

Bayesian Multi–Dipole Modeling of a Single Topography in MEG by Adaptive Sequential Monte Carlo Samplers

Alberto Sorrentino, Gianvittorio Luria and Riccardo Aramini

Dipartimento di Matematica – Università di Genova
Via Dodecaneso, 35 – 16146 Genova (Italia)

E-mail: sorrentino@dimma.unige.it

Abstract. In the present paper, we develop a novel Bayesian approach to the problem of estimating neural currents in the brain from a fixed distribution of magnetic field (called *topography*), measured by magnetoencephalography. Differently from recent studies that describe inversion techniques, such as spatio-temporal regularization/filtering, in which neural dynamics always plays a role, we face here a purely static inverse problem. Neural currents are modelled as an unknown number of current dipoles, whose state space is described in terms of a variable–dimension model. Within the resulting Bayesian framework, we set up a sequential Monte Carlo sampler to explore the posterior distribution. An adaptation technique is employed in order to effectively balance the computational cost and the quality of the sample approximation. Then, both the number and the parameters of the unknown current dipoles are simultaneously estimated. The performance of the method is assessed by means of synthetic data, generated by source configurations containing up to four dipoles. Eventually, we describe the results obtained by analyzing data from a real experiment, involving somatosensory evoked fields, and compare them to those provided by three other methods.

1. Introduction

Magnetoencephalography (MEG) is a non-invasive functional neuroimaging technique [12] that records the weak magnetic fields produced by neural currents in the brain. A modern MEG device, made of a helmet-shaped array of a few hundred of SQUID sensors, has two major advantages: it measures the most direct consequence of the brain electrical activity, and it does so with the outstanding temporal resolution of the order of the millisecond, only matched by its closest akin electroencephalography (EEG). By contrast, the more widely used functional Magnetic Resonance Imaging (fMRI) provides measures of the by-product of complex metabolic mechanisms at a rate of a second. The high temporal resolution is a crucial feature of MEG, since it enables to investigate the neural dynamics in a wide variety of conditions, both normal and pathological, where neural oscillations are thought to play a relevant role. Typical examples are Alzheimer’s disease [29], Parkinson’s disease [31] and epilepsy [34]. Moreover, compared to EEG, MEG is less influenced by the inhomogeneities and anisotropies inside the head, which have a high inter-individual variability and can hardly be taken into account within the models.

From a mathematical viewpoint, the estimation of neural currents from MEG data is known to be an ill-posed inverse problem [21], suffering from non-uniqueness of the solution [7, 10]. Source estimation from MEG data is therefore a challenging, yet worth investigating, problem.

Generally speaking, the above inverse problem may be addressed by two different approaches: the *dipolar* and the *distributed* source model [12]. In the latter framework, neural currents are modelled as continuous vector fields, discretized on a dense mesh, and the data depend linearly on the unknowns. Among the regularization techniques used, L^1 -penalty terms are becoming increasingly popular [5], as they provide sparser estimates that more closely reflect the neurophysiological evidence of the neural generators being rather focal. By contrast, this fact is naturally coded in dipolar source models, whereby neural currents are modelled as the superposition of a small number of point-like sources, called *current dipoles*. However, estimation of dipole parameters from MEG data is more difficult, since the relationship between the source positions and the generated magnetic field is non-linear. Traditionally, “fitting” algorithms are used [20] to produce maximum likelihood estimates of individual dipoles, but these methods need careful initialization and may get stuck in local minima. Recently, spatio-temporal analysis of MEG data has been performed by means of Bayesian Monte Carlo techniques, such as particle filters (PF) [25, 28] and Markov Chain Monte Carlo (MCMC) [15]. Given a time-series of MEG data, these methods approximate the posterior distribution for a time-varying set of current dipoles by exploiting the expected temporal continuity of the underlying generators. They have been shown to perform reasonably well with evoked responses, although PF exhibit relatively high variance at the appearance of a new source, since the importance distribution is too different from the posterior: this typically leads to slightly inaccurate localization of the onset, and a localization error

decreasing quickly in the very first time points, while particles concentrate in the high-probability region [26, 27].

In this paper, we take an apparent step back, and develop a new Bayesian approach to solve the seemingly simpler, static inverse problem of determining the number of dipoles and their parameters from a single spatial distribution of the magnetic field (*topography*). Importantly, such single topography may be obtained not only as a single time point of an MEG time-series, but also, for instance, by picking the magnetic field at a given frequency after Fourier transform, or a single spatial component after Independent Component Analysis (ICA); these data-processing techniques can be used, even in combination [14], for analyzing resting-state [1] and event-related [32] data.

The proposed method extends the applicability of Bayesian multi-dipole models to a variety of conditions where the only inversion methods currently available are linear inversion or the above mentioned dipole fit. It ought to be stressed that dipole estimation from a single topography can not be successfully performed by means of PF, not even the one designed to estimate static dipoles, described in [27]. Indeed, PF applied to a single topography would result in a single importance sampling step, with unsatisfactory results due to the difference between the prior and the posterior distribution. A seemingly reasonable option would be to re-iterate a PF with the same data (i.e., simulating a time-series which is constantly identical to the single topography). However, this would be rather questionable from a Bayesian point of view: from the second iteration onwards the posterior distribution of the previous iteration, convolved with the transition kernel, would be regarded as the prior for the current iteration, without being such in a Bayesian sense. In fact, while the absence of dynamics may breathe simplicity, the problem addressed in this paper and that solved by PF are comparably hard: although the state space is now smaller, the data are also fewer, and the number of sources has to be estimated at once, while in a typical MEG time-series dipoles will mostly appear one at a time.

Within our framework, the state space of current dipoles is described by making use of a variable-dimension model [4, 18]. We choose uninformative prior distributions for all parameters but the number of sources, which is assigned a Poisson distribution; this choice grounds on the consideration that dipole models implicitly assume a small number of sources. On the contrary, the likelihood is typically highly informative and naturally tends to high-dimensional configurations. The resulting posterior is therefore a complex, possibly multi-modal density in a high-dimensional space, and negligible in most of the state space: it is therefore not analytically tractable, nor easy to sample.

To explore a posterior distribution with such a complex structure, we apply a recent class of algorithms, called *Sequential Monte Carlo* (SMC) *samplers* [8]. The basic idea on which they are grounded is that of building a sequence of artificial distributions $\{\pi_i\}_{i=1}^{N_I}$ that smoothly moves from a tractable π_1 to the target $\pi \equiv \pi_{N_I}$. By sampling π_1 and letting the samples evolve gradually to account for the difference between two subsequent distributions, one eventually obtains a collection of points (*particles*) that is actually distributed according to π .

The choice of the sequence is clearly a key-point in the definition of an SMC sampler; in this paper, we describe an adaptive approach, called *Adaptive Sequential Monte Carlo* (ASMC) *sampler*, whereby the sequence is determined on-line, with an empirical measure of the distance between two subsequent distributions. In this way, we optimize the trade-off between the computational speed and the quality of the approximation.

We apply the resulting algorithm to a number of synthetic MEG topographies, constructed to assess its performances under a range of experimental conditions. We also show that the method performs reliably on real data, by making use of single time points selected from the responses to somatosensory stimulation, which is a well understood case.

The paper is organized as follows. In Section 2 we provide the mathematical formulation of the MEG problem and its statistical model. To improve readability, the detailed definition of the state-space is given in Appendix A. Section 3 recalls the basic concepts of generic SMC samplers and describes their specific implementation for MEG, as well as the corresponding adaptation technique. In Section 4 an extensive analysis of a large number of synthetic data sets is performed. Section 5 is concerned with the application of the same algorithm to a real MEG data set. Finally, we propose a discussion of the results and our conclusions in Section 6.

2. The MEG inverse problem

2.1. Mathematical model of MEG

The starting point for a formulation of the MEG problem is a mathematical model of the neural current densities that arise when a large number of neurons in one or more small brain regions undergo a simultaneous discharge process. In mathematical terms, this can be viewed as a set of point sources (i.e., current dipoles), each one representing one of the neuron populations. To implement this model, the brain volume is discretized by N_C cells, labelled by the index c . Inside each cell, one reference point is chosen and identified by its position vector $\mathbf{r}(c)$ with respect to an appropriate reference system. A single current dipole is then allowed to be active at each point of the grid $\{\mathbf{r}(c)\}_{c=1}^{N_C}$: this requires considering, for all c , a current density $\mathbf{p}(c) = \mathbf{q}(c) \delta(\mathbf{r} - \mathbf{r}(c))$, where $\mathbf{q}(c)$ is an applied vector representing the moment of the dipole and the Dirac delta results from a limit process of a current density whose support concentrates into the point $\mathbf{r}(c)$. The moment $\mathbf{q}(c)$, in turn, can be expressed as the product of the unit vector $\mathbf{u}(c)$, which describes the direction of the dipole, and the strength $q(c)$, which can be either a real or a complex number, depending on the kind of topography considered[‡].

Next, in order to formulate the equation linking the unknowns with the data of

[‡] For instance, $q(c)$ is a complex number whenever the topography is obtained by selecting a specific frequency of the Fourier transform of the measured field.

the MEG problem, the forward model of the head has to be taken into account. This involves the computation, for each cell, of the so-called *lead-field matrix* $[G(c)]_k^s$, which is defined as follows [12]. Given a Cartesian orthonormal basis $\{\mathbf{e}_k\}_{k=1}^3$, consider a unit current dipole $\mathbf{e}_k \delta(\mathbf{r} - \mathbf{r}(c))$ along each Cartesian direction: the entry $[G(c)]_k^s$ is then defined as the measure, made by the s -th sensor, of the magnetic field produced by the dipole $\mathbf{e}_k \delta(\mathbf{r} - \mathbf{r}(c))$. As a result, denoting by N_S the total number of sensors, we get a family of N_C matrices, each of them being of dimension $3 \times N_S$, which contains the information about the geometry and conductivity of the head model.

By referring the moment of a generic dipole to the basis $\{\mathbf{e}_k\}_{k=1}^3$, any dipole can be represented as the linear combination

$$\mathbf{p}(c) = \mathbf{q}(c) \delta(\mathbf{r} - \mathbf{r}(c)) = \sum_{k=1}^3 q(c) u^k(c) \mathbf{e}_k \delta(\mathbf{r} - \mathbf{r}(c)). \quad (2.1)$$

Let N_D be the number of dipoles and, for $N_D \geq 1$, let $c(d)$ denote the cell where the d -th dipole is located. Then, by exploiting the linearity of the Biot-Savart equation [12], the measure by the s -th sensor of the magnetic field produced by the N_D dipoles can be expressed as

$$b^s = \sum_{d=1}^{N_D} \sum_{k=1}^3 [G(c(d))]_k^s q(c(d)) u^k(c(d)) + \epsilon^s, \quad s = 1, \dots, N_S, \quad (2.2)$$

being ϵ^s the value of the noise affecting the measurement. Whenever $N_D = 0$, the sum on d in (2.2) is conventionally set to zero.

According to Eq. (2.2), the inverse problem of MEG can be formalized as follows: given the array $\mathbf{b} := (b^1, \dots, b^{N_S})$ of the measurements of the magnetic field, determine the state of the neural currents, namely the number N_D of dipoles, their positions (i.e., the cells $c(d)$ where they are located), the Cartesian components $u^k(c(d))$ of their directions, and their strengths $q(c(d))$.

The analysis made above suggests defining the state space \mathcal{X} of the neural currents as

$$\mathcal{X} := \bigcup_{N_D=0}^{N_D^{\max.}} \{N_D\} \times \mathcal{X}(N_D), \quad (2.3)$$

being $N_D^{\max.} (\ll N_C)$ the maximum number of dipoles allowed by the model and $\mathcal{X}(N_D)$ the state space for a fixed number N_D of dipoles. This is usually called a *variable-dimension model* [4, 18]. A detailed definition of the space $\mathcal{X}(N_D)$ can be found in [Appendix A](#). Here, we merely point out the following two features:

- (i) for the concept of “number of dipoles N_D ” to be well-posed, any ambiguous representation of a single dipole as a vector sum of dipoles applied at the same point should be avoided;
- (ii) as the *order* in any N_D -ple in $\mathcal{X}(N_D)$ has no actual physical meaning, any two N_D -ples differing only by a permutation of their components should be identified.

A point in the space \mathcal{X} shall be denoted by x and consists in a pair (N_D, j_{N_D}) , each j_{N_D} being an equivalence class of N_D -ples $(j^{(1)}, \dots, j^{(N_D)})$. In turn, for $d = 1, \dots, N_D$, each $j^{(d)}$ is of the form $(c^{(d)}, z^{(d)}, \varphi^{(d)}, q^{(d)})$, where $c^{(d)}, q^{(d)}$ are synonymous with $c(d), q(c(d))$ above and $z^{(d)}, \varphi^{(d)}$ are cylindrical coordinates which determine the components $u^k(c(d))$ by means of Eq. (A.1).

2.2. Statistical model

Generally speaking, the inverse problem of estimating neural currents from a single MEG topography is ill-posed because of non-detectable source configurations which make the kernel of the Biot-Savart operator non-trivial. We shall cope with this issue by resorting to a Bayesian approach, whereby the problem is recast in terms of statistical inference and the information content of any involved variable is coded into a probability distribution [24]. In particular, the solution is represented by a posterior density $\pi(x|\mathbf{b})$, combining the available information known before any measurement is made with the likelihood of the model, as expressed by Bayes theorem: $\pi(x|\mathbf{b}) \propto p(x) \mathfrak{L}(\mathbf{b}|x)$.

2.2.1. Prior distribution. The parameters of the prior probability density $p(x)$ are set in order to reflect neurophysiological knowledge. Moreover, as discussed above, no more than one dipole is allowed to be located in each cell. By Eq.(2.3), $p(x)$ can be decomposed into the product

$$p(N_D) p_{N_D}(j_{N_D}) = p(N_D) (N_D!) p_{N_D}(j^{(1)}) \cdots p_{N_D}(j^{(N_D)}),$$

where:

- $p_{N_D}(j_{N_D}) := p(j_{N_D} | N_D)$;
- $p(N_D)$ pertains to the number of sources. Since the use of a dipole model intrinsically entails an exiguous number of neural generators, we choose a Poisson distribution with a small parameter; simulations on synthetic data show that values around 0.3 provide good results;
- the factorial arises from the fact that the probability of the equivalence class j_{N_D} is here written as the sum of the probabilities of its representatives $(j^{(1)}, \dots, j^{(N_D)})$;
- each $p_{N_D}(j^{(d)})$ can itself be decomposed as

$$p_{N_D}(j^{(d)}) = p_{N_D}(c^{(d)} | c^{(1)}, \dots, c^{(d-1)}) \cdot p_{N_D}(z^{(d)}) \cdot p_{N_D}(\varphi^{(d)}) \cdot p_{N_D}(q^{(d)}).$$

As far as the single factors in the above equation are concerned, we make the following assumptions:

- since the cells have approximately the same extension, the prior distribution for the dipole location is uniform with respect to the set of vertices, namely $p_{N_D}(c^{(d)} | c^{(1)}, \dots, c^{(d-1)}) = \frac{1}{N_C - (d-1)}$;

- the prior distributions $p_{N_D}(z^{(d)})$ and $p_{N_D}(\varphi^{(d)})$ for the dipole orientation are uniform probability densities on the intervals $[0, 1]$ and $[0, 2\pi)$ respectively. This amounts to uniformly sampling from the half-sphere [23];
- the prior distribution for the strength of the dipole moment is log-uniform, i.e., $q^{(d)} = \pm 10^{3U} \cdot k$, where the sign has uniform distribution, U is uniformly distributed in the interval $(0, 1)$ and the order of magnitude of the constant k is chosen as 10^{-10} , so that the mean of the absolute value $|q^{(d)}|$ is a physiologically plausible value§.

2.2.2. Likelihood function. The analytical form of the likelihood is strictly related to the probability density function of the noise. According to our model, at each sensor $s = 1, \dots, N_S$, the true value of the magnetic field generated by the neural currents is corrupted by a realization ϵ^s of the random variable E^s , as expressed by Eq.(2.2). We assume each E^s to be Gaussian, with zero mean value and variance σ_{noise}^2 , and not correlated with its counterparts at different sensors.

This entails the likelihood to be the following multivariate Gaussian pdf:

$$\mathfrak{L}(\mathbf{b} | x) = \mathcal{N}_{N_S} \left(\mathbf{b} - \sum_{d=1}^{N_D} \sum_{k=1}^3 [G(c(d))]_k^s q(c(d)) u^k(c(d)); \mathbf{0}, \sigma_{\text{noise}}^2 \mathbf{I}_{N_S \times N_S} \right), \quad (2.4)$$

where $\mathbf{I}_{N_S \times N_S}$ denotes the $N_S \times N_S$ identity matrix.

3. An SMC for MEG source modeling

The posterior density described in the previous section is mostly a complicated function on a high-dimensional space, and is therefore an impractical solution to the problem. In this sense, point estimates should better be computed and thus a numerical approximation of $\pi(x | \mathbf{b})$ is needed.

In the next subsection we recall the key ingredients of SMC samplers, a recently developed class of methods for approximating a complex probability density of interest with a relatively small computational cost. The reader is referred to [8] for a full description of the subject. In Sec. 3.2 and Sec. 3.3 we specify our choices concerning the implementation of the SMC sampler for the MEG inverse problem. An adaptation technique, providing better and faster results, is then described in Sec. 3.4. Finally, in Sec. 3.5 we describe how point estimates of the dipole parameters can be computed from the approximated posterior distribution.

3.1. Sequential Monte Carlo Samplers

In most cases, $\pi(x | \mathbf{b})$ is complex enough to make direct importance sampling (IS) inapplicable. This is because any importance density will necessarily be too different

§ The dipole strength is measured in $[A] \cdot [m]$ (in S.I. units), and its typical values are around 10^{-8} .

from the target distribution, thus yielding a poor representation of the latter (see [19], p. 122).

Then, it becomes natural to resort to a sequential version of IS, whereby the solution is achieved gradually. To this end, a tempering sequence of distributions $\{\pi_i(x): \mathcal{X} \rightarrow \mathbb{R}\}_{1 \leq i \leq N_I}$ is built in such a way that $\pi_1(x) := p(x)$ be the prior distribution and $\pi_{N_I}(x) := \pi(x|\mathbf{b})$ be the posterior distribution. By construction, $\pi_1(x)$ can now be easily approximated by means of a sensible importance density $\eta_1(x)$; in our implementation, we shall make the choice $\eta_1(x) := p(x)$. For any i , the particle approximation $\{X_i^{(p)}\}_{p=1}^{N_P}$ of $\pi_i(x)$ is then moved to the subsequent step by means of a Markov kernel. However, it is typically impossible to compute the resulting importance weights analytically and therefore the standard IS estimate of $\pi_{i+1}(x)$ is not available. The SMC sampler technique overcomes this issue at the price of performing importance sampling in an increasingly larger space $\mathcal{X}^i = \underbrace{\mathcal{X} \times \dots \times \mathcal{X}}_{i \text{ times}}$.

Denoting by $x_{1:i}$ the point $(x_1, \dots, x_i) \in \mathcal{X}^i$, an extended probability distribution $\tilde{\pi}_i(x_{1:i}): \mathcal{X}^i \rightarrow \mathbb{R}$ is introduced for any i , in such a way that its marginal with respect to $x_{1:i-1}$ coincides with $\pi_i(x_i)$. This can be done by introducing artificial backward Markov kernels $\{L_k\}_{1 \leq k \leq N_I}$ such that $\int_{\mathcal{X}} L_k(X_{k+1}, x_k) dx_k = 1$ for any k , and setting

$$\tilde{\pi}_i(x_{1:i}) := \pi_i(x_i) \prod_{k=1}^{i-1} L_k(x_{k+1}, x_k).$$

On the other hand, the importance density $\eta_i(x_{1:i})$ on \mathcal{X}^i is naturally obtained from the initial distribution $\eta_1(x_1)$ by means of forward Markov kernels $\{K_k\}_{1 \leq k \leq N_I}$, i.e.,

$$\eta_i(x_{1:i}) := \eta_1(x_1) \prod_{k=2}^i K_k(x_{k-1}, x_k).$$

The corresponding importance weights are

$$\begin{aligned} w_i(x_{1:i}) &:= \frac{\tilde{\pi}_i(x_{1:i})}{\eta_i(x_{1:i})} = \frac{\pi_i(x_i) \prod_{k=1}^{i-1} L_k(x_{k+1}, x_k)}{\eta_1(x_1) \prod_{k=2}^i K_k(x_{k-1}, x_k)} = \\ &= w_{i-1}(x_{1:i-1}) \frac{\pi_i(x_i)}{\pi_{i-1}(x_{i-1})} \frac{L_{i-1}(x_i, x_{i-1})}{K_i(x_{i-1}, x_i)} = \\ &=: w_{i-1}(x_{1:i-1}) \tilde{w}_i(x_{i-1}, x_i), \end{aligned} \tag{3.1}$$

and therefore just the ‘‘incremental weights’’ \tilde{w}_i are to be evaluated.

The construction above makes it possible to define, by induction, the following algorithm. Introduce the normalized importance weights $W_i^{(p)} := w_i(X_{1:i}^{(p)}) / \sum_{j=1}^{N_P} w_i(X_{1:i}^{(j)})$; then, for $i = 1$, compute a particle approximation $\{X_1^{(p)}, W_1^{(p)}\}$ of $\pi_1(x_1)$. As already noticed, the weight function $w_1(x_1) = \pi_1(x_1) / \eta_1(x_1)$ can be computed exactly. Assuming that at iteration $i - 1$ the particle approximation $\{X_{1:i-1}^{(p)}, W_{i-1}^{(p)}\}$ of $\tilde{\pi}_{i-1}(x_{1:i-1})$ is available, let the path of each particle evolve by means

of the Markov kernel $K_i(x_{i-1}, x_i)$ and compute the incremental weight as in (3.1). The algorithm stops at $i = N_I$.

The variance of the (unnormalized) importance weights tends to increase with i and this may lead to a very poor representation of the target density by the sample set. This phenomenon, called *degeneracy*, is routinely measured by using the so-called Effective Sample Size (ESS), defined as $\text{ESS}(\{W_i^{(p)}\}) := (\sum_{p=1}^{N_P} (W_i^{(p)})^2)^{-1}$. If the degeneracy is too high (for example $\text{ESS}(\{W_i^{(p)}\}) < N_P/2$), it is advisable to perform a resampling step.

In order to put the SMC sampler into practice, three ingredients are to be properly chosen: the sequence $\{\pi_i\}$, the transition kernels $\{K_i\}$ and the artificial backward transition kernels $\{L_i\}$. We shall discuss in the following subsections which choices are more suited to the MEG inverse problem.

3.2. Sequence of distributions

In a Bayesian framework, a natural choice for the sequence of artificial distributions is

$$\pi_i := p \cdot \mathfrak{L}^{f(i)}, \quad \text{with } f(i) := \sum_{k=1}^i \delta_k; \quad \delta_1 := 0; \quad \delta_{N_I} := 1 - \sum_{k=1}^{N_I-1} \delta_k; \quad \delta_k \in (0, 1) \quad \forall k. \quad (3.2)$$

We thus start from the prior distribution p and go towards the posterior by increasing the exponent of the likelihood function \mathfrak{L} with the iterations; intuitively, this choice corresponds to embodying the information content of the data in the probability distributions step by step.

The specific choice of the increments δ_k , appearing in the definition of the function $f(i)$, will be done in Sec. 3.4. Here we just want to emphasize that Eq. (3.2) leads to an important consequence, which holds at least whenever the likelihood is a multivariate Gaussian density with covariance matrix $\sigma^2 \mathbf{I}$. In this case, $\mathfrak{L}^{f(i)}$ is still a multivariate Gaussian density with covariance matrix $\sigma_i^2 \mathbf{I}$, being $\sigma_i = \sigma / \sqrt{f(i)}$. This means that, as i varies, the π_i 's can all be interpreted as posterior distributions for the unknown, differing from one another by the estimate of the noise standard deviation. This sounds particularly interesting in those circumstances where the value of σ is unknown, as the SMC provides for free a comparison between the different estimates corresponding to different values of σ .

Moreover, such interpretation can be viewed in connection with the theory of regularization for inverse problems: the sequence $\{\log \pi_i\}$ closely reminds of the one-parameter family of functionals that underlie several regularization algorithms [24]; σ_i here plays the role of the regularization parameter, that tunes the balance between the penalty term (the prior) and the fit with the data (the likelihood). In this connection, the SMC samplers not only provide an approximation to the full posterior, rather than a single estimate as in regularization algorithms, but also explore a finite subset of the so-called regularization path.

3.3. Transition kernels

Within our variable-dimension model, states evolve by means of transition kernels $\{K_i\}$. We choose each K_i to be a Reversible Jump MCMC kernel [11] of invariant distribution π_i . More specifically, we shall assume

$$K_i(x_{i-1}, x_i) = K_i^{(b/d)}(x_{i-1}, x') \cdot K_i^{(mov)}(x', x_i), \quad (3.3)$$

where:

- $K_i^{(b/d)}$ is a Reversible Jump Metropolis-Hastings (RJMH) kernel that accounts for a possible *birth/death* move. One single birth is attempted with probability $P_{\text{birth}} = 1/3$; the location of the new dipole is uniformly distributed within the set of the (not already occupied) vertices, its orientation is uniformly distributed within $\mathbb{P}^2(\mathbb{R})$ and its strength has a log-uniform distribution. A death move is attempted with probability $P_{\text{death}} = 1/20$: in this case we propose to exclude one dipole, which is uniformly chosen among the existing ones;
- $K_i^{(mov)}$ is the part of the kernel which is in charge of the parameters evolution. We attempt to move each dipole separately and the possible changes in its location, orientation and strength are treated one at a time. This means that $K_i^{(mov)}$ is actually the product of $3 \cdot N_D$ Metropolis-Hastings (MH) kernels. As far as the proposals are concerned, we make the following assumptions. The new dipole location is drawn from a set of *neighbours*, i.e., grid points within a radius of 1cm, with probability proportional to a Gaussian centred at the starting position. The proposed orientation is obtained by means of a zero-mean isotropic perturbation of the original one, implemented in Cartesian coordinates. The new dipole strength is drawn from a Gaussian distribution of mean $q^{(d)}$ and standard deviation $q^{(d)}/6$, so that there is negligible probability of going below $q^{(d)}/2$.

As shown in [8], since K_{i+1} is an MCMC kernel of invariant distribution π_{i+1} , we can define the backward kernel L_i as

$$L_i(x_{i+1}, x_i) := \frac{\pi_{i+1}(x_i) K_{i+1}(x_i, x_{i+1})}{\pi_{i+1}(x_{i+1})}. \quad (3.4)$$

By replacing this into Eq. (3.1), we get

$$\tilde{w}_i(x_{i-1}, x_i) = \pi_i(x_{i-1}) / \pi_{i-1}(x_{i-1}). \quad (3.5)$$

As a consequence, the weights $\{W_i^{(p)}\}$ do not depend on the particles $\{X_i^{(p)}\}$ and can thus be already computed at iteration $i-1$. This opens the possibility of resampling the approximation $\{W_i^{(p)}, X_{i-1}^{(p)}\}$ of $\pi_i(x_{i-1})$ before the $\{X_i^{(p)}\}$ are sampled: heuristically, this enables to resample exactly those particles that are known to have a significant weight in the forthcoming iteration.

3.4. Adaptive SMC

The choice of the increments δ_k in Eq. (3.2) is a key-point for a good performance of the whole algorithm: the sequence $\{\pi_i(x)\}_{1 \leq i \leq N_I}$ should come to a compromise between its smoothness (i.e., π_i none too different from π_{i+1}) and an acceptable computational cost.

We first tried some of the possible choices suggested in [8], such as different kinds of geometric paths. However, we observed that they lacked in homogeneity in the sense described above: the difference between two adjacent distributions was unavoidably too little in some phases of the algorithm while too high in others, thus often compromising the quality of results.

Therefore, in the present work, we choose the function $f(i)$ so that the SMC sampler “adapts” to the data set. A similar technique is described in [9].

The construction grounds on the fact that, as explained above, our particular choice of the transition kernels makes it possible to evaluate the weights $w_{i+1}(x_{1:i+1})$ during iteration i . From Eqs. (3.1), (3.2) and (3.5), we get

$$w_{i+1}(x_{1:i+1}) = w_i(x_{1:i}) \cdot \frac{\pi_{i+1}(x_i)}{\pi_i(x_i)} = w_i(x_{1:i}) \cdot \mathfrak{L}^{\delta_{i+1}}. \quad (3.6)$$

We do not fix the value of δ_{i+1} beforehand but, on the contrary, let it vary within the closed set $[\delta_{\min}, \delta_{\max}]$. At each iteration, we first assume $\delta_{i+1} = \delta_{\max}$ and compute the ratio $\text{ESS}(\{W_{i+1}^{(p)}\}) / \text{ESS}(\{W_i^{(p)}\})$, which is an indicator of how much π_{i+1} will differ from π_i : if the value of such ratio falls into a fixed “reliability interval” $[I_{\min}, I_{\max}]$, we confirm the choice of δ_{i+1} , otherwise we propose a new value by bisection and so on. As a consequence, the total number of iterations N_I is unknown at the beginning of the analysis, with the constraint that $\sum_{k=1}^{N_I} \delta_k = 1$. In our algorithm, we set the values $\delta_{\min} = 10^{-5}$, $\delta_{\max} = 10^{-1}$, $I_{\min} = 0.9$ and $I_{\max} = 0.99$.

A schematic version of the algorithm is given below.

3.5. Point Estimates

From the approximation to the artificial distributions provided by the ASMC sampler, point estimates at iteration i are obtained as follows.

- The estimated number of active sources $(\hat{N}_D)_i$ is the mode of the marginal distribution for the number of dipoles, which can be computed from the approximation to the i -th density as

$$P_i(N_D = k | \mathbf{b}) = \sum_{p=1}^{N_P} W_i^{(p)} \delta(k, (N_D^{(p)})_i), \quad (3.7)$$

being $\delta(\cdot, \cdot)$ the Kronecker delta function and $(N_D^{(p)})_i$ the number of dipoles in the p -th particle at iteration i .

Algorithm 1 Adaptive SMC algorithm

Initialization: sample the prior

for $p = 1, \dots, N_P$ **do**

draw X_1^p from $p(x)$;

end for

Set $i = 1$ and $f(1) = 0$

Main cycle

while $f(i) \leq 1$ **do**

$i \rightarrow i + 1$

Possible resampling step

if $ESS(i) \leq N_P/2$ **then**

apply systematic resampling

end if

MCMC sampling

for $p = 1, \dots, N_P$ **do**

RJMH move: propose birth/death, then accept/reject

for $d = 1, \dots, N_D^p$ **do**

MH move for each dipole parameter: propose new value, then accept/reject

end for

compute tentative weights $\tilde{w}_{i+1}^p = \frac{\pi_{i+1}(X_i^p)}{\pi_i(X_i^p)}$

end for

Normalize weights and compute Effective Sample Size

for $p = 1, \dots, N_P$, **do** $w_i^p = \tilde{w}_i^p / W_i$, with $W_i = \sum_p \tilde{w}_i^p$ **end for**

Compute $ESS(i + 1)$

Adaptive determination of next exponent

while $ESS(i + 1)/ESS(i) \geq 0.99$ || $ESS(i + 1)/ESS(i) \leq 0.9$ **do**

increase/decrease δ_{i+1}

re-compute weights w_i^p and $ESS(i + 1)$

end while

Set $f(i + 1) = f(i) + \delta_{i+1}$

end while

- The estimated locations of the active sources are the $(\hat{N}_D)_i$ local modes of the intensity measure for the source locations, conditioned on the estimated number of sources and approximated as

$$P_i(c|\mathbf{b}, \hat{N}_D) = \sum_{p=1}^{N_P} W_i^{(p)} \delta\left((N_D^{(p)})_i, (\hat{N}_D)_i\right) \left(\sum_{d=1}^{(N_D^{(p)})_i} \delta(c, c^{(d)}(X_i^{(p)})) \right), \quad (3.8)$$

being $c^{(d)}(X_i^{(p)})$ the location of the d -th dipole in the p -th particle at iteration i . The fact that only those particles with $(\hat{N}_D)_i$ dipoles contribute to estimating the source locations has been introduced in order to avoid mis-localization effects that may appear with model averaging.

- The direction and the intensity of the estimated dipoles are the mean values of the conditional distribution, conditioned on the source location and on the estimated number of dipoles, approximated as

$$\mathbb{E}_i[\mathbf{q}(c)|\mathbf{b}, \hat{N}_D] = \sum_{p=1}^{N_P} W_i^{(p)} \delta\left((N_D^{(p)})_i, (\hat{N}_D)_i\right) \left(\sum_{d=1}^{(N_D^{(p)})_i} \mathbf{q}(c^{(d)}(X_i^{(p)})) \delta(c, c^{(d)}(X_i^{(p)})) \right). \quad (3.9)$$

4. Simulation experiments

Simulated data are used to validate and assess the performance of the proposed ASMC sampler. Sec. 4.1 describes how data are generated and which measures are used in order to appraise the discrepancy between the estimated and the true dipole configuration. The actual potential of the algorithm is then properly illustrated in Sec. 4.2, by means of an extensive treatment of the analysis of one single topography. Results coming from the entire amount of data sets are eventually summarized in Sec. 4.3.

4.1. Data generation and discrepancy measures

We designed synthetic data with the purpose of investigating the behaviour of the method under a range of different experimental conditions, such as the number of sources, their configurations (i.e., diverse locations, orientations and strengths), and noise levels.

The generation of the synthetic data was realized by using the same source grid and lead-field matrix that are used afterwards by the ASMC sampler. The geometry of the MEG device corresponds to that of a 306 channel Vectorview device, produced by Elekta Neuromag Oy, Helsinki, Finland. The computation of the lead-field matrix $[G(c)]_k^s$ has been carried out with a Boundary Element Method, starting from the geometry of the head of a real subject, extrapolated from MRI images with Freesurfer^{||}. Source points are regularly spaced at 5 mm distance.

^{||} <http://surfer.nmr.mgh.harvard.edu>

Overall we produced 1,200 topographies subdivided into 100 groups, in each of which the number of sources N_D and the noise level ν vary independently, with $N_D = 1, 2, 3, 4$ and $\nu = 0, \nu_{\text{low}}, \nu_{\text{high}}$, for a total of 12 different cases. The topography with N_D dipoles is obtained from that with $N_D - 1$ sources just by adding the signal of the N_D -th source to the pre-existing one.

Dipole locations are uniformly drawn from the (unoccupied) brain grid points. Dipole intensities are set to 7, 10, 5 and 8 nA·m respectively. The orientations are chosen to be those producing the strongest signal at the specified location. This choice is motivated by the will to avoid patently ill-posed conditions. Indeed, it is well-known that any head geometry characterizes a field of directions such that, at each location, all dipoles orientated along it produce a negligible magnetic field compared to that generated by dipoles lying on its perpendicular plane; as an example, radial dipoles produce no magnetic field outside of a perfectly spherical conductor.

In the low-noise condition, the noise standard deviation is set to 5% of the peak of the noise-free signal, corresponding to a Signal-to-Noise Ratio similar to that of evoked responses, while for the high noise level it is twice as large.

In the following, whenever the estimated quantities lack the subscript i , they are to be understood as evaluated at the final iteration N_I .

Now, let $(\hat{N}_D, \hat{j}_{N_D})$ and (N_D, j_{N_D}) be the estimated and the true dipole configuration respectively. In order to quantify the discrepancy between them, the following distances shall be used:

- Δ_{N_D} , which is the difference between the estimated and the true number of sources $\hat{N}_D - N_D$;
- Δ_r , which quantifies the localization error. This is a non-trivial task when the estimated number of dipoles differs from the true one. Here we use a modified version of the OSPA metric [22] with no penalty for cardinality errors, which are evaluated separately by Δ_{N_D} above.

$$\Delta_r := \begin{cases} \min_{\pi \in \Pi_{\hat{N}_D, N_D}} \frac{1}{\hat{N}_D} \sum_{d=1}^{\hat{N}_D} |\mathbf{r}(\hat{c}^{(d)}) - \mathbf{r}(c^{(\pi(d))})| & \text{if } \hat{N}_D \leq N_D, \\ \min_{\pi \in \Pi_{N_D, \hat{N}_D}} \frac{1}{N_D} \sum_{d=1}^{N_D} |\mathbf{r}(\hat{c}^{(\pi(d))}) - \mathbf{r}(c^{(d)})| & \text{if } \hat{N}_D > N_D, \end{cases} \quad (4.1)$$

where $\Pi_{k,l}$ is the set of all permutations of k elements drawn from l elements. When $\hat{N}_D = N_D$, Δ_r is the average distance between all the pairs of estimated and true source locations that is minimal with respect to all possible pairings between true and estimated sources. When $\hat{N}_D \neq N_D$, the farther elements of the more numerous set are simply ignored, and Δ_r is computed as above.

4.2. Illustrative description of the output

In the present section, the analysis of a single synthetic topography is taken as an example so as to illustrate the behaviour of the algorithm. We consider a case where the known underlying source distribution is composed by four dipoles. Figure 1 shows




Figure 1: Behaviour of the SMC at different iterations: (a) the posterior probability $P_i(N_D|\mathbf{b})$ and (b) the adaptively determined value of the exponent of the likelihood.

both the posterior probability $P_i(N_D|\mathbf{b})$ for the number of sources and the (adaptively determined) value of the exponent in the likelihood as functions of the iterations. The posterior probability mass smoothly moves from the zero-dipole scenario, which is preferred during the first iterations as the impact of the data is still small, to models with ever-increasing dimension. Eventually, around iteration 170, the probability of the measured magnetic field being generated by four dipoles is close to one. The exponent of the likelihood grows almost exponentially with the iterations⁺. Figure 2 shows snapshots of the intensity measure for the source locations, at iterations $i = 50, 80, 120, 170, N_I$. Black dots represent the grid $\{\mathbf{r}(c)\}_{c=1}^{N_C}$, while blue points represent those cells in which a dipole is located with intensity measure $P_i(c|\mathbf{b}, \hat{N}_D)$ exceeding 10^{-3} . True dipole locations are denoted by black diamonds. Red crosses identify the estimated ones. The first three snapshots portray the situation at $i = 50$ from the coronal, axial and saggital viewpoint respectively: the algorithm has localized the first dipole, the one producing the strongest signal. Then, at $i = 80$, a second source has also been recognized, while the uncertainty on the position of the first one has decreased. It is worth to observe the peculiar shape of the marginal probability map for the more recently estimated dipole: the left-hand panel shows a wide uncertainty on both the x - and the z -axis, while the middle panel indicates that the uncertainty in the direction of the depth is much larger than in its orthogonal one; this is indeed a well-known result in MEG literature [12]. At subsequent iterations the probability maps tend to concentrate around the true source locations; eventually, in the last iteration (which this time occurred for $i = 249$), the estimated dipoles perfectly coincide with the true ones.

4.3. Results

We analyzed the 1,200 topographies described in Sec. 4.1; the value of σ_{noise} in the likelihood function was set equal to the standard deviation of noise, with a lower threshold of 10^{-14} fT/m, to avoid overly peaked posterior distributions and bound the computational cost. We observed that moderate changes to the value of σ_{noise} do not

⁺ Note the logarithmic scale on the vertical axis.

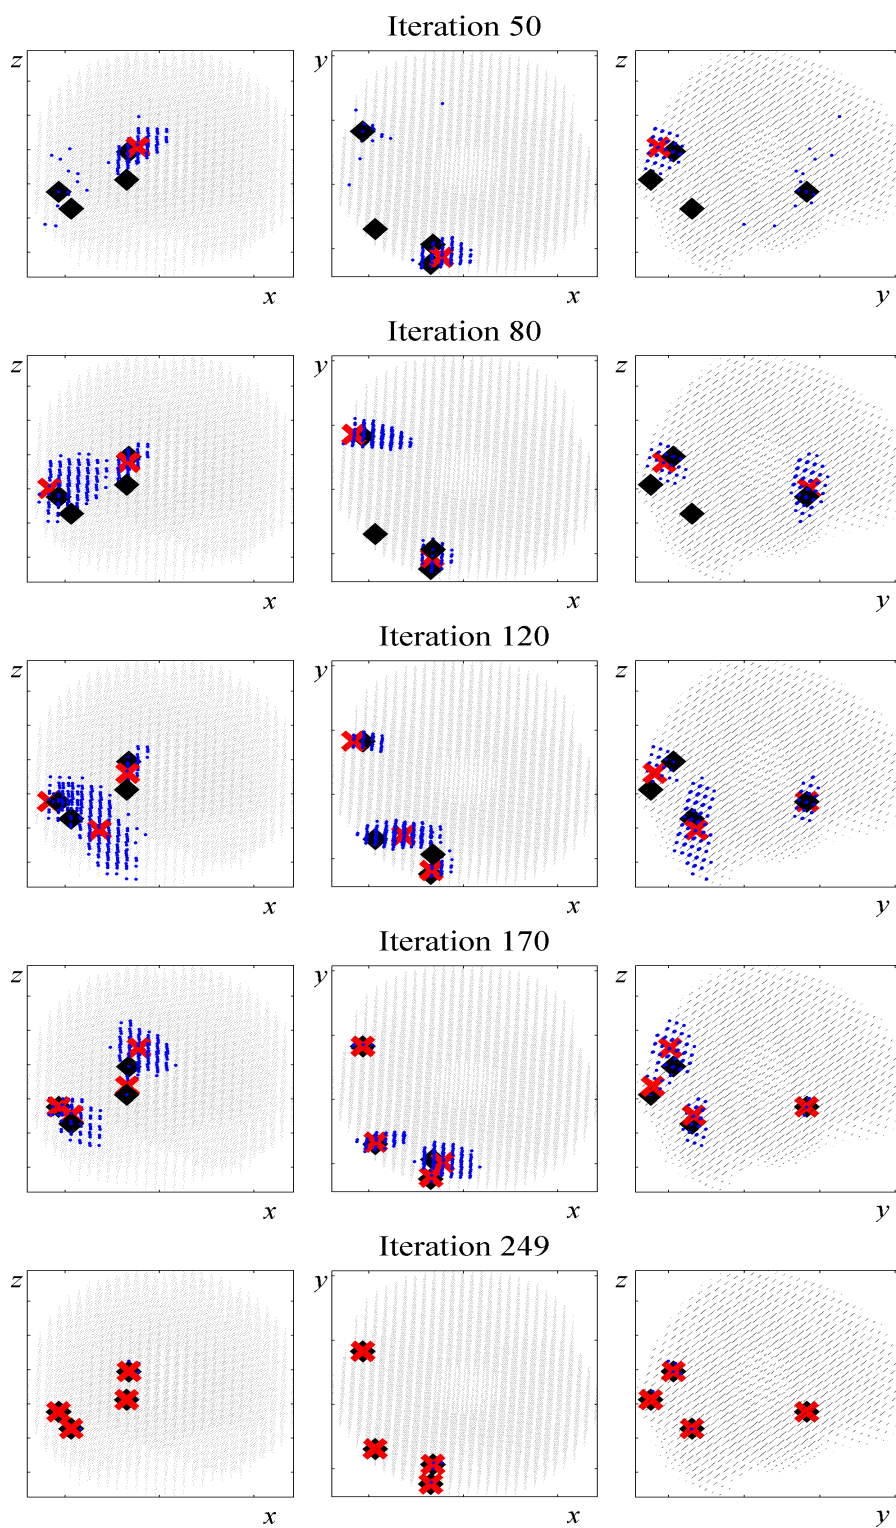


Figure 2: Coronal, axial and saggital views of the marginal probability for the dipole locations, taken at different iterations.

alter the results. In all simulations the number of particles has been set to $N_P = 10,000$, which seemed to provide a reasonable compromise between the quality of the results and the computational cost.

For each topography we computed the point estimate of the dipole configuration and calculated the discrepancy measures Δ_{N_D} and Δ_r described above. Average results along with standard deviations over the 100 realizations are shown in Table 1.

On average, the localization error as measured by Δ_r did not exceed 5 mm with the only exception of the high-noise / four-sources condition. It should be pointed out that the use of the very same grid for both the data generation and the inverse method leads to a slightly underestimated localization error. In the regular cubic grid used in this work, with a 5 mm spacing, such underevaluation amounts to the average distance between a point and the center of the cube the point belongs to, which is about 2.4 mm.

Expectedly, the discrepancy between the estimated and the true configuration tends to increase with increasing noise. This is due to two main effects. First, noise is a disturbance that may cause a small displacement of the peak of the posterior distribution; as a consequence, even in the single-dipole condition the algorithm sometimes fail to localize the source exactly; but such average displacement is extremely small. Second and more important, a higher noise level leads to a decreased resolving power: when two sources are very close to each other, it may become tolerable to replace them by a single source, whose dipole moment is the vector sum of the two. This is indeed reflected in the average Δ_{N_D} becoming negative and decreasing with noise for the three-dipole and four-dipole configurations, where it is more likely that two sources placed at random in the brain are very close to each other.

The impact of the number of sources on the discrepancy is also clearly visible: the localization error increases and Δ_{N_D} becomes negative as the true sources grow in numbers. Even though these average values remain well within the generally accepted localization errors (see for instance [30], where a source is considered as correctly estimated if the localization error is below 20 mm), care must be taken in their interpretation. In fact, many three-dipole and four-dipole configurations are localized very well (see for instance the case shown in Figure 2); in fewer cases, on the other hand, the estimated configuration can get rather far from the true configuration, thus contributing to increasing the average Δ_r as well as its standard deviation. However, this behaviour is not in general due to a poor particle approximation of the posterior distribution, but rather is mostly the consequence of two facts: the ill-posedness of the inverse problem and the limitations of the point estimation procedure.

As far as the former aspect is concerned, the presence of many dipoles may cause the data to be explained equally well by different configurations, which entails the multi-modality of the posterior distribution. This can be directly deduced from a visual inspection, e.g., of the marginal posterior distribution for the source locations, which shows a number of distinct peaks that is larger than the estimated number of dipoles.

Secondly, point estimates are computed as local maxima of the marginal probability (3.8) and, for several reasons, they may not coincide with the true source locations;

ν \backslash N_D	1	2	3	4
0	0.00 ± 0.00 (0.0 \pm 0.0) mm	0.01 ± 0.10 (0.8 \pm 3.1) mm	0.05 ± 0.32 (2.4 \pm 4.4) mm	-0.10 ± 0.46 (4.6 \pm 5.3) mm
ν_{low}	0.00 ± 0.00 (0.0 \pm 0.0) mm	0.01 ± 0.17 (1.4 \pm 4.1) mm	-0.01 ± 0.30 (2.9 \pm 5.3) mm	-0.10 ± 0.48 (5.0 \pm 5.5) mm
ν_{high}	0.00 ± 0.00 (0.2 \pm 1.0) mm	0.14 ± 0.17 (2.7 \pm 5.2) mm	-0.08 ± 0.42 (4.8 \pm 5.2) mm	-0.34 ± 0.57 (6.8 \pm 4.8) mm

Table 1: Discrepancy measures, averaged over 100 runs, for different noise conditions and different number of sources. In each single box, Δ_{N_D} (top) and Δ_r (bottom) are indicated.

however, the true positions (or at least their immediate neighbours) are always assigned a non-negligible probability even when the estimated source configuration is clearly wrong. Moreover, a bad estimation performance in the 3- or 4-dipoles scenario most often still yields a correct reconstruction of some of the sources (typically two).

5. Example on real data

To further test the ASMC sampler, we applied it to real data. Of course, direct validation is no longer possible, since the ground truth is not known exactly. To partially overcome this problem, we adopted the following strategy. Since the somatosensory response has been widely studied in the time domain [16] and is relatively well-understood, we chose our single topographies to be single time points of the evoked response elicited by median nerve stimulation in a healthy subject. Moreover, we compared the results of the ASMC sampler with the source estimate obtained by applying the PF described in [3] to the whole spatio-temporal recordings, and two widely used inverse methods, dynamic Statistical Parametric Mapping (dSPM, [6]) and sLORETA [17].

5.1. Experimental Data

Data from a Somatosensory Evoked Fields (SEFs) mapping experiment were acquired. The recordings were performed at Istituto Neurologico Carlo Besta, Milano (Italy), by means of a 306-channel MEG device (Elekta Neuromag Oy, Helsinki, Finland) comprising 204 planar gradiometers and 102 magnetometers in a helmet-shaped array. Informed consent and prior approval by the local ethics committee were obtained before the recording session. The position of the subject's head within the MEG helmet with respect to anatomical MRIs, obtained by using a 3-Tesla MRI device (General Electric, Milwaukee, USA), was determined by means of a 3D digitizer and four head position indicator coils. The left median nerve at wrist was electrically stimulated at the motor




Figure 3: Butterfly plot of the SEF data

threshold. The MEG signals were recorded at a sampling frequency of 2,000 Hz and then band-pass filtered in the window 0.1 – 100 Hz. Eye movements were monitored by Electrooculogram (EOG) in order to exclude artifacts from the MEG recordings: trials with EOG or MEG exceeding 150 mV or 3 pT/cm, respectively, were neglected and 69 clean trials were averaged. The signal space separation method described in [33] was used to reduce external interference.

A thorough analysis of the neural response to the stimulation of the median nerve has been described in [16], and we take it as a reference here. In that study, the authors modelled the brain activity with multiple current dipoles; the number of dipoles used to fit the response was established according to a subjective criterion, involving monitoring the difference between the predicted and the measured data. Dipole parameters were fitted individually, one at a time, at user-selected time points. While some variability among subjects was observed, the SEF response was shown to obey to the following general scheme: first, the activation of the primary somatosensory cortex, in the hemisphere contralateral to the stimulation, starting about 20 milliseconds after stimulus presentation; then, the activation of parietal sources, starting around 50 milliseconds after the stimulus; finally, activity in frontal areas, after 100 milliseconds.

The final, averaged data used in our study are plotted in Figure 3. The plot appears to confirm that there are three main temporal windows of interest: one between 20 and 35, one between 45 and 75, the last one starting around 110 milliseconds after stimulus.

5.2. Results

We applied the ASMC sampler to MEG topographies taken from the above recordings by selecting specific time points according to the previous analysis; the parameter values in the algorithm were the same as those used for the analysis of synthetic data, with the only exception of the noise standard deviation σ , here estimated from the pre-stimulus

interval. In order to validate the results of the ASMC, we also computed source estimates using three other methods: a PF, that approximates the posterior distribution for the current dipoles conditioned on the data *up to* the selected time point; dSPM, which is based on a distributed source model with an L^2 -prior, and consists in normalizing the Tikhonov regularized solution by the noise standard deviation; and sLORETA, which is similar to dSPM but is claimed to have a smaller localization bias. Figure 4 shows the results at $t = 30, 50$ and 120 ms after the stimulus onset. The results are visualized on a computer representation of the brain obtained by “inflating” the cortical surface: gray levels contain the anatomical information, light gray representing *gyri* and dark gray representing *sulci*; the activity estimate is coded in color scale, increasing from red to yellow. Importantly, this visualization allows activity in the sulci to be clearly visible; on the other hand, since neighbouring volumes may be moved apart by the inflation process, distinct activity regions are often due to underlying volumetric masses that are very close to each other.

Before describing the results, let us comment on the qualitative difference between the images produced by the ASMC sampler and the PF, on the one hand, and those produced by dSPM and sLORETA, on the other. First of all, we point out that all the quantities shown in the images of Fig. 4 are somehow related to the probability of activation at specific locations. Indeed, for both the ASMC and the PF we plot the approximation of the intensity measure (3.8); for any single grid point, this value can be interpreted as the probability of a dipole being at that location, while it integrates, over a given volume Σ , to the mean number of dipoles within Σ . As for dSPM, the represented quantity is a statistical value that is t -distributed under the null-hypothesis of zero activity; as a direct consequence, it also yields a probability of activation, which is however not constrained to be dipolar. Similar considerations apply to sLORETA, although with a different statistical distribution. Importantly, the representation of the results is clearly affected by the setting of the visualization threshold. Owing to the explained differences between the methods, it seems reasonable to use a different value for each method. At the same time, since the plotted quantity is a probability of activation, it seems right to use the same thresholding for different time points. In this connection, the thresholds in Fig. 4 have been chosen by hand following the guidelines just outlined.

Using the same thresholding and parameters at different time points makes the four methods respond differently to the diverse intensities of the different sources. Whenever a stronger source is active, both sLORETA and dSPM will tend to produce widespread estimates, while weaker sources will be represented as small active areas. The behaviour of the ASMC and of the PF is the opposite: a stronger signal will lead to a precise localization of a dipolar source, and then to a focal marginal distribution for the location; a weaker signal will translate to higher uncertainty on the source position, and therefore a more widespread posterior map.

The phenomenon just described is indeed clearly visible in Figure 4. At $t = 25$ ms, all the methods correctly identify the rather strong activation in the contra-lateral

primary somatosensory cortex: the ASMC and the PF provide very focal maps, while dSPM and sLORETA provide compatible widespread estimates; dSPM also exhibits a more posterior peak which does not fit with the commonly agreed models of the response to median nerve stimulation; this may be due to the formerly described brain inflation.

At $t = 50$ ms, the ASMC localizes two sources, one in the right hemisphere and a weaker one in the left hemisphere; these sources appear to be compatible, for timing and location, with the Posterior Parietal Contra and Parietal Opercular Ipsi described in [16]. The PF finds the very same source in the right hemisphere, but it estimates no dipoles in the left one. Such discrepancy between these two methods, which are based on the same source model, is possibly due to two facts: first, thanks to its iterative nature, the ASMC is likely to explore the state space more thoroughly and is therefore more suited to identify weak sources; second, the log-uniform prior on the strength of the dipole moment in the ASMC has a wider range than the corresponding Gaussian prior in the PF. Slightly different results are provided here by dSPM and sLORETA. The activity estimate computed by dSPM seems very similar to that of $t = 25$ ms, having the same spatial distribution with a weaker intensity. The estimate of sLORETA in the right hemisphere includes that of the ASMC and of the PF. Neither dSPM nor sLORETA find significant activity in the left hemisphere.

At $t = 120$ ms, the ASMC localizes again two sources: a stronger one in the right primary somatosensory cortex, and a weaker one in the left posterior frontal area. Both of them seem to be compatible, for timing and location, with the results in [16], the left dipole corresponding to the Frontal Ipsi source. Once again, the other three methods only recover the stronger source.

6. Discussion

In the present paper, we described the application of an SMC sampler for the estimation of brain activity, modelled as multiple current dipoles, from a single spatial distribution of magnetic field in MEG. We implemented an SMC sampler with a sequence of distributions built by exponentiating the likelihood to an increasing value in the interval $[0, 1]$. We made use of an adaptation technique that effectively tunes the speed of the algorithm, by monitoring the ESS at run time.

We applied the resulting ASMC sampler to a set of synthetic data, generated by source configurations containing up to four dipoles, and affected by different levels of noise. We computed discrepancy measures between the true sources and the point estimates provided by the approximated posterior distributions. The estimates are almost perfect for the one- and two-dipole configurations, and expectedly get worse with increasing noise in the three- and four-dipole configurations. In all conditions, the average localization error remains below or slightly above 5 mm, which is the grid spacing in our simulations. In addition, we observed that the fewer cases where results happen to be not completely satisfactory appear to be more a consequence of the multi-modality of the distribution, and of the intrinsic limitations of any point estimation




Figure 4: Source estimates for the Somatosensory Evoked Fields by ASMC, PF, dSPM and sLORETA in the first, second, third and fourth column respectively. The first row represents $t = 25$ ms; the second and third rows portray the output at $t = 50$ ms in the right and left hemisphere respectively; the fourth and fifth rows do the same for $t = 120$ ms.

technique, rather than pointing to an actual failure of the sampling process. Therefore, we suggest that further work is necessary to better investigate multi-modality.

We also tested the developed method against a set of somatosensory data recorded from a healthy subject. We chose to use topographies taken from single time points, so as to compare the ASMC sampler with PF; to further validate our analysis, we also computed source estimates using two well-known inverse methods, dSPM and sLORETA. Results on experimental data suggest that the ASMC can correctly identify neural sources in real scenarios. All four methods provided consistent estimates of the stronger sources; in addition, the ASMC finds tiny sources that are neglected by the other three methods. All the sources localized by the ASMC appear to be in accordance with the results of [16], where the response to median-nerve stimulation is analyzed by means of multi-dipole models.

Further work is clearly necessary to better assess the performance of the proposed method in real scenarios. For instance, it would be interesting to investigate the robustness of the solution with respect to (i) errors in the forward model and (ii) the number of particles used to approximate the posterior distribution. Further, a systematic validation of the method with experimental data, including data in the frequency domain, is being carried on and will be the topic of a future publication.

From a mathematical viewpoint, one may ask whether and how it would be possible to exploit the linearity of the forward equation with respect to the dipole moment, like it was done for the particle filter with Rao-Blackwellization [2]. Preliminary results have been obtained in this direction, and they look encouraging. In addition, the approach described in this paper could clearly be of interest in a number of different fields, where a sparse solution is desired and low-dimensional parametric models are available. Specifically, an application of the same methodology to an inverse problem in solar imaging [13] is currently ongoing.

More realistically, this work opens the possibility of applying Bayesian multi-dipole modeling to MEG data other than in the time domain. We reckon that our ASMC sampler represents a valuable alternative to dipole fitting algorithms, with the considerable advantages of providing an estimate of the number of sources and of not requiring careful initialization. We hope that our efforts can contribute to provide novel, powerful and statistically sound analysis tools to the MEG community.

Appendix A. Multi-dipole state space

The state space \mathcal{X} of the neural currents has been defined in eq. (2.3) as the variable dimension model

$$\mathcal{X} := \bigcup_{N_D=0}^{N_D^{\max.}} \{N_D\} \times \mathcal{X}(N_D).$$

The spaces $\mathcal{X}(N_D)$ are, in turn, set out as follows. First of all, $\mathcal{X}(N_D) := \emptyset$ for $N_D = 0$. Then $\mathcal{X}(N_D) := \mathcal{J}$ for $N_D = 1$, being \mathcal{J} the state space of a single dipole.

On account of Eq. (2.2), \mathcal{J} can be defined as the Cartesian product $\mathbb{N}_C \times \mathbb{P}^2(\mathbb{R}) \times K$, where $\mathbb{N}_C := [1, N_C] \cap \mathbb{N}$ is the space of cell numbers, the real projective plane $\mathbb{P}^2(\mathbb{R})$ is the space of dipole directions and K is the space of dipole strengths, which may be either \mathbb{R} or \mathbb{C} ; in the simulations of Sec. 4 and Sec. 5, $K := \mathbb{R}$. We recall that $\mathbb{P}^2(\mathbb{R})$ can be regarded as the quotient space of the 2-sphere by identifying any two antipodal points. Referring \mathbb{R}^3 to cylindrical coordinates (ρ, φ, z) , the Cartesian components of a point \mathbf{u} in $\mathbb{P}^2(\mathbb{R})$ can thus be expressed as

$$\begin{cases} u^1 = \sin(\arccos z) \cos \varphi \\ u^2 = \sin(\arccos z) \sin \varphi \\ u^3 = z, \end{cases} \quad (\text{A.1})$$

with $z \in [0, 1]$ and $\varphi \in [0, 2\pi)$ if $z \neq 0$, $\varphi \in [0, \pi)$ if $z = 0$. The two coordinates z and φ , in the ranges just specified, are used to express the Cartesian components (A.1) of the dipole direction \mathbf{u} . Summing up, the state of a single dipole is represented as the point $j = (c, z, \varphi, q) \in \mathcal{J}$.

To address the case $N_D > 1$, the first step is to consider the Cartesian product $\mathcal{J}^{N_D} := \mathcal{J} \times \dots \times \mathcal{J}$ of N_D copies of \mathcal{J} . A point of \mathcal{J}^{N_D} consists in a vector $(j^{(1)}, \dots, j^{(N_D)})$ where, for $d = 1, \dots, N_D$, each $j^{(d)}$ is the quadruple $(c^{(d)}, z^{(d)}, \varphi^{(d)}, q^{(d)})$. As pointed out in Sec. 2.1, care must be taken for the concept of “number of dipoles N_D ” to be well-posed: for this reason, the N_D dipoles are constrained to be applied at different points. This entails subtracting from \mathcal{J}^{N_D} the set

$$\mathcal{J}_{eq.}^{N_D} := \{(j^{(1)}, \dots, j^{(N_D)}) \in \mathcal{J}^{N_D} \mid \exists i, j \in \{1, \dots, N_D\} : c(i) = c(j)\} \quad (\text{A.2})$$

to obtain the new space $\mathcal{J}_{N_D} := \mathcal{J}^{N_D} \setminus \mathcal{J}_{eq.}^{N_D}$. Moreover, since the ordering in any N_D -ple of \mathcal{J}_{N_D} is physically meaningless, any two N_D -ples differing only by a permutation of their components $j^{(d)}$ are identified. This amounts to setting onto \mathcal{J}_{N_D} the equivalence relation induced by the action of the symmetric group S_{N_D} . The resulting quotient space \mathcal{J}_{N_D}/\sim is eventually the appropriate choice for the space $\mathcal{X}(N_D)$. An equivalence class in \mathcal{J}_{N_D}/\sim is denoted by j_{N_D} . Note that, for $N_D = 0$, $\mathcal{J}_{N_D}/\sim = \emptyset$ and, for $N_D = 1$, $\mathcal{J}_{N_D}/\sim = \mathcal{J}$, which is consistent with the previous settings.

As a result of the above investigation, the state space of neural currents is defined as the following disjoint union:

$$\mathcal{X} := \bigcup_{N_D=0}^{N_D^{\max.}} \{N_D\} \times \mathcal{J}_{N_D}/\sim. \quad (\text{A.3})$$

Acknowledgments

The authors acknowledge Dr. Dunja Duran, Dr. Ferruccio Panzica, Dr. Davide Rossi, Fabio Rotondi and Dr. Elisa Visani (Istituto Neurologico Carlo Besta, Milano, Italy), who provided the experimental data used in Sec. 5 and contributed to the

neurophysiological interpretation of the results. The authors also thank Prof. Michele Piana (Dipartimento di Matematica, Università di Genova, Italy), for useful comments on the manuscript and support to this research. A.S. kindly acknowledges Dr. Adam Johansen (Department of Statistics, University of Warwick, UK), for a gentle introduction and helpful discussions on the paper by Del Moral *et al.* [8] at the very beginning of this work.

References

- [1] M. Brookes, M. Woolrich, H. Luckhoo, D. Price, J.R. Hale, M.C. Stephenson, G.R. Barnes, S.M. Smith, and P.G. Morris. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *PNAS*, 108:16783–16788, 2011.
- [2] C. Campi, A. Pascarella, A. Sorrentino, and M. Piana. A Rao-Blackwellized particle filter for magnetoencephalography. *Inverse Problems*, 24:025023, 2008.
- [3] C. Campi, A. Pascarella, A. Sorrentino, and M. Piana. Highly automated dipole estimation. *Computational Intelligence and Neuroscience*, 2011:982185, 2011.
- [4] O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer, 2005.
- [5] W.T. Chang, S.P. Ahlfors, and F.H. Lin. Sparse current source estimation for meg using loose orientation constraints. *Human Brain Mapping*, 34:2190–2201, 2013.
- [6] A. Dale, A.K. Liu, B.R. Fischl, R.L. Buckner, J.W. Beldineau, J.D. Lewine, and E. Halgren. Dynamic statistical parametric mapping: Combining fmri and meg for high-resolution imaging of cortical activity. *Neuron*, 26:55–67, 2000.
- [7] G. Dassios, A.S. Fokas, and F. Kariotou. On the non-uniqueness of the inverse MEG problem. *Inverse Problems*, 21:L1–L5, 2005.
- [8] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society B*, 68:411–436, 2006.
- [9] P. Del Moral, A. Doucet, and A. Jasra. An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statistics and Computing*, 22:1009–1020, 2012.
- [10] A.S. Fokas, Y. Kurylev, and V. Marinakis. The unique determination of neuronal currents in the brain via magnetoencephalography. *Inverse Problems*, 20:1067–1082, 2004.
- [11] P.J. Green. Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82:711–732, 1995.
- [12] M. Hämäläinen, R. Hari, J. Knuutila, and O.V. Lounasmaa. Magnetoencephalography: theory, instrumentation and applications to non-invasive studies of the working human brain. *Reviews of Modern Physics*, 65:413–498, 1993.
- [13] G.J. Hurford, E.J. Schmahl, R.A. Schwartz, A.J. Conway, M.J. Aschwanden, A. Csillaghy, B.R. Dennis, C. Johns-Krull, S. Krucker, R.P. Lin, J. McTiernan, T.R. Metcalf, J. Sato, and D.M. Smith. The RHESSI imaging concept. *Solar Physics*, 210:61–86, 2002.
- [14] A. Hyvärinen, P. Ramkumar, L. Parkkonen, and R. Hari. Independent component analysis of short-time fourier transforms for spontaneouseeg/meg analysis. *NeuroImage*, 49:257–271, 2010.
- [15] S.C. Jun, J.S. George, J. Paré-Blagoev, S.M. Plis, D.M. Ranken, D.M. Schmidt, and C.C. Wood. Spatiotemporal bayesian inference dipole analysis for MEG neuroimaging data. *NeuroImage*, 28:84–98, 2005.

- [16] F. Mauguiere, I. Merlet, N. Forss, S. Vanni, V. Jousmaki, P. Adeleine, and R. Hari. Activation of a distributed somatosensory cortical network in the human brain. A dipole modelling study of magnetic fields evoked by median nerve stimulation. part I: location and activation timing of SEF sources. *Electroencephalography and Clinical Neurophysiology*, 104:281–289, 1997.
- [17] R.M. Pascual-Marqui. Standardize low resolution electromagnetic tomography (sLORETA: technical details. *Methods and Findings in Experimental and Clinical Pharmacology*, 24:5–12, 2002.
- [18] C.P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, second edition, 2004.
- [19] R.Y. Rubinstein. *Simulation and the Monte Carlo method*. John Wiley & Sons Inc., New York, 1981.
- [20] R. Salmelin. Multi-Dipole Modeling in MEG. In *MEG: An introduction to methods*, pages 124–155. Oxford University Press, 2010.
- [21] J. Sarvas. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys. Med. Biol.*, 32:11–22, 1987.
- [22] D. Schuhmacher, B.T. Vo, and B.N. Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 56:3447–3457, 2008.
- [23] M.-Z. Shao and N.I. Badler. Spherical sampling by archimedes’ theorem. Technical report, Department of Computer and Information Science, University of Pennsylvania, 1996.
- [24] E. Somersalo and J.P. Kaipio. *Statistical and computational inverse problems*. Springer Verlag, 2004.
- [25] E. Somersalo, A. Voutilainen, and J.P. Kaipio. Non-stationary magnetoencephalography by bayesian filtering of dipole models. *Inverse Problems*, 19:1047–1063, 2003.
- [26] A. Sorrentino. Particle Filters for Magnetoencephalography. *Archives of Computational Methods in Engineering*, 17:213–251, 2010.
- [27] A. Sorrentino, A.M. Johansen, J.A.D. Aston, T.E. Nichols, and W.S. Kendall. Dynamic filtering of static dipoles in Magnetoencephalography. *Annals of Applied Statistics*, 7:955–988, 2013.
- [28] A. Sorrentino, L. Parkkonen, A. Pascarella, C. Campi, and M. Piana. Dynamical MEG source modeling with multi-target bayesian filtering. *Human Brain Mapping*, 30:1911–1921, 2009.
- [29] C. J. Stam, W. de Haan, A. Daffertshofer, B. F. Jones, I. Manshanden, A. M. van Cappellen van Walsum, T. Montez, J. P. A. Verbunt, J. C. de Munck, B. W. van Dijk, H. W. Berendse, and P. Scheltens. Graph theoretical analysis of magnetoencephalographic functional connectivity in Alzheimer’s disease. *Brain*, 132:213–224, 2009.
- [30] L. Stenbacka, S. Vanni, K. Uutela, and R. Hari. Comparison of minimum current estimate and dipole modeling in the analysis of simulated activity in the human visual cortices. *NeuroImage*, 16:936–943, 2002.
- [31] D. Stoffers, J.L.W. Bosboom, J.B. Deijen, E. C. Wolters, H.W. Berendse, and C. J. Stam. Slowing of oscillatory brain activity is a stable characteristic of parkinson’s disease without dementia. *Brain*, 130:1847–1860, 2007.
- [32] M.T. Sutherland and A.C. Tang. Reliable detection of bilateral activation in human primary somatosensory cortex by unilateral median nerve stimulation. *NeuroImage*, 33:1042–1054, 2006.

- [33] S. Taulu, M. Kajola, and J. Simola. Suppression of interference and artifacts by the signal space separation method. *Brain Topography*, 4:269–275, 2004.
- [34] T. Uda, N. Tsuyuguchi, E. Okumura, S. Sakamoto, M. Morino, T. Nagata, H. Ikeda, N. Kunihiro, T. Takami, and K. Ohata. sLORETA-qm for interictal MEG epileptic spike analysis: Comparison of location and quantity with equivalent dipole estimation. *Clinical Neurophysiology*, 123:1496–1501, 2012.