Human Brain Mapping

Dynamical MEG source modeling with multi-target Bayesian tracking

Journal:	*Human Brain Mapping*
Manuscript ID:	draft
Wiley - Manuscript type:	Research Article
Date Submitted by the Author:	n/a
Complete List of Authors:	Sorrentino, Alberto; CNR-INFM LAMIA Parkkonen, Lauri; Helsinki University of Technology, Brain Research Unit, Low Temperature Laboratory Pascarella, Annalisa; Università di Verona, Dipartimento di Informatica; CNR-INFM LAMIA Campi, Cristina; Università di Genova, Dipartimento di Matematica; CNR-INFM LAMIA Piana, Michele; Università di Verona, Dipartimento di Informatica; CNR-INFM LAMIA
Keywords:	magnetoencephalography (MEG), random finite sets, Bayesian tracking, particle filter, inverse problem, source localization

Dynamical MEG source modeling with multi-target Bayesian tracking

A. Sorrentino[a,*], L. Parkkonen[b], A. Pascarella [a,c] , C. Campi [a,d] and M. Piana[a,c]

[a]*CNR - INFM LAMIA, Genova, Italy*

[b]*Brain Research Unit, Low Temperature Laboratory, Helsinki University of Technology, Finland*

[c]*Dipartimento di Informatica, Università di Verona, Italy*

[d]*Dipartimento di Matematica, Università di Genova, Italy*

[*]Correspondence to: Alberto Sorrentino, Via Dodecaneso 35, 16146 Genova (GE), Italy. Email: sorrentino@fisica.unige.it

Running title: Dynamical MEG source models

Key words: magnetoencephalography (MEG); random finite sets; particle filter; Bayesian tracking; source localization; inverse problem.

Abstract: We present a Bayesian filtering approach for automatic estimation of dynamical source models from magnetoencephalographic data. We apply multi-target Bayesian tracking and the theory of Random Finite Sets in an algorithm that recovers the life times, locations and strengths of a set of dipolar sources. The reconstructed dipoles are clustered in time and space to associate them with sources. We applied this new method to synthetic data sets and show here that it is able to estimate the source structure more accurately than either traditional multi-dipole modeling or minimum current estimation performed by uninformed human operators. We also show that from real somatosensory evoked fields the method reconstructs a source constellation comparable to that obtained by multi-dipole modeling.

INTRODUCTION

Magnetoencephalography (MEG) non-invasively measures, with excellent time resolution, the weak magnetic fields produced by the currents flowing in active neurons. Reconstructing the underlying neural currents from MEG measurements allows localizing the active brain regions with reasonable spatial accuracy and enables a variety of applications, both clinical and basic research.

MEG source reconstruction is a difficult task and several methods have been applied to this ill-posed inverse problem. The existing methods can be divided into two classes based on the mathematical model employed to describe the neural sources: "imaging" methods assume a continuous current density, discretized with a dense set of current dipoles, and often result in regularization algorithms such as Minimum Norm Estimation (Hämäläinen and Ilmoniemi, 1994), Minimum Current Estimate (MCE) (Uutela et al., 1999), and various types of beamformers (Van Veen et al., 1997; Sekihara et al., 2002); "parametric" methods employ a small set of current dipoles, and estimate their positions and magnitudes by non-linear optimization techniques (Hämäläinen et al., 1993; Mosher and Leahy, 1999; Uutela et al., 1998; Aine et al., 2000).

Both classes suffer from well-known shortcomings; among the "imaging" methods, the L2-norm estimates tend to be too wide-spread, and beamformers suppress temporally correlated sources. Although several automatic algorithms have appeared for the parametric approach, manually-assisted dipole fitting is still the most widely used method due to its simplicity.

Recently, owing to the increase in the available computational power, Bayesian methods have become feasible. They cast the problem in a more general setting and estimate the whole posterior probability density function (pdf) instead of searching for a single optimal solution. Jun et al. (2005) consider the dipole parameter estimation as a Bayesian inference problem and apply Markov Chain Monte Carlo methods for sampling the posterior density. Galka et al. (2004) and

Long et al. (2006), for EEG and MEG, respectively, consider the inverse problem as a dynamical one and apply Kalman filtering to a linear distributed source model.

In this paper, we consider MEG source estimation as a dynamical Bayesian inverse problem, and use a multi-dipole model for the source distribution. Contrary to most multi-dipole methods, we allow the number of sources, as well as their positions and orientations, to vary over time. Due to the non-linearity of the model, a particle filter (Doucet et al., 2000; Arulampalam et al., 2002) is employed to explore the posterior densities.

Random Finite Sets (RFSs) (Matheron, 1975; Molchanov, 2005) provide a suitable mathematical framework for dealing with a time-varying number of sources in a Bayesian setting. RFSs are a generalization of Random Variables (RV), and they do not constrain the unknowns to a dimensionality known *a priori*. In the present study we exploit the theory of RFSs to construct a particle filter which solves the MEG inverse problem in a completely general manner, allowing a genuinely dynamical model comprising multiple sources. The mathematical details of the theory will not be discussed in this paper; see, e.g., Mahler (2003) and Vo et al. (2005) for the background.

We apply the multi-dipole particle filter to both simulated and real data. A rather simple source constellation underlies the first synthetic data set whereas the next two sets reflect complex activation sequences that have been employed in a comparison of traditional dipole modeling and Minimum Current Estimate (Stenbacka et al., 2002). Thus, we can directly compare the performance of the new algorithm with that of these two other methods. We evaluate the performance of the particle filter also by analyzing somatosensory evoked fields recorded in one subject.

METHODS

Bayesian filtering for MEG source localization

In the Bayesian approach to inverse problems (Somersalo and Kaipio, 2004), the unknown and the measurements are modeled as Random Variables and the solution is the entire posterior probability density function of the unknown, obtained by the Bayes theorem. For dynamical inverse problems the sequential application of Bayes theorem, which requires a prior density at each time step, is mediated by the use of the Chapman-Kolmogorov equation; the result is a two-step algorithm known as Bayesian filtering.

In the MEG application, let j_t and b_t be the realizations of the Random Vectors $\mathbf{J}_t$ and $\mathbf{B}_t$, the primary current and the magnetic field at time t, respectively. The posterior density $\pi(j_t \mid b_{1:t})$ at time t, given the prior density $\pi(j_t \mid b_{1:t-1})$, is obtained by the Bayes theorem:

$$\pi(j_t \mid b_{1:t}) = \frac{\pi(b_t \mid j_t)\pi(j_t \mid b_{1:t-1})}{\pi(b_t \mid b_{1:t-1})} \quad . \tag{1}$$

where $\pi(b_t \mid j_t)$ is the likelihood function, determined by the forward model and the noise statistics; the denominator is the normalization constant $\pi(b_t \mid b_{1:t-1}) = \int \pi(b_t \mid j_t)\pi(j_t \mid b_{1:t-1}) dj_t$. The prior density at time $t+1$, given the posterior density at time t, is estimated by the Chapman-Kolmogorov equation:

$$\pi(j_{t+1} \mid b_{1:t}) = \int \pi(j_{t+1} \mid j_t)\pi(j_t \mid b_{1:t}) dj_t \quad . \tag{2}$$

where $\pi(j_{t+1} \mid j_t)$ is the transition kernel of the stochastic process underlying the data. Here, it reflects general assumptions, such as temporal continuity, on the dynamics of brain sources.

Equations (1) and (2) hold independently of the source and forward models; indeed, these equations have been applied to the MEG inverse problem in rather different frameworks: Long et al. (2006) utilized a Kalman filter for a distributed source model; Somersalo et al. (2003),

Sorrentino et al. (2007) and Campi et al. (2008) employed particle filters to obtain the parameters of a set of current dipoles.

Bayesian filtering of random finite sets of dipoles

Our aim is to apply Bayesian filtering for solving the MEG inverse problem in a truly dynamical multi-dipole framework, where neural sources may appear, strengthen, move, weaken and disappear in the course of time. This approach is conceptually different from that of most other source modeling methods which merely represent the data as a collection of sources whose amplitudes vary over time. These methods cannot readily estimate, *e.g.*, the number of sources at a particular time instant.

In general, the primary current at time t is a set of n_t dipoles:

$$J_t = \{d_t^1, ..., d_t^{n_t}\} = (R_t, Q_t); \quad (3)$$

where $d_t^k = (r_t^k, q_t^k)$ is an abbreviated notation for a single dipole at location r_t^k and with dipole moment q_t^k. Similarly,collections of dipole locations and dipole moments are denoted respectively as

$$R_t = \{r_t^1, \ldots, r_t^{n_t}\} \text{ and } Q_t = \{q_t^1, \ldots, q_t^{n_t}\}. \quad (4)$$

To apply Bayesian filtering, we model the primary current as a RFS of dipoles. Like RVs, RFSs have probability density functions. Furthermore, the integral of a set function is well-defined; Equations (1) and (2) of Bayesian filtering extend to RFSs, provided that the integrals involved are interpreted as set integrals (Vo et al., 2005).

We denote J_t the RFS of dipoles at time t by J_t, and a realization of J_t by J_t. The RFS J_t contains a random number $\mathbf{N}_t$ of random dipoles: $\mathsf{J}_t = \{\mathbf{D}_t^1, ..., \mathbf{D}_t^{\mathbf{N}_t}\}$, and $J_t = \{d_t^1, ..., d_t^{n_t}\}$.

The equations of Bayesian filtering now read:

$$\pi(J_t \mid b_{1:t}) = \frac{\pi(b_t \mid J_t)\pi(J_t \mid b_{1:t-1})}{\pi(b_t \mid b_{1:t-1})} \quad (5)$$

and

$$\pi(J_{t+1} \mid b_{1:t}) = \int \pi(J_{t+1} \mid J_t)\pi(J_t \mid b_{1:t})\delta J_t \,. \qquad (6)$$

Although these equations appear similar to Eqs (1) and (2), there are differences: here J_t are realizations of an RFS, $\pi(J_t \mid \cdot)$ are probability densities of an RFS and the integral in Eqs (5) and (6) is a set integral.

One may question the rationale of using RFSs instead of RVs; for example, Somersalo et al. (2003), Sorrentino et al. (2007) and Campi et al. (2008) modeled the multi-dipole problem with RVs. However, the representation of physical source constellations with RVs is not unique: for example, the state formed by the two dipoles d^1 and d^2 has two representative vectors $(d^1, d^2) \in \mathsf{D} \times \mathsf{D}$ and $(d^2, d^1) \in \mathsf{D} \times \mathsf{D}$, where D is the single-dipole space. As a consequence, the posterior density is permutation-invariant and (artificially) multi-modal: it will peak at (d^1, d^2) and at (d^2, d^1). This bi-modality would prevent us from using standard estimators such as the conditional mean (which would give erroneous estimates) or the Maximum A Posteriori (which would not be unique). This permutation symmetry disappears with RFSs, where each constellation has a unique representative set $\{d^1, d^2\} \equiv \{d^2, d^1\}$, since order of the elements is not relevant in a set. Furthermore, as shown later, a suitable estimator is available for RFSs.

Belief measures

The actual application of Bayesian filtering requires knowledge of three probability density functions: the very first prior density to initialize the algorithm, the likelihood function, and the transition kernel of Eqs (5) and (6). In the case of RFSs, the explicit form of the multi-target probability density function is often complicated; instead, it is preferable to use belief measures (Molchanov, 2005). The belief measure of a RFS is fully determined by a finite set of probability densities one defined on the single-dipole space D, another on the double-dipole space $\mathsf{D} \times \mathsf{D}$,

and so on. Since we assume that no more than $n_{\max}$ dipoles can be active simultaneously, $n_{\max}$ probability measures are enough to define the belief measure. Therefore, in our approach, the belief measures for the prior density, the likelihood function, and the transition kernel are computed starting from the more familiar densities over vector spaces. Now we introduce the three pdfs used in the algorithm.

Initial prior density

Since we have no *a priori* knowledge of the number of sources, we give a uniform distribution to the marginal probability $P(|J_1| = k) = 1/(n_{\max}+1),\ k = 0,1,...,n_{\max}$, where $|J_1|$ is the number of dipoles in the set J_1 and $n_{\max}$ is the maximum allowed number of simultaneous sources. Since we assume that simultaneous dipoles are independent, it is sufficient to define the prior density for a single dipole. For the locations r_1 we use a uniform distribution in the brain volume, and for the dipole moments q_1 we use a zero-mean Gaussian distribution $N(0,\sigma_q)$, where the standard deviation σ_q is of the order of magnitude of the expected sources.

Likelihood function

Since the dimension of the measurement vector is fixed, the data sequence can be modeled with Random Vectors $\mathbf{B}_t$. The forward operator, denoted as $F(\cdot)$, depends on the properties of the volume conductor. Additive noise $\mathbf{M}_t$ is included in the model:

$$\mathbf{B}_t = F(\mathrm{J}_t) + \mathbf{M}_t\,. \qquad (7)$$

Assuming a zero-mean Gaussian noise, the likelihood function is $N(b_t - F(J_t), \sigma_{\text{noise}})$ where σ_{noise} is the standard deviation of the noise. Such a noise model is an approximation; the statistical distribution of MEG noise is generally not known.

Transition kernel

The transition kernel in Eq (7) produces a prior density, and it is therefore *a priori* information inserted in the model: using a non-specific model can be interpreted as inserting less-informative priors. Therefore, we adopt a simple model for the dynamics of the sources; at each time point (i) a new dipolar source may appear with probability p_{new} and distribution $\pi_{\text{new}}(d_{t+1})$, (ii) an existing dipole may disappear with probability p_{dis}, or (iii) the existing dipole may survive, with probability $1 - p_{\text{dis}}$, and evolve according to a single-dipole evolution model $\pi(d_{t+1} \mid d_t)$. All these events are assumed to be independent. In terms of RFSs, the model equation is:

$$\mathsf{J}_{t+1} = \mathsf{S}(\mathsf{J}_t) \cup \mathsf{B}_{t+1} \qquad (8)$$

where B_{t+1} is the RFS of the new dipoles, and $\mathsf{S}(\mathsf{J}_t)$ is the RFS of the survived and evolved dipoles.

Assuming that dipoles evolve independently, it can be shown that the probabilities p_{new} and p_{dis} together with the single-dipole evolution model $\pi(d_{t+1} \mid d_t)$ and the density $\pi_{\text{new}}(d_{t+1})$ fully determine the belief measure and therefore the transition kernel in Eq. (6). The densities we use in the applications:

- the density $\pi_{\text{new}}(d_{t+1})$ of the new dipoles $d_{t+1} = (r_{t+1}, q_{t+1})$ is uniform for the location and Gaussian for the dipole moment, i.e. the same as the initial prior density;
- the probability for a dipole to disappear at time $t+1$ is $p_{\text{dis}} = 1/2n_t$, where $n_t = \mid J_t \mid$;
- the single-dipole evolution model is a random walk constrained within the brain volume; the transition kernel $\pi(d_{t+1} \mid d_t)$ for a single dipole is characterized by Gaussian densities $\pi(r_{t+1} \mid r_t) = N(r_{t+1} - r_t, \gamma_r)$ and $\pi(q_{t+1} \mid q_t) = N(q_{t+1} - q_t, \gamma_q)$ where the variances $\gamma_r = 1$ cm and $\gamma_q = 2$ nAm.

Multi-dipole particle filter

Since the MEG forward problem is non-linear with respect to dipole position, Eqs (5) and (6) can be solved only numerically. Here we apply a sequential Monte Carlo technique known as particle filtering (Doucet et al., 2000; Arulampalam et al., 2002), where a sample set ("particles") distributed according to the posterior density is obtained at each time step by discretizing Eqs (5) and (6), and introducing a resampling step which reduces the number of unlikely sample points. The general scheme of the particle filter we apply is as follows:

- Initialization: Draw a sample $\{J_1^i\}_{i=1,\dots,p}$ distributed according to the density $\pi(J_1)$. Each sample point J_1^i, or "particle", is a set of dipoles. Particles are sampled from a RFS distribution, and therefore different particles may contain a different number of dipoles.
- Observation: Apply the Bayes theorem, i.e. compute the forward solution for each particle J_t^i and its normalized weight (likelihood) $w_t^i = (1/k)\pi(b_t \mid J_t^i)$ with $k = \sum_{i=1}^{p} w_t^i$. The set of weighted particles $\{J_t^i, w_t^i\}$ is an approximation of the posterior density at time t;
- Resampling: Randomly select p particles from the set $\{J_t^i\}_{i=1,\dots,p}$ in such a way that the probability of extracting J_t^i is equal to its likelihood w_t^i. The set of uniformly weighted particles $\{\tilde{J}_t^i\}$ is a new approximation of the posterior density at time t;
- Evolution: Let each particle $\tilde{J}_t^i$ evolve according to the transition kernel by drawing a new particle J_{t+1}^i from each $\pi(J_{t+1} \mid \tilde{J}_t^i)$. Each surviving dipole of each particle evolves according to the single-dipole transition kernel described earlier. New dipoles may appear. The set of uniformly weighted particles $\{J_{t+1}^i\}$ is an approximation of the prior density at time $t+1$.

Source estimates

The posterior density of the current dipole set contains all the available information on the source constellation, but it is difficult to visualize as a whole; instead, different estimates can be computed in order to extract the relevant information.

The posterior density carries information on the number of sources: the marginal distribution

$$P(|J_t|=k)=\int_{\mathsf{D}(k)}\pi(J_t\,|\,b_{1:t})\delta J_t\,, \tag{9}$$

where $\mathsf{D}(k)$ is the set of finite subsets with k dipoles, provides a time-varying estimate of the number of active sources

$$\hat{n}_t=\operatorname{argmax}P(|J_t|=k)\,. \tag{10}$$

For estimating the source parameters we use the RFS-analogue of the first moment of a RV: the Probability Hypothesis Density (PHD). Here, PHD is a function in the single-dipole space D, its integral over region R provides an estimate of the number of dipoles in R, and the peaks of PHD can be used as estimates of the active dipoles. In a particle filter

$$PHD(d_t)=\sum_{i=1}^{p}w_t^i\left(\sum_{d\in J_t^i}\delta(d-d_t)\right). \tag{11}$$

The problem of estimating the source parameters is thus transformed to the problem of finding the local maxima of the PHD, which, in our case, is relatively easy since the number of peaks to be found is already given by the model selection (Eq. 10).

To further reduce the computational complexity, we compute $PHD(r_t)$ instead of $PHD(d_t)$. Therefore, the estimate of the dipole positions is

$$\hat{R}_t=\{\hat{r}_t^1,\ldots,\hat{r}_t^{\hat{n}_t}\}=\arg peaks(PHD(r_t)) \tag{12}$$

and the corresponding dipole moments are computed by standard least-squares techniques.

Implementation

Source-point grid

The computational cost of the algorithm is largely determined by the observation step which involves computing the MEG forward solution for each dipole of each particle at each time step. The large number of particles (about 100,000) may render the algorithm impractical even when resorting to the spherical conductor and its analytic forward solution. In order to reduce the computational cost, we constrain particles to a grid enclosing the brain volume, and pre-compute the lead field for each grid point. The grid is a 20 cm cube with 125,000 points in a 4 mm uniform lattice, requiring 375,000 forward solutions (the three orthogonal directions for each point). Employing 100,000 particles without the grid, assuming 2 simultaneous dipoles on average, and a typical 500-sample analysis window results in 100,000,000 forward solutions to compute; the cost thus reduces by a factor of 300, which makes the algorithm usable. On a standard PC (2 GHz CPU, 2 GB RAM), the time for analyzing a single time point with 100,000 constrained particles is about 1 s.

Such a reduction in computational cost could not be obtained by Rao-Blackwellization, which exploits the linear dependence of the measurements and the dipole moment (Campi et al., 2008). In Rao-Blackwellization, a covariance matrix has to be computed by inversion at each time step for each particle. Although statistical efficiency increases so that much fewer particles are needed to reach a given accuracy, the cost of the matrix inversions is still higher than that of the non-linear particle filter we use in this study.

Clustering

The multi-dipole particle filter provides dynamical estimates of the active dipoles. The estimated source locations at time t, $\{\hat{d}_t^1,...,\hat{d}_t^{\hat{n}_t}\}$ may be slightly different from those at time $t-1$, $\{\hat{d}_{t-1}^1,...,\hat{d}_{t-1}^{\hat{n}_{t-1}}\}$, and even the number of sources $\hat{n}_t$ may change. Thus, no particular dipole

is continuously bound to a given neural source and it is not possible to readily provide the amplitude waveform for each source.

To overcome this problem, we search for clusters within the set of all estimated dipoles $\{\hat{d}_1^1,...,\hat{d}_1^{\hat{n}_1},...,\hat{d}_T^1,...,\hat{d}_T^{\hat{n}_T}\}$, each cluster representing a neural source. The clustering is performed in a 6-dimensional space: 3 dimensions for the location, 2 for the orientation, and one for time t when the source dipole was present. This temporal parameter ensures continuity of the source waveforms. We apply k-means (Spath, 1980) clustering algorithm, however, it requires the number of clusters specified *a priori.* Since the total number of neural sources active during the analysis epoch is unknown, we apply the following iterative procedure:

1. overestimate the number of possible sources s;
2. cluster the dipoles in s clusters with k-means;
3. test each pair of clusters by the Wilcoxon test (Weerahandi, 1995) for statistical difference; if two or more clusters are not significantly different, decrement s and go to step 2; otherwise, stop.

RESULTS

We applied the multi-dipole particle filter to synthetic and real MEG recordings. We employed 3 different synthetic data sets: Simulation 1 comprises three spatially distinct but temporally correlated sources; Simulations 2 and 3 present more challenging situations with multiple sources mimicking the activations evoked by a complex visual stimulus.

First we show the estimates obtained by a single run of the particle filter. Then we use Simulation 2 and 3 to investigate the robustness and statistical reliability of the algorithm by performing multiple runs. We also apply the particle filter to somatosensory evoked fields recorded in one subject.

100,000 particles were utilized in all tests.

Simulation 1

Three dipolar sources comprised the first simulation. The 200-sample data set embedded two occipital dipoles with the same, gamma-function-like time course that peaked at the 20th sample, and a temporal dipole peaking at the 50th sample; the distance between the two occipital sources was 7 cm, and the temporal source was 10 cm from the nearest occipital source. A spherical conductor model was employed in the simulation, and Gaussian noise was added for a final SNR of about 10 dB.

The results are shown in Fig. 1. The top panel displays the marginal probability for the number of dipoles: each colored region represents the probability of a specific model. The other three panels show the original and reconstructed waveforms. The localization errors (averaged over the samples where the source is estimated to exist) are 5 and 4 mm for the two occipital sources and 3 mm for the temporal one. Part of the localization error is attributable to the 4-mm spacing of the source-point grid. The algorithm is able to reconstruct the three source waveforms except for the very low-SNR tails.

Simulations 2 and 3

We utilized synthetic data from a previous study (Stenbacka et al., 2002) designed to evaluate and compare traditional multi-dipole modeling to MCE when performed by human operators unaware of the source structure of the data. Re-using the data enabled a direct comparison of the particle filter to these other methods.

Stenbacka and colleagues employed four simulations of increasing complexity. Here we present the results from the two most complex data sets (Simulation 3 and Simulation 4 in Stenbacka et al. (2002), hereafter referred to as Simulation 2 and 3, respectively) with a variable number of temporally overlapping sources within a volume comparable to a lobe of the brain. Both data sets are crafted to approximate hypothetical neural responses to a complex visual stimulus. Table I summarizes the locations, orientations and peak latencies of the ten sources; six of them appeared in both Simulation 2 and 3, and the remaining four only in Simulation 3. The temporal waveforms of the sources are shown in Figures 2 and 3.

Realistically-shaped boundary element model of the brain was applied in the MEG forward calculation. Brain noise from a MEG experiment where the subject was silently resting was added to the simulated responses for a realistic signal-to-noise ratio. For details, see Stenbacka et al. (2002).

We evaluated the reconstructed sources using the same criteria as Stenbacka and colleagues; they regarded a source correctly estimated if its location was within 2 cm from the true source and its peak latency was within half the duration of the true source (Simulation 2), or within 10 ms from the peak of the true source or half of the time interval between the half maxima of the true source (Simulation 3).

The particle filter was able to reconstruct all of the 6 sources in Simulation 2. The average localization error was 9 mm. The estimated time courses (Fig. 2) never overlapped, *i.e.*, at each time point the algorithm recovered just a single dipole since the true overlapping sources were

within few millimeters from each other and thus a single dipole explained the measured field sufficiently well. The differing orientations of the sources enabled the algorithm to cluster them in 6 groups.

From Simulation 3 the particle filter was able to reconstruct 6 of the 10 sources (Fig. 3) with average localization error of 8 mm. The estimated source classified as $V2$ accounted for the activity of $V1$ and $V3$, whereas sources $V4$ and $V5_L$ were missed. The other 5 sources were correctly recovered. The apparently worse performance of the algorithm with respect to Simulation 2 is due to the different temporal behavior of the sources: $V1$ and $V2$ are shorter-lived and $V3$ is much weaker compared with Simulation 2.

Robustness and reliability

The particle filter requires *a priori* information: the initial prior density, the transition kernel and the likelihood function. The most relevant parameter is σ_{noise} in the likelihood function; it tunes the sensitivity of the algorithm. Furthermore, the algorithm relies on random extraction of samples from probability densities and therefore different runs, even with the same parameters, may yield different results. To investigate the statistical reliability and the robustness of the algorithm with respect to the parameter σ_{noise}, we performed 10 runs for each of 5 different values of σ_{noise} using the data of Simulations 2 and 3. Figure 4 shows the results; with both simulations the average number of recovered sources has a peak at $\sigma_{\text{noise}} = 4$; with higher values σ_{noise}, the algorithm considers only stronger sources ; at lower σ_{noise}, also weak sources are recovered but the algorithm is less stable and the localization error increases. Tuning σ_{noise} impacts also the variability across runs, which reaches a minimum around the same value $\sigma_{\text{noise}} = 4$. Finally, the number of false positives decreases quickly for increasing values of σ_{noise}: for too low values the algorithm tries to model also noise. As a result, the ratio between the number of correctly recovered sources and the number of false positives peaks clearly in

Simulation 2 and has a less pronounced but still visible peak in Simulation 3, both at $\sigma_{noise} = 4$; see Fig. 4.

Since σ_{noise} should reflect the noise level, it can be estimated from, e.g., the pre-stimulus baselines. Interestingly, the standard deviation of the baseline period is 4.3 in Simulations 2 and 3, *i.e.*, very close to the optimal $\sigma_{noise} = 4$.

Somatosensory responses

We applied the particle filter to somatosensory evoked fields (SEF) measured in one healthy human. The recordings were performed after informed consent and had a prior approval by the local ethics committee.

The SEFs were acquired with a 306-channel MEG device (Elekta Neuromag Oy, Helsinki, Finland) comprising 204 planar gradiometers and 102 magnetometers in a helmet-shaped array. The left median nerve at wrist was electrically stimulated at the motor threshold with an interstimulus interval randomly varying between 7.0 and 9.0 s. The MEG signals were filtered to 0.1–200 Hz and sampled at 600 Hz. Trials with EOG or MEG exceeding 150 μV or 3 pT/cm, respectively, were excluded and 84 clean trials were averaged. To reduce external interference, signal-space separation method (Taulu et al., 2004) was applied to the average.

A 3D digitizer and four head position indicator coils were employed to determine the position of the subject's head within the MEG helmet with respect to anatomical MRIs obtained with a 3-Tesla MRI device (General Electric Inc., Milwaukee, USA).

The SEFs (see e.g. Hari and Forss, 1999) were modeled with multiple dipoles whose anatomical locations were verified to be plausible: the N20m and P35m responses at 21 and 38 ms, respectively, localized in the assumed hand area of the S1 cortex contralateral to the stimulation, bilateral responses peaking around 90 ms in the S2 cortices, and a response at around 125 ms in the contralateral posterior parietal cortex (PPC). The N20m and P35m sources were

only about 6 mm apart and had antiparallel orientations. To avoid spurious interaction, the N20m dipole was omitted from the multi-dipole model. The goodness-of-fit of the four-dipole model with respect to the data from all 306 channels ranged from 77 to 97% at the response peaks. This model served as the reference for evaluating the models obtained by the particle filter.

The particle filter was able to recover 3 of the 4 sources found by dipole modeling (see Fig. 5) and missed the source in the contralateral S2 cortex (S2c) probably due to its relatively weak field pattern compared with the other sources. The 3 reconstructed sources satisfied the spatial criterion of Simulation 3 (the average location was within 2 cm from the reference location). Figure 6 presents the source amplitude waveforms estimated by both methods. The particle filter clearly separates the activations of the different sources whereas the time courses given by the multi-dipole model suffer from leakage of residual activity and noise thus showing a non-zero dipole moment throughout the analysis period. Interestingly, the particle filter reconstructs two sources, separated in time and by orientation, for the ipsilateral S2 cortex. The two peaks of opposite polarity in the dipole waveform (S2i in Fig. 6) likely reflect these two sources.

DISCUSSION

The particle filter described in this study represents a notable improvement with respect to the previous implementations (Somersalo et al., 2003; Sorrentino et al., 2007). The more sophisticated mathematical framework allowed a coherent description of the problem in terms of time-varying sets of dipoles, and provided a suitable estimator for multi-dipole states. The algorithm was able to localize a time-varying number of dipoles with no prior knowledge on their number. The use of a source-point grid made the algorithm fast enough to be practically usable.

The localization accuracy in time and space was investigated by testing the particle filter against two challenging data sets which were previously analyzed by Stenbacka et al. (2002) with dipole models and MCE. From the first data set (Simulation 2 in this study) the particle filter correctly localized 6 out of 6 sources, with an average localization error of 8 mm, thus outperforming dipole modeling done by uninformed users, who recovered on average 2.4 sources with dipole modeling and 3.2 with MCE. In the second data set (Simulation 3) the particle filter correctly localized 6 out of 10 sources, with an average localization error of 9 mm. Uninformed users obtained similar results, recovering on average 5.4 sources with dipole modeling and 6.0 sources with MCE, with average localization error of 7.2 mm in both cases. Since the particle filter allows the sources to move and turn during the analysis epoch, the clustering step, which exploits both position and orientation, is able to distinguish sources that dipole modeling unavoidably lumps together. This difference likely explains why the particle filter outperformed the other methods in Simulation 2. In Simulation 3, the results by the particle filter are comparable to those obtained by uninformed users with either dipole modeling or MCE; however, the particle filter algorithm operated automatically while the other two methods required subjective decisions on what to consider a true source. . Furthermore, the localization error of the particle filter is partly attributable to the 4-mm spacing of the source-point grid and could thus be improved by employing a denser grid. From somatosensory evoked fields,

previously analyzed by an informed human using multi-dipole modeling, the particle filter was able to reconstruct 3 out of the 4 sources found by the human operator, only missing the weak contralateral S2 response. The particle filter properly localized the sources also in time, avoiding contamination effects and cross-talk between the dipoles.

In this study we tested the particle filter under very general conditions, with as little *a priori* information as possible. However, such information is available and can be readily exploited. For example, the source space could be constrained to the cortical surface as the bulk of MEG responses are assumed to originate in the cortex. In addition, the noise covariance matrix could be estimated from the pre-stimulus intervals to further reduce the effect of non-white noise. The evolution model could also be replaced with a more realistic one. Optimal usage of the available prior information should be addressed in future work.

In conclusion, the algorithm presented here is a step towards an automatic MEG source modeling method which not only estimates source current distributions but provides a discrete set of significantly active sources for each time instant. Traditional multi-dipole modeling and minimum norm approaches require substantial post-processing to reach qualitatively comparable results. The multi-dipole particle filter we described directly provides instantaneous, time-varying estimates of the number of sources and of the dipole parameters.

REFERENCES

Aine, C., Huang, M., Stephen, J., Christner, R., 2000. Multistart Algorithms for MEG Empirical Data Analysis Reliably Characterize Locations and Time Courses of Multiple Sources. NeuroImage 12, 159–172.

Arulampalam, M., Maskell, S., Gordon, N., Clapp, T., 2002. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. IEEE Tr. Sig. Proc. 50 (2), 174–188.

Campi, C., Pascarella, A., Sorrentino, A., Piana, M., 2008. A Rao-Blackwellized particle filter for magnetoencephalography. Inverse Problems 24, 025023.

Doucet, A., Godsill, S., Andrieu, C., 2000. On sequential Monte Carlo sampling methods for Bayesian filtering. Statistics and Computing 10, 197–208.

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.

Hari, R. & Forss, N. 1999. Magnetoencephalography in the study of human somatosensory cortical processing. Philos Trans R Soc Lond B Biol Sci 354, 1145–1154.

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., 1994. Interpreting magnetic fields of the brain: minimum norm estimates. Med. & Biol. Eng. & Comput. 32, 35–42.

Jun, S., George, J., Paré-Blagoev, J., Plis, S., Ranken, D., Schmidt, D., Wood, C., 2005. Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data. NeuroImage 28, 84–98.

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. In: Proceedings of the 28thIEEE 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.

Matheron, G., 1975. Random Sets and Integral Geometry. Wiley.

Molchanov, I., 2005. Theory of Random Sets. Springer-Verlag.

Mosher, J., Leahy, R., 1999. Source Localization Using Recursively Applied and Projected (RAP) MUSIC. IEEE Transactions on Signal Processing 47, 332–340.

Sekihara, K., Nagarajan, S., Poeppel, D., Marantz, A., Miyashita, Y., 2002. Application of an MEG Eigenspace Beamformer to Reconstructing Spatio-Temporal Activities of Neural Sources. Human Brain Mapping 15, 199–215.

Somersalo, E., Kaipio, J., 2004. Statistical and Computational Inverse Problems. Springer Verlag.

Somersalo, E., Voutilainen, A., Kaipio, J., 2003. Non-stationary magnetoencephalography by Bayesian filtering of dipole models. Inverse Problems 19, 1047–1063.

Sorrentino, A., Parkkonen, L., Piana, M., 2007. Particle filters: a new method for re constructing multiple current dipoles from MEG data. In: Proceedings of the International Conference on Biomagnetism (BIOMAG 2006). Vol. 1300. pp. 173–176.

Spath, H., 1980. Cluster Analysis Algorithms for Data Reduction and Classification of Objects. Halsted Press.

Stenbacka, L., Vanni, S., Uutela, K., Hari, R., 2002. Comparison of Minimum Current Estimate and Dipole Modeling in the Analysis of Simulated Activity in the Human Visual Cortices. NeuroImage 16, 936–943.

Taulu, S., Kajola, M., Simola, J., 2004. Suppression of interference and artifacts by the signal space separation method. Brain Topography 4, 269–275.

Uutela, K., Hämäläinen, M., Salmelin, R., 1998. Global optimization in the localization of neuromagnetic sources. IEEE Transactions on Biomedical Engineering 45, 716–722.

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.

Vo, B., Singh, S., Doucet, A., 2005. Sequential Monte Carlo methods for multitarget filtering with random finite sets. IEEE Transactions on Aerospace and Electronic Systems 41, 1224–1245.

Weerahandi, S., 1995. Exact Statistical Methods for Data Analysis. Springer-Verlag, New York.

Acknowledgements

LP was supported by the Academy of Finland (National Programme for Centers of Excellence

2006–2011).

AS, AP, CC and MP were supported by a grant of the Fondazione CaRiVerona.

TABLES

TABLE I. Sources in Simulations 2 and 3

Area	Location / mm (x, y, z)	Orientation (x, y, z)	Sim 2 t_{peak} / ms	Sim 3 t_{peak} / ms
V1	(11.1, –53.4, 49.8)	$(\frac{1}{\sqrt{2}}, 0, \frac{1}{\sqrt{2}})$	70	60
V2	(13.6, –60.2, 55.9)	(1, 0, 0)	90	80
V3	(17.3, –59.4, 59.8)	(0, 0, 1)	110	100
V3 $_a$	(22.3, –54.8, 64.6)	$(-\frac{1}{\sqrt{2}}, 0, -\frac{1}{\sqrt{2}})$	130	120
V4	(23.1, –47.3, 35.8)	$(-\frac{1}{\sqrt{2}}, 0, \frac{1}{\sqrt{2}})$	150	220
V5 $_R$	(43.6, –36.8, 44.4)	$(-\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}, 0)$	170	160
V5 $_L$	(–33.7, –48.5, 48.1)	$(\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}}, 0)$	-	190
POS	(3.0, –40.0, 83.0)	$(0, \frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$	-	200
STS $_L$	(52.0, 0, 48.0)	(0, 0, –1)	-	280
STS $_R$	(–52.0, –4.0, 48.0)	(0, 0, –1)	-	220

V1–V5 refer to the visual cortices, POS to parieto-occpital sulcus, and STS to superior temporal sulcus; Locations and orientations in "head coordinates" (*x* from left to right preauricular point; *y* towards nasion and perpendicular to *x; z* normal to the *xy*-plane).

FIGURE LEGENDS

Figure 1. Simulation 1. The marginal probability for the number of dipoles plotted as a function of time (top). The original and reconstructed source waveforms (other panels).

Figure 2. Simulation 2. Time courses of the true (blue dashed lines) and estimated (red solid lines) sources.

Figure 3. Simulation 3. Time courses of the original (blue dashed lines) and estimated (red solid lines) sources.

Figure 4. Simulations 2 and 3. The average number of correctly estimated sources and the ratio of them and false positives. σ^{*}_{noise} is the pre-stimulus baseline variance.

Figure 5. Source reconstruction of somatosensory evoked fields. Axial (a) and coronal (b) views of the source models obtained by the particle filter (red squares) and dipole modeling (blue circles).

Figure 6. Somatosensory evoked fields: Time courses estimated by the particle filter (red solid lines) and multi-dipole model (blue dashed lines).

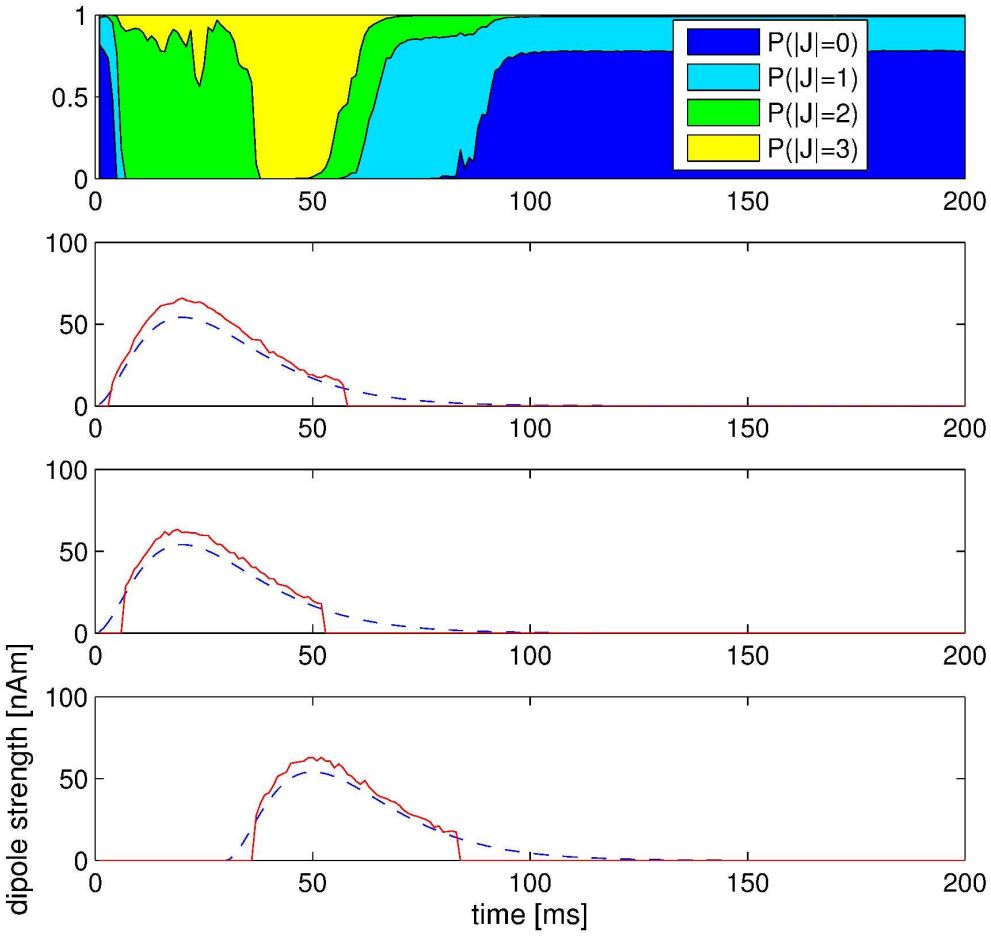

147x144mm (600 x 600 DPI)

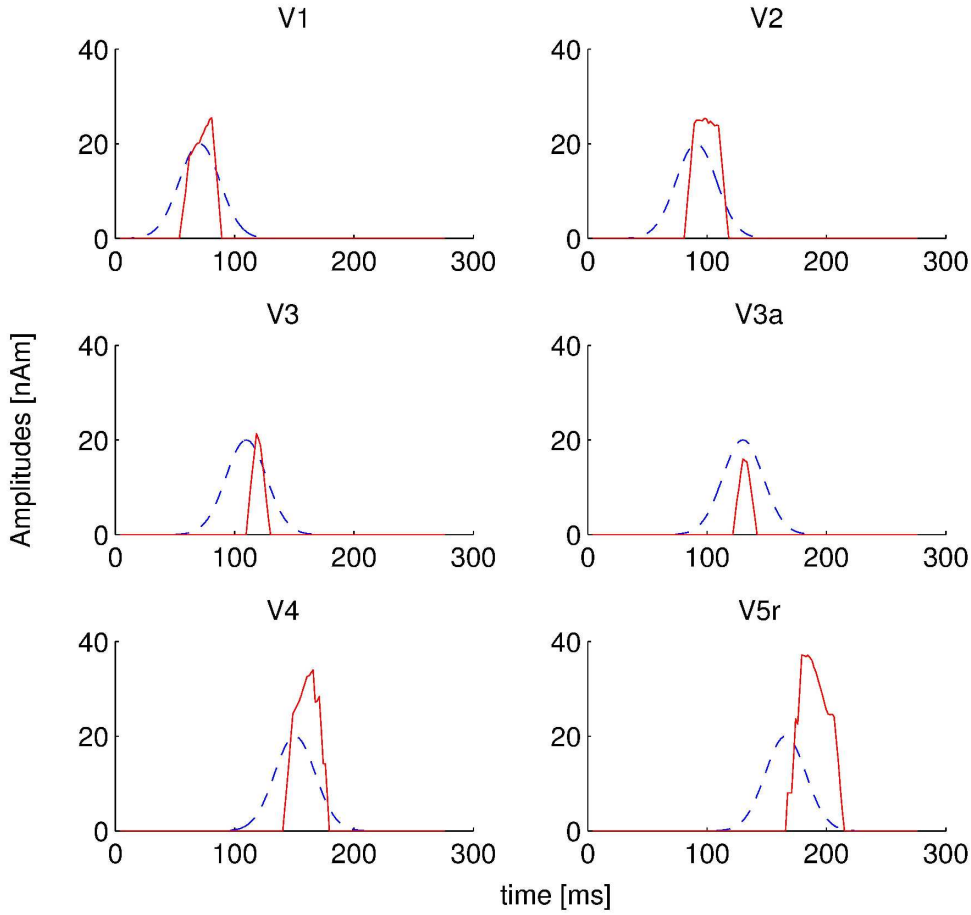

154x150mm (600 x 600 DPI)

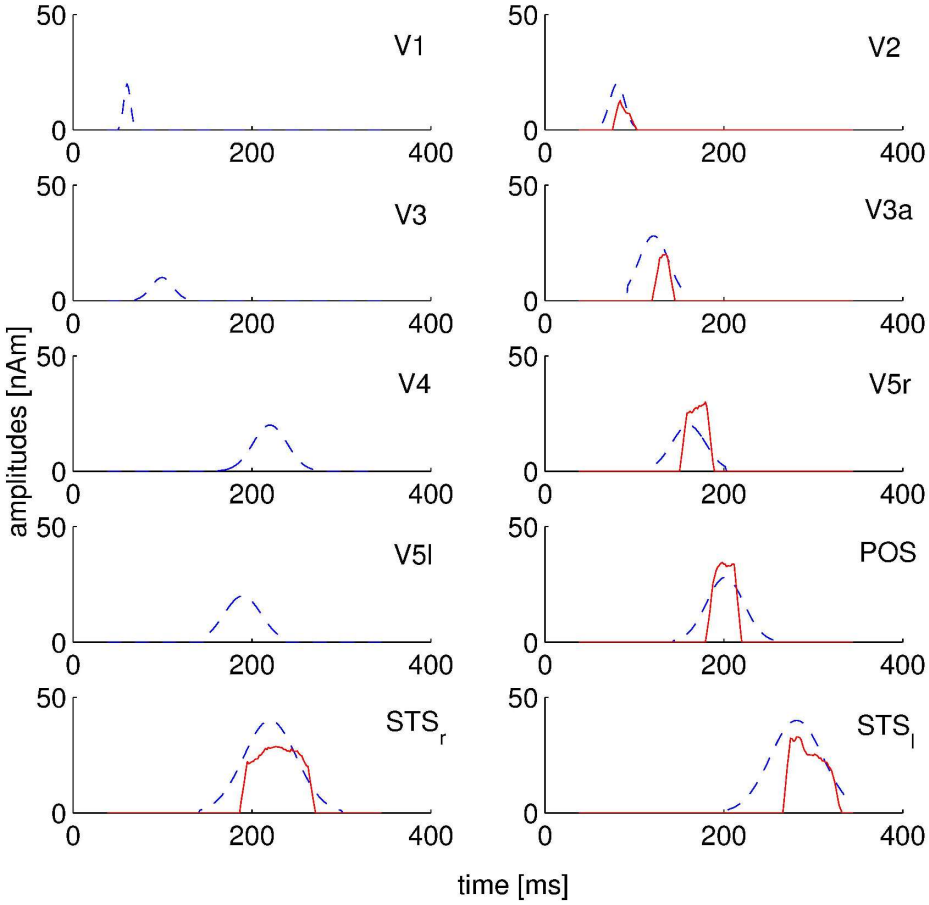

154x150mm (600 x 600 DPI)

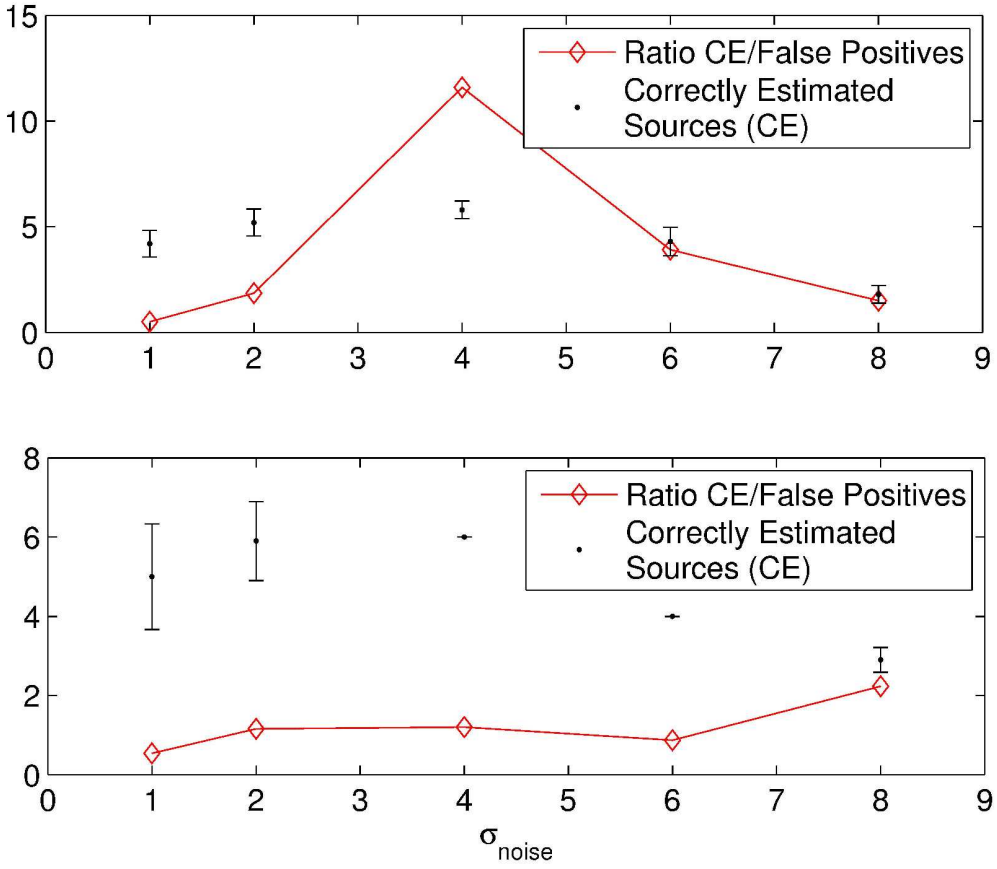

136x120mm (600 x 600 DPI)

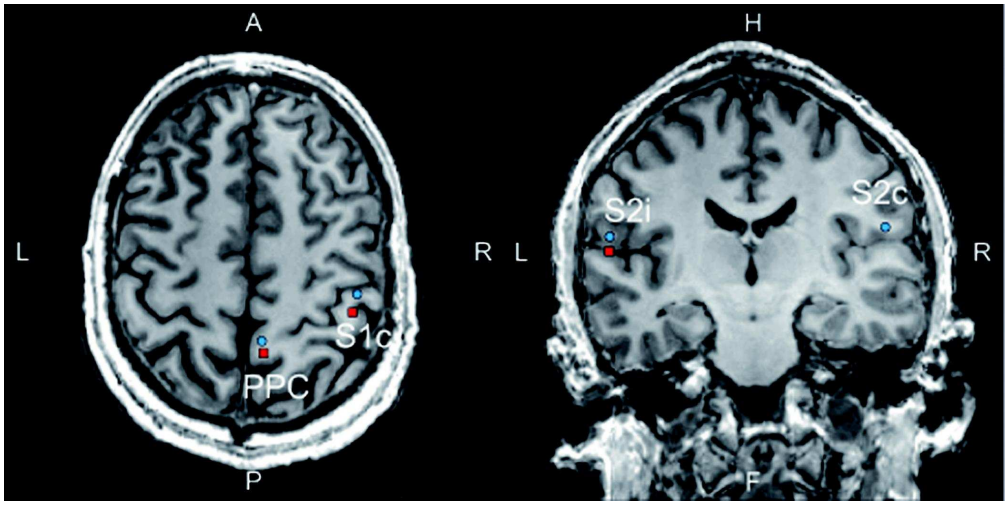

88x44mm (600 x 600 DPI)

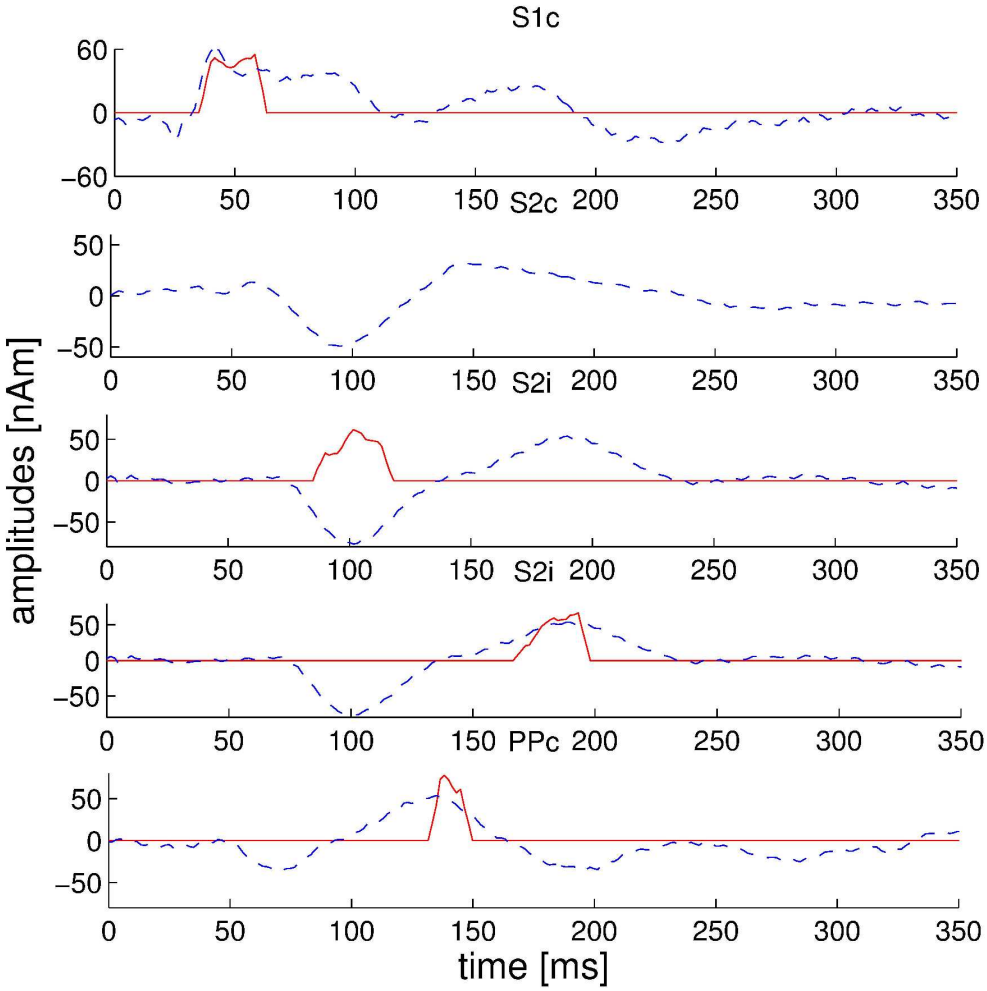

151x151mm (600 x 600 DPI)