SPATIAL MIXTURE MODELLING FOR THE JOINT DETECTION-ESTIMATION OF BRAIN ACTIVITY IN fMRI

Thomas Vincent[1,2], *Philippe Ciuciu*[1,2] *and Jérôme Idier*[3]

[1] Service Hospitalier Frédéric Joliot (CEA) 4, Place du Général Leclerc, 91406 Orsay, France
[2] IFR 49, Institut d'Imagerie Neurofonctionnelle, Paris, France
[3] IRCCyN (CNRS), 1 rue de la Noë, BP 92101 44321 Nantes cedex 3, France
[1] firstname.lastname@cea.fr, [3] Jerome.Idier@irccyn.ec-nantes.fr

ABSTRACT

Within-subject analysis in event-related functional Magnetic Resonance Imaging (fMRI) first relies on (i) a detection step to localize which parts of the brain are activated by a given stimulus type, and then on (ii) an estimation step to recover the temporal dynamics of the brain response. Recently, we have proposed a Bayesian detection-estimation approach that jointly addresses (i)-(ii) [1]. This approach provides both a spatial activity map and an estimate of brain dynamics. Here, we consider an extension that accounts for spatial correlation using a spatial mixture model (SMM) based on a binary Markov random field. It allows us to avoid any spatial smoothing of the data prior to the statistical analysis. Our simulation results support that SMM gives a better control of false positive (specificity) and false negative (sensitivity) rates than independent mixtures.

Index Terms— Bayes procedures, Biomedical signal detection, Magnetic resonance imaging.

1. INTRODUCTION

Since the first report of the BOLD effect in human [2], functional MRI (fMRI) has represented a powerful tool to non-invasively study the relation between cognitive task and the hemodynamic (BOLD) response. Within-subject analysis in fMRI essentially addresses two problems. The first one is about the *detection* or localization of activated brain areas in response to a given stimulus type or experimental tasks, while the second one concerns the *estimation* of the temporal dynamic of activated voxels, also known as the Hemodynamic Response Function (HRF). In [1], a novel detection estimation approach has been proposed to address both issues in a region-based analysis, that is on a set of prespecified regions of interest (ROI). Within the Bayesian framework, we first integrated physiological prior information to obtain a slow-varying time course as an estimate of the HRF in every ROI. We have also considered different two-class *independent* mixture models (IMM) as prior distribution on the response magnitude to accomodate the voxel and stimulus-dependent signal fluctuations within the ROI [3].

In this paper, we introduce an extension based on *spatial* mixture models (SMM) that accounts for spatial correlation between neighboring voxels in the brain volume (regular lattice in 3D) or on the cortical surface (irregular lattice in 2D) [4–6]. The ensued goal is to favor the detection of activating clusters rather than isolated voxels. This extension consists in modelling a priori the state of a given voxel (non-activating, activating) depending on its neighbors using a symmetric Ising random field. In this regards, our work is close to [5] but more general in the sense that the HRF is estimated in the same time. The parameter controlling the strength of the spatial correlation is set by hand, as the smoothing level used when spatially filtering the data. The combination of these prior distributions with the likelihood allows us to derive the target posterior distribution using Bayes' rule. We then resort to Gibbs sampling to draw realizations from this posterior law. The posterior mean (PM) estimates of the HRF, the Neural Response Levels (NRLs) and the corresponding labels are directly computed from the generated samples in the Markov Chain Monte Carlo (MCMC) procedure. Compared to [1], a significant gain is achieved in terms of sensitivity and specificity on artificial fMRI data.

2. REGION-BASED MODELLING OF fMRI DATA

2.1. Motivations

Hypothesis-driven approaches postulate a model of the HRF response and enable local inference at the voxel level. Such methods take place in the General Linear Model (GLM) framework. They have been popularized by the Statistical Parametric Mapping software (SPM, http://www.fil.ion.ucl.ac.uk/spm). In this formulation, the model chosen for the BOLD response is a crucial issue. SPM uses the same temporal model for the whole brain for simplicity and computational reasons. To help cognitive interpretations, we rather advocate for the necessity of a spatially adaptive GLM in which local estimation of the HRF would be performed. The latter does not need to be done at the voxel level, but rather at a coarser regional scale. To define this scale, we use a segregation of the brain volume constrained to the grey matter mask into a few hundreds of connected ROIs, called *parcels*. Any parcellation procedure can be used, as long as functional homogeinity is guaranteed within each parcel. In this respect, the assumption of a shape-invariant HRF is maintained.

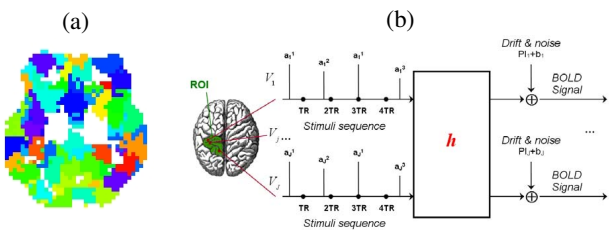

Fig. 1. (a): Slice of the color-coded parcellation at $z = -4$mm. (b): Parcel-based model of the BOLD signal.

Actually, we resort to [7] and Fig. 1(a) shows a slice resulting from such a parcellation.

2.2. Within-parcel formulation

Here, the parcel-based model of the BOLD signal introduced in [1] is adopted. As shown in Fig 1(b), it characterizes every parcel $\mathcal{P} = (V_j)_{j=1:J}$ by a single HRF shape and accounts for voxel-dependent and stimulus-related fluctuations of the magnitude of the BOLD signal. The BOLD fMRI time course measured in voxel V_j at times $(t_n)_{n=1:N}$ (where $t_n = n\,TR$, N being the number of scans and TR, the time of repetition) then reads

$$\boldsymbol{y}_j = \sum_{m=1}^{M} a_j^m \boldsymbol{X}^m \boldsymbol{h} + \boldsymbol{P}\boldsymbol{\ell}_j + \boldsymbol{b}_j, \quad \forall\, j,\, V_j \in \mathcal{P}, \qquad (1)$$

This model remains time-invariant while it incorporates voxel dependent and stimulus related amplitudes *Neural Response Levels* (NRLs) $\mathbb{a} = (\boldsymbol{a}_j^m)$. $\boldsymbol{X}^m = (x^m_{t_n - d\Delta t})_{n=1:N, d=0:D}$ is a $N \times (D+1)$ binary matrix that codes the arrival times of the mth stimulus. Δt is the sampling period of the HRF usually lower than TR. The stimulus onsets are put on the Δt-sampled grid by moving them to the nearest time point on this grid. Vector $\boldsymbol{h} = (h_{d\Delta t})_{d=0:D}$ represents the unknown HRF shape in parcel $\mathcal{P}$. A single HRF shape is sufficient provided that $\mathcal{P}$ is functionally homogeneous. Note also that $\boldsymbol{P}\ell_j$ models a low-frequency trend to account for physiological artifacts and that $\boldsymbol{b}_j \sim \mathcal{N}(0, \sigma^2_{\boldsymbol{b}_j})$ stands for the noise. For simplicity reasons, we have just considered a Gaussian white noise model, while more sophisticated modelling can be introduced using AR processes to account for serial correlation in the fMRI time series [8].

3. THE DETECTION-ESTIMATION PROBLEM

3.1. Non-spatial *vs* spatial mixture modelling

We propose to *estimate* the HRF shape $\boldsymbol{h}$ and the corresponding NRLs $\mathbb{a}$ in $\mathcal{P}$. Our aim is also to *classify* which voxels in $\mathcal{P}$ are involved in the experimental paradigm. Hence, in [1] we introduced *non-spatial* two-class mixture models for every condition m, in which class 0 describes non-activating voxels and class 1 models activating voxels in response to the mth stimulus type. In [3], we compared different types of mixture priors (two-class Gaussian (GGM) and Gamma-Gaussian (GGaM)). Importantly, we have shown that inhomogeneous GGaM better control the false positive rate in comparison with GGM, specifically in brain regions where most voxels are non-activating. Hence, they provide more reliable statistical maps. However, inference based upon GGa mixtures is computationally expensive. As a consequence, following [4–6, 9] we resort here to Gaussian *spatial* mixture modelling to control the specificity.

Our approach stands in the Bayesian framework, and so integrates prior knowledge on the sought objects, *i.e.* the HRF and the NRLs.

3.2. Priors

The HRF. According to [10], the HRF can be characterized as a causal slow-varying function which returns to its baseline after about 25 s. These assumptions lead to select a Gaussian prior on $\boldsymbol{h} \sim \mathcal{N}(\boldsymbol{0}, \sigma^2_{\boldsymbol{h}} \boldsymbol{R})$, where $\boldsymbol{R} = (\boldsymbol{D}_2^{\mathrm{t}} \boldsymbol{D}_2)^{-1}$ is a symmetric positive definite matrix and $\boldsymbol{D}_2$ is the *truncated* second-order finite difference matrix of size $(D-1)\times(D-1)$ such that $\|\partial^2 \boldsymbol{h}\|^2 = \boldsymbol{h}^{\mathrm{t}} \boldsymbol{R}^{-1} \boldsymbol{h}$.

The NRLs. We assume that different types of stimulus induce statistically independent NRLs *i.e.,* $p(\mathbb{a} \mid \boldsymbol{\theta}_{\boldsymbol{a}}) = \prod_m p(\boldsymbol{a}^m \mid \boldsymbol{\theta}^m)$ with $\mathbb{a} = (\boldsymbol{a}^m)_{m=1:M}$, $\boldsymbol{a}^m = (a_j^m)_{j=1:J}$ and $\boldsymbol{\theta}_{\boldsymbol{a}} = (\boldsymbol{\theta}^m)_{m=1:M}$. Vector $\boldsymbol{\theta}^m$ denotes the set of unknown hyperparameters related to the mth stimulus type. We define a spatial Bayesian model by introducing binary indicator variables q_j^m that states whether voxel V_j is activated ($q_j^m = 1$) or not ($q_j^m = 0$) in response to stimulus m, so that the NRL a_j^m is normally distributed according to $a_j^m \mid q_j^m = i \sim \mathcal{N}(\mu_{i,m}, v_{i,m})$, with $i = 0, 1$. We impose $\mu_{0,m} = 0$ for the mean of the NRLs in non-activating voxels, leading to $\boldsymbol{\theta}^m = [v_{0,m}, \mu_{1,m}, v_{1,m}]$. Note that a Bernoulli-Gaussian formulation has also been tested in fMRI in [5]. This modelling corresponds to a degenerate mixture in which $a_j^m = 0$ if $q_j^m = 0$.

In an IMM, the mixing probability $\Pr(q_j^m = 1) = \lambda_m$ is independent of j. Here, we introduce space-varying probabilities $\lambda_{j,m}$ through a spatially correlated Ising prior on the binary variables $\boldsymbol{q}^m$, while the NRLs remain independent conditionally to $\boldsymbol{q}^m$. An Ising model is a binary MRF commonly used in image analysis [9, 11]. It is of particular interest in fMRI analysis since it allows spatial correlation to be directly incorporated on the probabilities of activation. In what follows, we consider a *symmetric* Ising MRF defined by

$$\Pr(\boldsymbol{q}^m \mid \beta^m) = Z(\beta^m)^{-1} \exp\Big(-\beta^m \sum_{j \sim k} \omega_{jk} I(q_j^m = q_k^m)\Big), \quad (2)$$

where $I(A) = 1$ if A is true and $I(A) = 0$ otherwise. The notation $j \sim k$ means that the sum extends over all neighboring voxels, while ω_{jk} are prespecified constants that weight the interaction between voxels (V_j, V_k) according to the neighborhood system. It can be defined either in 3D in the brain volume intersecting parcel $\mathcal{P}$ or in 2D along the cortical surface. In this paper, we only consider the 3D case using 18 or 26 nearest neighbors. This means that $w_{jk} = 1$ for horizontal and vertical neighbors in a given slice and $w_{jk} = 1/\sqrt{2}$ for directly adjacent diagonal voxels. Note also that this MRF is *hidden* since $\mathbb{q} = (\boldsymbol{q}^m)_{m=1:M}$ are not observed in (1).

The parameter $\beta^m > 0$ in (2) controls the amount of spatial smoothing, with the elements of $\boldsymbol{q}^m$ being independent if $\beta^m = 0$. Large values of β^m associates higher probabilities to configurations containing clusters of like-valued neighboring binary variables. Function Z is a normalizing constant also called the *partition function*. Here, we do not estimate it since β^m is set by hand. For future adaptive spatial smoothing, Z could be evaluated on a discrete grid to make the sampling of β^m feasible (see [9, 11]). Combining all information, we get a spatial mixture model for every stimulus type:

$$p(\boldsymbol{a}^m \mid \boldsymbol{\theta}^m) = \sum_{\boldsymbol{q}^m} \Big[\prod_{j=1}^{J} p(a_j^m \mid q_j^m, \boldsymbol{\theta}^m)\Big] \Pr(\boldsymbol{q}^m \mid \beta^m) \qquad (3)$$

Other parameters. To complete the Bayesian model, priors are required for all the remaining parameters. For the drift parameters, we assume that $(\boldsymbol{\ell}_j)_j$ are independent of $\boldsymbol{h}$ and that $p((\boldsymbol{\ell}_j)_j \mid \sigma^2_\ell) = \prod_j p(\boldsymbol{\ell}_j \mid \sigma^2_\ell)$ with $\boldsymbol{\ell}_j \sim \mathcal{N}(\boldsymbol{0}, \sigma^2_\ell \boldsymbol{I}_Q)$ where Q defines the size of the orthogonal basis $\boldsymbol{P}$. Without informative prior knowledge, we define $p(\sigma^2_{\boldsymbol{h}}, \sigma^2_\ell) = (\sigma_{\boldsymbol{h}} \sigma_\ell)^{-1}$ and $p(\boldsymbol{\varepsilon}) = \prod_j \varepsilon_j^{-1}$. For variance $v_{0,m}$, we choose an improper Jeffreys' prior $p(v_{0,m}) = v_{0,m}^{-1/2}$ because we *do* expect non-activating voxels in any parcel. Hence, class 0 should never be empty a priori. However, for variance $v_{1,m}$, we resort to a conjugate prior (inverse Gamma pdf) denoted as $\mathcal{IG}(a_{v_1}, b_{v_1})$ since the class of activating voxels may be empty. We thus avoid degeneracy problem that could prevent its sampling. In the same way, we introduce a proper prior $\mathcal{N}(a_{\mu_1}, b_{\mu_1})$ on $\mu_{1,m}$.

3.3. The joint posterior distribution

Considering the constructed model and assuming no further prior dependence between parameters, Bayes' rule gives us:

$$p(\boldsymbol{h}, \mathbb{a}, (\boldsymbol{\ell}_j), \boldsymbol{\Theta} \,|\, \boldsymbol{y}) \propto p(\boldsymbol{y} \,|\, \boldsymbol{h}, \boldsymbol{a}, (\boldsymbol{\ell}_j), \boldsymbol{\epsilon}^2)\, p(\mathbb{a} \,|\, \boldsymbol{\theta_a})\, p(\boldsymbol{h} \,|\, \sigma_{\boldsymbol{h}}^2) \times \\ \times\, p((\boldsymbol{\ell}_j) \,|\, \sigma_\ell^2)\, p(\boldsymbol{\epsilon})\, p(\sigma_{\boldsymbol{h}}^2)\, p(\sigma_\ell^2) \prod_m p(\boldsymbol{\theta}^m).$$

Here, the nuisance variables $(\boldsymbol{\ell})_j$ are integrated out. This yields:

$$p(\boldsymbol{h}, \mathbb{a}, \boldsymbol{\Theta} \,|\, \boldsymbol{y}) \propto \left(\prod_{j=1}^{J} \epsilon_j^{-N-1+Q}\right) \sigma_{\boldsymbol{h}}^{-D} \exp\left(-\frac{\boldsymbol{h}^{\mathrm{t}} \boldsymbol{R}^{-1} \boldsymbol{h}}{2\sigma_{\boldsymbol{h}}^2}\right) \\ \exp\left(-\frac{\sum_{j=1}^{J} \widetilde{\boldsymbol{y}}_j^{\mathrm{t}} \boldsymbol{Q}_j \widetilde{\boldsymbol{y}}_j}{2}\right) \prod_{m=1}^{M} \Big(p(\boldsymbol{a}^m \,|\, \boldsymbol{\theta}^m) p(\boldsymbol{\theta}^m)\Big). \qquad (4)$$

with $\boldsymbol{Q}_j = \left(\boldsymbol{I}_N - \boldsymbol{P}\boldsymbol{P}^{\mathrm{t}}\right)/\epsilon_j^2$, $\widetilde{\boldsymbol{y}}_j = \boldsymbol{y}_j - \boldsymbol{S}_j \boldsymbol{h}$ and $\boldsymbol{S}_j = \sum_j a_j^m \boldsymbol{X}^m$. To get samples of the posterior pdf, we use a Gibbs sampler which consists in building a Markov chain, whose target distribution is (4), by sequentially generating random samples from the full conditional pdfs of all the unknown parameters and hyperparameters. Finally, posterior mean (PM) estimates are computed from the samples according to the following rule: $\widehat{x}^{\mathrm{PM}} = (K-I)^{-1} \sum_{k=I+1}^{K} x^{(k)}$, $\forall\, x \in \{\boldsymbol{h}, \mathbb{a}, \mathbb{q}, \boldsymbol{\Theta}\}$ where I stands for the length of the burn-in period. The sampling scheme for the posterior mixtures $(\mathbb{a}, \mathbb{q})$ is detailed in the next paragraph while for other quantities of interest $(\boldsymbol{h}, \boldsymbol{\Theta})$, the reader may refer to [1, Appendix A] since their sampling remains unchanged.

3.4. Computational details

Since the prior on the NRLs ($\mathbb{a}$) is a Gaussian mixture and the likelihood is Gaussian, then the full posterior pdf of $\mathbb{a}$ is a Gaussian mixture. From (4), it can be shown that each $a_j^m \in \mathbb{a}$ is obtained by sampling a 2-class posterior spatial Gaussian mixture in voxel V_j for the mth stimulus type. Letting $\mathrm{N}_j = \{V_k \,|\, k \sim j\}$, the latter reads:

$$p(a_j^m \,|\, \boldsymbol{y}_j, \boldsymbol{h}, \boldsymbol{\theta}^m, \epsilon_j^2, a_j^{m' \neq m}, q^m_{k \in \mathrm{N}_j}, \beta^m) = \sum_{i=0,1} \lambda_{i,j}^m \mathcal{N}\big(\mu_{i,j}^m, v_{i,j}^m\big)$$

which can be decomposed in three steps: (i) Identify the posterior parameters $(\lambda_{i,j}^m, \mu_{i,j}^m, v_{i,j}^m)$; (ii) Sample the binary label q_j^m according to $\lambda_{i,j}^m$ and (iii) Sample the NRL $a_j^m \,|\, q_j^m$ according to

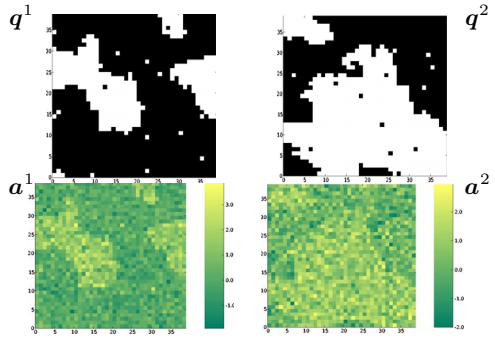

Fig. 2. True labels $\boldsymbol{q}^m$ (activating in white and non-activating in black) and true NRLs $\boldsymbol{a}^m$ for the two stimulus types involved in the simulated data.

$\mathcal{N}(\mu_{i,j}^m, v_{i,j}^m)$. As detailed in [1, Appendix A], we have for $i = 0, 1$:

$$v_{i,j}^m = \left(v_{i,m}^{-1} + \boldsymbol{g}_m^{\mathrm{t}} \boldsymbol{Q}_j \boldsymbol{g}_m\right)^{-1},\ \mu_{i,j}^m = v_{i,j}^m \left(\boldsymbol{g}_m^{\mathrm{t}} \boldsymbol{Q}_j \boldsymbol{e}_{m,j} + i\, \frac{\mu_{i,m}}{v_{i,m}}\right)$$

where $\boldsymbol{g}_m = \boldsymbol{X}^m \boldsymbol{h}$ and $\boldsymbol{e}_{m,j} = \boldsymbol{y}_j - \sum_{m' \neq m} a_j^{m'} \boldsymbol{g}_{m'} = \widetilde{\boldsymbol{y}}_j + \boldsymbol{g}_m$. The posterior probability $\lambda_{i,j}^m$ of the event $q_j^m = i$ reads:

$$\lambda_{i,j}^m = \left(1 + \frac{r_{1-i,j}^m}{r_{i,j}^m} \underbrace{\overbrace{\frac{\Pr(q_j^m = 1-i \,|\, q^m_{k \in \mathrm{N}_j}, \beta^m)}{\Pr(q_j^m = i \,|\, q^m_{k \in \mathrm{N}_j}, \beta^m)}}^{= \pi_{1-i,j}^m}}_{= \pi_{i,j}^m}\right)^{-1} \qquad (5)$$

with $r_{i,j}^m = \left(v_{i,j}^m / v_{i,m}\right)^{1/2} \exp\left((\mu_{i,j}^m)^2 / v_{i,j}^m - i(\mu_i^m)^2 / v_i^m\right)$. To calculate (5), we need to evaluate

$$\kappa_j^m = \pi_{1-i,j}^m / \pi_{i,j}^m = \exp\Big(\beta^m \sum_{k \in \mathrm{N}_j} w_{jk}(1 - 2q_k^m)\Big).$$

4. SIMULATION RESULTS

4.1. Artificial fMRI datasets

We simulated a random mixed sequence of indexes coding for $M = 2$ different stimuli (see Fig. 3(a)). These two set of trials (30 trials per stimulus) were then multiplied by stimulus-dependent and space-varying NRLs, which were generated according to (3) (*cf.* Fig. 2). To this end, we generated 2D slices composed of 40 x 40 binary labels $\boldsymbol{q}^m$ (activating and non-activating voxels) simulated from an Ising MRF for each stimulus type m. As shown in Fig. 2, we used a relatively large amount of spatial correlation ($\beta^m = \beta = 0.6$). Then, we simulated normally-distributed NRLs:

$$(a_j^1 \,|\, q_j^1 = 0) \sim \mathcal{N}(0, .3), \qquad (a_j^1 \,|\, q_j^1 = 1) \sim \mathcal{N}(1.5, .5),$$
$$(a_j^2 \,|\, q_j^2 = 0) \sim \mathcal{N}(0, .6), \qquad (a_j^2 \,|\, q_j^2 = 1) \sim \mathcal{N}(1, .5).$$

Since $\mu_{1,2} < \mu_{1,1}$, a lower signal-to-noise ratio (SNR) is obtained for stimulus type 2. While Eq. (3) tells that NRLs $\boldsymbol{a}^m$ are independent conditionally to $\boldsymbol{q}^m$, Fig. 2(c)-(d) illustrate the impact of the spatial correlation of activation probabilities $\boldsymbol{q}^m$ on the NRL maps. Following Fig. 1(b), the signal $\boldsymbol{S}_j \boldsymbol{h}$ was obtained after convolving the NRL-modulated stimuli sequence with a HRF $\boldsymbol{h}$, whose exact shape appears in Fig. 3 in blue. White Gaussian noise $\boldsymbol{b}_j$ and low-frequency drift $\boldsymbol{P}\boldsymbol{\ell}_j$[1] were then superimposed to $\boldsymbol{S}_j \boldsymbol{h}$ in voxel V_j.

4.2. Results with non-spatial *vs* spatial mixtures

We tested our method on this artificial fMRI dataset and compared it to previous work involving non-spatial mixtures [1]. First, note that both approaches provide very close HRF estimates as seen in Fig. 3(e). The marginal posterior activation probability maps (PPM) which are computed as $(\widehat{\boldsymbol{q}}^m)^{\mathrm{PM}}$ for $m = 1, 2$, are reported in Fig 4. It is shown that the proposed SMM provides as expected more reliable results than IMM in terms of false positive and negative rates, particularly in the low SNR situation arising for $m = 2$ (compare Fig. 4(a)-(b) with Fig. 4(c)-(d), respectively). The NRLs estimates are recovered more accurately using a SMM (result not shown).

In order to better assess the differences, we applied a variable threshold upon the marginal PPM and built classical Receiving Op-

[1] $\boldsymbol{P}$ was defined from a cosine transform basis and parameters $\boldsymbol{\ell}_j$ were drawn from a normal distribution.

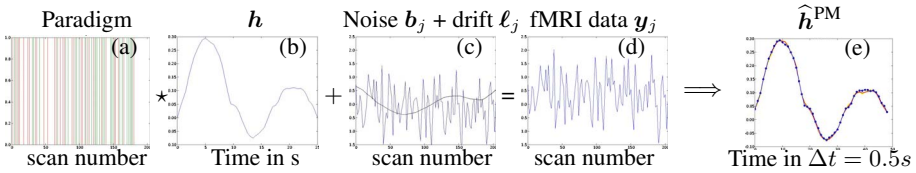

Fig. 3. (a)-(d) Simulation of artificial fMRI datasets. (e): True HRF (blue) and its SMM (red) and nSMM (yellow) estimates.

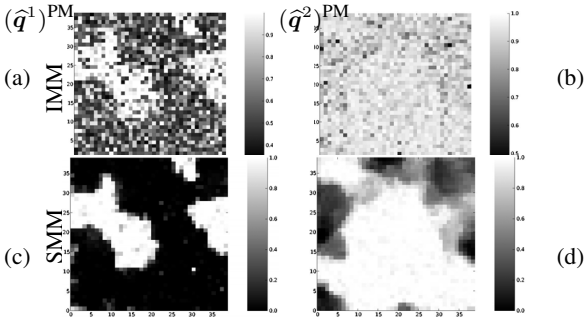

Fig. 4. Marginal PPM for the non-spatial (top) and spatial (bottom) mixture models. $m = 1$: (a)-(c); $m = 2$: (b)-(d).

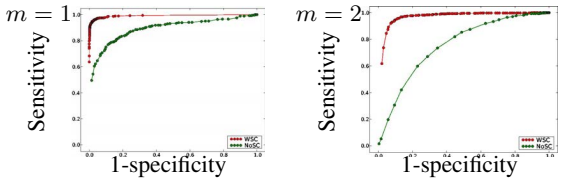

Fig. 5. ROC curves for non-spatial *vs* spatial mixture models.

erator Characteristics (ROC) curves. Fig. 5 illustrates that our SMM extension provides a more sensitive detection when specificity is fixed (vertical threshold) and conversely, that we obtain an improved specificity (or a more powerful statistical test) at a given sensitivity (horizontal threshold). The level of improvement depends on the SNR but is significant for both experimental conditions (red curve above green curve). Hence spatial correlation achieves a much more powerful discrimination at low SNR and succeeds well in recovering structural information from strongly altered signals.

5. CONCLUSION

The proposed contribution has shown promising results and has achieved the goal of handling very low SNR situations with an equally reliable dynamics estimation and an improved activation detection compared to the original approach [1]. We have statistically settled this result by highlighting more sensitive and specific tests on the activation detection. The proposed model needs to be further tested on simulated data which diverge from the priors used in the estimation. Moreover, assession on real fMRI data must of course be conducted.

Our current method could be enhanced with adaptive spatial correlation meaning an unsupervised setting of the correlation factors. One immediate approach would be to "hand tune" these factors by experimenting ranges of sensible values adapted to the BOLD data. A more skillfull approach would consist in sampling the parameters β^m in a fully Bayesian approach. This requires the precise computation of the partition function $Z(\beta^m)$ prior to running the Gibbs sampler. Fortunately, this topic has been strongly investigated [9,11].

6. REFERENCES

[1] S. Makni, P. Ciuciu, J. Idier, and J.-B. Poline, "Joint detection-estimation of brain activity in functional MRI: a multichannel deconvolution solution," *IEEE Trans. Signal Processing*, vol. 53, no. 9, pp. 3488–3502, Sep. 2005.

[2] S. Ogawa, T. Lee, A. Kay, and D. Tank, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation," *Proc. Natl. Acad. Sci. USA*, vol. 87, no. 24, pp. 9868–9872, 1990.

[3] S. Makni, P. Ciuciu, J. Idier, and J.-B. Poline, "Bayesian joint detection-estimation of brain activity using MCMC with a Gamma-Gaussian mixture prior model," in *Proc. 31th Proc. IEEE ICASSP*, May 2006, vol. V, pp. 1093–1096.

[4] N. Vaever Hartvig and J. Jensen, "Spatial mixture modeling of fMRI data," *Hum. Brain Mapp.*, vol. 11, no. 4, pp. 233–248, 2000.

[5] M. Smith, B. Pütz, D. Auer, and L. Fahrmeir, "Assessing brain activity through spatial Bayesian variable selection," *Neuroimage*, vol. 20, pp. 802–815, 2003.

[6] M. Woolrich, T. Behrens, Ch. Beckmann, and S. Smith, "Mixture models with adaptive spatial regularization for segmentation with an application to fMRI data," *IEEE Trans. Med. Imag.*, vol. 24, no. 1, pp. 1–11, Jan. 2005.

[7] B. Thirion, G. Flandin, P. Pinel, A. Roche, P. Ciuciu, and J.-B. Poline, "Dealing with the shortcomings of spatial normalization: Multi-subject parcellation of fMRI datasets," *Hum. Brain Mapp.*, vol. 27, no. 8, pp. 678–693, Aug. 2006.

[8] M. Woolrich, B. Ripley, M. Brady, and S. Smith, "Temporal autocorrelation in univariate linear modelling of fMRI data," *Neuroimage*, vol. 14, no. 6, pp. 1370–1386, Dec. 2001.

[9] D. Smith and M. Smith, "Estimation of binary Markov random fields using Markov Chain Monte Carlo," *J. Comput. and Graph. Stats.*, vol. 15, no. 1, pp. 207–227, 2006.

[10] G. Marrelec, H. Benali, P. Ciuciu, M. Pélégrini-Issac, and J.-B. Poline, "Robust Bayesian estimation of the hemodynamic response function in event-related BOLD MRI using basic physiological information," *Hum. Brain Mapp.*, vol. 19, no. 1, pp. 1–17, May 2003.

[11] D. M. Higdon, "Auxiliary variable methods for Markov chain Monte Carlo with applications," *J. Amer. Statist. Assoc.*, vol. 93, no. 442, pp. 585–595, June 1998.