Particle filters for the magnetoencephalography inverse problem: increasing the efficiency through a semi-analytic approach (Rao-Blackwellization)

Alberto Sorrentino
Dipartimento di Fisica, Università di Genova, Italy
E-mail: `sorrentino@fisica.unige.it`

Annalisa Pascarella
Dipartimento di Matematica, Università di Genova, Italy

Cristina Campi
Dipartimento di Matematica, Università di Genova, Italy

Michele Piana
Dipartimento di Informatica, Università di Verona, Italy

Abstract. We consider the problem of dynamically estimating the parameters of point-like neural sources from magnetoencephalography data. Since the problem is non-linear, we apply the sequential Monte Carlo algorithms known as particle filters for solving the Bayesian filtering problem. We suggest that the linear dependence of the data on a subset of the parameters allows the analytic computation of the posterior density for these parameters, i.e. Rao-Blackwellization; this considerably improves the accuracy of the method and its statistical efficiency.

1. Introduction

In the magnetoencephalography (MEG) inverse problem the aim is to recover the time-varying neural currents from high-frequency (1000 Hz) measurements of the magnetic field distribution around the subject's head [4]. The mathematical model is accounted for by the Biot-Savart equation, which linearly relates the magnetic field ditribution outside the head to the electrical current distribution inside the brain volume. Standard regularization techniques, both with L^2 [5] and L^1 [10] constraints, have been applied for solving this linear inverse problem, however, the reconstructed current densities often appear to be unrealistically widespread (with respect to the known neurophysiological contraints) in the case of L^2 solutions, and the number of false positives provided by L^1 regularization is too high. To overcome these limits, it is common practice in the MEG community to use low-dimensional parametric models for the neural currents: it is possible to assume that the neural current is a superposition of a small ($\sim$10) set of point-wise currents, named current dipoles, whose positions and moments have to be estimated. In this

Journal of Physics: Conference Series **124** (2008) 012046 doi:10.1088/1742-6596/124/1/012046

case the inverse problem becomes non-linear: a considerable amount of non-linear optimization methods have been applied, but the non-uniqueness of the solution leads to a number of local minima which become difficult to handle when the Signal-to-Noise Ratio (SNR) is not very high. In 2003 Somersalo et al. proposed to recast the MEG inverse problem as a Bayesian filtering (or Bayesian tracking) problem, and proposed the use of particle filters for the non-linear problem [8]. The advantage of using a Bayesian method is clear: instead of a difficult search of local-global minima of some potential, one can study the whole probability density on the solution space, conditioned on the measurements. This approach has been validated in [9] where the authors show that the approach is actually feasible and apply the method to a real data set and in [6], where the approach is compared to a state-of-the-art MEG inversion algorithm.
Here we briefly review the basic concepts of particle filtering and suggest that, due to the linear substructure which still remains in the non-linear problem, a semi-analytic approach, known as Rao-Blackwellization in the statistical literature, is feasible and can increase the statistical efficiency of the algorithm. We present numerical results in a 2-dimensional, single-source simplified setting, borrowed from Somersalo et al.; the results show that Rao-Blackwellization actually improves the performances of the algorithm.

2. Particle filters and Rao-Blackwellization

We first present the basic ideas of bayesian filtering and particle filtering, then consider the special case of a linear substructure in the model equation.

2.1. Bayesian filtering and particle filters

Bayesian filtering is a suitable framework for dynamic inverse problems: let $\{X_t\}_{t=1,...}$ and $\{Y_t\}_{t=1,...}$ be two stochastic processes, we will call state process and measurements process respectively, related by

$$X_{t+1} = g(X_t) + W_t^x \tag{1}$$
$$Y_t = f(X_t) + W_t^y \tag{2}$$

where $f(\cdot)$ and $g(\cdot)$ are possibly non-linear functions, and W_t^x and W_t^y represent the process noise and the measurements noise, respectively. In Bayesian filtering, one aims at recovering information about the state process from observations of the measurements process. We denote by $\pi(x)$ the probability density function (pdf) of the random vector (RV) X and by $\pi(x|y)$ the pdf of the RV X conditioned on the realization y of the RV Y. With this notation, the Bayesian filtering problem can be defined as the problem of computing the so-called posterior pdf

$$\pi(x_t|y_{1:t}) \qquad \text{with } y_{1:t} = \{y_1, y_2, ..., y_t\} \tag{3}$$

sequentially in time. Under the assumption of Markovian processes [7], and provided that the transition kernel of the state process $\pi(x_{t+1}|x_t)$, the likelihood function $\pi(y_t|x_t)$ and the first prior density $\pi(x_1|y_0) = \pi(x_1)$ are available, the Bayesian filtering problem can be solved by sequentially applying:

$$\pi(x_t|y_{1:t}) = \frac{\pi(y_t|x_t)\pi(x_t|y_{1:t-1})}{\pi(y_t|y_{1:t-1})} \tag{4}$$
$$\pi(x_{t+1}|y_{1:t}) = \int \pi(x_{t+1}|x_t)\pi(x_t|y_{1:t})dx_t \quad . \tag{5}$$

It is well known that in a linear Gaussian environment the Kalman filter, which sequentially updates the first and second moment of the densities, provides the exact solution of the Bayesian

filtering problem. However this is also the only case where the exact solution is easy to compute. For the general case, numerical approximations are needed: in the last decade, the class of sequential Monte Carlo methods known as particle filters have received great attention, as they allow computation of (4),(5) in a general non-linear non-Gaussian setting [1, 3].
In the Monte Carlo approach, one wants to compute integrals of some target density $\pi(x)$ and therefore tries to obtain samples $\{x^i\}_{i=1,\ldots,N}$ distributed according to $\pi(x)$, because the Law of Large Numbers guarantees

$$\frac{1}{N}\sum_{i=1}^{N} f(x^i) \stackrel{N\to\infty}{\longrightarrow} \int f(x)\pi(x)dx \quad . \tag{6}$$

In Importance Sampling, the impossibility of drawing samples directly from the target density $\pi(x)$ is overcome by a weighting procedure: one draws samples x_q^i from a so-called proposal density $q(x)$ and uses

$$\sum_{i=1}^{N} f(x_q^i)w(x_q^i) \stackrel{N\to\infty}{\longrightarrow} \int f(x)\pi(x)dx \qquad \text{with } w^i = \frac{\pi(x_q^i)}{q(x_q^i)} \quad . \tag{7}$$

The simplest particle filter, known as Sampling Importance Resampling (SIR) particle filter, is indeed the sequential application of an importance sampling strategy; in fact, the two following operations are iterated:

- apply an importance sampling strategy for the posterior density, using the prior density as proposal: draw a sample $\{x_t^i\}_{i=1,\ldots,N}$ from the prior density and exploit the Bayes formula for computing the weights $w_t^i = \pi(y_t|x_t^i)$;
- resample the posterior density: sample N particles $\tilde{x}_t^i$ from the set $\{x_t^i\}_{i=1,\ldots,N}$, with replacement, in such a way that the probability of drawing x_t^i is equal to w_t^i. The new sample set $\{\tilde{x}_t^i\}$ is a random sample from the posterior density.

A few remarks. (i) Drawing points from the prior density is straightforward at the beginning, when the prior density is chosen by the user. For a general t, assume we have a sample $\{x_{t-1}^i\}_{i=1,\ldots,N}$ distributed according to $\pi(x_{t-1}|y_{1:t-1})$; then (6) guarantees that $\sum_{i=1}^N \pi(x_t|x_{t-1}^i) \to \int \pi(x_t|x_{t-1})\pi(x_{t-1}|y_{1:t-1})dx_{t-1}$. Therefore, one can draw samples from this mixture; the easiest way consists in drawing one point for each i. (ii) The resampling procedure, which makes particle filters a powerful method (because here they loose track of unlikely states), introduces correlation among the particles, which makes standard convergence results not valid; however, convergence results for particle filters exist as well, under some additional assumptions (see for example [2]). (iii) Importance sampling is a generalization of random sampling, and the performances are increasingly better as the proposal density resembles the target density. Therefore, the variance of the weights is a measure of the goodness of the proposal density.

2.2. A semi-analytic approach: Rao-Blackwellization

Next we consider a special subset of models, which turns out to be of particular interest for our applied problem. We assume that the state vector can be decomposed $X_t = (P_t, Q_t)$ in such a way that the model equations (1),(2) can be written as

$$P_{t+1} = P_t + \Delta P_t \tag{8}$$
$$Q_{t+1} = Q_t + \Delta Q_t \tag{9}$$
$$Y_t = g(P_t)\cdot Q_t + W_t^y \tag{10}$$

with $g(\cdot)$ a possibly non-linear function; then, under the assumption of Gaussian distribution for W_t^y, Q_1 and ΔQ_t, the model for Q_t conditional on P_t is a linear Gaussian model to which Kalman filter can be applied. To state it differently, the fact that the posterior density naturally splits

$$\pi(p_t, q_t|y_t) = \pi(q_t|p_t, y_t)\pi(p_t|y_t) \tag{11}$$

can be properly exploited and, given a random sample p_t^i for the density $\pi(p_t|y_t)$, one can use a set of Kalman filters for computing the first and second moment of $\pi(q_t|p_t^i, y_t)$, $i = 1, ..., N$. This procedure is known in statistical literature as Rao-Blackwellization; it can be proven that it provides a reduction of the variance of the importance weights [3].

3. The MEG inverse problem

Magnetoencephalography measurements consist of a high-frequency ($\sim$ 1000 Hz) sampling of the magnetic field produced by neural currents; the magnetic field is collected by an helmet shaped array of sensors around the subject's head. Due to the temporal scale of the detectable neural events, the quasi-static approximation holds [4] and the MEG inverse problem is therefore described by the Biot-Savart equation:

$$y_t(r_i) = \int_\Omega x_t(r') \times \frac{r_i - r'}{|r_i - r'|^3} dr' \tag{12}$$

where $y_t(r_i)$ is the magnetic field measured by the i-th sensor, located in r_i, and produced by the current distribution $x_t(r')$ inside the brain volume Ω.

The so called dipolar assumption, commonly used in the MEG community, consists in approximating the neural current distribution inside the head as a set of point-wise currents (current dipoles) $x_t(r) = \sum q_t^i \delta(r - p_t^i)$, located in p_i and of moment q_i. For a single dipolar source, and neglecting the contribution of volume currents, the Biot-Savart equation reads

$$y_t(r_i) = \frac{\mu_0}{4\pi}\left(q_t \times \frac{p_t - r_i}{|p_t - r_i|^3}\right) \quad . \tag{13}$$

We remark that this equation implies a strong non-linear relationship between the dipole position p_t and the measurements. On the other hand, measurements depend linearly on the dipole moment q_t. In the following we denote by $y_t = (y_t(r_1), ..., y_t(r_N))$ the collection of measurements at time t and by $y(x_t)$ the measurement set produced by a current dipole x_t.

4. A Rao-Blackwellized particle filter for MEG

In our applied problem, the state process $\{X_t\}_{t=1,...}$ is the neural current process and the measurements process $\{Y_t\}_{t=1,...}$ is the sequence of magnetic measurements; the two processes are related by (13), a strongly non-linear equation. Following [8], we apply a particle filter for solving such a problem, however, we apply Rao-Blackwellization for improving the statistical performances of the algorithm.
The implementation of a particle filter requires to define three probability densities: the likelihood function, the transition kernel and the initial prior density. The first prior density $\pi(x_1)$ is chosen to be a uniform density for the position p_1, and a Gaussian density with high variance for the dipole moment q_1; the rationale for this choice is that we assume to have no prior information about the source, and thus use non-informative priors. The likelihood function is constructed on the basis of equation (13) and of a Gaussian measurement noise: its explicit form is therefore $\pi(y_t|x_t) = \mathcal{N}(y_t|y(x_t), \Sigma)$, where $\mathcal{N}(\cdot|\mu, \Psi)$ denotes the Gaussian distribution of mean μ and covariance matrix Ψ, and Σ is the covariance matrix of noise, which is assumed

to be stationary. The choice of the transition kernel implies an assumption on the underlying dynamics of the neural current, which is basicly unknown. Therefore we assume no specific dynamics, and treat the neural dynamics as a random walk in the parameter space; the explicit form of the transition kernel is then $\pi(x_{t+1}|x_t) = \mathcal{N}(x_{t+1}|x_t, \Gamma)$, where Γ controls the step size of the random walk.

5. Numerical results: performance improvement and variance reduction

Following [8] we consider a 2D limit of the MEG inverse problem: sources, current dipoles, move on the plane $z = 0$, while the sensors are placed on a square grid on the plane $z = 1$. We consider the synthetic data produced by a single dipolar source, which moves clockwise on the plane, and whose orientation changes as well with time, and add a resonable quantity of Gaussian noise (about 5%).

We run both a SIR particle filter and a Rao-Blackwellized particle filter on the same data set and compare the results provided by the two algorithms. In Fig. 1 we show the reconstructions provided by the standard SIR particle filter and by the R-B particle filter, when both run with 500 particles. The superiority of the Rao-Blackwellized version is clear in this figure, as the positions of the original source and of the reconstructed source coincide almost always, which does not happen for the SIR filter.

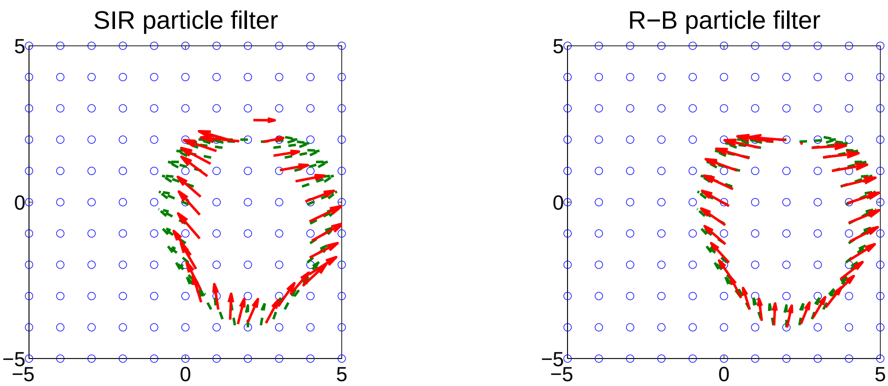

Figure 1. The sequence of source dipoles (dashed) together with the reconstructions (solid) provided by the SIR filter (left) and the R-B filter (right). The improvement obtained by Rao-Blackwellization is clear.

Figure 2 provides a more quantitative information about the difference perfomances of the two methods by plotting the localization error as a function of time. The localization error of the Rao-Blackwellized version is considerably smaller and more constant throught the sequence of measurements with respect to the error produced by the SIR filter.

Finally, we consider the variance of the importance weights, i.e.

$$\mathrm{var}_t = \frac{1}{N}\sum_{i=1}^{N}(w_t^i - \bar{w}_t)^2 \qquad \text{with } \bar{w}_t = \frac{1}{N}\sum_{i=1}^{N} w_t^i \quad . \tag{14}$$

Rao-Blackwellization is expected to provide a reduction of this quantity [3]. Indeed, Figure 3 shows that the variance of the importance weights of the SIR filter is always higher than the variance of the importance weights of the R-B filter. From the statistical viewpoint, this

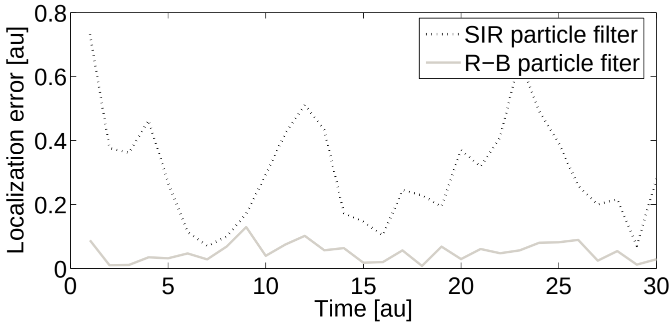

Figure 2. The localization error, computed as the euclidean distance between the position of the source dipole and the position of the reconstructed dipole.

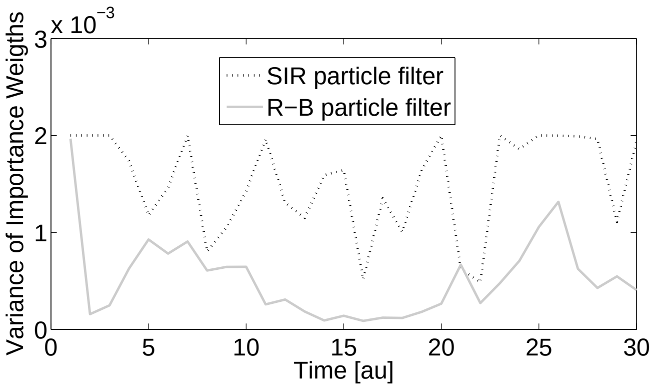

Figure 3. The variance of the importance weights for the standard SIR filter and the R-B version as a function of time.

variance reduction represents an improvement because it means that the importance density is more similar to the target density.

6. Conclusions

We considered the application of particle filters for the solution of the MEG inverse problem in a dipolar approximation. We showed that the linear dependence of the measurements on the dipole moment allows the use of Rao-Blackwellization, in other words, the analytic computation of the posterior density for the dipole moment. This leads to improved results as far as the reconstruction accuracy is concerned, and to a statistical improvement as far as the variance of the importance weights is concerned. On the other hand, we observe that the computational cost of the Rao-Blackwellized version, for a given number of points, is definitely higher than that of the standard SIR filter. A more accurate analysis, together with examples in 3D and applications to real data, will be part of future work.

7. References

[1] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Tr. Sig. Proc.*, 50(2):174–188, 2002.

[2] D. Crisan and A. Doucet. A survey of convergence results on particle filtering methods for practitioners. *IEEE Tr. Sig. Proc.*, 50:736–746, 2002.

[3] A. Doucet, S. Godsill, and C. Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and Computing*, 10:197–208, 2000.

[4] M. Hamalainen, 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.

[5] M. Hamalainen and R.J. Ilmoniemi. Interpreting magnetic fields of the brain: minimum norm estimates. *Med. & Biol. Eng. & Comput.*, 32:35–42, 1994.

[6] A. Pascarella, A. Sorrentino, M. Piana, and L. Parkkonen. Particle filters and rap-music in meg source modelling: a comparison. In *Proceedings of the 15th International Conference on Biomagnetism, Vancouver, BC, Canada, 20-26 August 2006*, volume 1300 of *International Congress Series*. Elsevier, 2007.

[7] E. Somersalo and J.P. Kaipio. *Statistical and computational inverse problems*. Springer Verlag, 2004.

[8] E. Somersalo, A. Voutilainen, and J.P. Kaipio. Non-stationary magnetoencephalography by bayesian filtering of dipole models. *Inverse Problems*, 19:1047–1063, 2003.

[9] A. Sorrentino, L. Parkkonen, and M. Piana. Particle filters: a new method for reconstructing multiple current dipoles from meg data. In *Proceedings of the 15th International Conference on Biomagnetism, Vancouver, BC, Canada, 20-26 August 2006*, volume 1300 of *International Congress Series*. Elsevier, 2007.

[10] K. Uutela, M. Hamalainen, and E. Somersalo. Visualization of magnetoencephalographic data using minimum current estimates. *NeuroImage*, 10:173–180, 1999.