NEUROIMAGE **8,** 340–349 (1998)
ARTICLE NO. NI980372

fMRI Signal Restoration Using a Spatio-Temporal Markov Random Field Preserving Transitions

Xavier Descombes, Frithjof Kruggel, and D. Yves von Cramon

Max Planck Institute of Cognitive Neuroscience, 22-26 Inselstrasse, 04103, Leipzig, Germany

Received November 14, 1997

In fMRI studies, Gaussian filtering is usually applied to improve the detection of activated areas. Such lowpass filtering enhances the signal to noise ratio. However, undesirable secondary effects are a bias on the signal shape and a blurring in the spatial domain. Neighboring activated areas may be merged and the high resolution of the fMRI data compromised. In the temporal domain, activation and deactivation slopes are also blurred. We propose an alternative to Gaussian filtering by restoring the signal using a spatiotemporal Markov Random Field which preserves the shape of the transitions. We define some interaction between neighboring voxels which allows us to reduce the noise while preserving the signal characteristics. An energy function is defined as the sum of the interaction potentials and is minimized using a simulated annealing algorithm. The shape of the hemodynamic response is preserved leading to a better characterization of its properties. We demonstrate the use of this approach by applying it to simulated data and to data obtained from a typical fMRI study. © 1998 Academic Press

INTRODUCTION

Signal detection in the presence of random noise plays a crucial role in functional imaging as it represents the first step of the experimental analysis. With efficient signal detection techniques it is possible to diminish the number of task repetitions and to better describe low-level activations. Both points are of great importance when studying cognitive processes by fMRI. The analysis of fMRI data is usually achieved by performing a voxelwise (parametric or nonparametric) statistical test in the time-series to detect areas with signal changes related to the experimental design. The statistical measure (say, a *z*-score) is then compiled in a statistical parameter map (SPM). Significantly activated areas are found by thresholding the SPM (Friston *et al.*, 1994; Worsley and Friston, 1995; Xiong *et al.*, 1996). The threshold is established by testing the null-hypothesis (H_0-hypothesis). A P value which prevents false alarms is set to a given value (usually 0.05) and the threshold is derived by limiting the probability of false alarms to P. In such a binary decision process we can have four different outcomes: true positive ($P(1|H_1)$), false positive ($P(1|H_0)$), true negative ($P(0|H_0)$), false negative ($P(0|H_1)$). A spatial Gaussian filter is often used in fMRI studies to reduce the noise before computing the SPM and thus to prevent false alarms. The introduced smoothness is then taken into account when computing the P value (Friston *et al.*, 1995; Forman *et al.*, 1995). However, the signal itself is markedly affected by this lowpass filter, yielding blurring and possible displacement of activated areas. Moreover, signals of low amplitude are completely removed. Thus, the detection of activation is improved but the characteristics of the signal and the accuracy of the activated areas localization are lost. We propose to overcome this problem by restoring the data instead of filtering them. The restoration process smooths the noise but at the same time preserves the signal shape.

Restoration has been widely investigated in signal and image processing (Andrews and Hunt, 1977). Here, we focus only on the problem of noise reduction. Filtering techniques and in particular the application of adaptive filters have provided solutions to the restoration problem. However, if noise reduction is considered as an inverse problem, the regularization theory which takes into account some *a priori* knowledge on the solution has been demonstrated to be more powerful. The Bayesian approach plays a key role in the regularization theory for its ability to integrate both a data attachment term (goodness of fit with respect to data) and some general and flexible priors (probability distribution modeling the expected signal/image). A popular approach is to use Markov random fields (MRFs) (Besag, 1974; Geman and Geman, 1984) to define the prior. The *a priori* knowledge induced by MRFs is general enough not to overly constrain the solution but providing interesting regularizing properties. MRFs were introduced in the engineering sciences by Besag (1974). Since then, they have been widely used, especially in image processing, for various problems such as

1053-8119/98 $25.00

image segmentation (Derin and Elliott, 1987; Geiger and Yuille, 1990), image restoration (Chellappa *et al.*, 1988; Geman and Reynolds, 1992; Charbonnier *et al.*, 1992), texture analysis (Descombes, 1995; Gimel'farb, 1996; Descombes, 1997) or movement estimation (Konrad and Dubois, 1992; Heitz and Bouthemy, 1993). Here we consider MRFs in the context of signal and image restoration. The basic idea is to consider contextual information by defining interactions between neighboring voxels. The result of the restoration on a given voxel depends not only on the data originating from that voxel but also on the values of its neighbors. The optimization scheme provides a global solution which represents a major advantage in comparison to local filtering. The decisive advantage in the restoration process is that it preserves the transition shape (and more generally discontinuities) of the underlying signal to avoid blurring and distortions. Moreover, fine structures must be differentiated from noise to be preserved. Using MRFs, a first approach involves defining a line process on the dual lattice to inhibit interactions between voxels belonging to different objects (Geman and Geman, 1984; Chalmond, 1988; Zerubia and Geiger, 1991; Charbonnier *et al.*, 1992), however, at the cost of an increase in complexity of the model. The second approach, adopted in this paper, defines interactions to avoid smoothing of edges and transitions too drastically (Geman and Reynolds, 1992; Nikolova, 1996).

Our approach involves a 3D spatiotemporal model. The two first dimensions correspond to the spatial dimensions of an fMRI slice, whereas the third one relates to time. We first provide a quantitative measure of the restoration effect using synthetical signals (a square wave, a sine wave, and a prototypical hemodynamic response function). We then apply this restoration model to functional magnetic resonance images (fMRI) obtained by echoplanar imaging (EPI) and compare the results with (i) data modified by a spatial Gaussian filtering and (ii) data for which no preprocessing step was undertaken.

The paper is organized as follows. Under Materials and Methods we derive the basics of MRFs and the Bayesian approach for image restoration and detail the proposed model. Under Results and Discussion we present the results obtained from a functional study. We assess the quality of the approach by comparing results using a correlation coefficient between the temporal signal and a reference box-car. We also compare the signal shape before and after restoration. Finally, we draw conclusions.

MATERIALS AND METHODS

Background on Markov Random Fields

A MRF is a random field defined by local conditional probabilities. Consider a set of sites $S \times T = \{(s = (i, j), t)\}$ defined by the time samples for each voxel in the context of fMRI signals (s represents the spatial coordinates and t the temporal coordinate) and a state space Λ (possible values for the samples). A fMRI temporal signal is then an element of the configuration space $\Omega = \Lambda^{S \times T}$ denoted $Y = \{y_s(t), t \in T, s \in S\}$. Consider a random field P defined on Ω. P is said to be a MRF if and only if:

$$\forall (s, t) \in S \times T, P(y_s(t)) > 0, \quad \forall (s, t) \in S \times T, \forall y_s(t), y_{s'}(t') \in \Lambda, \tag{1}$$

$$P(y_s(t) | y_{s'}(t'), (s', t') \in S \times T - \{s, t\}) = P(y_s(t) | y_{s'}(t'), (s', t') \in N_{(s,t)}), \tag{2}$$

where $N_{(s,t)}$ is a finite subset of $S \times T$ called the neighborhood of (s, t) such that:

$$(s, t) \notin N_{(s,t)}, \tag{3}$$

$$(s', t') \in N_{(s,t)} \Rightarrow (s, t) \in N_{(s',t')}. \tag{4}$$

This property states that if the value of the voxels are known in a neighborhood of (s, t), denoted $N_{(s,t)}$ (containing spatial and temporal neighbors), then we know the probability law for $y_s(t)$ and this probability law does not depend on the pixel values outside $N_{(s,t)}$. The Hammersley–Clifford theorem establishes that a MRF can be written as a Gibbs Field:

$$\forall Y \in \Omega, P(Y) = \frac{1}{Z} \exp -U(Y) = \frac{1}{Z} \exp -\sum_{c \in C} V_c(y_s(t), (s, t) \in c), \tag{5}$$

where U is the energy function, C is the set of cliques (a clique is a finite subset of sites), and V_c is the potential associated with the clique c, Z being the normalization constant also referred to as the partition function. Therefore, a MRF is defined by a set of cliques and their associated potentials. Gibbs Fields, especially the Ising model, have been widely investigated in statistical physics to model ferromagnetism phenomena. The cliques define atoms which interact with each other, whereas the potentials represent the strength of the interactions between the spins of atoms. For instance, the Ising model consists of interactions which tend to give the same orientation to the spin of interacting atoms, creating a magnetic field. In image processing, atoms correspond to voxels and spins to the grey level value.

Generally, fMRI data are given as a 4-dimensional matrix, where 3 dimensions correspond to the space

and one to the time. Because most of our fMRI data are highly anisotropic with respect to the third spatial dimension, we drop one spatial dimension and consider in this paper a 3-dimensional MRF. There is no limitation in formulating this restoration framework in 4D.

The Bayesian Framework for Restoration

We consider that the data X are corrupted by an additional noise η:

$$X = Y + \eta, \tag{6}$$

where Y represents the underlying signal.

To restore the signal we have to maximize the *a posteriori* probability $P(Y|X)$. Therefore we have an inverse ill-posed problem. To invert the problem we apply the Bayes rule which provides the following result:

$$P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)} \propto P(X|Y)P(Y), \tag{7}$$

where $P(X|Y)$ refers to the likelihood (or data driven term or goodness of fit) and $P(Y)$ to the prior model. $P(X|Y)$ is defined by the noise model, whereas $P(Y)$ is defined by the proposed Markov random field.

To optimize our model we consider a Bayesian estimation framework. Consider a probability distribution $P(\omega)$ on a universe (configuration space) denoted Ω. We then define a cost function between two configurations $R(\omega, \omega')$. Bayesian estimation involves minimizing the Bayesian risk, namely finding the configuration which minimizes the expectation of the cost function:

$$\hat{\omega} = \arg\min_{\omega'} \int_{\omega} R(\omega, \omega')P(\omega)\, d\omega. \tag{8}$$

Solving a problem in the Bayesian framework is then equivalent to defining:

- A distribution $P(\omega)$ on the universe Ω
- A cost function $R(\omega, \omega')$
- A minimization algorithm

Here, we optimize the distribution given by $P(X|Y) \cdot P(Y)$ by reaching the maximum *a posteriori* (MAP) criterion. The MAP criterion is obtained by minimizing the cost function defined by:

$$R_{\mathrm{MAP}}(\omega, \omega') = 1 - \delta_{\omega' = \omega}, \tag{9}$$

where δ is the Kronecker symbol ($\delta_x = 1$ if x is true and 0 otherwise).

This criterion can be reached using the simulated annealing (SA) algorithm (see Geman and Geman, 1984, for a proof of convergence).

The Φ-Models

In this subsection, we introduce the Φ-model, proposed by Geman and McClure (1987) and studied by Geman and Reynolds (1992), to define the prior $P(Y)$.

The aim of the restoration prior is to add some constraints of homogeneity and smoothness on the solution. Lowpass filters smooth the data oblivious of the underlying signal and for this reason result in some loss of high frequency information. Using a MRF as a prior in a Bayesian framework improves the restoration of fine structures and transitions around edges. However, the choice of the potentials is crucial to avoid blurring. Our approach involves deriving potentials which preserve transitions while reducing noise in homogeneous areas. The main issues which guide the choice of a prior are edge recovery, stability, bias over large transition, and resolution (Nikolova, 1996). The Φ-models studied in (Geman and Reynolds, 1992) support edge recovery without introducing any bias over the large transitions. In this paper, we consider a Φ-model to preserve the transitions between activated and nonactivated areas in the spatial domain and between activated periods and baseline periods in the temporal domain.

The general formulation given in (Geman and Reynolds, 1992) for the Φ-function is the following:

$$\Phi(u) = \frac{-\beta}{1 + (|u|/\delta)^p}. \tag{10}$$

The Φ-function depends on three parameters: β, δ, and p. The β parameter is a scaling factor on the Y-axis. It represents the strength of the interaction. The two remaining parameters define the shape of the Φ-function. The Φ-function has a minimum equal to $-\beta$ at 0. At infinity, the X-axis is an asymptote. The value of the parameter β corresponds to the cost of an edge, whereas the behavior around the minimum defines the smoothing into homogeneous regions. The δ parameter allows us to define the values of the differences between neighboring voxels which have the highest probability of being transitions. The p parameter defines the shape of the transition between nonedge and edge values. As p increases, the transition is sharper (see Fig. 2).

After having tested different values for p, we consider only $p = 2$, which has provided the best results. Indeed, with $p = 2$ the Φ-function behaves like a quadratic

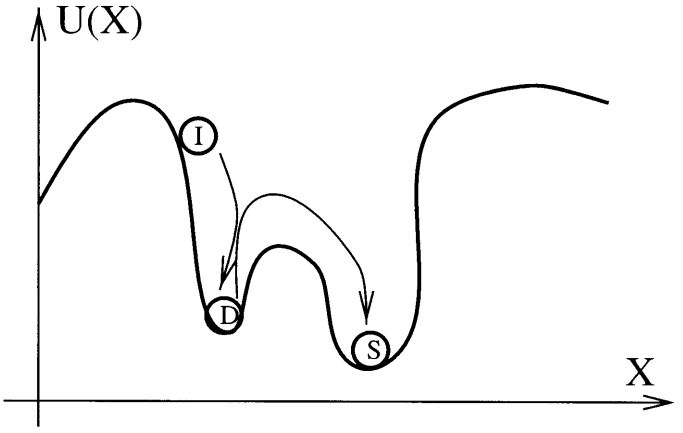

FIG. 1. Convergence of the simulated annealing. U(X), energy function; I, Initial configuration; D, Configuration obtained with a deterministic algorithm; S, Configuration obtained with a simulated annealing.

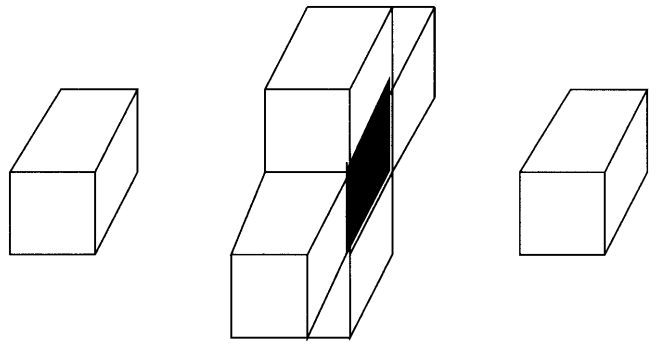

FIG. 3. Spatio-temporal neighborhood of the proposed model: four spatial neighbors and two temporal neighbors.

function near the origin and therefore the Φ-model can be approximated by a Gaussian model for low deviations while preserving the transitions because of the asymptotic shape of the Φ-function. We analyze the restoration for different values of the two remaining parameters.

A Spatio-Temporal MRF

To define the model we consider a 4-connectivity in the spatial domain and the nearest neighbors in the temporal domain. Therefore a given pixel has six neighbors (see Fig. 3). To preserve transitions between activated and nonactivated areas (in the spatial domain) and transitions between baseline and activation (in the temporal domain), we define Φ-models on pairwise interactions.

In the spatial domain we expect two kind of transitions: (i) transitions between activated and nonactivated areas and (ii) transitions between different anatomical structures (grey and white matter, . . .). To be insensitive to anatomical transitions we normalize the data set subtracting the output of a moving average filter in the temporal domain, thus assigning a zero mean to the time course of each voxel. Then we expect the same regularization in both the spatial and the temporal domain. Since we have twice as many neighbors in the spatial domain, we consider stronger interactions for the temporal potentials: $\beta_{temp} = 2\beta_{spat} = \beta$. The height of the transitions is the same in spatial and temporal domains as they represent the difference between activation and nonactivation. Therefore, we take the same value for the δ parameter in both domains: $\delta_{temp} = \delta_{spat} = \delta$. Finally, in some imaging protocols, voxels are not isotropic, i.e., the size on the I-axis (lines) and on the J-axis (columns) differ. We compensate for this anistropy by reducing the spatial interactions along the I-axis by the

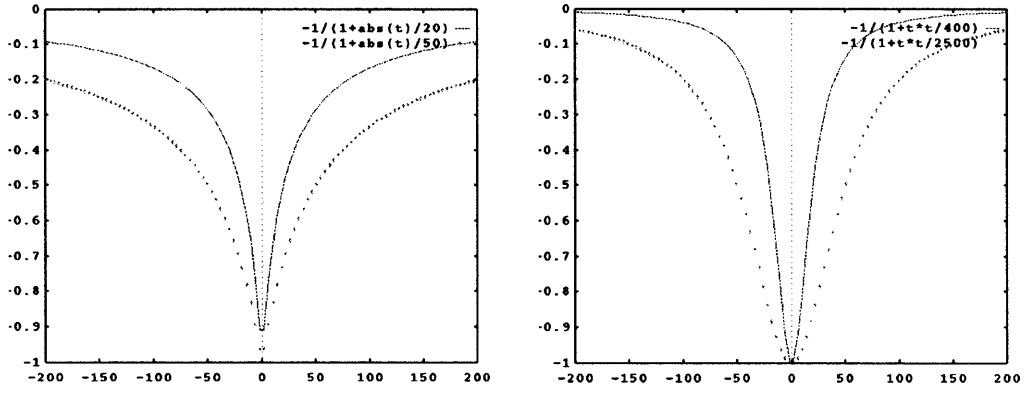

FIG. 2. Shape of the Φ-function for $p = 1$ (a) and $p = 2$ (b).

same ratio. Thus, the proposed model is written as follows:

$$P(Y) = \frac{1}{Z} \exp - \left[\sum_t \sum_{(i,j)} V_T(y_{i,j}(t), y_{i,j}(t+1)) + \sum_{(i,j)} \sum_t V_I(y_{i,j}(t), y_{i+1,j}(t)) + \sum_{(i,j)} \sum_t V_J(y_{i,j}(t), y_{i,j+1}(t)) \right], \quad (11)$$

where

$$V_T(y_{i,j}(t), y_{i+1,j}(t+1)) = \frac{-2\beta}{1+(y_{i,j}(t) - y_{i,j}(t+1))^2/\delta^2}, \quad (12)$$

$$V_I(y_{i,j}(t), y_{i+1,j}(t)) = \frac{-A\beta}{1 + (y_{i,j}(t) - y_{i+1,j}(t))^2/\delta^2}, \quad (13)$$

$$V_J(y_{i,j}(t), y_{i,j+1}(t)) = \frac{-\beta}{1 + (y_{i,j}(t) - y_{i,j+1}(t))^2/\delta^2}, \quad (14)$$

A denotes the voxel anisotropy ratio in the spatial domain.

Therefore, the prior model depends on two parameters: β and δ.

A Likelihood Taking Outliers into Account

To complete the model, we now define the likelihood function. This term represents the noise model. We first make an assumption concerning the conditional independency of voxels, which is written as follows:

$$P(X|Y) = \prod_{(i,j),t} p(x_{(i,j)}(t)|y_{(i,j)}(t)). \quad (15)$$

By this we assume that the noise is not correlated. This assumption made for simplicity purpose, is only an approximation in practice (Zarahn *et al.*, 1997). However, the main information provided by the data on a site s is given by the value of the voxel x_s itself. The interactions of the prior model the correlation between adjacent voxels.

In the absence of any precise information about the noise in fMRI data, a first approach is to consider a Gaussian model. However, the Gaussian model underestimates outliers corresponding to scanner artifacts, physiological artifacts, etc. . . To enhance the tail of the distribution and to normalize the likelihood function with respect to the prior model we also consider a Φ-model for the likelihood function:

$$p(x_{(i,j)}(t)|y_{(i,j)}(t)) = \frac{1}{Z} \exp - V_L(y_{(i,j)}(t)), \quad (16)$$

where

$$V_L(y_{(i,j)}(t)) = \frac{-\beta_L}{1 + (y_{(i,j)}(t) - x_{(i,j)}(t))^2/\delta^2}. \quad (17)$$

The Global Energy

The *a posteriori* probability is written as follows:

$$\begin{aligned} P(Y|X) &\propto P(X|Y)P(Y) \\ &\propto \prod_{(i,j),t} p(x_{(i,j)}(t)|y_{(i,j)}(t)) \exp - U_p(Y) \\ P(Y|X) &\propto \exp - \sum_{(i,j),t} -\log p(x_{(i,j)}(t)|y_{(i,j)}(t)) \\ &\quad \cdot \exp - U_p(Y) \propto \exp - U(Y) \end{aligned} \quad (18)$$

where U is called the energy and is defined as follows:

$$\begin{aligned} U(Y) = &\sum_{(i,j),t} V_L(y_{(i,j)}(t)) \\ &+ \sum_{c=|t,t+1|} \sum_{(i,j)} V_T(y_{i,j}(t), y_{i,j}(t+1)) \\ &+ \sum_{c=|(i,j),(i+1,j)|} \sum_t V_I(y_{i,j}(t), y_{i+1,j}(t)) \\ &+ \sum_{c=|(i,j),(i,j+1)|} \sum_t V_J(y_{i,j}(t), y_{i,j+1}(t)). \end{aligned} \quad (19)$$

The global minimum is defined up to a multiplicative constant. Therefore we can impose $\beta_L = 1$ without loss of generality. The proposed model depends on the two remaining parameters, β and δ.

Optimization

To find the configuration minimizing the global energy corresponding to the MAP criterion, we have implemented a simulated annealing algorithm. The energy function will inevitably possess some local minima. Therefore, deterministic algorithms are highly dependent on the initial configuration and in general do not lead to the global minimum. The idea of simulated annealing is to accept some transitions which increase

the energy, i.e., which temporally provide a configuration with a lower probability. This allows the process to jump away from local minima of the energy. A temperature parameter is introduced which controls the jumps. From a given configuration, configurations with a lower probability are accessible with a certain probability defined by the temperature parameter T (cf Fig. 1). The temperature decreases during iterations to reduce the amplitude of these jumps. Therefore, during iterations, the algorithm provides in average configurations which increase the probability. Geman and Geman (1984) have proven the convergence of the simulated annealing algorithm to the global minimum (i.e., the configuration which maximizes the *a posteriori* probability) in case of MRFs. Moreover, the result does not depend on the initialization. The algorithm is written as follows:

1. Compute a random initialization $Y^0 = (y_s^0(t))$, set $T = T_0$ and $n = 0$

2. For each voxel (s, t) of the time series:

2a. Select a new value randomly: *new*

2b. Compute the local energy $U(y_s(t)|N_{(s,t)}^n)$ for the current value $y_s(t)=y_s^n(t)$ and the new value $y_s(t) = new$.

2c. If *Unew* <= *Ucur* then assign *new* to $y_s^{n+1}(t)$

2d. Else assign *new* to $y_s^{n+1}(t)$ with probability $\exp[Unew - Ucur]/T$ and keep the current value $y_s^n(t)$ with probability $1 - \exp[Unew - Ucur]/T$

3. If not converged decrease T, set $n = n + 1$ and go to 2

Note that the Markov property allows us to consider only the local energy $U(y_s(t)|N_{(s,t)})$. This reduces the computation time and leads to a highly parallelizable algorithm.

RESULTS AND DISCUSSION

Validation on Synthetical Data

The main advantage of the proposed approach is to preserve the signal shape during the noise removal operation. We also have to check that the prior knowledge included in the MRF does not introduce artifacts in the signal. We therefore consider different synthetic waveforms and study their restoration.

We modulated three test functions (a full sine wave, a half square wave, and a prototypical hemodynamic response) onto a patch from a fMRI experiment, which showed no significant response in a standard evaluation. These artificial datasets were corrected for baseline fluctuations (Kruggel *et al.*, 1998), and either filtered spatially using a Gaussian filter (gauss), filtered by the proposed spatiotemporal Markov–Random Field (mrf), or left unfiltered (native). We computed the z-scores of the activation peak, as well as the recovery rates for the three test functions. We define the signal recovery c of the model function $m(t)$ from the preprocessed time series $y_s(t)$ by:

$$c = \frac{\sum_t^N (y_s(t) - (am(t) + b))^2}{\sum_t^N y_s(t)^2}, \qquad (20)$$

where we maximize c by variation of a and b. Thus, $c = 1$ denotes a perfect recovery, while $c = 0$ corresponds to a complete loss of the signal. The results are compiled in the Tables 1 and 2.

For Gaussian filtering we used a sigma of 0.8, for image restoration we selected $\beta = 0.6$ and $\delta = 60$.

TABLE 1

Waveform	Native	Gauss	mrf
Sine	0.456	0.581	0.872
Hemo	0.456	0.579	0.804
Square	0.458	0.556	0.631

Note. Recovery rates of test functions after preprocessing using baseline correction only (native), Gaussian filtering in the spatial domain, and image restoration.

Description of the Functional Experiment

Data for the evaluation of this algorithm were taken from an fMRI study performed in our institute.

Task. During an 18-s period a standard tone ($f = 600$ Hz), a deviant tone ($f = 660$ Hz) or a unique novel sound (a crashing car, a barking dog) were presented binaurally via air conduction through plastic tubing connected to insert ear plugs. Phases were presented in a pseudo-randomized order. Subjects were requested to count the deviant phases. A total of 215 timesteps were recorded.

Imaging. Four T2-weighted axial images (image matrix 128×64; slice thickness = 5 mm; gap = 2 mm; echo time TE = 40 ms; repetition time TR = 2.0 s; flip angle = 40°) were acquired using a gradient echo, echo planar imaging sequence (Bruker 3.0 Tesla MEDSPEC System).

TABLE 2

Waveform	Native	Gauss	mrf
Sine	8.0	10.6	14.3
Hemo	12.9	14.7	19.1
Square	12.8	14.0	18.2

Note. Z scores of test functions after preprocessing using baseline correction only (native), Gaussian filtering in the spatial domain, and image restoration.

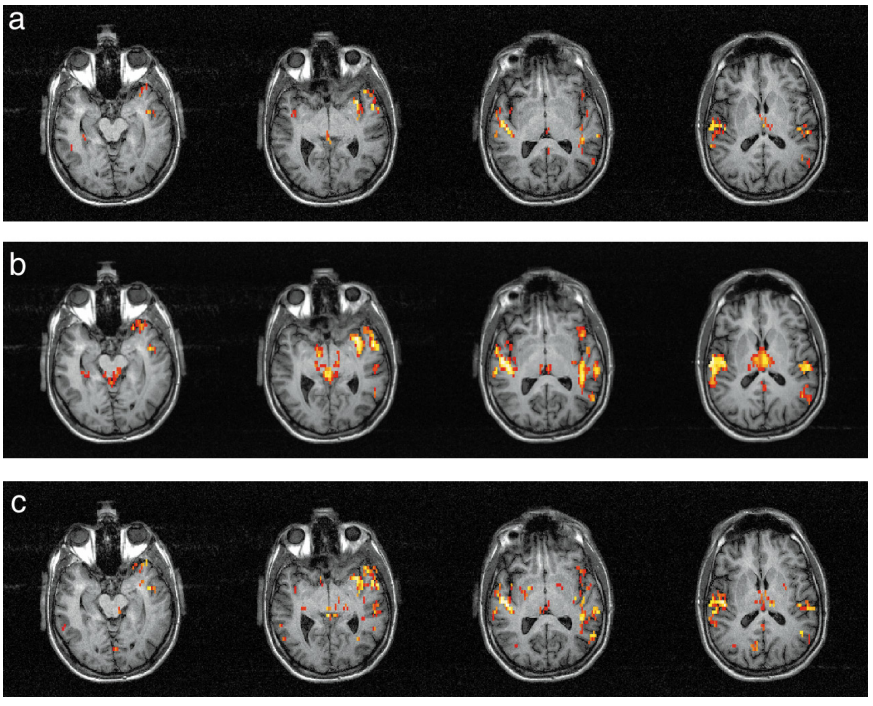

FIG. 4. Statistical evaluation of fMRI experiment: correlation coefficient with a threshold equal to 2.75 on original data (a), data filtered with a spatial Gaussian kernel $\sigma = 0.8$ (b) and data restored with the proposed model ($\delta = 4$, $\beta = 0.4$) (c).

Processing. Raw data were preprocessed for the exclusion of artifact-containing slices, corrected for subject motion, and filtered using this image restoration scheme. Voxels activated during the task were detected by the computation of a Pearson correlation coefficient with a standard box-car function shifted the characteristic lag between stimulus and fMRI response of 6 s. Only the standard vs the novel phases were evaluated. Correlation coefficients were converted into z-scores, thresholded ($z \geq 2.75$), color-coded, and overlaid onto T1-weighted anatomical images acquired at the same location.

To validate the restoration model we consider a correlation coefficient (see Eq. 20). Herein, we compare the results before, after restoration and with Gaussian filtering. To address the analysis of different activated areas, we also compare the signal shape by averaging the time courses of voxels within a given activated area.

Results

When comparing the results of different fMRI preprocessing techniques it is essential to recognize that there is no "ground truth" by which the correctness of the result or the improvement achieved by a filter can be measured.

However, experience with results from functional experiments, careful evaluation of the raw data, and neuroanatomical arguments enables us to distinguish between "real" cortical activations, activations in venous compartments, and false positives due to noise, artifacts, etc. Activations of interest are expected to (i) follow closely the cortical band, (ii) be found in brain locations which are connected with the task, and (iii) exhibit a certain shape (i.e., activation strength vs activation extent). However, especially with low-level activations, it is often difficult to make a distinction between true and false positives. A preprocessing scheme is aimed at enhancing signals that match certain characteristics while suppressing noise.

Figure 4 shows the thresholded z-map ($t = 2.75$) using raw data, after Gaussian spatial filtering ($\sigma^2 = 0.8$) and after restoration ($\delta = 4$, $\beta = 0.4$), respectively.

As expected, highly significant bilateral activations are found in the superior temporal gyrus, and, to a lesser extent, in the thalamus on both sides.

In comparison with the native data in the top row, Gaussian filtering (middle row) leads to a general enhancement of the signal but also to a considerable spatial blurring. In contrast, the signal restoration scheme proposed here (bottom row) surpasses the signal enhancement of a Gaussian filter while retaining anatomical details.

Figure 5 shows the thresholded z-map after restoration from top to bottom for different parameter values (respectively, $\delta = 2$ and $\beta = 0.3$, $\delta = 2$ and $\beta = 0.6$, $\delta = 6$ and $\beta = 0.3$, $\delta = 6$ and $\beta = 0.6$). $\beta = 0.3$ corresponds to a weak contribution of the prior model to the

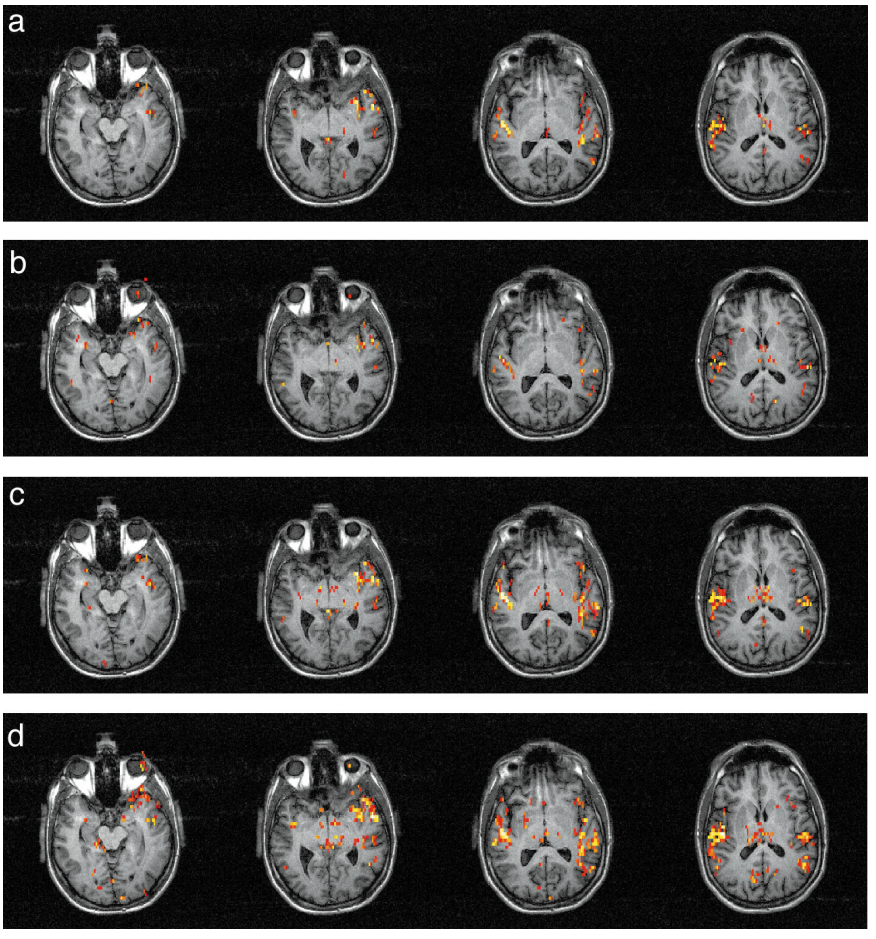

FIG. 5. Parameters influence in the restoration: (a: $\delta = 2$, $\beta = 0.3$; b: $\delta = 2$, $\beta = 0.6$; c: $\delta = 6$, $\beta = 0.3$; d: $\delta = 6$, $\beta = 0.6$.)

global energy. Therefore, the results are very similar to those obtained without restoration. Increasing the value of the β parameter allows us to study the influence of the δ parameter. A small value ($\delta = 2$) leads to a too large definition of transitions. In this instance, some noise will be considered as signal and with a high value of β noise appears in the result. Excessively high value for $\delta(\delta = 6)$ yields a too strict definition of edges. Transitions between activated and nonactivated voxels are then smoothed, which results in unduly blurring. Indeed, a very high value for δ leads to the same result as Gaussian filtering. A good compromise seems to be obtained with $\delta = 4$ and $\beta = 0.4$ (see Fig. 4).

In Fig. 6 we have plotted the time course of the fMRI signal during one period of the box-car waveform (the three curves correspond to raw data, data after spatial Gaussian filtering and data after restoration). These signals have been averaged over the activated area in the right auditory cortex (i.e., the area on the left side of the images in column 3 of Fig. 4). The restoration provides a better interpretation of the activation due to the regularization in the time domain. Noise has been removed but the slopes of activation and deactivation have not been blurred. Using a temporal Gaussian filter would have removed the noise but also deformed the activation and deactivation slopes.

The counterpart of the proposed model is an increase of the computation time in comparison with Gaussian filtering. Indeed, the defined energy function has local

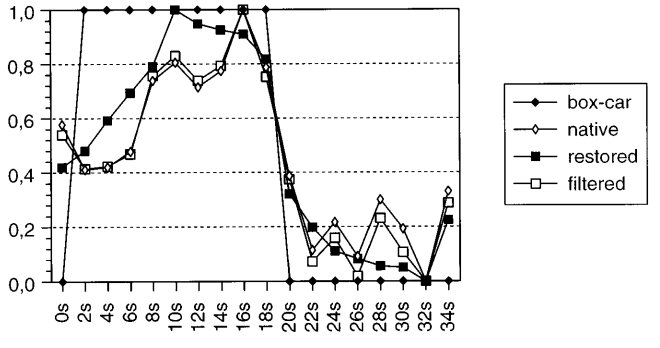

FIG. 6. Time course of the signal in the right auditory cortex: (shifted) box-car waveform as model, the raw data, data filtered with a spatial Gaussian kernel $\sigma = 0.8$, and data restored with the proposed model.

minima. Therefore, we have to use a stochastic algorithm to reach the global minimum. Some deterministic algorithms adapted to MRFs have been proposed but they rely on a good initialization of the solution. Such an initial solution is not easily available in our problem leading us to apply simulated annealing. On a typical workstation, our program requires 2 s per timestep for a 128 × 64 voxels slice. For the example dataset of 4 slices and 912 timesteps this corresponds to 2 h of computation time.

Spatial filtering as well as temporal filtering improves the signal to noise ratio as reported by Zarahn *et al.* (1997) and Aguirre *et al.* (1997). In the temporal domain, best results are obtained with a filter matching the hemodynamic function (Friston *et al.*, 1995). However, these filters are lowpass filters and spoil the high frequencies of the signals. With the proposed approach the prior model makes the distinction between noise and signal: noise is suppressed without spoiling the signal. This results in a better delineation of the activated areas and a better definition of the hemodynamic response. Moreover, an *a priori* model of the hemodynamic response is not required.

Conclusion

Spatial Gaussian filtering is generally used to enhance activation in fMRI studies. However this results in blurring of activated areas and a loss of low-level activations. In this paper, we propose to restore the signal by using a Markov random field embedded in a Bayesian framework. This leads to a noise reduction while restoring the shape of signals corresponding to activated voxels. The advantage of this approach has been demonstrated using fMRI data, improving the detection of activation with respect to the results obtained either on initial data or on filtered data. The restoration of the shape of the temporal signal provides a better characterization of the activation, especially for the slopes of activation and deactivation.

In the classical SPM approach, within the threshold of the statistical test a P value is computed to assess the significance of the detected activation. The effect of the Gaussian filtering can be taken into account to compute this probability. The extension of this mathematical derivations to the case of data restored with a MRF has not been achieved and seems very problematic. Indeed, to preserve transitions we have considered some Φ-functions which introduce nonlinearities. Some Monte Carlo simulations could provide a way to derive empirically a P value for analyzing the statistical test performed on restored data. The correlation coefficient performed in this paper allows us to detect signals matching the reference waveform. The statistical significance of this coefficient after the restoration process is under study.

ACKNOWLEDGMENTS

The authors thank Dr. David Norris and Dr. Christopher Wiggins for providing the data sets, Bertram Opitz for designing and conducting the fMR experiments, and Rosie Dymond for editorial comments.

REFERENCES

Aguirre, G., Zarahn, E., and D'Esposito, M. 1997. Empirical analysis of BOLD fMRI statistics. ii. Spatially smoothed data collected under null-hypothesis and experimental conditions. *Neuroimage* **5:**199–212.

Andrews, H., and Hunt, B. 1977. *Digital Image Restoration.* Prentice Hall, Englewood Cliffs, NJ.

Besag, J. 1974. Spatial interaction and statistical analysis of lattice systems. *Acad. R. Stat. Soc. Series B* **36:**721–741.

Chalmond, B. 1988. Image restoration using an estimated Markov model. *Signal Proc.* **15:**115–129.

Charbonnier, P., Blanc-Feraud, L., and Barlaud, M. 1992. Noisy image restoration using multiresolution Markov random fields. *Journal of Visual Communication and Image Representation,* **3**(4):338–346.

Chellappa, R., Simchony, T., and Jinchi, H. 1988. Relaxation algorithms for MAP restoration of gray level images with multiplicative noise. Technical report, Univ. S. Calif., Signal Image Processing Inst.

Derin, H., and Elliott, H. 1987. Modelling and segmentation of noisy and textured images using Gibbs random fields. *IEEE Trans. Pattern Anal. Machine Intelligence* **9**(1):39–55.

Descombes, X. 1995. A fission and fusion Markovian approach for multichannel segmentation. In *Proc. IGARSS'95,* Firenze, Italy, pp. 124–127.

Descombes, X. 1997. A dense class of Markov Random Fields and associated parameter estimation. *J. Vis. Commun. Image Rep.* **8**(3):299–316.

Forman, S., Cohen, J., Fitzgerald, M., Eddy, W., Mintun, M., and Noll, D. 1995. Improved assessment of significant activation in fMRI: Use of a cluster-size threshold. *Magn. Reson. Med.* **33:**636–647.

Friston, K., Frackowiak, K. W. R., Mazziotta, J., and Evans, A. 1994. Assessing the significance of focal activations using their spatial extent. *Human Brain Mapping* **1:**210–220.

Friston, K., Holmes, A., Poline, J., Grasby, P., Williams, S., Frackowiak, R., and Turner, R. 1995. Analysis of fMRI time-series revisited. *NeuroImage* **2:**45–53.

Geiger, D., and Yuille, A. 1990. A common framework for image segmentation. In *Proc. Int. Conf. Patt. Recog. ICPR-90* (A. City, Ed.), pp. 502–507.

Geman, S., and Geman, D. 1984. Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intelligence* **6**(6):721–741.

Geman, S., and McClure, D. 1987. Statistical methods for tomographic image reconstruction. In *Proc. of the 46th session of the ISI,* pp. 22–26.

Geman, S., and Reynolds, G. 1992. Constrained restoration and recovery of discontinuities. *IEEE Trans. Pattern Anal. Machine Intelligence* **14**(3):367–383.

Gimel'farb, G. 1996. Texture modeling by multiple pairwise pixel interactions. *IEEE Trans. Pattern Anal. Machine Intelligence* **18**(11):1110–1114.

Heitz, F., and Bouthemy, P. 1993. Multimodal estimation of discontinuous optical flow using Markov Random Fields. *IEEE Trans. Pattern Anal. Machine Intelligence* **15**(12):1217–1232.

Konrad, J., and Dubois, E. 1992. Bayesian estimation of motion

vector fields. *IEEE Trans. Pattern Anal. Machine Intelligence* **14**(9):910–927.

Kruggel, F., Descombes, X., and von Cramon, D. 1998. Preprocessing of fmr datasets. In *IEEE Workshop on Biomedical Image Analysis, Santa Barbara*.

Nikolova, M. 1996. Regularisation functions and estimators. In *Proc. IEEE Int. Conf. Ima. Proc., ICIP-96,* 457–460.

Worsley, K., and Friston, K. 1995. Analysis of fMRI time-series revisited again. *NeuroImage* **2:**173–181.

Xiong, J., Gao, J., Lancaster, J., and Fox, P. T. 1996. Assessment and optimization of functional MRI analyses. *Human Brain Mapping* **4:**153–167.

Zarahn, E., Aguirre, G., and D'Esposito, M. 1997. Empirical analysis of BOLD fMRI statistics. i. spatially unsmoothed data collected under null-hypothesis conditions. *NeuroImage* **5:**179–197.

Zerubia, J., and Geiger, D. 1991. Segmentation d'image et propagation de lignes. In Lyon-Villeurbanne, editor, *8eme congres de Reconnaissance des Formes et Intelligence Artificielle,* 661–668.