We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

4,800
Open access books available

122,000
International authors and editors

135M
Downloads

Our authors are among the

154
Countries delivered to

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

CLARIVATE ANALYTICS
BOOK CITATION INDEX
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

EEG-fMRI Fusion:
Adaptations of the Kalman
Filter for Solving a High-Dimensional
Spatio-Temporal Inverse Problem

Thomas Deneux[1,2]
[1]Centre National de Recherche Scientifique,
Institut de Neurosciences Cognitives de la Méditerranée, Marseille,
[2]Institut National de la Recherche en Informatique et Automatique, Sophia-Antipolis,
France

1. Introduction

Recording the dynamics of human brain activity is a key topic for both neuroscience fundamental research and medicine. The two main techniques used, Electro- and Magneto-encephalography (EEG/MEG) on the one hand, functional Magnetic Resonance Imaging (fMRI) on the other hand, measure different aspects of this activity, and have dramatically different temporal and spatial resolutions. There is an important literature dedicated to the analysis of EEG/MEG (REF) and fMRI data (REF). Indeed, the both techniques provide partial and noisy measures of the hidden neural activity, and sophisticated methods are needed to reconstruct this activity as precisely as possible. Adaptive filtering algorithms seem well-adapted to this reconstruction, since the problem can easily be formulated as a dynamic system, but it is only recently that such formulations have been proposed for EEG analysis (Jun et al., 2005), fMRI analysis (Johnston et al., 2008; Murray & Storkey, 2008; Riera et al., 2004), or EEG-fMRI fusion (Deneux & Faugeras, 2006b; Plis et al., 2010).

In this chapter, we focus on the so-called "EEG-fMRI fusion", i.e. the joint analysis of EEG/MEG and fMRI data obtained on the same experiment. For more than a decade, EEG-fMRI fusion has become a hot topic, because it is believed that both techniques used together should provide higher levels of information on brain activity, by taking advantage of the high temporal resolution of EEG, and spatial resolution of fMRI. However, the two modalities and their underlying principles are so different from each other that the proposed solutions were often *ad hoc*, and lacked a common formalism. We show here how the use of dynamic system formulation and adaptive filter algorithms appears to be a natural way to achieve the EEG-fMRI fusion.

However, not only do adaptive filtering techniques offer new possibilities for the EEG-fMRI fusion, but also this specific problem brings new challenges and fosters the development of new filtering algorithms. These challenges are mostly a very high

dimensionality, due to the entanglement between the temporal and spatial dimensions, and high levels of non-linearity, due to the complexity of the physiological processes involved. Thus, we will present some new developments that we issued, in particular, the design of a variation of the Kalman filter and smoother which performs a bi-directional sweep, first backward and second forward. And we will show directions for the development of new algorithms.

The results presented in this chapter have already been published in (Deneux and Faugeras, 2010). Therefore, we focus more here on explaining in detail our comprehension of the EEG-fMRI fusion problem, and of its solution through the design of new algorithms. In this introduction, we pose the problem, its specific difficulties, and advocate the use of adaptive filters to solve it. In a second part, we will tackle a simplified, linear, problem: we present our Kalman-based fusion algorithm, discuss its characteristics and prove that it is more suitable to estimate smooth activities, while the estimation of sparse activities would rather necessitate the development of new algorithms based on the minimization of a L^1-norm. In a third part, we will address the problem of strong nonlinearities: we present a modification of the Kalman-based algorithm, and also call for the development of new, more flexible, methods based for example on particle filters.

1.1 Physiological basis of EEG/MEG and fMRI

Figure 1(A) briefly explains how the cerebral activity gives raise to the EEG/MEG and fMRI signals. EEG and MEG measure directly the electrical activity in the brain. In the case of EEG, a set of electrodes (up to 300 in the most advanced EEG helmets) are positioned on the head of the subject, in electric contact with the skin, and measure an electric potential. In the case of MEG, a set of coils are positioned around the head but without touching it, and measure the magnetic field generated by the currents circulating inside the head. These currents themselves are the consequence of the electric activity of a large number of neurons which are activated together. EEG and MEG have an excellent temporal resolution, since the propagation of currents is instantaneous at this temporal scale. They also provide some spatial information, since it is possible to model the current propagation and then solve an inverse problem to localize the activity which generated the specific pattern observed over the different sensors (Hämäläinen et al., 1993). The spatial resolution of this localization however is poor (error range of ~1cm); even, this inverse problem is ill-posed since some sources configurations can generate no signal on the sensors.

fMRI measures secondary effects of the electrical activity, called the hemodynamic response. Indeed, the increased energy consumption in an activated brain region leads to a chain of events, in particular a higher O_2 extraction from the blood, followed by an increase in the blood flow. This impacts the magnetic resonance signals recorded by the MRI scanner, because of the changes in the concentration of the deoxyhemoglobine molecule. Indeed, the magnetic properties of the hemoglobin molecule change after it delivered the oxygen molecule it was carrying, which induces higher decays of the magnetic resonance signals. All in one, a cerebral activity leads to a smooth increase in the MRI signal, also called blood-oxygen level dependent (BOLD) signal; this increase lasts for a few (3~4) seconds, and is usually followed by a small undershoot (Ogawa et al., 1993). This BOLD signal is localized with a millimeter or sub-millimeter precision but, obviously, lacks temporal resolution.
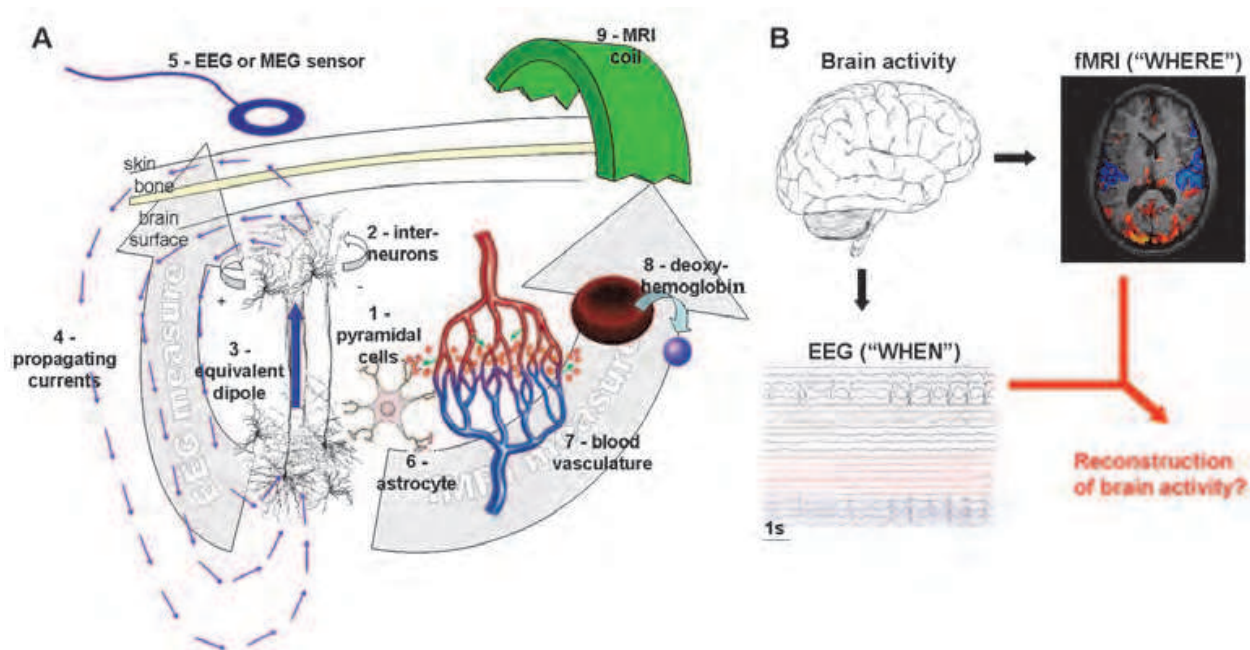
Fig. 1. (A) Physiological basis: this figure briefly summarizes the main effects giving rise to the measured signals. For the EEG/MEG: the brain gray matter is organized in cortical columns where large number of cells work synchronously; in particular, the large pyramidal cells (1), which have a characteristic parallel vertical organization, but also other cellular types such as the smaller interneurons (2); when synchrony is sufficiently large, the electrical activities of the neurons sum up together, and can be represented by an equivalent dipole (3), which generate circulating currents (4) through the brain and even outside of it; EEG sensors touching the skin, or MEG sensors placed close to it (5) can detect voltage differences, or currents, generated by the neural activity. For the functional MRI: neuronal activity (1,2) consumes energy, which is provided by the astrocytes (6), which themselves extract glucose and oxygen from the blood, and regulate the blood flow in the cerebral vasculature; this affects the concentration of deoxyhemoglobin, i.e. hemoglobin which delivered its oxygen molecule; deoxyhemoglobin itself, due to its paramagnetic properties, perturbs the local magnetic field and modifies the magnetic resonance signals recorded by the MRI coil (9). (B) EEG-fMRI fusion: EEG or MEG capture mostly temporal information about the unknown brain activity: the electric signals measured by the sensors on the scalp. On the contrary, fMRI captures mostly spatial information, which leads to precise maps showing which brain regions are active. Ideally, EEG-fMRI fusion should produce an estimation of the activity which takes into account these complimentary information.

It thus appears that EEG/MEG and fMRI recordings are complimentary (Figure 1(B)), the first providing more information on the timing of the studied activity, the second, on its localization. Therefore, many experimental studies combine acquisitions using the two modalities. Such acquisitions can be performed separately, by repeating the same experimental paradigm under both modalities: in such case, averaging over multiple repetitions will be necessary to reduce the trial-to-trial variability. Also, simultaneous acquisition of EEG and fMRI is possible and found specific application in the study of epilepsy (Bénar et al., 2003; Gotman et al., 2004; Lemieux, et al., 2001; Waites et al., 2005)

and resting states (Goldman et al., 2002; Laufs et al., 2003) (note that, since on the contrary MEG cannot be acquired simultaneously with fMRI, we focus here on EEG-fMRI fusion; however, our results will remain true for MEG-fMRI fusion in the context of separate average acquisitions). Apart from a few trivial cases, the joint analysis of the two dataset in order to best reconstruct the observed neural activity presents several specific challenges

1.2 Challenges of EEG-fMRI fusion
1.2.1 Concepts

Many different approaches and strategies have been proposed for EEG-fMRI fusion. Reviews such as (Daunizeau et al., 2010; Rosa et al., 2010) propose different criteria to classify them, two important ones being (i) whether the method is symmetric or not, and (ii) what information do EEG and fMRI share: spatial information, temporal information, or both?

Here, we use schematic examples to help explaining about these different methods. Figure 2(A) shows a very simple example where brain activity only consists of a unique event, which occurs at a specific location and with a specific dynamic. Then, this activity can be fully reconstructed as the cross-product of the spatial information provided by fMRI and the temporal information provided by EEG. On the contrary, in figure 2(B), several events occur at distinct instants and distinct locations, and then more information is needed to determine which part of the signals corresponds to which event. For example, if only spatial information is extracted from fMRI (find 2 regions which are activated), then the weak spatial resolution of the EEG must be used to determine which of the signals it records are likely to originate from the first region or from the second: this is the principle of non-symmetric fMRI-guided EEG reconstructions (Ahlfors et al., 2004). Conversely, one could only extract dynamics from the EEG data, and then use the weak temporal resolution of fMRI to determine which regions in the brain match those dynamics: this is the principle of non-symmetric fMRI analysis based on EEG regressors used in epilepsy, for example (Grova et al., 2008) and rest studies (Laufs et al., 2003). In fact, these two examples are very useful to help understand any EEG-fMRI method, even when additional levels of complexity are added, for example in region-based algorithms which rely on a known parcellation of the cortex (Daunizeau et al., 2007; Ou et al. 2010). And they rather call for the use of symmetric methods, which extract both spatial and temporal information from both the EEG and fMRI.

Figure 2(C) sketches a more complex pattern of activity, and the corresponding EEG and fMRI measures. EEG has the same temporal resolution as the neural activity, but its spatial dimension is smaller, indicating loss of information; and the opposite is true for fMRI. Then, each modality could be used alone to estimate the original activity, while obviously the best estimate should be obtained when using the two datasets.

1.2.2 High dimensionality

Figure 2(C) also introduces a Bayesian formalism to describe EEG-fMRI fusion. If we note u the neural activity, y^{EEG} and y^{fMRI} the EEG and fMRI measures, the aim of fusion is to estimate u given y^{EEG} and y^{fMRI}, or even better, its a posteriori distribution $p(u \mid y^{EEG}, y^{fMRI})$. It would then be also possible to compute a posteriori distribution when considering only one of the modalities, $p(u \mid y^{EEG})$ and $p(u \mid y^{fMRI})$.
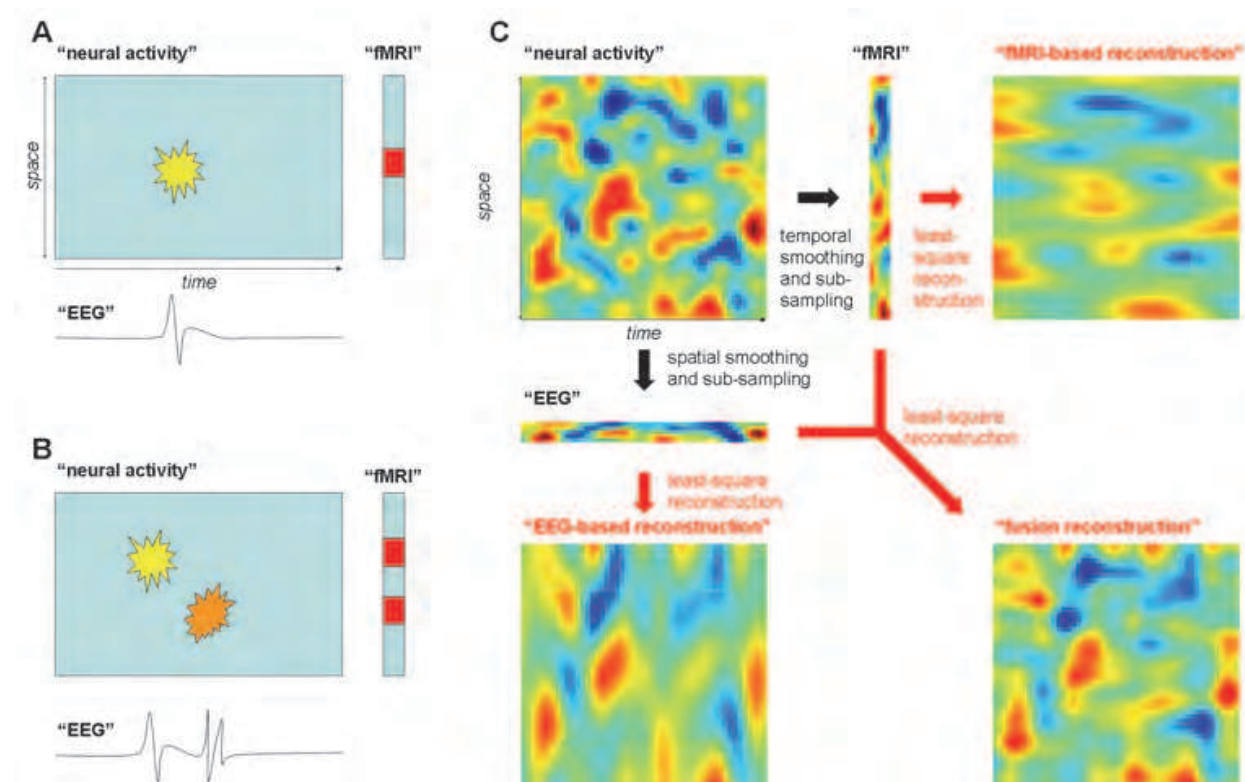
Fig. 2. Schematic representations of EEG-fMRI fusion. (A) The neural activity dynamics are represented by a 2-dimensional array (x-axis is time, y-axis is space). A yellow pattern inside this array marks a single event; the location of this event can be identified precisely by fMRI (red mark), and its exact dynamic can be recorded by the EEG (time courses). In such case, these only information, i.e. spatial information from the fMRI and temporal information from the EEG are sufficient to fully describe the event. (B) The same schematics are used, but two events occur now, at two different locations and with different dynamics (see the yellow and orange patterns). In such case, the sole spatial information from fMRI and temporal information from EEG is not sufficient to fully describe the events since it is not possible to determine which part of these information correspond to which event. (C) Now, a similar spatio-temporal array features a complex activity. Both the temporal and spatial dimensions of EEG and fMRI are considered: the fMRI [EEG, resp.] measure is represented by a narrow vertical [horizontal] array to indicate a reduced temporal [spatial] resolution; more precisely, these measures were obtained by low-pass filtering and sub-sampling the neural activity along the appropriate dimension. The "EEG-fMRI fusion" problem consists in estimating the neural activity given the two measures, and should result in a better reconstruction than when using only the EEG or the fMRI measures (it is indeed the case here, since fMRI-only reconstruction lacks spatial precision, and EEG-only reconstruction lacks temporal resolution).

As we noticed above, the temporal dimension and spatial dimension are highly entangled in the EEG-fMRI fusion problem. Indeed, one EEG measure at time t on a specific sensor is influenced by neural activity at time t in a large part of the cortex if not all; and conversely, one fMRI measure at time t and at a specific spatial location x is influenced by neural activity during the last ~10 seconds before t at location x. Therefore, traditional approaches estimate $p(u \mid y^{EEG})$ independently for each time point, and $p(u \mid y^{fMRI})$ independently for

each spatial location. But it is not possible to "cut the problem in smaller pieces" in order to estimate p(u | y^{EEG}, y^{fMRI}). This results in an inverse problem of very high dimensionality: the dimension of u, which surely depends on the temporal and spatial resolution used. If we choose for u the spatial resolution of fMRI, and the temporal resolution of EEG, then its dimension is the product of several thousands of spatial locations (number of cortical sources) by up to one thousand of time instants per second of experiment (for an EEG sampling rate at 1kHz).

If the experimental data is averaged over repetitions of a specific paradigm, then its total temporal length can be limited to a few seconds, such that the temporal and spatial sizes of u are in the same range. On the contrary, if the interest is in estimating the neural activity without any averaging, the length can be of several minutes or more, and the temporal size of u becomes extremely large. It is in this case that adaptive filter techniques are particularly interesting. Also, it is obvious that, although fMRI measures depend on activity which occurred in the last ~10s, they are independent from earlier activities; therefore, the adaptive filter will need to keep in memory some information in order to link the delayed detections of the activity by EEG and fMRI, but this memory does not need either to cover the full extent of the experiment.
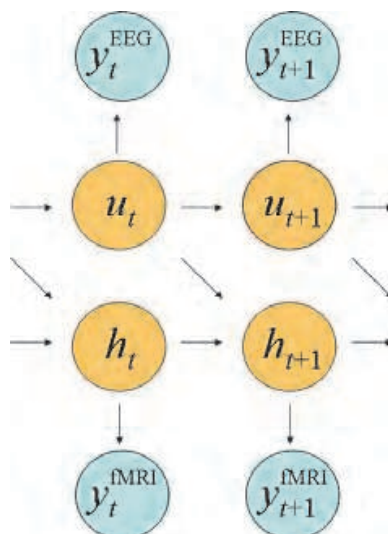

Fig. 3. Graphical representation of the forward model. This model features the evolution of the neural activity u, and of the hemodynamic activity h (driven by the neural activity), the EEG measure y^{EEG} (which depends only on the neural activity), and the fMRI measure (which depends directly only on the hemodynamic state, and hence depends indirectly on the neural activity). Blue background indicate measures, which are known, while orange background indicate hidden states, which have to be estimated from the measures.

The graph in figure 3 represents the forward model which will guide us in designing algorithms for EEG-fMRI fusion. The neural activity at time t, u_t, has its own evolution, and is driving the evolution of metabolic and hemodynamic variables, such as oxygen and glucose consumption, blood flow, volume and oxygenation, represented altogether by the variable h_t. At time t, the EEG measure is a function only of neural activity at the same time, while the fMRI measure is a function of the hemodynamic state. Note that the acquisition rate of fMRI is in fact much lower than that of EEG, but it can be useful for the sake of simplicity to re-interpolate it at the rate of EEG, without loss in the algorithm capabilities (Plis et al., 2010).

1.2.3 Nonlinearity

The second challenge of EEG-fMRI fusion is the high level of nonlinearity which can exist in the forward model depicted in figure 3. To make it clear, let us first show how this graph corresponds to the physiological models depicted in figure 1. Clearly, EEG measure y_t^{EEG} is the signal recorded by the set of EEG sensors (5) and fMRI measure y_t^{fMRI} is the MRI image reconstructed after resonance signals have been recorded by the MRI coil (9). The metabolic/hemodynamic state h_t is the state of all elements involved in the hemodynamic response (6,7,8). The case of the so-called "neural activity" u_t is more delicate, since it can be either a complex description of the actual activities of different types of neurons (1,2), or simply the electric dipole that averages electrical activities in a small region (3).

And here resides the main source of nonlinearity. Indeed, different types of neuronal activities can lead to the same energy consumption and hence to similar fMRI signals, and yet, average into very different equivalent dipoles of current. For example, some activity can result in a dipole with an opposite orientation, or even can be invisible to the EEG. This explains in particular that some activity can bee seen only by fMRI (when electrical activities do not have a preferred current orientation), or only by EEG (when a large area has a low-amplitude but massively parallel activity).

Besides, authors have also often emphasized the nonlinearity of the hemodynamic process (transition $h_t \rightarrow h_{t+!}$), but in fact these nonlinearities, for example those found in the "Balloon Model" modeling (Buxton et al., 2004), are less important, and linear approximations can be used as long as the error they introduce does not exceed the level of noise in the data (Deneux et al., 2006a). Note also that the spatial extent of the hemodynamic response to a local activity can be larger than that of the activity itself, since changes can be elicited in neighboring vessels; however, such distant hemodynamic effects can still be a linear function of the local activity. In any case, such small discrepancies in the spatial domain are usually ignored, mostly because of a lack of knowledge.

As a summary, EEG-fMRI fusion is a difficult problem, because of its high dimensionality, where space and time are intrinsically entangled, and because of nonlinear relations – and surely a lack of robust knowledge – at the level of the underlying link between electric and hemodynamic activities. Therefore, different approaches are proposed to tackle these difficulties. The aforementioned non-symmetric methods are efficient in specific experimental situations. Region-based methods decrease the dimensionality by clustering the source space according to physiologically-based. Here, we do not attempt to decrease the dimensionality, or simplify the inverse problem, but we propose the use of adaptive filters to solve it.

2. Kalman-based estimation under the linear assumption

We first prove the feasibility of EEG-fMRI fusion using adaptive filters under the linear assumption, i.e. we ignore the last comments about nonlinearity and rather assume a straight, linear, relation between the electric activity (modeled as the amplitude of a current dipole directed outward the brain surface) and the energy consumption and hemodynamic changes. The inverse problem can then be solved using the Kalman filter and smoother. We show estimation results on simulated dataset using a neural activity spread on 2000 cortical sources, sampled at 200Hz, during one minute. Since we observe that the method performs better when the activity is smooth rather than sparse, we use schematic examples to prove that indeed, this is the consequence of the minimization of a L^2-norm by the Kalman

algorithm, while new algorithms which would minimize a L1-norm would be more adapted to the estimation of sparse activities.

2.1 Methods
2.1.1 Kalman filter and smoother
The forward model summarized in figure 3 can be simply described in the dynamic model formalism:

$$\begin{cases} \dot{x}(t) = F(x(t)) + \xi(t) \\ y(t) = G(x(t)) + \eta(t) \end{cases}'$$
(1)

where the $x(t)$, the *hidden state*, is the combination of the neural activity $u(t)$ and the hemodynamic state $h(t)$, and $y(t)$, the *measure*, is the combination of $y^{\text{EEG}}(t)$ and $y^{\text{fMRI}}(t)$, $\xi(t)$ is a white noise process, and $\eta(t)$ a Gaussian noise. Once time is discretized and the evolution and measure equations are linearized, it yields:

$$\begin{cases} x_1 = x^0 + \xi^0, & \xi^0 \sim N(0, Q^0) \\ x_{k+1} = A x_k + \xi_k, & \xi_k \sim N(0, Q) \\ y_k = D x_x + \eta_k, & \eta_k \sim N(0, R) \end{cases},$$
(2)

where A and D are the evolution and measure matrices, obtained by linearization of F and G, $N(0, Q^0)$ is the Gaussian initial a priori distribution of the hidden state, and $N(0, Q)$ and $N(0, R)$ are the Gaussian distributions of the evolution and measure noises.

Estimating the hidden state given the measures is performed with the 2 steps of the Kalman filter and smoother (Kalman, 1960; Welch & Bishop, 2006; Welling, n.d.). The first step runs forward and successively estimates the distributions $p(x_k \mid y_1, ..., y_k)$ for increasing values of k. The second runs backward and estimates $p(x_k \mid y_1, ..., y_n)$ for decreasing values of k (n being the total number of measure points).

We recall here the equations for this estimation, for which we introduce the following notation (note that all distributions are Gaussian, therefore they are fully described by their mean and variance):

$$\begin{aligned} \hat{x}_k^l &= E(x_k \mid y_1, ..., y_l) \\ P_k^l &= V(x_k \mid y_1, ..., y_l) \end{aligned}.$$
(3)

First, the Kalman filters starts with the a priori distribution of x_1:

$$\begin{aligned} \hat{x}_1^0 &= x^0 \\ P_1^0 &= Q^0 \end{aligned}'$$
(4)

then repeats for $k = 1, 2, ..., n$ the "measurement update":

$$\begin{aligned} K &= P_k^{k-1} D^T (D P_k^{k-1} D^T + R)^{-1} \\ \hat{x}_k^k &= \hat{x}_k^{k-1} + K(y_k - D\hat{x}_k^{k-1}) \\ P_k^k &= (I - KD) P_k^{k-1} \end{aligned}'$$
(5)

and the "time update":

$$\hat{x}_{k+1}^k = A\hat{x}_k^k$$
$$P_{k+1}^k = AP_k^k A^T + Q \tag{6}$$

Second, the Kalman smoother repeats for $k = n\text{-}1, n\text{-}2, …, 1$:

$$J = P_k^k A^T (P_{k+1}^k)^{-1}$$
$$\hat{x}_k^n = \hat{x}_k^k + J(\hat{x}_{k+1}^n - \hat{x}_{k+1}^k) \tag{7}$$
$$P_k^n = P_k^k + J(P_{k+1}^n - P_{k+1}^k)J^T$$

Specific details about the implementation of the algorithm will be mentioned in the discussion sub-section below.

2.1.2 Physiological models
Now, we give detail on the evolution and measure functions and noises, by modeling explicitly each transition in figure 3.
The neural evolution is modeled by an auto-regressive Ornstein-Uhlenbeck process:

$$\dot{u}(t) = -\lambda u(t) + \xi_u(t) , \tag{8}$$

where λ is a positive parameter which controls the temporal autocorrelation of sources time courses ($<u(t)u(t+\Delta t)>/<u(t)u(t)> = \exp(-\lambda \Delta t)$), and ξ_u is a Gaussian innovation noise. This noise is white in time, but can be made smooth in space and thus make u itself smooth in space, by setting its variance matrix such that it penalizes the sum of square differences between neighboring locations.
Since u represents here the amplitudes of equivalent dipoles of current at the sources locations, the EEG measure is a linear function of it:

$$y^{EEG}(t) = Bu(t) + \eta^{EEG}(t) , \tag{9}$$

where B is the matrix of the EEG forward problem, constructed according to the Maxwell equations for the propagation of currents through the different tissues and through the skull (Hamalainen et al. 1993). The measure noise $\eta^{EEG}(t)$ is Gaussian and independent in time and space.
Finally, the hemodynamic state evolution and the fMRI measure are modeled using the Balloon Model introduced by R. Buxton (Buxton et al., 2004):

$$\begin{cases} \ddot{f}(t) & = & \varepsilon u - \kappa_s \dot{f}(t) - \kappa_f (f(t) - 1) \\ \dot{v}(t) & = & \dfrac{1}{\tau}(f(t) - v(t)^{1/\alpha}) \\ \dot{q}(t) & = & \dfrac{1}{\tau}(f(t)\dfrac{1 - (1 - E_0)^{1/f(t)}}{E_0} - v(t)^{1/\alpha}\dfrac{q(t)}{v(t)}) \\ y^{fMRI}(t) & = & V_0(a_1(v(t) - 1) + a_2(q(t) - 1)) + \eta^{fMRI}(t) \end{cases} \tag{10}$$

where the hemodynamic state $h(t)$ is represented by four variables: the blood flow $f(t)$, its time derivative, the blood flow $v(t)$ and the blood oxygenation $q(t)$. ε , κ_s , κ_f , τ , α , E_0 , V_0 , a_1 and a_2 are physiological parameters. The measure noise $\eta^{fMRI}(t)$ is Gaussian and independent in time and space.
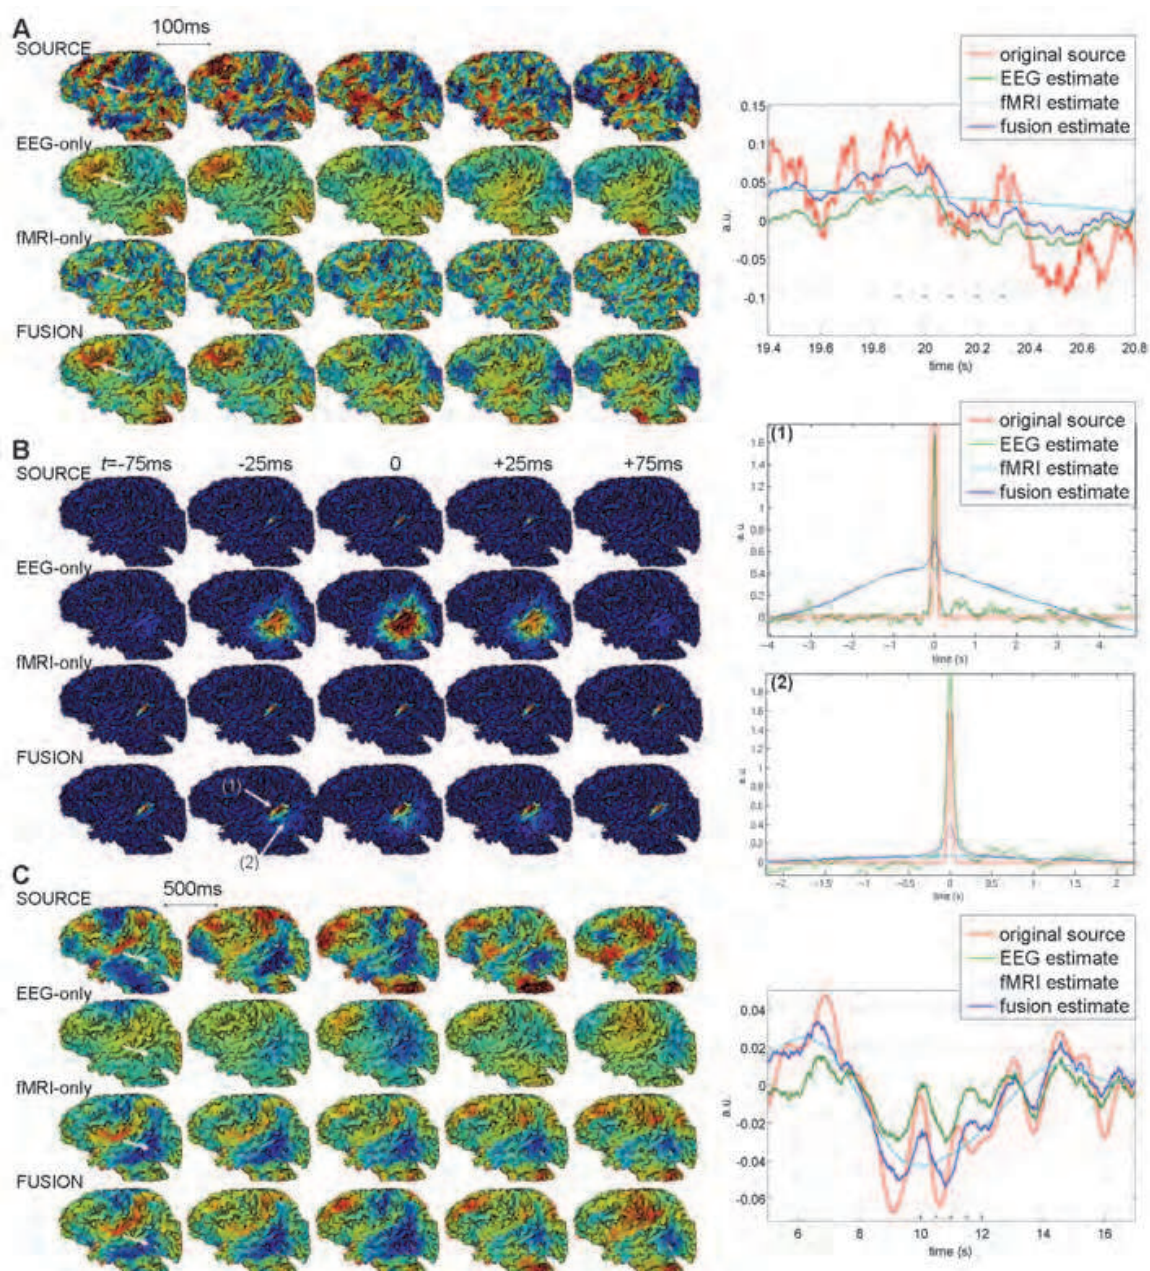
Fig. 4. EEG-fMRI fusion using the linear Kalman filter and smoother. (A) A neural activity has been generated on 2,000 sources spread on the cortical surface according to the model, then EEG and fMRI measures of this activity were generated. 5 snapshots of the activity at intervals of 100ms are shown (row 'SOURCE'). Below are shown the Kalman reconstruction using only EEG measures ('EEG-only'), only fMRI measures ('fMRI-only'), or both together ('FUSION'). Noticeable, EEG estimate is very smooth in space, fMRI estimate is varies very slowly in time, and the fusion estimate is the closest to the true activity. The white arrows indicate one particular source location, and the graph on the right compares the time courses of the true activity of this source with its three estimations. We can also observe that the fMRI estimate is much smoother than the EEG estimate, and that the fusion estimate is the closest to the true signal. (B) In a second simulation, neural activity was chosen to represent a brief (50ms) activation of a few neighboring sources. Similar displays show that EEG finds a brief but spread activity, fMRI find a precisely localized but slow activity, and the fusion estimates finds

EEG-fMRI Fusion: Adaptations of the Kalman
Filter for Solving a High-Dimensional Spatio-Temporal Inverse Problem
243

a trade-off between the two of them. Two locations are shown with arrows: the center of the true activity, and the center of the activity found by EEG alone: again, the two graphs on the right show how the fusion algorithm finds a trade-off between the two previous estimates. (C) In a third simulation, the generated activity is smooth in both time and space. The results are similar to (A), and it appears even more clearly that the fusion algorithm efficiently combines information provided by EEG and fMRI.

Note that the above physiological parameters, as well as the variances of the different evolution and measure noise, are fixed to some physiologically-meaningful values (Friston et al., 2000) for both simulation and estimation.

2.2 Results

We used the forward model (equation (8)) to generate a random neural activity at 2000 source locations spread on the cortical surface, for a time duration of 1 min, at a 200Hz sampling rate (Figure 4(A), first row). And to generate EEG (equation (9)) and fMRI measures (equation (10)) of this activity. Then, the Kalman filter and smoother (equations (4)-(7)) were used to estimate the initial activity, based either on the EEG measure alone, fMRI measure, or both together. Figure 4 shows these estimation results and compares them to the true initial activity: the left part compares snapshots of the activity map (both true and estimated) at different time instants, and the right part compares the true and estimated time courses of one selected source.

Specific characteristics of the estimation results can be observed: the EEG-alone estimations are very smooth spatially, but change fast in time; inversely, the fMRI-alone estimations have more spatial details, but vary very slowly in time; finally, the fusion estimations are the most similar to the true activity. All this is in accordance with the idea that EEG-fMRI should combine the good temporal resolution of EEG and good spatial resolution of fMRI. More quantitatively, table 1 (first column) shows indeed that the fusion estimate is the one which best correlates the true activity. Even, the part of variance explained by the fusion estimate is almost equal to the sum of the parts of variance explained by the EEG and fMRI estimates, indicating that the two modalities capture complimentary rather than redundant information, and that the fusion algorithm efficiently combines these information.

% correlation (and % of variance explained)	Activity generated by the forward model	Sparse activity	Smooth activity
EEG estimate	29.7 (8.8)	7.2 (-14.3)	54.5 (27.6)
fMRI estimate	33.9 (11.5)	7.7 (0.4)	63.6 (40.4)
Fusion estimate	43.1 (18.6)	11.0 (1.1)	74.1 (54.6)

Table 1. **Quantification of estimation accuracies in the case of three different simulations.** We use two different measurements of how good an estimate $\hat{u}$ fits the real signal u, both expressed in percentages. The first is the correlation $< u, \hat{u} > / \| u \|_2 \| \hat{u} \|_2$; it is important to note that correlation does not inform whether signals were estimated with the correct amplitude (the correlation is 1 if the signals are only proportionals). Thus we also use the percentage of variance of source u explained by estimate $\hat{u}$, defined as $1 - \| u - \hat{u} \|_2^2 / \| u \|_2^2$; note that it can be negative even if the correlation if positive, in the case where subtracting $\hat{u}$ to u does not decrease the variance of u (i.e. when the part of the variance of $\hat{u}$ which really accounts for some variance present in u is less than the part of pure estimation error).

This is shown with even more details in figure 5, where true and estimated activities have been decomposed as sums of activities in specific temporal and spatial bandwidths. Thus, the goodness of fit could be computed independently for different temporal and spatial frequencies. We can observe then that fMRI estimates accurately activities with low temporal frequencies (blue surface), while EEG best estimates activities with low spatial frequencies (green surface). And the EEG-fMRI fusion (yellow surface) efficiently combines information provided by both measures, since it provides a better fit than both measures in any frequency domain.

As a result, it appears that EEG-fMRI fusion gives the best performances on activities which are smooth both in time and space, as opposed to fast-varying activities. We illustrated this result by producing two additional simulations. First, we generated a very sparse and focused activity, by activating only a small number of neighboring sources during a short period of time (50ms). Second, we generated a smooth activity, by performing a temporal and a spatial low-pass filtering of activity generated by the model. In both cases, we used the model to generate corresponding EEG and fMRI measures, and then estimated the activity: results are shown in figure 4, parts B and C, and the goodness of fit are displayed in table 1, two rightmost columns. In both cases EEG and fMRI succeed in collaborating to produce a fusion estimation better than they do alone. As expected from the previous considerations, the fusion is very satisfying on the smooth dataset, and more than 50% of the variance of the true sources activity is explained. On the contrary, its result on the sparse activity is disappointing: it seems to give an average of the EEG and fMRI estimates rather than recovering the focused activity.
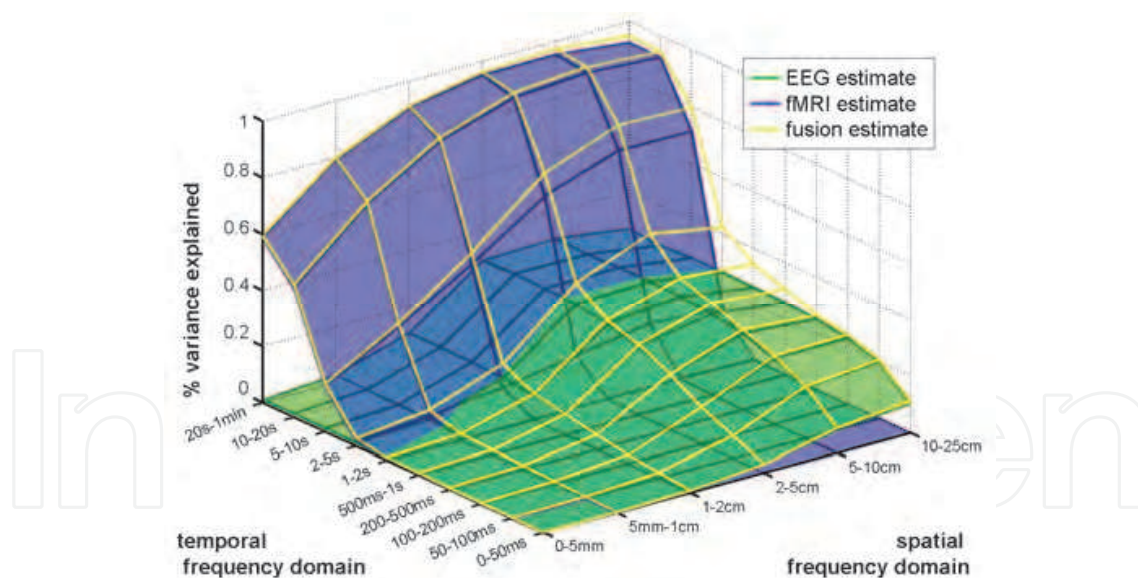

Fig. 5. Spectral analysis of EEG-fMRI fusion efficiency. The correspondence between the true activity and estimated activity are quantified independently in separate temporal and spatial frequency domains, using the percentage of variance explained as in table 1. We observe that fMRI estimation is accurate at low time frequencies, while EEG estimation is accurate at low spatial frequencies. The fusion estimation is efficient, in the sense that its accuracy is superior to both EEG and fMRI estimations in all frequency domains, and in some specific domains where EEG and fMRI do bring complementary information (at temporal scale of 1~10s and spatial scale >2cm), it yields even higher values by combining these information.

2.3 Discussion
2.3.1 Computational cost

Our simulations prove the feasibility of EEG-fMRI fusion at a realistic scale despite its high dimensionality, at least in the case of a linear forward model, using the Kalman filter and smoother algorithms. We would like to stress here the specific aspects of their implementation which made it possible, and further improvements which are desirable.

First of all, it is important to note that the fact that we used a linear model was critical for keeping a reasonable computational cost, because the Kalman filter and smoother exhibit some particular properties which would have been lost if we had used the extended Kalman filter to run estimations using the exact nonlinear equations (in particular, equation (10)).

Indeed, the most costly operations are the matrix computations in equations (5)-(7). The P_k^{k-1} and P_k^k variance matrices describe the degree of certitude about the hidden state estimation just before and after the measure update; they obviously depend on the characteristics of the noise (matrices Q^0, Q and R), but, in the case of the linear Kalman filter, they do not depend on the values of the measurements y_k. Moreover, they converge to some limits (Welling, n.d.), which we note $\bar{P}$ and $\hat{P}$, when $k \to +\infty$; these limits are the solutions of the system:

$$\begin{cases} K = \bar{P}D^T(D\bar{P}D^T + R)^{-1} \\ \hat{P} = (I - KD)\bar{P} \\ \bar{P} = A\hat{P}A^T + Q \end{cases} \qquad . \tag{11}$$

Therefore, the variance matrices can be pre-computed, and for values of k sufficiently large, it is possible to use $\bar{P}$ and $\hat{P}$ instead, which in the case of long experiments decreases dramatically the number of times that the matrix operations in (5) and (6) must be performed. We can even go one step further by choosing $Q^0 = \bar{P}$: then, for all k we have $P_k^{k-1} = \bar{P}$ and $P_k^k = \hat{P}$. This choice for Q^0 is equivalent to saying that at time 0 the system has been running for a long time already according to its natural evolution equations, and that we have an information on its state from previous measures; since in fact such previous measures do not exist, this a priori variance is too small, which could have the consequence that the beginning of the estimated sequence underfits the measures; however we did not observe important estimation errors in the beginning of our simulations; on the other hand this choice is particularly convenient in terms of computational and memory cost.

The EEG-fMRI fusion algorithm can then be performed the following way:

- Apply iteratively the matrix computations in (5) and (6) until P_k^{k-1} and P_k^k converge to their limits $\bar{P}$ and $\hat{P}$, starting with $P_1^0 = Q$ and stopping after a number of iterations determined heuristically

- Compute $K = \bar{P}D^T(D\bar{P}D^T + R)^{-1}$ and $J = \hat{P}A\bar{P}^{-1}$

- Apply the vector operations in equations (5), (6) and (7) to compute the estimation of the hidden states, using for all k $P_k^{k-1} = \bar{P}$ and $P_k^k = \hat{P}$. Note that it is not necessary to compute the values of the P_k^n variance matrices.

The total computational cost is determined by the cost of the first step, the matrix convergence. The time complexity of this convergence is O($n^3 f_s$), where n is the number of cortical sources, and f_s the sampling frequency: indeed, the limiting factor is the

multiplication of matrices of size O(n), repeated O(f_s) times to reach the convergence. Since memory limitation can also occur for large values of n, it is also important to know the space complexity: it is equal to O(n^2), the number of elements of the variance matrices. For our simulations, the convergence of the variance matrices (300 iterations) took 24 hours on an AMD Opteron 285, 2.6GHz, and occupied 5Gb RAM. Since we used n = 2,000 sources, and since the size of the hidden state $x(t)$ = {$u(t),h(t)$} is 5n, it involved multiplications and inversion of square matrices of size 10,000.

One should note that such size is in fact quite small, compared to the total size of the neural activity to be estimated. For example, in our simulation the number of elements of u is $n*f_s*T$ = 2,000*200*60 = 24,000,000. As mentioned in the introduction, decreasing this dimensionality was made possible by the use of adaptive filtering.
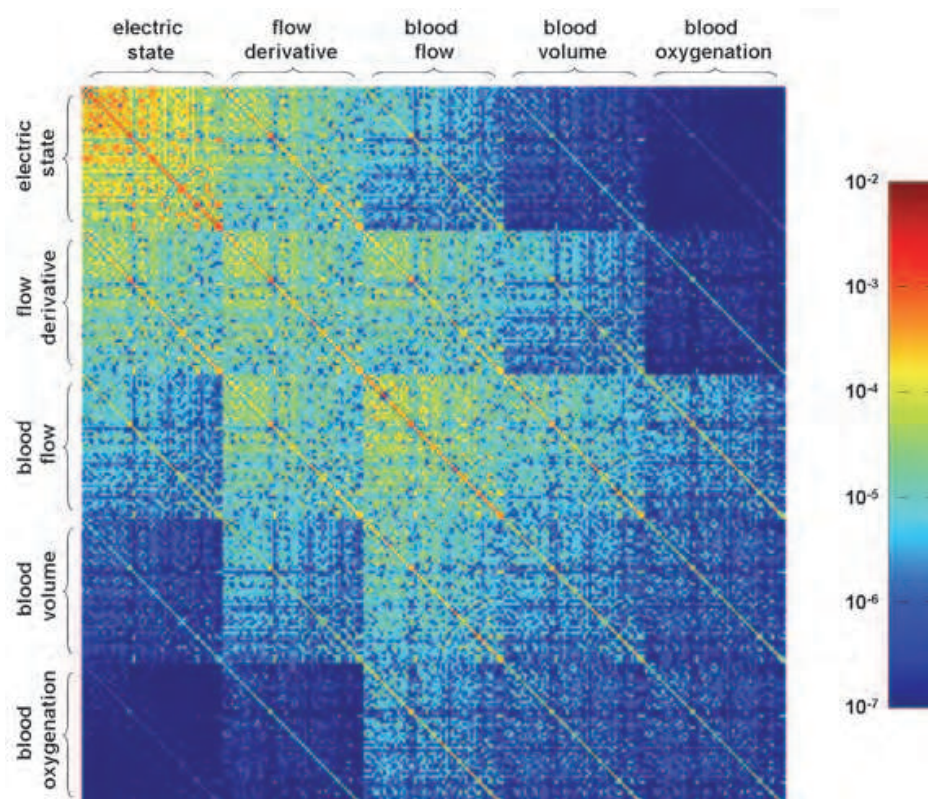

Fig. 6. Structure of the Kalman covariance matrices, displayed at a logarithmic scale. An example variance matrix is shown: the $\hat{P}$ matrix obtained after convergence of the P_k^k, in the case of a low-dimension simulation with 50 sources. Since the hidden state $x(t)$ is the concatenation of the electric state $u(t)$ and the four hemodynamic variables in $h(t)$, this matrix, which expresses the a posteriori variance of $x(t)$ given all measures at $t'<t$ is organized by blocks: the first block is the a posteriori variance of $u(t)$, the next diagonal blocks are the variances of the hemodynamic states, and the other blocks are covariances between different states. Within one block, the diagonal contains the variances for individual sources, while the other elements are covariance between different sources. Although the matrix is dense, it is highly structured, and a large number of values are very small (note the logarithmic scale: 6 orders of magnitudes are displayed): we suggest that future work will study this structure and the convergence process, and look for ways to decrease its computational cost.

A promising direction for decreasing the computational cost would be to describe the variance matrices in a more compact way than full symmetric matrices. Indeed, although the matrices are not sparse, they exhibit some specific structure, as shown in figure 6, and studying carefully this structure should enable the description of the matrices using a lower dimension representation. Maybe even these matrices can be only approximated: indeed, they contain covariance information between all pairs of hidden states, for example between the electric state at a first location and the hemodynamic state at second distant location, which is not necessarily critical for the estimation. However, it must be ensured that approximations do not affect the numerical stability of the algorithm.

Another approach would be to use specialized numerical methods in order to reach faster the convergences of matrices $\bar{P}$ and $\hat{P}$, unless an analytic expression as a function of Q and R can be derived!

2.3.2 Algorithm

Our simulations showed that using the Kalman filter and smoother to perform EEG-fMRI fusion is efficient when estimating an activity which is smooth in space and time. In such case indeed, the effects of the disparity between the resolutions of both modalities are attenuated, therefore the information they provide on the underlying neural activity can be combined together in an efficient way.

On the contrary, when estimating a sparse and focused activity as in figure 4(B), results are more disappointing. The EEG-only estimation is focused in time, but spread and imprecise in space, while the fMRI-only estimation is focused in space, but spread in time. Thus, an ideal fusion algorithm should be able to reconstruct an activity focalized both in time and space. But instead, our fusion estimate looks more like an average between the EEG-only and fMRI-only estimates.

In fact, this is a direct consequence of using the Kalman filter; even, this is a consequence of using Bayesian methods which assume Gaussian distributions! Indeed, when using the Kalman filter or any such method, the a priori distribution of the neural activity u is a Gaussian (fully described in our case by the initial distribution $p(u_1)$ and the Markov chain relation $p(u_{k+1}|u_k)$), and estimating u will rely on a term which minimizes a L^2-norm of u (i.e. a weighted sum of square differences). We show using schematic examples that the minimization of a L^2-norm results in solutions which are smooth rather than sparse, and in the case of two measures of the same activity u which have very different temporal and spatial resolution, estimation results show similar pattern as observed above with two components, one smooth in time, the other smooth in space.

These examples are as in figure 2(C): a 2-dimensional array represents u, one dimension standing for time and the second for space. We use a simple forward model: "EEG" and "fMRI" measures are obtained by low-pass filtering and subsampling u in the spatial and temporal dimension, respectively, resulting in lower resolution versions of u. The corresponding inverse problem – estimate u given the two measures – is then strongly under-determined, as is the real EEG-fMRI fusion inverse problem, and some constraint can be imposed on the estimated $\hat{u}$ in order to decide between an infinity of solutions. We compared minimizing the L^2-norm of $\hat{u}$, i.e. its mean square, and minimizing the L^1-norm, i.e. its mean absolute value. Figure 7 shows several different estimations, using various patterns for the true activity u.
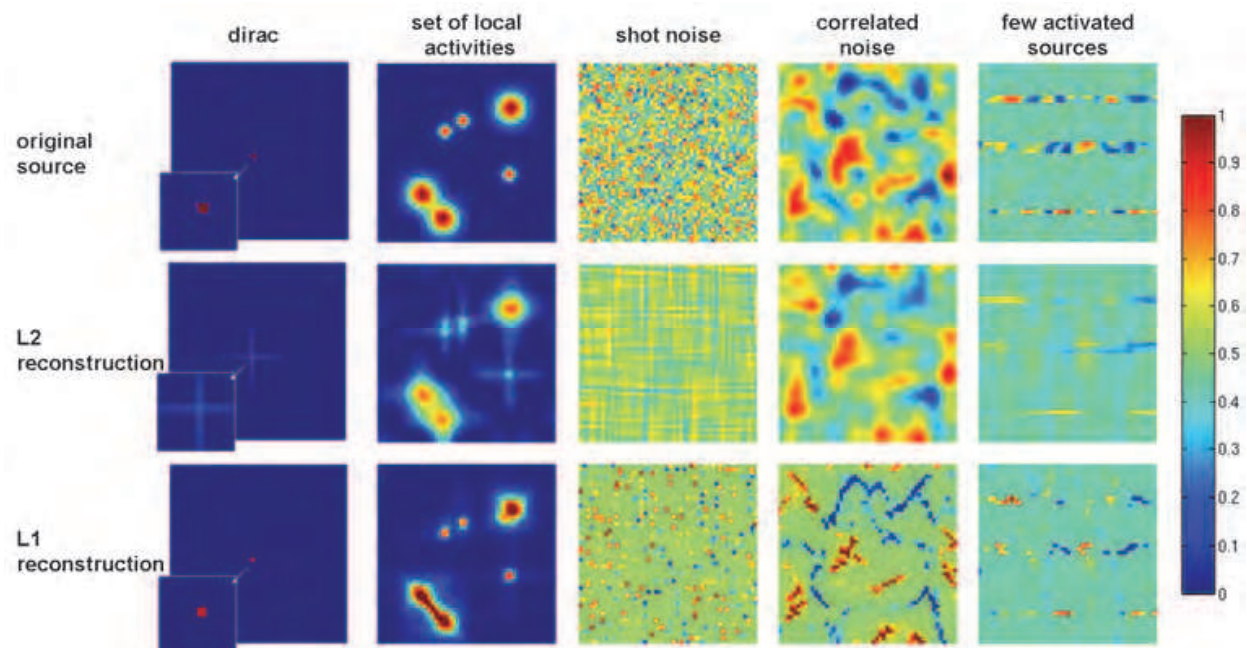
Fig. 7. Schematic reconstructions illustrate the different behavior of L²-norm based and L¹-norm based algorithms. Different patterns of activity are generated, one per column, to produce 2-dimensional array as in figure 2(C), representing the neural activity. Schematic EEG and fMRI measures are generated, again as in figure 2(C). The first row displays these activity. The second and third row display reconstructions of the activity from the 2 measures, based either on the minimization of the L²-norm or the L¹-norm of the source. We can notice that L¹-norm method produces sparse estimates (which is preferable in the case of a Dirac, a set of local activities, or when only a few sources are activated). The L²-norm method produces smoother estimates, more precisely, the sum of two components, one smooth in time, the second smooth is space; it is more appropriate to the case of a smooth activity (the 'correlated noise' column).

It happens that, very clearly, the L¹-norm leads to sparse estimates, while the L²-norm leads to smooth estimates, which even more precisely show two component, one smooth in time, the second smooth in spread. For example, a Dirac is perfectly estimated by the L¹-norm algorithm, and very poorly by the L²-norm algorithm. A set of focal activities is also better estimated using the L¹-norm. On the contrary, a smooth activity is better estimated using the L²-norm.

In fact, all this is quite well-known in the image processing domain (Aubert and Kornprobst 2006; Morel and Solimini 1995). However, the consequences are important: algorithms based on L²-norm such as the Kalman filter and smoother are appropriate for EEG-fMRI fusion whenever the neural activity to be estimated is known to be smooth in time and space (one can think for example about spreading waves of activity). Inversely, if the activity is known to be sparse, the inverse problem should be solved through the minimization of other norms, such as the L¹-norm, which means that any Bayesian method relying on Gaussian distribution would not be appropriate. Rather this calls for the development of new algorithms, if possible still performing a temporal filtering, which would be based for example on the L¹-norm.

EEG-fMRI Fusion: Adaptations of the Kalman
Filter for Solving a High-Dimensional Spatio-Temporal Inverse Problem
249

3. Handling nonlinear electric-metabolic relation

As we explained in the introduction, assuming a linear relation between the electric activities that yield the EEG measure and the metabolic and hemodynamic activities that yield the fMRI BOLD signal is a coarse approximation, and can even be a non-sense for a wide range of activities. For example, applied on oscillating activities, where the orientation of the equivalent dipole is repeatedly inverted, a linear model would generate no hemodynamic activity (since temporal averaging of the oscillations would be null). It is thus necessary to use a nonlinear modeling: for example, one would think to consider energy consumption not any more proportional to the electric activity, but to its absolute value, and regardless of his sign.

However, in the case of a nonlinear model, the EEG-fMRI fusion inverse problem becomes much harder to solve. Though we still advocate the use of adaptive filter techniques, the choice of the Kalman filter, or here its nonlinear version, the extended Kalman filter (EKF), becomes more problematic. First, the computational cost advantage of pre-computing the variance matrices cannot apply any more.
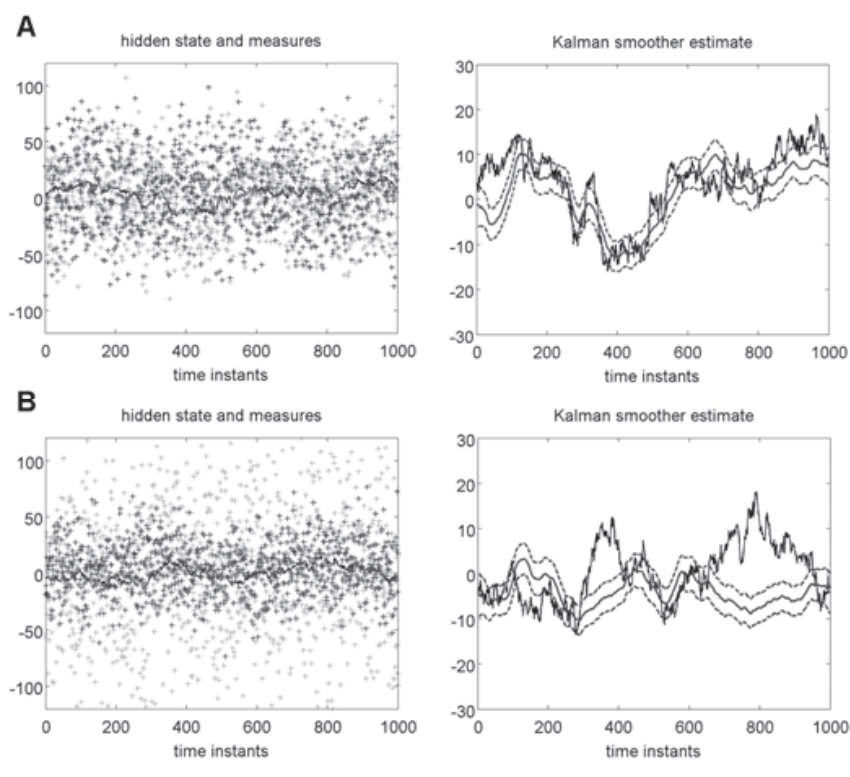

Fig. 8. Limits of the extended Kalman filter and smoother. (A) Left: the time course of a scalar hidden state is generated as a Brownian motion (solid line), and a 2-elements noisy measure is generated, the first measure being obtained by the linear identity function (dark gray crosses), and the second, by the non-linear absolute value function (light gray crosses). Right: reconstruction of the hidden state by the Kalman filter and smoother; black solid line: true time courses; gray solid line: estimated time courses; gray dashed lines: representation of the a posteriori variance, the true time courses indeed lies within this variance. (B) More weight is given to the nonlinear measure (the linear measure is noisier, while the nonlinear measure is less noisy). The estimation now fails is gets trapped in local minima where the sign of the hidden state is misestimated.

Second, EKF can handle only small nonlinearities, and not strong ones such as the use of an absolute value. Simulations in figure 8 demonstrate this fact: we use a simple dynamical model, the hidden state is a one-dimensional Brownian motion, and the measure consists of two scalars, the first equal to the hidden state plus noise, the second equal to the absolute value of hidden state plus noise (left part of the figure). In figure 8(A), the noise in the first measure is sufficiently low so that the estimation using the Kalman filter and smoother is successful. On the contrary, in figure 8(B), the noise in the first measure is higher, so that information on the sign of the hidden state is weaker: then the estimation diverges strongly from the true value. More precisely, the estimation gets trapped in a local minimum, and the local linearizations and approximations with a Gaussian performed by the EKF become totally erroneous compared to the real a posteriori distribution.

As a consequence, we would like to encourage the development of new adaptive filter algorithms, based for example on particle filtering. However, we will show in this section on low-dimensionality examples that it is still possible to use the extended Kalman filter, thanks to (i) a modeling of electric-metabolic relations which minimizes nonlinearities and can handle the strong disparity between the fast temporal variations of the EEG signals and slow variations of the fMRI signals, and (ii) a new backward-forward estimation scheme that minimizes estimation errors.

3.1 Methods
3.1.1 Physiological model
Several studies attest that fMRI signals correlate particularly well oscillating EEG activities in specific frequency domains (Kilner et al., 2005; Laufs et al., 2003; Martinez-Montes et al., 2004; Riera et al., 2010). In such cases, the hemodynamic activity depends linearly, not on the current intensities themselves, but rather on the power of the currents in these specific frequencies. We propose here a simple model for the current intensity that favors oscillating patterns and makes it easy to link hemodynamics to the oscillation amplitude. It consists in describing the current intensity at a given source location as:

$$u(t) = \chi(t)\cos(2\pi f_c t + \varphi(t)) , \qquad (16)$$

where $\chi(t)$ is the envelope of the signal, f_c is the preferred frequency, and $\varphi(t)$ a phase shift. $\cos(2\pi f_c t)$ reflects the intrinsic dynamic of the region, while $\chi(t)$ reflects a modulation of this activity in amplitude, and $\varphi(t)$ a modulation in frequency (if $\dot{\varphi}(t) > 0$, the frequency is higher, and if $\dot{\varphi}(t) < 0$, the frequency is lower) . Note that although $u(t)$ varies fast, $\chi(t)$ and $\varphi(t)$ can change at a slower pace (figure 9).
The evolution of χ and φ is modeled as follows:

$$\begin{cases} \dot{\chi}(t) = -\lambda_\chi \chi(t) + \xi_\chi(t), & \xi_\chi \sim N(0, Q_\chi) \\ \dot{\varphi}(t) = -\lambda_\phi(\varphi(t) - \tilde{\varphi}^{\sigma\varphi}(t)) + \xi\varphi(t), & \xi_\varphi \sim N(0, Q_\varphi) \end{cases}. \qquad (17)$$

χ is driven by the same Ornstein-Uhlenbeck process as in our earlier model (8), except that we did not allow χ to become negative (at each iteration of the generation of the simulated data, whenever it would become negative, it was reset to zero). φ is driven by a similar innovative process, but the phase shift at a given source location, instead of being pulled back towards zero, is pulled back towards the average phase shift over neighboring locations, $\tilde{\varphi}^{\sigma\varphi}$: this ensures a spatial coherence between sources, controlled by the model parameter σ_φ.
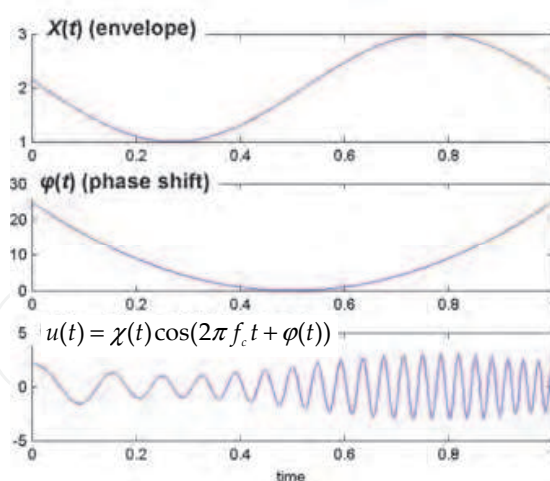
Fig. 9. Hidden states of the nonlinear model for neural activity. The signal u is defined by two hidden states χ and φ which can vary much slower than u itself.

Then we model the EEG measure as in equation (9), i.e. it is a linear function of $u(t)$. Note however that, since now u itself is represented by χ and φ, the EEG measure has become a nonlinear function of the hidden state: in fact, the transition $u_t \rightarrow y_t^{EEG}$ is now nonlinear. Besides, fMRI measure now depends linearly on the signal envelope $\chi(t)$: the first equation in (10) is replaced by

$$\ddot{f}(t) = \varepsilon \chi(t) - \kappa_s \dot{f}(t) - \kappa_f(f(t) - 1) . \tag{18}$$

3.1.2 Backward-forward scheme for the extended Kalman filter

We now have a nonlinear dynamical system:

$$\begin{cases} x_1 = x^0 + \xi^0 , & \xi^0 \sim N(0,Q^0) \\ x_{k+1} = F(x_k) + \xi_k , & \xi_k \sim N(0,Q) \\ y_k = G(x_x) + \eta_k , & \eta_k \sim N(0,R) \end{cases} \tag{19}$$

where F and G are nonlinear evolution and measure functions. Whereas it would be possible to linearize again F without major harm, the measure function G on the contrary is highly nonlinear because of the $\cos(2\pi f_c t + \varphi(t))$ in the EEG measure.

We would like to use the extended Kalman filter and smoother (Arulampalam et al., 2002; Welch & Bishop, 2006) to estimate the hidden state. However, as observed in figure 9, small errors of the estimation can lead to erroneous linearizations, and therefore to increased errors in the next filtering steps, causing divergence. Also, as explained earlier, the classical Kalman filter estimates $p(x_k | y_1,...,y_k)$ for k = 1…n, and the Kalman smoother estimates $p(x_k | y_1,...,y_n)$ for k = n-1…1. The problem is that $p(x_k | y_1,...,y_k)$ does not use any fMRI information to estimate x_k, since the fMRI measure occurs only several seconds later. Therefore, we propose to first run a backward sweep during which we estimate $p(x_k | y_k,...,y_n)$ for k = n…1, and only after a forward sweep, to estimate $p(x_k | y_1,...,y_n)$ for k = 2…n. Indeed, $p(x_k | y_k,...,y_n)$ will make less estimation error since $\{y_k,...,y_n\}$ contains both EEG and fMRI information about activity x_k.

We introduce the new notation:

$$\begin{aligned}\breve{x}_k^l &= E(x_k \mid y_l,...,y_n) \\ \breve{P}_k^l &= V(x_k \mid y_l,...,y_n)\end{aligned}. \tag{19}$$

First, it is important to point out that, in the absence of any measurement, all states x_k have the same a priori distribution $N(x^0, P^0)$, where P^0 can be obtained by repeating the Kalman filter time update step until convergence:

$$P^{(0)} = 0, \quad P^{(k+1)} = AP^{(k)}A^T + Q, \quad P^0 = \lim_{k \to \infty} P^{(k)}. \tag{20}$$

First, the backward step starts with the a priori distribution of x_n:

$$\begin{aligned}\breve{x}_n^{n+1} &= x^0 \\ \breve{P}_n^{n+1} &= P^0\end{aligned}', \tag{21}$$

then, for $k = n, n\text{-}1, ..., 1$, repeats the "measurement update":

$$\begin{aligned}K &= \breve{P}_k^{k+1}D^T(D\breve{P}_k^{k+1}D^T + R)^{-1} \\ \breve{x}_k^k &= \breve{x}_k^{k+1} + K(y_k - G(\breve{x}_k^{k+1})) \\ \breve{P}_k^k &= (I - KD)\breve{P}_k^{k+1}\end{aligned}, \tag{22}$$

where D is the derivative of the measure function at $\breve{x}_k^{k+1}$: $G(x) \approx G(\breve{x}_k^{k+1}) + D(x - \breve{x}_k^{k+1})$, and the "backward time update":

$$\begin{aligned}J &= P^0 A^T(AP^0 A^T + Q)^{-1} \\ \breve{x}_k^{k+1} &= x^0 + J(\breve{x}_{k+1}^{k+1} - x^0) \\ \breve{P}_k^{k+1} &= J\breve{P}_{k+1}^{k+1}J^T + (I - JA)P^0\end{aligned}, \tag{23}$$

where A is the derivative of the evolution function at $\breve{x}_k^{k+1}$: $F(x) \approx F(\breve{x}_k^{k+1}) + A(x - \breve{x}_k^{k+1})$.

Second, the forward smoothing steps repeats for $k = 2, 3, ..., n$:

$$\begin{aligned}L &= (Q^{-1} + (P_{k+1}^{k+1})^{-1})^{-1}Q^{-1} \\ \breve{x}_{k+1}^1 &= \breve{x}_{k+1}^{k+1} + L(F(\breve{x}_k^1) - \breve{x}_{k+1}^{k+1}) \\ \breve{P}_{k+1}^1 &= (Q^{-1} + (P_{k+1}^{k+1})^{-1})^{-1} + LA\breve{P}_k^1(LA)^T\end{aligned}. \tag{24}$$

We derived these equations in a similar way to how the measure update step is derived in (Welling, n.d.):

- Our measure update is the same as for the Kalman filter.
- The backward time update is obtained by first proving that $p(x_k \mid x_{k+1}) = N(x^0 + J(x_{k+1} - x^0), (I - JA)P^0)$, and then by calculating $p(x_k \mid y_{k+1},...,y_n) = \int_{x_{k+1}} p(x_k \mid x_{k+1})p(x_{k+1} \mid y_{k+1},...,y_n)dx_{k+1}$.
- The forward smoothing step is obtained by first proving that $p(x_{k+1} \mid x_k, y_1,...,y_n) = N(\breve{x}_{k+1}^{k+1} + L(F(\breve{x}_k^1) - \breve{x}_{k+1}^{k+1}), (Q^{-1} + (P_{k+1}^{k+1})^{-1})^{-1})$, and then by calculating $p(x_{k+1} \mid y_1,...,y_n) = \int_{x_{k+1}} p(x_{k+1} \mid x_k, y_1,...,y_n)p(x_k \mid y_1,...,y_n)dx_{k+1}$.

Since we observed some divergences when applying this second smoothing step due to accumulating errors in the estimation of the phase, we applied it only on the estimation of the envelope (and kept the result from the first step only for the phase estimate).

3.2 Results

We have run this algorithm on a highly reduced dataset where the cortical surface was downsampled to 10 cortical patches, and only 3 EEG electrodes were considered (so as to preserve the underdetermination of the EEG backward problem). We generated 1min long patterns of activity, sampled at 200Hz, with an average frequency of the oscillations f_c=10Hz. First we generated a cortical activity and measures according to the model (17); then we estimated the activity from these measures: figure 10(A) shows estimation results, which are quantified in table 2, first two columns. We compare estimations obtained by applying the filtering technique to either the EEG alone, fMRI alone, or both together (in the case of fMRI alone, since there is no information on the phase φ, only the estimated envelope is displayed): we see that the envelope of the signals is best estimated using fMRI alone, whereas the EEG-fMRI fusion estimates the signals better than EEG-only. We also generated a one-second width pulse of activity and noisy fMRI and EEG measures, and the estimation results (figure 10(B) and two last columns in table 2) are even more significant: the fusion estimate clearly provides the best estimate of the signals and of their envelopes.

% correlation	Activity generated by the forward model		Sparse activity	
	Original source	Envelope of source	Original source	Envelope of source
EEG estimate	29.0	17.6	42.2	36.2
fMRI estimate		50.7		49.6
Fusion estimate	24.	34.1	23.2	68.5

Table 2. Quantification of estimation accuracies for the nonlinear model in the case of two different simulations using 10 sources. For every estimation, we quantify using correlation the fit between the estimated sources and the original simulated sources ($u = \chi \cos \varphi$), as well as between the envelopes (hidden state χ). Note that in the case of fMRI-only estimation, there is no information on φ, so no number is given for the fit to original source and for the fit to EEG.

3.3 Discussion
3.3.1 Algorithm

We briefly discuss here the specific ideas presented in this section. First, the way we proposed to describe the electric activity in term of envelope and phase is very specific, and should be used only in the analysis of experiments where it is well-founded (i.e. experiments where the fMRI signals are known to correlate with EEG activity in a given frequency range). However, the key idea of this modeling could still be used in different models: it is to set apart the information which is accessible to fMRI, in our case χ, the envelope of the activity, which somehow measures "how much the source is active" in total, but is ignorant of the exact dynamics of the currents u. In such a way, EEG and fMRI measures collaborate in estimating χ, but the phase shift φ is estimated only by the EEG (at least at a first approximation, since in fact in the fusion algorithm, all the different states are linked through their covariance). Furthermore, χ is easier to estimate for the fMRI since it can evolve slower that u, which corresponds more to its weak temporal resolution; and similarly, φ can be easier to estimate for the EEG if there is a high phase coherence among the sources of a same region, because then less spatial resolution is needed.
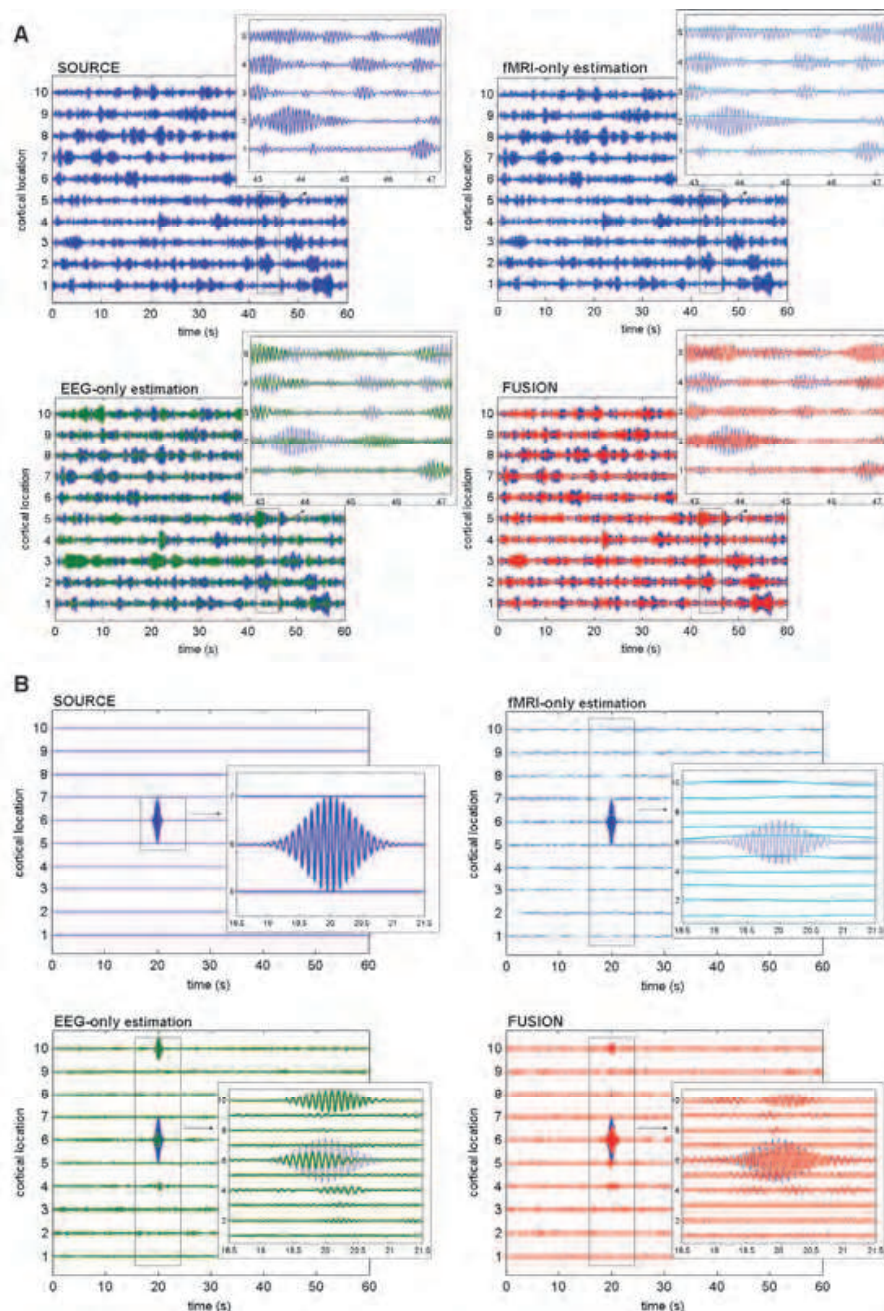
Fig. 10. EEG-fMRI fusion using the nonlinear Kalman algorithm. (A) A neural activity has been generated on 10 sources according to the nonlinear model, then EEG and fMRI measures of this activity were generated (panel 'SOURCE'). We can observe in the structure of this activity the modulations of χ, the envelope of the signals, the intrinsic oscillatory activity at frequency f_{c}, and in the magnified area, the dephasing between different sources due to the variations of the state φ. EEG and fMRI measures were generated according to the model, and estimations were performed on either EEG, fMRI or both (the 3 next panels). Note that the fMRI estimate is not oscillating since φ is not estimated. We can observe the improved estimation in the fusion case compared to the EEG case. (B) In a second simulation, a single pulse of activity was generated. We can observe very well that the EEG estimate finds the accurate dynamics, the fMRI estimate finds the accurate localization, and the fusion produces a very satisfying estimate, with the correct dynamic and localization.

Indeed, we could run estimation for ten (as shown here) to a few hundred sources, as shown in (Deneux & Faugeras, 2010) spread on the cortical surface, but to estimate the activity of several thousands sources would be too demanding, because the simplifications which apply in the linear case (section 2.3.1) do not any more. However, there is a large room for further improvements.
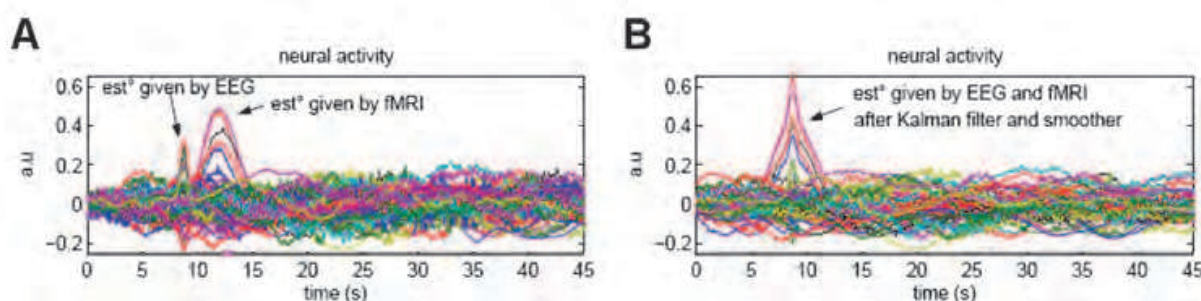

Fig. 11. Details of the Kalman filter and smoother steps. Estimation obtained using the linear model on a dataset with 1,000 sources with a focused activity. (A) After only the Kalman filter has been applied, the neural activity estimate is not accurate, and one can recognize two peaks, which 'imprint' the information given by the EEG and fMRI measure (the second could not be placed correctly in time since the filter encounters the fMRI measures too late). (B) After the Kalman smoother has been applied, this 'fMRI-informed peak' has been put back at the correct time location.

Second, we want to stress that the adaptation we propose for the extended Kalman filter, where the backward sweep is performed before the forward sweep, might be well-adapted for many other situations where adaptive filters are used on a nonlinear forward model, but where the measure of the phenomenon of interest are delayed in time. We already explained above that the order between these two sweeps matters because accumulating errors cause divergence, therefore it is preferable that the first sweep will be as exact as possible. Note that this is due exclusively to the errors made in the local linearizations performed by the extended Klaman filter; on the contrary, the order of the sweeps does not matter in the case of a linear model, and all the derivations for the means and variances of $p(x_k|y_1,...,y_l)$ are exact. To illustrate this fact, we show in figure 11 the neural activity estimated in the previous part by the linear Kalman filter and smoother: although the Kalman filter estimate is not accurate, it somehow "memorizes" the information brought late by fMRI data (figure 10(A)); therefore, this information is later available to the Kalman smoother, which can then "bring it back to the good place" (figure 10(B)).

3.3.2 Future improvements

In the previous part of this chapter we had proven that adaptive filters can be used to solve the EEG-fMRI fusion inverse problem on a dataset of a realistic size, but we had to assume a linear forward model of the measures. Here we have shown that this linear assumption can be overcome, and that more physiologically accurate nonlinear models can be used and new algorithms designed that handle these nonlinearities, although at a significantly increased computational cost.

First, Kalman methods can probably be improved to reduce the computational cost, as we already mentioned above (for example, by using low-dimensional descriptions of the variance matrices). But it is probably even preferable to develop novel algorithms that can

specifically deal with the strong nonlinearities of the model. For example, (Plis et al., 2010) proposed a particle filter algorithm (Arulampalam et al., 2002) to estimate the dynamics of a given region of interest based on EEG and fMRI measures. Unfortunately, the fact that they reduced the spatial dimension to a single ROI reduced the interest of their method, since the fMRI could not bring any additional information compared to the EEG, and they admit that increasing the spatial dimension would pose important computational cost problems, as the number of particles used for the estimation should increase exponentially with the number of regions. However, particle filters seem to be a good direction of research, all the more since they do not assume specific Gaussian distributions, and hence – as we saw in the previous part – could be more efficient in estimating sparse and focused activities.

Besides, there are interesting questions to ask about the dimension of the hidden neural activity. On the one hand, we would like it to have a temporal resolution as good as that of the EEG, and a spatial resolution as good as that of the fMRI. On the other hand, this leads to a very high total dimension for u, which is the main cause of high computational cost, and it is tempting to rather decrease this dimension. Several works have proposed approaches based on regions of interests (Daunizeau et al., 2007; Ou et al., 2010), where the spatial dimension is reduced by clustering the sources according to anatomical or physiological considerations. It is possible however that new approaches will be able in the same time to keep high temporal and spatial resolution, and decrease the total dimension of the source space, by using multi-level, or any other complex description of the neural activity relying on a reduced number of parameters.

4. Conclusion

We have proposed two algorithms to perform EEG-fMRI fusion, both based on the Kalman filter and smoother. Both algorithms aim at estimating the spatio-temporal patterns of a hidden neural activity u responsible for both EEG and fMRI measurements, the relation between them being described by a physiological forward model. The first algorithm assumes a linear model and in such case, fusion appeared to be feasible on data of a realistic size (namely, activity of 2,000 sources spread on the cortical surface, sampled at 200Hz during several minutes). The second algorithm relies on a more accurate nonlinear model, and though its increased computational cost led to estimations on only 100 sources, it gives a proof of principle for using adaptive filters to perform EEG-fMRI fusion based on realistic nonlinear models.

We would like to stress however that these progresses on the methodological front should not hide the fundamental physiological questions regarding whether EEG and fMRI measure the same phenomena at all, as we discussed in our introduction, and as recent reviews (Rosa et al., 2010) pinpoint to still be a problematic question. Nevertheless, if the question should be answered negatively *in general*, it remains that many simultaneous EEG-fMRI studies proved to be successful in specific contexts such as epilepsy or α-rythms activity. In such case, the development of new fusion algorithms which target the estimation of the neural activity spatio-temporal patterns should focus on such specific applications and on the specific links which exist between the two measures in theses contexts. This is what we plan to do in our future works with real data: in fact, our nonlinear algorithm, where we assumed a very specific oscillatory activity and a linear relation between the fMRI BOLD signals and the envelope of this activity is a preliminary for such development.

As we mentioned in the introduction, this work is not only a contribution to the human brain imaging field, but also to the algorithmic field, since a new modification of the Kalman

filter was proposed that filters first backward, and second forward, and since it fosters the development of new algorithms and filters. In particular, we call for new methods that can handle nonlinear models and non-Gaussian distributions. We advocate in particular the minization of L_1-norm rather than L_2-norm in order to estimate accurately sparse and focused activities, and the reduction of the computational cost of either Kalman methods or particle filter methods through the use of low-dimensional representations of high-dimensional distributions.

5. References

Ahlfors SP, Simpson GV. (2004). Geometrical interpretation of fMRI-guided MEG/EEG inverse estimates, *Neuroimage* 1:323–332

Arulampalam, M. Sanjeev; Maskell, Simon; Neil Gordon & Clapp, Tim. (2002). A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking *IEEE transactions on signal processing*, 50, 174-188

Bénar, C.; Aghakhani, Y.; Wang, Y.; Izenberg, A.; Al-Asmi, A.; Dubeau, F. & Gotman, J. (2003). Quality of EEG in simultaneous EEG-fMRI for epilepsy *Clin Neurophysiol*, 114, 569-580T

Buxton, R. B.; UludaNg, K.; Dubowitz, D. J. & Liu, T. (2004). Modelling the hemodynamic response to brain activation *NeuroImage*, 23, 220 - 233

Daunizeau, J.; Grova, C.; Marrelec, G.; Mattout, J.; Jbabdi, S.; Pélégrini-Issac, M.; Lina, J. M. & Benali, H. (2007). Symmetrical event-related EEG/fMRI information fusion in a variational Bayesian framework *Neuroimage.*, 36, 69 - 87

Daunizeau J, Laufs H, Friston KJ. (2010). EEG-fMRI information fusion: Biophysics and data analysis, In: *EEG-fMRI-Physiology, Technique and Applications*, Mulert L (eds.) Springer DE

Deneux, T. & Faugeras, O. (2006a). Using nonlinear models in fMRI data analysis: model selection and activation detection *Neuroimage*, 32, 1669-1689

Deneux, T. & Faugeras, O. (2006b). EEG-fMRI fusion of non-triggered data using Kalman filtering *International Symposium on Biomedical Imaging*, 2006, 1068-1071

Deneux, T. & Faugeras, O. (2010). EEG-fMRI Fusion of Paradigm-Free Activity Using Kalman Filtering *Neural Comput.*, 22, 906 – 948

Friston, K. J.; Mechelli, A.; Turner, R. & Price, C. J. (2000). Nonlinear Responses in fMRI : the Balloon Model, Volterra Kernels, and Other Hemodynamics, *NeuroImage*, 12, 466-477

Goldman, R. I.; Stern, J. M.; Engel, J. J. & Cohen, M. S. (2002). Simultaneous EEG and fMRI of the alpha rhythm, *Neuroreport*, 13, 2487-2492

Gotman, J.; Bénar, C. & Dubeau, F. (2004). Combining EEG and fMRI in epilepsy: Methodological challenges and clinical results, *Journal of Clinical Neurophysiology*, 21

Grova C, Daunizeau J et al. (2008). Assessing the concordance between distributed EEG source localization and simultaneous EEG-fMRI studies of epileptic spikes, *Neuroimage* 39:755–774

Hämäläinen, M.; Hari, R.; Ilmoniemi, R. J.; Kunuutila, J. & Lounasmaa, O. V. (1993). Magnetoencephalography: Theory, instrumentation, and applications to noninvasive studies of the working human brain, *Rev. Modern Phys.*, 65, 413-497

Johnston, L. A., Duff, E., Mareels, I., and Egan, G. F. (2008). Nonlinear estimation of the bold signal. *Neuroimage* 40, 504–514.

Jun, S.-C., George, J. S., Pare-Blagoev, J., Plis, S. M., Ranken, D. M., Schmidt, D. M., and Wood, C. C. (2005). Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data. *Neuroimage* 28, 84–98.

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems, *Transactions of the ASME--Journal of Basic Engineering*, 82, 35-45

Kilner, J. M.; Mattout, J.; Henson, R. & Friston, K. J. (2005). Hemodynamic correlates of EEG: a heuristic, *Neuroimage*, 28, 280-286

Laufs, H.; Kleinschmidt, A.; Beyerle, A.; Eger, E.; Salek-Haddadi, A.; Preibisch, C. & Krakow, K. (2003). EEG-correlated fMRI of human alpha activity, *Neuroimage*, 19, 1463-76

Lemieux, L.; Krakow, K. & Fish, D. R. (2001). Comparison of spike-triggered functional MRI BOLD activation and EEG dipole model localization, *Neuroimage*, 14, 1097-1104

Martinez-Montes, E.; Valdes-Sosa, P. A.; Miwakeichi, F.; Goldman, R. I. & Cohen, M. S. (2004). Concurrent EEG/fMRI analysis by multiway Partial Least Squares, *Neuroimage*, 22, 1023-1034

Murray, L., and Storkey, A. (2008). Continuous time particle filtering for fMRI, In: *Advances in Neural Information Processing Systems*, Vol. 20, eds J. C. Platt, D. Koller, Y. Singer, and S. Roweis (Cambridge, MA: MIT Press), 1049–1056.

Ogawa, S.; Menon, R. S.; Tank, D. W.; Kim, S.; Merkle, H.; Ellerman, J. M. & Ugurbil, K. (1993). Function brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging: a comparison of signal characteristics with a biophysical model, *Biophys. J.*, 64, 803-812

Ou, W.; Nummenmaa, A.; Ahveninen, J. & Belliveau, J. W. (2010). Multimodal functional imaging using fMRI-informed regional EEG/MEG source estimation *Neuroimage.*, 52, 97-108

Plis, S. M., Ranken, D. M., Schmidt, D. M., and Wood, C. C. (2005). Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data. *Neuroimage* 28, 84–98.

Riera, J.; Watanabe, J.; Kazuki, I.; Naoki, M.; Aubert, E.; Ozaki, T. & Kawashima, R. (2004). A state-space model of the hemodynamic approach: nonlinear filtering of BOLD signals, *NeuroImage*, 21, 547-567

Riera J, Sumiyoshi A. (2010). Brain oscillations: Ideal scenery to understand the neurovascular coupling, *Curr Op Neurobiol* 23:374–381, 2010.

Rosa, M. J.; Daunizeau, J. & Friston, K. J. (2010). EEG-fMRI integration: A critical review of biophysical modeling and data analysis approaches, *J Integr Neurosci.*, 2010, 9, 453 - 476

Waites, A. B., Shaw, M. E., Briellmann, R. S., Labate, A., Abbott, D. F., & Jackson, G. D. (2005). How reliable are fMRI-EEG studies of epilepsy? A nonparametric approach to analysis validation and optimization. *Neuroimage*, 24(1), 192–199.

Welch, G. & Bishop, G. (2006). *An Introduction to the Kalman Filter*

Welling, M., Max Welling's classnotes in machine learning. (N.d.). Available online at http://www.cs.toronto.edu/~welling/classnotes/papers_class/KF.ps.gz.

Adaptive Filtering

Edited by Dr Lino Garcia

Adaptive filtering is useful in any application where the signals or the modeled system vary over time. The configuration of the system and, in particular, the position where the adaptive processor is placed generate different areas or application fields such as prediction, system identification and modeling, equalization, cancellation of interference, etc., which are very important in many disciplines such as control systems, communications, signal processing, acoustics, voice, sound and image, etc. The book consists of noise and echo cancellation, medical applications, communications systems and others hardly joined by their heterogeneity. Each application is a case study with rigor that shows weakness/strength of the method used, assesses its suitability and suggests new forms and areas of use. The problems are becoming increasingly complex and applications must be adapted to solve them. The adaptive filters have proven to be useful in these environments of multiple input/output, variant-time behaviors, and long and complex transfer functions effectively, but fundamentally they still have to evolve. This book is a demonstration of this and a small illustration of everything that is to come.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Thomas Deneux (2011). EEG-fMRI Fusion: Adaptations of the Kalman Filter for Solving a High-Dimensional Spatio-Temporal Inverse Problem, Adaptive Filtering, Dr Lino Garcia (Ed.), ISBN: 978-953-307-158-9, InTech, Available from: http://www.intechopen.com/books/adaptive-filtering/eeg-fmri-fusion-adaptations-of-the-kalman-filter-for-solving-a-high-dimensional-spatio-temporal-inve

INTECH
open science | open minds