



MEG source localization under multiple constraints: An extended Bayesian framework

J r mie Mattout,^{a,*} Christophe Phillips,^b William D. Penny,^a
Michael D. Rugg,^c and Karl J. Friston^a

^aWellcome Department of Imaging Neuroscience, 12 Queen Square, WC1N 3BG London, UK

^bCentre de Recherches du Cyclotron, Li ge, Belgium

^cCenter for the Neurobiology of Learning and Memory, Irvine, CA 92697-3800, USA

Received 7 April 2005; revised 19 October 2005; accepted 31 October 2005

To use Electroencephalography (EEG) and Magnetoencephalography (MEG) as functional brain 3D imaging techniques, identifiable distributed source models are required. The reconstruction of EEG/MEG sources rests on inverting these models and is ill-posed because the solution does not depend continuously on the data and there is no unique solution in the absence of prior information or constraints. We have described a general framework that can account for several priors in a common inverse solution. An empirical Bayesian framework based on hierarchical linear models was proposed for the analysis of functional neuroimaging data [Friston, K., Penny, W., Phillips, C., Kiebel, S., Hinton, G., Ashburner, J., 2002. Classical and Bayesian inference in neuroimaging: theory. *NeuroImage* 16, 465–483] and was evaluated recently in the context of EEG [Phillips, C., Mattout, J., Rugg, M.D., Maquet, P., Friston, K., 2005. An empirical Bayesian solution to the source reconstruction problem in EEG. *NeuroImage* 24, 997–1011]. The approach consists of estimating the expected source distribution and its conditional variance that is constrained by an empirically determined mixture of prior variance components. Estimation uses Expectation-Maximization (EM) to give the Restricted Maximum Likelihood (ReML) estimate of the variance components (in terms of hyperparameters) and the Maximum A Posteriori (MAP) estimate of the source parameters. In this paper, we extend the framework to compare different combinations of priors, using a second level of inference based on Bayesian model selection. Using Monte-Carlo simulations, ReML is first compared to a classic Weighted Minimum Norm (WMN) solution under a single constraint. Then, the ReML estimates are evaluated using various combinations of priors. Both standard criterion and ROC-based measures were used to assess localization and detection performance. The empirical Bayes approach proved useful as: (1) ReML was significantly better than WMN for single priors; (2) valid location priors improved ReML source localization; (3) invalid location priors did not significantly impair performance. Finally, we show how model selection, using the log-evidence, can be used to select the best combination of priors. This

enables a global strategy for multiple prior-based regularization of the MEG/EEG source reconstruction.

  2005 Elsevier Inc. All rights reserved.

Keywords: MEG; Inverse problem; Multiple constraints; Empirical Bayes; Restricted maximum likelihood (ReML); Model selection; ROC; Multivariate Source Prelocalization (MSP)

Introduction

Magnetoencephalography (MEG) and Electroencephalography (EEG) both provide a non-invasive and instantaneous measure of whole brain activity. These measures reflect synchronous post-synaptic potentials of cortical populations of neurons (Nunez and Silberstein, 2000). Unfortunately, localizing those electromagnetic sources is an ill-posed inverse problem that, in the absence of constraints, does not admit a unique solution (Baillet et al., 2001). Consequently, deriving a realistic and unique solution rests on prior knowledge, in addition to the observed measurements.

Any source reconstruction approach is characterized by three components. The first relates to the definition of the solution space and a parametric representation of the sources. The second embodies the information about the physical and geometrical properties of the head. The latter is needed for modeling the propagation of the sources electromagnetic field through various tissues. Together, these two components constitute a generative or forward model of the MEG/EEG data that can also be used for data simulation (see Synthetic MEG data). Finally, given a forward model, the third component is an inverse operator which, according to some criterion, defines a unique source distribution. For instance, when based on a probabilistic approach (Baillet and Garnero, 1997; Schmidt et al., 1999; Phillips et al., 2002; Amblard et al., 2004), the unique inverse solution corresponds to

* Corresponding author. Fax: +44 207 807 1420.

E-mail address: jmmattout@fil.ion.ucl.ac.uk (J. Mattout).

Available online on ScienceDirect (www.sciencedirect.com).

the most likely solution according to a predefined criterion formulated in terms of the source probability distribution.

Two types of inverse method can be distinguished by their respective source models: the equivalent current dipole (ECD) and the distributed model (DM). Although other source models have been used, such as multipoles (Jerbi et al., 2004) or continuous current densities (Riera et al., 1998), both approaches usually rely upon a dipolar representation of cortical sources, which are parameterized in terms of location, orientation and intensity. An ECD models the activity of a large cortical area. MEG or EEG data are then explained by few ECDs (usually less than five). Distributed models consider a large number (typically $\sim 10,000$) of dipoles distributed at fixed locations over the cortical surface. Although the underlying parametric models are the same, the parametrization of the solution space is very different, calling for different forward calculations as well as different inverse operators and solutions.

ECDs are fitted using iterative algorithms that estimate the source parameters in order to explain the data as accurately as possible. In the iterative process, the source parameters are modified to minimize the residual error (Scherg and von Cramon, 1986; Koles, 1998). The solution is very sensitive to the number of sources and initial parameters (dipole locations and orientations), which need to be specified a priori. Indeed, ECD models require non-linear optimization with the possibility of local minima. Moreover, determining the optimal number of ECD (model complexity) is a non-trivial issue (Waldorp et al., 2005); some simulation studies have shown that, even with the right number of sources, ECD approaches are less reliable than distributed ones, when dealing with more than one source (Yao and Dewald, 2005). Finally, unlike distributed methods, ECD models do not address the anatomical deployment of an active region.

In contradistinction to ECD approaches, a DM uses the subject's anatomy derived from high resolution anatomical Magnetic Resonance Images (MRI) (Dale and Sereno, 1993). The solution space and associated forward models can then be made as realistic as allowed by computational constraints and the precision of head tissue conductivity measures. Moreover, due to the use of fixed dipoles, the forward solution only needs to be computed once, prior to any inverse operation. The DM represents a highly under-determined but linear system (see Notation). This (general) linear model, although under-determined, is formally similar to those encountered in signal and image processing and can be treated in a Bayesian way, using priors to furnish a unique solution.

In this paper, we focus on distributed source models and explore the usefulness of Bayesian model selection for determining the best combination of constraints on the inverse solution. To establish the face validity of the ensuing model selection, we also evaluated performance using conventional criteria based on detection and localization error. To assess localization error, we used simulations with quite focal sources. It is possible that ECD models would have been better than the distributed models for these focal responses. In principle, one could use Bayesian model selection to disambiguate between distributed source and ECD models for the same data. Furthermore, the application of Bayesian model selection to ECD models provides a principled way of finding the optimum number of ECDs. We are currently exploring this in the context of dynamic causal models for ERPs. In this paper, we introduce

model selection and illustrate it in the context of selecting constraints (as opposed to sources).

In the context of distributed approaches, priors based on mathematical, anatomical, physiological and functional heuristics have been considered (Hämäläinen and Ilmoniemi, 1994; Pascual-Marqui et al., 1994; Gorodnitsky et al., 1995; Baillet and Garnero, 1997; Dale et al., 2000; Phillips et al., 2002; Mattout et al., 2003; Babiloni et al., 2004). Although these approaches involve different constraints and inverse criteria, they all obtain a unique solution by optimizing a goodness of fit term and a prior term in a carefully balanced way. Most can be framed in terms of a Weighted Minimum Norm criterion (WMN), which represents the classical and most popular distributed approach (see Classical regularization: single prior) (Hauk, 2004).

However, a critical outstanding issue lies in the relative weighting of the accuracy and regularization criteria upon which the solution depends. Usually, in the context of Tikhonov regularization or WMN solutions, this weighting is fixed arbitrarily, or by using the L-curve heuristic (see Classical regularization: single prior). The latter case, which we will refer to as the (classical) WMN, is limited because it can only accommodate a single constraint on the source parameters. This means that multiple constraints (e.g., spatial and temporal; Baillet and Garnero, 1997) have to be mixed into a single prior term, using ad hoc criteria.

In this paper, we generalize the WMN approach, using a hierarchical (general) linear model that embraces, under the assumption of Gaussian errors, multiple constraints specified in terms of variance components (see Empirical Bayes: multiple priors). These priors can be formulated in sensor or source space. The optimal weight associated with each constraint is estimated from the data following an empirical Bayesian approach and is computed iteratively using Expectation Maximization (EM) (Friston et al., 2002). These weights are equivalent to restricted maximum likelihood (ReML) estimates of the prior covariance components.

In a companion paper (Phillips et al., 2005), we addressed the face validity of empirical Bayes in this context. In the present paper, the proposed framework is applied to simulated event related field (ERF) data with realistic noise. Our investigation focused on the comparison between the ReML approach and the classical WMN and on the comparison between single vs. multiple priors when solving the MEG inverse problem. The main contribution of this paper is the introduction of a second level of inference using Bayesian model selection. Because each model is defined by its prior covariance component, we can compare different combinations of priors in a principled way. We will illustrate this by showing that model selection can identify invalid priors and point to the optimum number of valid or useful priors.

The paper is organized as follows. In the Method section, we review the classical weighted minimum norm approach and present the ReML scheme that enables a principled and unique incorporation of multiple priors. In the Application section, we describe the simulations we have used to compare quantitatively the classical and ReML inverse approaches. The priors have been chosen to emphasize the role of ReML in the context of multimodal integration. In addition to the conventional localization error criterion, two complementary evaluation procedures are introduced. The first refers to the notion of detection power and is based upon Receiver Operating Characteristic (ROC) curve analysis. The second is based on Bayesian model selection and the evidence for different models with different prior covariance

components. The results are presented in the final section and commented in the Discussion.

Method

Notation

Consider a t -sample-long window of MEG measurements acquired on n sensors. A distributed source model, involving p dipoles with fixed position and orientation (Dale and Sereno, 1993), can be expressed as the linear matrix equation

$$\mathbf{M} = \mathbf{K}\mathbf{J} + \mathbf{E}, \quad (1)$$

where \mathbf{M} is the $n \times t$ data matrix, \mathbf{K} is the $n \times p$ forward operator defining the propagation of the magnetic field in head tissues and \mathbf{J} is the $p \times t$ matrix of dipole magnitudes to be estimated. Data are corrupted by additive measurement noise \mathbf{E} .

The columns of \mathbf{K} are called the “forward fields” and describe the measurements observed across sensors, induced by a particular dipole. The rows of \mathbf{K} are the “lead fields” and describe the flow of current for a given sensor through each dipole location (Ermer et al., 2001). \mathbf{K} is obtained by solving the forward problem for each dipole location and orientation of the given DM. It relies only upon the geometry and conductivity of head tissues (Mosher et al., 1999).

Classical regularization: single prior

A classical approach to the inverse problem is the weighted minimum norm (WMN) solution (Tikhonov and Arsenin, 1977). Simple and convenient, it has become a standard solution for the MEG inverse problem as stated by Eq. (1) (Hämäläinen and Ilmoniemi, 1994; Dale and Halgren, 2001; Hauk, 2004). It entails minimizing the quadratic energy function

$$\mathcal{U}(\mathbf{J}) = \|\mathbf{M} - \mathbf{K}\mathbf{J}\|_{\mathbf{W}_n}^2 + \lambda \|\mathbf{J}\|_{\mathbf{W}_p}^2, \quad (2)$$

where $\|\cdot\|_{\mathbf{W}}$ indicates the L_2 norm associated with metric \mathbf{W} and λ is a hyperparameter which tunes the relative importance of the two terms to be minimized, the accuracy and prior term, respectively.

The unique minimum of Eq. (2) can be expressed in two ways using the matrix inversion lemma

$$\begin{aligned} \mathbf{J}_{\text{wmn}} &= \left[\mathbf{K}^T (\mathbf{W}_n^T \mathbf{W}_n) \mathbf{K} + \lambda (\mathbf{W}_p^T \mathbf{W}_p) \right]^{-1} \mathbf{K}^T (\mathbf{W}_n^T \mathbf{W}_n) \mathbf{M} \\ &= (\mathbf{W}_p^T \mathbf{W}_p)^{-1} \mathbf{K}^T \left[\mathbf{K} (\mathbf{W}_p^T \mathbf{W}_p)^{-1} \mathbf{K}^T + \lambda (\mathbf{W}_n^T \mathbf{W}_n)^{-1} \right]^{-1} \mathbf{M}. \end{aligned} \quad (3)$$

Under Gaussian assumptions, this relates directly to the Bayesian estimate of the source posterior density whose maximum a posteriori (MAP) or conditional expectation is given by

$$\begin{aligned} \mathbf{E}[\mathbf{J}|\mathbf{M}] &= \left[\mathbf{K}^T \mathbf{C}_n^{-1} \mathbf{K} + \mathbf{C}_p^{-1} \right]^{-1} \mathbf{K}^T \mathbf{C}_n^{-1} \mathbf{M} \\ &= \mathbf{C}_p \mathbf{K}^T [\mathbf{K} \mathbf{C}_p \mathbf{K}^T + \mathbf{C}_n]^{-1} \mathbf{M}, \end{aligned} \quad (4)$$

with $\mathbf{E} \sim \mathcal{N}(0, \mathbf{C}_n)$ and $\mathbf{J} \sim \mathcal{N}(0, \mathbf{C}_p)$.

This formulation shows how the constraints can be expressed in terms of prior source covariance matrices such that $\mathbf{C}_p = (\lambda \mathbf{W}_p^T \mathbf{W}_p)^{-1}$. The lower the prior variance at some dipole location, the more the dipole amplitude will be “shrunk” to zero.

Several different constraints have been considered within the WMN framework. The most simple, so-called (unweighted) minimum norm, consists of setting $\mathbf{C}_p = \mathbf{I}_p$ where \mathbf{I}_p is the $p \times p$ identity matrix (Hämäläinen and Ilmoniemi, 1994). When λ tends to zero, one obtains the maximum likelihood (ML) solution. The well-known LORETA approach entails setting \mathbf{W}_p equal to a spatial Laplacian operator (Pascual-Marqui et al., 1994). This is also known as the maximum smoothness solution. Finally, external functional constraints such as derived from fMRI have also been considered (Liu et al., 1998; Phillips et al., 2002).

The WMN solution depends upon the hyperparameter λ , which is usually evaluated using the L-curve approach (Hansen, 1992; Gorodnitsky et al., 1995; Babiloni et al., 2001) (cf. Appendix A). A major drawback of this heuristic is that it cannot readily be extended to the estimation of multiple hyperparameters. Therefore, only one single constraint can be considered. This is why the above priors have never been evaluated in conjunction.

Empirical Bayes: multiple priors

The ReML scheme relies upon rewriting Eq. (1) as a 2-level hierarchical model

$$\begin{aligned} \mathbf{M} &= \mathbf{K}\mathbf{J} + \mathbf{E}_1 \\ \mathbf{J} &= \mathbf{0} + \mathbf{E}_2, \end{aligned} \quad (5)$$

where $\mathbf{E}_1 \sim \mathcal{N}(0, \mathbf{C}_n)$ and $\mathbf{E}_2 \sim \mathcal{N}(0, \mathbf{C}_p)$.

In this framework, errors at both levels are Gaussian variables with zero mean but unknown variance. The matrices \mathbf{C}_n and \mathbf{C}_p are modeled as linear combination of variance components such that

$$\begin{aligned} \mathbf{C}_n &= \mu_1 \mathbf{Q}_{n,1} + \mu_2 \mathbf{Q}_{n,2} + \dots \\ \mathbf{C}_p &= \lambda_1 \mathbf{Q}_{p,1} + \lambda_2 \mathbf{Q}_{p,2} + \dots \end{aligned} \quad (6)$$

where $\mu = [\mu_1, \mu_2, \dots]$ and $\lambda = [\lambda_1, \lambda_2, \dots]$ are the unknown hyperparameters that balance the various variance components introduced, either at the first (sensor) or second (source) level. At the second level, \mathbf{C}_p acts as a prior covariance (a shrinkage prior).

In the context of parametric empirical Bayes (PEB), the hyperparameters are estimated from the data, and multiple priors can be incorporated. After \mathbf{C}_n and \mathbf{C}_p have been estimated, the inverse solution \mathbf{J} is given by the corresponding MAP estimator (cf. Eq. (4)).

In practice, hyperparameter estimation is performed using Expectation Maximization (EM) which iteratively calculates the parameters (E-step) and hyperparameters of the model (M-step) until convergence. Because the system is linear, the E-step can be absorbed in the M-step. The M-step returns the Restricted Maximum Likelihood (ReML) solution, which differs from the standard ML as it accounts for the loss of degrees of freedom due to conditional uncertainty about the parameters (Friston et al., 2002). The objective function associated with the ReML solution is given in Appendix B.

Heuristically, the closer the prior variance component \mathbf{Q}_k and the true source spatial covariance, the higher the corresponding hyperparameter λ_k . Consequently, the hyperparameters quantify the relative importance of each constraint.

Bayesian model selection and comparison

The quality of each prior model can be assessed in terms of its evidence. Each model is defined by \mathbf{P}_k , representing a set of constraints ($Q_1, \dots, Q_{m_k}; Q_1, \dots, Q_{q_k}$) that comprises m_k (resp. q_k) prior variance components in sensor (resp. source) space. From Bayes rule, the parameter's posterior density is the likelihood times the prior divided by the evidence

$$p(\mathbf{J}/\mathbf{M}, \mathbf{P}_k) = \frac{p(\mathbf{M}/\mathbf{J}, \mathbf{P}_k)p(\mathbf{J}/\mathbf{P}_k)}{p(\mathbf{M}/\mathbf{P}_k)}. \quad (7)$$

Given a model \mathbf{P}_k , our first level of inference is based on the posterior, which is maximized using ReML, while treating the evidence as a normalization constant.

We introduce here a second level of inference which identifies the most plausible model given the data. This is achieved by computing the evidence given model \mathbf{P}_k

$$p(\mathbf{M}/\mathbf{P}_k) = \int p(\mathbf{M}/\mathbf{J}, \mathbf{P}_k)p(\mathbf{J}/\mathbf{P}_k)d\mathbf{J} = \int p(\mathbf{J}, \mathbf{M}/\mathbf{P}_k)d\mathbf{J}. \quad (8)$$

Since \mathbf{P}_k is defined fully by the set of prior variance components, i.e., the set of hyperparameters (λ, μ), the log-evidence corresponds to the final estimate of the ReML objective function (cf. Appendix B) which is maximized in the first level. The best model has the maximum log-evidence. We used the log-evidence to assess the four priors and their various combinations.

Furthermore, Bayes factors \mathcal{B}_{12} defined as the ratio of the evidences associated with model \mathbf{P}_1 and \mathbf{P}_2 enable us to compare two prior models when applied to the same data (Kass and Raftery, 1995)

$$\mathcal{B}_{12} = \frac{p(\mathbf{M}/\mathbf{P}_1)}{p(\mathbf{M}/\mathbf{P}_2)}. \quad (9)$$

When $\mathcal{B}_{12} > 1$, the data favor model \mathbf{P}_1 over model \mathbf{P}_2 , and when $\mathcal{B}_{12} < 1$, the data favor model \mathbf{P}_2 . To make decisions based on Bayes factors, some cut-off value is required. Kass and Raftery (1995) proposed an interpretation of the Bayes factor based on the use of P values in classical statistics. This has been recently applied to neuroimaging data and we adopt here the same interpretation, which is summarized in Table 1 in Penny et al. (2004). It specifies that, given candidate hypotheses encoded by \mathbf{P}_1 and \mathbf{P}_2 , $\mathcal{B}_{12} \geq 20$ ($\mathcal{B}_{12} \geq 150$) corresponds to a belief of at least 95% (resp. 99%) in the statement ' \mathbf{P}_1 is true'. This corresponds to 'strong evidence' (resp. 'very strong evidence') in favor of \mathbf{P}_1 , relative to \mathbf{P}_2 .

Note that unlike the LE and ROC-based criteria (see Evaluation procedure), the log-evidence does not refer to any true or reference source distribution. It is a relative measure of the data and how well they are explained by a model. This makes log-evidence and associated Bayes factors especially useful in an empirical setting.

Application

Synthetic MEG data

Numerical phantom

Since MEG/EEG sources are believed to be restricted to pyramidal cells (Nunez and Silberstein, 2000), a common approach, within the distributed model framework, is to constrain the dipoles to the cortical surface extracted from a structural Magnetic Resonance Image (MRI) (Dale and Sereno, 1993).

Following segmentation of the MRI volume, dipoles are typically located at each node of a triangular mesh covering the white/grey matter interface (Mangin, 1995). Furthermore, since the apical dendrites of these cortical neurons are organized perpendicularly to the surface, the corresponding dipoles are often constrained to this normal orientation.

To simulate MEG data, a 3D high resolution (voxel size: 0.9375 mm \times 0.9375 mm \times 1.5 mm) MRI volume from a healthy volunteer was segmented. The boundary between white and grey matter was approximated with small triangles whose vertices provided 7081 dipole locations spread uniformly over the cortex. The spatial resolution of this numerical phantom was a sufficient representation of cortical topology, since the mean distance between two neighboring dipoles was about 3 mm. We calculated the forward operator \mathbf{K} , for this dipole mesh, using a single-shell spherical head model (Sarvas, 1987).

Simulations

MEG data were simulated over 130 sensors spread uniformly over the head, by activating two extended sources (cf. Fig. 1). Each source was a cluster comprising one randomly chosen dipole and its four nearest neighbors. The extent of each simulated source was about 5 mm in radius. The activation was modeled with a half-period sine function (over 15 time bins). A delay of two time bins was applied to waveforms of the two sources. After projection onto sensor space, white Gaussian noise was added (SNR = 20 dB, cf. Fig. 1).¹

Five hundred different source configurations were simulated to:

- compare the ReML approach with the classical WMN estimation,
- study the performance of the ReML scheme when considering various combinations of priors.

Regularization constraints

At the sensor level, for the sake of simplicity and to focus on constraints in source space, we only considered a single measurement noise component defined by $\mathbf{C}_e = \mathbf{I}_n$, i.e., independent measurement noise on each sensor with identical variance. At the source level, three types of constraint were considered, either individually or together. These constraints provide complementary information about the source locations. They were chosen to illustrate the possible types of constraint one might employ for regularizing MEG source reconstruction. Moreover, such priors can be expressed easily in terms of covariance components.

Smoothness constraint

The prior covariance component Q_{sc} (see Eq. (6)) associated with this constraint is defined by

$$Q_{sc}(i,j) = \exp\left(\frac{-d_{ij}^2}{2s^2}\right), \quad (10)$$

where d_{ij} is the euclidian distance between dipoles i and j . The higher d_{ij} , the lower the correlation between activity at

¹ SNR stands for Signal to Noise Ratio and is here expressed in decibels, i.e. $\text{SNR} = 20 \log_{10}(A_s/A_n)$, where A_s (resp. A_n) refers to the maximum absolute signal (resp. noise) value. A SNR of 20dB thus corresponds to a 10% noise level.

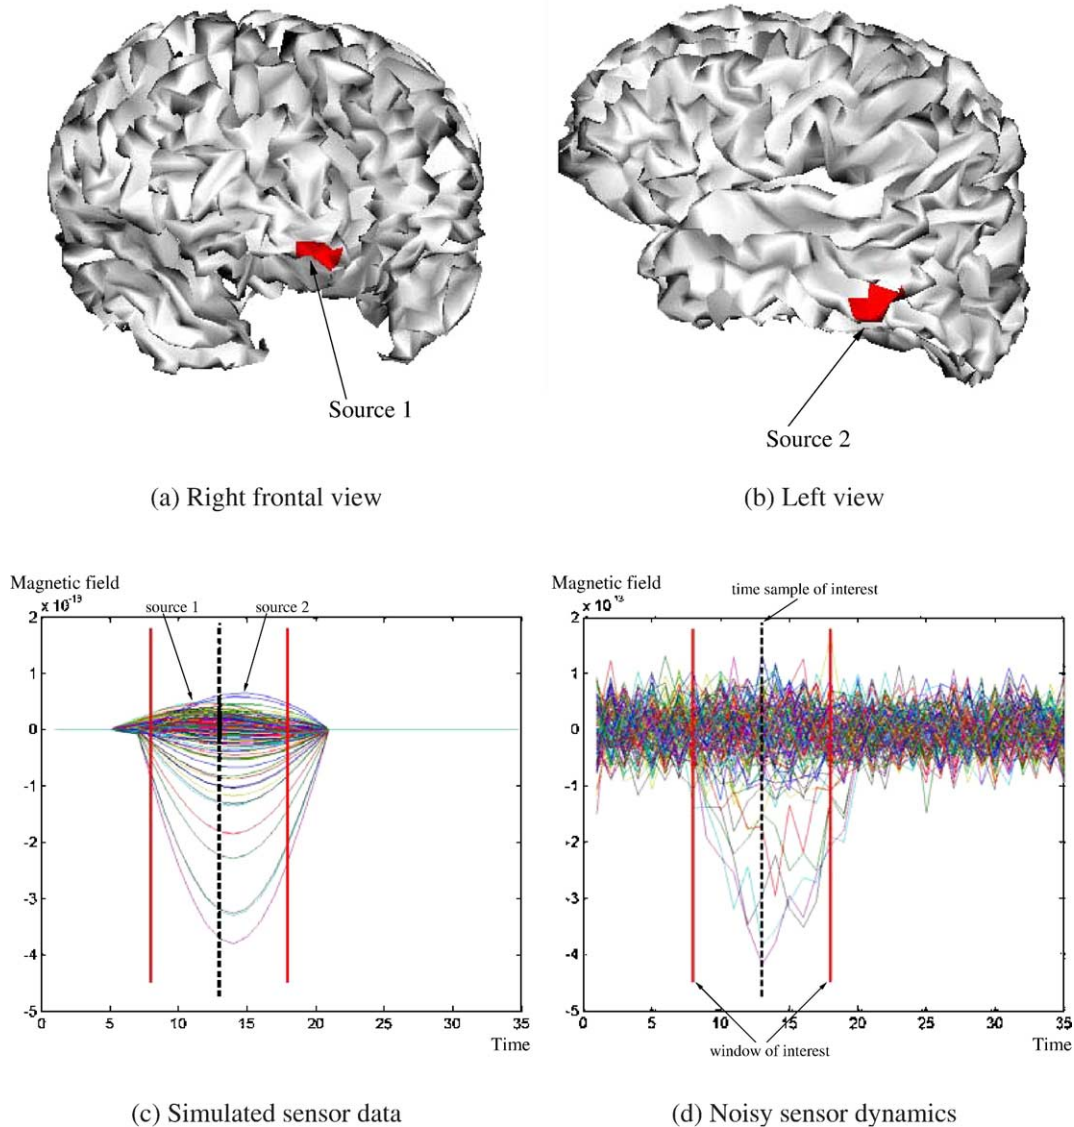


Fig. 1. Example of a simulated MEG data set: the two extended underlying sources (a, b) and their corresponding dynamics in sensor space, before (c) and after adding white Gaussian noise to render SNR = 20 dB (d).

the dipoles i and j .² The spatial smoothness parameter s was set to 8 mm. It specifies the standard deviation of the implicit smoothness. Like LORETA (Pascual-Marqui et al., 1994), this prior enforces correlation among neighboring sources.

Intrinsic functional constraint

It has been shown recently that an efficient regularization prior can be derived from the MEG data itself (Mattout et al., 2005). This is known as Multivariate Source Prelocalization (MSP) which provides, at each source location, a coefficient of activation. This coefficient refers to the spatial support of activity

but does not encode any estimated intensity. The MSP process focuses on where sources are expressed, without considering their amplitude.

MSP entails a multivariate comparison of the observed normalized magnetic field scalp topology, $\bar{\mathbf{M}}$, with all possible linear combinations of the normalized forward fields, $\bar{\mathbf{K}}$. This process returns a coefficient that quantifies the correspondence of a normalized forward field (the putative contribution of a given dipole) with the normalized observations.

The larger the prelocalization coefficient α_i associated with dipole i , the higher the affinity between forward field \mathbf{k}_i and the filtered MEG data. These coefficients can be introduced as quantitative priors into a regularization scheme as shown here. They can also enable a substantial reduction of the inverse solution space, by only considering dipoles that are most likely to be active (high α_i). Regularizing and restricting the solution space in this way have been shown to improve localization (Mattout et al., 2005).

² Note that this smoothness operator might be improved by using the geodesic distance between nodes rather than the Euclidian distance.




Fig. 2. Illustration of the two functional external priors considered (hatched regions) pertaining to the location of the two simulated sources (a): right location (b) and wrong location (c) priors.

For each source configuration, we restricted the solution space to the 1500 dipoles with the highest coefficients.³ Within this subset, we defined the intrinsic prior covariance component Q_{int} as

$$\begin{cases} Q_{\text{int}}(i,i) = \alpha_i, \\ Q_{\text{int}}(i,j) = 0 \quad \text{when } i \neq j. \end{cases} \quad (11)$$

Extrinsic functional constraints

The final sort of constraint we considered is either based on previous experience or on data from other imaging modalities, typically fMRI. This constraint was simply designed as a binary mask, distinguishing qualitatively between a priori active and non-active cortical areas. The corresponding prior source variance component is defined by the diagonal matrix

$$\begin{cases} Q_{\text{ext}}(i,i) = 1 & \text{in active regions,} \\ = 0 & \text{elsewhere.} \end{cases} \quad (12)$$

We modeled two sorts of extrinsic priors; valid ($Q_{\text{ext}} = Q_{\text{ext}}^v$) and invalid ($Q_{\text{ext}} = Q_{\text{ext}}^i$) (cf. Fig. 2). We were interested in the impact of invalid priors, particularly in the context of multiple priors where there is an opportunity to discount invalid priors in favor of valid priors.

Evaluation procedure

Since the classical WMN approach (i.e., the L-curve approach) can only accommodate one constraint, single and multiple prior cases are considered separately. Single prior simulations enabled us to compare the proposed Bayesian estimation to classical WMN, while the second set of analyses allowed us to study the behavior of the ReML approach under multiple constraints.

Using ReML, each of the 500 simulated data sets were analyzed under all combinations of one, two, three or four priors.⁴ The evaluation of those reconstructions relied upon the two comple-

mentary measures described next. The model comparison metric introduced here, based on the model evidence, constitutes the third and only data-driven evaluation criterion.

Localization error

Each estimated source distribution was first compared to the true distribution in terms of Localization Error (LE). The LE provides a measure of the localization accuracy. It corresponds to the distance between the true source and the estimated source with the greatest absolute amplitude. To obtain a single value per inverse solution, the LE was set to the maximum of the LE associated with the two simulated sources (i.e., an upper bound).

ROC analysis

ROC techniques enabled us to characterize the inverse methods in terms of correctly classifying each dipole, as either active or not. This complements performance measures based solely on localization.

A ROC curve can be constructed for each estimated source distribution, which represents the true positive rate (*sensitivity*) as a function of the false positive rate ($1 - \textit{specificity}$). To assess the false positive rate, we chose 10 random dipoles among the non-active ones.

The area under the curve (AUC) quantifies the detection power of the method (Metz, 1998). The AUC ranges between 0 and 1 and indicates the probability of correct separation of an active source from a non active one.

Comparing the AUC of different inverse models enables one to assess the relative performance in terms of detection power. We analyzed the AUC formally using Analysis of Variance (ANOVA).

Results

Empirical Bayes vs. classical estimation

Fig. 3 shows the distribution of the localization error (LE) derived from the simulations, for the WMN (i.e., L-curve) and ReML estimations, when considering each of the four priors independently. The valid location prior (Q_{ext}^v) gave the best results, whatever the inverse method, with an exact localization of the two active sources in more than 90% of the simulations.⁵ Conversely,

³ Although a statistical approach has been proposed for this dimension reduction (Mattout et al., 2005), we kept the size of the solution space constant (1500 dipoles) so that each simulation would entail the same number of parameters. Critically, the span of the solution space remained identical whatever the set of priors. By modeling the prior covariance as a mixture of components, we can use improper (i.e., non-invertible components). For example, a prior component of zero does not imply that the mixture is zero. This would not be the case if we had modeled the precision as a mixture. Here, a prior component of infinity means the mixture would also be precise.

⁴ The four priors are the four variance components Q_{sc} , Q_{int} , Q_{ext}^v and Q_{ext}^i described in Regularization constraints. This gave 19 inverse solutions per dataset: [WMN + ReML] * [4 single-priors] + [ReML] * [62-priors + 4 3-priors + 1 4-priors] = 19 source reconstructions.

⁵ Although a valid location prior was used here, the correct localization was obtained for most but not all simulations. This could be explained by the improper (i.e. non-invertible) form of the prior. It is possible that adding another prior (e.g. smoothness constraint) would yield a proper mixture of covariance components and a better inverse solution. This speaks to the importance of model selection as provided by an extended Bayesian framework (cf. Bayesian model selection and comparison).

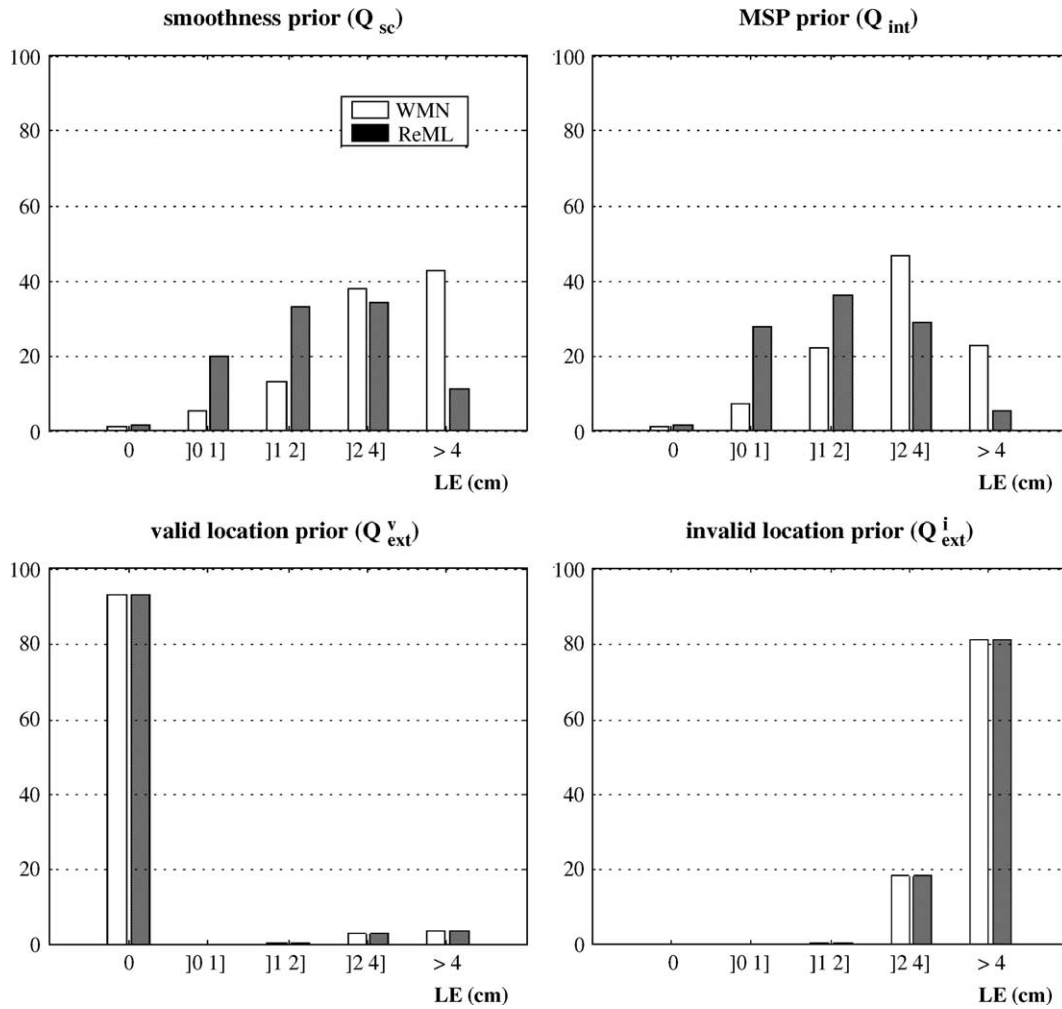


Fig. 3. Single-prior LE histogram: obtained for the WMN (L-curve) and ReML approach, considering each of the four constraints: smoothness, MSP, valid and invalid location priors.

the invalid location prior (Q_{ext}^i) gave the worse results, with a LE over 4 cm in more than 80% of the simulations.

When considering either the smoothness or intrinsic prior, the ReML localization is better than WMN. Indeed, under the smoothness prior (resp. intrinsic prior), ReML localizes the two sources with a LE of less than 2 cm in 55% (resp. 65%) of the simulations, while the WMN affords the same spatial precision in only 20% (resp. 30%) of the same simulations.

This difference between the two approaches was confirmed by the statistical comparison of the AUC for each inverse method and prior (cf. Table 1). A two by four ANOVA was performed, whose factors were the inverse method and the prior. The main effect of method (ReML vs. WMN) proved highly significant ($F(1,499) = 81,01; P < 0.001***$), implying a much better source detection profile with ReML.

Since the ReML and WMN approaches differ only in the way they estimate the hyperparameters, these results suggest that the ReML estimate of the balance between the constraint and data fit affords more precise estimates than obtained with the traditional L-curve approach.

For the toy example described in Fig. 1, we show in Fig. 4 the reconstructed sources obtained with WMN and ReML under a single smoothness constraint. In both cases, the source distribution

is rather distributed and smooth and the first source is not seen. However, the second source location is better estimated with ReML than WMN.

Table 1
Averaged values of the AUC for the WMN and ReML approaches and the four different prior models

	AUC	ReML	WMN
	Q_{sc}	0.7883	0.7770
	Q_{int}	0.7944	0.7746
	Q_{ext}^v	0.8560	0.8560
	Q_{ext}^i	0.4994	0.4994
2 constraints	Q_{sc}, Q_{int}	0.7999	
	Q_{sc}, Q_{ext}^v	0.8211	
	Q_{sc}, Q_{ext}^i	0.7931	
	Q_{int}, Q_{ext}^v	0.8211	
	Q_{int}, Q_{ext}^i	0.7962	
	Q_{ext}^v, Q_{ext}^i	0.8536	
3 constraints	$Q_{sc}, Q_{int}, Q_{ext}^v$	0.8211	
	$Q_{sc}, Q_{int}, Q_{ext}^i$	0.7972	
	$Q_{sc}, Q_{ext}^v, Q_{ext}^i$	0.8211	
	$Q_{int}, Q_{ext}^v, Q_{ext}^i$	0.8211	
4 constraints	$Q_{sc}, Q_{int}, Q_{ext}^v, Q_{ext}^i$	0.8206	

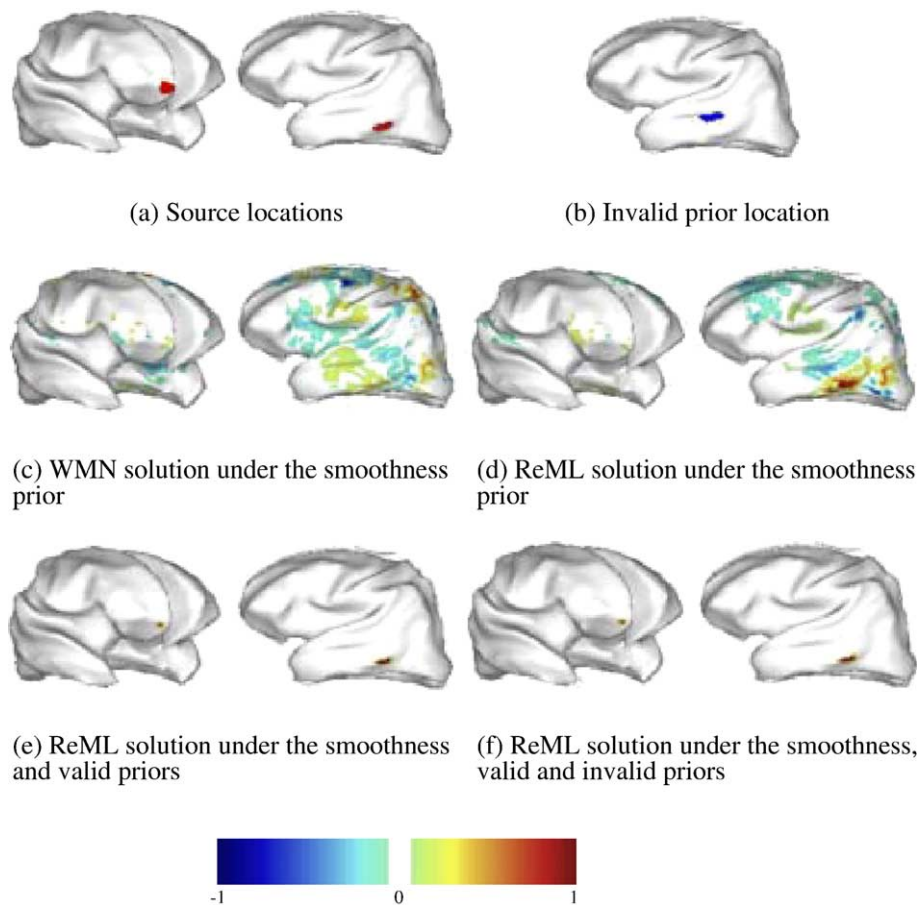


Fig. 4. Sample of simulation results (see also Fig. 1): inflated cortical representation of the two source locations (a), the invalid prior location (b) and some typical reconstructions using classical WMN and ReML (c–f). To allow a qualitative comparison of the different localizations, the source distributions have been normalized between -1 and 1 .

Regularization under multiple constraints

The results presented in this section pertain only to ReML analyses. All the different prior models were considered and the corresponding values of the AUC are provided in Table 1.

Fig. 4 illustrates the effect of valid and invalid location priors on ReML solution using the same example as in Fig. 1. It is noticeable that the valid prior, when added to the single smoothness prior, enables ReML to identify the two source locations and to greatly reduce spurious activity. Moreover, the ReML solution is not affected by the introduction of an invalid location prior. These two anecdotal observations were confirmed and quantified by the evaluation below.

Effect of the valid location prior

Fig. 5 illustrates the effect of the valid location prior on the LE of the ReML-based source reconstruction. It shows that, whatever the set of priors, if this set includes the valid location constraint, then the LE decreases dramatically. This finding generalizes the equivalent result obtained for the single prior case.

The importance of valid location priors was also observed in terms of the ROC analysis. A two by seven ANOVA was performed whose factors were the inclusion or not of the valid location constraint and the seven possible prior models. The main effect on the valid prior was highly significant ($F(1,499) = 2565.272$; $P < 0.001^{***}$).

Effect of the invalid location prior

Fig. 6 illustrates the effect of the invalid location prior on the LE of the ReML-based source reconstruction. It shows that, whatever the set, if it includes the invalid location constraint, then the LE is not compromised, showing only a slight increase. Indeed, any deterioration in detectability was insignificant, as assessed by a two by seven ANOVA on the ROC AUC ($F(1,499) = 0.140$; $P = 0.708$).

The key conclusion here is that the ReML scheme is robust to misspecified or redundant priors.

Selecting and comparing prior models

Finally, the Log-evidence and Bayes factors allowed us to select and compare the different prior models.

Bayesian model selection. For a given data set, Bayesian model selection enables one to select the best model, based on its Log-evidence. The higher the Log-evidence, the better the model. As an example, Table 2 gives the Log-evidences obtained for each prior model when applied to the particular data set described in Fig. 1. In this case, the single prior model comprising the valid location prior had the maximum Log-evidence.

More generally, Fig. 7 presents the distribution, over all simulations, of the prior models selected according to their Log-evidence. For comparison, the histogram of prior models leading

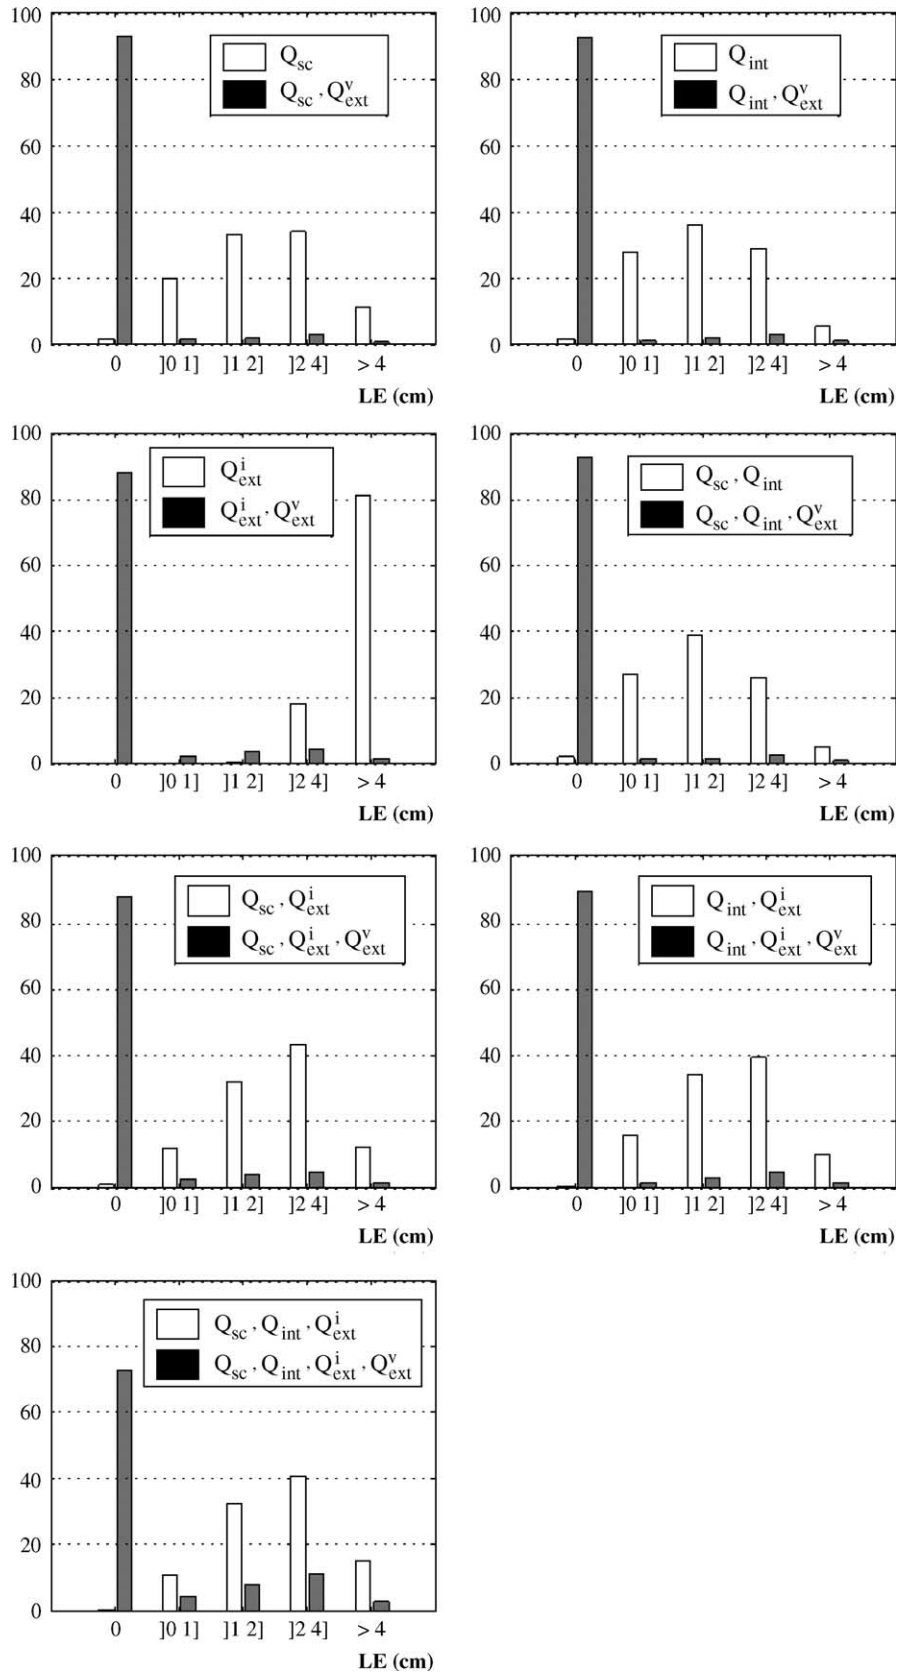


Fig. 5. Effect of valid location prior on the ReML Localization Error distribution, for each prior model, with and without the valid location prior.

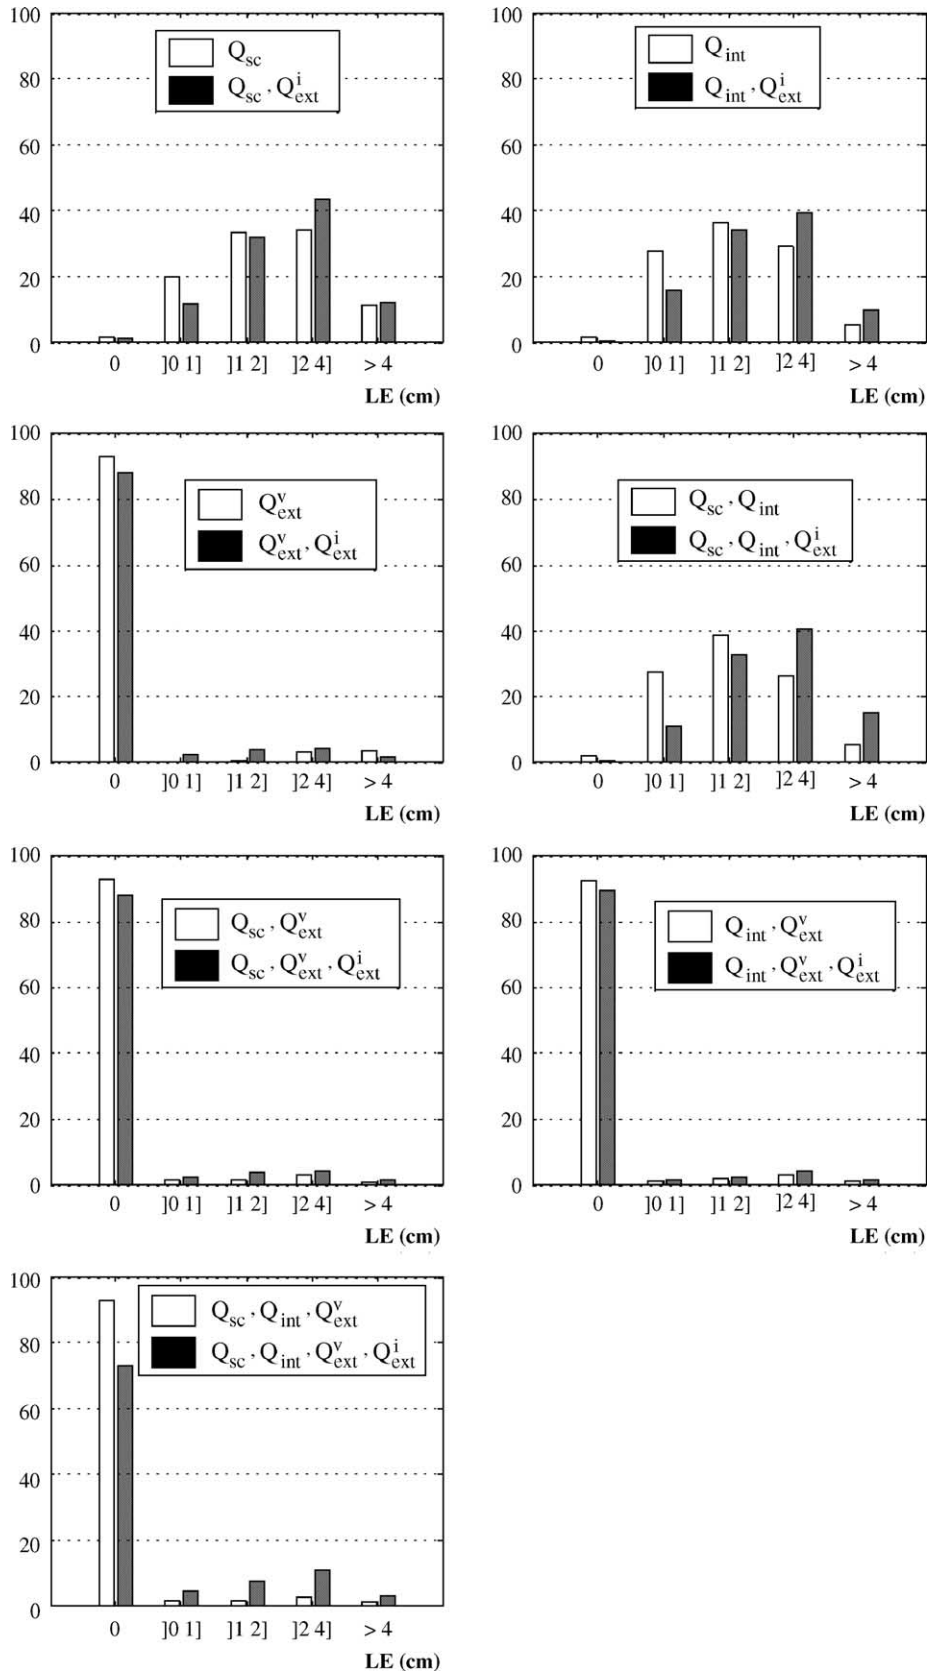


Fig. 6. Effect of invalid location prior on the ReML Localization Error distribution, for each prior model, with and without the invalid location prior.

Table 2

Log-evidence calculated for each prior model, when applying ReML to the simulation sample described in Fig. 1

Log-evidence		
1 constraint	Q_{sc}	205.2
	Q_{int}	208.4
	Q_{ext}^v	215.6
	Q_{ext}^i	131.5
2 constraints	Q_{sc}, Q_{int}	207.4
	Q_{sc}, Q_{ext}^v	214.1
	Q_{sc}, Q_{ext}^i	204.9
	Q_{int}, Q_{ext}^v	214.9
	Q_{int}, Q_{ext}^i	207.4
	Q_{ext}^v, Q_{ext}^i	213.2
3 constraints	$Q_{sc}, Q_{int}, Q_{ext}^v$	211.5
	$Q_{sc}, Q_{int}, Q_{ext}^i$	207.2
	$Q_{sc}, Q_{ext}^v, Q_{ext}^i$	214.7
	$Q_{int}, Q_{ext}^v, Q_{ext}^i$	212.7
4 constraints	$Q_{sc}, Q_{int}, Q_{ext}^v, Q_{ext}^i$	211.3

to the lowest Localization Error is also shown. In 87% of the 500 simulated data sets, the model with the largest Log-evidence was the same as the one with the lowest LE.

Note, however, that contrary to the Log-evidence, the LE criterion does not account for the goodness of fit, which explains the minor differences between the two histograms.

Bayesian model comparison. In Table 2, one can see, on one hand, that any model which contains the valid location prior has a high Log-evidence and, on the other hand, that any model which contains both valid and invalid location priors does not show a dramatic decrease in Log-evidence, compared to the same model without the invalid location prior. However, to fully assess this observation and quantitatively confirm the results in Effect of the valid location prior and Effect of the invalid location prior, one

needs to compute and interpret the corresponding Bayes factors (see Bayesian model selection and comparison). Table 3 illustrates, on the simulation toy example of Fig. 1, the effect of valid and invalid location priors as quantified using Bayes factors. It shows significantly enhanced evidence in favor of models including valid location priors. On the other hand, it shows weak evidence in favor of models that do not contain the invalid prior compared to the ones that do. As expected, this suggests that the valid location prior is the best constraint to include in the model. It further suggests that the smoothness and intrinsic location priors are valid priors that can ameliorate the misleading effect of invalid priors.

Discussion

In Phillips et al. (2005), we described and evaluated, in the context of EEG, a variance component estimation framework based upon parametric empirical Bayes (PEB) as introduced in Friston et al., (2002). The key aspect of this generalized regularization approach is the data-driven estimation of the various covariance components of the EEG/MEG source parameters. Since many constraints, for example, anatomical, functional, physiological or mathematical, can be expressed in terms of variance components, the proposed methodology affords a precise and realistic estimate of the source covariance structure. Furthermore, the contribution of each prior or component is quantified by its corresponding hyperparameter. In Phillips et al. (2005), using a simplified geometrical model and EEG synthetic data, the approach was shown to accommodate different levels of noise while accounting for various location priors on the sources. The use of accurate location priors led to negligible localization errors. When both valid and invalid location priors were introduced, the solution was not degraded by the invalid priors.




Fig. 7. Histograms, over all simulations, of the best prior models as indicated by their Log-evidence and the LE, respectively (prior models are ordered according to the rows in Tables 1 and 2).

Table 3

Three different models \mathbf{P}_1 are compared to associated models \mathbf{P}_2 (resp. \mathbf{P}_3) which only differ by also including the valid (resp. invalid) location constraint

Bayes factor		
\mathbf{P}_1	\mathcal{B}_{21}	\mathcal{B}_{31}
$(\mathbf{I}_n; Q_{sc})$	7047	0.8
$(\mathbf{I}_n; Q_{int})$	655	0.4
$(\mathbf{I}_n; Q_{sc}, Q_{int})$	60	0.8

In the present paper, we extend this Bayesian framework so that the contribution of various priors can be assessed formally, using Bayesian model selection with the log-evidence metric. We focused on MEG, using Monte-Carlo simulations based on a realistic cortical model. Our evaluation considered the ReML approach in relation to hyperparameter estimation using the L-curve heuristic. We then proceeded with a quantitative study of multiple prior-based regularization in source space.

Three complementary metrics were used in this evaluation. First, a standard criterion was used to compare the source reconstructions in terms of Localization Error (LE). Second, ROC curve analysis was introduced to assess the detection power of the different models. Meta-analysis of the area under the curves (AUC) enabled us to evaluate performance in terms of correctly classifying active and non-active sources. This type of analysis might complement evaluation procedures for EEG/MEG source reconstruction. Unlike LE, ROC techniques enable one to quantify the accuracy of the location of active sources as well as their spatial extent.

Finally, we introduced a supra ordinate level of inference. This was based on the log-evidence or data likelihood given a particular model (Trujillo-Barreto et al., 2004; Penny et al., 2004; Daunizeau et al., 2005). We used this probability-based criterion to identify the best prior model and, more generally, the best strategy for introducing priors. Unlike the LE and ROC criteria, the log-evidence can be calculated for real data. No reference to any true source distribution is needed.

The application of the three evaluation criteria led to the following conclusions.

First, using only a single prior on the sources, the ReML hyperparameter estimation led to a similar or better reconstruction than the classical WMN, whose hyperparameter was estimated by the L-curve approach. This confirms the construct validity of the proposed method. However, when considering either valid or invalid location priors, WMN and ReML gave very similar results. ReML proved slightly but significantly better only when considering smoothness and intrinsic functional priors. This can be explained by the fact that the extrinsic priors were either fully valid or fully invalid. The L-curve approach maximizes a mixture of the prediction error and departure of the conditional estimates from their prior expectations (see Appendix A). This is a poor approximation to the evidence $p(\mathbf{M}|\lambda, \mu)$ which is maximized by ReML (see Appendix B). The L-curve approach therefore provides a suboptimal solution.

Under multiple priors, within the ReML scheme, the findings of Phillips et al. (2005) were replicated in the context of our cortical source-based MEG simulations. Indeed, whatever the dimension of the prior model:

- valid location priors significantly improved the source reconstruction,

- invalid location priors did not compromise the results given at least one other informative prior.

The second finding reflects the flexible and adaptive nature of the ReML approach and suggests that ReML could be of particular interest for multimodal integration, namely the introduction of fMRI-derived priors for constraining the MEG inverse solution. A major concern about these data fusion approach is the different time scale and nature of the physiological processes measured with MEG and fMRI techniques: as noted by many authors (Nunez and Silberstein, 2000; Gonzalez Andino et al., 2001; Pflieger and Greenblatt, 2001), an adaptative method that could distinguish between valid and invalid location priors is particularly relevant for multimodal fusion.

Note that other approaches have been proposed to estimate hyperparameters. For instance, the Generalized Cross-Validation method (GCV) (Golub et al., 1979) could be extended to estimate several regularization parameters. However, contrary to the ReML approach, the GCV method does not provide any estimate of (multiple) noise components. Moreover, in the context of a single hyperparameter, better results are sometimes obtained with the classical Tikhonov/L-curve approach (Babiloni et al., 1998).

Finally, we demonstrated the effectiveness of the extended Bayesian framework for data-driven model selection and comparison. Based on model evidence, Bayesian model selection enables one to identify the best prior model for reconstructing the sources of a given data set. With our linear hierarchical models, under Gaussian assumptions, the model evidence can be computed exactly and does not require any approximations implicit in surrogate criteria; like the Akaike Information Criterion (AIC) or the Bayesian Information Criterion (BIC) (Kass and Raftery, 1995). Our results show that the log-evidence reflects faithfully the LE criterion. This suggests that, with real data, when LE and ROC curves cannot be established, the log-evidence affords a reliable criterion for model selection. Note that to perform Bayesian model selection, all prior models, corresponding to all possible combinations of available constraints, have to be estimated with ReML. It is important to emphasize that ReML computation is not time consuming,⁶ since only low dimensional matrices ($n \times n$) need to be inverted.

Model evidence ratios or Bayes factors can be used to compare two prior models. Importantly, Bayes factors can be interpreted quantitatively, as described in Bayesian model selection and comparison and illustrated here on a simulation toy-example. In that example, although the prior model defined by the smoothness (or functional intrinsic) prior yielded a higher log-evidence than when including invalid location priors, the Bayes factor did not show any significant evidence in favor of omitting invalid priors. This result endorses the above conclusions about ReML and the ability of Bayesian model comparison to evaluate the contribution of a given prior.

Note that this paper addresses Bayesian model comparison as opposed to Bayesian model averaging. The formalism developed in this paper can be extended easily to include averages over models (Trujillo-Barreto et al., 2004). We have restricted ourselves to model selection because we wanted to focus on the quantitative evaluation of different priors, using the log-

⁶ Here, the estimation of the 19 prior models associated with each simulation required less than 1 min CPU time on a standard workstation.

evidence. In principle, it would be possible to average over all combinations of priors. However, practically speaking, this might be a little redundant, given that the different priors are weighted optimally within the optimum combination. Having said this, it would be interesting to compare the best model with the Bayesian model average.

Conclusion

The proposed framework for the MEG/EEG inverse problem relies upon a linear hierarchical model, typically provided by distributed source models. It estimates the source parameters as well as their spatial covariances (hyperparameters) that can be expressed as linear combination of independent prior components. Any putative informative prior on the sources can be introduced independently and can be subject to a quantitative evaluation of its contribution. We focussed on smoothness, functional intrinsic and extrinsic priors but other components of different forms may be useful, such as constraints on source orientation (Phillips et al., 2005).

The extended framework also affords a data-driven estimation of the model hyperparameters as well as a data-driven evaluation of each constraint's contribution. Indeed, Bayesian model comparison proved useful for evaluating the relative contribution of constraints. This is of particular interest in the context of multimodal integration and speaks to a 'Leave-one-out-strategy' where different priors, typically those derived from fMRI data analysis (each prior corresponding to an activated area), could be evaluated by being successively introduced or excluded from the ReML scheme according to the log-evidence. Further studies will evaluate this procedure on real multimodal data. This approach may be useful for characterizing empirically the overlap and discrepancy between brain functional mapping as revealed by EEG/MEG and fMRI, respectively.

In summary, the extended framework affords a general methodology for solving the EEG/MEG inverse problem, accounting for multiple independent priors and for evaluating quantitatively their relative contributions.

So far, this framework has been applied to static data averaged over narrow time-windows. Since the source covariance might change with time, the model selection could be extended dynamically to estimate temporal dynamics as proposed in Phillips et al. (2004). In multimodal integration, this might contribute to the better understanding of the dynamics of neural networks revealed by fMRI data.

Finally, the same framework can be extended to estimate induced responses which are not phased-locked to the stimulus or task manipulation. This extension is described in a companion paper (Friston et al., in press). Note also that to account for the full model complexity, one should model uncertainty about the hyperparameters. This could be achieved using Laplace or Variational approximations, or indeed an adjustment to the log-evidence as described in Trujillo-Barreto et al. (2004). In its current form, the ReML approach does not account for the increase in model complexity with increasing numbers of hyperparameters (i.e., covariance components). Quantitatively, this additional complexity is small relative to the accuracy components of the Log-evidence. However, we have pursued this issue by augmenting the ReML objective function to provide a better evaluation of the Log-evidence (Friston et al., Variational Free energy and the Laplace approximation, in preparation).

Acknowledgments

The Wellcome Trust funded this work and Jeremie Mattout is funded by an EC Marie Curie fellowship.

Appendix A. The L-curve approach

This heuristic involves estimating the WMN solution for various values of a hyperparameter λ . Plotting the norm of the prior term against the norm of the data fit term leads to an L-shape curve whose inflection point indicates an optimal hyperparameter value. It amounts maximizing the following log-likelihood

$$\ln p(\mathbf{M}/\lambda) = -\frac{1}{2}(\mathbf{M} - \mathbf{K}\hat{\mathbf{J}})^T \mathbf{C}_n^{-1}(\mathbf{M} - \mathbf{K}\hat{\mathbf{J}}) - \frac{1}{2}\hat{\mathbf{J}}^T \mathbf{C}_p^{-1}\hat{\mathbf{J}}, \quad (13)$$

where $\hat{\mathbf{J}}$ indicates the current estimate of the parameters.

The more exhaustive the discrete scanning of hyperparameter space, the more precise the estimation of this optimal value. An important drawback of this approach thus lies in the need for a large number of estimations to find an appropriate level of regularization.

In this paper, we used this approach as a reference for the estimation of the single hyperparameter in the WMN simulations. Practically, we used

$$\lambda = \beta \times \frac{\|\mathbf{K}\mathbf{K}^T\|}{n}, \quad (14)$$

and considered the following [30] values for β : [0.00005 0.000075 0.0001 0.0002 0.0003 0.0004 0.0005 0.0006 0.0007 0.0008 0.0009 0.001 0.002 0.003 0.004 0.005 0.006 0.007 0.008 0.009 0.01 0.02 0.03 0.04 0.05 0.075 0.1 0.5 1 50].

This sampling of hyperparameter space was chosen to ensure an L-shape curve with a relatively fine sampling in the vicinity of the inflection point.

Appendix B. The ReML objective function

The Restricted Maximum Likelihood (ReML) scheme allows the simultaneous estimation of parameters \mathbf{J} and hyperparameters (λ, μ) while accounting for the loss of degrees of freedom in the model incurred from estimating \mathbf{J} when (λ, μ) are calculated (Harville, 1977). The ReML estimation consists of maximizing the likelihood of the observed data $p(\mathbf{M}/\lambda, \mu)$, conditional on the hyperparameters, in the presence of parameters \mathbf{J} . This is equivalent to maximizing the log-likelihood

$$\ln p(\mathbf{M}/\lambda, \mu) = \ln \int p(\mathbf{J}, \mathbf{M}/\lambda, \mu) d\mathbf{J} \geq F(q(\mathbf{J}), \lambda, \mu), \quad (15)$$

where F corresponds to the negative free energy in statistical thermodynamics and is given by the sum of potential and entropy terms

$$F(q(\mathbf{J}), \lambda, \mu) = \int q(\mathbf{J}) \ln p(\mathbf{J}, \mathbf{M}/\lambda, \mu) d\mathbf{J} - \int q(\mathbf{J}) \ln q(\mathbf{J}) d\mathbf{J}, \quad (16)$$

where $q(\mathbf{J})$ is any distribution over the model parameters (Neal and Hinton, 1998).

The log-likelihood is implicitly maximized through the maximization of F . This is achieved using an EM algorithm which alternates between maximizing F with respect to the distribution $q(\mathbf{J})$ (E-step) and the hyperparameters (λ, μ) (M-step) until convergence. It can be shown that the maximum in the E-step is obtained when $q(\mathbf{J}) = p(\mathbf{J}/\mathbf{M}, \lambda, \mu)$ at which point the log-likelihood is strictly equal to F . The M-step finds the Maximum Likelihood (ML) estimate of the hyperparameters by integrating $p(\mathbf{J}, \mathbf{M}/\lambda, \mu)$ over the parameters using the current estimate of their conditional distribution.

In practice, the ReML objective function is obtained by embedding the E-step into the M-step, substituting the posterior variance estimate $\mathbf{C}_{\mathbf{J}/\mathbf{M}} = (\mathbf{K}^T \mathbf{C}_n^{-1} \mathbf{K} + \mathbf{C}_p^{-1})^{-1}$ into the expression for the negative energy so that

$$F = -\frac{1}{2}(\mathbf{M} - \mathbf{K}\hat{\mathbf{J}})^T \mathbf{C}_n^{-1}(\mathbf{M} - \mathbf{K}\hat{\mathbf{J}}) - \frac{1}{2}\hat{\mathbf{J}}^T \mathbf{C}_p^{-1}\hat{\mathbf{J}} - \frac{1}{2}\ln|\mathbf{C}_n| - \frac{1}{2}\ln|\mathbf{C}_p| - \frac{1}{2}\ln|\mathbf{K}^T \mathbf{C}_n^{-1} \mathbf{K} + \mathbf{C}_p^{-1}| + \text{const.} \quad (17)$$

We refer the reader to [Friston et al. \(2002\)](#) for a detailed description of the ReML objective function and the implementation of EM⁷ and to [Phillips et al. \(2005\)](#) for a detailed formulation in the context of EEG/MEG. In brief, this substitution eliminates the need for an explicit E-step and ReML estimates of the hyperparameters require only iterations of the M-step. An explicit E-step can be performed after convergence to obtain \mathbf{J} according to Eqs. (3) and (4).

References

- Amblard, C., Lapalme, E., Lina, J.M., 2004. Biomagnetic source detection by maximum entropy and graphical models. *IEEE Trans. Biomed. Eng.* 51, 427–442.
- Babiloni, F., Carducci, F., Babiloni, C., Urbano, A., 1998. Improved realistic laplacian estimate of highly-sampled EEG potentials by regularization techniques. *Electroencephalogr. Clin. Neurophysiol.* 106, 336–343.
- Babiloni, F., Carducci, F., Cincotti, F., Del Gratta, C., Pizzella, V., Romani, G.L., Rossini, P.M., Tecchio, F., Babiloni, C., 2001. Linear inverse source estimate of combined EEG and MEG data related to voluntary movements. *Hum. Brain Mapp.* 14, 197–209.
- Babiloni, F., Babiloni, C., Carducci, F., Romani, G.L., Rossini, P.M., Angelone, L.M., Cincotti, F., 2004. Multimodal integration of EEG and MEG data: a simulation study with variable signal-to-noise ratio and number of sensors. *Hum. Brain Mapp.* 22, 52–62.
- Baillet, S., Garnero, L., 1997. A Bayesian approach to introducing anatomic-functional priors in the EEG/MEG inverse problem. *IEEE Trans. Biomed. Eng.* 44, 374–385.
- Baillet, S., Mosher, J.C., Leahy, R.M., 2001. Electromagnetic brain mapping. *IEEE Signal Process. Mag.* 18, 14–30.
- Dale, A.M., Halgren, E., 2001. Spatiotemporal mapping of brain activity by integration of multiple imaging modalities. *Curr. Opin. Neurobiol.* 11, 202–208.
- Dale, A.M., Sereno, M., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI surface reconstruction: a linear approach. *J. Cogn. Neurosci.* 5, 162–176.
- Dale, A.M., Liu, A.K., Fischl, B.R., Buckner, R.L., Belliveau, J.W., Lewine, J.D., Halgren, E., 2000. Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron* 26, 55–67.
- Daunizeau, J., Grova, C., Mattout, J., Marrelec, G., Clonda, D., Goulard, B., Péligrini-Issac, M., Lina, J.M., Benali, H., 2005. Assessing the relevance of fMRI-based prior in the EEG inverse problem: a Bayesian Model Selection approach. *IEEE Trans. Sign. Proc.* 53, 3461–3472 (Special issue).
- Ermer, J.J., Mosher, J.C., Baillet, S., Leahy, R.M., 2001. Rapidly recomputable EEG forward models for realistic head shapes. *Phys. Med. Biol.* 46, 1265–1281.
- Friston, K., Penny, W., Phillips, C., Kiebel, S., Hinton, G., Ashburner, J., 2002. Classical and Bayesian inference in neuroimaging: theory. *NeuroImage* 16, 465–483.
- Friston, K., Henson, R., Phillips, C., Mattout, J., in press. Bayesian estimation of evoked and induced responses. *Hum. Brain Mapp.*
- Golub, G.H., Heath, M., Wahba, G., 1979. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* 21, 215–223.
- Gonzalez Andino, S.L., Blanke, O., Lantz, G., Thut, G., Grave de Peralta Menendez, R., 2001. The use of functional constraints for the neuroelectromagnetic inverse problem: alternatives and caveats. *Int. J. Bioelectromagn.* (<http://www.ijbem.org/volume3/number1/gravedeperalta/index.htm>).
- Gorodnitsky, I.F., George, J.S., Rao, B.D., 1995. Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. *Electroencephalogr. Clin. Neurophysiol.* 95, 231–251.
- Hämäläinen, M.S., Ilmoniemi, R.J., 1994. Interpreting magnetic fields of the brain—minimum norm estimates. *Med. Biol. Eng. Comput.* 32, 35–42.
- Hansen, P.C., 1992. Analysis of discrete ill-posed problems by means of L-curve. *SIAM Rev.* 34, 561–580.
- Harville, D.A., 1977. Maximum likelihood approaches to variance component estimation and related problems. *J. Am. Stat. Assoc.* 72, 320–338.
- Hauk, O., 2004. Keep it simple: a case for using classical minimum norm estimation in the analysis of EEG and MEG data. *NeuroImage* 21, 1612–1621.
- Jerbi, K., Baillet, S., Mosher, J.C., Nolte, G., Garnero, L., Leahy, R.M., 2004. Localization of realistic cortical activity in meg using current multipoles. *NeuroImage* 22, 779–793.
- Kass, R.E., Raftery, A.E., 1995. Bayes factors. *J. Am. Stat. Assoc.* 90, 773–795.
- Koles, Z.J., 1998. Trends in EEG source localization. *Electroencephalogr. Clin. Neurophysiol.* 106, 127–137.
- Liu, A.K., Belliveau, J.W., Dale, A.M., 1998. Spatiotemporal imaging of human activity using functional MRI constrained magnetoencephalography data: Monte Carlo simulations. *Proc. Natl. Acad. Sci. U. S. A.* 95, 8945–8950.
- Mangin, J.F., 1995. From 3D magnetic resonance images to structural representations of the cortex topography using topology preserving deformations. *J. Math. Imaging Vis.* 5, 297–318.
- Mattout, J., Péligrini-Issac, M., Bellio, A., Daunizeau, J., Benali, H., 2003. Localization Estimation Algorithm (LEA): a supervised prior based approach for solving the EEG/MEG inverse problem. In: Taylor, C.J., Noble, J.A. (Eds.), *Information Processing in Medical Imaging (IPMI 2003)*, LNCS 2732. Springer-Verlag, Berlin, pp. 536–547.
- Mattout, J., Péligrini-Issac, M., Garnero, L., Benali, H., 2005. Multivariate Source Prelocalization (MSP): use of functionally informed basis functions for better conditioning the MEG inverse problem. *NeuroImage* 26, 356–373.
- Metz, C.E., 1998. Statistical comparison of two ROC-curve estimates obtained from partially paired datasets. *Med. Decis. Mak.* 18, 110–121.
- Mosher, J.C., Leahy, R.M., Lewis, P.S., 1999. EEG and MEG: forward solutions for inverse methods. *IEEE Trans. Biomed. Eng.* 46, 245–259.
- Neal, R.M., Hinton, G.E., 1998. A view of the EM algorithm that justifies incremental, sparse and other variants. *Learning in Graphical Models*. Kluwer Academic Press, pp. 355–368.

⁷ This algorithm is implemented (spm_reml.m) in the current version of the SPM software (Wellcome Department of Imaging Neuroscience, <http://www.fil.ion.ucl.ac.uk/spm> (2002)).

- Nunez, P.L., Silberstein, R.B., 2000. On the relationship of synaptic activity to macroscopic measurements: does co-registration of EEG with fMRI make sense? *Brain Topogr.* 13, 79–96.
- Pascual-Marqui, R.D., Michel, C.M., Lehmann, D., 1994. Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. *IEEE Trans. Biomed. Eng.* 418, 49–65.
- Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004. Comparing dynamic causal models. *NeuroImage* 22, 1157–1172.
- Pfliederer, M.E., Greenblatt, R.E., 2001. Nonlinear analysis of multimodal dynamic brain imaging data. *Int. J. Bioelectromagnetism* 3 (<http://www.ijbem.org/volume3/number1/greenblatt/index.htm>).
- Phillips, C., Rugg, M.D., Friston, K.J., 2002. Systematic regularisation of linear inverse solutions of the EEG source localisation problem. *NeuroImage* 17, 287–301.
- Phillips, C.L., Mattout, J., Maquet, P., Friston, K.J., 2004. Dynamic reml solution for the EEG source reconstruction problem. *Hum. Brain Mapp. (HBM) Conference*, WE 280.
- Phillips, C., Mattout, J., Rugg, M.D., Maquet, P., Friston, K., 2005. An empirical Bayesian solution to the source reconstruction problem in EEG. *NeuroImage* 24, 997–1011.
- Riera, J.J., Fuentes, M.E., Valdes, P.A., Oharriz, Y., 1998. Eeg distributed inverse solutions for a spherical head model. *Inverse Problems* 14, 1009–1019.
- Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys. Med. Biol.* 32, 11–22.
- Scherg, M., von Cramon, D., 1986. Evoked dipole source potentials of the human auditory cortex. *Electroencephalogr. Clin. Neurophysiol.* 65, 344–360.
- Schmidt, D., George, J., Wood, C., 1999. Bayesian inference applied to the electromagnetic inverse problem. *Hum. Brain Mapp.* 7, 195–212.
- Tikhonov, A.N., Arsenin, V.Y., 1977. *Solutions of Ill-Posed Problems*. John Wiley, New York.
- Trujillo-Barreto, N.J., Aubert-Vazquez, E., Valdes-Sosa, P.A., 2004. Bayesian model averaging in EEG/MEG imaging. *NeuroImage* 21, 1300–1319.
- Waldorp, L.J., Huizenga, H.M., Nehorai, A., Grasman, R.P., Molenaar, P.C., 2005. Model selection in spatio-temporal electromagnetic sources analysis. *IEEE Trans. Biomed. Eng.* 52, 414–420.
- Yao, J., Dewald, J.P.A., 2005. Evaluation of different cortical source localization methods using simulated and experimental EEG data. *NeuroImage* 25, 369–382.