = \sum_{k=1}^{K} \alpha_k \mathbf{j}_k$ of basis currents $\mathbf{j}_k$, $1 \le k \le K$, leads to a discrete model

$$
\begin{aligned}
u_\ell &= u(x_\ell) = \sum_{k=1}^{K} \left(\int_D \mathcal{G}_\mathrm{N}(x_\ell, y) \nabla \cdot \mathbf{j}_k(y) dy \right) \alpha_k \\
&= \sum_{k=1}^{K} M_{\ell k}^{\mathrm{e}} \alpha_k, \quad 1 \le \ell \le L,
\end{aligned}
$$

where the matrix $M^{\mathrm{e}} \in \mathbb{R}^{L \times K}$ is the *electric lead field matrix*. Similarly, assume that outside the head at points x_n, $1 \le n \le N$, the projection of the magnetic field in given directions $\mathbf{e}_n$ is measured. Substituting the expression (3) in the Biot-Savart law and representing the impressed current in the basis $\mathbf{j}_k$, we have the discrete model

$$
\begin{aligned}
v_n &= \mathbf{e}_n \cdot \mathbf{B}(x_n) \\
&= \sum_{k=1}^{K} \alpha_k \mathbf{e}_n \frac{\mu_0}{4\pi} \int_D \left\{ \mathbf{j}_k(y) - \sigma(y) \nabla \int_D \mathcal{G}_\mathrm{N}(y, z) \nabla \cdot \mathbf{j}_k(z) dz \right\} \times \frac{x_n - y}{|x_n - y|^3} dy \\
&= \sum_{k=1}^{K} M_{nk}^{\mathrm{m}} \alpha_k, \quad 1 \le n \le N,
\end{aligned}
$$

where the matrix $M^{\mathrm{m}} \in \mathbb{R}^{N \times K}$ is the *magnetic lead field matrix*.

The EEG inverse source problem is to estimate the vector $\alpha \in \mathbb{R}^K$ from the noisy observations of the voltage potential $u \in \mathbb{R}^L$, while in the MEG inverse source problem the data consist of the noisy observations of the magnetic field component vector $v \in \mathbb{R}^N$.

In our model, we assume that the current elements $\mathbf{j}_k$ that constitute the basis for the distributed current are either dipoles or dipole-like vector-valued elements.

3. IMAGING IN THE BAYESIAN FRAMEWORK

In the Bayesian framework, inverse problems are recast in the form of statistical inference [6, 15]. The lack of information about any of the quantities appearing in the formulation of the problem is expressed by modelling them as random variables, and the available information is encoded in the probability densities.

In the electromagnetic inverse source problem, the goal is to estimate the coefficient vector α from the observations

$$b = M\alpha + e,$$

where M is either the electric or magnetic lead field matrix and e is noise that, for simplicity, is modelled here as additive. Although not necessary, we assume that the noise is white Gaussian with known variance σ^2, which we assume known, leading to a likelihood model of the form

$$\pi(b \mid \alpha) \propto \exp\left(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2\right).$$

We consider prior models that are conditionally Gaussian and of the form

$$\pi_{\mathrm{prior}}(\alpha \mid \theta) \propto \exp\left(-\frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2 - \frac{1}{2}\sum_{j=1}^{K}\log\theta_j\right),$$

where D_θ is a diagonal matrix, $D_\theta = \mathrm{diag}(\theta_1, \ldots, \theta_K)$, and the logarithmic term comes from normalizing of the prior density by the determinant of $D_\theta^{-1/2}$. The posterior density conditional on θ is, by Bayes' formula,

$$\begin{aligned}
\pi(\alpha \mid b, \theta) &\propto \pi_{\mathrm{prior}}(\alpha \mid \theta)\pi(b \mid \alpha) \\
&\propto \exp\left(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2 - \frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2 - \frac{1}{2}\sum_{j=1}^{K}\log\theta_j\right).
\end{aligned}$$

Assuming the variance vector θ known and fixed, the MAP estimate for α,

$$\alpha_{\mathrm{MAP}} = \mathrm{argmin}\left(\frac{1}{2\sigma^2}\|b - M\alpha\|^2 + \frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2\right),$$

is the classical Tikhonov regularized solution with a penalty defined by the diagonal matrix D. It is known that if θ has equal entries, this solution is smeared out even if the data corresponds to a focal input. To improve the localization, non-quadratic penalties, for example the ℓ^p-norm, $p < 2$, of the coefficient vector α, have been proposed. Here, we take a different approach assuming instead that the variance vector θ is unknown, and thus making its estimation a part of the inverse problem. The variance vector θ is modelled as a random variable, and available a priori information concerning it is expressed by a hyperprior $\pi_{\mathrm{hyper}}(\theta)$. The prior probability density of the pair (α, θ) is then

$$\pi_{\mathrm{prior}}(\alpha, \theta) = \pi_{\mathrm{hyper}}(\theta)\pi_{\mathrm{prior}}(\alpha \mid \theta),$$

and, according to Bayes' formula, the posterior probability density, conditioned on the observation b alone, becomes

$$\pi(\alpha, \theta \mid b) \propto \pi_{\mathrm{hyper}}(\theta)\pi_{\mathrm{prior}}(\alpha \mid \theta)\pi(b \mid \alpha).$$

This implies that in the present formulation we need to estimate both α and its prior variance vector θ.

4. HYPERMODELS: MCE, MINIMUM ℓ^p AND BEYOND

Prior densities require quantitative information about the unknown, e.g., an estimate for its mean and its dynamical range, which is in turn related to the prior variance. The flexibility of hypermodels lie in their ability to import *qualitative* information into the estimation process, see [5, 7, 4, 6]. In this article, we assume that the only a priori information concerning the impressed current is that it should consist of few focal sources. In statistical terms, such information can be expressed by the following three statements:

(1) Nearby source current elements, should not be, a priori, mutually dependent, to favor the focality;
(2) No location preference for the activity should be given a priori;
(3) Most of the dipole-like sources should be silent, while few of them could have large amplitude.

To encode these conditions into the hyperprior, we observe first that stochastic dependence among the variances θ_k would couple the dynamical ranges of the corresponding coefficients α_k. Therefore, it is reasonable to assume that the variances θ_k are mutually independent. On the other hand, without prior knowledge about the location of the active sources, it is also reasonable to assume a priori that the variances are equally distributed. Furthermore, we want to allow the distribution of the variances to favor small values while permitting rare large outliers which correspond to large amplitude of the source. Among the wealth of distributions that meet these requirements, we choose the parametric family of distributions, known as the *generalized gamma distribution*, $\theta_k \sim \mathrm{GenGamma}(r, \beta, \theta_0)$, defined as

$$
\begin{aligned}
\pi_{\mathrm{hyper}}(\theta) \quad &= \quad \pi_{\mathrm{hyper}}(\theta; r, \beta, \theta_0) \\
(4) \qquad &\propto \quad \prod_{k=1}^{K} \theta_k^{r\beta-1}\exp\left(-\frac{\theta_k^r}{\theta_0^r}\right) = \exp\left(-\sum_{k=1}^{K}\frac{\theta_k^r}{\theta_0^r} + (r\beta-1)\sum_{k=1}^{K}\log\theta_k\right).
\end{aligned}
$$

In particular, we remark that by choosing $r = 1$, we have the *gamma distribution*, $\theta_k \sim \mathrm{Gamma}(\beta, \theta_0) = \mathrm{GenGamma}(1, \beta, \theta_0)$,

$$\pi_{\mathrm{hyper}}(\theta) \propto \prod_{k=1}^{K} \theta_k^{\beta-1}\exp\left(-\frac{\theta_k}{\theta_0}\right) = \exp\left(-\sum_{k=1}^{K}\frac{\theta_k}{\theta_0} + (\beta-1)\sum_{k=1}^{K}\log\theta_k\right),$$

while with $r = -1$, we obtain $\theta_k \sim \mathrm{InvGamma}(\beta, \theta_0) = \mathrm{GenGamma}(-1, \beta, \theta_0)$ which is the *inverse gamma distribution*, and

$$\pi_{\mathrm{hyper}}(\theta) \propto \prod_{k=1}^{K} \theta_k^{-\beta-1}\exp\left(-\frac{\theta_0}{\theta_k}\right) = \exp\left(-\sum_{k=1}^{K}\frac{\theta_0}{\theta_k} - (\beta+1)\sum_{k=1}^{K}\log\theta_k\right).$$

For gamma and inverse gamma distributions, the parameters β and θ_0 are referred to as *shape parameter* and *scaling parameter*, respectively. A discussion of the similarities and differences of these two distributions, in particular with regard

to the frequency of occurrence and value of outliers, can be found in [5], where we also propose the following Iterative Alternating Sequential (IAS) algorithm for computing the MAP estimate, $(\alpha_{\mathrm{MAP}}, \theta_{\mathrm{MAP}}) = \mathrm{argmax}\{\pi(\alpha, \theta \mid b)\}$:

   IAS MAP estimation algorithm:

   (1) Initialize $\theta = \theta^0$ and set $i = 1$;
   (2) Update α by defining $\alpha^i = \mathrm{argmax}\{\pi(\alpha \mid b, \theta^{i-1})\}$;
   (3) Update θ by defining $\theta^i = \mathrm{argmax}\{\pi(\theta \mid b, \alpha^i)\}$;
   (4) Increase i by one and repeat from 2. until convergence.

In the algorithm, following Bayes' formula, the conditional posterior probabilities are obtained as $\pi(\alpha \mid b, \theta^{i-1}) \propto \pi(\alpha, \theta^{i-1} \mid b)$ and $\pi(\theta \mid b, \alpha^i) \propto \pi(\alpha^i, \theta \mid b)$, i.e., alternatingly the vector α or the vector θ in the expression for the posterior density set to the current estimate.

The IAS algorithm has been previously applied to image and signal deblurring [6, 5, 7], and it is found to give a fast and stable algorithm that is easy to implement.

Before discussing how to organize the computations for different choices of the hyperprior parameter r in (4), pointing out connections with known algorithms as appropriate, we want to point out a difference between our approach and previously proposed ones.

In the literature of Bayesian hierarchical models, the gamma distribution is often suggested as a hyperprior for the precision, or inverse of the variance, because of its conjugacy property. This corresponds to the inverse gamma distribution for the variance. The conjugacy property is useful, e.g., when variational Bayes methods, or variants of the closely related EM algorithm, are used and analytic marginal integrals are desired (Rao-Blackwellization), see, e.g., [17]. Relevant references for the MEG problem are [26, 33]. Interestingly, in [26] the gamma hyperprior for the precision was suggested but the connection with the regularization methods was not pointed out. We emphasize that our approach neither needs the conjugacy, nor does it take advantage of it. In fact, as we will show below, the IAS algorithm yields fast, efficient and explicit estimators with a large range of parameter values corresponding to non-conjugate models. The generalized gamma distribution is chosen here solely on the basis that it allows rare outliers.

Gamma distribution and Minimum Current Estimate. Consider the posterior density of the pair (α, θ) when the hyperprior is the gamma distribution

$$\pi(\alpha, \theta \mid b) \propto \exp\Bigg(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2$$
$$-\frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2 - \frac{1}{\theta_0}\sum_{k=1}^K \theta_k + \Big(\beta - \frac{3}{2}\Big)\sum_{k=1}^K \log\theta_k \Bigg).$$

To solve the first maximization problem in the IAS MAP estimation algorithm, set $\theta = \theta^{i-1}$ and let the updated α^j be the minimizer of the negative of the log-posterior,

$$(5) \qquad \alpha^i = \mathrm{argmin}\left(\frac{1}{2\sigma^2}\|b - M\alpha\|^2 + \frac{1}{2}\|D_{\theta^{i-1}}^{-1/2}\alpha\|^2 \right),$$

which is the least squares solution of the linear system

$$(6) \qquad \begin{bmatrix} \sigma^{-1}M \\ D_{\theta^{i-1}}^{-1/2} \end{bmatrix} \alpha = \begin{bmatrix} \sigma^{-1}b \\ 0 \end{bmatrix}.$$

Subsequently, setting $\alpha = \alpha^i$, the updated value of θ can be found by differentiating the log-posterior with respect to θ and setting the derivative equal to zero. The resulting equation is a second order equation

$$\frac{1}{2}\alpha_j^2/\theta_j^2 - 1/\theta_0 + \eta/\theta_j = 0, \quad \eta = \beta - 3/2, \quad \alpha_j = \alpha_j^i,$$

for the positive root, whose analytic expression is

$$\theta_j^i = \frac{1}{2}\theta_0\left(\eta + \sqrt{\eta^2 + 2\alpha_j^2/\theta_0}\right).$$

In particular, if we let $\eta = 0$, we have $\theta_j^i = |\alpha_j^i|\sqrt{\theta_0/2}$, and, by substituting this in (5), we obtain

$$\alpha^i = \operatorname{argmin}\left(\frac{1}{2\sigma^2}\|b - M\alpha\|^2 + \frac{1}{\sqrt{2\theta_0}}\sum_{k=1}^{K}\frac{\alpha_k^2}{|\alpha_k^{i-1}|}\right),$$

which is also a fixed point iterate of the minimization problem whose solution is the Minimum Current Estimate (MCE) [31]

$$\alpha_{\mathrm{MC}} = \operatorname{argmin}\left(\|b - M\alpha\|^2 + \delta\sum_{k=1}^{K}\frac{\alpha_k^2}{|\alpha_k|}\right), \quad \delta = \sqrt{\frac{2}{\theta_0}}\sigma^2.$$

Observe that, by choosing $\beta > 3/2$ in the hyperprior, we avoid the problem of dividing by components α_k^{i-1} near or equal to zero. This hence provides a natural regularization method for solving the Minimum Current Estimate problem.

Generalized gamma distribution and ℓ^p–estimates. The choice of the hyperprior from the family of generalized gamma distributions leads to the posterior model

$$\pi(\alpha, \theta \mid b) \propto \exp\Bigg(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2$$

$$-\frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2 - \frac{1}{\theta_0^r}\sum_{k=1}^{K}\theta_k^r + \left(r\beta - \frac{3}{2}\right)\sum_{k=1}^{K}\log\theta_k\Bigg).$$

The updating of α given the current value θ^{i-1} requires solving the system (6) as in the case of gamma hyperprior, while the updating formula for the variance parameter changes. Setting the derivative of the logarithm of the posterior density equal to zero leads to the algebraic equation

$$\frac{1}{2}\alpha_j^2/\theta_j^2 - r\theta_j^{r-1}/\theta_0^r + (r\beta - 3/2)/\theta_j = 0, \qquad \alpha_j = \alpha_j^i.$$

This equation does not have, in general, a closed form solution, although when $r\beta = 3/2$ the solution is simply

$$\theta_j^i = (\theta_0^r\alpha_j^2/2r)^{1/(r+1)}.$$

As in the case of the gamma distribution, after substituting this solution into the objective function in (5) we notice that the updated α^i is a fixed point iterate of the ℓ^p–penalized regularization problem,

$$\alpha_p = \operatorname{argmin}\left(\|b - M\alpha\|^2 + \delta\sum_{k=1}^{K}|\alpha_k|^p\right), \quad \delta = 2\sigma^2\left(\frac{2r}{\theta_0^r}\right)^{1/(r+1)},$$

with $r = p/(2-p)$. It is known that when $0 < p < 1$, i.e., $0 < r < 1$, this solution, like the MCE solution, tends to minimize the support of the estimated current, thus yielding a good localization of the focal activity. On the other hand, letting $r \to \infty$, or, equivalently, $p \to 2$, the MAP estimate approaches the Tikhonov regularized solution with a quadratic penalty. The intermediate case $1 < p < 2$ is related to the analysis presented in [1].

Inverse gamma distribution and Minimum Support Estimate. The posterior density for the inverse gamma hypermodel is of the form

$$\pi(\alpha, \theta \mid b) \propto \exp\left(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2\right.$$
$$\left. -\frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2 - \theta_0\sum_{k=1}^{K}\frac{1}{\theta_k} - \left(\beta + \frac{3}{2}\right)\sum_{k=1}^{K}\log\theta_k\right).$$

Once again, the updated value α^i is found by solving (6) keeping θ fixed to the current value θ^{i-1}, while θ_j is the zero of the derivative of the logarithm of the posterior which satisfies

$$\frac{1}{2}\alpha_j^2/\theta_j^2 + \theta_0/\theta_j^2 - \kappa/\theta_j = 0,$$

with $\kappa = \beta + 3/2$, and $\alpha_j = \alpha_j^i$, and can be expressed as

$$\theta_j^i = \left(\frac{1}{2}\alpha_j^2 + \theta_0\right)/\kappa.$$

By interpreting this algorithm as a fixed point step of a regularization scheme with a nonlinear penalty term, we can reformulate it as the following minimization problem

$$\alpha = \operatorname{argmin}\left(\|b - M\alpha\|^2 + \delta\sum_{k=1}^{K}\frac{(\alpha_k)^2}{(\alpha_k)^2 + 2\theta_0}\right), \quad \delta = 4\kappa\sigma^2,$$

whose solution is the Minimum Support Estimate (MSE) [20]. In [6, 5, 7], the authors have shown that, in the context of traditional image processing, the corresponding penalty is related to the Perona-Malik functional [23].

Higher order inverse gamma distributions. The gamma and inverse gamma distributions are standard models in hierarchical Bayesian methods and the MAP estimates correspond to known regularizing schemes. The combination of the generalized gamma distribution and the IAS algorithm provide models that are computationally tractable and lead to new estimators. As an example, if we choose

$r = -q$, where $q > 1$ is an integer, the posterior distribution is

$$\pi(\alpha, \theta \mid b) \propto \exp\Bigg(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2$$
$$-\frac{1}{2}\|D_\theta^{-1/2}\alpha\|^2 - \theta_0^q \sum_{k=1}^K \frac{1}{\theta_k^q} - \left(q\beta + \frac{3}{2}\right)\sum_{k=1}^K \log\theta_k \Bigg).$$

The algebraic equation for updating θ_k is

$$q\theta_0^q + \frac{1}{2}x_k^2\theta_k^{q-1} - \left(q\beta + \frac{3}{2}\right)\theta_k^q = 0,$$

i.e., θ_k is a positive root of a polynomial of order q, and can be computed in a stable way by considering the companion matrix.

5. Hyperpriorconditioners

Direct computation of the least squares solution of the linear system (6), which yields the updated α, becomes prohibitively slow when the dimensionality of the problem is large, as is the case in the applications that we are considering here. Iterative methods are the method of choice for the solution of large scale linear systems. In the case of linear discrete ill-posed problems, this approach is particularly attractive because of the inherent regularizing properties of iterative solvers when equipped with suitable stopping criteria. Iterative linear system solvers start from a given approximate solution and proceed to determine a sequence of improved approximate solutions.

Assume, for the moment, that in the process of updating α we introduce the change of variable $w = D_{\theta_c}^{-1/2}\alpha$ where θ_c is the current value for θ, and express the solution of the optimization step in the form

$$(7) \qquad w_+ = \operatorname{argmin}\left(\frac{1}{2\sigma^2}\|b - MD_{\theta_c}^{1/2}w\|^2 + \frac{1}{2}\|w\|^2\right), \quad \alpha_+ = D_{\theta_c}^{1/2}w_+.$$

We remark that while in the statistical framework a suitable choice of the matrix D_{θ_c} makes the change of variable equivalent to whitening the random variable α, in the context of iterative linear systems solvers, this transformation amounts to preconditioning. Since in our problem the choice of the matrix $D_{\theta_c}^{1/2}$ is dictated by the selection of the prior, following [6] we refer to it as priorconditioner, to emphasize the connection between the numerical performance and the statistical setting.

After the change of variables, the solution of the least squares problem (7) for updating α coincides with the standard Tikhonov regularized solution for solving the preconditioned linear system

$$(8) \qquad\qquad \sigma^{-1}MD_{\theta_c}^{1/2}w = \sigma^{-1}b, \quad \alpha = D_{\theta_c}^{1/2}w.$$

It has been shown in the literature that iterative Krylov subspace methods such as CGLS (see [24]) with early stopping of the iterations may give results of comparable quality to Tikhonov regularization but are computationally much more efficient.

The introduction of a suitable right priorconditioner, which in the Bayesian framework is related to the prior of the unknown of primary interest, to bias the iterates towards a desirable subspace, has been shown to improve the quality of the

computed solution, in particular when the number of iterations is limited by either high computational costs or a large noise level in the data [6].

In general, statistically inspired preconditioners, which convey into the linear system solver prior beliefs about the solution, are constructed from the covariance matrix of the prior. In the application considered here, however, the prior is a function of unknown parameters, whose distribution is, in turn, given by the hyperprior. Since the parameters of the prior are themselves part of the estimation problem, they are updated at each iteration step. To ensure that the solution of the minimization problem (5) uses the most up to date information about the problem, as soon as θ_c becomes available, we update the priorconditioning matrix $D_{\theta_c}^{1/2}$, then proceed to compute a new estimate of α.

We remark that the idea of updating the preconditioner to take advantage of newly acquired information about the linear systems was already proposed in the context of a Flexible Generalized MINimal RESidual (FGMRES) scheme [24], although the motivation for the updating were quite different. In the Bayesian paradigm we can view priorconditioning in the context of hypermodels as priorconditioning *conditioned on* the present estimate of the prior parameters.

Finally, it is of interest to notice that the alternating updating scheme where the least squares problem is solved with preconditioned iterative methods (8) includes as a special case an algorithm previously proposed in the literature, based on adaptive weighting of the ℓ^2–norm. In fact, when the hyperprior is the gamma distribution with $\beta = 3/2$, the IAS is essentially the FOCUSS algorithm discussed, e.g. in [10]. While in the FOCUSS algorithm the regularization is obtained by passing from (8) to the normal equation corresponding to (7), here we advocate regularization by truncated iteration. The connections between the minimum current estimate and FOCUSS from the empirical Bayesian point of view have been pointed out also in [33].

6. MCMC and regions of interest

A great advantage of the Bayesian approach over different regularization schemes is that starting from the posterior density, we can compute a number of different estimates and furthermore quantify their reliability. The uncertainty quantification, however, usually requires sequential sampling techniques which are computationally considerably more intensive than optimization based computation of single estimates, in particular when applied to a detailed three dimensional model.

Various dimension reduction methods have been proposed in the literature to make the MCMC sampling viable. A common approach is to restrict the source sampling either to the surface of the brain or to a thin cortical layer, see, e.g., [26]. In MEG, a further reduction of the dimensionality of model may be achieved by restricting the sampling to the cortical regions with a non-radial normal vector, see, e.g., [1]. These model reductions are not applicable for us, since we consider both EEG and MEG and the possibility of recovering deep sources with the MCMC sampling.

Fortunately, in applications where we are interested in local sources, it is often sufficient to restrict the sampling to a much smaller Region Of Interest (ROI), around the potentially active area. The selection of the ROI can be based on prior information about the expected activity: for example, the primary response to a visual stimulus is expected to occur in the occipital lobe. Alternatively, the ROI

can be selected around an estimated focus of activity. Note that the concept of ROI does not exclude the possibilities of restricting the sampling to a portion of the cortical layer, as is done in the cited articles, or of sampling over the whole brain, if the computing resources are not an issue.

Once the ROI has been identified, we collect the indices of the source basis vectors $\mathbf{j}_k$ whose support is in the ROI in the index vector I_{ROI}, and the remaining ones in the vector I_0. We then partition the vectors α and θ accordingly, introducing the notation $\alpha_{\mathrm{ROI}} = \alpha(I_{\mathrm{ROI}})$, $\alpha_0 = \alpha(I_0)$, $\theta_{\mathrm{ROI}} = \theta(I_{\mathrm{ROI}})$ and $\theta_0 = \theta(I_0)$.

We can now perform MCMC sampling over the ROI, fixing the outside current values and prior variance to prescribed values α_0, θ_0, using the conditional distribution,

$$\pi(\alpha_{\mathrm{ROI}}, \theta_{\mathrm{ROI}} \mid \alpha_0, \theta_0, b) \propto \pi([\alpha_{\mathrm{ROI}}, \alpha_0], [\theta_{\mathrm{ROI}}, \theta_0] \mid b).$$

Evidently, $\alpha_0 = 0$ is the most natural choice, corresponding to the assumption that no activity outside the ROI appears. The MCMC algorithm that we propose is an independence sampling method, where α_{ROI} and θ_{ROI} are updated sequentially by a procedure analogous to that of the IAS estimation algorithm. The updating of α_{ROI} is done using the conditional normality of the posterior, while the updating of θ_{ROI} is done via a full scan Gibbs sampler [6, 15, 17].

MCMC sampling over ROI:

(1) Initialize α^0_{ROI}, θ^0_{ROI} and set $i = 1$. Define M, the desired sample size.

(2) Draw α^i_{ROI} from the Gaussian density

$$\pi(\alpha_{\mathrm{ROI}} \mid \theta^{i-1}_{\mathrm{ROI}}, \alpha_0, \theta_0, b) \propto \pi([\alpha_{\mathrm{ROI}}, \alpha_0], [\theta^{i-1}_{\mathrm{ROI}}, \theta_0] \mid b).$$

(3) Draw θ^i_{ROI} componentwise with a Gibbs sampler from the density

$$\pi(\theta_{\mathrm{ROI}} \mid \alpha^i_{\mathrm{ROI}}, \alpha_0, \theta_0, b) \propto \pi([\alpha^i_{\mathrm{ROI}}, \alpha_0], [\theta_{\mathrm{ROI}}, \theta_0] \mid b).$$

(4) If $i = M$, stop; otherwise increase i by one and repeat from 2.

In the practical implementation of the algorithm we update α_{ROI} by defining a matrix G and its partitioning,

$$G = \begin{bmatrix} \sigma^{-1} M \\ D^{-1/2}_{\theta^{i-1}} \end{bmatrix} = \begin{bmatrix} G_{\mathrm{ROI}} & G_0 \end{bmatrix},$$

where G_{ROI} contains the columns of G with indices in I_{ROI} and G_0 the remaining columns. The updated α^i_{ROI} is obtained by solving, in the least squares sense, the linear system

$$G_{\mathrm{ROI}} \alpha_{\mathrm{ROI}} = \begin{bmatrix} \sigma^{-1} b \\ 0 \end{bmatrix} - G_0 \alpha_0 + w, \quad w \sim \mathcal{N}(0, I).$$

The updating of θ_k, $k \in I_{\mathrm{ROI}}$, is performed by drawing from the one-dimensional probability density

$$\pi_k(\theta_k) \propto \exp\left(-\frac{\alpha^2_k}{2\theta_k} - \left(\frac{\theta_k}{\theta_0}\right)^r + \left(r\beta - \frac{3}{2}\right) \log \theta_k \right)$$

by the inverse cumulative distribution method [6, 15]. Hence, the sampling technique just outlined takes advantage of the conditional normality of the prior and of the mutual independence of the variances, similarly to the IAS algorithm.
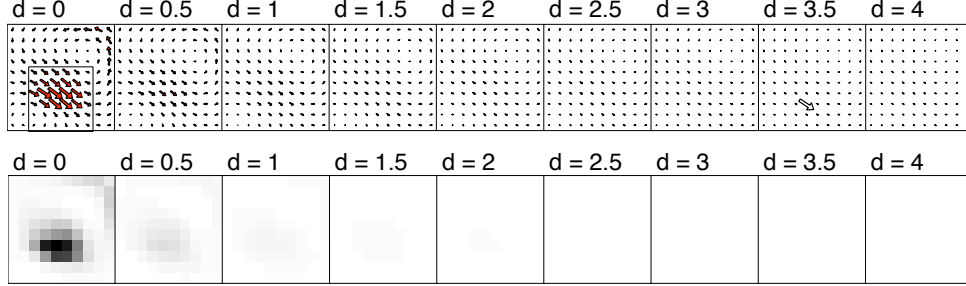
FIGURE 1. MAP estimates of the current (top row) and of the variance of the prior (bottom row) at different depth in the ROI. The hyperprior is the gamma distribution. The true current dipole used for generating simulated data is shown as a hallow arrow, and the ROI is marked on the superficial layer.
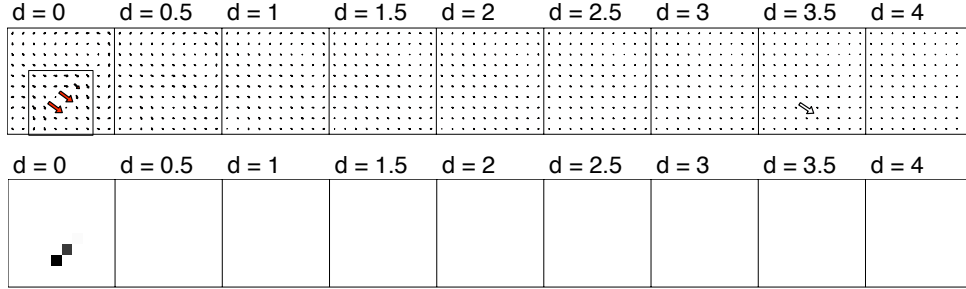

FIGURE 2. MAP estimates of the current (top row) and of the variance of the prior (bottom row) at different depth in the ROI. The hyperprior is the inverse gamma distribution. The true current dipole used for generating simulated data is shown as a hallow arrow, and the ROI is marked on the superficial layer.

7. COMPUTED EXPERIMENTS

In this section we apply the methodology derived above to inverse source problems by first considering an example with a simplified planar geometry using a traditional singular dipole model, then applying it to a realistic conductivity model for the human head. In the latter case, both the MEG and EEG modalities are considered. Since the finite element method (FEM) is needed to solve the potential distribution, we use FEM basis functions also for representation of the current density.

MEG in planar geometry. Consider a half space as a local model for the human head. The half space model is particularly appropriate to illustrate the depth
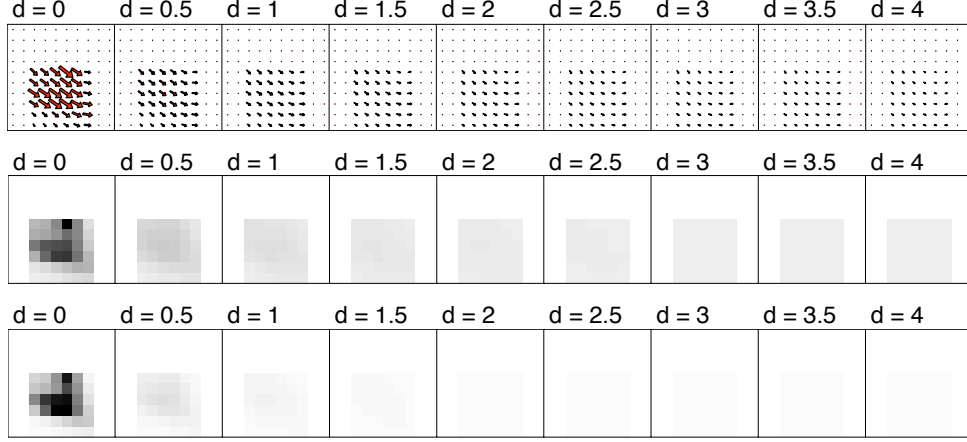
FIGURE 3. Posterior mean estimates of the current (top row), of the variance of the prior (center row), and of the variance of the norm of the current estimated over the sample (bottom row). The hyperprior is the gamma distribution.
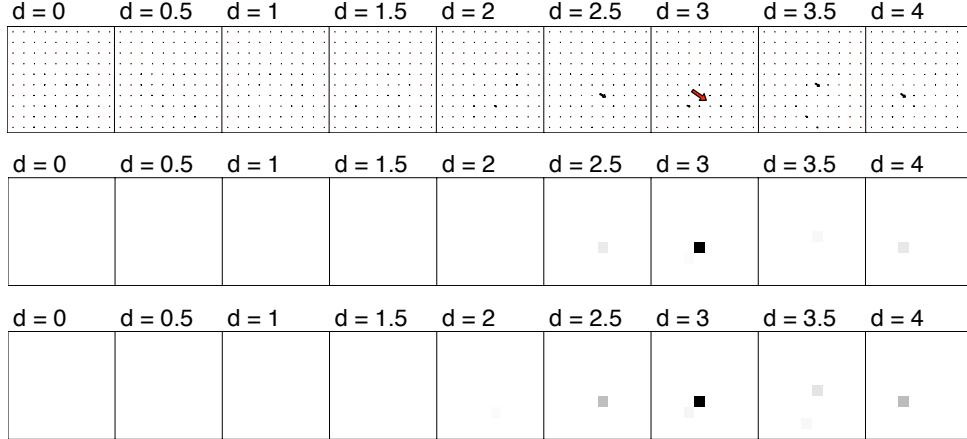

FIGURE 4. Posterior mean estimates of the current (top row), of the variance of the prior (center row) and the variance of the norm of the current estimated over the sample (bottom row). The hyperprior is the inverse gamma distribution.

resolution with different hypermodel parameters and with different statistical estimators.

The magnetic field component perpendicular to the surface is recorded in a rectangular array of observation points x_ℓ above the surface $z = 0$. We represent
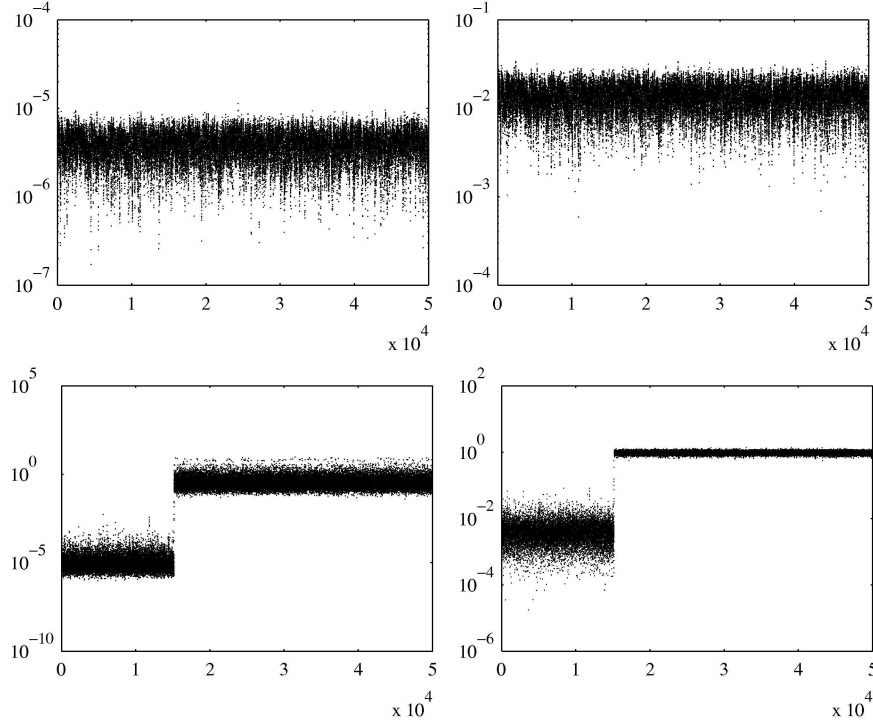
FIGURE 5. Sample histories of the components θ_j and $\|\mathbf{q}_j\|$ for the gamma hyperprior (top row) and the inverse gamma (bottom row), where j is the index of the component of maximum value of θ_{CM}.

the current density as a linear combination of point-wise current dipoles,

$$\mathbf{J}(y) = \sum_{j=1}^{K} \delta(y - y_k)(\alpha_k^1 \mathbf{e}_1 + \alpha_k^2 \mathbf{e}_2) = \sum_{j=1}^{K} \delta(y - y_k)\mathbf{q}_k,$$

where $\mathbf{e}_1$ and $\mathbf{e}_2$ are orthogonal basis vectors parallel to the plane $z = 0$. In this example we ignore the volume currents, which is tantamount to setting $\sigma = 0$. More generally, assuming a conductivity density that depends only on the depth [25], we arrive at a particularly simple magnetic lead field model,

$$b_\ell = \frac{\mu_0}{4\pi} \sum_{k=1}^{K} \sum_{j=1}^{2} \frac{\mathbf{e}_3 \cdot (\mathbf{e}_j \times (x_\ell - y_k))}{|x_\ell - y_k|^3} \alpha_k^j, \quad \text{or} \quad b = \begin{bmatrix} M_1^{\mathrm{m}} & M_2^{\mathrm{m}} \end{bmatrix} \begin{bmatrix} \alpha^1 \\ \alpha^2 \end{bmatrix},$$

with $\alpha = [\alpha^1; \alpha^2] \in \mathbb{R}^{2K}$. When writing the conditional prior, we identify the variances θ_k^1 and θ_k^2 of the two mutually perpendicular dipoles whose locations coincide, and the unknown variance $\theta_k^1 = \theta_k^2 = \theta_k$ is a vector of length K. The

posterior density then becomes

$$\pi(\alpha, \theta \mid b) \propto \exp\left(-\frac{1}{2\sigma^2}\|b - M\alpha\|^2 \right.$$
$$\left. -\frac{1}{2}\sum_{j=1}^{K}\frac{\|\mathbf{q}_k\|^2}{\theta_k} - \sum_{k=1}^{K}\left(\frac{\theta_k}{\theta_0}\right)^r + \left(r\beta - \frac{3}{2}\right)\sum_{k=1}^{K}\log\theta_k \right),$$

with $\|\mathbf{q}_k\|^2 = (\alpha_k^1)^2 + (\alpha_k^2)^2$. We use the IAS algorithm to calculate the MAP estimate, and MCMC sampling over the ROI to estimate the CM and obtain a measure of its uncertainty.

The geometry of our model consists of a rectangular array of 10×10 vertical magnetometers 2 cm above the half space, with a distance between adjacent magnetometers of 1 cm. The dipoles are located below the magnetometers in nine horizontal layers, each containing a 10×10 rectangular array of dipoles. The depth of these layers varies from zero (superficial sources) to 4 cm, with a distance between the layers of 0.5 cm.

Since the MAP estimation algorithm is tantamount to fixed point iteration with localizing penalties, we expect good performance at detecting focal sources of known depth. It is well known [11, 16, 31] that when the depth of the source is unknown, the minimum current and minimum norm estimates due to their tendency to bias towards superficial sources may lead to gross misplacements of the deep focal sources. We expect the same behavior from the MAP estimates of our hypermodel. We test this by generating synthetic data in which a single dipole source is placed 3.5 cm below the surface of the half space. The standard deviation in the likelihood model was assumed to be 5% of the maximum noiseless signal. In this simulation, we did not add artificial noise to the data, since we are only interested in the model bias, not in the noise sensitivity.

The MAP estimates for the dipole fields as well as for the prior variance θ with model parameter values $r = 1$ (gamma) and $r = -1$ (inverse gamma) are shown in Figures 1 – 2. When $r = 1$, we use the values $\theta_0 = 10^{-7}$ and $\beta = 3$ for the scaling parameter and the shape parameter, while when $r = -1$, we set $\theta_0 = 10^{-5}$ and $\beta = 3$. In both cases we perform 15 iterations with the IAS algorithm.

As expected, both hypermodels favor superficial sources, with a relatively good localization in the horizontal direction, i.e., the MAP estimate of the activity is above the true source. The major differences between the two hypermodels are in the convergence rate and in the focality of the MAP estimate. With the inverse gamma hypermodel the iterative algorithm converges faster than with the gamma hypermodel, in particular for non-superficial sources, and seeks to explain the data with fewer active superficial dipoles.

To reduce the dimensionality of the sampling space, we select the ROI to be a cylinder with a 6×6 cm^2 square base around the estimated superficial focal activity, shown in Figure 1 – 2, containing $9 \times 6 \times 6 = 324$ dipoles. For each hyperparameter model we generate a sample of size $M = 50\,000$, conditional on the currents vanishing outside the ROI, and calculate estimates of the posterior means of the vectors α and θ, $\alpha_{\mathrm{CM}}^j = \frac{1}{M}\sum_{i=1}^{M}\alpha^{j,i}$, $j = 1, 2$, and $\theta_{\mathrm{CM}} = \frac{1}{M}\sum_{i=1}^{M}\theta^i$, and of the posterior variance of the dipole amplitudes, $\mathrm{Var}(\|\mathbf{q}_k\|) = \frac{1}{M}\sum_{i=1}^{M}\{(\alpha^{1,i} - \alpha_{\mathrm{CM}}^1)^2 + (\alpha^{2,i} - \alpha_{\mathrm{CM}}^2)^2\}$.

TABLE 1. The domains of the head model and respective conductivities.

Layer	Conductivity
scalp	0.33
skull	0.0042
cerebrospinal fluid	1
brain	0.33

Figure 3 displays the plots of the posterior mean of the current and the estimates of the prior variance, and the posterior variance of the amplitude with the gamma hyperprior, while Figure 4 shows the analogous results for the inverse gamma hyperprior. The results with the two hyperpriors are qualitatively very different. The posterior mean of the current for the gamma hyperprior model is biased towards the surface and is very similar to the MAP estimate, while the inverse gamma hypermodel has a good depth resolution, with an error between the true and the estimated source depth of 0.5 cm. The observation that the gamma hyperprior leads to a posterior density that is qualitatively closer to the Gaussian prior, whose mean and maximum coincide, is in line with the fact that, as $r \to \infty$, the MAP estimate approaches the minimum norm, or Tikhonov regularized, solution. The latter corresponds to an ℓ^2 prior model.

The sample histories of single components $\|\mathbf{q}_j\| = ((\alpha_j^1)^2 + (\alpha_j^2)^2)^{1/2}$ and θ_j reveal an interesting feature of the posterior density. Figure 5 shows the sample histories of the components corresponding to the maximum value of θ_{CM}. The sample histories with the gamma hyperprior exhibit good mixing, while those corresponding to the inverse gamma distribution seem to suggest a bimodal posterior density.

Note that the posterior mean of the prior variance θ is in good agreement with the posterior variance of the current amplitude, implying that the reliability of the posterior mean current in this geometry could be assessed directly from the mean of the variance parameter θ.

EEG/MEG in realistic geometry. To validate the hyperprior models in a more realistic geometry, we performed EEG and MEG tests with a realistic human head model, based on MRI data. We assume that both electric and magnetic fields are measured outside the skull at 31 different locations, as shown in Figure 6. This model is partitioned into four domains of constant electric conductivity: scalp, skull, cerebrospinal fluid (CSF), and brain, as illustrated in Figure 6. Anisotropies in the brain as well as possible electroconductive differences between the gray and the white matter are ignored.

The electric and magnetic lead field matrices for this setup were constructed using the complete electrode model and a Finite Element Method (FEM), described in the Appendix, under the assumption that all electrode contact impedances are equal to one. The FEM mesh was generated in two stages. The meshes for the skull and the brain domains were generated first using triangular elements. The meshes for the scalp and CSF domains were then generated by positioning prism elements, each to be subsequently divided into three tetrahedral elements, between the brain and the skull surface. The total number of tetrahedral elements in the head mesh is 108 914. The electric conductivities of the domains are given in Table 1. These values are as in [32].
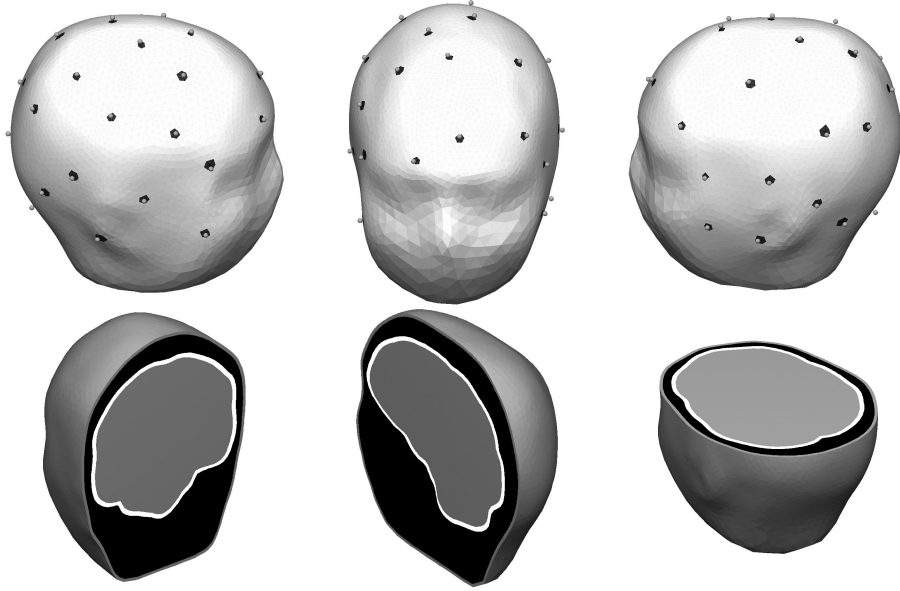
FIGURE 6. Sagittal, coronal, and axial projections (left to right) of the head model (top row) and the conductivity distribution inside it (bottom row). The electric field is measured using 31 contact electrodes marked by dark grey surface patches. The 31 magnetic field measurement locations are indicated by the lighter dark spheres over the patches.
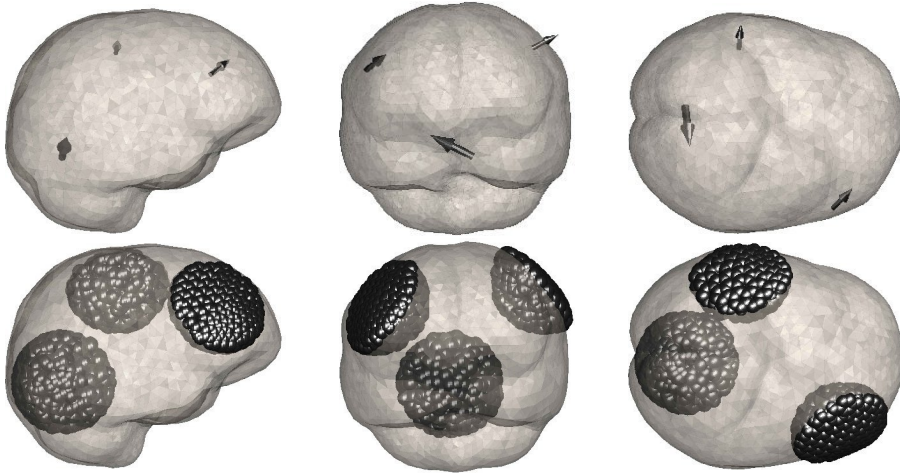

FIGURE 7. Sagittal, coronal, and axial projections (left to right) of the reference current density (top row) and of the region of interest (bottom row).

Lowest order $H(\mathrm{div})$-conforming Raviart-Thomas elements were applied for the source current density. In these elements, basis functions are linear over the tetrahedron, vanish at one of the vertices, and have a constant direction normal to the
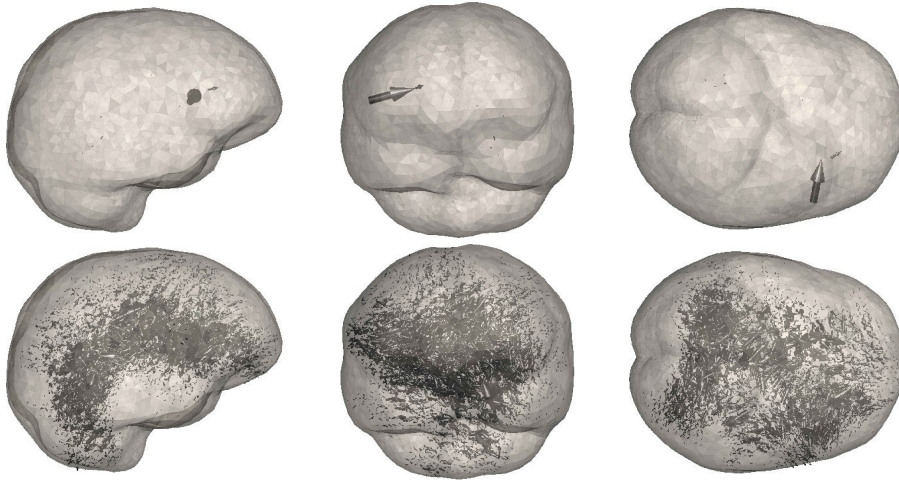
FIGURE 8. Sagittal, coronal, and axial projections (left to right) of the MAP estimate from EEG data of the current for the gamma hypermodel (top row) and of the inverse gamma hypermodel (bottom row). To improve the readability of the three dimensional plots, only the current elements whose amplitude is above 5% of the maximum of the amplitudes in the estimate were plotted.
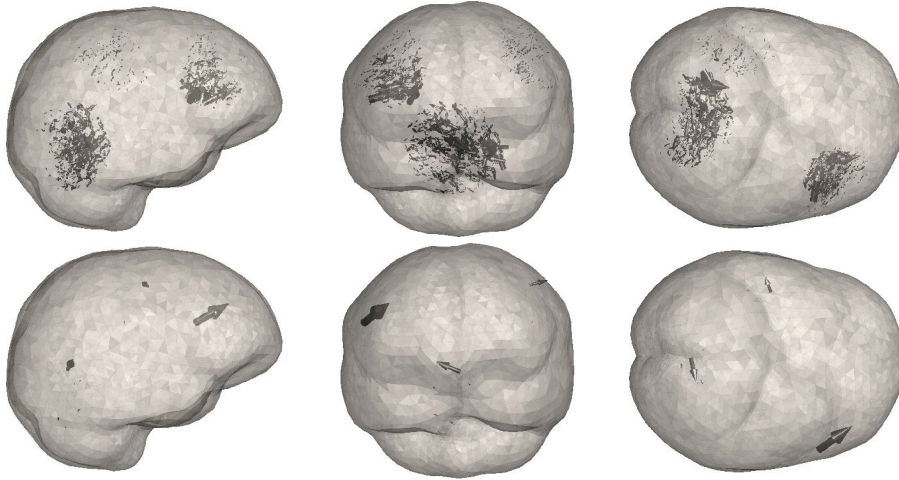

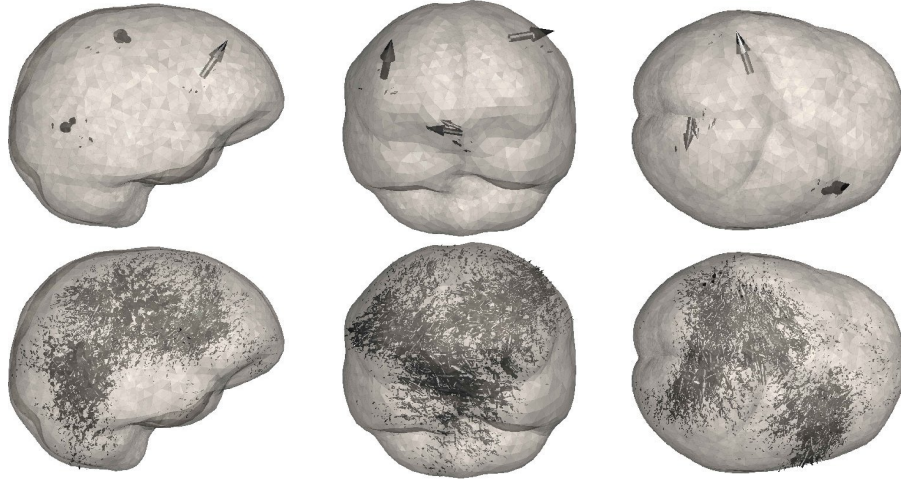FIGURE 9. Sagittal, coronal, and axial projections (left to right) of the CM estimate from EEG data of the current for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row).

face opposite to the vertex where they vanish [3]. For a curl-free current density, e.g. a dipole source, it is necessary to use $H(\mathrm{div})$-conforming elements [18], i.e. elements in which basis functions and their divergences are square integrable. Basis functions

FIGURE 10. Sagittal, coronal, and axial projections (left to right) of the MAP estimate from MEG data of the current for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row).
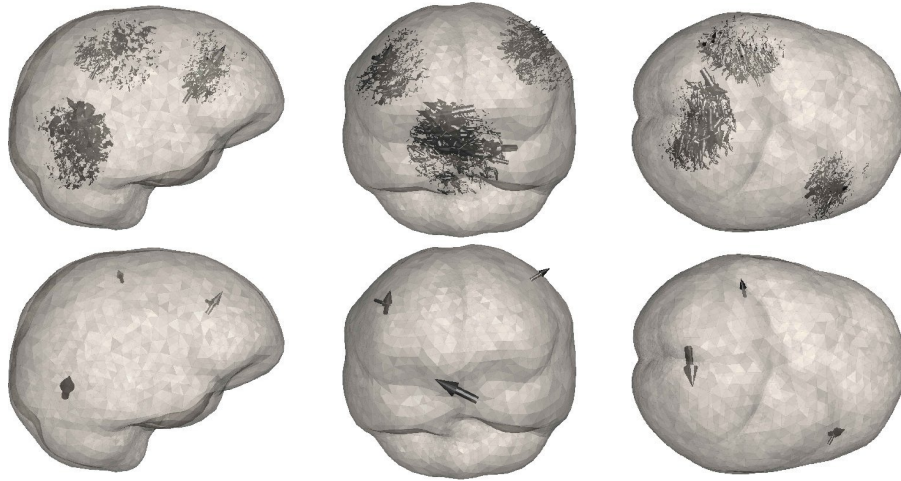

FIGURE 11. Sagittal, coronal, and axial projections (left to right) of the CM estimate from MEG data of the current for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row).

for $H(\mathrm{div})$-conforming elements can be interpreted to represent dipole-like currents [29].

The electric potential u was modelled using Lagrange elements [3]. Quadratic Lagrange elements and a fifteen point Gauss quadrature rule [3] were employed to generate simulated, noiseless reference data, while the exploration of the posterior was done using linear Lagrange elements and a four point Gauss quadrature to
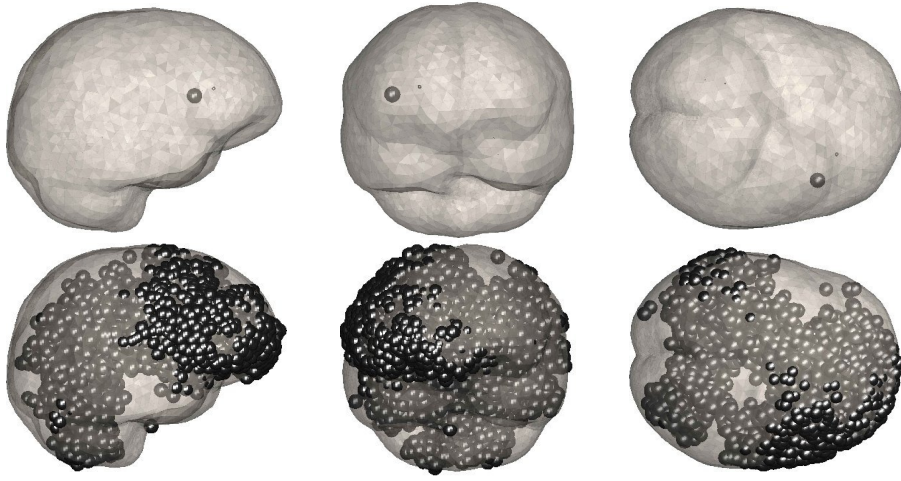
FIGURE 12. Sagittal, coronal, and axial projections (left to right) of the MAP estimate from EEG data of the variance of the prior for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row). In the top row, location of the components of θ larger than 5% of the maximum value are marked with a bubble. In the bottom row only elements whose values were greater than or equal to 85% of the maximum were plotted, indicating that the estimated variance is nearly maximal over the entire brain.
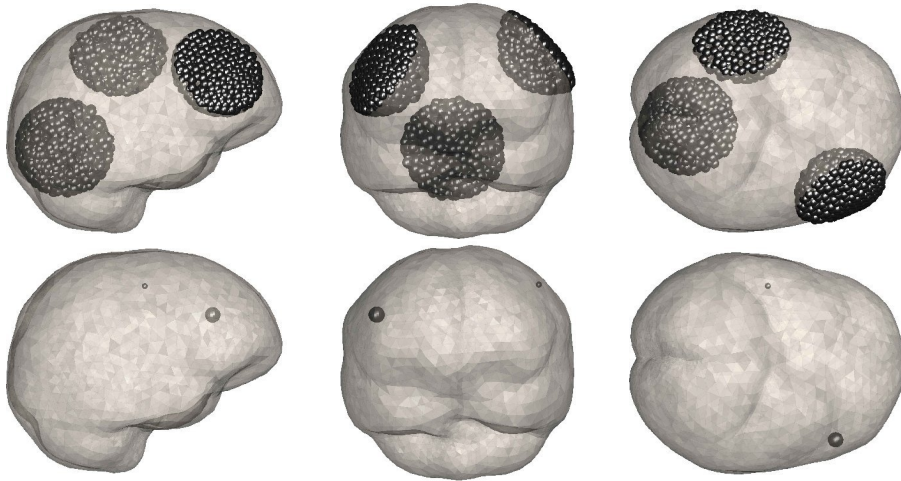

FIGURE 13. Sagittal, coronal, and axial projections (left to right) of the CM estimate from EEG data of the variance of the prior for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row). The thresholding level is set to 5% in this plot.

speed up the computation. The resulting forward modelling error for the electric
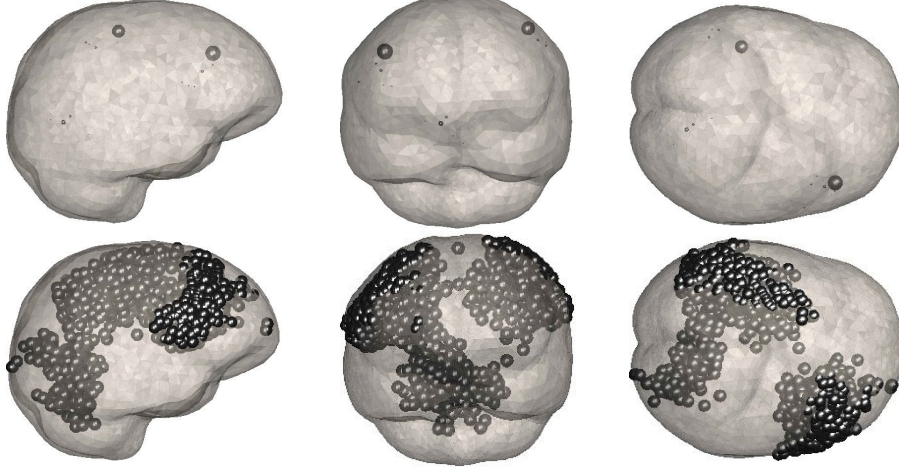
FIGURE 14. Sagittal, coronal, and axial projections (left to right) of the MAP estimate from MEG data of the variance of the prior for the gamma hypermodel (top row) and inverse gamma hypermodel (bottom row). The thresholding in the visualization is as in Figure 12.
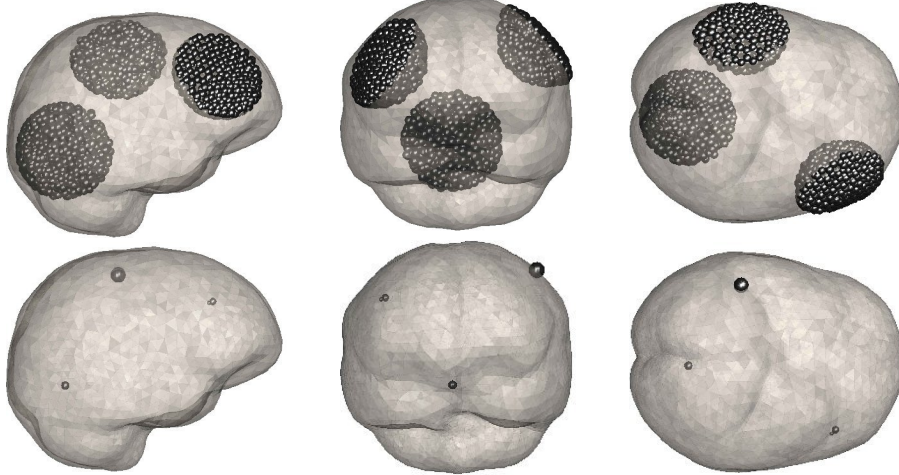

FIGURE 15. Sagittal, coronal, and axial projections (left to right) of the CM estimate from MEG data of the variance of the prior for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row). Here, the visualization threshold is set to 5%.

lead field matrix in the brain domain was approximately 3 % in the Frobenius norm [3], and approximately 2 % for the magnetic lead field matrix.

The reference current density used to generate the reference data, shown in Figure 7, corresponds to the case where we have three dipole-like source currents, one positioned deeply, approximately 2.5 cm under the occipital lobe, and the other
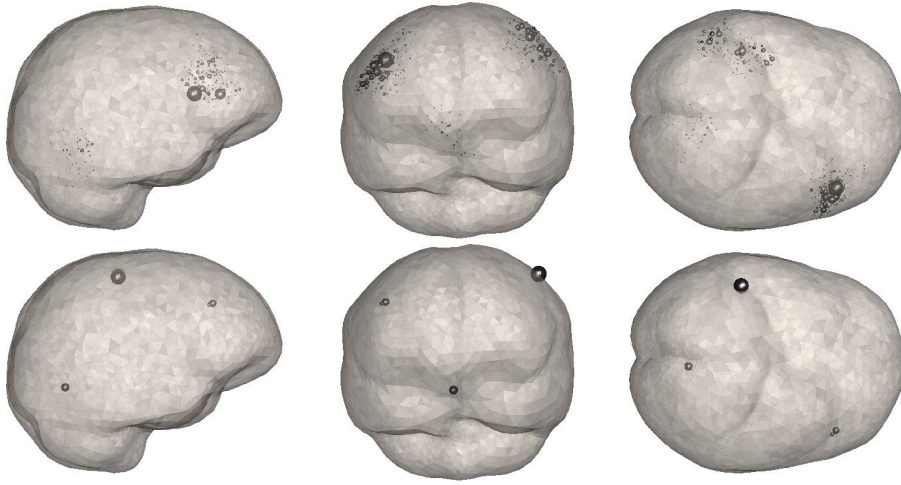
FIGURE 16. Sagittal, coronal, and axial projections (left to right) of the sample based estimates of the posterior variance from EEG data for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row). The visualization threshold is set at 5% here.

FIGURE 17. Sagittal, coronal, and axial projections (left to right) of the sample based estimates of the posterior variance from MEG data for the gamma hypermodel (top row) and for the inverse gamma hypermodel (bottom row). The visualization threshold is set to 5%.
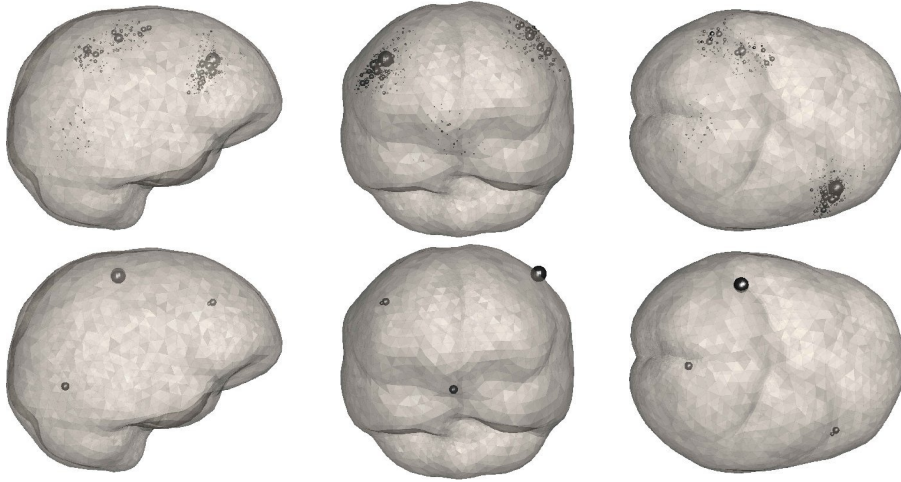
two on the surface layer modelling the cortex. The source on the right of the frontal lobe is almost tangential to the surface, and the source close to the left central sulcus is almost normal to the surface.
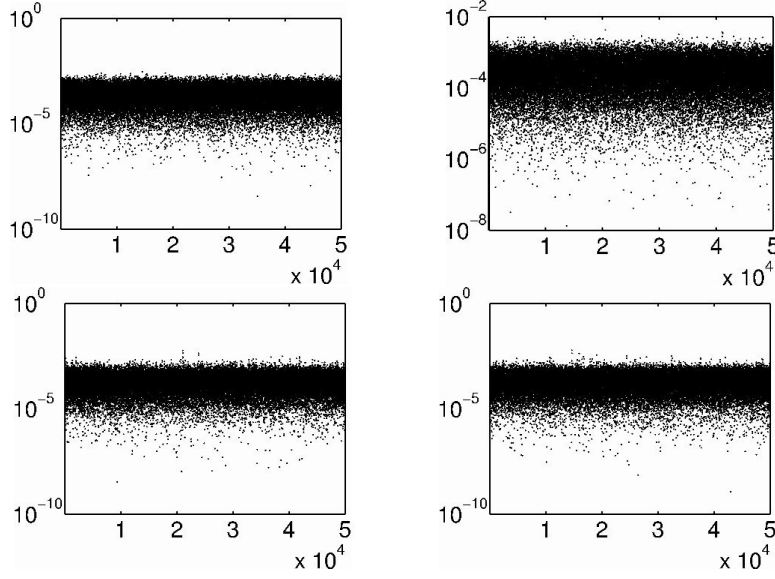
FIGURE 18. Sample histories of the component corresponding to the maximum of the posterior mean estimate of the current estimate from the EEG data and the gamma hypermodel (top left) and with the inverse gamma hypermodel (bottom left). In the right column we show the corresponding plots obtained from the MEG data and the gamma distribution (top right) or the inverse gamma distribution (bottom right).

As in the planar geometry example, the standard deviation in the likelihood model is assumed to be 5% of the maximum noiseless signal. In this case, to simulate real situations, the noise is added to the generated data.

The MAP estimate with the IAS algorithm and posterior mean by MCMC sampling over the ROI are computed, corresponding to two different hyperprior models (gamma, inverse gamma) and to the EEG and MEG recording modalities. In these examples, the scaling and shape parameter values are held constant at $\theta_0 = 10^{-7}$ and $\beta = 1.55$ and up to 20 iterations of the IAS methods were allowed to compute the MAP estimates. The ROI consisted of three disjoint sets containing together $5\,982$ elements, see Figure 7. Corresponding to each combination of hyperprior model and recording modality, an MCMC sample of size $M = 50\,000$ is generated, assuming that impressed currents differ from zero only inside the ROI.

In Figures 8 – 9, the MAP and posterior mean estimates of the current vector α using the EEG data are superimposed with three different projections (sagittal, coronal, and axial) of the brain. The top row corresponds to the gamma hyperprior and the bottom row to inverse gamma. Interestingly, the MAP estimate with the gamma distribution is more focal than with the inverse gamma, while with the posterior mean, the inverse gamma yields more focal estimates. The same behavior is also seen in the estimates based on the MEG data, see Figures 10 – 11.

Although focal, the estimate obtained from the EEG data is unable to locate all the sources. The preference for finding the right cortical source while missing

the left cortical source may be related to the orientation of the dipoles and the positioning of the electrodes in the simulation.

Unlike in the case of the planar geometry, when using a realistic head model the deep source in the MAP estimates is not driven to the surface, indicating that there is enough geometric information in the three dimensional positioning of the electrodes or magnetometers to resolve the depth.

The MAP and posterior mean estimates of the variance vector θ of the prior are shown in Figures 12 – 13 using the EEG data and in Figures 14 – 15 for the MEG data. The results are in full agreement with the corresponding estimates for the currents.

Finally, based on the MCMC sampling, we estimate the posterior variances $\mathrm{Var}(\alpha)$ of the current vector. The estimated variances are visualized in Figures 16 – 17. Interestingly, even when the posterior estimates are blurry, the posterior variances are relatively focal, giving additional information of the likely positions of the sources.

To monitor the performance of the MCMC sampler, in Figure 18 we have plotted the sample histories of the components of α and θ corresponding to the index of maximum value of the conditional mean estimate. In the realistic geometry, we do not encounter effects that would suggest multimodality of the distribution, and the mixing and convergence seem, by visual inspection, satisfactory.

8. Discussion

In this article, we consider the MEG/EEG inverse problems of localizing few focal sources in the framework of Bayesian hierarchical models. It is shown that by using a conditionally Gaussian prior combined with generalized gamma distributions as hyperpriors, a rich family of posterior densities ensues, and it is possible to generate numerous estimates, some of which are closely related to previously proposed, regularization based estimators.

We propose a simple and effective numerical algorithm, the Iterative Alternating Sequential (IAS) algorithm for computing the MAP estimate simultaneously for the current density and its variance. The versatility of this approach is confirmed by its ability of producing an efficient fixed point implementation of several well known focal reconstruction methods based on non-quadratic penalties or non-Gaussian prior distributions. Particular instances, that correspond to a choice of few scalar parameters in the hypermodel, include the minimum current estimate, the ℓ^p-regularized estimate and the limited support functional estimate. Furthermore, the efficient numerical implementation using iterative solvers gives also a natural interpretation in the statistical framework for the FOCUSS algorithm when applied to the MEG/EEG imaging problem. Compared to the empirical Bayes methods proposed in the literature such as evidence maximization, or Automatic Relevance Determination (ARD), Expectation Maximization (EM) and variational methods [26, 33], the IAS algorithm is very fast, easy to implement and does not require sophisticated minimization methods, nor does it lead to intractable integral expressions. As shown in this article, explicit expressions for the iterative minimization steps can be found with numerous choices of the hyperprior parameters without requiring conjugacy property of the hyperprior.

The different choices of the hypermodel parameter lead to algorithms that behave qualitatively similarly with respect to the MAP estimation: when the depth

resolution due to the measurement geometry is poor, as is the case in the local half space model, superficial focal sources are localized well, while deep sources are biased towards the surface. The MCMC analysis of the posterior distributions, however, reveals a qualitative difference between the hypermodels with different values of the hypermodel parameter r. In the case where $r = 1$, corresponding to the gamma hypermodel, the MAP estimate coincides with the minimum current estimate and, like the posterior mean estimate, is biased towards superficial sources. With $r = -1$, yielding the inverse gamma hypermodel, the MAP estimate is also biased towards superficial sources, but the posterior mean is not. This qualitative difference may be interpreted by saying that the smaller the parameter r, the less Gaussian the model, and in the limit as $r \to \infty$ we obtain the minimum norm solution. Note that for Gaussian distributions, the posterior mean and the MAP estimates coincide.

From the point of view of Bayesian modelling paradigm, the qualitatively different behavior of solutions corresponding to different hyperpriors may seem strange, since in none of the cases any a priori preference to superficial sources is given. The explanation is related to the parameter values in the hyperpriors. The scaling parameter θ_0 in the gamma distribution was chosen very small in our example with planar geometry to obtain good localization in the tangential direction. Since the mean of the gamma distribution is $\theta_0 \beta$, this choice of θ_0 favors small current dipoles in the conditional mean, and the energetically easiest way to achieve this is to place all the currents on the surface. The mean of the inverse gamma distribution, $\theta_0/(\beta - 1)$, is also a small number for this choice of θ_0, but since this distribution allows significantly larger outliers, it has no difficulty in letting few large dipoles in the lower layers explain the data. For completeness, we also performed MCMC runs with the gamma distribution with considerably larger scaling parameter values. The result suggest that when the scaling parameter is large enough to allow dipoles of the correct size in the lower layers where the true dipole lies, the focal properties of the posterior mean are lost. We conclude that the findings are in line with the Bayesian philosophy and seem to suggest that the inverse gamma prior is more suitable for deep sources.

When using a three-dimensional more realistic geometry, the question of depth resolution becomes more delicate since the geometry starts to play a significant role, and measurements from different directions possibly contribute to the depth resolution more than the prior.

The computed results obtained with the realistic head model are in good agreement with the planar case results, suggesting that the posterior mean estimate is most effective in combination with the inverse gamma prior. The MAP estimate, on the other hand, is most effective when applied in connection with the gamma prior. In general, for the inverse gamma hypermodel the posterior mean estimation produced better results with larger values of the scaling parameter θ_0 than MAP estimation, while for the gamma hypermodel, MAP estimates were superior to CM estimates with larger θ_0 values. In these model both estimate types were found to be more sensitive to the choice of the shape parameter β than in the case of the planar geometry, with large value of β leading to blurred estimates and ultimately to invisibility of the deep source current. The hyperparameter values $\theta_0 = 10^{-7}$ and $\beta = 1.55$, used in the realistic case, were chosen based on visual inspection.

Future extensions of this work include a hierarchical extension of the model where the values of the hypermodel parameters will be chosen based on the data, and the extension of the formalism to include time dependent sources with a longitudinal correlation structure.

9. Appendix: Complete electrode model and FEM for EEG/MEG

This appendix describes briefly how the electric and magnetic lead field matrices can be constructed through the complete electrode model [28] and the FEM for a realistic three dimensional head, denoted here by Ω.

The complete electrode model assumes that a set of contact electrodes $e_1, e_2, \ldots, e_L$ with contact impedances $z_1, z_2, \ldots, z_L$ is attached to the boundary $\partial\Omega$. The electrode potential values are collected into a vector $U = (U_1, U_2, \ldots, U_L)$ and the electric potential field u satisfies the equation

$$(9) \qquad \nabla \cdot (\sigma \nabla u) = \nabla \cdot \mathbf{J}, \quad \text{in } \Omega,$$

as well as the boundary conditions

$$(10) \qquad \sigma \frac{\partial u}{\partial n}\Big|_{\partial\Omega \setminus \cup_\ell e_\ell} = 0, \quad \int_{e_\ell} \sigma \frac{\partial u}{\partial n}\, ds = 0, \quad \left(u + z_\ell \sigma \frac{\partial u}{\partial n}\right)\Big|_{e_\ell} = U_\ell,$$

with $\ell = 1, 2, \ldots, L$. Additionally, the Kirchoff's voltage law $\sum_{\ell=1}^{L} U_\ell = 0$ is assumed to hold. The weak form of (9) and (10) can be formulated by requiring that $u \in H^1(\Omega) = \{ w \in L^2(\Omega) : \partial w/\partial r_i \in L^2(\Omega), i = 1, 2, 3 \}$ and $\mathbf{J} \in \mathbf{H}(\text{div}; \Omega) = \{ \mathbf{w} \in L^2(\Omega)^3 : \nabla \cdot \mathbf{w} \in L^2(\Omega) \}$. These function spaces are thoroughly discussed e.g. in [18].

The finite element discretized fields corresponding to $u \in H^1(\Omega)$ and $\mathbf{J}^p \in H(\text{div}; \Omega)$ can be defined as $u_{\mathcal{T}} = \sum_{i=1}^{N_u} \zeta_i \psi_i$ and $\mathbf{J}_{\mathcal{T}} = \sum_{i=1}^{N_J} \alpha_i \mathbf{w}_i$, respectively. Here, the functions $\psi_1, \psi_2, \ldots, \psi_{N_u} \in H^1(\Omega)$ and $\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_{N_J} \in H(\text{div}; \Omega)$ are, respectively, scalar and vector valued finite element basis functions, defined on a shape regular finite element mesh $\mathcal{T}$ [3], and the coefficients form the coordinate vectors $\zeta = (\zeta_1, \zeta_2, \ldots, \zeta_{N_u})^T$ and $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_{N_J})^T$. Furthermore, since the sum of the electrode potentials is assumed to be zero in the complete electrode model, it is required that $U = (U_1, U_2, \ldots, U_L)^T = R\tilde{\zeta}$, where $\tilde{\zeta}$ is a vector and $R \in \mathbb{R}^{L \times (L-1)}$ is a matrix with entries $R_{1,j} = -R_{j+1,j} = 1$ for $j = 1, 2, \ldots, L-1$, and otherwise $R_{i,j} = 0$.

The vectors α, ζ and $\tilde{\zeta}$ are linked through a symmetric and positive definite linear system of the form

$$(11) \qquad \begin{bmatrix} B & C \\ C^T & G \end{bmatrix} \begin{bmatrix} \zeta \\ \tilde{\zeta} \end{bmatrix} = \begin{bmatrix} F\alpha \\ 0 \end{bmatrix},$$

in which the submatrix entries are given by

$$
\begin{aligned}
F_{i,k} &= \int_\Omega (\nabla \cdot \mathbf{w}_k)\psi_i \, d\Omega, \\
B_{i,j} &= \int_\Omega \sigma \nabla \psi_i \cdot \nabla \psi_j \, d\Omega + \sum_{\ell=1}^{L} \frac{1}{z_\ell} \int_{e_\ell} \varphi_i \varphi_j \, dS, \\
C_{i,j} &= -\frac{1}{z_1} \int_{e_1} \varphi_i \, dS + \frac{1}{z_{j+1}} \int_{e_{j+1}} \varphi_i \, dS, \\
G_{i,j} &= \frac{1}{z_j} \int_{e_j} dS + \frac{\delta_{i,j}}{z_{j+1}} \int_{e_{j+1}} dS,
\end{aligned}
$$

where $\delta_{i,j}$ denotes the Kronecker delta. The system (11) arises from the Ritz-Galerkin discretization [3] of the weak form of (9) and (10). Similarly, a discretized version of the Biot-Savart law (2) can be expressed as $B = W\alpha - V\zeta$, where B is a vector containing the magnetic field values at the measurement locations, and the matrices are defined by

$$
\begin{aligned}
W_{i,3(j-1)+k} &= \int_\Omega \frac{\mathbf{e}_k \cdot \mathbf{w}_j \times (\mathbf{r}_i - \mathbf{r})}{|\mathbf{r}_i - \mathbf{r}|^3} \, d\mathbf{r}, \quad \text{and} \\
V_{i,3(j-1)+k} &= \int_\Omega \frac{\mathbf{e}_k \cdot \sigma \nabla \psi_j \times (\mathbf{r}_i - \mathbf{r})}{|\mathbf{r}_i - \mathbf{r}|^3} \, d\mathbf{r},
\end{aligned}
$$

with $\mathbf{r}_j$ denoting the jth measurement location and $\mathbf{e}_k$ denoting the kth natural basis vector.

The dependences of U and B on the vector α are described by the electric and magnetic lead field matrices M^{e} and M^{m}, respectively. These matrices are given by

$$
\begin{aligned}
M^{\mathrm{e}} &= R(C^T B^{-1} C - G)^{-1} C^T B^{-1} F, \\
M^{\mathrm{m}} &= W - V(B - CG^{-1}C^T)^{-1}F,
\end{aligned}
$$

as can be verified through straightforward linear algebra manipulations. Note that these expressions are valid only if a set of contact electrodes is attached to the head during the magnetic field measurement.

References

[1] T. Auranen, A. Nummenmaa, M. S. Hämäläinen, I. P. Jääskeläinen, J. Lampinen, A. Vehtari, and M. Sams, *Bayesian analysis of the neuromagnetic inverse problem with ℓ^p-norm priors*, NeuroImage, 26 (2005), pp. 870–884.

[2] S. Baillet and L. Garnero, *A Bayesian approach to introducing anatomo-functional priors in the EEG/MEG inverse problem*, IEEE Trans. BME, 44 (1997), pp. 374–385.

[3] D. Braess, *Finite elements*, Cambridge University Press, Cambridge, UK, 2001.

[4] D. Calvetti and E. Somersalo, *A unified Bayesian framework for algorithms to recover blocky signals*, (2007). Advanced Signal Processing Algorithms, Architectures and Implementations XVII, edited by Franklin T. Luk, Proc. of SPIE Vol. 6697 (SPIE 2007), paper no. 6697–4 (10 pages).

[5] ——, *Gaussian hypermodels and recovery of blocky objects*, Inverse Problems, 23 (2007), pp. 733–754.

[6] ——, *Introduction to Bayesian scientific computing – ten lectures on subjective computing*, Springer Verlag, New York, 2007.

[7] ——, *Hypermodels in the Bayesian imaging framework*, (2008). To appear in Inverse Problems.

[8] J. C. DE MUNCK, B. W. VAN DIJK, AND H. SPEKREJSE, *Mathematicaldipoles are adequate to describe realistic genrators of human brain activity*, IEEE Trans. BME, 35 (1988), pp. 960–966.

[9] N. G. GENÇER AND C. E. ACAR, *Sensitivity of EEG and MEG measurements to tissue conductivity*, Phys. Med. Biol., 49 (2004), pp. 701–717.

[10] I. F. GORODNITSKY AND B. D. RAO, *Sparse signal reconstruction from limited dat using FOCUSS: a re-weighted minimum norm algorithm*, IEEE Trans. Signal Proc., 45 (1997), pp. 600–616.

[11] M. S. HÄMÄLÄINEN, R. HARI, R. ILMONIEMI, J. KNUUTILA, AND O. V. LOUNASMAA, *Magne-toencephalography – Theory, instrumentation and applications to nonivasive studies of the working human brain*, Rev. Mod. Phys., 65 (1993), pp. 413–498.

[12] M. S. HÄMÄLÄINEN AND R. ILMONIEMI, *Interpreting measured magnetic fields of the brain: estimates of current distributions*, (1984). Report TKK-F-A559 Low Temperature Laboratory, Helsinki University of Technology, SF-02150 Espoo, Finland.

[13] K. JERBI, S. BAILLET, J. C. MOSHER, G. NOLTE, L. GARNERO, AND R. M. LEAHY, *Localization of realistic cortical activity in MEG using current multipoles*, NeuroImage, 22 (2004), pp. 779–793.

[14] S. C. JUN, J. S. GEORGE, J. PARÉ-BLAGOEV, S. M. PLIS, D. M. RANKEN, D. M. SCHMIDT, AND C. C. WOOD, *Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data*, NeuroImage, 28 (2005), pp. 84–98.

[15] J. KAIPIO AND E. SOMERSALO, *Statistical and computational inverse problems*, Springer Verlag, New York, 2004.

[16] F. H. LIN, T. WITZEL, S. P. AHLFORS, S. M. STUFFLEBEAM, J. W. BELLIVEAU, AND M. S. HÄMÄLÄINEN, *Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates*, NeuroImage, 31 (2006), pp. 160–171.

[17] J. S. LIU, *Monte Carlo strategies in scientific computing*, Springer Verlag, New York, 2001.

[18] P. MONK, *Finite Element Methods for Maxwell's Equations*, Clarendon Press, Oxford, UK, 2003.

[19] J. C. MOSHER, P. S. LEWIS, AND R. M. LEAHY, *Multiple dipole modelling and localization from spatiotemporal MEG data*, IEEE Trans. BME, 39 (1992), pp. 541–557.

[20] S. NAGARAJAN, O. PORTNIAGUINE, D. HWANG, C. R. JOHNSON, AND K. SEKIHARA, *Controlled Support MEG imaging*, NeuroImage, 33 (2006), pp. 878–885.

[21] A. NUMMENMAA, T. AURANEN, M. S. HÄMÄLÄINEN, I. P. JÄÄSKELÄINEN, J. LAMPINEN, M. SAMS, AND A. VEHTARI, *Hierarchical Bayesian estimates of distributed MEG sources: Theoretical aspects and comparison of variational MCMC methods*, NeuroImage, 35 (2007), pp. 669–685.

[22] R. D. PASCUAL-MARQUI, *Review of methods for solving the EEG inverse problem*, Int. J. Bioelectromagnetism, 1 (1999), pp. 75–86.

[23] P. PERONA AND J. MALIK, *Scale-space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., 12 (1990), pp. 629–639.

[24] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2003.

[25] J. SARVAS, *Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem*, Phys. Med. Biol., 32 (1987), pp. 11–22.

[26] M. SATO, T. YOSHIKA, S. KAJIHARA, K. TOYAMA, N. GODA, K. D. KAWATO, AND M. KAWATO, *Hierarchical Bayesian estimation for MEG inverse problem*, NeuroImage, 23 (2004), pp. 806–826.

[27] D. M. SCHMIDT, J. S. GEORGE, AND C. C. WOOD, *Bayesian inference applied to the electromagnetic inverse problem*, Human Brain Mapping, 7 (1999), pp. 195–212.

[28] E. SOMERSALO, M. CHENEY, AND D. ISAACSON, *Existence and uniqueness for electrode models for electric current computed tomography*, SIAM J. Appl. Math., 52 (1992), pp. 1023–1040.

[29] O. TANZER, S. JÄRVENPÄÄ, J. NENONEN, AND E. SOMERSALO, *Representation of bioelectric current sources using Whitney elements in finite element method*, Phys. Med. Biol., 50 (2005), pp. 3023–3039.

[30] N. J. TRUJILLO-BARRETO, E. AUBERT VÁSQUEZ, AND P. A. VALDÉS-SOSA, *Bayesian model averaging in EEG/MEG imaging*, NeuroImage, 21 (2004), pp. 1300–1319.

[31] K. UUTELA, M. S. HÄMÄLÄINEN, AND E. SOMERSALO, *Visualization of magnetoencephalographic data using minimum current estmates*, NeuroImage, 10 (1999), pp. 173–180.

[32] P. Wen, F. He, and K. Sammut, *A pseudo-conductivity inhomogeneous head model for computation of EEG*, Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 20 (1998).

[33] D. Wipf, R. Ramirez, J. Palmer, S. Makeig, and B. Rao, *Analysis of empirical bayesian methods for neuroelectromagnetic source localization*, in Advances in Neural Information Processing Systems 19, B. Sch ed., 2007.

[34] C. H. Wolters, A. Anwandter, X. Tricoche, D. Weinstein, M. A. Koch, and R. S. McLeod, *Influence of tissue conductivity anisotropy on EEG/MEG field and return current computation is a realistic head model: A simulation and visualization study using high-resolution finite element modeling*, NeuroImage, 30 (2006), pp. 813–826.

D. Calvetti: Department of Mathematics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, USA
  E-mail address: daniela.calvetti@case.edu

H. Hakula: Institute of Mathematics, Box 1100, FI-02015 Helsinki University of Technology, Finland
  E-mail address: harri.hakula@tkk.fi

S. Pursiainen: Institute of Mathematics, Box 1100, FI-02015 Helsinki University of Technology, Finland
  E-mail address: sampsa.pursiainen@tkk.fi

E. Somersalo: Department of Mathematics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, USA
  E-mail address: erkki.somersalo@case.edu