Chapter 5

Article 1: Modelling of Post-surgical Brain and Skull Defects in the EEG Inverse Problem with the Boundary Element Method

Context

The possibility of describing the sources of EEG potentials as a set of current dipoles is not recent. Traditionally, the volume conductor into which the currents propagate was modelled as a set of concentric spheres, which permitted an analytic solution. More recently, the advent of fast computers allowed modelling arbitrary closed surfaces by meshes of triangles (“boundary element method”, or BEM), thereby giving a better representation of the shape of the head; this is the “realistic head modelling”.

This is a topic in EEG where anatomical MRI images can be of great help. Meshes can be obtained that permit to delineate the boundaries between brain and skull or skull and skin. Also, the MRI images can help defining the shape of the cortex, which can be used in turn to restrict the sources in the distributed sources technique.

In this study, we will concentrate on a particular aspect of head modelling, namely the case of patients with skull and brain defects resulting from brain surgery. Such situation arises when a patient is not seizure free after a first operation and is considered for a second intervention. This is not an uncommon situation, as approximately 20-30 % of patients are not completely seizure free after surgery.

This article was published as [Bénar 01].

Abstract

Objective: In order to obtain accurate EEG inverse solutions in patients subjected to surgery, we have studied the feasibility and influence of incorporating brain and skull defects in realistic head models.
Methods: We first measured the conductivity of the methacrylate used for cranioplasty. Then, we designed realistic Boundary Element Method head models with a skull burr hole, a methacrylate plug or a temporal lobe resection. We simulated the potentials that would be produced at 71 electrode locations (10/10 system) by dipoles located near the defects. Then, we fitted dipoles on these potentials using a defect-free head model. We also ran simulations in a noisy situation and with higher skull and CSF conductivity .
Results: The largest errors were found for burr holes, with a localization error up to 20 mm for a radial dipole located 30 mm below the hole and an amplification factor of 8. Methacrylate plugs lead to errors up to 5 mm and 0.5; the resection only lead to errors of 2 mm and 1.3. Results obtained with noise were consistent with those obtained without noise. Doubling the skull conductivity led to errors that were reduced by 10 to 20 % while doubling CSF conductivity increased the errors by up to 31 %.
Conclusions: We have shown that it is important to incorporate skull defects in realistic head models when sources are located near the defects and precision is sought. Brain cavities of the size of a typical anterior temporal lobe resection may be omitted without a significant impact on dipole localization.

5.1 Introduction

Techniques solving the inverse problem of electroencephalography can be used in an attempt to localize generators of brain activity such as evoked potentials or epileptic spikes. These methods use the theory of electromagnetic fields in order to calculate the potential distribution that is generated at the surface of a volume conductor by internal current sources. They need a model of neuronal activity as well as a mathematical representation of the head that takes into account the different tissue conductivities. Indeed, the electrical activity of a limited area of cortex can be represented as that of a current dipole inside a volume conductor [Nunez 81]. A classical head model comprises a set of concentric spheres representing the brain, skull and skin interfaces [Rush 69, Kavanagh 78]. A more recent approach makes use of a “realistic head model” consisting of meshes of discreet elements based on real anatomical data. In the boundary element method (BEM) [Barnard 67, Hämäläinen 89, Meijs 89], only boundary surfaces are tessellated, while in the finite element method (FEM) [Yan 91] the whole head volume is meshed. The FEM can represent more heterogeneous configurations than the BEM and can take into account anisotropy, but is more demanding computationally.

Patients who have undergone brain surgery challenge head modelling techniques, presenting skull and brain defects that affect the conduction of electrical currents. However, source reconstruction may be needed for patients who are not seizure-free after the surgical removal of an epileptogenic lesion. We therefore propose in this study to investigate the feasibility and effectiveness of accounting for post-surgical defects in dipole modelling. We will limit ourselves to the BEM.

5.2 Epilepsy Surgery and its Effects on the EEG

In a standard procedure of brain surgery, burr holes of approximately 10 mm in diameter are drilled into the skull as a first step in the removal of the bone flap. During surgery, parts of the brain are resected, leaving a space that will be filled later with cerebrospinal fluid (CSF) (for a description of temporal lobectomy, see [Olivier 87]). At the end of the surgical procedure, burr holes are filled with a methacrylate-based paste, except one that is left open for drainage. In the period following the operation, the open burr hole fills naturally with soft tissue, which creates a path of least electrical resistance through the skull. We address here three types of defects: open burr holes, methacrylate plugs and brain resections. Whereas the resistivity of CSF and soft tissue is known, that of the methacrylate-based substance used in brain surgery has not to our knowledge been published. We present here a test of the electrical properties of this substance.

The fact that skull defects have a large effect on the electroencephalogram (EEG) has been known from the beginning of electroencephalography [Berger 33]. Cobb et al. used the word “breach rhythm” for the mu-like activity that can arise close to skull holes; they also reported an enhancement of the amplitude of alpha and frontal fast rhythms “over or near unilateral posterior and frontal defects respectively” [Cobb 79]. Thevenet et al.

studied the effect of a large hole in the skull on the localization of tangential dipoles in a 3-sphere FEM model [Thevenet 92]. They placed tangential dipoles in a sphere with a hole and fitted dipoles on the resulting potential without taking the hole into account. They found that the recovered dipoles were closer to the hole, and that "the closer the dipole is from the hole (or turn towards it), the more important is the error". Cuffin has done a similar type of study, with local variations in the thickness of the skull and scalp layers [Cuffin 93]. In the case of a depression in the skull-scalp surface, he found a localization error of about 5 mm and an amplification factor of about 1.2 for a radial dipole 10 mm below the skull. Van den Broek et al. modelled holes of about 20 mm in diameter in a FEM realistic head model [van den Broek 98]. They quantified the effect of the holes on the scalp potential with the relative-difference measure (RDM). They found the largest influence for a radial dipole located just below the hole, which resulted in an RDM of 600 %. It is also interesting to note that they found an influence of hole location. Ollikainen et al. studied the influence on dipole models of an inhomogeneity with a radius of 5 mm and small conductivity [Ollikainen 99]. They generated scalp potentials in the model with the inhomogeneity, added noise to the simulated potentials and fitted a dipole onto the result. For 64 electrodes, they found a mean localization error of about 3.5 mm with the non-uniform model and about 10 mm with the uniform model. The error in magnitude was 9 % and 19 % respectively. Van Burik and Peters injected currents in intracerebral electrodes in order to create artificial current dipoles [van Burik 00a]. Dipole models were then fitted on the data, with a fixed location and orientation corresponding to the true ones. Different head models were used, with or without the two trephine holes of diameter 23 mm and their bone plugs. They found that potential maps were greatly distorted by holes, even for deep sources. Vanrumste et al. studied the impact of omitting a hole of diameter 20 mm on dipole localization [Vanrumste 00]. They used simulations and the finite difference method, and found errors up to 5.2 mm in a 53-electrode setting.

The influence of CSF cavities or zones of low conductivity like methacrylate plugs are less documented. Van den Broek et al. found that inclusion of the ventricles in the head model had a large effect on the RDM for dipoles within a few centimetres of these cavities [van den Broek 98]. Also, dipoles located behind a zone of very low conductivity inside the brain could present a RDM as low as 10 %. Vanrumste et al. found that neglecting the ventricular system could lead to errors of up to 6.1 mm for dipoles in its vicinity [Vanrumste 00]. We can infer from the results presented above that post-operative defects are likely to have an influence on dipole modelling for sources located near the defects. It is however not totally clear how large is the brain area influenced by a typical burr hole of a diameter of 10 mm. There is also a need in our clinically oriented approach to study the influence of burr holes filled with methacrylate or of an extended brain resection. We will first present our measurements on the resistivity of methacrylate, then the main study.

5.3 Measurement of Methacrylate Resistivity

5.3.1 Method

We built a $4.9 \times 4.9 \times 0.2$ cm plate of methacrylate (Howmedica International Inc., Ireland) and placed it in the center of a $9.8 \times 4.9 \times 4.9$ cm container, thereby dividing it in two equal parts. The edges of the plate were sealed using a silicone gel. The two parts of the container were filled with a solution of artificial CSF (aCSF) in order to simulate physiological conditions. One circular electrode of diameter 2.5 cm was placed at each end of the container. We then measured the resistance and capacity of the system, as well as those of the container with aCSF alone. The meter used was a Hewlett Packard 4261A LCR, set at a frequency of 1kHz. The resistivity ρ_{acryl} (Ω·cm) of the plate was calculated as:

$$\rho_{acryl} = \frac{R_{meas} \cdot Surf}{l} \tag{5.1}$$

with R_{meas} (Ω) the measured resistance, Surf (cm^2) the surface of the immersed plate, l (cm) the plate thickness. The plate was far enough from the electrodes so as to consider the potential uniform on its immersed surface; aCSF resistivity was neglected.

5.3.2 Results

The meter reached its saturation level, which means that the actual resistance was over 20 MΩ. This corresponds to a resistivity of more than 2500 MΩ·cm, or a conductivity of less than $4 \cdot 10^{-8}$ S·m^{-1}. This is several orders of magnitude higher than skull resistivity, which is in the range of 1 kΩ·cm to 20 kΩ·cm [Law 93]. The measured resistance of aCSF alone was 79 Ω, which confirmed that it is negligible compared to that of methacrylate.

5.4 Main study: Methods and Materials

5.4.1 Visualization of defects

A burr hole that has been left open after surgery produces a high intensity signal characteristic of soft tissues on T1-weighted magnetic resonance imaging (MRI). A methacrylate plug is more difficult to discriminate with MRI, because it leads to a low intensity signal that can be confused with that of skull. However, it produces a rupture in the bone marrow high intensity line when that is visible. Also, some assumptions can be made on the skull defects locations based for example on the skin flap, which is visible on a 3D skin reconstruction. Another helpful way to track skull defects is to perform a curvilinear reconstruction [Bastos 99] of the skull. Of course, skull defects are also very visible on computed tomography (CT) scans or plain X-Ray films, but these are not always available. See figure 5.1 for an overview of the burr holes seen by MRI. A resected brain area is clearly seen with MRI: it is filled with CSF and is of low intensity on T1-weighted images (Fig. 5.2).

5.4.2 Defect-free model

All the realistic BEM models in our study were created with Curry V4 (Neuroscan Labs). They were based on the T1-weighted MRI scan of a patient presenting skull defects and a temporal-lobe resection. We used Curry's automatic segmentation of scalp, skull and brain boundaries, based on gray level thresholds. We modified manually these boundaries in order to build a model that does not take into account the defects. The lengths of the triangle sides were 12 mm, 10 mm and 8 mm for scalp, skull and brain meshes respectively. This produced meshes with 1742, 1668 and 2062 triangles, leading to a model with 2779 nodes. Enclosed conductivity was 0.33 $S \cdot m^{-1}$, 0.0042 $S \cdot m^{-1}$, 0.33 $S \cdot m^{-1}$ for scalp, skull and brain respectively [Geddes 67].

5.4.3 Modelling of defects

We then modified the defect-free model in order to account for skull and brain defects. It is important to note that the BEM is limited in that it can only model closed surfaces. Also, there is a generally accepted "rule of thumb" stating that two BEM surfaces should not be closer than half their mesh triangle size: this implies that a local mesh refinement is needed when surfaces are too close to one another.

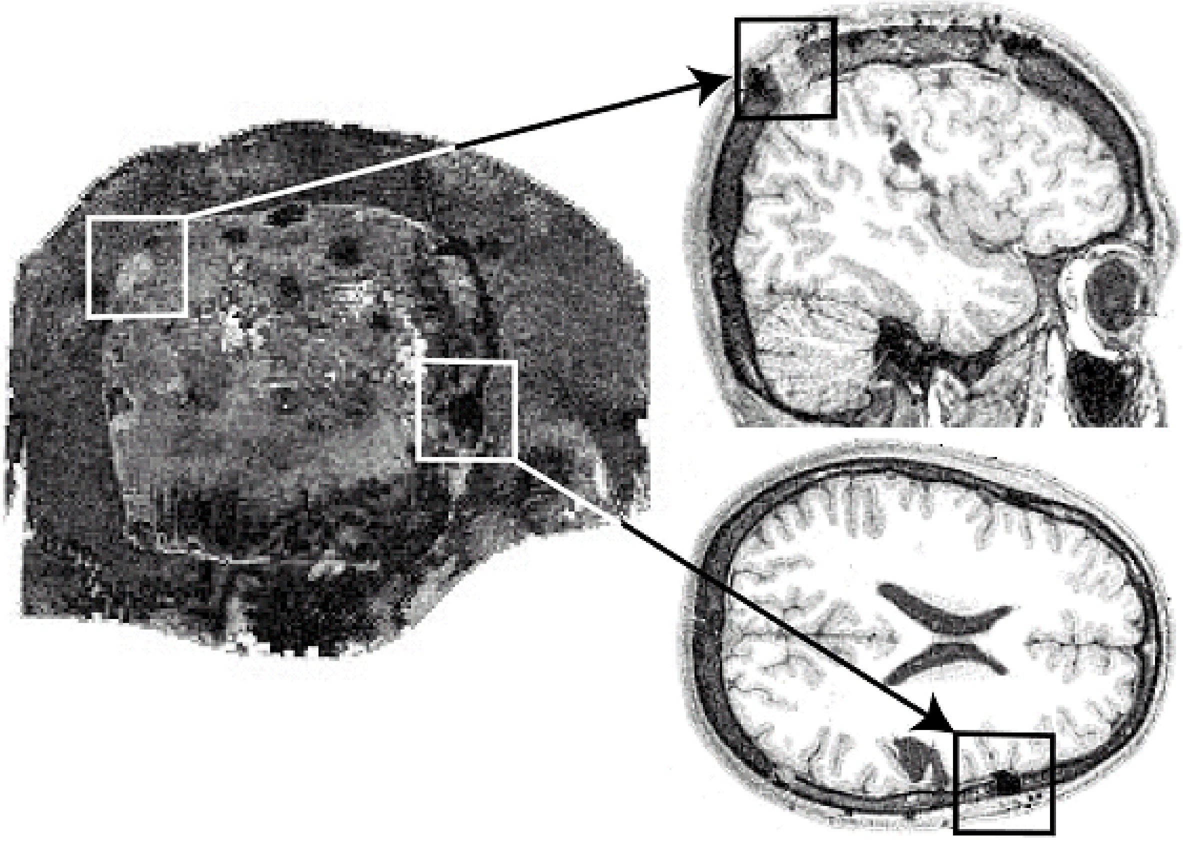

Fig. 5.1 Burr holes as seen with T1-weighted MRI: (left) reconstruction of an image along a curved surface following the skull (Brainsight software); (upper right) open burr hole; (lower right) burr hole that has been filled with methacrylate.

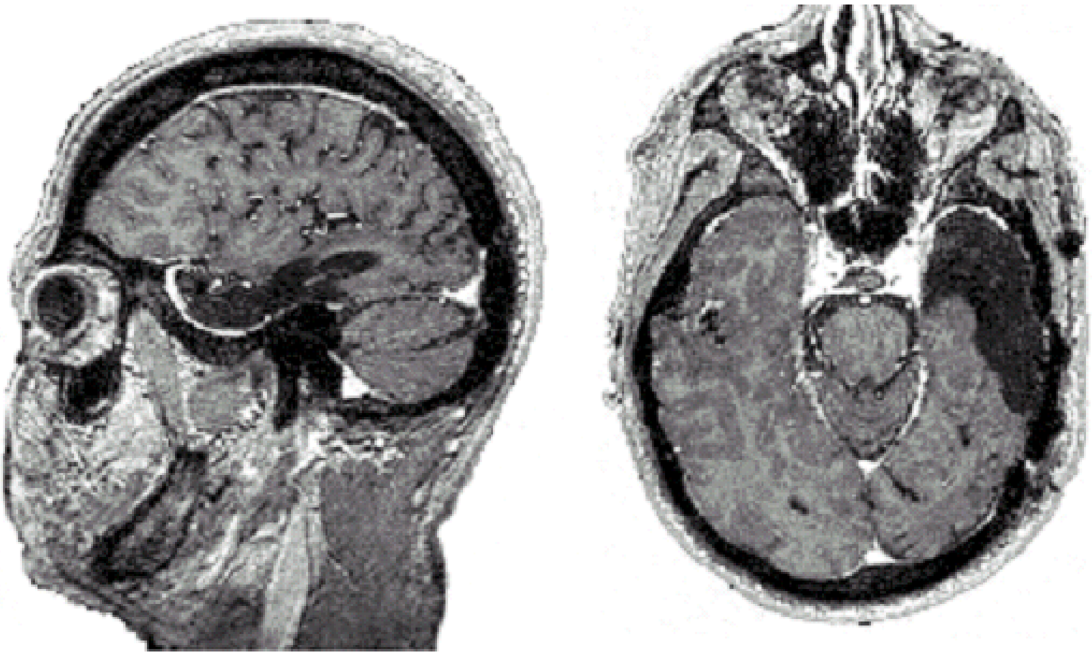

Fig. 5.2 Sagittal (left) and transverse (right) views of a temporal brain resection seen with T1-weighted MRI.

In the first category of burr hole model, we joined the brain and skull BEM surfaces at the level of the hole. We obtained one BEM surface only, which enclosed the remaining intact skull. The hole was placed at the vertex of the head in order to have a simple geometry in terms of skull conformation and electrode positions around the hole. This was the "open burr hole" model (see fig. 5.3(a) for a schematic representation of an equivalent spherical model, fig. 5.4(a) for a view of the mesh). We also created a model with a hole at the location of the real burr hole in our reference patient in order to quantify the influence of the location of the hole. This was the "parietal burr hole" model. In a second category of models, the hole was represented by a depression in the skull boundary. This type of model is a classical 3-layer model that was used as a reference in the validation of the less typical full hole model. Care was taken that skull and brain boundaries did not get closer that 1 mm. This was insured by expanding the brain BEM surface with morphological tools in order to delineate the base of the depression. The skull thickness was approximately 7 mm in the region of the depression, the depth of which was 6 mm as a consequence. This was the "partial burr hole" model (Fig. 5.3(b)). For all models, the skull BEM mesh was generated with a 10mm triangle size and a 1.5 mm refinement in the area of the hole. Brain and scalp meshes were also refined around the defect.

The methacrylate plug was modelled as a small cylinder with near zero conductivity (10^{-10} S·m^{-1}). It was placed at the vertex of the head, between the skull and scalp surface (Fig. 5.3(c)). In a similar way as for the open burr hole, we used morphological tools to depress the scalp surface by 1 mm and expand the brain surface by 1mm in order to delineate the top and bottom surfaces of the cylinder, the height of which was consequently around 5 mm. We used a 3 mm mesh. The temporal lobe resection was modelled with 6mm-sided triangles, producing a surface enclosing 30 ml of CSF (conductivity of 1 S·m^{-1}) (Fig. 5.3(d) and 5.4(b)).

5.4.4 Quantifying the influence of the defects

We placed a set of radial dipoles in the vicinity of each defect: below and lateral to the burr holes (Fig. 5.5(a)), behind and above the resection (Fig. 5.5(b)). For each dipole, we simulated the potentials that it would produce at 71 locations (10/10 system) on the surface of the realistic head model including the defect (see above "modelling of defects"). Then, a dipole was fitted on these potentials in the defect-free head model, using Curry's simplex method and constraining the sources to be at least 3 mm away from the innermost BEM surface. We measured the resulting error in location and amplitude.

In order to validate our hole models, we also reproduced the experimental settings of [van den Broek 98] and [Thevenet 92]. These studies used the FEM and were therefore not subject to some of the limitations of the BEM that are discussed below. In the first setting, we used a model with a full hole at a 60° elevation and dipoles on the z-axis (90° elevation). We computed the relative difference measure (RDM) [van den Broek 98] between the potentials computed with and without the hole. In the second setting, the hole was central and tangential dipoles were placed at different elevations; we simulated tangential dipoles in a full hole model and localized them without the hole. In the two experiments, meshes were refined around the hole and the dipoles.

5.4.5 Influence of noise on localization error

The methods we presented in the previous sections assume a noise-free environment. A more realistic configuration needs to take noise into account. We wanted to test a) if the noise-free localization error is a good measure of the mean error in a noisy situation and

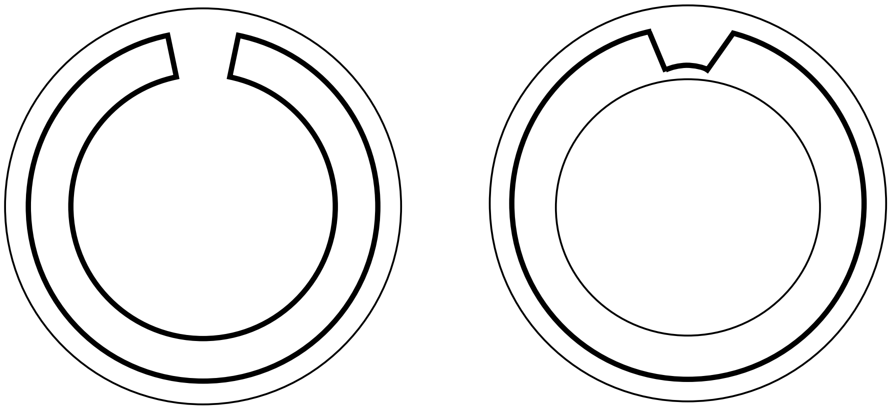

(a) Full hole. (b) Partial hole.

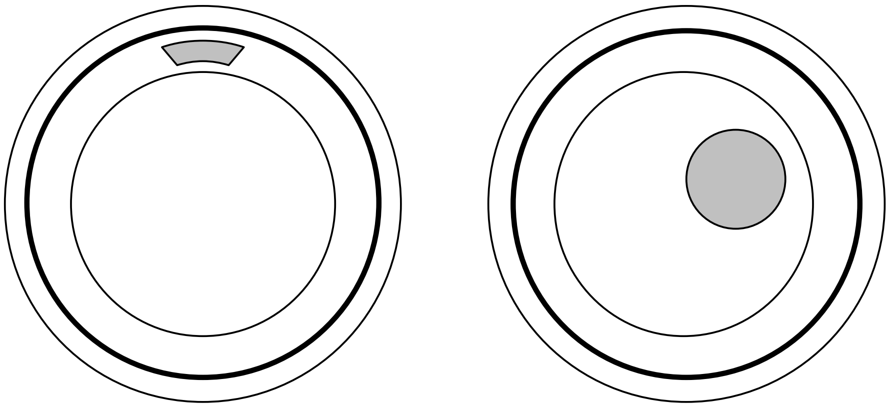

(c) Methacrylic plug (d) Temporal resection.

Fig. 5.3 Schematic views of the different types of models. The outer circle represents the skin surface, the thick line represents the skull , the inner circle the brain.

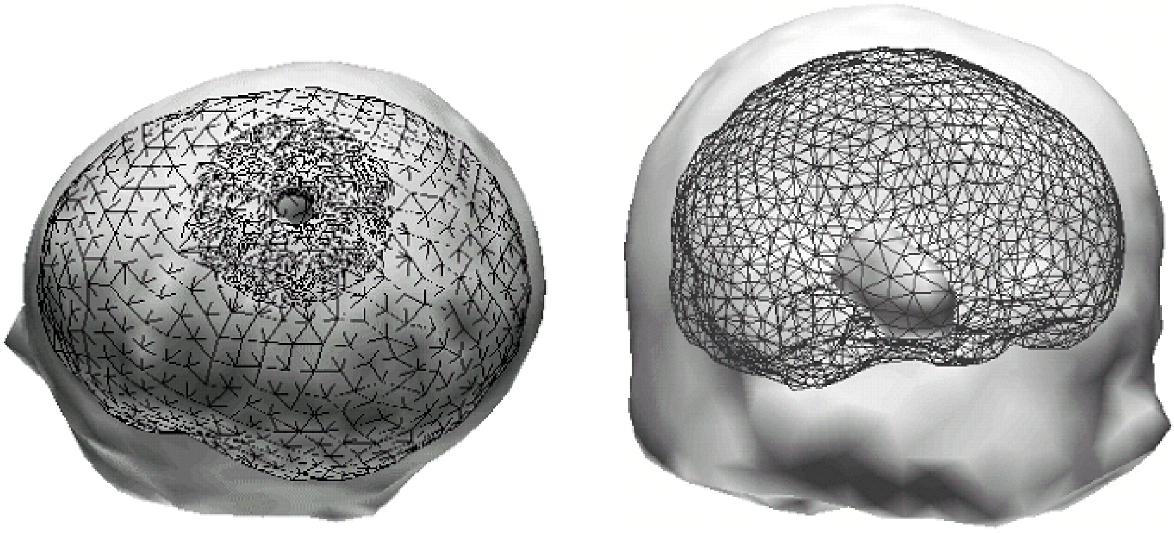

(a) Central hole.

(b) Temporal resection.

Fig. 5.4 BEM models.

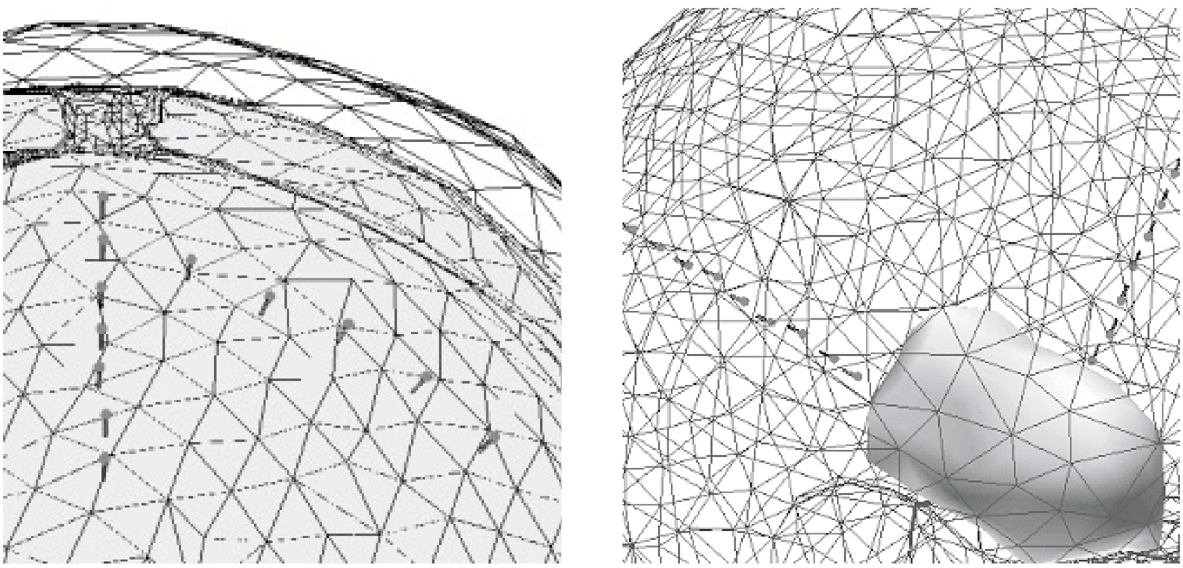

(a) Dipoles below and lateral to hole.

(b) Dipoles behind temporal resection.

Fig. 5.5

b) if the variance of the localization error is similar when using either defect or defect-free model. We used the potentials created by a dipole located 20 mm below the skull in the central partial hole model. We added to the signal at each electrode one realization of a gaussian noise process. The amplitude of the noise was chosen so as to produce a signal to noise ratio (SNR) of 10. Then a dipole was fitted on these potentials with the defect-free head model, as in section 5.4.3. We repeated the simulation 100 times, and computed the mean and standard deviation of the localization error, in both defect and defect-free configurations.

5.4.6 Use of different conductivities

Some authors have suggested that skull conductivity may be higher than the classical value of 1/80 times the scalp conductivity [van Burik 00b, Cuffin 01b]. We performed again the burr hole simulations with a skull conductivity of 1/40 times the scalp value in order to assess the influence of skull conductivity on the reconstruction error. The CSF conductivity at body temperature may also be higher than the classical value of 1 $S{\cdot}m^{-1}$: Baumann et al. recommend a value of 1.79 $S{\cdot}m^{-1}$ [Baumann 97]. We ran again the resection simulations with a value of 2 $S{\cdot}m^{-1}$.

5.5 Results

5.5.1 Open burr hole

Radial dipoles below the full hole were all reconstructed closer to the defect. For an original dipole between 5 mm and 20 mm below the hole, the fitted dipole was found as close to the hole as was permitted by our settings. For dipoles between 25 mm and 35 mm deep, the location error was of the order of 20 mm. The relative amplitude reached a maximum for dipoles just below the hole: 8 for a dipole 5 mm deep. It decreased exponentially with depth, and went below 1.5 at a depth of 30 mm. See Fig. 5.6(a) for an overview of the results. Radial dipoles lateral to the full hole were reconstructed closer both to the hole and to the surface. Location errors were of the order of 6 mm to 15 mm. The relative amplitude went below 1.5 at a distance of 20 mm from the hole (Fig. 5.6(b)).

The use of a partial central hole model lead to results with the same trend, but with much less extent. Indeed, the maximum location error was around 7 mm for a dipole 20

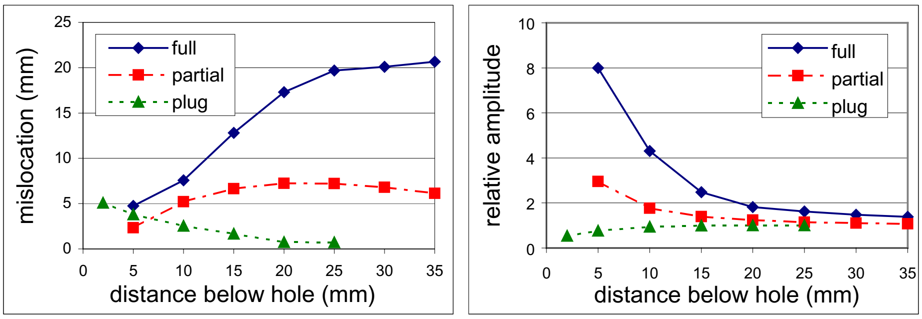

(a) Radial dipoles below central burr hole.

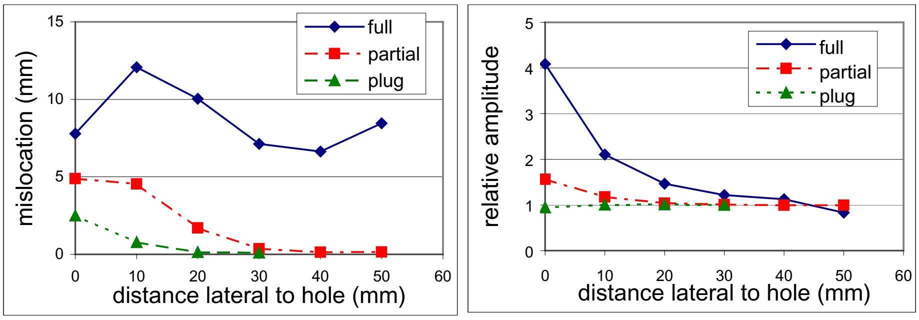

(b) Radial dipoles lateral to central burr hole.

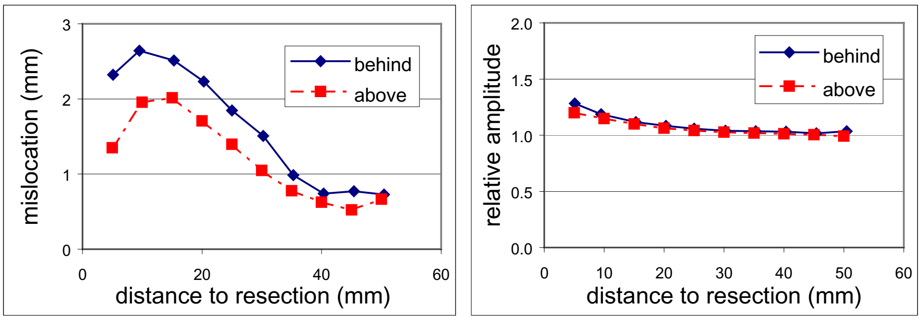

(c) Dipoles behind and above temporal resection.

Fig. 5.6 Charts of errors in localization and amplitude mm deep and the relative amplitude was of the order of 3 for a dipole 5 mm deep (Fig. 5.6(a) and 5.6(b)). In the partial parietal hole case, the amplitude factor was very similar to the central hole results, ranging from 1.1 to 2.7. However, the effect on location error was somewhat smaller: errors ranged from 3.5 mm to 5.5 mm, with a peak when the dipole is at a depth of 15 mm only (results not shown).

In the two sets of simulations similar to [van den Broek 98] and [Thevenet 92], the results were in most cases very close to the corresponding FEM results obtained by these authors (Fig. 5.7 and 8 respectively.). In the first study, though, the RDM values for the radial dipoles close to the surface was much higher than in [van den Broek 98]: for the dipole at 3 mm below the hole, we obtained 180 % instead of 40 % (Fig. 7b). We also noticed in the second study a tendency for dipoles to be reconstructed closer to the surface than in [Thevenet 92] (Fig. 8); this was strongly reduced by a local refinement of the mesh around the dipoles.

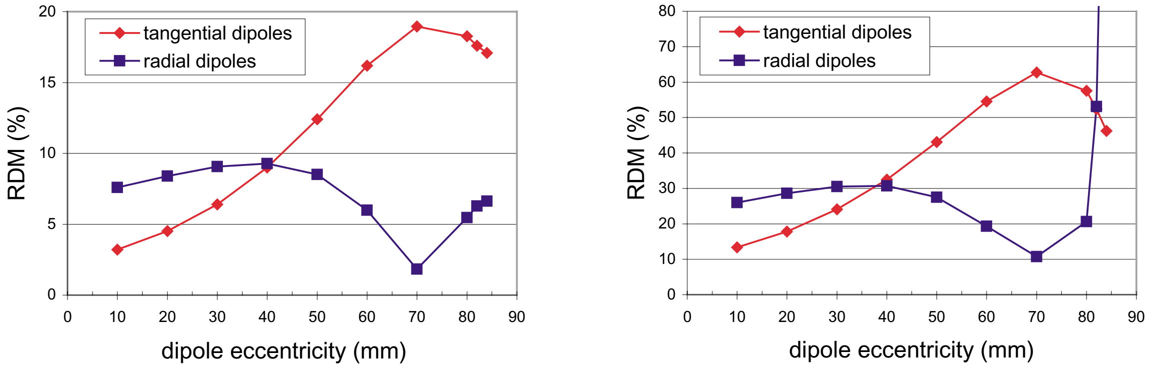

(a) Partial hole.

(b) Full hole (the value at eccentricity 84 mm is 181 %, not shown).

Fig. 5.7 Relative difference measure (RDM) between potentials in the central hole models and the defect-free model for dipoles at given locations.

5.5.2 Burr hole filled with methacrylate

In the case of the methacrylate plug model, radial dipoles were reconstructed further from the holes than their original location, with a location error and relative amplitude of 5 mm and 0.5 for a dipole 3 mm below the plug (Fig. 5.6(a)). For the dipoles lateral to the defect,

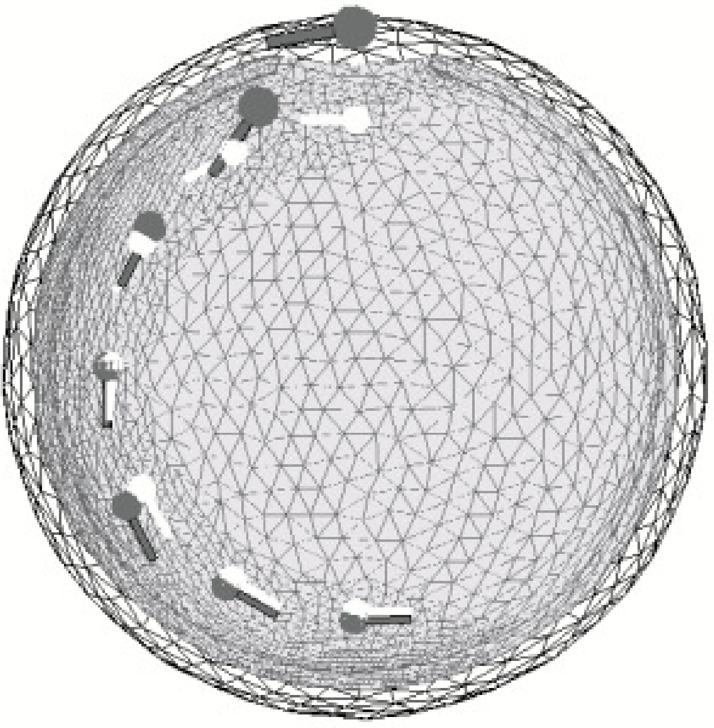

Fig. 5.8 Original dipoles in a full hole model (light grey) and dipoles reconstructed with the defect-free model (dark grey)

the effect was significant (more than 3 mm of location error) only in a radius of about 10 mm around the center of the plug (Fig. 5.6(a)).

5.5.3 Brain resection

The resection had a small effect on the dipoles placed behind and above it: the maximum localization error was less than 3 mm and the relative amplitude less than 1.3 (Fig. 5.6(c)).

5.5.4 Simulations with noise

The mean and standard deviation of localization error for the 20 mm deep dipole simulated with the partial hole model and reconstructed using the same model were 2.7 ± 1.3 mm (see fig. 5.9 for a histogram of the location error). This corresponds to the effect of the noise alone. When the defect free model was used for localization, they were 7.2 ± 1 mm, reflecting the effect of the partial hole in the context of noise (to be compared to the value of 7.24 mm in a noiseless situation, cf. fig. 5.6(a)). We have to note the presence of outliers that probably correspond to instabilities in our models and were not included in our statistics (Fig 5.9).

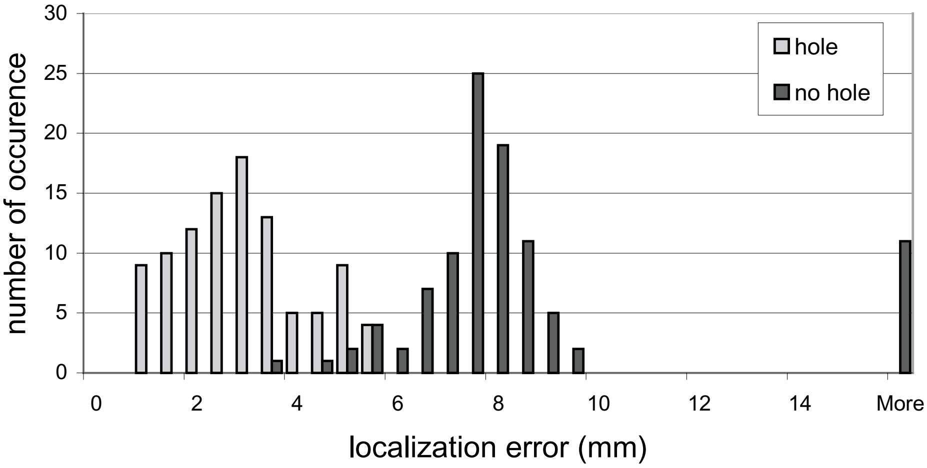

Fig. 5.9 Histograms of localization error for dipoles simulated in the partial hole model and reconstructed in a noisy environment with a model that includes the hole (light bars) or not (dark bars). Note the outliers at the right of the graph when the hole is not included.

5.5.5 Use of different conductivity values

When full hole models with higher skull conductivity were used, we obtained location errors of 78 ± 17 % and a relative amplitude of 73 ± 9 % of the ones with a classical conductivity. In the partial hole models, the ratios were 88 ± 7 % and 94 ± 5 % respectively. Errors caused by burr holes were therefore smaller in the context of higher bone conductivity.

In the resection simulations, the use of a CSF conductivity of 2 $\mathrm{S{\cdot}m^{-1}}$ lead to a maximum localization error of 3.5 mm for a dipole 10 mm behind the resection and a maximum amplitude factor of 1.4 for a dipole 5 mm behind the resection. This corresponds respectively to 31 % and 8 % more than the results for the classical value.

5.6 Discussion

We have shown in this paper that postsurgical burr holes have a significant effect on EEG and dipole localization, in concordance with previous studies [Thevenet 92, van den Broek 98, Vanrumste 00]. If holes were not included, large errors in location and amplitude were obtained when reconstructing sources located near the defects, as expected. A quantitative evaluation of the impact of omitting the full hole lead to values generally similar to the ones of previous FEM studies [Thevenet 92, van den Broek 98], with some important variations though. Indeed, the impact of omitting the hole was much larger than in [van den Broek 98] for dipoles very close to the surface. This is probably due to the typical instability for dipoles close to BEM surfaces, and could possibly be reduced by an even finer mesh. Also, dipoles very far from the hole are reconstructed closer to the surface, contrary to [Thevenet 92]. This could be due to the well-known numerical error resulting from a low-conductivity skull layer in the BEM, which cannot be minimized here by the isolated problem approach (IPA), because this latter requires concentric layers (see for ex. [Meijs 89]). This effect seems to be more pronounced in the lower part of the sphere - probably because we are dealing with smaller values that are more affected by numerical errors. Nevertheless, it is important to note that BEM results were close to the FEM results for dipoles in the regions where currents sources are likely to be located (more than 5mm away from the skull and in the middle and upper parts of the sphere). The fact that errors were reduced by a higher mesh refinement emphasizes again the importance of refining the meshes in the region of the dipoles. We found an influence of the hole location, parietal or central, as in [van den Broek 98], but it was quite limited.

We demonstrated that the more numerous methacrylic plugs also had an influence on dipole localization, suggesting that they should be taken into account when sources are near the plug and precision is sought. The situation was quite different for the temporal CSF-filled cavity: it produced little effect on dipole modelling despite its large size and can therefore be omitted without much consequence. This is in discordance with [van den Broek 98] and [Vanrumste 00] who had both studied the impact of the CSF-filled ventricular system. This discrepancy may be due to the relative position of the dipoles, the CSF cavity and the electrodes. Indeed, in [Vanrumste 00] the dipoles that are more affected by the cavity are the ones located beneath it, i.e. the ones for which the cavity stands between the dipole and a large number of electrodes. In our simulations the cavity lies on the side of the dipoles and possibly produce less interference on the measured potentials, as it has an important effect on a small number of electrodes.

We studied the localization errors in a noisy environment, and verified that the mean dipole was close to the one reconstructed in a noise-free setting. This suggests that we can use the error value encountered in a noise-free environment as a measure of a “mean” localization error that would be obtained with noise. The error variance was similar in the defect and defect-free situations. We have to note the presence of instabilities in the fitted dipoles in the noisy situation (outliers). This suggests again that one has to be particularly cautious when using complex BEM models. We ran again the noise-free simulations studies in the context of doubled skull and CSF conductivity and found only a small impact on the resulting errors.

All our models comprised 71 electrodes, which represent a high spatial sampling. In a clinical setting, the lower number of electrodes is typically of the order of 30. This will almost certainly lead to localization errors due to spatial aliasing on top of the errors coming from not considering the defects. As the potential gradients are stronger in the case of a skull hole (i.e. higher spatial frequency content), the aliasing is expected to be stronger, emphasizing the need for high spatial sampling around the hole [Bénar 01].

We used BEM models in all our EEG studies, which can only handle closed surfaces that are not in contact with one another. This is a limitation in modelling defects, for example with methacrylate plugs or multiple holes. Some investigation should therefore be conducted in order to assess the advantage in our context of using the more computationally demanding FEM, which does not have this limitation, and in the context of magnetoencephalography, which is not significantly influenced by the skull. Nevertheless, we have found the largest effect for the open burr hole, which can be easily modelled when unique (the most common case). Therefore, EEG modelling using BEM remains a valuable tool in the study of post-operative epileptic patients.

Acknowledgments

We would like to thank the reviewers for their very useful comments, as well as Dr. Aviva Abosh for her input on epilepsy surgery and Dr. André Leduc for his help in measuring methacrylate conductivity. This research was supported in part by grant MT-10189 of the Canadian Institutes of Health Research and in part by a Natural Science and Engineering Research Council studentship.

Chapter 6

Article 2: Quality of EEG in Simultaneous EEG-fMRI for Epilepsy

Context

In the preceding article, we have discussed the possibility of greatly improving the inverse problem of EEG by incorporating anatomical information in the computation of the potentials created by a given source. Still, even with a very fine modelling of the volume conductor and of the cortex geometry, the inverse problem remains ill-defined. Moreover, the EEG is very little sensitive to sources that are deep or with complex source geometry, such as hippocampus, amygdala and thalamus - structures that play an important role in epilepsy.

The recent technique of fMRI permits to acquire images in the whole head volume, and is based on a well-defined reconstruction problem. As the epileptic activity is usually spontaneous, it is only the advent of systems for recording the EEG inside the scanner that allowed the study of spikes with fMRI.

The following article concerns the technique of simultaneous recording of EEG and functional MR images. In particular, we shall examine whether one can obtain an EEG of sufficient quality for visualizing the epileptic spikes.

This article was published as [Bénar 02a]. The set-up of the recording equipment is described in Annex A.

Abstract

It is now possible to record the EEG continuously during fMRI studies. This is a very promising methodology that combines knowledge about neuronal activity and its metabolic response. The EEG recorded inside the fMRI scanner is, however, heavily contaminated by artefacts caused by the high intensity magnetic field and rapidly changing field gradients. Methods have been reported in the literature to reduce or eliminate these artefacts, in particular the ballistocardiogram and the artefact caused by currents induced by rapidly changing magnetic gradients. Nevertheless, recording the EEG simultaneously with fMRI remains an extremely delicate operation. In addition the use of artefact-removal methods has only been reported by the laboratories in which they were developed. We report here the practical procedures we developed to reduce artefacts in a series of 10 epileptic patients, in the context of the visualization of epileptic spikes. We illustrate the effectiveness of methods designed to remove the scanning artefact and present new methods for removing the ballistocardiographic artefact.

6.1 Introduction

The possibility of recording the EEG inside an MR scanner was first reported by Ives et al. [Ives 93], and later by Huang-Hellinger et al. [Huang-Hellinger 95] and Lemieux et al. [Lemieux 97]. These reports established the safety of recording a good quality EEG inside the scanner, as well as the possibility to obtain high quality MR images despite the presence of EEG electrodes and equipment. These findings opened the way to functional MR imaging studies of spontaneous EEG events such as alpha activity or epileptic spikes.

The MR scanner is a very difficult environment in which to record EEG. Any movement of the electrode wires inside the large static magnetic field or any variation of the field around the wires will induce currents that manifest as an artefact in the EEG. A common movement artefact is the ballistocardiogram, caused by the slight motion of the head that occurs after each heart beat [Ives 93, Allen 98]. An example of a field variation artefact is the very large signal caused by gradient switching during fMRI image acquisition, which renders the EEG apparently non-interpretable.

Subtraction techniques have been presented to eliminate the ballistocardiogram or reduce its effect on EEG quality [Allen 98, Goldman 00]. Another approach in removing this artefact involves the application of a spatial filter optimizing the output signal to noise ratio [Bonmassar 99].

The first method used to circumvent the gradient artefact relied on reading the undistorted EEG between image acquisitions. For example, the “spike-triggered” fMRI paradigm takes advantage of the lagging nature of the BOLD hemodynamic response. In this technique, one whole-head EPI frame is acquired a few seconds after each epileptic spike [Warach 96, Seeck 98, Krakow 99]. It has recently become clear, however, that it is possible to remove this large artefact and see the EEG during an EPI sequence, thus allowing a continuous fMRI recording with visualization of the EEG. The continuous fMRI recording allows the collection of several images for each spike, thereby increasing statistical power. It also made possible the plotting of the actual BOLD response to epileptic events [Lemieux 01b, Bénar 02b]. However, continuous recording does require amplifiers with special characteristics, in particular a very large dynamic range, and artefact removal techniques. One such technique is the averaging of the artefact followed by subtraction [Allen 98], whereas the other one is based on the Fourier transformation and relies on the basic rhythmic nature of the EPI sequence [Hoffmann 00].

The field of combined EEG and fMRI studies is still very new, and the practical results obtained using the above techniques outside the laboratory where they were developed have hardly been reported. Baudewig et al. reported a single case where they performed a “semi-continuous” recording [Baudewig 01]. They scanned during one half-second every two seconds in a patient with bursts of spike-and-wave, and relied on reading the EEG between scans rather than removing scanner artefact to see the EEG under the artefact. We describe here our results with several of the techniques aimed at reducing or eliminating artefacts, and present important practical considerations in obtaining a good EEG recording during continuous fMRI studies. We also introduce spatial filtering schemes based on Principal Component Analysis (PCA) and Independent Component Analysis (ICA), for removing the ballistocardiogram. We finally present, for the removal of the gradient artefact, a variation of the subtraction technique that does not rely on a triggering signal and operates at a 1kHz sampling rate only, which we compare to the Fourier filtering method of [Hoffmann 00].

6.2 Subjects and Methods

6.2.1 Subjects

We studied 10 patients. Patients were selected on the basis of having an active interictal epileptiform discharge (one where we could expect at least 10 to 20 spikes per hour). Consent was obtained in accordance with regulations of the Research Ethics Board of our institution.

6.2.2 MRI recording

A Siemens 1.5 T Magnetom Vision whole body scanner was used. We acquired frames of 25 BOLD mosaic 64×64 images, with a voxel size of $5 \times 5 \times 5$ mm, a TE of 50 ms and an inter-frame interval of 3 s. The acquisition of the 25 images took 2.5 seconds. There was therefore a gap of 0.5 seconds between series of images.

6.2.3 EEG recording

We used the EMR32 amplifier system (Schwarzer GmbH, Munich, Germany). The amplifier is specially designed for use within the MR scanner. It allows recording of 32 channels with a common reference. This reference is the average of two electrodes (as often used for the average between ears). We selected two reference electrodes remote from the subject's main epileptic spikes. We typically used the F3-F4 electrode pair for subjects having temporal or posterior discharges and the O1-O2 pair for frontal activity. We recorded from the standard 19 electrodes of the 10/20 system, with Ag/AgCl electrodes, as well as from two ECG electrodes.

Long electrode wires (3 metres) were used, thus allowing the amplifier to be outside the bore of the magnet. We immobilized the wires by two procedures: placing a tight bandage on the subject's head and installing heavy sand bags over the wires between the subject's head and the amplifier (inside the magnet behind the patient's head and outside the magnet up to the amplifier).

EEGs were digitized at 1000 Hz, filtered with a pass band filter between 0.1 and 70 Hz, and transmitted to a recording computer outside the magnet room via an optic cable. The linear range of amplification was ±3.2 mV.

The subject's head, as well as most of the electrode wires near the head, were immobilized with a plastic bag (50 × 70 cm, 10 L fill) filled with very small polystyrene spheres, in which a vacuum was obtained by air suction (S & S X-Ray products, Brooklyn, NY, USA). This proved extremely effective in preventing head movement. It also rendered the procedure more comfortable, removing the strain of lying directly on electrodes. Subjects stayed in the magnet for approximately 90 minutes and seldom complained of discomfort.

Electrode wires were set on the patient's head to be as tight as possible and without any loops, arranged in as direct a line as possible from the electrode position to the amplifier outside of the magnet. Since we started with electrode wires of equal length, some loops were inevitable at the amplifier end, but this is less important since the amplifier is outside the magnet. A more convenient alternative would be to use an electrode cap. However such a system could not be adjusted to all head sizes as well as the system we used.

6.2.4 Testing of EEG setting

We ran a series of tests of the EEG recording setup on one healthy subject (table 6.1). The reference configuration was our routine setup, against which we tested the other configurations. First, we assessed the influence of each major element (sand bags, vacuum in head cushion, head bandage) by removing that element. Tests were performed in a sequence such that only one element at a time was altered (i.e. B, A, C and D). Second, we recorded the EEG of the subject in a room remote from the scanning room during a separate session (test E). For each test (A to E), we recorded at least 2 minutes of background EEG. The subject was asked to keep his eyes closed during each recording.

For each test, we measured the average power of the EEG traces in each frequency band. This was done by splitting 100 s of EEG into segments of 2.048 s (i.e. 2048 samples). We performed a fast Fourier transform (FFT) on each segment and computed the power at each frequency bin as twice the mean FFT square amplitude. No windowing was applied. For a given frequency band, we averaged the power of all bins of this band. Bands were defined as: Delta from 1 to 4 Hz, Theta from 4 to 8 Hz, Alpha from 8 to 13 Hz, Beta from 13 to 35 Hz. All signal processing in this paper - apart from the Fourier filtering method - was done with the Matlab software (Mathworks, Natick, MA, USA), on a Pentium IV processor.

Test	Description		
	Sandbags	Vacuum in cushion	Bandage
A	Yes	Yes	Yes
B	No	Yes	Yes
C	Yes	No	Yes
D	Yes	Yes	No
E	outside of scanner		

Table 6.1 Summary of the tests performed to assess the influence of various parameters on quality of the EEG in the scanner

6.2.5 Removal of ballistocardiogram

Methods that have been used to remove the ballistographic artefact rely on the identification of several artefacts in order to average them, and then subtract the average from the signal [Allen 98]. We present two alternative methods consisting of building a spatial filter, based on either Principal Component Analysis (PCA) or Independent Component Analysis (ICA) [Comon 94, Bell 95]. This type of filtering has already been applied in removing EEG artefacts such as ocular artefact (Lagerlund et al. 1997 with PCA and Jung et al. 2000 with ICA). We chose a reference page on which we performed PCA and ICA.

The PCA was obtained by Singular Value Decomposition on raw data, i.e.:

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \tag{6.1}$$

with $\mathbf{X}$ the data matrix for one page (n time points $\times$ m channels), $\mathbf{U}(n \times m)$ the matrix of normalized time courses of the PCA components (one per column), $\mathbf{S}(m \times m)$ the diagonal matrix of eigenvalues or amplitudes of the components, $\mathbf{V}(m \times m)$ the matrix of spatial distributions of each components (one per column).

The ICA was computed with the Matlab toolbox from the Computational Neurobiology Laboratory (Salk Institute, La Jolla, CA, USA), and resulted in the following decomposition:

$$\mathbf{X}^T = \mathbf{W}^{-1}\mathbf{U}, \tag{6.2}$$

with $\mathbf{W}^{-1}(m \times m)$ the matrix of spatial distributions of the component (one per row) ($\mathbf{W}$ is usually referred as the "unmixing" matrix) and $\mathbf{U}(m \times n)$ the matrix of component time courses (one per row).

For each method, we visually selected the components with a time course corresponding to ballistocardiographic activity, using the ECG as a guide. We constructed two spatial filters, one for PCA and one for ICA, by excluding the components corresponding to ballistocardiogram in the spatial matrices. This resulted in

$$\mathbf{A} = \mathbf{V}\mathbf{I}_0\mathbf{V}^T \tag{6.3}$$

for the PCA filter and

$$\mathbf{A} = \mathbf{W}^T\mathbf{I}_0(\mathbf{W}^T)^{-1} \tag{6.4}$$

for the ICA filter, with $\mathbf{I}_0$ diagonal matrix with elements set to zero when the component was retained and zero when the component was excluded. Once computed on the reference page, these fixed filters could then be applied to any page in the recording by:

$$\mathbf{Y}' = \mathbf{Y}\mathbf{A} \tag{6.5}$$

with $\mathbf{Y}$ matrix of original data and $\mathbf{Y}'$ filtered data.

In order to test the spatial filters, we selected for each patient two 10 s pages of EEG with clear ballistocardiographic artefact and without any other artefact nor any epileptic event. We used the first page to build the filters, as indicated above. The second page was used to test ballistocardiogram removal. In this page, we marked the peaks of each artefact and compared their amplitudes before and after filtering in the channel with maximum artefact. In addition, we repeated this procedure on several spikes outside the gradient artefact in order to measure the impact of the spatial filters on epileptic activity.

6.2.6 Removal of gradient artefact

The first method to remove the artefact from the EPI sequence was based on the technique of [Hoffmann 00], as implemented in the FEMR program provided by Schwarzer. In summary, the method consists of selecting a few 10-second pages of EEG that were recorded inside the magnet but without scanning artefact, and using them as baseline. Ideally, these pages should contain epileptic activity, which is the signal we want to preserve, and no movement-related artefact. The frequency spectrum of this baseline is then obtained. In order to filter a page with scanning artefact, the FFT of that 10 s page is computed and frequencies for which the FFT amplitude differs from the baseline spectrum by more than a certain factor (typically 2 or 3) are noted. Presumably these frequencies correspond to the fundamental and harmonics of the artefact. Their corresponding Fourier coefficients are set to zero and an inverse Fourier transformation is obtained, yielding a signal from which the artefact should be removed.

As a second method, we implemented a subtraction filter that consisted in building a model of the artefact and subtracting it from each frame. Because the artefact onset is not synchronous with EEG sampling, each repetition of the artefact is sampled at slightly different time points with respect to EEG samples: for instance, the EPI frame could start one time 0.1 ms before the nearest EEG sample and the next time 0.4 ms before the nearest EEG sample (there is 1 ms between EEG samples). We assume here that the artefact is identical from frame to frame, as the EPI sequence is fixed and the wires immobilized. In a typical recording, we have of the order of 1000 frames. Therefore, one can divide the sampling period in 10 bins of 0.1 ms and obtain around 100 realizations of the artefact for each bin. All realizations for each bin can then be averaged. This results in an estimator of the artefact sampled at 10 kHz.

To implement the above algorithm, we found the rough location of each frame by setting a threshold on the EEG traces. This allowed for the division of the recording into *blocks*, each containing one frame only. It should be noted that each frame was not positioned at the exact same location inside the blocks. We used the first frame of the first run as a reference, and found the location of the following frames inside their corresponding blocks by correlating each block with the reference. We interpolated the reference and each frame by a factor of 10 to refine our measure. Frames with the same lag with respect to the reference were grouped and averaged. We finally subtracted for each frame the average artefact corresponding to its bin.

In order to compare the Fourier filter to the subtraction filter, EEG traces obtained on epileptic patients were reviewed by a trained electroencephalographer, and the epileptic discharges marked. The two methods (Fourier filtering and subtraction) were used in parallel. For the Fourier method, the threshold on the ratio between current page and background spectrum was tuned to two or three. For each event within a gradient artefact, the reviewer assessed the quality of filtering by each method with three marks: 1, if the event seen with the given filter was clearly an epileptiform discharge; 2, if the event was ambiguous; 3, if the event was almost not visible or seemed clearly non epileptiform (such a situation could only occur when an event appeared clearly with one filter but not with the other). The rate of spiking outside and within the gradient artefact were also compared to assess if the action of filtering in general enabled a fair recovery of the expected epileptic activity.

6.3 Results

6.3.1 EEG setting

Figure 6.1 presents pages of EEG obtained inside the scanner during Test *A* (i.e. inside the scanner with all restraints on, Fig. 6.1(a)) and Test *E* (i.e. outside the scanner, Fig. 6.1(b)), after low pass filtering at a cut-off frequency of 45 Hz. Alpha activity can be clearly seen on both recordings. Inside the scanner (Fig. 6.1(a)), all channels were affected by the scanning environment, but at different degrees - the occipital and midline channels being the least corrupted on visual inspection. The results of average root power for the different tests are listed per frequency band in Table 6.2. We assume that the power of the signal (brain waves) remained approximately constant and that the variations originated mainly from variations in the noise level. By comparing test *E* (outside the scanner) with test *A* (inside the scanner with all restraints on), one can quantify the impact of the scanning environment on EEG quality. The delta band was most affected, with a ratio between root power measures of 0.76, followed by theta, with a ratio of 0.78. In distinction to the other bands, there was the same power in the alpha band in test *E* and test *A* (ratio of 1.00). This latter finding could reflect the large amount of alpha activity recorded in the session outside the scanner rather than equality in noise levels.

Test *B*, test *C* and test *D* measure in turn the impact of each restraining procedure. The variable that caused the main differences in EEG power was the presence or absence of sandbags. Considering all bands, the ratio between test *B* (no sandbag) and test *A* (sandbags on) was 1.29. This effect was more pronounced in the beta band (ratio of 2.02). The second most important variable was the presence or absence of a bandage: the ratio between test *D* and test *A* was 1.10. The presence or absence of vacuum in the head bag came third (ratio of 1.06 between test *C* and test *A*). All these ratios correspond to

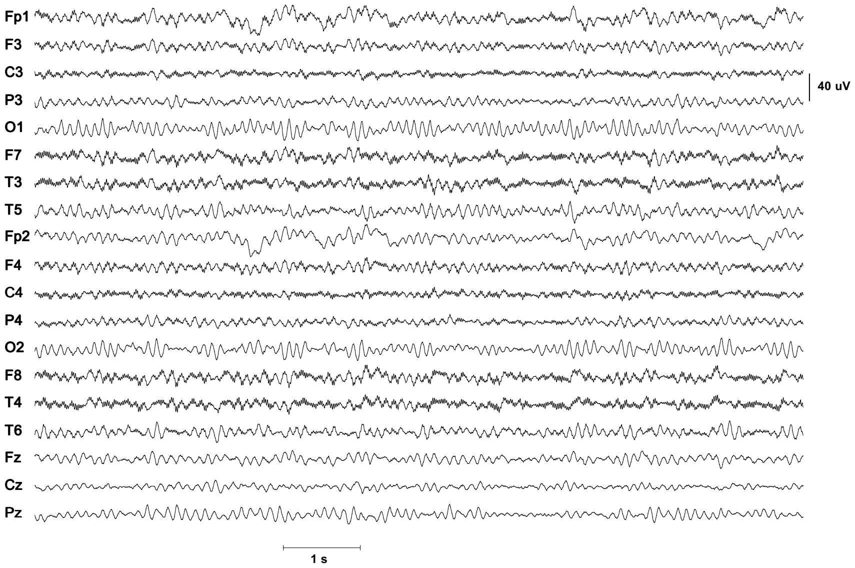

(a) One 10 s page of EEG during Test *A* (i.e. inside the scanner with all restraints on).

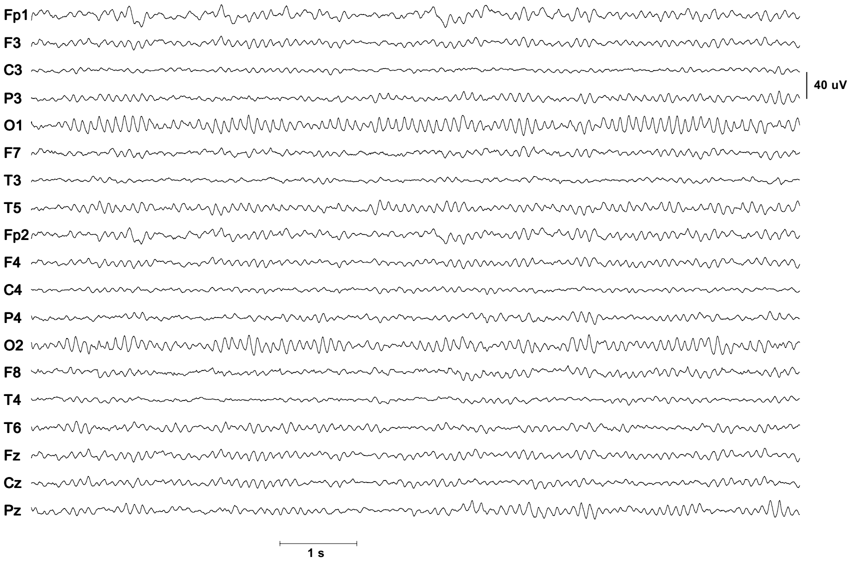

(b) One 10s page of EEG during Test *E* (i.e. outside the scanner).

Fig. 6.1 Examples of EEG obtained during the tests of recording settings. Low pass filtering at a cut-off frequency of 45 Hz was applied to both pages. The reference was the average of F3 and F4.

differences of mean log power that are significant at p=0.05 in a Student t-test (the log transform is needed to render the data more Gaussian).

Test	Root Power (V) (ratio to test A)				
	Delta	Theta	Alpha	Beta	All bands
A	1.40 – 0.20 (1.00)	1.63 – 0.39 (1.00)	2.46 – 0.27 (1.00)	0.46 – 0.03 (1.00)	1.25 – 0.14 (1.00)
B	1.57 – 0.21 (1.12*)	2.14 – 0.60 (1.31*)	2.94 – 0.34 (1.20*)	0.92 – 0.04 (2.02*)	1.62 – 0.15 (1.29*)
C	1.44 – 0.16 (1.03)	1.86 – 0.43 (1.14)	2.62 – 0.28 (1.07)	0.49 – 0.04 (1.07*)	1.33 – 0.13 (1.06*)
D	1.43 – 0.19 (1.02)	1.75 – 0.54 (1.07)	2.77 – 0.28 (1.12)	0.51 – 0.04 (1.12*)	1.38 – 0.15 (1.10*)
E	1.06 – 0.13 (0.76*)	1.28 – 0.27 (0.78*)	2.45 – 0.22 (1.00)	0.44 – 0.04 (0.96*)	1.17 – 0.09 (0.94*)

Table 6.2 Root mean power (RMP) of the EEG during the different tests, for each frequency band. The ratio between RMP of each test and that of test *A* (reference setup) is shown inside brackets. A star indicates that the corresponding difference between measures of power is significant at p=0.05 (t-statistics are performed on log-transformed power). The level of noise for each test can be assessed by comparing the RMP of each band to test *A*.

6.3.2 Ballistocardiogram

Table 6.3, table 6.4 and table 6.5 summarize ballistocardiogram studies for all patients. When present, the average ballistocardiogram amplitude in the most affected channel ranged between 30.5 μV and 54.2 μV, with an average of 40.7 μV (table 6.4). Two cases (patients 1 and 8) had no visible artefact. In another case (patient 5), there was no page free of epileptic events, and we therefore had to choose a background page with both spikes and ballistocardiogram. Among all subjects, the ballistocardiogram was usually most prominent in the frontal and anterior temporal channels (table 6.3). The polarity was typically reversed between left and right channels, and the artefact often had biphasic morphology.

When performing PCA decomposition of the reference page, ballistocardiographic activity was usually clearly visible in one or two components. When there were two components, the best results in term of balance between removing the ballistocardiogram while preserving spikes was obtained by rebuilding the signal after removing one component only. For ICA, there was up to five components with clear ballistocardiographic activity (typically three or four), which could all be removed while preserving the spikes. See figure 6.2 for examples of PCA and ICA decomposition performed on the same EEG data from patient 3.

Up to twenty spikes were marked for each patient. In two cases (patients 4 and 7), the epileptic events consisted mainly of slow waves with very small spikes. The slow waves were then marked for analysis. For the other patients, up to four different types of spikes were analyzed.

Spatial filtering based on PCA reduced the artefact to an amplitude ranging from 4 % to 64 % of its original amplitude, with a mean of 32.9 % (table 6.4). The ratios for the spike amplitudes in the same patients ranged from 42.6 % to 108.5 % with an average of 86.4 %. When ICA was used, the relative amplitude of the filtered artefact ranged from 9.4 % to 46.1 % (average of 22.6 %), while the spikes amplitudes in filtered pages ranged from 46.5 % to 132.1 % (average of 96.5 %, table 6.5). In five cases, ICA filtering resulted unexpectedly in an increase of spike amplitude. In addition, in the patients with slow waves, there was a strong change of morphology in some channels. In patient 5, in whom there was no page free of epileptic events, ICA performed much better that PCA in leaving the spikes intact (relative amplitude of 99 % vs 43 %). This suggests that, in this case, ICA was more efficient in separating the spike from the ballistocardiogram.

Figure 6.3(a) shows one page of EEG with three epileptic spikes. There are clear ballistocardiographic artefacts, with moderate amplitudes. Figure 6.3(b) and 6.3(c) show the results of PCA and ICA filtering, respectively, on the same data. In this example, ICA filtering removed entirely the ballistocardiogram (on visual inspection), whereas PCA left remnant artefact in the right-hemisphere channels. In both cases, the spike was left intact. The PCA results are due to the fact that the component corresponding to right side ballistocardiographic activity was not removed in the corresponding filter. Removing this latter component filtered out all ballistocardiogram - as for ICA - but also removed most of the spikes in the process.

		Ballistocardiogram		IEDs		
Patient Number	Recording reference	Maximum Channel	Other channels of high amplitude	Type	Maximum channel	Other channels of high amplitude
1	(Fz+Cz)/2	none		1	T6	
				2	F7	T3, T5
2	(O1+O2)/2	T4	T3, F8, T6, T5, F7	1	F7	F3, Fp1, T3
				2	F7	F7, Fp1, F3
3	(O1+O2)/2	T3	T5, F8, T4	1	F3	F4, Fz
				2	F3	F4, Fz, Fp1, Fp2
				3	Fz	F3, Fp1
4	(O1+O2)/2	T3	F7, T5, F8, T4	1	F3	F7
5	(P4+O2)/2	Fp2	F4, F8, T4	1	P3	C3, Cz, Pz
				2	F7	Fp1, C3, P3
6	(O1+O2)/2	F8	T5, Fp2, T4, T3	1	F3	Fp1, Fz, Cz
7	(O1+O2)/2	T3	Fp1, F7, T4, T5, O1, F8	1	F4	F8, T4
8	(O1+O2)/2	none		1	F7	T3
9	(F3+F4)/2	F8	T3, T5, F7, T4, F4	1	O1	T5
				2	O1	O2,T6,T6
				3	O1	O2
				4	O2	T6
10	(F3+F4)/2	Fp2	F4, F8, T4, F3, C3, T3	1	O2	T6

Table 6.3 Ballistocardiogram removal results. Channels of maximum amplitude for ballistocardiogram and spikes.

		Ballistocardiogram			Spikes			
Patient number	Number of PCA components	Number of artifacts averaged	Average maximum amplitude (V)	Relative amplitude after filtering (%)	Spike Type	Number of spikes averaged	Average maximum amplitude (V)	Relative amplitude after filtering (%)
2	1	9	30.5	36.0	1	16	93.3	66.5
					2	4	56.3	104.4
3	1	11	37.0	4.0	1	5	31.8	95.3
					2	6	52.1	96.9
					3	5	35.2	97.1
4	1	11	54.2	15.7	1	10	83.1	84.3
5	1	13	42.8	40.6	1	16	59.2	42.6
					2	3	41.0	79.7
6	1	18	47.1	54.5	1	6	30.2	87.0
7	1	12	39.8	33.0	1	20	139.4	85.3
9	1	11	34.1	15.0	1	9	91.0	74.6
					2	3	83.4	79.8
					3	1	66.2	108.5
					4	1	71.0	100.6
10	1	9	40.1	64.0	1	6	41.0	93.7
Mean	1	11.8	40.7	32.9		7.4	64.9	86.4

Table 6.4 Ballistocardiogram removal results. Results of PCA filtering. Amplitudes are the absolute values of the EEG signal.

Patient number	Number of ICA components	Ballistocardiogram			Spikes			
		Number of artifacts averaged	Average maximum amplitude (V)	Relative amplitude after filtering (%)	Spike Type	Number of spikes averaged	Average maximum amplitude (V)	Relative amplitude after filtering (%)
2	3	9	29.5	46.1	1	16	93.3	68.1
					2	4	58.9	84.4
3	4	11	37.0	28.8	1	5	46.2	132.1
					2	6	59.2	119.6
					3	4	38.0	110.5
4	5	11	54.6	9.4	1	10	91.2	95.8
5	3	13	49.2	17.0	1	16	55.0	98.8
					2	3	37.7	46.5
6	4	18	47.1	11.7	1	5	31.1	91.9
7	3	12	39.8	25.5	1	20	135.9	86.3
9	3	11	34.5	14.1	1	8	97.2	103.5
					2	3	83.4	100.0
					3	1	69.6	119.9
					4	1	62.6	91.5
10	2	8	39.1	28.0	1	5	56.6	98.9
Mean	3.4	11.6	41.4	22.6		7.1	67.7	96.5

Table 6.5 Ballistocardiogram removal results. Results of ICA filtering. Amplitudes are the absolute values of the EEG signal. The amplitudes of events can differ slightly between PCA and ICA tables because these were obtained in two different sessions and were sometimes taken at different time points.

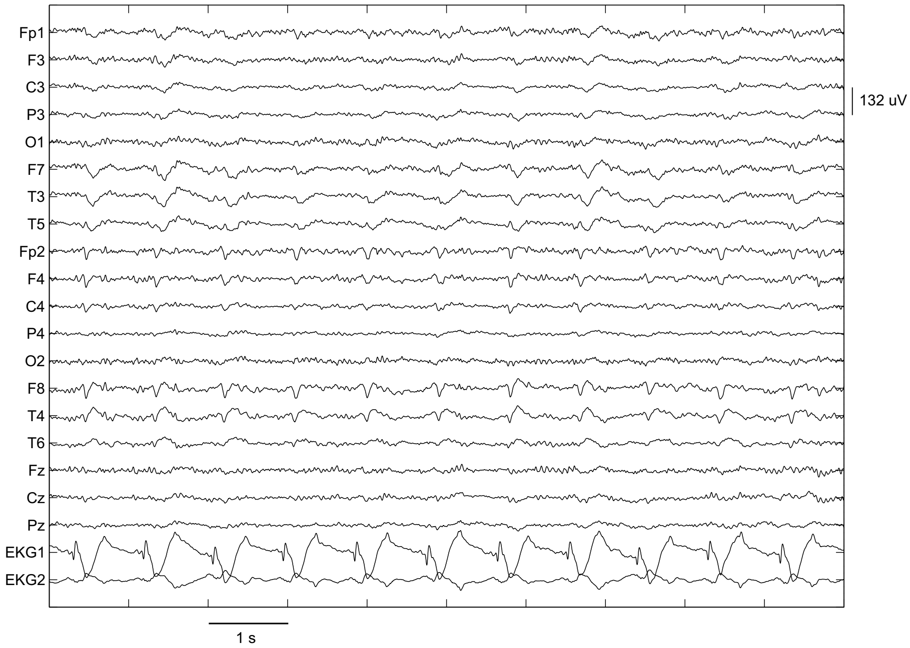

(a) Page used in building spatial filters for patient 3. Note the ballistocardiographic activity in channels Fp2, F4, F8, T4.

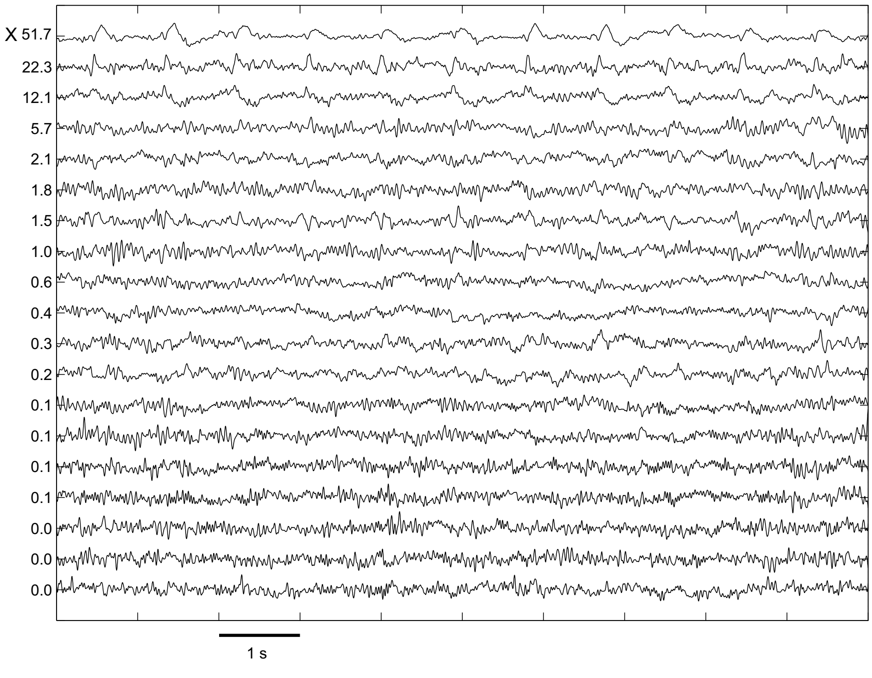

(b) PCA temporal components of this page.

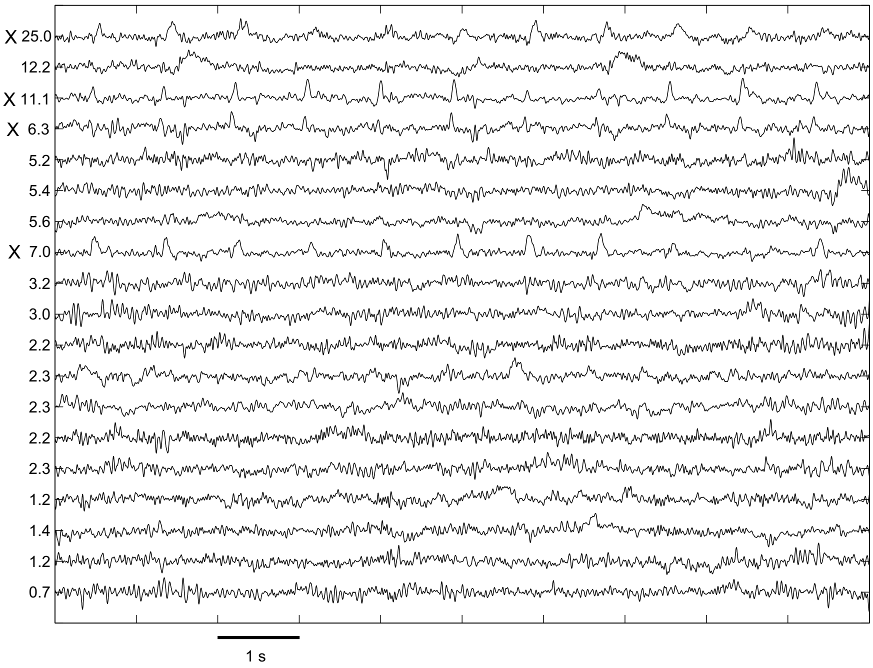

(c) ICA temporal components.

Fig. 6.2 Decomposition of the EEG into its components by spatial filtering. For each component, the ratio of projected variance versus total variance is displayed. A cross indicates that the component was chosen in the corresponding spatial filter. The second PCA component in (b) was not included as this resulted in significant spike distortion. Low pass filtering at a cutoff frequency of 35Hz was applied to the data prior to decomposition.

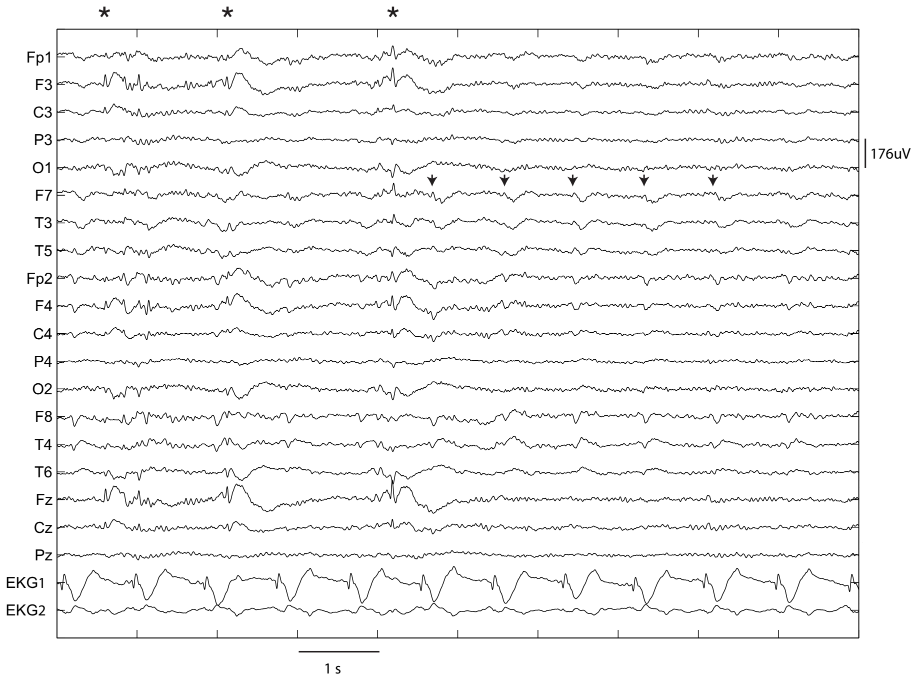

(a) Original recording. The timing of spikes is marked with stars. The timing of some ballistocardiographic artefacts is marked with arrowheads.

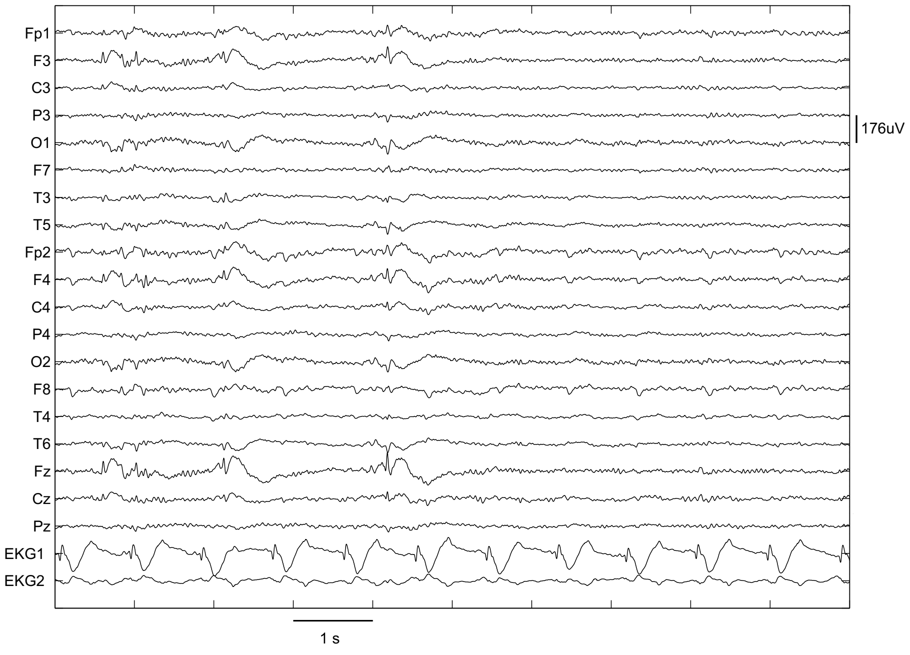

(b) PCA filtering.

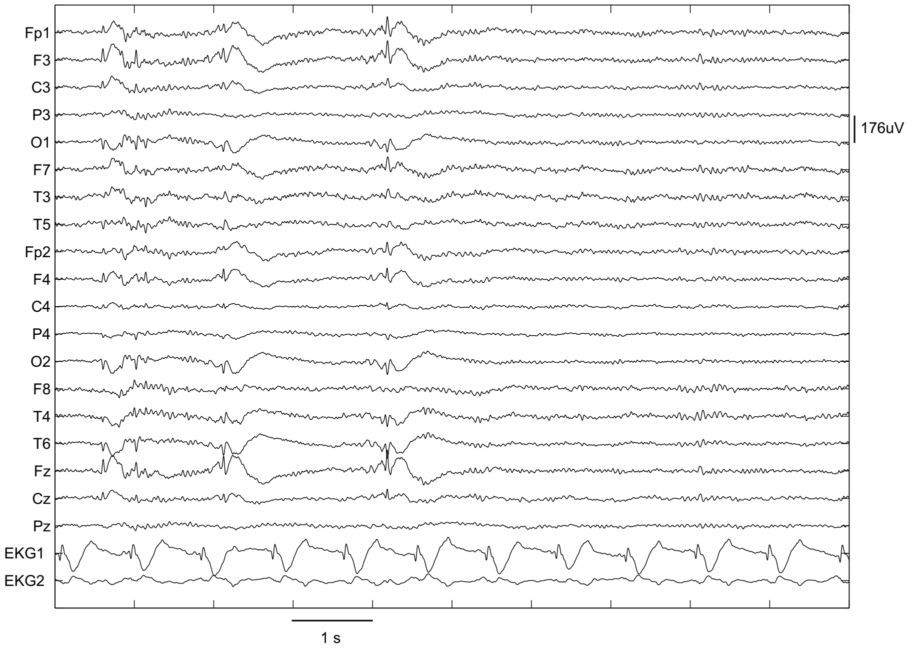

(c) ICA filtering.

Fig. 6.3 Results of Ballistocardiogram removal (patient 3). The epileptic spikes were left intact with both PCA and ICA filters; the ICA performed better in terms of removing the ballistocardiogram.

6.3.3 Gradient artefact

The subtraction and Fourier filters usually produced an EEG of reasonable quality on visual assessment. This allowed the identification of a significant number of epileptic events. In one patient (patient 4) one run was interrupted because the patient had a seizure. This run had to be excluded from the subtraction procedure because the movement artefacts contaminated the calculation of the average artefact. Typically, the Fourier filtering left an artefact at the frame boundary, possibly due to the typical "ringing" effect of this class of filter. The subtraction filter also often left some remaining artefact inside frames. Indeed, the middle part of each slice (i.e. the actual image acquisition) was subject to an apparently random offset from frame to frame. This variability was only present 50ms every 100ms, and was approximately constant in level for each frame, which resulted in a square-wave like artefact. We obtained some results in removal of this artefact by fitting on each frame a computed square wave by least-square regression. This model was then subtracted from the data (not shown).

We counted the spikes outside the scanning periods (i.e. between runs of frames). We also counted the spikes occurring during scanning that were detected with any of the two artefact removal methods. The reviewer stopped marking either at the end of the session or when enough events had been marked to enable assessing the filters. In this process, no measure was taken to filter out the ballistocardiogram, as this was not deemed necessary to identify epileptic events. The average rate of detected events across patients was lower during scanning (4.1 vs 6.5 spike/min), though there was a large variability between patients. Indeed, four patients have a higher rate of spikes during scanning than between scans.

The majority of detected events were tagged equally with both filtering technique: 36.1 % were ranked as clear spikes (type 1) and 34.9 % as ambiguous events (type 2) with both filters, summing up to 71 % of events. A total of 12.3 % of events were marked differently for the two methods, most of them being ranked as type 1 by subtraction filter and type 2 with the Fourier filter (11.7 %), while only 0.6 % were marked as type 2 by the subtraction filter and type 1 by the Fourier filter. 15.6 % of events were marked as not visible with the Fourier filter while only 1.2 % were invisible with the subtraction filter. Again, there was a large variability across patients, but for all cases the subtraction filter performed better than the Fourier filter (8 subjects) or equally (2 subjects).

Figure 4a shows one page of EEG before artefact removal. Figure 4b presents the same page after Fourier filtering and figure 4c after subtraction filtering. Table 4 is a summary of the results in gradient artefact removal.

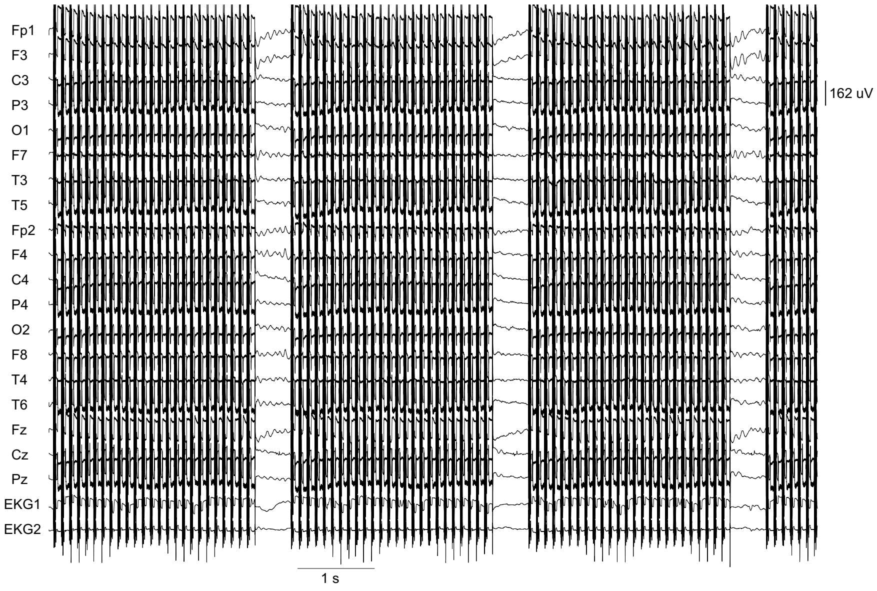

(a) Before removal.

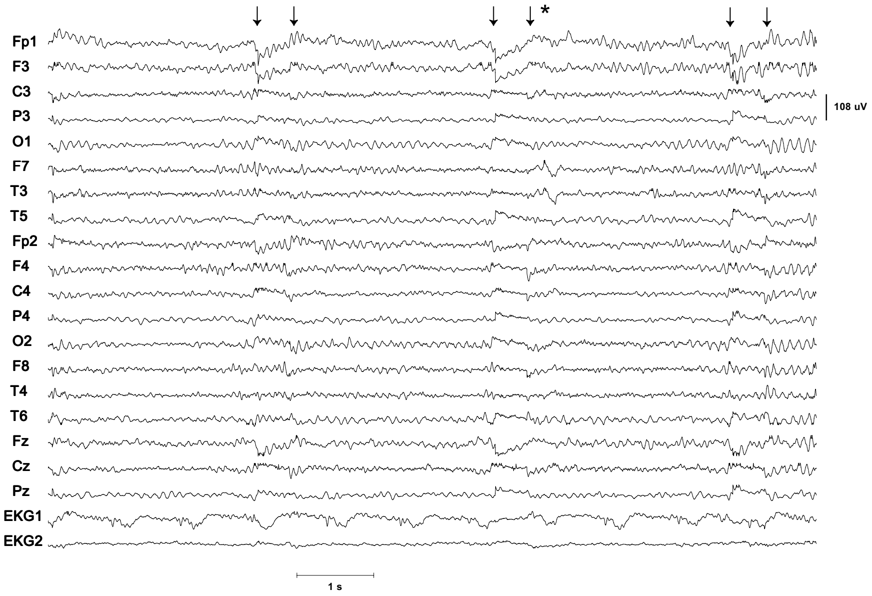

(b) Same page, after Fourier filter.

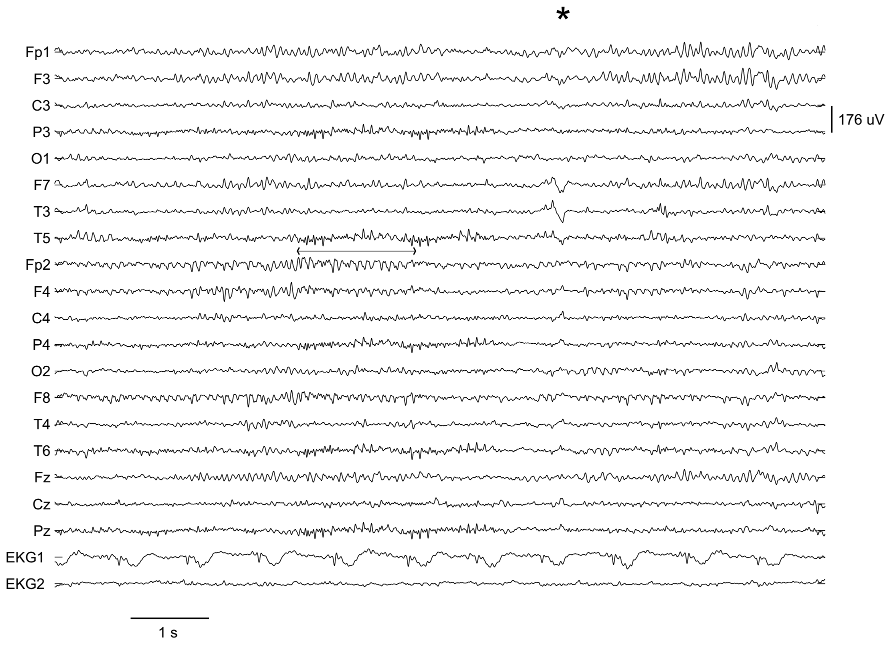

(c) Same page, after subtraction filter.

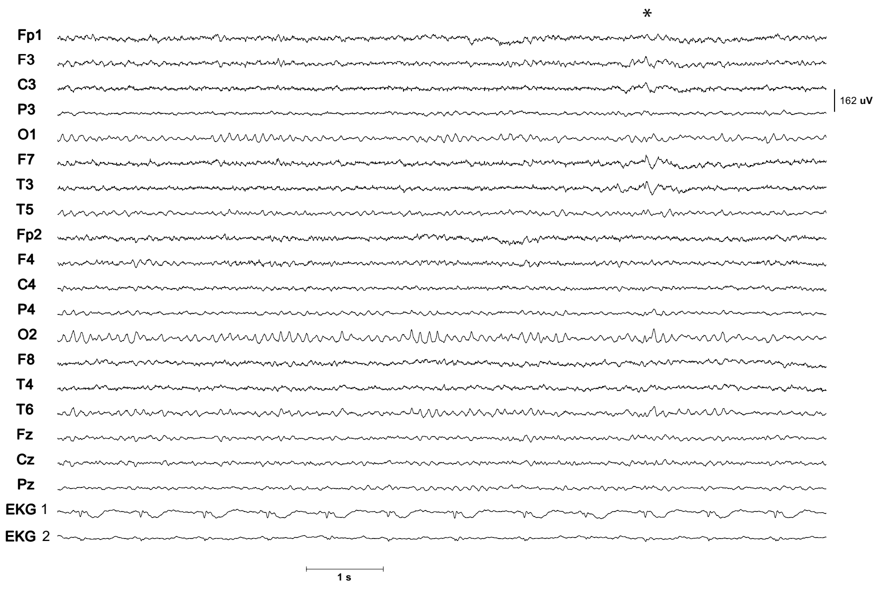

(d) A different page than a-c, located outside scanning.

Fig. 6.4 Gradient artefact removal. The time of epileptic spikes is marked by a star at the top of each graph. In the Fourier filter, there is some distortion of the traces at the boundaries of fMRI frames. The subtraction filter was not able to fully subtract the artefact (e.g. frontal channels of second frame).

6.4 Discussion

The technique of continuous recording of the EEG during EPI scanning is new. The results reported in the laboratories where the techniques were developed do not necessarily apply fully to other laboratories, particularly because some of the technical problems depend on the specific equipment used and on practical considerations that are difficult to replicate.

We found that the most important variable that determines the quality of the EEG and MRI recording is the immobilization of the patient's head and electrode wires. Although this may seem trivial, we struggled for a long time with methods to immobilize the head *comfortably* enough to allow a scanning time of 90 minutes or more. It is particularly important that the patient feels comfortable because any pressure or irritation point will cause the desire to move, and movement results in important problems in analyzing the fMRI data. We experimented with various padding schemes, but found that the plastic bag filled with small polystyrene spheres in which a vacuum is established was an ideal solution. It allows not only immobilization of the head but also of the electrode wires as they leave the head. It is also critical to immobilize the electrode wires on the way to the amplifiers. MRI-compatible sand bags are well suited for this.

We also found that removing the ballistocardiographic artefact was not absolutely necessary to identify epileptic events. This is in distinction to several studies [Allen 98, Bonmassar 99, Goldman 00]. Our generally low level of ballistocardiogram could be due to the use of our head restraints, or to our positioning of the wires inside the magnet. It is obvious, though, that suppressing the ballistocardiogram improves the quality and readability of the EEG. This should be particularly useful when trying to identify subtle events, such as small epileptic discharges. In this context, we propose the well-known methods of PCA and ICA, which appear to be well suited for this purpose. They make use of time-invariant filters that do not rely on the ECG signal or automatic detection of ECG waveforms, contrary to the existing methods. The two methods performed well in eliminating ballistocardiographic artefacts while preserving the spikes. ICA usually produced a clearer separation between artefact and non-artefact components and performed slightly better in artefact removal and in keeping spikes intact. This was sometimes at the expense of occasional unexpected distortion manifesting as a spike of increased amplitude. This could be due to the use of different pages for defining the filter and for applying it. If this were the case, then ICA would be more sensitive to this fact. The selection of the components to be excluded relied on the visual examination of the PCA or ICA decomposition. The identification of components with activity resembling the ballistocardiogram was usually a relatively easy task. The selection of the subset to be excluded was more difficult as excluding more components increased the chances of distorting the spikes. This was especially true for PCA, possibly because the orthogonality constraint on the components lead to a separation between artefact subspace and signal subspace that was not optimal.

When eliminating the EPI artefact with the Fourier filter, we tried to optimize the filter parameters, such as the spectrum amplitude threshold. This resulted in an EEG that was not of the quality illustrated by Hoffmann et al. [Hoffmann 00], but that was of sufficient quality to recognize epileptic spikes in the majority of cases. A major difference with the EEG published by Hoffman et al. lies in the presence of oscillations (or “ringing”) at the border of frames. These oscillations originate from the abrupt changes occurring at the borders of the frames, which we have not been able to suppress during acquisition. In some cases, spikes after filtering were not immediately recognizable because of distortion from the filter, but practice and a priori knowledge of their spatial distribution allowed their identification.

We implemented a subtraction technique for gradient artefact removal that does not use a triggering signal and operates with a 1 kHz sampling frequency. This is in contrast to [Sijbers 99, Allen 00, Goldman 00] who used a trigger signal from the scanner to help post-processing. This is also in contrast to [Allen 00] who used a sampling frequency of 5 kHz and [Cohen 01] who triggered sampling on image acquisition. The use of the subtraction filter produced less distortion than the Fourier filter in a significant number of events. This is particularly important if one is interested in the precise shape of the events. Moreover, a significant proportion of the events were identifiable mainly with the subtraction method - in one patient, this was the case for most events. Nevertheless, the subtraction method often left large remaining artefacts inside frames, lowering the readability of EEG. The most sensitive point in our implementation of the subtraction filter was the automatic identification of the frame time. Indeed, correlating each frame with a reference sampled at different time points can fail to identify the actual lag when there is too much jitter. Also, the method is sensitive to the large noise introduced by movement artefacts. Therefore, it would be valuable to record a trigger signal from the scanner. The sampling frequency of 1 kHz seems appropriate if one uses the inter-frame jitter in frame acquisition versus sampling time to sample the artefact shape with a finer resolution than the actual sampling period.

The advantage of Fourier filtering is its relative simplicity of implementation. Its drawback is that there is some overlap between the spectrum of the artefact and that of epileptic discharges, leading to some signal distortion. Fourier filtering also tends to respond to abrupt changes - such as the beginning and end of gradient artefact - with oscillations in the filtered signal. The advantage of subtracting the artefact is that the shape of epileptic events could potentially be better preserved. The drawback of this method is that it requires knowledge of the precise timing of each artefact and is more computationally demanding (in particular the interpolation step). In fact, both filtering techniques presented some remnant artefacts, which is a direction for future work. Goldman et al. have shown good results in this field with adaptive filtering [Goldman 00]. We believe that using a priori knowledge on the shape of remnant artefact is also a promising path.

6.5 Conclusions

We obtained inside the scanner a background EEG that was clearly distorted compared to a regular EEG, but that was still definitely usable for our purposes. During the scanning periods, the filtering methods enabled to identify a good amount of epileptic events, even though the lower rate in some patients suggests that we may have missed some spikes. There is obviously room for improvement, both in the recording setup and apparatus and in the post-processing techniques, but the results we obtained are very encouraging and confirm the feasibility of simultaneous EEG and fMRI recordings. Our study was focused on the applications in epilepsy, where recovering the precise shape of events may not be so crucial as long as they can be identified with confidence. However, some of our results could be useful to other fields of EEG such as evoked potentials or studies of EEG rhythms. In this case, more studies would be needed to assess the impact of the filtering procedures on the EEG signal of interest.

Acknowledgements

We thank Anne Alex, Mike Ferreira and Vincent Navarro for their contribution in data gathering. This work was supported by the Canadian Institutes of Health Research (CIHR) under grant MOP-38079. C.G. Bénar is funded by a CIHR doctoral research award.

Chapter 7

Article 3: The BOLD Response to Interictal Epileptiform Discharges

Context

The preceding article dealt with the methods for recording EEG in the scanner and interpreting the resulting signals. In the present study, we will show fMRI responses to epileptic spikes, results that were obtained thanks to this simultaneous recording of EEG and fMRI.

We will consider this response both in space (activated fMRI areas and their relation to the epileptic condition of the patients) and in time (average temporal course of the BOLD response).

This article was published as [Bénar 02b].

Abstract

We studied single-event and average BOLD responses to EEG Interictal Epileptic Discharges (IEDs) in four patients with focal epilepsy, using continuous EEG-fMRI during 80-minute sessions. The detection of activated areas was performed by comparing the BOLD signal at each voxel to a model of the expected signal. Since little is known about the BOLD response to IEDs, we modelled it with the response to brief auditory events (Glover 1999). For each activated area, we then obtained the time course of the BOLD signal for the complete session and computed the actual average haemodynamic response function (HRF) to IEDs.

In two of four patients, we observed clear BOLD responses to single IEDs. The average

response was composed of a positive lobe peaking between 6 s and 7 s in all patients and a negative undershoot in three patients. There were important variations in amplitude and shape between average HRFs across patients. The average HRF presented a wider positive lobe than the Glover model in three patients and a longer undershoot in two. There was a remarkable similarity in the shape of the HRF across areas for patients presenting multiple activation sites. There was no clear correlation between the amplitude of individual BOLD responses and the amplitude of the corresponding EEG spike.

The possibility of a longer HRF could be used to improve statistical detection of activation in simultaneous EEG-fMRI. The variability in average HRFs across patients could reflect in part different pathophysiological mechanisms.

7.1 Introduction

Interictal epileptiform discharges (IEDs) are studied routinely in the evaluation of patients with epilepsy. The classical technique used is electroencephalography (EEG), on which IEDs produce marked and stereotyped trace deviations. Recently, the recording of EEG in the magnetic resonance (MR) scanner [Ives 93, Huang-Hellinger 95] has enabled the study of electrographic events with functional MR imaging (fMRI), making use of the Blood Oxygen Level Dependent (BOLD) effect (Kwong et al. 1992, Ogawa et al. 1992). This is a challenging task because of the artifacts caused by the changing magnetic field on the EEG and by the EEG equipment on MR images. It is, however, an exciting possibility because it may allow accurate anatomical localization of the generator of IEDs by correlating changes in fMRI signal and the presence of IEDs recorded with surface EEG. It may also provide information on the mechanisms underlying IEDs, since the metabolic change measured by fMRI is a consequence of the abnormal neuronal activity generating the IED [Buxton 98, Hoge 99]. Images can be acquired either immediately after each event ("EEG-triggered fMRI", [Warach 96, Seeck 98, Allen 98]) or continuously ("continuous EEG-fMRI", [Lemieux 01b, Baudewig 01]. The temporal characteristics of the BOLD response to IEDs constitute a very important aspect of this technique. Indeed, the statistical detection of brain activation relies on assumptions about this signal that have not yet been verified. Also, it may be of physiological interest to understand better the determinants of the magnitude of the BOLD response.

We propose to study in a group of selected patients the relationships between IEDs and

BOLD response, using continuous EEG-fMRI.

7.2 Subjects and Methods

7.2.1 Recording

We recorded within a 1.5 T MR scanner (Siemens Vision, Siemens, Erlangen, Germany) 21 channels of EEG on patients presenting focal epileptiform discharges during their clinical scalp EEG evaluation. EEG recordings were made with an EMR32 amplifier (Schwarzer, Munich, Germany). In each subject we acquired between 7 and 10 runs of 120 frames per run. Each frame consisted of 25 BOLD EPI 64×64 axial images (i.e. 2D slices) acquired sequentially and covering the entire brain with a voxel size of $5 \times 5 \times 5$ mm. The TE was 50 ms and the flip angle 90 degrees. The inter-slice time interval within each frame was approximately 100 ms and the total inter-frame interval was 3 s. Each run therefore lasted approximately 6 minutes. Given the one to two minute inter-run gap, the complete session lasted up to 80 minutes.

7.2.2 EEG analysis

The EEG was processed offline in order to filter out the scanner artifact (FEMR software, Schwarzer, [Hoffmann 00]), the maximum value of which is of the order of 50 times the average amplitude of background EEG. Although there was also some ballistocardiogram artifact, we found that it was sufficiently low to allow identification of epileptic spikes and we did not use any special procedure to remove it. Each filtered EEG signal was reviewed by two experienced electroencephalographers (D.W.G. and J.G.) who marked the times of epileptiform discharges. They found that, although traces were distorted by the filtering process, these were good enough to identify the epileptic events.

7.2.3 Image analysis

We performed statistical processing of the images using the methods and software of Worsley et al. [Worsley 96, Worsley 02] in order to find the areas that were activated in response to the IEDs. A model of the expected BOLD response to the IEDs was created based on the timing of the epileptic discharges and a model of the haemodynamic response (HRF) to brief auditory events [Glover 99]. The HRF model consists of the sum of two gamma functions. The parameters of the two functions (time to peak, approximate FWHM) are (5.4, 5.2) and (10.8, 7.35) respectively; the second function is multiplied by (-0.35). Images were motion corrected and smoothed with a gaussian filter (6 mm FWHM). We obtained statistical maps of activation indicating at each voxel the level of correlation between the fMRI signal and the model (t-stat maps). The first three frames of each run were excluded from the analysis to let the system reach a steady state. Model and signals were prewhitened with an AR(1) model. Slow fluctuations in the fMRI signal were taken into account by including in the linear regression a third order polynomial fitted to each run. We computed for each patient the significance level for an average brain volume (1200 cc), which depends on the total number of frames. The probability that the maximum value in the brain passes the level by chance is 0.05 when there is no effect (no signal). We used two methods in computing the level of significance. The first method is based on the theory of random fields [Worsley 96], and takes into account the smoothness of the field (the 6 mm FWHM filtering). The resulting values ranged between 5.04 and 5.09. The second method was a Bonferroni correction for sample size (considering we have approximately $1200.10^3/(5 \times 5 \times 5) = 9600$ independent variables). The resulting values ranged between 4.41 and 4.44. We used 4.4 as a threshold on the t-stat maps. We defined regions of interest (ROI) as clusters of contiguous voxels (4-neighbour connectivity) with a t-stat score above the threshold

7.2.4 BOLD time course

We obtained for each ROI the time course of the fMRI signal during each run by averaging the signals of all the voxels within the ROI. This resulted in 7 to 10 120-sample BOLD time courses for each ROI, with a 3 s sampling interval.

7.2.5 Individual and mean haemodynamic response

In order to estimate the response to IEDs with a good signal to noise ratio, we needed to average the fMRI signal following several IEDs. To avoid contamination by neighbouring IEDs, we averaged only the response to isolated IEDs, defined as spikes or bursts of spikes with no other epileptiform activity in a window around the event. This window spanned 30 s preceding and 30 s following the event, except for patient C where this constraint resulted in a very poor signal to noise ratio. Reducing this constraint to 15 s increased the number of IEDs that could be averaged and resulted in a better signal to noise ratio, at the expense of the length of the averaging window. Prior to averaging, we fitted a third order polynomial to every run of fMRI signal (excluding the first three frames) and subtracted it from the data to remove slow fluctuations, thus emulating the polynomial used during statistical processing.

In order to average the individual responses, we had to define for each response the point corresponding to $t = 0$ s. As IEDs occur at a random time within a 3 s image frame, it is important to take into account this variability to average points that correspond to the same part of the HRF. The method is described in figure 7.1. The timing of each fMRI signal sample was taken as the time of the slice containing the maximum score of the ROI. We used spline interpolation to obtain a value every second. For each individual BOLD response, the time origin for averaging was taken as the point in the oversampled signal that was closest to the corresponding IED (Fig. 7.1). As we mentioned earlier, one fMRI signal sample is the average signal on a ROI. Considering that the 25 slices are taken sequentially with 100 ms between two slices, and that ROIs span only a few slices, the impact of averaging few neighbouring slices is small.

7.2.6 Correlation IED/BOLD response

We attempted to determine if there was a relationship between the amplitude of the EEG during the IED and the amplitude of the corresponding BOLD response. The amplitudes of isolated IEDs and single BOLD responses were measured as the square root of the energy of these signals in the time range of the event. For IEDs this included the spike and the following slow wave. For the BOLD response, this included the first positive lobe and the undershoot, except for patient A where we only considered the positive lobe (no undershoot). Isolated IEDs were defined as events without another IED 30 s before or 30 s after. We compared the amplitude of each BOLD response either to the amplitude of the corresponding single EEG spike or to the sum of amplitudes of all the spikes in the burst when applicable.

7.2.7 Correlation between BOLD responses in different regions

As patients B and D presented several distinct areas of activation, we measured the correlation in the time course of the fMRI signal between pairs of ROIs as a function of time.

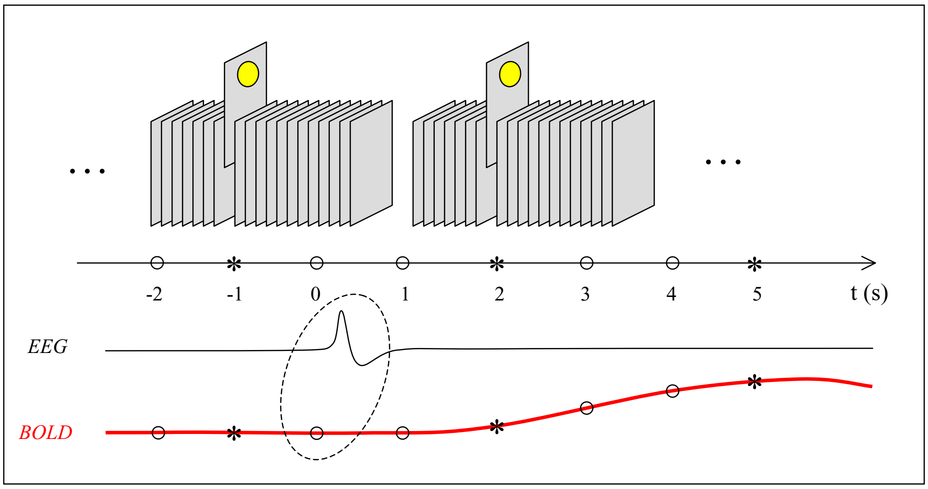

(a) Schematic representation of the selection of the fMRI signal time point corresponding to the EEG event for the purpose of averaging. Middle line: relative time. Above the time line: two fMRI frames (consisting each of 25 slices; the slice containing the maximum of the ROI is raised). Below the time line: EEG and fMRI signals. Star: fMRI signal sample, circle: interpolated point. The origin of time is taken as the point in the fMRI signal that is closest to the EEG event. Note that the EEG discharge is close to an interpolated point.

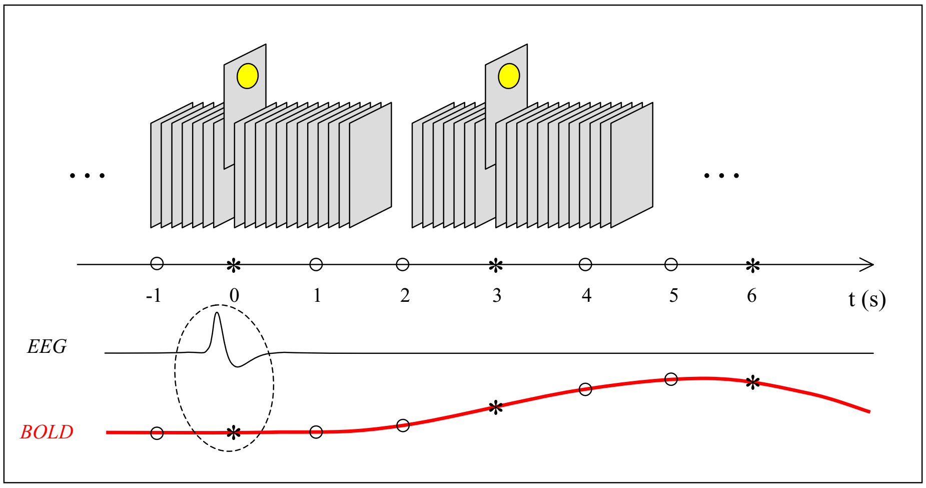

(b) Same as (a), but this time the EEG discharge is close to a fMRI signal sample.

Fig. 7.1

We used a running correlation of 10 samples (i.e. 30 s) with overlapping windows.

7.2.8 Subjects

Patients were selected on the basis of having focal and frequent interictal spiking; their informed consent was obtained in accordance with regulations of the Research Ethics Board of the Montreal Neurological Institute and Hospital. We included in this study only the small subgroup of patients with clear fMRI activations (i.e. statistically significant clusters of activated voxels) that were spatially consistent with EEG findings.

Patient A is a 19 year old right-handed male. He presents partial complex seizures (PCS), with an age of seizure onset of 1 and a seizure frequency of 2 to 4 per month. He has had a left temporal lobectomy that was performed because of temporal sclerosis. He has multiple small calcified lesions scattered in the brain parenchyma secondary to neurocysticercosis. His EEG findings are left infero-mesial interictal spikes (max. T3-Sp1, F7-T3) and left mid-posterior temporal ictal activity. His treatment consisted in carbamazepine, lamotrigine and dilantin. Carbamazepine was reduced at the time of the study.

Patient B is a 34 year old right-handed male. He presents PCS with secondary tonic-clonic generalization (SGTC); the age of onset was 9, and the seizure frequency is 1-5 per month. He has no imaging findings. The EEG interictal and ictal activity is in the temporal region (max. at T3, F7, T9). His treatment consisted in valproic acid and carbamazepine. Valproic acid was stopped at the time of the study.

Patient C is an 18 year old right-handed male. He presents PCS, rarely SGTC, an age of onset of 10 and a seizure frequency of 3-4 per month. He has no imaging findings. The interictal and ictal activity is right fronto-temporal (max. Fp2, F4, Fz, F8, Sp2). His treatment consisted in carbamazepine and lamotrigine. Doses were not changed at the time of the study.

Patient D is a 46 year old right-handed male. He presents PCS, with an age of onset of 17. He has bilateral occipital periventricular nodular heteretopia. His EEG interictal activity is left hemispheric with maximum middle and posterior left temporal (T3,T5,T9). Ictally, he has no change or left temporal activity. His treatment consisted in topiramate, oxcarbazepine and clobazam. Doses were not changed at the time of the study.

7.3 Results

Table 7.1 presents a summary of EEG and fMRI findings for the four patients in the study. Typical IEDs are shown for each patient on Figure 7.2. Patient A presented only single spikes, patient B only bursts of one to three seconds. Patient C presented mainly isolated spikes with a few consecutive spikes, patient D mainly single discharges with a few cases of two or three consecutive discharges. In all four patients, artifact filtering introduced only minor distortion in the waveforms, which were quite similar in temporal shape and spatial distribution to the waveforms outside the scanning period (including the atypical waveform of patient D). Figure 7.3 illustrates the activated areas for each patient. All but one (third ROI of patient D) are within a few centimeters of the EEG electrodes with maximum signal amplitude. The size of the ROIs for patients A, B and C ranged between 20 and 58 voxels, which correspond to volumes between 2.5 and 7.25 cm^3. Patient D had only 1 or 2 voxels above the threshold in each of the three ROIs. It is important to remember that this threshold was determined for *individual* voxel activation, which is a statistically demanding definition [Worsley 96]. If one considers clusters of contiguous voxels, the threshold can be lowered and still result in significant regions - defined as *sets* of connected voxels [Cao 99]. We verified that lowering the threshold to 3 for patient D lead to sets of connected voxels that were significant as a group (any set of three connected voxels above this threshold is significant).

Patient	**EEG activity**	**# runs**	**# IED's**	**# IED's in average HR**	**fMRI activation**	**Max. t-score**	**# significant voxels**
A	L posterior temporal	8	7	6	L posterior temporal	7.6	20
B	L temporal	7	34	14	L posterior temporal	14	32
					L anterior temporal	8.1	58
C	R frontal	9	218	27	R superficial frontal	8.7	39
D	L temporal	10	98	9	L temporal	4.7	2
					L parieto-temporo-occipital	4.9	1
					L parietal	4.9	2

Table 7.1 Comparison of EEG and fMRI findings.

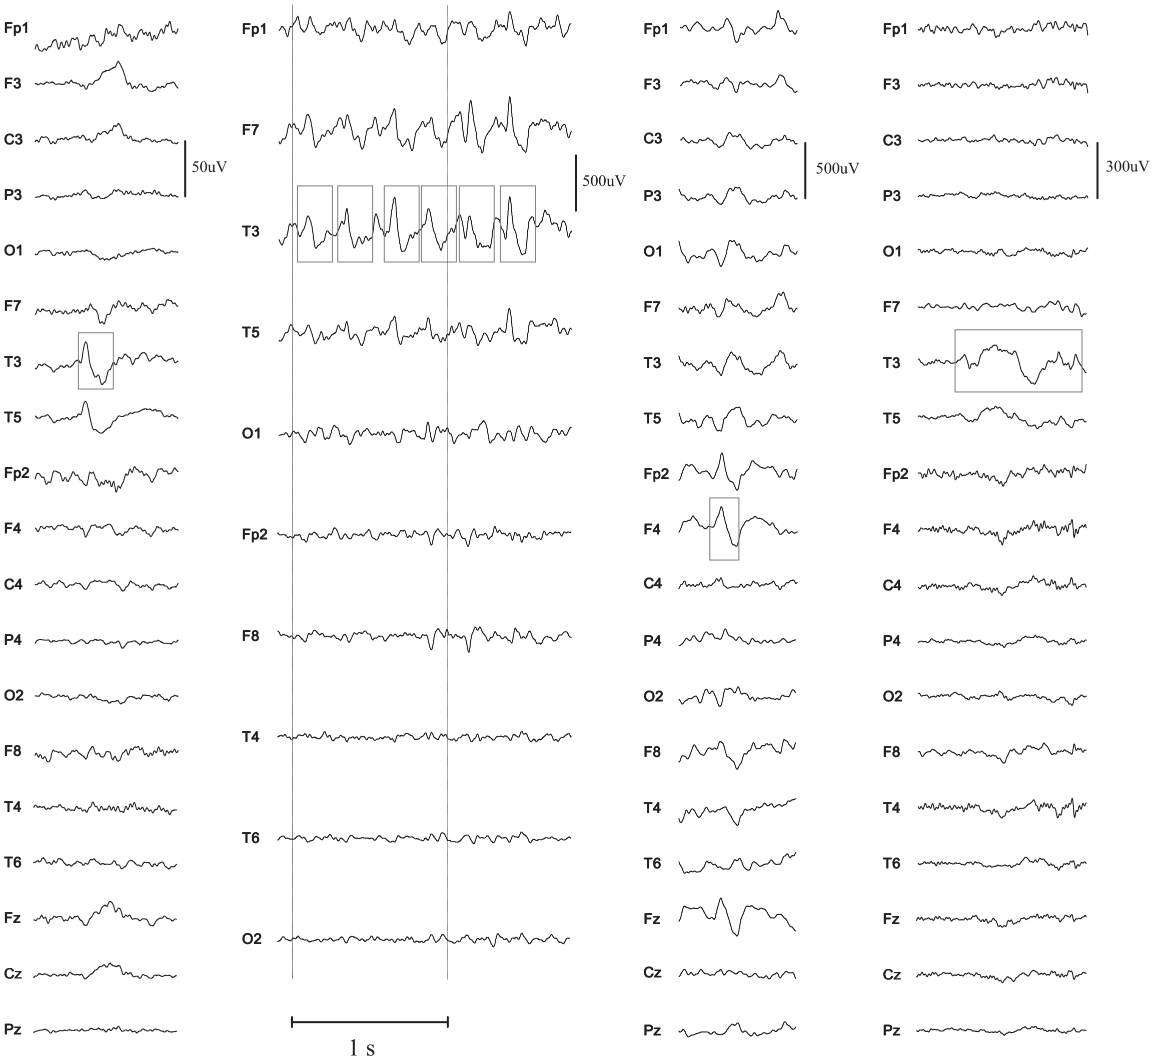

Fig. 7.2 Examples of IEDs recorded during scanning and subsequently filtered (gray boxes: time windows and channel used in computing the IED energy). Left to right: patient A to D. Reference electrodes: (F3+F4)/2. Time lines are separated by 1s.

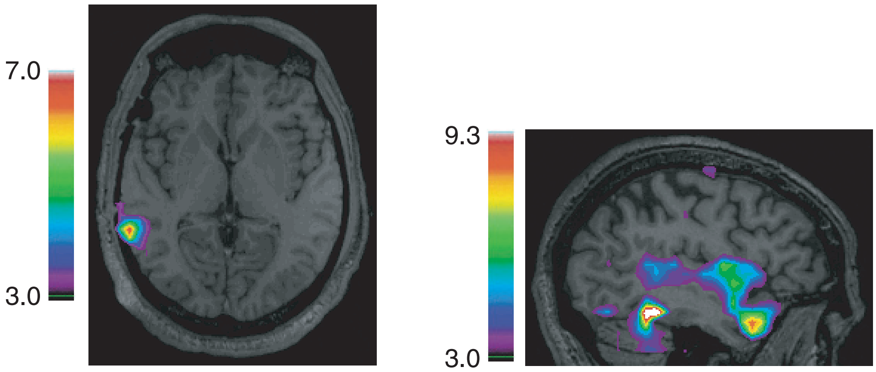

(a) Patient A (b) Patient B

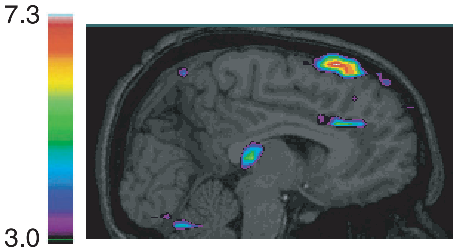

(c) Patient C

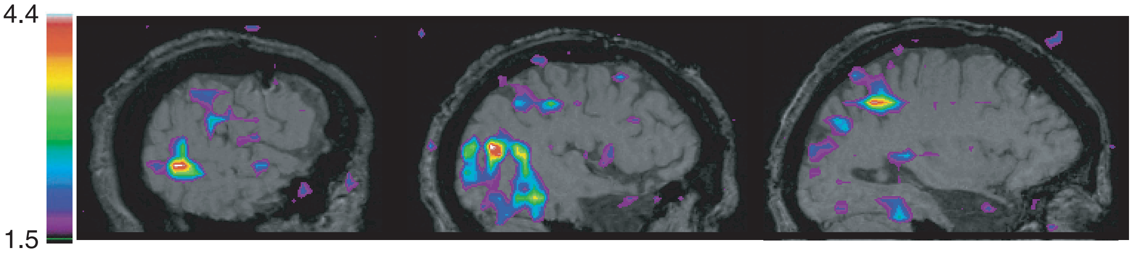

(d) Patient D

Fig. 7.3 Functional MRI activations (t-stat thresholds: 3 for Patients A, B, C, 1.5 for patient D; the color range spans between the threshold and the maximum in the image). One slice each is shown for patients A, B and C. Patients A and B have only one region of activation. In patient C, both regions of activation are visible on the same slice. In patient D, the three regions of activation are shown on three different slices.

7.3.1 Bold time course

Figure 7.4 illustrates the BOLD time course of one representative 6-minute run for each patient. For patients A and B, every IED (marked by arrows) was followed by a clear increase of the BOLD signal, lasting several samples. Single responses were more difficult to see for patients C and D. In some cases, series of consecutive IEDs were followed by a larger response than individual ones (e.g. around frame 100 for patient C and around frame 30 for patient D), but this was not always the case (e.g. before frame 40 in patient C). We noted the presence of a few possible BOLD events that resembled the BOLD increase following an IED but did not correspond to any visible scalp epileptiform EEG activity (asterisk in Fig. 7.4(b)). We also observed large slow fluctuations with an relative amplitude of up to 1 % in the course of a run, and a period of up to 100 s (e.g. frames 31 to 61 in Fig. 7.4(a)).

7.3.2 Average haemodynamic response

Table 7.1 shows the number of IEDs that were used to obtain the average haemodynamic response for each patient, given the constraint that IEDs are isolated. The average HRFs were relatively similar to the model of Glover: they all had a positive lobe that peaks between 6 and 7 seconds after the IED, and three presented a negative response (undershoot) following the main lobe (Fig. 7.5). There were notable differences though: patients A, B and D had a wider main lobe, patients B and D had a longer and more pronounced undershoot than the model - especially patient D whose response lasted up to 45 s. There was no undershoot at all in patient A. In patient C, the standard error was particularly high, but the morphology of the response was relatively similar to the model. Across patients, the maximum signal change varied significantly, ranging from 0.4 % to 1.5 % (relative to baseline, defined as the value of the third order polynomial that models slow fluctuations). For patients with several areas of activation, the average HRFs were remarkably similar between ROIs (Fig. 7.5(b) and 7.5(d)).

7.3.3 Correlation IED/BOLD response

The time range and channel in which the signal energy was computed is shown on the traces on figure 7.2. We can note a large variability in the root energy of EEG and fMRI signals: there is a factor of up to 3 between the smallest values and the largest values for

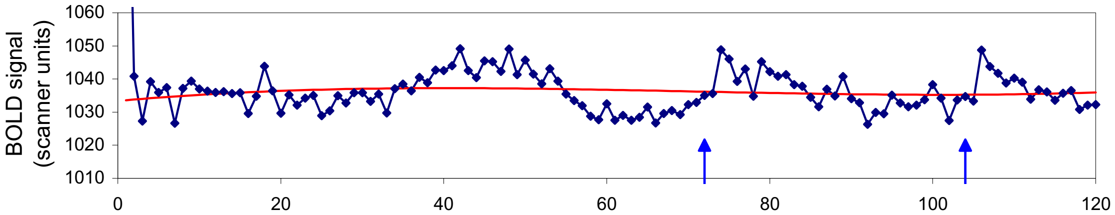

(a) Patient A

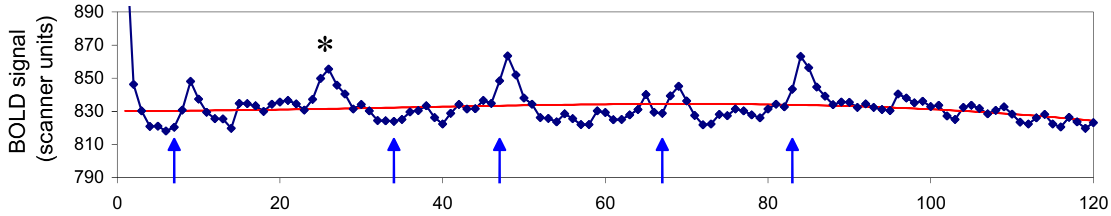

(b) Patient B

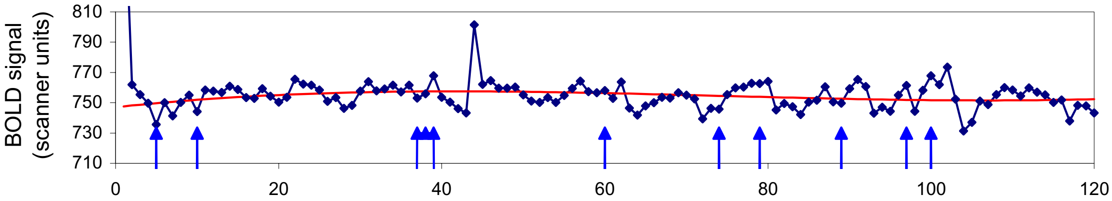

(c) Patient C

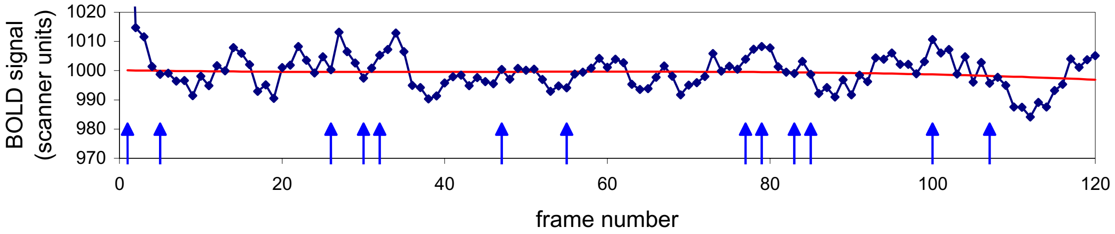

(d) Patient D

Fig. 7.4 The fMRI signal in the ROI of one run for each patient. Arrows: IEDs, star: event that resembles the HRF but not preceded by a scalp IED. Each point corresponds to one 3 s frame.

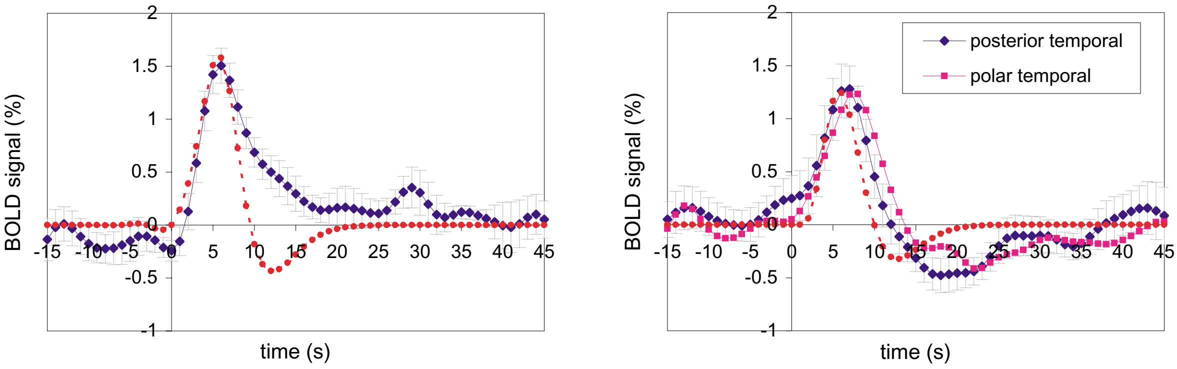

(a) Patient A (b) Patient B

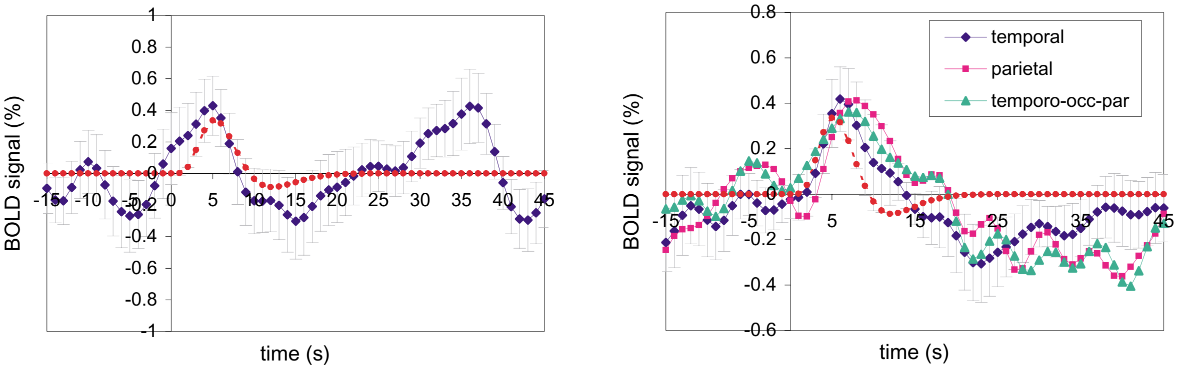

(c) Patient C (d) Patient D

Fig. 7.5 Average haemodynamic response and model (red dotted lines). The model has been scaled in amplitude for each patient to enable direct comparison. Signals are presented in percentage of activation relative to the baseline. For patient B, the model has been convolved with a 5 s window to take into account the average burst duration.

both signals (figure 7.6). The correlation coefficients are all very low, but there could be a tendency for a positive correlation in patient A and C.

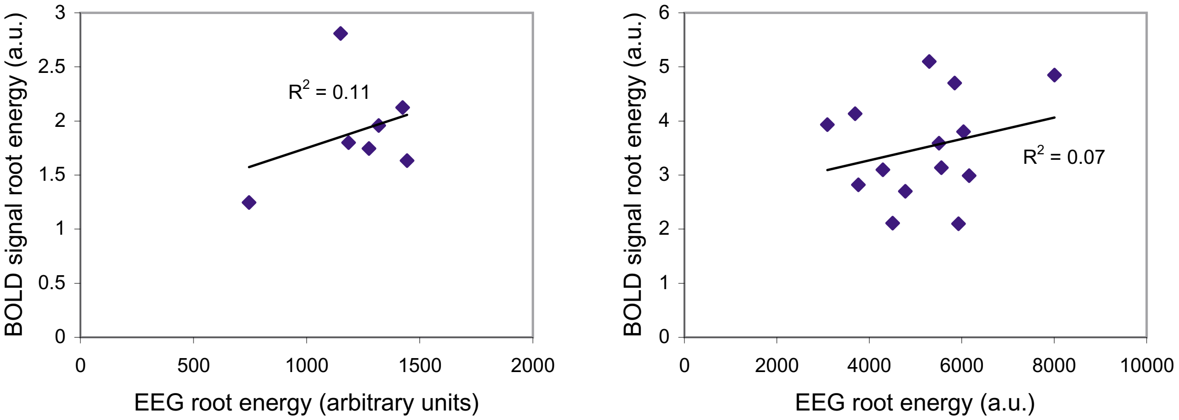

(a) Patient A

(b) Patient B

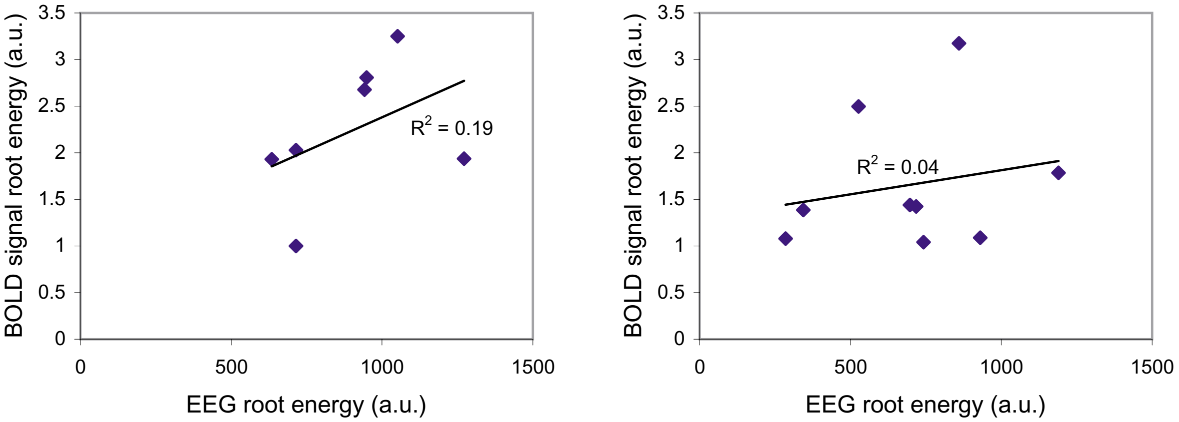

(c) Patient C

(d) Patient D

Fig. 7.6 Plots of root energy of the BOLD response versus root energy of the EEG signal for the set of isolated IEDs.

7.3.4 Correlation between BOLD signals in different regions

The running correlation between the BOLD signals of the two ROIs of patient B showed that their activity is in general correlated around the timing of IEDs and not elsewhere (fig. 7.7). There are some notable exceptions though, where correlation was high with no visible activity on the scalp or, on the contrary, when there is was a low correlation even at the time of IEDs. For some events there was a BOLD response in the posterior temporal lobe and not in the anterior part, but not the reverse. The results of patient D are more difficult to interpret due to the high density of events but show the same tendency (i.e. high correlation around IEDs and not elsewhere, not shown).

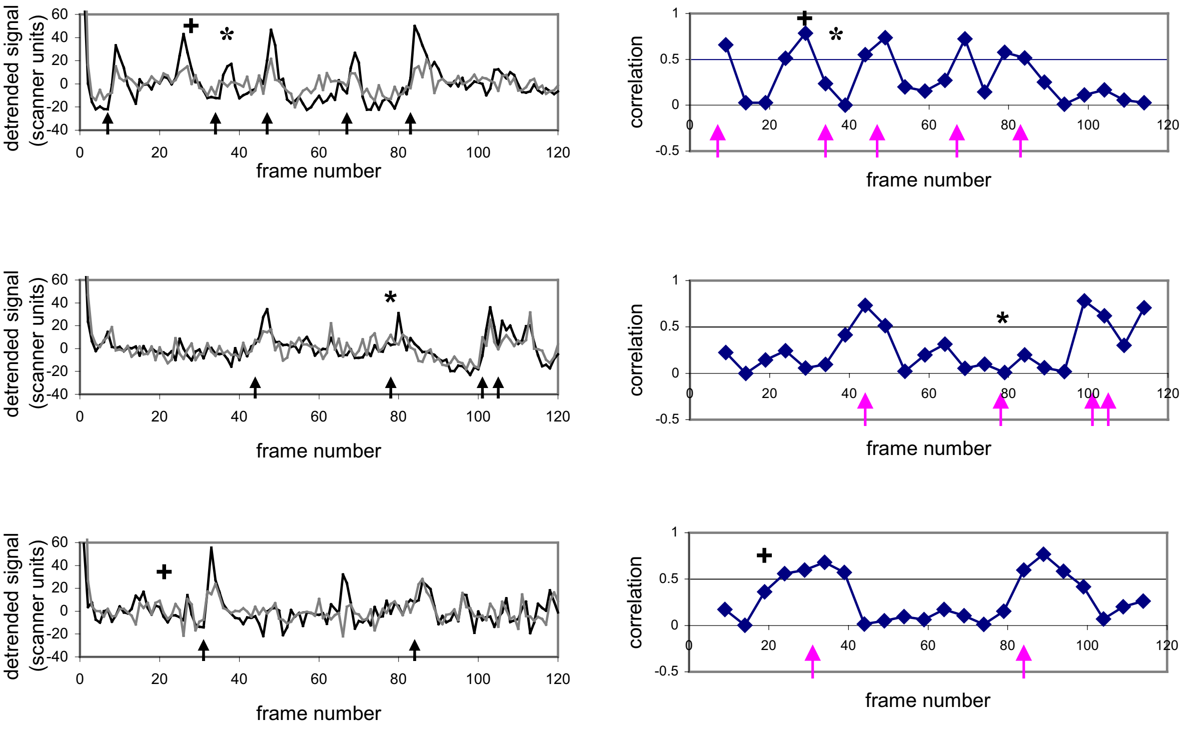

Fig. 7.7 Correlation between time courses of fMRI signals in different brain regions. Left column: three runs of patient B (black trace: posterior temporal ROI, gray trace: anterior temporal). Right column: corresponding correlation coefficient between the two areas with a 30 s-sliding window (arrows: IEDs). Plus sign: High correlation between fMRI signals in the absence of scalp IED; Star: low correlation between fMRI signals in the presence of scalp IED.

7.4 Discussion

We recorded simultaneously and continuously the EEG and BOLD activity in patients with focal epilepsy. We have performed statistical detection of the areas of the brain where the fMRI activity was correlated with the timing of epileptic activity in the EEG. We used in the detection a model of the expected HRF, as part of our standard protocol. This is a limitation in our study, as we may miss BOLD responses that are too different from the model. In the subset of patients in whom we found brain regions that were significantly activated in relation to epileptic activity, we examined the time course of the fMRI signal throughout the experiment. In two of four patients, we observed a clear fMRI signal increase following most individual IEDs and little change in activity away from IEDs. Considering that the active regions are compatible with the EEG field and that the amplitude of the fMRI signal increase is in the expected range for increases in cortical venules [Ogawa 98], we feel reasonably confident that this represents a true response to the epileptic events. In the other two patients, we encountered a less convincing increase in fMRI signal following individual IEDs, and a larger amount of fluctuations outside of events. Nevertheless, the high statistical score on the whole data set and the spatial compatibility with EEG suggests again that this is a true albeit small response. The increases in fMRI signal activity away from IEDs could be related to epileptic activity that is not apparent on the scalp, or to normal brain fluctuations that are of a magnitude similar to that of epileptic responses.

We observed the typical slow fluctuations in the fMRI signal, but these were not always modelled correctly by the third order polynomial used in our analysis. This is a sensitive point because of the scarcity of the events and the low signal amplitude in some cases. This would call for a more efficient preprocessing in order to improve statistical detection, such as high pass filtering [Kruggel 99b]. Also, the fact that slow fluctuations are of the same order of magnitude as the signal could pose problems to spike-triggered methods that only consider samples widely spaced in time and cannot completely take into account such baseline fluctuations.

There was a high variability in the amplitude of individual BOLD responses within patients. The amplitude of BOLD signal was not correlated with the EEG amplitude in a statistically significant way, though there was some hints of a correlation in some patients. It would be interesting to perform more measures of this type to increase our sample size. Also, the use of low-pass filtering to better remove the slow fluctuations would result in better estimates of BOLD signal energy. A correlation could be explained by the fact that the number of neurons that are taking part in the excessive epileptic discharge in part determines the amplitude of the EEG. Presumably this same number of neurons will determine the magnitude of the metabolic response. On the other hand, one does not expect a simple relationship between the scalp EEG amplitude and the number of neurons involved, as spatial and temporal synchrony is also a very important factor. This may be why our correlation coefficients are so low. Another explanation could be that the metabolic requirement of a burst of spikes might not be the sum of the metabolic requirements of single spikes. Also, the two patients that have the higher number of spikes (C and D) are the ones with smaller HRF amplitudes, possibly suggesting different mechanisms in the metabolic demand for different types of epileptic activity.

The overall shape of the average HRF was consistent with the one obtained in classical event-related fMRI experiments, and with the one reported for one patient by Lemieux et al. [Lemieux 01b]. This was expected because the detection method we used implies that the response is similar to the model, but there is still some latitude for shape changes. Indeed, the HRFs we observed are wider than the model and two patients present undershoots much longer than the model. Patient C shows some late activity, but one has to remember that the window with no other IED was shortened in this case, and therefore the data after 15 s could be contaminated by other IEDs. It is remarkable that the HRFs of epileptic spikes, events that last 100 to 200 ms, seem to last over 30 seconds, in one patient up to 45 seconds. The variability in observed response from patient to patient could reflect normal inter-subject variability [Aguirre 98] or variable types of epileptic processes. In order to evaluate this possibility, it will be necessary to study many patients with different pathologies. The variability in HRF morphology contrasts with a remarkable intra-patient similarity in the mean HRFs across areas, in those patients with more than one active area.

The variability in the amount of fMRI signal change across patients could be due to different densities of active neurons. It could be also that in some cases the EEG spike simply results from a neuronal synchronization in a region that is also active away from spikes, but less synchronized: in this context, the synchronization would probably not imply a much higher metabolic demand. A third possibility is that different lesions may produce different patterns of activation - it is a point to be noted especially because the blood supply of some lesions may be abnormal.

In patient B, we noticed that for some events there was a BOLD response in the posterior temporal lobe and not in the anterior part, but not the reverse. This suggests that the posterior region may be a “primary” region and the anterior a place of non-systematic propagation. This is compatible with EEG interictal and ictal findings in this patient. Indeed, the coherence function [Gotman 83] computed on 4 epochs of bursts between bipolar EEG channels T3-T5 (posterior) and Fp1-F7 (anterior) shows a slope of -3.6 deg/Hz, which corresponds to a delay of 10 ms, with the posterior region leading. Also, on ictal EEG, the activity seems to start at the posterior temporal electrode (T3) and propagate to other regions, including anteriorly (F7).

The statistical method used to analyze the relationships between epileptic events and BOLD response is derived from the methods used for classical event-related behavioral studies. Some of our findings raise the possibility that this may not be the optimum approach. Three factors are particularly worth noticing. (1) Some transient fMRI signal increases away from IEDs resemble the HRF and could correspond to deep epileptic discharges that are not visible on scalp EEG. The time course model can obviously not take into account such events. (2) We have seen that the HRF shape may differ significantly from the standard model. (3) Epileptic events are variable in intensity and result in variable amplitude single BOLD responses. All these factors could be taken into account in the statistical analysis and would probably result in better detection power. Point (1) and (3) could be addressed by fitting a model for each event instead of the whole run. This would not take into account the activity remote from event, and allow the amplitude of the fitted coefficient to vary from one event to the other. The fact that the correlation between the energy of neuronal events and the energy of the bold response is very weak implies that taking into account the amplitude of each individual EEG event in the analysis would probably not be very helpful. Point (2) could be addressed by using an average HRF across patients as a model. We assessed the influence of the HRF shape by repeating the statistical processing for patient A with his average HRF as a model, and found a maximum t-stat score of 13.1 instead of 7.6 with the Glover model. Another possibility would be to use methods that do not rely on a specific shape of response [Josephs 97].

We obtained statistically significant regions that correspond only to small areas of cortex. We know from previous studies that this is probably only the “tip of the iceberg”, as such areas are too small to produce a visible EEG deviation on the scalp [Cooper 65, Merlet 99]. Improving the detection method as outlined above would probably increase the extent of significant areas, but one has to keep in mind that EEG and fMRI are measuring different phenomena and that one cannot expect a one-to-one relationship between them [Nunez 00].

Epileptic spikes are often considered an important marker of the epileptogenic region, but they are usually considered of no clinical significance in themselves. It has been demonstrated, however, that they can be associated with brief cognitive impairment [Shewmon 88, Binnie 92]. We show here that in some patients each spike results in a clear haemodynamic response lasting 30 to 40 seconds, whereas in others the response is of much smaller amplitude and duration. This opens the door to a new differentiation between spikes possibly having different mechanisms and significance. We have also shown that the HRF for epileptic spikes can be somewhat different from the standard model and is also different from patient to patient. This will result in more complex analysis procedures, since there does not appear to be a standard response to all epileptic spikes.

Acknowledgements

The authors want to thank Dr. Abdullah Al-Asmi for the clinical data and the EEG technicians of the MNI for their precious contribution. This work was supported by the Canadian Institutes of Health Research (CIHR) under grant MOP-38079. C.G. Bénar is funded by a CIHR doctoral research award.

Chapter 8

Article 4: Statistical Maps for EEG Dipolar Source Localization

Context

The preceding studies dealt with techniques for localizing the brain regions responsible for generating epileptic spikes, first with EEG (chapter 5), then with fMRI (chapter 6 and 7) - this latter was rendered possible by the simultaneous recording of EEG and fMRI. Once this localizing data is available, there comes the issue of integrating the results for gaining a better spatial estimation of the activity in the one hand and optimally combine the spatial capacities of fMRI and the temporal resolution of EEG in the other hand.

We propose to build statistical maps for EEG localization, in a somehow similar manner to fMRI maps. The two types of maps can then be compared in order to establish correspondences between activated areas. Areas activated in the two modalities will greatly increase the confidence that we are not dealing with spurious results, as the results originate from independent measurements. Areas activated in one modality only may point to complementary roles for the two techniques. Once a correspondence is made, dipoles located in the centre of gravity of regions activated in the EEG maps can help timing the activity across the regions.

This article was submitted.

Abstract

We present a method that estimates three-dimensional statistical maps for EEG source localization. The maps assess the likelihood that a point in the brain contains a dipolar source, under the hypothesis of one, two or three activated sources. This is achieved by examining all combinations of one to three dipoles on a coarse grid and attributing to each combination a score based on an F statistic. The probability density function of the statistic under the null hypothesis is estimated non-parametrically, using bootstrap resampling. A theoretical F distribution is then fitted to the empirical distribution in order to allow correction for multiple comparisons.

The maps allow for the systematic exploration of the solution space for dipolar sources. They permit to test whether the data support a given solution. They do not rely on the assumption of uncorrelated source time course. They can be compared to other statistical parametric maps such as those used in fMRI.

Results are presented for both simulated and real data. The maps were compared with LORETA and MUSIC results. For the real data consisting of an average of epileptic spikes, we observed good agreement between the EEG statistical maps, intracranial EEG recordings and functional MRI activations.

8.1 Introduction

The electroencephalogram (EEG) is a recording of electrical potentials at a set of electrodes placed on the scalp. Though it is a measurement at the surface of the head, the EEG can still be used in an attempt to infer the location of the neural electrical sources within the brain that produce given EEG potentials, for example epileptic spikes or somatosensory evoked potentials. This localization problem is commonly referred to as the inverse problem of electroencephalography; a similar problem is found for the magnetoencephalogram (MEG). For a review of the electromagnetic inverse problem, see [Baillet 01b].

The inverse problem is ill-posed as there is an infinite number of source configurations that can produce the exact same potential at the surface of the head. However, if one assumes that the sources are made of current dipoles, with a small number of dipoles [Wax 89] and a sufficient number of electrodes [Srinivasan 96, Bénar 01], then the solution is unique [El Badia 00]. The dipolar model [Brazier 49] is a reasonable approximation for focal sources [de Munck 98] (for a review of dipolar methods, see [Cuffin 98]). Unfortunately, even with the dipolar assumption, different combinations of sources can produce very similar potential patterns. As a consequence, the inverse problem of EEG becomes even more ambiguous in the presence of background noise that can be higher than the difference between the source potentials.

Several groups have investigated the ambiguity of the electromagnetic inverse problem. For a given number of sources, the uncertainty in the source location parameters was quantified by defining confidence intervals. This has been done using the Cramer-Rao bound [Mosher 93, Radich 95, Muravchik 01], simulated data [Kobayashi 03] or implanted sources [Cuffin 01a]. Determining the number of sources is a difficult topic though, and this additional ambiguity has also been a subject of much interest. For example, classical goodness of fit and model order selection criteria can help in selecting which model is better suited to the data [Supek 93, Waldorp 02]. A widely used approach is based on principal component analysis (PCA, [Joliffe 86]). Indeed, for uncorrelated source time courses, the number of sources is given by the number of singular values significantly greater than zero. These non-zero singular values define a signal subspace that can be scanned using one dipole only; peaks in the scan indicate possible locations of the sources; this is the MUSIC method [Mosher 92]. Drawbacks of the PCA decomposition are that it can be difficult to define the dimension of the signal subspace, and that PCA fails to find the number of sources when the time courses are fully correlated (even though a refinement of the MUSIC method has been proposed for handling fully correlated sources [Mosher 98]). Another way to explore the brain volume is to use spatial filtering, or “beamforming” [Sekihara 01]. The efficiency of beamforming is also sensitive to the level of correlation of the sources, although it has been shown to be robust for a medium level of correlation [Sekihara 02].

A sensible way to reduce the ambiguity in the location parameters is to incorporate constraints originating from anatomical considerations, such as forcing sources to be distributed on the cortical surface and oriented perpendicular to it [Dale 93]. Another type of constraint involves the use of data from other functional imaging modalities, such as functional magnetic resonance imaging (fMRI) [Liu 98, Trujillo-Barreto 01].

As the electromagnetic inverse problem is ambiguous, some authors have advocated an approach that avoids considering only a single “best” solution, but rather estimates a probability distribution of solutions [Clarke 89, Schmidt 99]. They used a Bayesian framework that allows the incorporation of *a priori* knowledge in a formal way. In [Schmidt 99], the posterior distribution is sampled with the Markov Chain Monte Carlo (MCMC) technique; in [Kincses 03] this approach is implemented within a maximum likelihood framework. The beamforming technique can be also be used in order to built statistical parametric maps [Barnes 03].

In this paper, we investigate the ambiguity in both number and localization of dipolar sources, by exploring all the combinations of one to three dipolar sources on a coarse grid. This exhaustive search is made possible by limiting the resolution and the number of sources. We give to each combination of sources a score based on a statistical measure that reflects how well it describes the data. This is to be contrasted with the methods that first establish a best-fit solution and then estimate a confidence interval on the parameters of the solution [Mosher 93, Radich 95, Muravchik 01]. We then build a three-dimensional map assessing for each elementary volume (or voxel) the likelihood that it contains a source by summing the scores of all combinations containing this voxel[1].

The systematic exploration of the solution space is usually referred to as "dipole scanning". The advantage of a scan with multiple sources is to permit the recovery of sources with perfect correlation, and also to avoid the potential local minima that exist in minimization-based approaches. The integration of the results when there is more than one source has not to our knowledge been reported. We propose a method for integrating at each point the results over all possible combinations, whereas the multiple-source scanning of [Mosher 98] is searching for an optimum combination. The attribution of a score at each point of the head resembles the MUSIC scan [Mosher 92]. However, we perform statistical measures that allow for the assessment of the significance of the scores, contrary to the MUSIC measure.

Our statistical approach is related to that of [Schmidt 99], but differs in several aspects. First, we use simple statistical tests in a frequentist framework. Second, it is an easier method to implement as it does not require MCMC computations. Third, we introduce in our scoring scheme a model order criterion that helps to reject additional sources that merely describe noise. Furthermore, we do not rely solely on a parametric description of the noise distribution that would make the statistical tests very dependent on the (difficult) estimation of the noise covariance matrix. Instead, we use empirical techniques in order to

---

[1] The term "likelihood" is used here in a general sense; this is not to be confused with the more strict statistical meaning, i.e. the probability of the observations given the parameters, seen as a function of the parameters.

determine the distributions of our statistics.

In section 8.2, we introduce the statistical framework and concepts. We present in detail our method in section 8.3. In section 8.4, we introduce the data used for validating the method. In section 8.5, we present the results of the validation with both simulation and real data, and of a comparison with the LORETA [Pascual-Marqui 94] and MUSIC [Mosher 92] methods. The maps obtained for real data are compared to fMRI activations and intracranial EEG recordings belonging to the same patient.

8.2 Statistical Framework

8.2.1 The Linear Model

We assume that the EEG signal that we observe can be modelled by a set of s dipolar sources at locations $\theta_i = (x_i, y_i, z_i)$, $i \in \{1, ..., s\}$. Each source is represented by a basis of three colocalized orthonormal dipoles describing any possible dipolar activity inside a given region [Mosher 92, Scherg 86]. For each source, we define $\mathbf{x}(\theta_i)$ as the (n channels $\times$ 3 dipoles) matrix whose columns are the potentials created by the corresponding three orthonormal dipoles.

The classical linear model for the EEG signal in the presence of noise is

$$\mathbf{Y} = \mathbf{X}(\Theta)\mathbf{B} + \mathbf{E}, \tag{8.1}$$

where $\mathbf{Y}$ is the spatiotemporal data matrix (n channels $\times$ p time points); in order to enhance the SNR, $\mathbf{Y}$ is generally resulting from the average of k identical events. $\mathbf{X}(\Theta) = [\mathbf{x}(\theta_1) \ ... \ \mathbf{x}(\theta_s)]$ is the sensor array matrix (n channels $\times$ $3s$ dipoles), $\Theta = [\theta_1 \ ... \ \theta_s]$ contains the dipole location parameters, $\mathbf{B}$ contains the dipole waveforms ($3s$ dipoles $\times$ p time points) and $\mathbf{E}$ is the error matrix (n channels $\times$ p time points). In the following sections, we will generally drop the reference to Θ to simplify the notations.

Each channel of EEG data is recorded with respect to a reference signal; this means that the rank of $\mathbf{Y}$ is $min\{n-1, p\}$; this is also reflected in the degrees of freedom for the statistical tests in 8.2.3.

The error matrix $\mathbf{E}$ can be broken down into "pure EEG noise" $\mathbf{E}_{eeg}$, i.e. the error arising from background EEG activity superimposed on the signal, and "model error"

$\mathbf{E}_{model}$, i.e. the part of the signal that has not been explained by the model [Waldorp 02]:

$$\mathbf{E} = \mathbf{E}_{eeg} + \mathbf{E}_{model}. \tag{8.2}$$

Typically, the EEG error $\mathbf{E}_{eeg}$ is assumed to be multivariate Gaussian. The model error $\mathbf{E}_{model}$ is the term we seek to minimize by selecting the correct number and location of sources. The time course $\mathbf{B}$ of the dipoles can be further projected onto temporal basis functions [Dogandžić 00]; this has the advantage of taking into account known temporal properties such as smoothness [Baillet 97] and of reducing the number of parameters to estimate.

The classical statistical tests for the general linear model (8.1) require that the error $\mathbf{E}$ be white and Gaussian [Srivastava 02]. The Gaussian assumption is reasonable for averaged EEG data with a high number of events k. However, the hypothesis of whiteness is strongly challenged as there is a high level of both temporal and spatial correlation in EEG data. An option is to preprocess the data to render the noise as uncorrelated as possible, using the covariance matrix of the noise $\mathbf{\Sigma}$ ("prewhitening" step, [Beucker 97, de Munck 00]). Equivalently, one can incorporate $\mathbf{\Sigma}$ directly into the model and use the generalized least squares method [Sekihara 99, Lutkenhoner 98b]. Using the information from $\mathbf{\Sigma}$ is important both during the fitting of the model, to avoid overfitting the noise, and during the inference step, to be able to model parametrically the noise properties (typically as a normal distribution). However, the estimation of $\mathbf{\Sigma}$ from data with a high level of spatio-temporal correlation is a difficult operation. Indeed, a full $\mathbf{\Sigma}$ describing all possible covariances has a dimension of $np \times np$ and has $np(np + 1)/2$ free elements. It is possible to reduce considerably the number of free elements in $\mathbf{\Sigma}$ by assuming that it is the Kronecker product of an $n \times n$ spatial covariance matrix $\mathbf{\Sigma}_S$ and a $p \times p$ temporal covariance matrix $\mathbf{\Sigma}_T$ [de Munck 00]. Still, in order to estimate reliably the spatial covariance, one needs at least $n(n+1)/2$ independent $n \times 1$ time points. This can be quite demanding, especially given the constraint of having a stationary EEG for the estimation procedure. An option is to perform more "robust" estimation by assuming a further constrained structure, such as a covariance decaying with distance [Huizenga 02].

8.2.2 Estimation of parameters

We assume as in [de Munck 00] that the prewhitening process can be separated into a spatial matrix $\mathbf{W}_S$ and a temporal matrix $\mathbf{W}_T$. The prewhitened signal is then $\tilde{\mathbf{Y}} = \mathbf{W}_S\mathbf{Y}\mathbf{W}_T$. The spatial transformation must also be applied to the array matrix: $\tilde{\mathbf{X}} = \mathbf{W}_S\mathbf{X}$.

Under the assumption that $(\tilde{\mathbf{X}}^T\tilde{\mathbf{X}})$ is full rank, an estimate $\hat{\mathbf{B}}$ of the (temporally prewhitened) source amplitude parameters can then be computed by ordinary least squares [Srivastava 02]:

$$\hat{\mathbf{B}} = (\tilde{\mathbf{X}}^T\tilde{\mathbf{X}})^{-1}\tilde{\mathbf{X}}^T\tilde{\mathbf{Y}}. \tag{8.3}$$

The matrix of residuals is formed by projecting the prewhitened data onto the subspace orthogonal to that spanned by the columns of $\tilde{\mathbf{X}}$:

$$\mathbf{R} = \tilde{\mathbf{Y}} - \tilde{\mathbf{X}}\hat{\mathbf{B}} = (\mathbf{I} - \tilde{\mathbf{X}}(\tilde{\mathbf{X}}^T\tilde{\mathbf{X}})^{-1}\tilde{\mathbf{X}}^T)\tilde{\mathbf{Y}}. \tag{8.4}$$

8.2.3 Model Testing

Goodness of fit

One may want to assess if the model accounts properly for the signal, i.e. if residuals are only noise. Typically, the residuals $vec(\mathbf{R})$ are hypothesized to be distributed with a multivariate normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, with $\mathbf{0}$ $np \times 1$ vector of zeros and $\mathbf{I}$ $np \times np$ identity matrix.

The residuals sum of squares, or "sum of squared errors", is the square distance between the model and the data:

$$SSE = tr(\mathbf{R}^T\mathbf{R}). \tag{8.5}$$

Under the hypothesis of normally distributed residuals, (8.5) has a χ^2 distribution with d.o.f. $df(s) = (n-1-3s)p$.

Model Order

Increasing the order of the model, i.e. the number of sources, is likely to result in a reduction of the sum of squared errors (8.5), as the residuals are formed by projecting the data onto a smaller subspace in (8.4). It is therefore important to test whether this reduction is significant, i.e. if the reduction is larger than that expected if the additional source was only explaining noise. For a given combination of s sources, and a reduced model with $(s-1)$ sources, the null hypothesis $H_0^{(order)}$ is that the portion of the sum of squares explained by the additional source has a χ^2 distribution. A simple test of model order uses the following statistic:

$$F = \frac{df(s)}{df(s-1) - df(s)} \frac{SSE^{s-1} - SSE^{s}}{SSE^{s}}, \tag{8.6}$$

where SSE^s is the sum of squared errors for the fit with s sources and SSE^{s-1} is the sum of squared errors for the reduced model. Under the null hypothesis $H_0^{(order)}$, (8.6) has an F distribution with d.o.f. $(df(s-1) - df(s)) = 3p$ and $df(s) = (n-1-3s)p$.

8.3 Construction of Statistical Maps of Activation

In order to construct a map for s sources, we consider all the combinations of s sources on a grid. For each combination, we compute an F test using spatially prewhitened data, that assesses if the data support the hypothesis of s sources versus $(s-1)$. We then compute a score at a given point of the grid by integrating the scores of all the combinations containing this point that pass a threshold of significance. The thresholds are computed using bootstrap resampling.

8.3.1 Prewhitening

We perform spatial prewhitening using an estimate of the spatial covariance matrix,

$$\hat{\mathbf{\Sigma}}_S = \mathbf{Z}\mathbf{Z}^T, \tag{8.7}$$

where $\mathbf{Z}$ is a (n channels $\times$ L time points) spatiotemporal matrix containing averaged background EEG. We assume that the averaged background noise was stationary across the matrix $\mathbf{Z}$. The size L of the background window is a compromise: it needs to be large in order to perform a correct estimation of the noise properties, but choosing too large a window increases the chances of incorporating spikes in the background or having a non-stationary window [Cohen 77].

The matrix $\hat{\mathbf{\Sigma}}_S$ is broken down based on a Singular Value Decomposition,

$$\hat{\mathbf{\Sigma}}_S = \mathbf{U}\Lambda\mathbf{U}^T, \tag{8.8}$$

with

$$\Lambda = \begin{pmatrix} \lambda_1 & & (0) \\ & \ddots & \\ (0) & & \lambda_n \end{pmatrix} \tag{8.9}$$

matrix of singular values and $\mathbf{U}$ the matrix of eigenvectors. The spatial prewhitening filter is then defined as

$$\mathbf{W}_S = \begin{pmatrix} \frac{1}{\sqrt{\lambda_1}} & & (0) & 0 \\ & \ddots & & \vdots \\ (0) & & \frac{1}{\sqrt{\lambda_{n-1}}} & \vdots \\ 0 & \dots & \dots & 0 \end{pmatrix} \mathbf{U}^T. \tag{8.10}$$

We set the last diagonal element in (8.10) to zero as the rank of $\mathbf{Z}$, and thereby that of $\hat{\mathbf{\Sigma}}_S$, is (n-1). Indeed, each channel of EEG data is always recorded versus the signal of the same reference electrode. We do not perform any temporal prewhitening, i.e. $\mathbf{W}_T = \mathbf{I}_p$.

8.3.2 Testing the Combinations

The statistical maps are constructed on a spatial grid of N points $P_j = (x_j, y_j, z_j), j \in \{1...N\}$. In order to obtain a map for s sources, $s \in \{1, 2, 3\}$, we consider all combinations of s grid points with parameters Θ_k^s, $k \in \{1...C_N^s\}$, where $C_N^s = (N!)/(s!(N-s)!)$. For each combination, we create an sensor array matrix $\mathbf{X}(\Theta_k^s)$ (cf. 8.4.1) and perform a least square fit with equation (8.3). For this combination, we compute the sum of squares errors $SSE(\Theta_k^s)$ with (8.5) and the F statistic $F(\Theta_k^s)$ with (8.6). We define Θ_{min}^s as the parameters that give the lowest sum of squared errors across all combinations of s sources, i.e. the best fit:

$$\Theta_{min}^s = \arg\min_{\Theta_k^s}(SSE(\Theta_k^s)). \tag{8.11}$$

For a given combination of s sources with location Θ_k^s, the F value (8.6) is defined versus the combination with one less source that gave the lowest sum of squared errors: $SSE^{s-1} = SSE(\Theta_{min}^{s-1})$. For the one-dipole scan, we use the data total sum of squares as

SSE^{s-1}. *We are thereby testing that the combination with* s *sources is an improvement on the best solution with* $(s-1)$ *sources.*

8.3.3 Computing the Scores

The score for the combination Θ_k^s is defined as:

$$\sigma_k = \begin{cases} F(\Theta_k^s) & \text{if } F(\Theta_k^s) \geq thr_F(s) \\ 0 & \text{otherwise.} \end{cases}, \ k \in 1...C_N^s, \tag{8.12}$$

with $thr_F(s)$ the significance threshold at the order s (cf. 8.3.4).

At each point P_j of the grid, $j \in \{1...N\}$, we integrate the scores of all the combinations containing this point and normalize by the number of combinations:

$$map(P_j, s) = \frac{\sum\limits_{\{k|P_j \in \Theta_k^s\}} \sigma_k}{C_{N-1}^{s-1}}, \tag{8.13}$$

which reflects the likelihood that this grid point contains a source under the hypothesis of s sources.

8.3.4 Computation of Thresholds

In order to compute the thresholds for the statistical tests, we estimate empirically the distribution of the F statistic (8.6) under the null hypothesis (i.e. the additional source is only explaining noise). To do so, we construct a set of 5000 realizations $\mathbf{E}_b^*, b \in \{1...5000\}$ of realistic averaged EEG noise $\mathbf{E}_{eeg}$.

Each realization $\mathbf{E}_b^*$ of the noise is obtained by drawing at random with replacement sections of 20 consecutive points (100 ms) from the averaged background $\mathbf{Z}$ (i.e. one could obtain in one realization several copies of the same section). Enough sections are joined together in order to obtain a matrix $\mathbf{E}_b^*$ with the same number of columns as the original signal of length p. This is the "moving block bootstrap" for serially correlated data, as described in [Efron 93]. Our assumptions are that the averaged background is stationary and sufficiently long to represent well the variability of the underlying stochastic process, and that the length of the sections is sufficient to preserve the temporal structure of the process. This implies that each realization $\mathbf{E}_b^*$ has the same statistical properties as the original averaged EEG background.

We fit the model to each realization as in (8.3):

$$\hat{\mathbf{B}}_b^{s,*} = [\tilde{\mathbf{X}}(\Theta_{min}^s)^T \tilde{\mathbf{X}}(\Theta_{min}^s)]^{-1} \tilde{\mathbf{X}}(\Theta_{min}^s)^T \tilde{\mathbf{E}}_b^*, \tag{8.14}$$

where $\tilde{\mathbf{E}}_b^* = \mathbf{\Sigma}_S \mathbf{E}_b^*$ is the prewhitened noise, and $\Theta = \Theta_{min}^s, s \in \{1, 2, 3\}$ are the parameters corresponding to the combination of sources that produced the lowest sum of squared errors when fitted to the data (8.11). We then compute the statistics $SSE_b^*(s)$ using (8.5) and $F_b^*(s)$ using (8.6), with $SSE^{s-1} = SSE_b^*(s-1)$. For $s = 1$, we use the sum of squares of $\tilde{\mathbf{E}}_b^*$ as SSE^{s-1}. At each s, we compute the empirical distributions (histograms) of $F_b^*(s), b \in \{1...5000\}$.

In a second step, the estimated thresholds need to be corrected for multiple comparison. However, the empirical distributions of $F_b^*(s)$ typically have an insufficient sampling of the tails for this procedure. Therefore we fit to each empirical distribution a theoretical F distribution by varying the corresponding degrees of freedom; we use a non-linear minimization method (simplex algorithm). We then compute the thresholds $thr_F(s)$ corresponding to the $p = 0.05$ percentile corrected with the Bonferroni method, i.e. $p = 0.05/C_N^s$.

8.4 Evaluation of the Method

We analyzed three data sets in order to evaluate the capacities of our method for localizing sources and assessing dipolar models. The first set was an idealized configuration consisting of two simulated dipoles located on the grid used for scanning, with potentials corrupted by Gaussian noise correlated in space but not in time. The second data set was a more realistic situation, where we simulated extended patches of cortex, with real EEG noise added. The third set was a real average of epileptic spikes, obtained on a patient for whom we had functional MRI and intracerebral EEG results.

For each data set, we computed on a 10 mm grid the statistical maps for one to three dipoles, as well as a MUSIC scan and a LORETA current density reconstruction. The MUSIC and LORETA methods are the versions implemented in the Curry software (Neuroscan Labs). For MUSIC, we included in the signal subspace the eigenvectors corresponding to the singular values above the level of the noise.

8.4.1 Model Computations

We computed all possible components of the array matrix **X** with the Curry 4.5 software (Neuroscan, El Paso, TX). The computations used a boundary element method (BEM) realistic head model based on the subject's own MRI scan, with BEM surfaces corresponding to the brain (7 mm mesh), skull (10 mm mesh) and skin (12 mm mesh) [Hämäläinen 89]. Conductivities were set to $0.33\ \mathrm{S} \cdot \mathrm{m}^{-1}$, $0.0083\ \mathrm{S} \cdot \mathrm{m}^{-1}$ and $0.33\ \mathrm{S} \cdot \mathrm{m}^{-1}$ respectively (ratio of skull to brain of 1/40). We created a uniform square grid inside the brain volume with 10 mm spacing; points corresponding to deep brain structures were not included. To enhance computation speed for the scan for $s = 3$ sources, a low-resolution 15 mm spacing grid was also created. For each point of the grids, we computed the potentials generated by three unit orthogonal dipoles. These potentials were also referenced to the average.

8.4.2 Idealized Simulation

We computed the potentials of two simulated radial dipoles located symmetrically in the left and right central regions. We used the head model built for the patient data (cf 8.4.1). The time course of both dipoles was a half-period of a sine wave, lasting 40 samples; the two dipoles were perfectly synchronous. The background noise (1000 samples) was generated by assigning to each unit dipole on the 10 mm grid a pseudo-random amplitude following a Gaussian distribution, similarly to [de Munck 92]. This ensured spatial correlation and temporal whiteness of the noise. The potentials created by the sources were scaled in order to obtain a signal to noise ratio of 10 (ratio of total sum of squares across time and channels). Figure 8.1(a) shows a butterfly plot of the data.

8.4.3 Realistic Simulation

We created an $n \times p$ $(n = 43, p = 34)$ spatiotemporal data matrix by adding the potentials generated by two simulated sources to real EEG background noise sampled at $F_s = 200$ Hz. The first source was placed in the lateral part of the temporal lobe. The second source was in the frontal region. The time course of the first source was obtained by adding three consecutive gamma functions peaking at 40 ms, 110 ms and 250 ms respectively. The parameters of the gamma functions were obtained by fitting the real average spike presented in the next section. The time course of the frontal source was obtained by stretching that

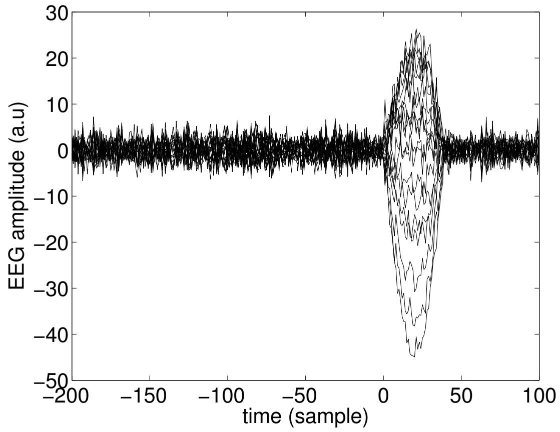

(a) EEG data for the idealized simulation (two central radial dipoles on the grid); only the 10-20 channels and 200 samples of background are shown. The SNR is 10 (ratio of sum of squares across all channels and time points).

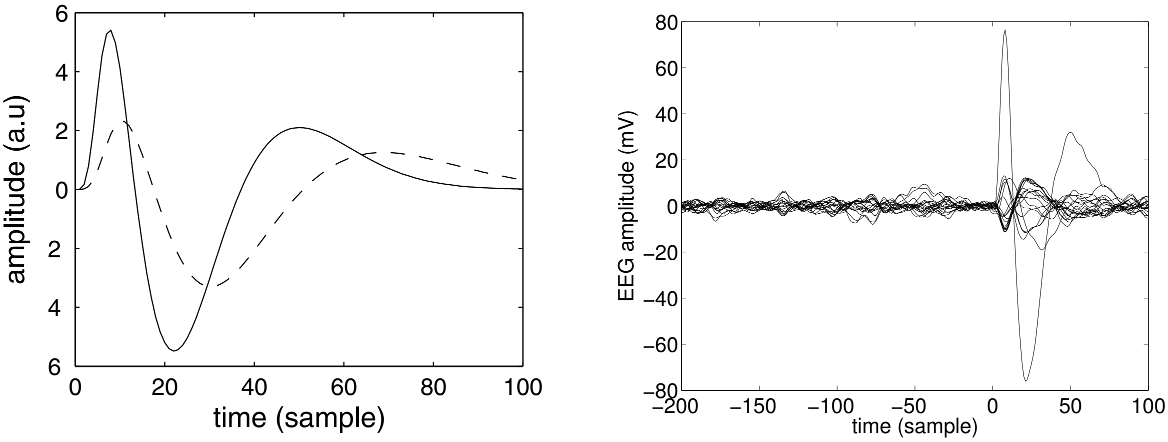

(b) Source time courses for the realistic simulation (one temporal patch and one frontal patch). Solid line: temporal source, dashed line: frontal source. The first peak of the frontal source is delayed by 3 samples (15 ms). The time courses were scaled to have the same energy.

(c) Data for the realistic simulation; only the 10-20 channels and 200 samples of background are shown. The SNR is 50.

Fig. 8.1 Simulated data: temporal characteristics. The sampling frequency is 200 Hz.

of the first source in order to have the first peak delayed by 15 ms, which is a reasonable value for neuronal propagation (cf. Fig. 8.1(b)).

We simulated each realistic source as a patch of cortex. We used a 2 mm mesh of the cortical surface, obtained from the MRI of the patient presented in next section, and placed a dipole at each vertex in the direction normal to the mesh. A patch was defined as a set of dipoles within a sphere of radius 1.3 cm (i.e. of section $\pi \cdot 1.3^2 \simeq 5.3$ cm^2). The background noise was obtained by adding 20 sections of real EEG with no spike obtained from the patient described in next section. The sum of the source potentials was scaled so as to have a signal to noise ratio of 100 (ratio of total sum of squares across time and channels). Figure 8.1(c) shows a butterfly plot of the data.

8.4.4 Patient Data

The EEG data consisted of a 43-channel spatio-temporal matrix resulting from the average of $k = 23$ epileptic spikes sampled at a frequency $F_s = 200$ Hz, from a patient with focal epilepsy. The spikes were chosen to be as similar as possible, in terms of spatial distribution and temporal waveform. The spikes were aligned by cross-correlation before averaging. The EEG was filtered with a band-pass filter set from 1.6 Hz to 35 Hz. The averaging and filtering was performed with the BESA software (Megis, Munich, Germany). The data was "referenced to the average", i.e. the average of all channels was subtracted from the data at each time point. We defined a signal window of $p = 34$ samples (170 ms), i.e. the portion of interest of the averaged EEG spike, and a background noise window of $L = 1000$ samples (5 s), i.e. an average of $k = 23$ EEG sections preceding each spike and not containing spikes. The first point of the background window was set to 2000 samples (10 s) before the spike, as the background immediately preceding the spike was contaminated with spikes. The signal window was chosen to encompass the spike itself and the slow-wave that follows (Fig. 8.2). The signal to noise ratio was 296.5 (ratio of total sum of squares across time and channels).

For this patient, we also had intracranial EEG recordings, anatomical and functional MRI data. The intracranial recordings were performed during presurgical evaluation. The fMRI protocol consisted of recording 19 EEG channels (10/20 system) inside the MR scanner (details in [Bénar 03]). The timing of the EEG spikes was used for the statistical analysis of the fMRI images. We then recorded the EEG outside the scanner, adding 24 electrodes (10/10 system) for a total of 43 electrodes. It is the recording outside the scanner that we used for computing the data matrix $\mathbf{Y}$. As the 10/20 electrodes were visible on the anatomical scan recorded during the fMRI protocol, we could mark them onto the realistic head model. The remaining 10/10 electrodes were placed manually on a 3D rendering of the head surface.

8.4.5 Computation Time

The computation time for each scan is a function of C_N^s. In order to reduce the computation time for the three-sources scan, we used a multi-resolution approach to restrict the search to a subset of the grid. We first computed the tests on a 15 mm grid, and retained the points in the 10 mm grid whose nearest neighbour in the 15 mm grid had a non-zero value.

The total computing time for the scans on the 15 mm grid (311 points) was 1 h 04 min. For the idealized simulation (cf. 8.4.2), the number of selected points in the 10 mm grid was 251 out of 836 (30 %). The total computing time for the 10 mm scans, with the three-sources scan on the restricted grid, was 1 h 37 min. For the realistic simulation (cf. 8.4.3), 68 points (8.1 %) were preselected. The computation time for the scans on the 10 mm grid was 4 min 59 s. For the patient data, the number of selected points was 82 out of 836 (9.8 %). The total computing time for the 10 mm scans was 4 min 37 s.

We used a Pentium M laptop with a processor speed of 1.4 GHz and 512 Mb of RAM. The algorithm was implemented with the Matlab software (Mathworks, Natick, MA).

8.5 Results

8.5.1 Empirical Distributions of F

Fig. 8.3 displays the empirical distributions of $F_b^*(s)$ (8.6) for one to three sources, for both simulated and patient data, along with the theoretical F distributions (corresponding to the assumption of white gaussian residuals) and the fitted F distributions (cf. section 8.3.4).

For the idealized simulation, the empirical distributions are very close to the theoretical distribution, showing that the spatial prewhitening was efficient in removing the spatial correlations. For both realistic simulation and the patient data (fig. 8.3(b) and 8.3(c)), the empirical distributions are broader than the theoretical F distribution. This is likely to

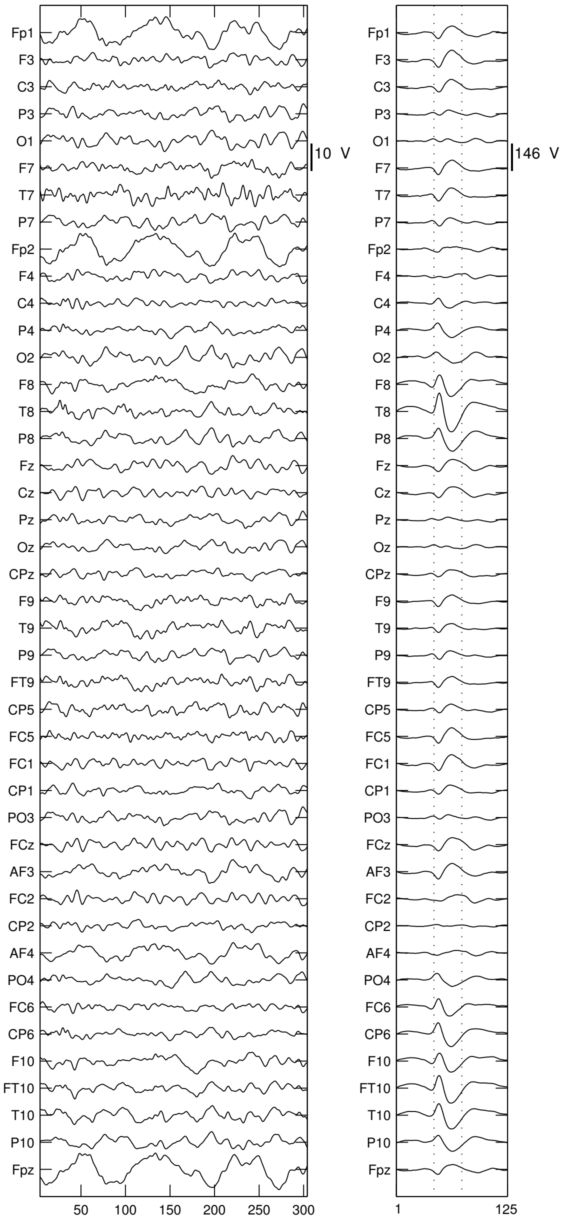

Fig. 8.2 Two sections of 43-channel EEG, created by averaging 23 epileptic spikes (patient data). Left panel: a section of averaged background. Right panel: the average spike, with maximum amplitude in the right temporal region (channels T8, T10, FT10); the dotted lines mark the section selected for computing the maps. Time is in samples (sampling frequency: 200 Hz). The SNR is 296.5. Note that the amplitude scale is different between panels.

reflect the fact that there is still some correlation left in the data - thereby decreasing the effective degrees of freedom. Indeed, we did not perform temporal prewhitening in order to preserve the signal [de Munck 00] and the spatial covariance was estimated on a limited number of time points in order to ensure stationarity and avoid contamination by spikes.

The fitted distributions generally represent well the empirical distributions. However, for the real background EEG, there is a tendency for the empirical distribution to have a thicker tail than the fitted distribution, which is more prominent for 3 sources (fig. 8.3(b) and 8.3(c)). This may suggest that in our examples the real EEG background is not perfectly Gaussian.

8.5.2 Maps for Idealized Simulation

The maps resulting from processing the idealized simulated data (cf. 8.4.2) are shown in Fig. 8.4. The grid points with non-zero values permit to assess the extent of the regions where the hypothesis of a dipolar source is supported by the data.

The one-source map, which is simply the significant F scores, displays the grid points where one dipole explains a significant portion of the data. There is a large number of points with a high score in both left and right regions, with a right predominance that likely reflects the fact that this source contributed more to the signal (as the two sources are not perfectly symmetrical).

For the two-source map, which represents the integration of the scores of all combinations of two dipoles, the maps present two local peaks at the correct dipole locations ($z = 187$). The maps are more focused around the true dipole locations, which could be because the criterion (significance of two dipoles versus one) is stricter than for the one-source map (significance of one dipole versus none). Another effect is that the two-source map is created by averaging over a large number of combinations, which favours points that are present in a large number of combinations.

The maximum F score for two sources is 153, which shows that the data strongly supports the hypothesis of two dipolar sources versus one dipolar source. In contrast, the map for three sources is blank, which shows that the method correctly rejected the hypothesis of three dipolar sources.

The MUSIC scan (signal subspace of dimension one) finds one peak in the middle of the simulated dipoles, as expected because of the perfect correlation of the sources. The

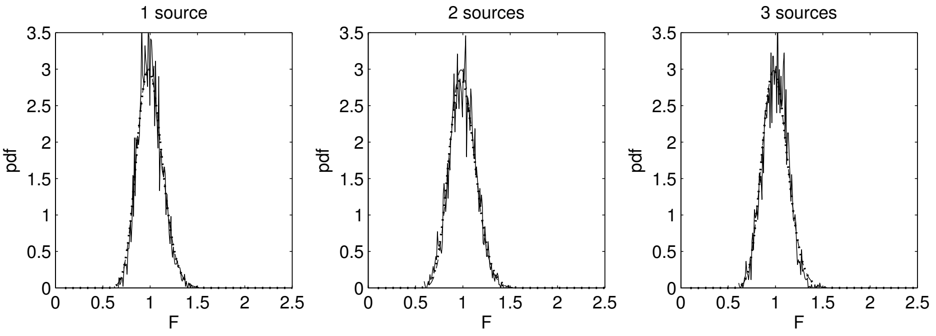

(a) Idealized simulation.

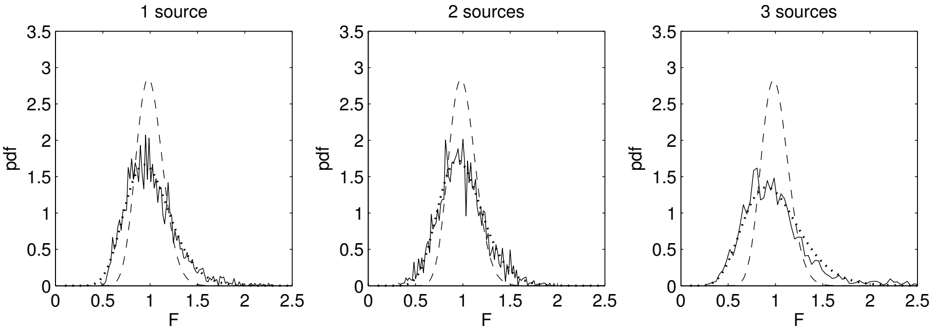

(b) Realistic simulation.

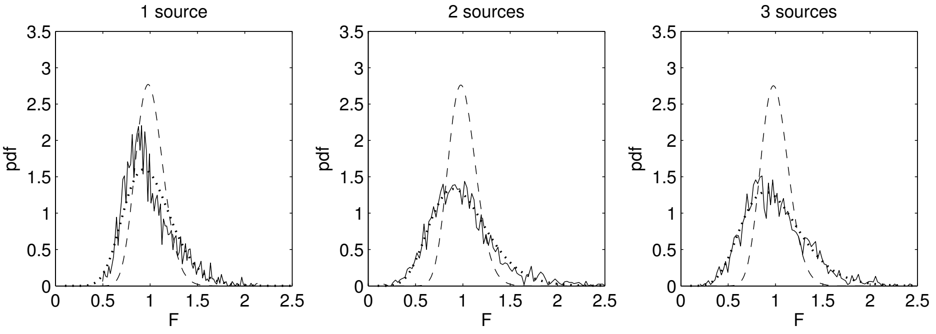

(c) Patient data.

Fig. 8.3 Empirical distributions of the F statistics for one to three sources. On each graph, the theoretical distributions for white gaussian noise is shown in dashed line. The fitted F distribution are in dotted lines. For the realistic simulation and the patient data, the empirical distribution are broader than the theoretical distributions, due to the remaining correlations in the EEG data.

LORETA method correctly identifies two peaks, but at a lower z-value. This lower value of z could be a consequence of the smoothness constraint or of the partial prewhitening (only the diagonal of the covariance matrix was used).

We present in figure 8.5 the time courses reconstructed by considering a dipolar source at each local maximum of the two-source map. The time courses correspond very well to the simulated waveform.

8.5.3 Maps for Realistic Simulation

The maps resulting from processing the realistic simulated data (cf. 8.4.3) are shown in Fig. 8.6. In the maps for one and two sources, a large number of points have a non-zero value, including the areas corresponding to the sources but also the region lying between the two simulated sources. The points that stand out are all in the left temporal region. This reflects the fact that the temporal source contributes much more to the data than the frontal source (the SNRs for the temporal and frontal sources considered separately are 46.7 and 3.3 respectively).

The map for three sources shows that there are significant combinations of three dipolar sources that can explain the data, although only two sources have been simulated. This can be explained by the fact that the simulated sources are not single dipoles lying perfectly on the grid as in the idealized simulation, and therefore require more than two dipoles to explain fully the signal. Nevertheless, the points with higher values are in the correct regions. The match is very good for the temporal source, as the global maximum lies around the location of the center of gravity of the patch. The points are displaced to a lower z-value for the frontal source. The maximum F score is much higher for the two-sources maps than for the three-sources map (48.3 versus 5.8), showing that the impact of going from one source to two sources is higher that when going from two sources to three sources.

Both MUSIC (signal subspace of dimension two) and LORETA identify two regions of activation. For both, the global maxima corresponding to the temporal source ($z = 127$) are displaced to a lower z value with respect to the centre of gravity of the patch. Only MUSIC was able to detect the frontal source at the exact z location ($z = 187$), whereas both LORETA and our method found this frontal source displaced to a lower z-value ($z = 167$ and $z = 177$ resp.). The MUSIC scan is the method that gives the best contrast for the

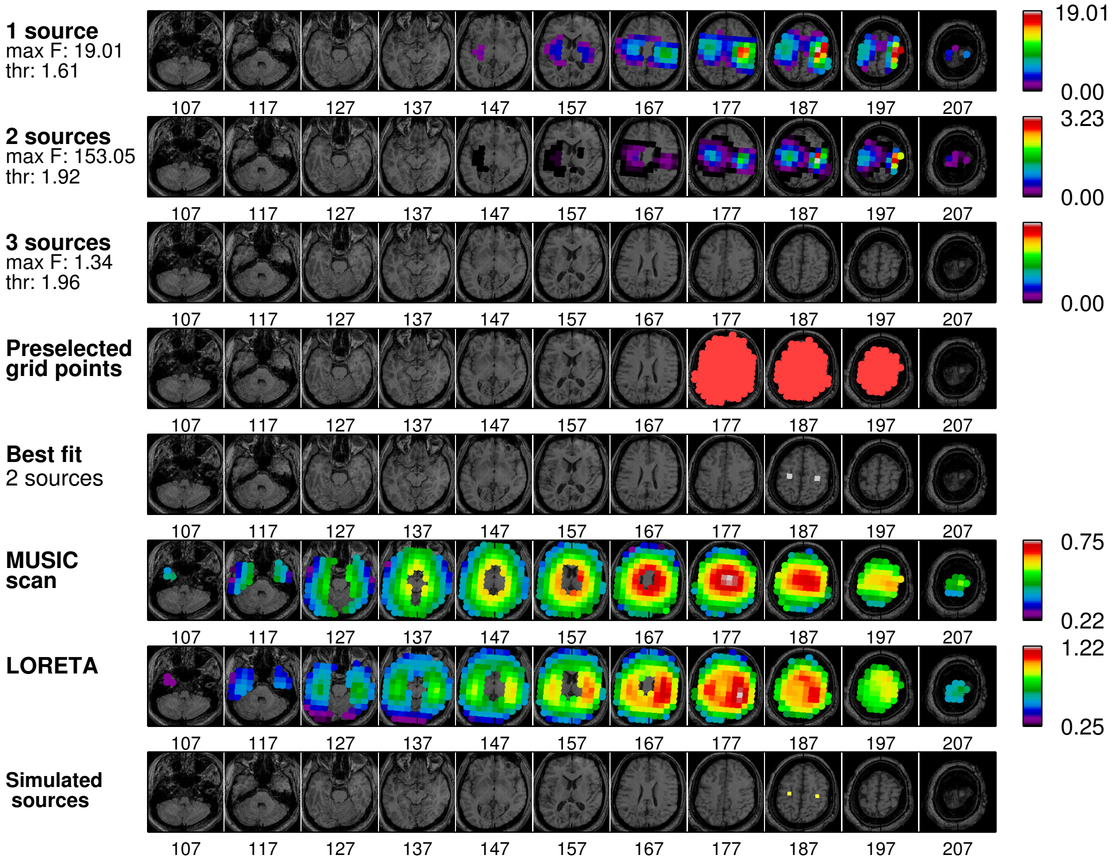

Fig. 8.4 Statistical maps $Map(P_j, s)$ for one, two and three sources for the idealized simulation. The F threshold and the F score corresponding to the combination with the best fit are displayed on the left of each map. In the fourth row, the (manually) preselected grid points are shown. In the last row, the actual locations of the two simulated dipoles are shown. Abscissae: elevation in mm (z axis). Voxels with a value of zero are not shown; those in black correspond to very low non-zero values. The locations of the dipoles are detected as local maxima in the maps for two sources ($z = 139$ and $z = 187$). No combination of three dipoles was significant. The MUSIC 1-dipole scan was unable to separate the sources. LORETA correctly identified two peaks, but at $z = 177$.

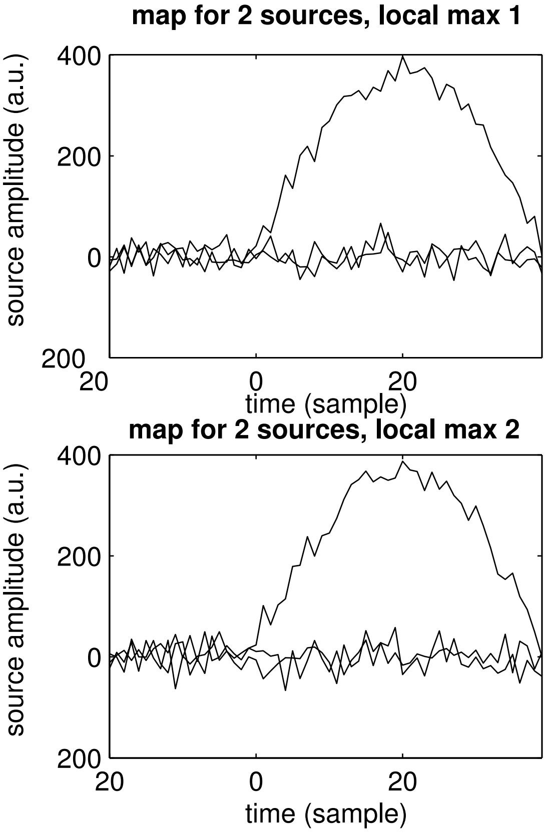

Fig. 8.5 Reconstructed time course for the idealized simulation, for sources placed at the local maxima of the map for two sources. For each source, a basis of three orthogonal dipoles is used. The data displayed is reconstructed using a singular value decomposition of the source time course (this is equivalent to having one dipole of the basis oriented along the direction of higher variance). The time courses of the main component correspond very well to the simulated waveforms for both sources (half sine wave of period 40 samples).

frontal source.

We present in figure 8.7 the time courses reconstructed by considering a dipolar source at each local maximum of the two-sources map. For the temporal source, the time course corresponds very well to the simulated waveform. The waveform of the frontal source is more difficult to retrieve because of the lower SNR of this source. However, it is still possible to see that the source is activated with a delay with respect to the temporal source.

8.5.4 Maps for patient data

Fig. 8.8 presents the maps for the patient data, as well as the intracranial EEG (SEEG) electrodes and the functional MRI t-stat map. The threshold of $t = 3$ for the fMRI map implies that five contiguous points need to be above the threshold in order to have a cluster significant at $p = 0.05$ (corrected, [Cao 99]).

In the map for one source, the right temporal region ($z = 114$) presents the highest values. The activated region extends all the way up to the right superior parietal region. In the maps for two and three sources, the global maximum is still temporal at $z = 124$ and additional local maxima appears in the inferior frontal region ($z = 144$ and $z = 154$ for the two- and three-sources respectively) and in the right parietal region ($z = 164$ for both).

The MUSIC scan (subspace dimension of three) finds local maxima only in the temporal region; however the three regions found with the statistical maps present high values. The LORETA method finds local maxima in the temporal ($z = 127$) and parietal regions, but none in the lower frontal region (even though there are still relatively high values in this region).

There is no SEEG electrode at the level of the temporal global maximum at $z = 124$, but it is located just below activated SEEG electrodes at $z = 134$. The inferior frontal local maximum in the three-sources map ($z = 154$) is at the level of an activated SEEG electrodes; this local maximum corresponds to one dipole of the best fitting combination.

The activated SEEG electrodes at $z = 144$ also correspond to a high value in the three-source map. The SEEG electrode closer to the parietal local maximum is active ($z = 164$). This local maximum does not correspond to any dipole of the best fitting combination. Moreover, it is located at the level of the fMRI cluster ($z = 154 - 164$).

Given the fact that the spatial resolutions of the different techniques (EEG dipole

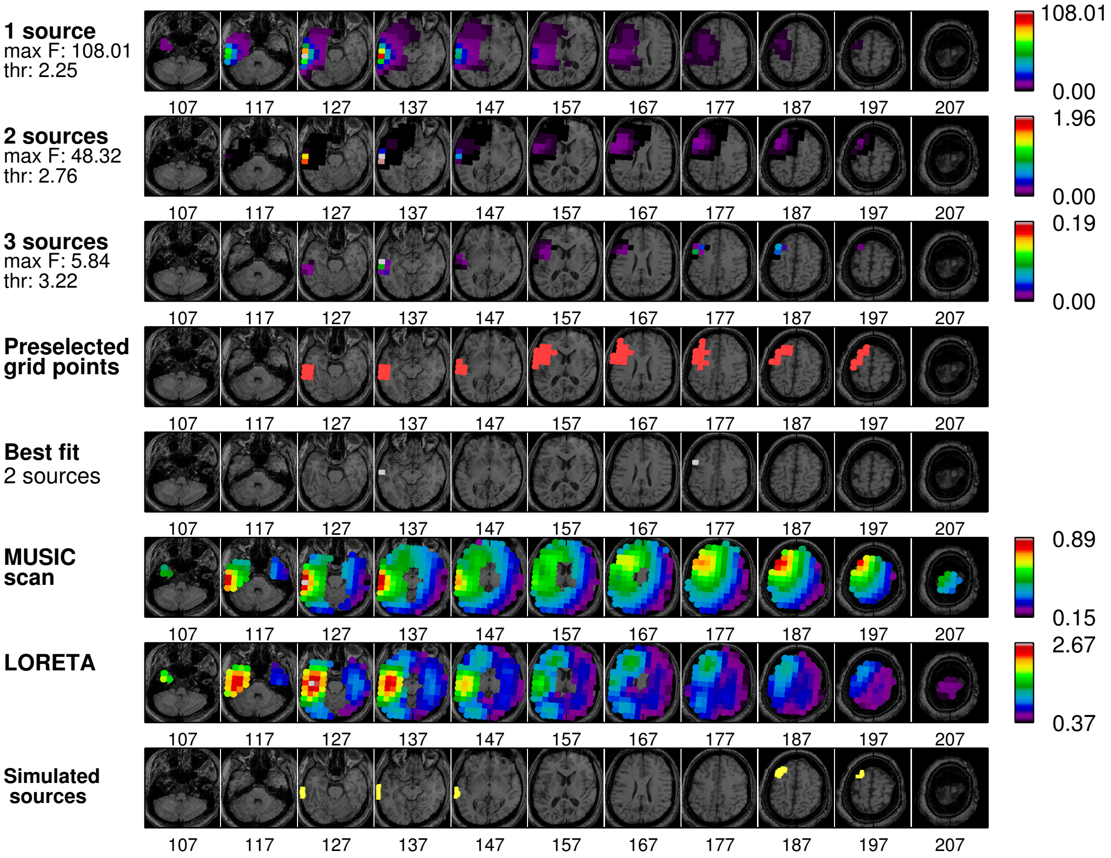

Fig. 8.6 Statistical maps $Map(P_j, s)$ for one, two and three sources for the realistic simulation. The F threshold and the F score corresponding to the combination with the best fit are displayed on the left of each map. In the last row, the actual locations of the two simulated patches are shown. Abscissae: elevation in mm (z axis. Voxels with a value of zero are not shown; those in black correspond to very low non-zero values. The locations of the patches correspond to grid points with high values in the maps three sources ($z = 137$ and $z = 177 - 187$). As the number of sources increase, the maps become more focused towards the center of gravity of the patches. Both MUSIC (signal subspace of dimension two) and LORETA identify two regions of activation. For both, the global maxima corresponding to the temporal source ($z = 127$) are displaced to a lower z value with respect to the centre of gravity of the patch.

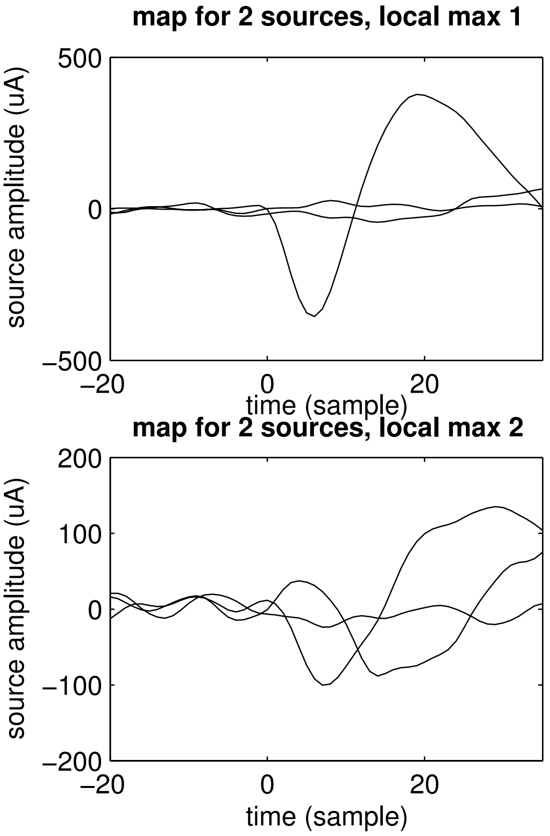

Fig. 8.7 Reconstructed time course for the realistic simulation, for sources placed at the local maxima of the map for two sources. For each source, a basis of three orthogonal dipoles is used. The data displayed is reconstructed using a singular value decomposition of the source time course (this is equivalent to having one dipole of the basis oriented along the direction of higher variance). For the temporal source (top), which has the higher SNR, the waveform is well reconstructed in the main component. For the frontal source (bottom), the waveform is more difficult to visualize because of the lower SNR of this source. The first peak of the main component permits nevertheless to see that this source is activated with a delay.

localization, fMRI at 1.5 T and SEEG) are of the order of 1 cm [Cuffin 01a, Turner 02], our cross-modality results are compatible with one another. The peak in the EEG map that is closest to the fMRI activation is that with the lowest score, and the other peaks do not correspond to any fMRI activation. In particular, this is the case for the main EEG peak in the inferior temporal region, possibily because of fMRI signal loss in this region due to a magnetic susceptibility effect. This suggests that one should be cautious when considering the possibility of a one-to-one correspondence between EEG and fMRI results, as pointed out in [Nunez 00].

8.6 Conclusion

We have presented a method for building statistical maps for EEG source localization. These maps assess for each point in the brain the likelihood that it contains a source by giving it a score based on a statistical test. They allow for the exploration of the solution space in a systematic way. This permits to assess the range of plausible solutions, but can also be useful in order to avoid the local minima that plague the minimization-based approaches. Contrary to a PCA decomposition, there is no need to assume that source time courses are uncorrelated. The maps also permit a comparison of the results obtained under the hypothesis of one, two or three dipolar sources. They are three-dimensional, contrary to a classical dipole solution and can be compared to other statistical parametric maps, such as those used in functional MRI. Indeed, a region that is not a local maximum in the EEG and the fMRI maps but still leads to a significant score in both maps would increase the confidence that this region is active.

We use a test of model order at each level that assesses if the data support the addition of another source. The use of a threshold permits to restrict the result to the combinations where all the sources contribute to the model. Summing up the scores of the combinations at each point enables to visualize the results in a condensed manner. It also enhances the points contained in several plausible combinations, which we assume are more likely to contain a true source. Indeed, less weight should be given to spurious detections that appear only in a few combinations, whereas true sources should be part of a large number of combinations.

The information on the number of sources that is supported by the data is very important in EEG dipole modelling. Our results for the realistic simulation point to a limitation

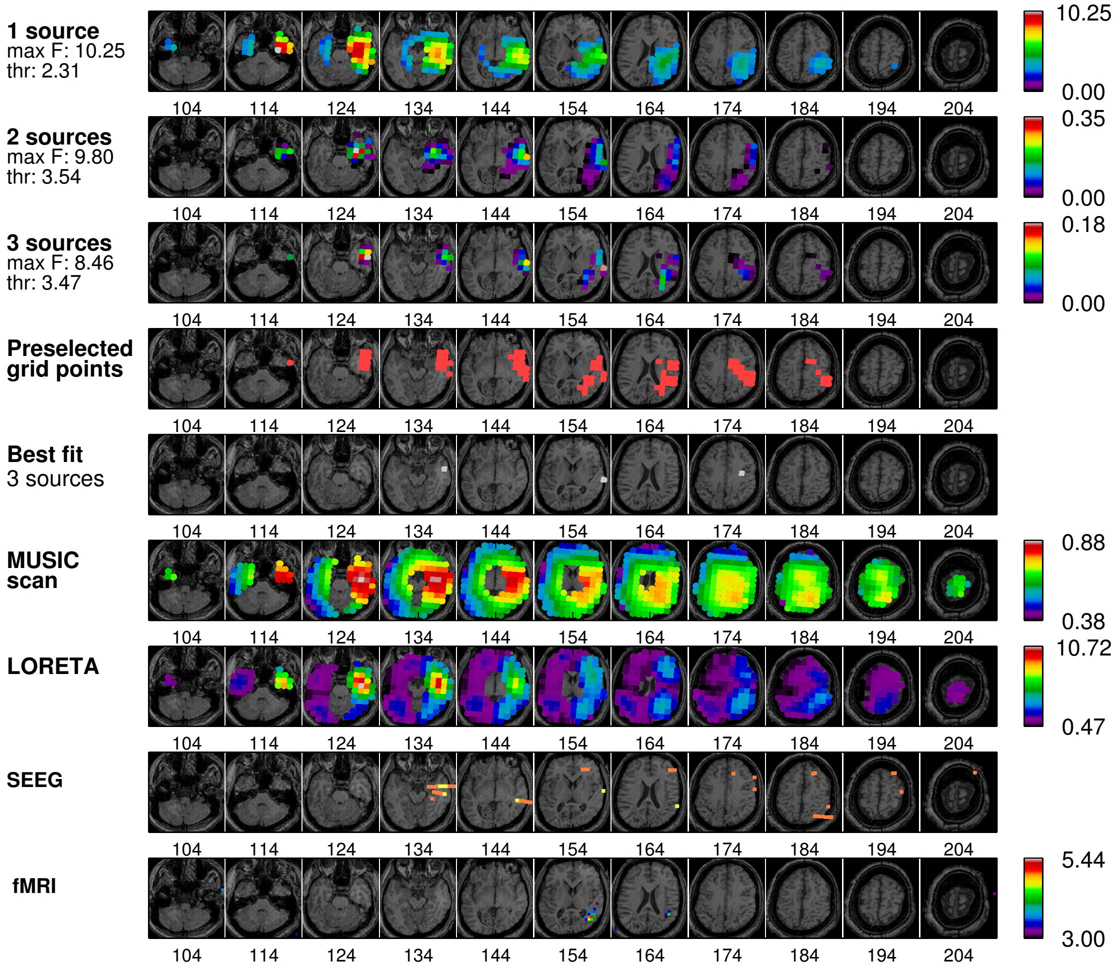

Fig. 8.8 Statistical maps $Map(P_j, s)$ for one, two and three sources for the patient data, compared with MUSIC and LORETA results, depth EEG recordings(SEEG, electrodes with spiking activity in yellow, other electrodes in orange) and functional MRI (t-stat map). The F value corresponding to the combination with the best fit and the corresponding threshold are displayed on the left of each map. In the fourth row, the grid points that have been preselected are shown. Abscissae: elevation in mm (z axis). For the EEG maps, the value at each point is the mean score across all combinations containing the point. For the fMRI map, the values are t-statistics. In the three-source map, the global maximum at $z = 124$ in the anterior temporal region is located below activated intracerebral electrodes at $z = 134$. The local maximum at $z = 154$ in the inferior frontal region corresponds well to an activated subdural electrode at $z = 154$. The local maximum in the parietal region at $z = 164$ is just above the fMRI main cluster of activation ($z = 154$). The MUSIC scan finds local maxima only in the temporal region. The LORETA method finds local maxima in the temporal ($z = 127$) and parietal region, but none in the lower frontal region.

of a dipolar scan in that sense. Indeed, as real sources are not lying on a grid and are not point sources, but rather extended patches of cortex, the F test of model order results in a number of dipoles larger than the number of activated patches. We expect this limitation to have also a strong impact on the sum of squares (8.5), as a large number of sources will be needed to bring the residuals below the level of the noise.

The localization results we obtained in both simulated and real data are very encouraging. In the simulated data, the maps presented peaks consistent with the locations of the sources. In the real data, the peaks in the scan corresponded well with both depth EEG and fMRI results. Quantitative evaluation on real data remains difficult, though, due to the limited spatial sampling of SEEG. Further simulations will be required with different combinations of sources and noise levels, as well as more tests on patient data.

We have compared our localization results to those obtained with two well-known methods, namely a MUSIC scan and LORETA, which also produce three-dimensional maps of activity. On the data presented, our method compares well with MUSIC and LORETA in terms of detection capacities. Only our method was able to detect the fully correlated sources. The variation of MUSIC for correlated sources proposed in [Mosher 98] would probably have been able to detect these sources too; however we also display with our method the three-dimensional extent of plausible dipolar solutions. Our statistical approach could be compared to the noise-normalized implementations of the distributed sources methods [Dale 00, Pascual-Marqui 02], which could be more suited for extended sources.

Generally speaking, all these methods should not be seen as antagonist, but as asking different questions with respect to the data. MUSIC and LORETA aim at finding the best solution in a given sense, the sources most correlated to the signal subspace for MUSIC and the spatially smoothest current density for LORETA. We aim at assessing whether a given combination of sources is supported by the data by obtaining a statistical threshold. In addition, we combine all significant solutions to create a map of the likelihood of having a dipolar source at a given point.

We have used an empirical method for the computation of thresholds. This permits to take into account the fact that noise is strongly correlated, both spatially and temporally, and to adapt the tests accordingly. The other option would have been to rely solely on prewhitening of the data, which has two drawbacks: the robust estimation of the covariance matrix with highly correlated data is difficult, and temporal prewhitening can reduce the signal [de Munck 00].

The models fitted in order to estimate the empirical distributions were based on the best solutions at level s and $(s-1)$, mainly for computational reasons. The thresholds we obtain are expected to be conservative, as we use the lowest SSE at level s, i.e. that giving the highest F value. Also, the information on the significance of the best combination at level s with respect to the best at level $(s-1)$ is valuable in the context of model-order selection. A refinement of the method could be to estimate the distributions by including different sources combinations in the bootstrap method.

When estimating the distributions under H_0 using bootstrap resampling, we chose to fit realizations of noise only (cf. (8.14)) and not to include signal. An option would have been to include in the bootstrap realizations the potentials generated by the best $(s-1)$ sources. We do not expect this to have a high impact though. Indeed, by fitting the best combinations of $(s-1)$ and s sources, most of the signal in the data generated by $(s-1)$ sources should be removed, and any remaining signal should be at the level of the noise.

We have corrected our statistics for the multiple comparison problem. We have used a simple Bonferroni correction, which is likely to be conservative because of the dependence between the tests on different combinations. An empirical approach that would control the false positive rate by performing a large amount of scans on background noise (as in [Pantazis 03]) would be very computationally demanding, especially for the three-source scan. An alternative could be to define “resolution elements”, i.e. sets of grid points for which the tests are highly correlated, as was done in fMRI [Worsley 96]. However, we have not found in our simulated data that the Bonferroni correction was particularly conservative. This is possibly because the conservatism of the Bonferroni threshold is compensated by the fact that sources are not ideal dipoles, which leads to elevated F tests.

We have used a multiresolution approach, similar to [Gavit 01]. This allowed us to obtain maps for three sources with a resolution of 10 mm, a reasonable performance in the context of spike localization for presurgical evaluation. Nevertheless, further restrictions on the scanned volume could allow for a finer resolution or a higher number of sources if this is deemed necessary.

Acknowledgment

C.-G. B. would like to thank Dr. Keith Worsley, Dr. Sylvain Baillet and Dr. Sunil Kukreja for fruitful discussions, Dr. Yahya Agha Khani for his help in reviewing the EEG data and

Dr. Andrew Bagshaw for proofreading the manuscript.