

Statistical Approaches to the Inverse Problem

A. Pascarella¹ and A. Sorrentino²

¹*Dipartimento di Neuroscienze, Università di Parma*

²*Department of Statistics, University of Warwick*

¹*Italy*

²*United Kingdom*

1. Introduction

Magnetoencephalography (MEG) can be considered as one of the most powerful instruments for non-invasive investigation of the brain functions (Hämäläinen et al., 1993). Its good temporal resolution (sampling frequency can reach several thousands of Hertz), only comparable to the sampling frequency of electroencephalography (EEG), allows following the neural dynamics on a millisecond time scale. This implies that, with respect to other functional imaging modalities, MEG provides an enormous amount of information about the brain activity. On the other hand, localization of brain activity from MEG data requires to solve a highly ill-posed inverse problem (Sarvas, 1987): exact reconstruction of the neural current is not possible even with noise-free data, and the debate about the actual spatial resolution of MEG is still open. Despite this drawback, the appeal of MEG data is such that a large number of methods have been developed for localizing brain activity from MEG recordings. Starting from the middle Eighties, when the first Minimum Norm Estimate (Hämäläinen & Ilmoniemi, 1984) appeared, to the recent developments in spatio temporal regularization (Ou et al., 2009), state space models (Long et al., 2006; Sorrentino et al., 2009) and Markov Chain Monte Carlo (Jun et al., 2005), the attempts to insert enough a priori information to actually provide stable and reliable solutions have been rising rather than diminishing. The goal and the scope are clear: automatic and reliable source localization can be of practical use in several fields, possibly including clinical applications (pre-surgical evaluation, foci localization in epilepsy), new technology development (e.g. brain computer interfaces), and scientific research (e.g. for analysis of large datasets).

Despite having been presented in rather different frameworks and with largely different names and techniques, most methods appeared so far in the MEG inverse literature can be described rather straightforwardly in a statistical setting, which helps unifying and clarifying the assumptions and the limitations of the methods. In this chapter, we attempt a brief review of the most largely known methods, as well as of some of the most recently appeared ones, highlighting the statistical features. We first provide a brief description of the statistical approach to inverse problems in Section 2; then we describe the main features of the MEG inverse problem in Section 3; in Section 4 and 5 we review methods developed by other authors, while in Section 6 we discuss our work on Bayesian filtering for MEG.

2. Statistics and inverse problems

2.1 Inverse problems

In applied mathematics the term inverse problem denotes a class of problems, often originating from physics equations, where the unknown is an unobserved quantity which has to be recovered from indirect measurements. Rather often, there is a causal relationship between the unknown and the data, and the inverse problem involves moving backwards from a measured effect to an unknown cause. An excellent source of inverse problems is the field of medical imaging, where one tries to recover information about quantities within the body from external measurements.

Inverse problems are usually represented in functional form as

$$g = A(f) \quad (1)$$

where $g \in H_2$, $f \in H_1$, and $A : H_1 \rightarrow H_2$ is an operator between the two possibly distinct functional spaces (e.g. Banach spaces, Hilbert spaces, etc.) H_1 and H_2 . From a mathematical viewpoint, the inverse problem consists in finding the function f given the function g and the operator A . In most cases of practical interest, however, the solution of the inverse problem cannot be found exactly, because the operator A is pathological: either its kernel is non-empty - leading to non-uniqueness of the inverse solution - or the range of A is not the whole H_2 - leading to non-existence problems - or the inverse operator, when it exists, is not continuous - leading to stability problems. In applied mathematics, the typical solution to these pathologies is provided in the celebrated regularization framework for linear inverse problems (Bertero & Boccacci, 1998), where the ill-posed problem is replaced by a one-parameter family of well-posed ones. The most common regularization algorithm, Tichonov method, consists in finding the minimizer \hat{f} of

$$\|g - A \cdot f\|_{H_2}^2 + \lambda \|f\|_{H_1}^2 \quad (2)$$

which is rather straightforwardly

$$\hat{f} = (A^*A + \lambda I)^{-1}A^*g \quad (3)$$

where A^* is the adjoint of A and λ is the regularization parameter, which tunes the balance between stability of the solution and its capability to explain the measured data. More in general, the recipe of any regularization algorithm is to minimize a functional which combines the misfit - the difference between the actual data and the model - with some type of norm of the solution - implicitly encouraging some stability property for the solution. The choice of this norm, together with a criterion for choosing the value of the regularization parameter, are key ingredients which highly depend on the specific problem.

2.2 Statistics for inverse problems

Given the stochastic nature of noise affecting the data, it is natural to view inverse problems as statistical inference problems. The core of this approach is to model the data as a Random Variable (RV). Then, there exist two main frameworks for dealing with the inverse problem: in the Maximum Likelihood (ML) framework, the unknown is viewed as a set of parameters which characterize the probability distribution of the data, and the inverse problem is re-phrased as a parameter estimation problem; in the Bayesian framework, the unknown is viewed as a Random quantity as well, and the inverse problem is re-phrased as the problem of calculating the posterior distribution of the unknown, given the data.

More specifically, let y be the realization of the Random Variable Y , representing the measured data. In the ML framework, x represents the unknown parameters; the distribution of the data $L(y|x)$ is called likelihood function and contains both the physical model - the operator A mapping the unknown onto the data - and the assumptions on the statistical distribution of noise. Typically one wants to find the x which maximizes L for the observed noisy data y ; such

$$\hat{x} = \arg \max(L(y|x)) \quad (4)$$

is naturally called Maximum Likelihood estimate of x .

In the Bayesian framework, on the other hand, x is viewed as the realization of the Random Variable X rather than as a parameter, and it is given a prior distribution; this prior is then combined with the likelihood through Bayes theorem to produce the so called posterior distribution, i.e. the probability distribution for the unknown, conditioned on the realization of the data. The solution of the inverse problem is here defined to be the posterior distribution of the unknown, which can then be used to compute point estimates and bounds. More specifically, let $p(x)$ be the prior probability distribution for the unknown, embodying all the information which is known a priori, before the data are collected; let $L(y|x)$ be the likelihood function, the same function which is used in the ML framework; the posterior distribution is given by Bayes theorem as

$$p(x|y) = \frac{L(y|x)p(x)}{p(y)} \quad (5)$$

where $p(y)$ is the normalizing constant and is given as $p(y) = \int L(y|x)p(x)dx$.

2.3 The linear Gaussian case

Let us briefly consider a linear inverse problem with additive zero-mean Gaussian noise N , with standard deviation σ ; in the ML framework, the model equation is

$$Y = Ax + N \quad (6)$$

and the likelihood function is

$$L(y|x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\|y - Ax\|^2}{2\sigma^2}\right). \quad (7)$$

The ML estimate in this example is the solution of the least squares problem.

In Bayesian terms, the model equation becomes

$$Y = AX + N \quad (8)$$

and the posterior distribution under zero-mean Gaussian prior for the unknown turns out to be:

$$p(x|y) \propto \exp\left(-\frac{\|y - Ax\|^2}{2\sigma^2}\right) \exp\left(-\frac{\|x\|^2}{2\gamma^2}\right) \quad (9)$$

where γ is the standard deviation of the prior. Comparison of equations (9) and (2) suggests that the minimizer of the Tichonov functional is the mode of the posterior distribution, provided that the regularization parameter is identified with the ratio between the noise variance and the variance of the prior distribution.

As this simple example suggests, in the statistical framework regularization is interpreted as providing a priori information to solve the ill-posed problem. Furthermore, the statistical framework is an elegant and unifying framework where linear as well as non-linear problems can be formally described. Rather than providing computational recipes, this framework provides a theoretical background which is general enough for different kinds of problems, for different kinds of distributions of the noise, and always provides a meaningful interpretation of the variables and parameters in the game.

3. Basic facts about the MEG inverse problem

3.1 Ill posedness

The magnetic field \mathbf{b} produced by a given electrical current distribution $\mathbf{j}(\mathbf{r})$ inside a volume V and measured at a point $\mathbf{r} \notin V$ is given by the Biot-Savart equation

$$\mathbf{b}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_V \mathbf{j}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} d\mathbf{r}' . \quad (10)$$

It has been known for 150 years that the problem of estimating the electrical current from measurements of the magnetic field is an ill posed problem. Recently a few rigorous proofs have appeared in the literature, which are summarized in the following.

Let $\mathcal{B}_1 : (C(V))^3 \rightarrow (C(D))^3$ be the operator described by (10), mapping the electrical current distribution, belonging to the space $(C(V))^3$ of continuous function within the brain volume V , onto the magnetic field, belonging to the space $(C(D))^3$ of continuous functions on a surface D which contains the boundary of V . In Kress et al. (2002) it has been proved that the kernel of \mathcal{B}_1 contains the linear subspace

$$M = \{\mathbf{j} | \mathbf{j} = \Delta m, m \in (C_0^2(V))^3\}. \quad (11)$$

This implies that the inverse problem has no unique solution; in particular, given a solution \mathbf{j}^* of the inverse problem (10), any current distribution of the form $\tilde{\mathbf{j}} = \mathbf{j}^* + \lambda \mathbf{j}'$, with $\mathbf{j}' \in M$ and $\lambda \in \mathbb{R}$ is also a solution of the problem.

Let $\mathcal{B}_2 : (L^2(V))^3 \rightarrow (L^2(V))^3$ be the operator described by (10), mapping the electrical current distribution in the volume V onto the magnetic field within the same volume V . In this case, all three components of both the electrical current and the magnetic field are assumed to be L^2 functions. In Cantarella et al. (2001), it has been proved that \mathcal{B}_2 is a compact operator. This result implies that the generalized inverse is not a continuous operator.

3.2 The inverse model

The Biot Savart equation (10) is only a partial description of the MEG inverse problem: there are two main issues which need to be discussed in order to provide computational solutions to the inverse problem. The first issue, only mentioned here because it is not relevant in the rest of the Chapter, concerns the geometrical and physical model used to approximate the head; it is usually referred to as the forward problem, because it amounts to accurate calculations of the magnetic field produced by a known neural current distribution. Due to the non-zero conductivity of the tissues inside the head, the electrical currents elicited by brain activity are only a portion of the electrical currents flowing in the brain; this means that $\mathbf{j}(\mathbf{r})$ in (10) contains both neural activity and passively flowing currents, which are generated

by the neural activity itself and contribute to the measured field as well. The forward problem accounts for the mechanisms by which the neural currents generate the passive currents, and how these passive currents flow in the brain. In the simplest model, available since the Eighties, the head is modelled as a homogeneous spherical conductor, and the effect of the passive currents can be calculated analytically (Sarvas, 1987). A more realistic model, nowadays largely used, takes into account the actual shape of the brain, while maintaining a piecewise homogeneous conductivity, and is solved with numerical integration techniques known as Boundary Element Methods (Gramfort et al., 2011). More recently the use of Finite Element Methods (Pursiainen et al., 2011), far more computationally demanding, allows implementation of completely non-homogeneous and even anisotropic conductivity models.

The second relevant point is the discretization of equation (10). The standard approach comprises discretization of the brain volume with a finite set of \mathcal{N}_{points} points, possibly distributed along the cortical surface. The neural current is then represented as a vector field on these support points. The single current element, with support in a single point of the brain, is usually referred to as *current dipole*. It is characterized by a location \mathbf{r} and a dipole moment \mathbf{q} ; mathematically speaking it is an applied vector, while from the physical viewpoint it represents the first order approximation of any current distribution. The whole neural current is then viewed as the superposition of thousands of current dipoles:

$$\mathbf{j}(\mathbf{r}) = \sum_{k=1}^{\mathcal{N}_{points}} \mathbf{q}^k \delta(\mathbf{r} - \mathbf{r}^k) \quad (12)$$

where \mathbf{q}^k is the dipole moment of the current dipole at location \mathbf{r}^k and $\delta(\cdot)$ is the Dirac delta. The Biot-Savart operator is discretized in a matrix $\mathbf{G} = [\mathbf{G}(\mathbf{r}^1), \dots, \mathbf{G}(\mathbf{r}^{\mathcal{N}_{points}})]$ of size $\mathcal{N}_{sensors} \times 3\mathcal{N}_{points}$ where each column of the matrix $\mathbf{G}(\mathbf{r}^k)$ represents the magnetic field produced by a unit current dipole at a given position and oriented according to one of the three orthogonal directions; this matrix is usually referred to as the leadfield matrix. The discretized mathematical model becomes here

$$\mathbf{b} = \mathbf{G} \cdot \mathbf{j} \quad (13)$$

where the vector \mathbf{b} is $\mathcal{N}_{sensors} \times 1$ and contains the collection of data measured by all the sensors, while the vector $\mathbf{j} = [\mathbf{j}^1, \dots, \mathbf{j}^{\mathcal{N}_{points}}]^T$ is $3\mathcal{N}_{points} \times 1$ and contains the dipole moments of the current dipoles located at the fixed locations contained in the leadfield \mathbf{G} . In the remainder of the Chapter we will denote by \mathbf{b}_t and \mathbf{j}_t the measured data and the neural current at time t ; the leadfield matrix is usually assumed to be not time dependent.

3.3 Classification of inverse methods

Equation (13) is the main equation of practical interest for the solution of the inverse problem. It reflects the fact that the inverse problem is linear, in the sense that the measured data are the linear superpositions of the contributions of individual current dipoles. The physical and geometrical model discussed in the previous section is embodied in the matrix \mathbf{G} , which is assumed to be known when solving the inverse problem. However, talking about *the* MEG inverse problem may be misleading for the following reason. While in the general case one will model the neural current as a vector field with support in the whole brain, and the inverse problem amounts to inversion of the (non-invertible) matrix \mathbf{G} , in an alternative approach one can assume that the brain activity is highly localized in few brain regions; if these regions

are small compared to their distance from the sensors, each region can be approximated as a single current dipole; in this case, the inverse problem consists in estimating the number of sources, their location, orientation and strength. The mathematical properties of the two inverse problems are rather different. In the former case, one will face a linear equation resulting in a highly under-determined problem due to the ratio between the typical number of sensors ($\simeq 100$) and the typical number of points ($\simeq 10,000$). In the following, we will call *imaging methods* the methods based on a continuous current distribution model. In the second case, one will want to estimate the intensity, the orientation, as well as the number and the locations of the point sources; the dependence of the magnetic field on the two latter parameters is highly non-linear, thus resulting in a genuinely non-linear problem. In the following, we will call *parametric methods* the methods based on a point source model.

The classification of inverse methods in imaging and parametric methods is rather common practice in many reviews and dates back at least to Hämäläinen et al. (1993), where dipole methods are separated from regularization methods. In this chapter, we introduce a new - to the best of our knowledge - classification of inverse methods, which accounts for a recent trend in the MEG inverse problem community. We will call *static methods* the inverse methods which do not give a priori information on the temporal behaviour of the neural sources, and *dynamic methods* those which do provide a prior on it. The recent appearance of dynamic methods is indeed a consequence of the increased availability of cheap computational power. Dynamic methods are far more computationally demanding for the simple reason that they perform regularization in a much higher-dimensional space. At the same time, prior information in the temporal domain can really help in stabilizing the solution of this highly ill-posed problem.

4. Imaging methods

Imaging methods solve the linear inverse problem defined by

$$\mathbf{B}_t = \mathbf{G}\mathbf{J}_t + \mathbf{N}_t \quad (14)$$

where: $\mathbf{B}_t \in \mathbb{R}^{\mathcal{N}_{sensors}}$ is the Random Vector representing the measured data at time t ; \mathbf{G} is the leadfield matrix as in (13); $\mathbf{J}_t \in \mathbb{R}^{3\mathcal{N}_{points}}$ is the Random Vector representing the neural current distribution at time t ; $\mathbf{N}_t \in \mathbb{R}^{\mathcal{N}_{sensors}}$ explicitly accounts for the presence of noise in the data. The output of an imaging method is an estimate of \mathbf{J}_t for every t ; as such, it is a sequence of current distributions which can be visualized as a dynamic intensity map representing the flow of brain activity.

From a statistical perspective, all the imaging methods we will describe can be interpreted as Bayesian estimation methods; this is due to the ill-posedness of the inverse problem which requires rather strong regularization algorithms, i.e. a priori information, to get stable solutions. Due to the fact that most of these methods have been first described in a regularization/optimization framework, the estimate they provide is always easily interpreted as the MAP of the posterior distribution. All methods share basically the same likelihood function: since noise is assumed to have a zero-mean Gaussian distribution and to be uncorrelated with the sources, the likelihood function is:

$$\pi(\mathbf{b}_t | \mathbf{j}_t) = \frac{1}{\sqrt{(2\pi)^M |\boldsymbol{\Sigma}|}} \exp\left(-(\mathbf{b}_t - \mathbf{G}\mathbf{j}_t)^T \boldsymbol{\Sigma}^{-1} (\mathbf{b}_t - \mathbf{G}\mathbf{j}_t)\right) \quad (15)$$

where $\mathbf{b}_t, \mathbf{j}_t$ are realizations of the data and the unknown RV, $\boldsymbol{\Sigma}$ is the noise covariance matrix and $M = \mathcal{N}_{sensors}$.

What really makes each method unique is the specific choice of the prior distribution for the spatio/temporal distribution of the neural current.

4.1 Static imaging methods - Minimum Norm Estimate

Minimum Norm Estimate (MNE) is perhaps the most popular method for solving the MEG inverse problem. It has been introduced almost thirty years ago and is the straightforward implementation of the Tichonov algorithm. As seen in Section 2, this choice corresponds to a Gaussian prior distribution

$$\pi(\mathbf{j}_t) = \frac{1}{\sqrt{2\pi}\gamma} \exp\left(-\frac{\|\mathbf{j}_t\|^2}{2\gamma^2}\right). \quad (16)$$

The explicit form of the MNE solution, i.e. the MAP of the posterior distribution under the given prior, is given by

$$\hat{\mathbf{j}}_t = (\mathbf{G}^T \boldsymbol{\Sigma}^{-1} \mathbf{G} + \lambda \mathbf{I})^{-1} \mathbf{G}^T \boldsymbol{\Sigma}^{-1} \mathbf{b}_t \quad (17)$$

where the regularization parameter is $\lambda \propto \frac{1}{\gamma^2}$. Due to the Gaussian spatial prior, solutions provided by MNE tend to be rather smooth, a feature that makes MNE less suitable than other methods for estimating focal activity. Furthermore, another limitation of MNE is that the zero-mean Gaussian prior gives an implicit preference to superficial sources, which need lower intensity to produce the measured field. This results in an estimation bias which is stronger for deeper sources. Despite these drawbacks, MNE is a very attractive method thanks to its simplicity and to the very low computational cost.

4.2 Static imaging methods - weighted MNE, LORETA and LCMV beamformer

The class of weighted minimum norm solutions has been developed with the aim of overcoming the depth bias introduced by the simple prior used by MNE. The main difference with respect to standard MNE is that the prior is still Gaussian but the source covariance matrix is no longer proportional to the identity matrix; instead, the source covariance matrix contains a weighting factor for each source, aimed at removing the bias towards superficial solutions. In Lin et al. (2006), for instance, the $3\mathcal{N}_{points} \times 3\mathcal{N}_{points}$ source covariance matrix \mathbf{R} is diagonal with

$$\mathbf{R}_{j,j} = \|\mathbf{G}(\mathbf{r}^k)\|^{-p}, j = 3k - 2, 3k - 1, 3k, k = 1, \dots, \mathcal{N}_{points} \quad (18)$$

where $\mathbf{G}(\mathbf{r}^k)$ is the lead field corresponding to the k -th source point, $\|\cdot\|$ is the euclidean norm and p is the depth weighting parameter. With this choice, deeper sources have a Gaussian prior with larger variance with respect to shallower sources. The final solution of WMNE is given by the two following equivalent formulations

$$\begin{aligned} \hat{\mathbf{j}}_t &= \mathbf{R} \mathbf{G}^T (\mathbf{G} \mathbf{R} \mathbf{G}^T + \boldsymbol{\Sigma})^{-1} \mathbf{b}_t = \\ &= (\mathbf{G}^T \boldsymbol{\Sigma}^{-1} \mathbf{G} + \mathbf{R}^{-1})^{-1} \mathbf{G}^T \boldsymbol{\Sigma}^{-1} \mathbf{b}_t. \end{aligned} \quad (19)$$

Low Resolution Electromagnetic Tomography (LORETA) (Pascual-Marqui et al., 1994) is a WMNE with the further addition of a Laplacian operator. The prior for LORETA is given by

$$\pi(\mathbf{j}_t) \propto \exp\left(-\mathbf{j}_t^T \mathbf{R}^T \mathbf{L}^T \mathbf{L} \mathbf{R} \mathbf{j}_t\right) \quad (20)$$

where \mathbf{R} is as in (18) and \mathbf{L} is the discrete spatial laplacian operator (Pascual-Marqui, 1999). Such operator \mathbf{L} produces a smoother estimate of the current distribution since the inverse

matrix \mathbf{L}^{-1} implements a discrete spatial smoothing operator. LORETA has been shown to provide accurate localization even for deep sources.

The term *beamformer* encompasses a broad range of algorithms which have become popular in the last fifteen years. In one way or the other, most or all of them can be proved to be particular choices of WMNE (Moshier et al., 2003). Here we only consider the Linearly Constrained Minimum Variance Beamformer (Van Veen et al., 1997). The derivation of this method is made in an engineering framework and is based on a constrained minimization. The idea behind beamforming is that of constructing a spatial filter which isolates the activity produced in a selected location, blocking the activity originated elsewhere. The main working assumption of the method described in (Van Veen et al., 1997) is that all the sources are uncorrelated. The result of the constrained minimization turns out to be

$$\hat{\mathbf{j}}_t^k = \left[(\mathbf{G}(\mathbf{r}^k))^T \mathbf{R}_b^{-1} \mathbf{G}(\mathbf{r}^k) \right]^{-1} (\mathbf{G}(\mathbf{r}^k))^T (\mathbf{G} \mathbf{R} \mathbf{G}^T + \mathbf{\Sigma})^{-1} \mathbf{b}(t) \quad (21)$$

where \mathbf{R}_b is the data covariance matrix and $\hat{\mathbf{j}}_t^k$ is the estimated current density at k -th location. We remark that this solution corresponds to the WMNE solution (19) if the source covariance matrix \mathbf{R} is a block diagonal matrix, composed by the 3×3 matrices:

$$\mathbf{R}(\mathbf{r}^k) = \left[(\mathbf{G}(\mathbf{r}^k))^T \mathbf{R}_b^{-1} \mathbf{G}(\mathbf{r}^k) \right]^{-1} \quad (22)$$

and zero elsewhere. Importantly, the assumption that the neural sources are not correlated in time is a strong assumption, often not verified in experimental data. As a consequence of this assumption, however, beamformers face severe difficulties when the data actually contain signals produced by correlated sources. An improved version of the beamforming technique has been proposed, which tries to overcome this limitation (Sekihara et al., 2002).

4.3 Static imaging methods - dSPM and sLORETA

Dynamic Statistical Parametric Mapping (dSPM) and Standardized LORETA (sLORETA) combine Bayesian inference for the electrical current with noise normalization. The output of these methods is no longer an estimate of the neural current distribution, but rather a statistical measure of brain activity.

Instead of imaging the estimated electrical current intensity, dSPM (Dale et al., 2000) produces maps of the following neural activity index

$$z_t^k = \frac{\|\hat{\mathbf{j}}_t^k\|_2^2}{\|(\mathbf{W}(\mathbf{r}^k))^T \mathbf{\Sigma} \mathbf{W}(\mathbf{r}^k)\|_F^2} \quad (23)$$

where $\mathbf{W}(\mathbf{r}^k)$ is the regularized inverse operator in (19)

$$\mathbf{W}(\mathbf{r}^k) = (\mathbf{G} \mathbf{R} \mathbf{G}^T + \mathbf{\Sigma})^{-1} \mathbf{G}(\mathbf{r}^k) \mathbf{R}(\mathbf{r}^k) \quad (24)$$

$\mathbf{\Sigma}$ is the noise covariance matrix, and $\|\cdot\|_F$ is the Frobenius norm; in practice, the output of a weighted MNE $\hat{\mathbf{j}}_t$ at the k -th location is divided by an estimate of the portion of the source variance, at the same location, which is due to the noise in the data. Under the null hypothesis of zero brain activity, the resulting quantity z_t has a Student's t distribution.

sLORETA (Pascual-Marqui, 2002) consists in a similar approach, but the standardization is obtained considering the variance of the estimated sources (19), instead of using just the variance due to the noise component. SLORETA produces maps of the quantity

$$z_t^k = \frac{\|\hat{\mathbf{j}}_t^k\|_2^2}{\|(\mathbf{W}(\mathbf{r}^k))^T(\mathbf{G}\mathbf{R}\mathbf{G}^T + \boldsymbol{\Sigma})^{-1}\mathbf{W}(\mathbf{r}^k)\|_F^2}. \quad (25)$$

The claim of the authors is that this method produces unbiased, zero-error localization of point sources. sLORETA is largely used in EEG source analysis, where source localization is made particularly difficult by the strongest impact of the passive currents on the measured data.

4.4 Static imaging methods - minimum current estimate

The MCE algorithm (Uutela et al. (1999)) is substantially different from all the imaging methods presented so far in that it does not use a Gaussian prior for the current distribution, but an exponential prior instead:

$$\pi(\mathbf{j}_t) = \exp(-\|\mathbf{j}_t\|_1) \quad (26)$$

where $\|\mathbf{j}_t\|_1 = \sum_{k=1}^{3\mathcal{N}_{points}} |\mathbf{j}_t^k|$.

This is usually phrased in the regularization framework as using an L^1 prior, and the main consequence of this choice is that sparsity of the solution is promoted, so that estimates provided by MCE are typically highly focused.

In contrast to the MNE, MCE is therefore well suited for estimating focal activity. As a drawback, it has been suggested that MCE is in fact not capable to recover smooth current distributions, always shrinking to few points the estimated active area. From a computational viewpoint, the minimization with the L1 norm cannot be expressed any more in a closed form. Linear programming algorithms can be used to solve the minimum problem, but the computational cost is remarkably higher compared to that of methods based on a Gaussian prior.

4.5 Dynamic imaging methods - L1L2 regularization

From the conceptual point of view, the L1L2 approach (Ou et al., 2009) consists in replacing the time-varying Random Vectors in equation (14) with the full random matrices containing the whole spatio-temporal pattern of the data and of the brain activity:

$$\mathbf{B} = \mathbf{G}\mathbf{J} + \mathbf{N} \quad (27)$$

where $\mathbf{B} = [\mathbf{B}_1, \dots, \mathbf{B}_T]$, $\mathbf{J} = [\mathbf{J}_1, \dots, \mathbf{J}_T]$ and $\mathbf{N} = [\mathbf{N}_1, \dots, \mathbf{N}_T]$; the functional to be minimized is now

$$\|\mathbf{B} - \mathbf{G}\mathbf{J}\|_F^2 + \|\mathbf{J}\|_2^1 \quad (28)$$

with

$$\|\mathbf{J}\|_2^1 = \sum_{k=1}^{3\mathcal{N}_{points}} \sqrt{\sum_{t=1}^T (\mathbf{J}_t^k)^2}. \quad (29)$$

In Bayesian terms, this can be viewed as using a Gaussian prior distribution in time and a Laplacian prior distribution in space. This type of regularization promotes both sparsity in

the spatial domain, in order to produce good estimates of focal sources, and smoothness in the temporal domain. From the computational viewpoint, this method is realized by means of a projection of the data and of the source space onto a subset of basis functions, to reduce the dimensionality; a gradient method is used to minimize the functional, i.e. to calculate the approximate MAP of the posterior distribution.

4.6 Dynamic imaging methods - Kalman filter

In Kalman filtering, one explicitly models the neural currents and the measurements as two stochastic processes, $\{\mathbf{J}_1, \dots, \mathbf{J}_T\}$, $\{\mathbf{B}_1, \dots, \mathbf{B}_T\}$. Under Markovian assumptions for these processes, the posterior distribution $\pi(\mathbf{j}_t | \mathbf{b}_{1:t})$ can be computed at any time-point t applying sequentially the following two equations:

$$\pi(\mathbf{j}_t | \mathbf{b}_{1:t-1}) = \int \pi(\mathbf{j}_t | \mathbf{j}_{t-1}) \pi(\mathbf{j}_{t-1} | \mathbf{b}_{1:t-1}) d\mathbf{j}_{t-1} \quad (30)$$

$$\pi(\mathbf{j}_t | \mathbf{b}_{1:t}) = \frac{\pi(\mathbf{b}_t | \mathbf{j}_t) \pi(\mathbf{j}_t | \mathbf{b}_{1:t-1})}{\pi(\mathbf{b}_t | \mathbf{b}_{1:t-1})} \quad (31)$$

where $\{\mathbf{b}_{1:t}\}$ denotes the set of measurements up to time t and $\pi(\mathbf{j}_t | \mathbf{j}_{t-1})$ is the transition kernel which describes the evolution of the neural currents in time. Under Gaussian noise, Gaussian prior assumptions, the posterior distribution remains Gaussian at any time point, and these two equations have closed form solution through the so-called Kalman filter: at every time point, the mean and the covariance matrix of the Gaussian distribution are updated. The main difference between this approach and all the previously described approaches is that here the data are analyzed in sequence. The solution at time t depends on all the past history of observations; the prior density in the Kalman filter is time dependent and is given by

$$\pi(\mathbf{j}_{1:t}) = \pi(\mathbf{j}_0) \prod_{n=0}^t \pi(\mathbf{j}_{n+1} | \mathbf{j}_n). \quad (32)$$

The main practical issue in the application of a Kalman filter to the MEG inverse problem is that the Kalman filter requires to invert the source covariance matrix at each time step. In MEG, this covariance matrix is typically very large, being $3\mathcal{N}_{points} \times 3\mathcal{N}_{points}$. In Galka et al. (2004), the very high dimensional problem is decomposed in a collection of low dimensional problems, exploiting the local nature of the covariance matrix; in Long et al. (2006), a network of computers is instead used to solve the full problem.

5. Parametric methods

Parametric methods solve the non-linear inverse problem defined by

$$\mathbf{B}_t = \sum_{i=1}^{\mathcal{N}_{dipoles}} \mathbf{G}(\mathbf{R}_t^i) \mathbf{Q}_t^i + \mathbf{N}_t \quad (33)$$

where $\mathcal{N}_{dipoles}$ is the unknown number of active dipoles at time t , \mathbf{R}_t^i and \mathbf{Q}_t^i are the Random Vectors representing the location and the dipole moment of the i -th source at time t , respectively, and $\mathbf{G}(\mathbf{r})$ is the lead field of a unit dipole at location \mathbf{r} . Remarkably, the number of active dipoles is not expected to be higher than, say, ten; this implies that the number of unknowns in the parametric approach is considerably lower than the number of unknowns in the imaging approach. However, as already observed, the data depend now non-linearly

on the parameters, namely, on the number of sources and on the source locations. This non-linearity causes in fact several problems if one attempts to solve the full problem. For this reason, several methods for dipole estimation work under two important assumptions: first that the number of sources is constant in time, and second that the source locations are fixed. Only recently, Bayesian methods have appeared which try to overcome these artificial constraints. The authors have personally developed one of these methods, namely particle filtering; to this method is devoted Section 6. In the present Section, we give a brief overview of the other main parametric methods.

5.1 Static parametric methods - dipole fitting

Dipole fitting is one of the oldest and most largely used methods for solving the MEG inverse problem. In dipole fitting, current dipoles are estimated one at a time, at fixed time points; the user typically chooses the best time point, the one at which the data exhibit a dipolar field pattern, and then starts a non-linear optimization algorithm - typically a Levenberg-Marquardt algorithm - which maximizes the likelihood for a single current dipole at that particular time point

$$\exp\left(-(\mathbf{b}_t - \mathbf{G}(\mathbf{r}) \cdot \mathbf{q})^T \boldsymbol{\Sigma}^{-1} (\mathbf{b}_t - \mathbf{G}(\mathbf{r}) \cdot \mathbf{q})\right) \quad (34)$$

\mathbf{r} and \mathbf{q} being the location and dipole moment of one single source. The single dipole fitting procedure is typically iterated, usually at different time points, until a satisfactory set of different dipoles is found. To produce an estimate of the temporal behaviour of the estimated sources, a multiple dipole model is constructed by combining the estimated dipole locations and - possibly - orientations; the dipole strengths are then estimated linearly from the data, by optimization of the likelihood

$$\exp\left[-\left(\mathbf{b} - \sum_{i=1}^{\mathcal{N}_{dipoles}} \mathbf{G}(\mathbf{r}^i) \cdot \mathbf{q}^i\right)^T \boldsymbol{\Sigma}^{-1} \left(\mathbf{b} - \sum_{i=1}^{\mathcal{N}_{dipoles}} \mathbf{G}(\mathbf{r}^i) \cdot \mathbf{q}^i\right)\right] \quad (35)$$

A key problem with dipole fitting is that it strongly relies on the personal experience of the person who performs the analysis; in some sense then, despite it never formally makes use of a prior density, it may be considered as a Bayesian method due to the strong prior embodied in the human experience. A second relevant problem of dipole fitting is that, involving decision making from the user, it is usually quite time consuming.

5.2 Static parametric methods - MUSIC and RAP-MUSIC

The Recursively Applied and Projected Multiple Signal Classification (RAP MUSIC) algorithm is perhaps the most largely known method for automatic estimation of current dipoles from MEG data. It was first introduced in Mosher et al. (1992) and further developed in order to deal efficiently with correlated sources (Mosher & Leahy, 1999).

MUSIC is based on a simplified version of the model described in equation (33), as it assumes that the number of dipoles and their locations are fixed in time; thus the model becomes now

$$\mathbf{B}_t = \sum_{i=1}^{\mathcal{N}_{dipoles}} \mathbf{G}(\mathbf{R}^i) \mathbf{U}^i \cdot S_t^i + \mathbf{N}_t = \mathbf{A} \cdot \mathbf{S}_t + \mathbf{N}_t \quad (36)$$

where $\mathbf{A} = [\mathbf{G}(\mathbf{R}^1)\mathbf{U}^1, \dots, \mathbf{G}(\mathbf{R}^{\mathcal{N}_{dipoles}})\mathbf{U}^{\mathcal{N}_{dipoles}}]$ is the field generated by $\mathcal{N}_{dipoles}$ dipoles with locations $\{\mathbf{R}^i\}$ and orientations $\{\mathbf{U}^i\}$; S_t^i is the strength of the i -th dipole and $\mathbf{S}_t = [S_t^1, \dots, S_t^{\mathcal{N}_{dipoles}}]$ the amplitudes vector; the dipole moment is therefore $\mathbf{Q}^i = \mathbf{U}^i \cdot S^i$. If one further assumes that the neural sources are not correlated, simple calculations yield for the data covariance matrix \mathbf{R}_b

$$\mathbf{R}_b = \mathbf{A}^T \mathbf{R}_s \mathbf{A} + \sigma_{noise}^2 \mathbf{I} \quad (37)$$

where \mathbf{R}_s is the covariance matrix of the sources. Due to the assumption of uncorrelated sources, the source covariance matrix has full rank, and the data covariance matrix contains $\mathcal{N}_{dipoles}$ eigenvalues considerably larger than the others; this number represents an estimate of the number of dipoles.

Furthermore the corresponding eigenvectors span the so called signal subspace $\hat{\phi}_s$ and an estimate of the unknown parameters $\{\mathbf{r}^i, \mathbf{u}^i\}$ is obtained by maximizing the subspace correlation (Golub & Loan, 1984) between the dipole field and the signal subspace:

$$subcorr(\mathbf{G}(\mathbf{r})\mathbf{u}, \hat{\phi}_s). \quad (38)$$

Once the locations and orientations are estimated, the source amplitudes are found by a least square approach. A recursive procedure deals with the problem of local minima.

From a statistic view point it can be shown that the estimated parameters represent a Maximum Likelihood estimate of the Random Vector in the model (33) (Stoica & Nehorai, 1989).

5.3 Dynamic parametric methods - MCMC of current dipoles

A more recently proposed strategy for solving the parametric MEG problem is that of using Bayesian inference for approximating the posterior density in the parameter space (Jun et al., 2005). In these studies, the authors assume a fixed but unknown number of spatially stationary dipolar sources; the modeling parameters are then the Random Variable of the number of sources $\mathcal{N}_{dipoles}$ and the Random Vectors of the source locations $\mathbf{R} \in \mathbb{R}^{3\mathcal{N}_{dipoles}}$, of the source time courses $\mathbf{S} \in \mathbb{R}^{\mathcal{N}_{dipoles} \times T}$ and of the dipole orientations $\mathbf{U} \in \mathbb{R}^{3\mathcal{N}_{dipoles}}$. The authors seek an approximation of the posterior density through Bayes theorem (5):

$$\begin{aligned} \pi(n_{dipoles}, \mathbf{r}, \mathbf{s}, \mathbf{u} | \mathbf{b}) = & \\ \pi(\mathbf{b} | n_{dipoles}, \mathbf{r}, \mathbf{s}, \mathbf{u}) \pi(\mathbf{u} | \mathbf{r}, n_{dipoles}) \times & \\ \times \pi(\mathbf{s} | n_{dipoles}) \pi(\mathbf{r} | n_{dipoles}) \pi(n_{dipoles}) , & \end{aligned} \quad (39)$$

where the probability distributions for the source parameters are (necessarily) conditioned on the number of sources. Uniform distributions are used for $\mathcal{N}_{dipoles}$, \mathbf{U} and \mathbf{R} , while a Gaussian distribution is assumed for \mathbf{S} . In order to explore the high-dimensional posterior density, they resort to reversible jump Markov Chain Monte Carlo techniques after marginalization of the posterior density with respect to the noise covariance matrix and the source time courses. In Jun et al. (2006), the multi-dipole model is refined by introducing a source-dependent temporal activity window: in other words, each dipole is active only for a finite time interval, and is totally silent for the rest of the sequence, thus avoiding possible cross-talk phenomena. The implementation is again performed with a Reversible Jump Markov Chain Monte Carlo algorithm.

6. Highly Automated Dipole Estimation

In this section we present the formulation of the multi-dipole particle filter we have been developing in collaboration with our colleagues, and which we have published as an online open source code under the name of HADES (Highly Automated Dipole ESTimation). The idea behind the method is that of using the Bayesian filtering equations (43)-(44) in a multi-dipole setting rather than in the imaging framework. In practice, the algorithm tracks in time the probability distribution of the dipole parameters; the main difficulties here are related to the non-linear dependence of the data on the dipole parameters, and on the necessity of estimating the number of dipoles as well. The particle filtering approach has been first proposed in a methodological paper, (Somersalo et al., 2003). In a subsequent set of publications, with our colleagues we have investigated the use of particle filtering for estimation of current dipoles from experimental MEG data (Sorrentino, 2010; Sorrentino et al., 2009), and compared particle filtering with other source estimation methods in Pascarella et al. (2010). Particle filtering has also been used to solve the EEG inverse problem in Antelis & Minguez (2010) and Mohseni et al. (2008); however, in these studies a simpler setting was used where the number of dipoles is fixed a priori.

The main advantages of the approach outlined below are that (i) it is a very general approach to multiple dipole modeling, as it explicitly allows the number of dipoles to be time varying and provides dynamically a posterior probability for the number of sources; (ii) it can be used online in real time, at least in principle, because data are processed in sequence and not as a whole.

In the following, we first describe the mechanism of the most basic particle filter; we then consider the practical modeling of multiple dipole with Random Finite Sets; we outline briefly the use of semi-analytic techniques, and eventually describe the open source software with Graphical User Interface for use of particle filtering.

6.1 Particle filtering

Bayesian filtering relies on Markovian assumptions for the data and the unknown processes: if \mathbf{J}_t and \mathbf{B}_t represent the neural current and the measured data at time t , respectively, these assumptions can be formulated in terms of conditional distributions, and amount to

$$\pi(\mathbf{b}_t | \mathbf{j}_t, \mathbf{j}_{t-1}, \mathbf{j}_{t-2}, \dots, \mathbf{j}_1) = \pi(\mathbf{b}_t | \mathbf{j}_t) \quad (40)$$

$$\pi(\mathbf{j}_t | \mathbf{j}_{t-1}, \mathbf{b}_{t-1}, \dots, \mathbf{b}_1) = \pi(\mathbf{j}_t | \mathbf{j}_{t-1}) \quad (41)$$

$$\pi(\mathbf{j}_t | \mathbf{j}_{t-1}, \dots, \mathbf{j}_1) = \pi(\mathbf{j}_t | \mathbf{j}_{t-1}) . \quad (42)$$

Given these premises, and assuming knowledge of the two conditional distributions at the right hand side of (40)-(42), the posterior distribution at time t can be obtained from the posterior distribution at time $t - 1$ with the two following equations

$$\pi(\mathbf{j}_t | \mathbf{b}_{1:t-1}) = \int \pi(\mathbf{j}_t | \mathbf{j}_{t-1}) \pi(\mathbf{j}_{t-1} | \mathbf{b}_{1:t-1}) d\mathbf{j}_{t-1} \quad (43)$$

$$\pi(\mathbf{j}_t | \mathbf{b}_{1:t}) = \frac{\pi(\mathbf{b}_t | \mathbf{j}_t) \pi(\mathbf{j}_t | \mathbf{b}_{1:t-1})}{\pi(\mathbf{b}_t | \mathbf{b}_{1:t-1})} . \quad (44)$$

in (43) a prior density at time t is obtained by convolution of the posterior density at time $t - 1$ with the transition kernel; in (44), the posterior density at time t is obtained by combining the prior with the likelihood in the standard Bayes theorem. In addition to the two conditional

distributions mentioned above, the density $\pi(\mathbf{j}_0|\mathbf{b}_0) = \pi(\mathbf{j}_0)$ at time $t = 0$ is needed for initialization.

Exact computation of these two equations is possible under rather restrictive assumptions, such as the ones leading to Kalman filtering and described in Section 4. For more general cases, such as the one considered here which involves a non-linear relationship between the data and the unknowns, numerical techniques have to be used to get reasonable approximations of the posterior densities. Particle filters are one of the most used numerical techniques for solving this type of problems. They are basically a sequential Monte Carlo sampling technique. In their simplest form, the SIR (Sampling Importance Resampling) particle filter, they amount to sequential application of an importance sampling - resampling mechanism.

The main steps of the SIR particle filter are:

- 0 [initialization] draw $\mathcal{N}_{particles}$ sample points $\{\tilde{\mathbf{j}}_0^i\}_{i=1,\dots,\mathcal{N}_{particles}}$ distributed according to $\pi(\mathbf{j}_0|\mathbf{b}_0)$; the sampled points represent therefore an approximation of the initializing density:

$$\pi(\mathbf{j}_0) \simeq \sum_{i=1}^{\mathcal{N}_{particles}} \frac{1}{\mathcal{N}_{particles}} \delta(\mathbf{j}_0 - \tilde{\mathbf{j}}_0^i) ; \quad (45)$$

- 1 [Evolution] for $t \geq 0$, let $\{\tilde{\mathbf{j}}_t^i\}_{i=1,\dots,\mathcal{N}_{particles}}$ be a sample distributed according to $\pi(\mathbf{j}_t|\mathbf{b}_{1:t})$; then exploiting the Chapman-Kolmogorov equation we may approximate the next prior probability density function as follows:

$$\begin{aligned} \pi(\mathbf{j}_{t+1}|\mathbf{b}_{1:t}) &= \int \pi(\mathbf{j}_{t+1}|\mathbf{j}_t) \pi(\mathbf{j}_t|\mathbf{b}_{1:t}) d\mathbf{j}_t \simeq \\ &= \frac{1}{\mathcal{N}_{particles}} \sum_{i=1}^{\mathcal{N}_{particles}} \pi(\mathbf{j}_{t+1}|\tilde{\mathbf{j}}_t^i) ; \end{aligned} \quad (46)$$

sample the so obtained prior density, by drawing a single particle \mathbf{j}_{t+1}^i from $\pi(\mathbf{j}_{t+1}|\tilde{\mathbf{j}}_t^i)$ for each i ; we have now an approximation of the prior at time $t + 1$:

$$\pi(\mathbf{j}_{t+1}|\mathbf{b}_{1:t}) \simeq \sum_{i=1}^{\mathcal{N}_{particles}} \frac{1}{\mathcal{N}_{particles}} \delta(\mathbf{j}_{t+1} - \mathbf{j}_{t+1}^i) ; \quad (47)$$

- 2 [Observation] apply the Bayes theorem, i.e. compute the relative weights of the particles:

$$w_{t+1}^i = \pi(\mathbf{b}_{t+1}|\mathbf{j}_{t+1}^i) , \quad (48)$$

and then normalize the weights so that $\sum_{i=1}^{\mathcal{N}_{particles}} w_{t+1}^i = 1$; an approximation to the posterior probability density function is given by

$$\pi(\mathbf{j}_{t+1}|\mathbf{b}_{1:t+1}) \simeq \sum_{i=1}^{\mathcal{N}_{particles}} w_{t+1}^i \delta(\mathbf{j}_{t+1} - \mathbf{j}_{t+1}^i) ; \quad (49)$$

- 3 [Resampling] resample the sample set representing the posterior density: extract a new set of particles $\{\tilde{\mathbf{j}}_{t+1}^i\}_{i=1}^{N_{particles}}$ from the old set $\{\mathbf{j}_{t+1}^i\}_{i=1}^{N_{particles}}$, in such a way that the probability of drawing the i -th particle is w_{t+1}^i (see Arulampalam et al. (2002) for the details); we get

$$\pi(\mathbf{j}_{t+1} | \mathbf{b}_{1:t+1}) \simeq \sum_{i=1}^{N_{particles}} \frac{1}{N_{particles}} \delta(\mathbf{j}_{t+1} - \tilde{\mathbf{j}}_{t+1}^i) ; \quad (50)$$

go to Step 1.

Remarkably, the combination of Step 1 and 2 can be viewed as an importance sampling.

6.2 Filtering of multiple dipoles with Random Finite Sets

In order to deal with a time-varying and unknown number of sources, in Sorrentino et al. (2009) we have proposed the use of Random Finite Sets as a mathematical tool to model the neural current.

Random Finite Sets (Molchanov, 2005) are generalizations of Random Vectors to a non-fixed dimension case. As the name suggests, realizations of a Random Finite Set are sets containing a random number of random objects. The use of Random Finite Sets for multiple target tracking has been proposed in (Mahler, 2003) for the first time.

In the RFS approach to the MEG problem, the neural current at time t is assumed to be a RFS of current dipoles, denoted by $\mathcal{J}_t = \{\mathbf{D}_t, \dots, \mathbf{D}_t^{N_t}\}$, where each \mathbf{D}_t^i is a Random Vector containing the parameters of a single current dipole and the Random Variable N_t represents the number of sources at time t . A realization of \mathcal{J}_t will be denoted as J_t . The model now amounts to:

$$\mathcal{J}_{t+1} = \mathcal{E}(\mathcal{J}_t) \cup \mathcal{B}_{t+1} \quad (51)$$

$$\mathbf{B}_t = \mathbf{B}(\mathcal{J}_t) + W_t \quad (52)$$

In (51), the dipole set at time $t + 1$ is given by the union of the set $\mathcal{E}(\mathcal{J}_t)$, containing the subset of dipoles active at time t which have survived and evolved, and the set of new born dipoles \mathcal{B}_{t+1} . In (52), the Random Finite Set of dipoles \mathcal{J}_t produces the data; since the dimension of the data vector is fixed to the number of sensors in the MEG device, the data is represented as a standard Random Vector.

Despite the formulation in terms of Random Finite Sets, single dipole probabilities can be used to describe most of the models as long as individual dipoles are assumed to behave independently on each other. In the following we describe the main features of the model in terms of single dipole probability density functions.

6.2.1 Initializing density

HADES adopts an uninformative initializing density for the multi-dipole states. Let v_{max} be the maximum number of dipoles in a RFS (which we need to fix for computational reasons); then the prior probability of different models is uniform, $P(n_0) = 1/(v_{max} + 1)$ for $n_0 = 0, \dots, v_{max}$. Dipoles are assumed to be independent and identically distributed with the same density used for the single-dipole case; the probability density function of each dipole $\mathbf{D}_0 \in \mathcal{J}_0$ is

$$\pi(\mathbf{d}_0 | \mathbf{b}_0) \propto \frac{1}{V} \times \exp\left(-\mathbf{q}_0^T \Xi \mathbf{q}_0\right) \quad (53)$$

where we have defined separately the uniform density for the dipole location in the volume V and the zero-mean Gaussian density with covariance matrix Ξ for the dipole moment.

6.2.2 Likelihood function

Due to the linearity of the Biot-Savart equation, the magnetic field produced by a set of sources is given by the linear combination of the magnetic fields produced by the individual dipoles contained in the set. If we denote by $\mathbf{b}(J_t)$ the magnetic measurement produced by the dipole set J_t , under the standard assumption of additive white Gaussian noise, with covariance matrix Σ , the likelihood function is

$$\pi(\mathbf{b}_t|J_t) \propto \exp\left(-[\mathbf{b}_t - \mathbf{b}(J_t)]^T \Sigma^{-1} [\mathbf{b}_t - \mathbf{b}(J_t)]\right) . \quad (54)$$

6.2.3 Transition kernel

Lack of knowledge on the true dynamics of the neural sources is coded in a transition kernel based on the random walk. The transition probabilities between sets with different number of dipoles are uniform. The evolution of a dipole set containing n dipoles at time t is governed by the following rules:

- the probability that the dipole set contains n dipoles at time $t + 1$ is $1/3$; the probability that it contains $n + 1$ dipoles is $1/3$; the probability that it contains $n - 1$ dipoles is $1/3$;
- if the dipole set at time $t + 1$ contains $n - 1$ dipoles, the dipole to be removed is selected randomly from the set;
- if the dipole set at time $t + 1$ contains $n + 1$ dipoles, the new dipole is drawn from the same distribution of the initializing prior, eq. (53);
- all the surviving dipoles of a dipole set evolve independently according to the single-dipole transition kernel:

$$\pi(\mathbf{d}_{t+1}|\mathbf{d}_t) \propto \mathcal{I}_{B(\mathbf{r}_t)}(\mathbf{r}_{t+1}) \times \exp\left(-[\mathbf{q}_{t+1} - \mathbf{q}_t]^T \Xi_\delta^{-1} [\mathbf{q}_{t+1} - \mathbf{q}_t]\right) . \quad (55)$$

where $\mathcal{I}_{B(\mathbf{r})}(\cdot)$ is the indicator function of a ball with center in \mathbf{r} and Ξ_δ is the covariance matrix for the Gaussian random walk.

6.2.4 Source estimates

Point estimates of the dipole parameters are provided in a three-step automatic procedure. First, the number of sources is estimated from the marginal distribution

$$P(|J_t| = k) = \int_{S(k)} \pi(J_t|\mathbf{b}_{1:t}) \delta J_t , \quad (56)$$

where $S(k)$ is the set of all subsets with k elements; the maximum of the marginal distribution can be used as an estimator:

$$\hat{n}_t = \arg \max_k P(|J_t| = k) . \quad (57)$$

Second, source locations are estimated as the maxima of the marginal quantity often referred to as Probability Hypothesis Density (PHD) in the RFS framework. The PHD is defined as

$$PHD(\mathbf{d}_t) = \int \pi(\{\mathbf{d}_t\} \cup Y) \delta Y , \quad (58)$$

where the integral is a set integral, $\{\mathbf{d}_t\}$ is a single-dipole set and $\pi(\cdot)$ is the probability density function of the Random Finite Set \mathcal{J}_t . In a particle filter implementation, the PHD can be

approximated by

$$PHD(\mathbf{d}_t) = \sum_{i=1}^{\mathcal{N}_{particles}} w_t^i \left(\sum_{\mathbf{d} \in \mathcal{J}_i^t} \delta(\mathbf{d} - \mathbf{d}_t) \right) . \quad (59)$$

The PHD can also be viewed as the intensity measure of a point process.

Third and last step, once the dipole locations have been estimated the dipole moments are estimated with a linear least squares.

6.3 Rao-Blackwellization

The SIR particle filter described in the previous subsection is very general as it allows approximation of non-Gaussian densities which arise due to the non-linear dependence between the data and the parameters. However, in the MEG case such non-linearity is limited to the parameters describing the location of the sources; on the other hand, the data depend linearly on the dipole moment; this linearity can be in fact exploited using a technique known as Rao-Blackwellization. The main idea is that the posterior density can be split (Campi et al. (2008))

$$\pi(\mathbf{r}_t, \mathbf{q}_t | \mathbf{b}_{1:t}) = \pi(\mathbf{q}_t | \mathbf{r}_t, \mathbf{b}_{1:t}) \pi(\mathbf{r}_t | \mathbf{b}_{1:t}) \quad (60)$$

At the right hand side, the density for \mathbf{q}_t $\pi(\mathbf{q}_t | \mathbf{r}_t, \mathbf{b}_{1:t})$ conditional on \mathbf{r}_t is a Gaussian density, whose parameters can be tracked with a Kalman filter; on the other hand, the locations parameters must be sampled, which leads to an algorithm that realizes a set of Kalman filters for the dipole moment, one for each sampled location.

6.3.1 Clustering

Due to its dynamic nature, the particle filter described above tends to provide slightly moving estimates of static dipoles. In order to produce a compact description of the estimated dipoles, it can be useful to apply a clustering algorithm to the whole set of estimated dipoles $\{\hat{\mathbf{d}}_1^1, \dots, \hat{\mathbf{d}}_1^{\hat{n}_1}, \dots, \hat{\mathbf{d}}_t^1, \dots, \hat{\mathbf{d}}_t^{\hat{n}_t}\}$. In Sorrentino et al. (2009) we have applied an iterative algorithm based on the k -means (Spath, 1980) to obtain a small set of neural sources. The iterations are needed because the number of clusters is unknown, while the k -means works with a fixed number of clusters: hence, we first cluster the dipoles in a large number of clusters $\mathcal{N}_{clusters}$; then apply the Wilcoxon test (Weerahandi, 1995) and check whether all clusters are significantly diverse; if not, we set $\mathcal{N}_{clusters} = \mathcal{N}_{clusters} - 1$ and try again the clustering procedure; if they are all diverse, the algorithm stops. A similar idea, implemented in a more sophisticated strategy, has been proposed in Clark & Bell (2007).

6.4 HADES

The particle filter described above has been implemented in Matlab, and is currently distributed with a Graphical User Interface, named HADES (Highly Automated Dipole ESTimation) (Campi et al. (2011)). The software can be downloaded from <http://hades.dima.unige.it>. The software is open source and is released under the standard GPL license. Through the MNE toolbox (http://www.nmr.mgh.harvard.edu/martinos/userInfo/data/MNE_register/index.php) it supports input/output in NeuroMag *.fif*, FreeSurfer *.w* and MNE *.stc* formats, further than the standard Matlab *.mat* format.

7. Conclusion

Localization of brain activity from MEG recordings is a challenging problem which has been solved in a variety of frameworks and using many different algorithms. In this Chapter we have attempted a short review of some of these methods, the ones which seem to be most used in the MEG community, with the aim of describing them in the unifying framework offered by statistics. Interestingly, when described in such framework some clear connections between many of these methods appear naturally. Indeed, it emerges clearly from this review that the only difference between most of these methods, except for the type of source model, relies in the choice of the prior density. It also emerges that the main difference between more recent and previously developed methods is that the recent methods tend to exploit the dynamic nature of the inverse problem by performing regularization in the whole spatio-temporal domain; in statistical terms, this corresponds to providing a priori information on the temporal behaviour of the sources in addition to the a priori information on the spatial distribution. Among these new dynamic methods, we have given particular emphasis to the description of HADES, the multi-dipole particle filter we have been developing in collaboration with our colleagues. We look forward to see how far this new generation of spatio-temporal solvers can reach, and what the next generation will be.

8. Acknowledgements

AS has been supported by the European Union Seventh Framework Programme (FP7/2007-2013) under the grant agreement PIEF-GA-2009-252440.

9. References

- Antelis, J. & Minguez, J. (2010). DYNAMO: dynamic multi-model source localization method for EEG and/or MEG, *IEEE Engineering in Medicine and Biology Society Conference Proceedings*, pp. 5141–5144.
- Arulampalam, M., Maskell, S., Gordon, N. & Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking, *IEEE Transactions on Signal Processing* 50(2): 174–188.
- Bertero, M. & Boccacci, P. (1998). *Introduction to inverse problems in imaging*, IoP, Institute of Physics Publishing.
- Campi, C., Pascarella, A., Sorrentino, A. & Piana, M. (2008). A rao-blackwellized particle filter for magnetoencephalography, *Inverse Problems* 24: 025023.
- Campi, C., Pascarella, A., Sorrentino, A. & Piana, M. (2011). Highly automated dipole estimation, *Computational Intelligence and Neuroscience* 2011: 982185.
- Cantarella, J., Turck, D. D. & Gluck, H. (2001). The biot - savart operator for application to knot theory, fluid dynamics, and plasma physics, *Journal of Mathematical Physics* 42: 876–905.
- Clark, D. & Bell, J. (2007). Multi-target state estimation and track continuity for the particle PHD filter, *IEEE Transactions on Aerospace and Electronic Systems* 43: 1441–1453.
- Dale, A., Liu, A., Fischl, B., Buckner, R., Belliveau, J., Lewine, J. & Halgren, E. (2000). Dynamic statistical parametric mapping: Combining fmri and meg for high-resolution imaging of cortical activity, *Neuron* 26: 55–67.
- Galka, A., Yamashita, O., Ozaki, T., Biscay, R. & Valdes-Sosa, P. (2004). A solution to the dynamical inverse problem of eeg generation using spatiotemporal kalman filtering, *NeuroImage* 23: 435–453.

- Golub, G. & Loan, C. V. (1984). *Matrix Computation*, The Johns Hopkins University Press.
- Gramfort, A., Papadopoulos, T., Olivi, E. & Clerc, M. (2011). Forward field computation with openmeeg, *Computational Intelligence and Neuroscience* 2011: 923703.
- Hämäläinen, M., Hari, R., Knuutila, J. & Lounasmaa, O. (1993). Magnetoencephalography: theory, instrumentation and applications to non-invasive studies of the working human brain, *Reviews of Modern Physics* 65: 413–498.
- Hämäläinen, M. & Ilmoniemi, R. (1984). Interpreting measured magnetic fields of the brain: estimates of current distributions.
- Jun, S., George, J., Paré-Blagojev, J., Plis, S., Ranken, D., Schmidt, D. & Wood, C. (2005). Spatiotemporal bayesian inference dipole analysis for MEG neuroimaging data, *NeuroImage* 28: 84–98.
- Jun, S., George, J., Plis, S., Ranken, D., Schmidt, D. & Wood, C. (2006). Improving source detection and separation in a spatiotemporal bayesian inference dipole analysis, *Physics in Medicine and Biology* 51: 2395–2414.
- Kress, R., Kuhn, L. & Potthast, R. (2002). Reconstruction of a current distribution from its magnetic field, *Inverse Problems* 18: 1127–1146.
- Lin, F., Witzel, T., Ahlfors, S., Stufflebeam, S., Belliveau, J. & Hamalainen, M. (2006). Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates, *NeuroImage* 31: 160–171.
- Long, C., Purdon, P., Temeranca, S., Desai, N., Hämäläinen, M. & Brown, E. (2006). Large scale kalman filtering solutions to the electrophysiological source localization problem - a meg case study, *Proceedings of the 28th IEEE EMBS Annual International Conference*, Vol. 5, pp. 4532–4535.
- Mahler, R. (2003). Multitarget bayes filtering via first-order multitarget moments, *IEEE Transactions on Aerospace and Electronic Systems* 39: 1152–1178.
- Mohseni, H., Wilding, E. & Sanei, S. (2008). Sequential monte carlo techniques for EEG dipole placing and tracking, *2008 IEEE Sensor array and multichannel signal processing workshop*, pp. 95–98.
- Molchanov, I. (2005). *Theory of Random Sets*, Springer-Verlag.
- Mosher, J., Baillet, S. & Leahy, R. (2003). Equivalence of linear approaches in bioelectromagnetic inverse solutions, *2003 IEEE Workshop on Statistical Signal Processing* pp. 294–297.
- Mosher, J. & Leahy, R. (1999). Source localization using Recursively Applied and Projected (RAP) MUSIC, *IEEE Transactions on Signal Processing* 47: 332–340.
- Mosher, J., Lewis, P. & Leahy, R. (1992). Multiple dipole modeling and localization from spatio-temporal MEG data, *IEEE Transactions on Biomedical Engineering* 39: 541–557.
- Ou, W., Hämäläinen, M. & Golland, P. (2009). A distributed spatio-temporal EEG/MEG inverse solver, *NeuroImage* 44: 932–946.
- Pascarella, A., Sorrentino, A., Campi, C. & Piana, M. (2010). Particle filtering, beamforming and rap-music in the analysis of magnetoencephalography time series: a comparison of algorithms, *Inverse Problems and Imaging* 4: 169–190.
- Pascual-Marqui, R. (1999). Review of methods for solving the eeg inverse problem, *Int. J. Bioelectromagnetism* 1: 75–86.
- Pascual-Marqui, R. (2002). Standardize low resolution electromagnetic tomography (sLORETA: technical details, *Methods and Findings in Experimental and Clinical Pharmacology* 24: 5–12.

- Pascual-Marqui, R., C. Michel & D. Lehman (1994). Low resolution electromagnetic tomography: A new method for localizing electrical activity in the brain, *Int. J. Psychophysiol.* 18: 49–65.
- Pursiainen, S., Sorrentino, A., Campi, C. & Piana, M. (2011). Forward simulation and inverse dipole localization with lowest order Raviart-Thomas elements for electroencephalography, *Inverse Problems* p. in press.
- Sarvas, J. (1987). Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem, *Phys. Med. Biol.* 32: 11–22.
- Sekihara, K., Nagarajan, S., Poeppel, D. & Marantz, A. (2002). Performance of an meg adaptive beamformer technique in the presence of correlated neural activities: effects on signal intensity and time course estimates, *IEEE Transactions on Biomedical Engineering* 49: 1534–1546.
- Somersalo, E., Voutilainen, A. & Kaipio, J. (2003). Non-stationary magnetoencephalography by bayesian filtering of dipole models, *Inverse Problems* 19: 1047–1063.
- Sorrentino, A. (2010). Particle Filters for Magnetoencephalography, *Archives of Computational Methods in Engineering* 17: 213–251.
- Sorrentino, A., Parkkonen, L., Pascarella, A., Campi, C. & Piana, M. (2009). Dynamical MEG source modeling with multi-target bayesian filtering, *Human Brain Mapping* 30: 1911–1921.
- Spath, H. (1980). *Cluster Analysis Algorithms for Data Reduction and Classification of Objects*, Halsted Press.
- Stoica, P. & Nehorai, A. (1989). MUSIC, Maximum Likelihood and Cramer-Rao bound, *IEEE Transactions on acoustics, speech and signal processing* 37: 720–741.
- Uutela, K., Hämäläinen, M. & Somersalo, E. (1999). Visualization of magnetoencephalographic data using minimum current estimates, *NeuroImage* 10: 173–180.
- Van Veen, B., van Drongelen, W., Yuchtman, M. & Suzuki, A. (1997). Localization of brain electrical activity via linearly constrained minimum variance spatial filtering, *IEEE Transactions on Biomedical Engineering* 44: 867–880.
- Weerahandi, S. (1995). *Exact Statistical Methods for Data Analysis*, Springer-Verlag, New York.