Application and validation of spatial mixture modelling for the joint detection-estimation of brain activity in fMRI.

Thomas Vincent, Philippe Ciuciu, Jérôme Idier

Application and validation of spatial mixture modelling for the joint detection-estimation of brain activity in fMRI

Thomas Vincent, Philippe Ciuciu and Jérôme Idier

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 second on (ii) an estimation step to recover the temporal dynamics of the brain response. Recently, a Bayesian detection-estimation approach that jointly addresses (i)-(ii) has been proposed in [1]. This work is based on an independent mixture model (IMM) and provides both a spatial activity map and an estimate of brain dynamics. In [2], we accounted for spatial correlation using a spatial mixture model (SMM) based on a binary Markov random field. Here, we assess the SMM robustness and flexibility on simulations which diverge from the priors and the generative BOLD model and further extend comparison between SMM and IMM on real fMRI data, focusing on a region of interest in the auditory cortex.

I. INTRODUCTION

Since the first report of the BOLD effect in human [3], functional MRI (fMRI) has stood for a powerful tool to non-invasively study the relation between cognitive task and cerebral activity through the analysis of this hemodynamic BOLD signal. 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 task, 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, physiological prior information has been integrated to obtain a slow-varying time course as an estimate of the HRF in every ROI. Different two-class *independent* mixture models (IMM) as prior distribution on the response magnitude were also tested to accomodate the voxel and stimulus-dependent signal fluctuations within the ROI [4]. Due to the computational complexity of inhomogeneous mixtures, we introduced a two-class Gaussian *spatial* mixture model (SMM) as a relevant alternative in [2]. It accounts for spatial correlation between neighboring voxels and aims at favoring the detection of activating clusters rather than isolated voxels. We used a symmetric Ising random field to model a priori the state of a given voxel

(activating or non-activating). 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 a Markov Chain Monte Carlo (MCMC) procedure. In [2], this SMM approach was shown to give a significant gain in terms of sensitivity and specificity on artificial fMRI data, compared to the IMM counterpart. Since our primary interest was the comparison of SMM with IMM, our simulations followed the generative model and did not truly reflect a "realistic" BOLD data structure. Here, we assess the robustness against putative misspecifications from the assumed model and present first results on real fMRI data obtained during an event-related paradigm.

II. REGION-BASED MODELLING OF FMRI DATA

A. 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 and is actually refined to allow variation of the HRF at the voxel level through additional regressors: HRF derivatives [5] or half-cosine parametrisation [6]. Although such approach yields better fit, increasing the number of regressors implies fewer effective degrees of freedom in any subsequent statistical
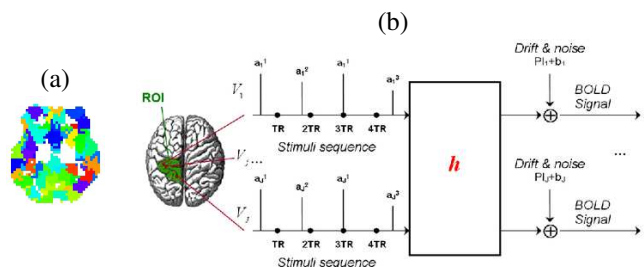

Fig. 1.   (a): Slice of the color-coded parcellation at $z = -4$mm. (b): Parcel-based model of the BOLD signal.

test. Moreover, a Fisher statistic has to be computed between experimental conditions and across all HRF-related estimated effects, making interpretation of activation maps delicate as signed comparison can not be performed. In this respect, we claim the necessity of a unique forward HRF model in order to keep a *single* regressor for each condition, and thus allow direct statistical comparison. Still, to conduct the analysis in an efficient and reliable manner, local estimation is performed at a regional scale coarser than the voxel level. 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 homogeneity is guaranteed within each parcel. In this respect, the assumption of a shape-invariant HRF is maintained. Actually, we resort to [7] and Fig. 1(a) shows a slice resulting from such a parcellation procedure.

B. *Within-parcel formulation*

Here, the parcel-based model of the BOLD signal introduced in [1] is adopted: a linear time-invariant (LTI) system. 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 fMRI time course measured in voxel V_j at times $(t_n)_{n=1:N}$ (where $t_n = nTR$, N being the number of scans and TR, the time of repetition) then reads

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

This model remains time-invariant while it incorporates spatially-varying and stimulus-related magnitudes, called *Neural Response Levels* (NRLs) (a_j^m) in the following. $\boldsymbol{X}^m = (x_{t_n - d\Delta t}^m)_{n=1:N, d=0:D}$ denotes the $N \times (D+1)$ binary matrix that codes the arrival times of the mth stimulus which are approximated to fit a Δt-sampled grid, where Δt is the sampling period of the HRF ($\Delta t < TR$). Vector $\boldsymbol{h} = (h_{d\Delta t})_{d=0:D}$ represents the unknown HRF shape in parcel $\mathcal{P}$. Note also that $\boldsymbol{P}\boldsymbol{\ell}_j$ models a low-frequency trend to account for physiological artifacts and that $\boldsymbol{b}_j \sim \mathcal{N}(0, \sigma_{\boldsymbol{b}_j}^2)$ stands for the noise. For simplicity reasons, we have only considered a Gaussian white noise model.

III. THE DETECTION-ESTIMATION PROBLEM

A. *Spatial mixture modelling*

We propose to *estimate* the HRF shape $\boldsymbol{h}$ and the corresponding NRLs $(\boldsymbol{a}^m)_{m=1:M}$ in $\mathcal{P}$. Our aim is also to *detect* which voxels in $\mathcal{P}$ are involved in the experimental paradigm. In [2], we extended an originally IMM approach [1] to account for spatial correlation by expressing a prior SMM relying on a binary Markov random field (MRF). As our approach stands in the Bayesian framework, other prior knowledge is formulated upon every other sought object in the LTI model. The reader is referred to [1] for their expressions. However, for the sake of clarity we recall details on the NRL SMM from [2] in what follows.

B. *NRLs prior*

We assume that different types of stimulus induce statistically independent NRLs *i.e.,* $p((\boldsymbol{a}^m) \,|\, \boldsymbol{\theta_a}) = \prod_m p(\boldsymbol{a}^m \,|\, \boldsymbol{\theta}^m)$ with $(\boldsymbol{a}^m) = (\boldsymbol{a}^m)_{m=1:M}$, $\boldsymbol{a}^m = (a_j^m)_{j=1:J}$ and $\boldsymbol{\theta_a} = (\boldsymbol{\theta}^m)_{m=1:M}$. Vector $\boldsymbol{\theta}^m$ denotes the set of unknown hyper-parameters 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 activating ($q_j^m = 1$) or not ($q_j^m = 0$) in response to stimulus m, so that the NRL a_j^m is normally distributed according to $(a_j^m \,|\, q_j^m = i) \sim \mathcal{N}(\mu_{i,m}, v_{i,m})$, with $i = 0, 1$. We impose $\mu_{0,m} = 0$ for the mean of the NRLs in non-activating voxels, leading to $\boldsymbol{\theta}^m = [v_{0,m}, \mu_{1,m}, v_{1,m}]$. Here, we introduce space-varying probabilities $\Pr(q_j^m = 1) = \lambda_{j,m}$, $\Pr(q_j^m = 0) = 1 - \lambda_{j,m}$ through a spatially correlated Ising prior on the binary variables $\boldsymbol{q}^m$, while the NRLs remain independent conditionally to $\boldsymbol{q}^m$. More precisely, a *symmetric* Ising MRF is considered:

$$\Pr(\boldsymbol{q}^m \,|\, \beta^m) \propto \exp\big(-\beta^m \sum_{j \sim k} \omega_{jk} I(q_j^m = q_k^m)\big),$$

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): $\omega_{jk} \propto 1/d(V_j, V_k)$ where $d(V_j, V_k)$ is the distance between V_j and V_k. Note also that this MRF is *hidden* since $(\boldsymbol{q}^m)_{m=1:M}$ are not observed in (1).

Combining all information, we get a prior spatial mixture model for every stimulus type:

$$p(\boldsymbol{a}^m \,|\, \boldsymbol{\theta}^m) = \sum_{\boldsymbol{q}^m} \Big(\prod_{j=1}^{J} p(a_j^m \,|\, q_j^m, \boldsymbol{\theta}^m)\Big) \Pr(\boldsymbol{q}^m \,|\, \beta^m). \quad (2)$$

Parameter $\beta^m > 0$ in $\Pr(\boldsymbol{q}^m \,|\, \beta^m)$ which controls the amount of spatial smoothing is presently set by hand. Its estimation is fairly conceivable and currently under study. This first requires the estimation of the partition function (the normalisation constant of the MRF) which is computationaly intensive, and then the inclusion of a Metropolis step into the Gibbs Sampling procedure to propose suitable values of the correlation factor. Actually this field is strongly investigated in statistical physics [8], [9].

C. *The joint posterior distribution*

Considering the constructed model and assuming no further prior dependence between parameters, Bayes' rule yields [1]:

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

The nuisance variables $\boldsymbol{\ell}_j$ can be integrated out leading to:

$$p(\boldsymbol{h}, (\boldsymbol{a}^m), \boldsymbol{\Theta} \,|\, \boldsymbol{y}) \propto \left(\prod_{j=1}^{J} \epsilon_j^{-N-1+Q} \right) \sigma_{\boldsymbol{h}}^{-D} \exp\left(-\frac{\boldsymbol{h}^{\mathrm{t}} \boldsymbol{R}^{-1} \boldsymbol{h}}{2\sigma_{\boldsymbol{h}}^2} \right)$$

$$\exp\left(-\frac{\sum_{j=1}^{J} \widetilde{\boldsymbol{y}}_j^{\mathrm{t}} \boldsymbol{Q}_j \widetilde{\boldsymbol{y}}_j}{2} \right) \prod_{m=1}^{M} \Big(p(\boldsymbol{a}^m \,|\, \boldsymbol{\theta}^m) p(\boldsymbol{\theta}^m) \Big). \quad (3)$$

with $\boldsymbol{Q}_j = (\boldsymbol{I}_N - \boldsymbol{P}\boldsymbol{P}^{\mathrm{t}}) / \epsilon_j^2$, $\widetilde{\boldsymbol{y}}_j = \boldsymbol{y}_j - \boldsymbol{S}_j \boldsymbol{h}$ 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 (3), by sequentially generating random samples from the full conditional pdfs of all the unknown parameters and hyper-parameters. Finally, posterior mean (PM) estimates are approximated using the samples according to: $\widehat{x}^{\mathrm{PM}} = (K - I)^{-1} \sum_{k=I+1}^{K} x^{(k)}, \ \forall x \in \{\boldsymbol{h}, (\boldsymbol{a}^m), (\boldsymbol{q}^m), \boldsymbol{\Theta}\}$ where I stands for the length of the burn-in period. The sampling scheme for the posterior mixtures $((\boldsymbol{a}^m), (\boldsymbol{q}^m))$ is detailed in the next paragraph while for other quantities of interest $(\boldsymbol{h}, \boldsymbol{\Theta})$, the reader may refer to [1, Appendix A].

D. NRLs a posteriori

Since the prior on the NRLs $(\boldsymbol{a}^m)$ is a Gaussian mixture and the likelihood is Gaussian, the full posterior pdf of $(\boldsymbol{a}^m)$ is also a Gaussian mixture. From (3), it can be shown that each $a_j^m \in (\boldsymbol{a}^m)$ 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:

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

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

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

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

$$\lambda_{i,j}^m = \left(1 + \frac{r_{1-i,j}^m}{r_{i,j}^m} \frac{\pi_{1-i,j}^m}{\pi_{i,j}^m} \right)^{-1} \quad (4)$$

with $r_{i,j}^m = \big(v_{i,j}^m / v_{i,m}\big)^{1/2} \exp\big((\mu_{i,j}^m)^2 / v_{i,j}^m - i(\mu_i^m)^2 / v_i^m\big)$ and $\pi_{i,j}^m = \Pr(q_j^m = i \,|\, q_{k \in \mathrm{N}_j}^m, \beta^m)$. To calculate (4), we need to evaluate:

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

which only depends on the labels in the neighborhood N_j.

IV. SIMULATION RESULTS

To date, no experiment explained the causality of the BOLD response in terms of neural activation and vascular properties. Thus the estimation-detection task induced by

fMRI analysis is a completely *blind* estimation process. In this respect, a linear time-invariant model is admittedly a simplified approach. It is therefore relevant to test such an estimation procedure on data embedding a more "realistic" structure than the one used in the generative model. In this respect, we assess the robustness of our approach against spatial variability in the shape of the HRF and temporal non-stationarity in the NRLs.

A. Artificial fMRI datasets

We generated a random impulse sequence for $M = 2$ different stimuli (or conditions). These two sets of trials (30 trials per stimulus with 30 "blank" stimuli) produced a paradigm sequence of length $\simeq 180\mathrm{sec}$ with a mean inter-stimulus interval of 2 seconds.

2D binary label $\boldsymbol{q}^m$ slices of size 20x20 voxels for every conditions were hand-drawn with "plausible" shapes for activation clusters. This enables a first deviation from the assumed model in our estimation procedure, since these maps are not generated from the Ising MRF. Conditionally to $\boldsymbol{q}^m$, normally-distributed "basal" NRLs were simulated:

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

with a better contrast-to-noise ratio (SNR) for $m = 1$. Hence, in what follows we focus on the more degraded situation ($m = 2$). Simulated NRL and label maps for condition $m = 2$ are presented in Fig. 2.
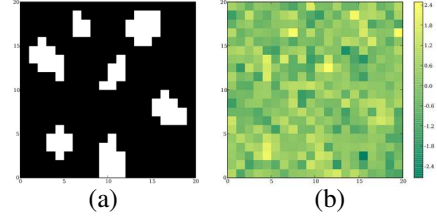

Fig. 2. Condition $m = 2$: (a) labels (white-activating) (b) Basal NRLs .

HRFs were simulated using Bezier curves parametrisation (which do not derive from the Gaussian HRF prior) where time-to-peak (TP), time-to-undershoot (TU) and undershoot amplitude (UA) values are spatially varied (see Fig. 3(a)-(b) for TU and UA respectively). At each voxel, simulated BOLD signal was computed according to (1) with $\sigma_{b_j}^2 = 0.5$ and significant polynomial low-frequency drift Pl_j.

B. Hemodynamic Response Function shape variability

Functional homogeneity of the considered ROI through a shape-invariant HRF is a strong assumption in our model. However, some invalidating physiological factors might occur, such as local heterogeneity in the vascular system or partial volume effect (voxels composed of different tissue types). We simulated HRF Bezier curves with spatially variable parameters, that is the TP, TU, and UA parameters. Here, TP variability can be interpreted as a delay in the neural response to a given stimulus type. Spatial distribution of these parameters follows a 2D Gaussian field as shown

in Fig. 3(a) for TP and Fig. 3(b) for UA. Fig. 3(c) depicts the HRF estimation results against simulated situation: the estimated shape tends towards the average over all different simulated HRFs. As a comparison, Fig. 3(d) shows that an adaptive HRF estimation outperforms the SPM-like approach in terms of detection, with more sensitivity and specificity. Fig. 3(e)-(f) shows the same comparison on detection maps. Here some clusters are not detected by the SPM-like approach whereas they are with an adaptive HRF (see lower right of maps).
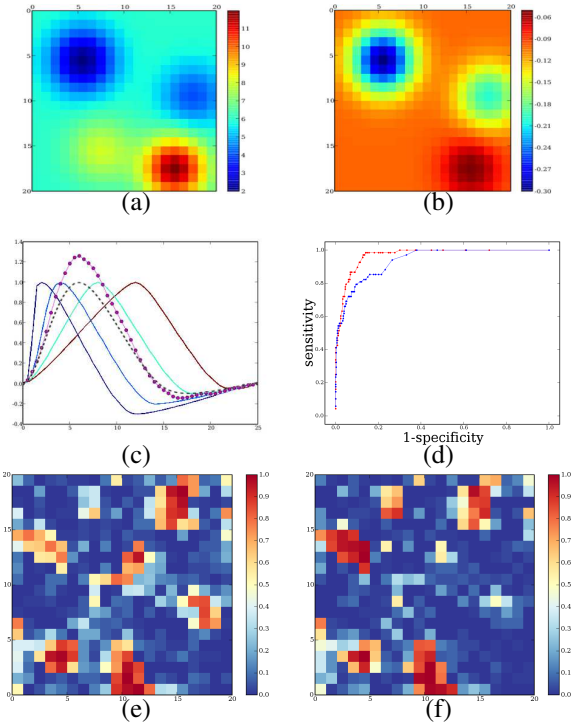

Fig. 4. ROC curves over $\widehat{(q^2)}^{\mathrm{PM}}$ when increasing the temporal deviation of NRLs, values of which are shown in the rectangular insert on the lower right.

Fig. 3. (a) Times to Peak 2D map. (b) Undershoot Amplitudes 2D map. (c) $\widehat{h}^{\mathrm{PM}}$ (dotted line), scale adjusted by $\overline{(\widehat{a}^m)}^{\mathrm{PM}}$ in activating class - h_j^{true} (solid lines) from different voxels colored as in TP map - $\overline{h}_j^{true}$ in dashed black line. (d) ROC curves over $\widehat{q^2}^{\mathrm{PM}}$. Red line: using adaptive HRF estimation, blue line: using a fixed HRF shape. (e) Label estimates $\widehat{q^2}^{\mathrm{PM}}$, with adaptive HRF estimation. (e) Label estimates $\widehat{q^2}^{\mathrm{PM}}$, with a fixed HRF.

C. Inter-trial Neural Response Level variability

Temporal invariance assumed for NRLs is not physiologically intuitive since metabolism can fluctuate or patient attention may decrease as fMRI experiments often last long (up to 40 mins). To account for such deviation from the generative model, we introduced inter-trial variability by injecting variance upon basal NRLs: $a_{j,\tau}^m \sim \mathcal{N}(a_j^m, v_T)$, for $\tau \in (t_n)_{n=1:N} | x_\tau^m = 1$. Our estimation-detection procedure is tested with 15 v_T incremental values so that: $a_{j,\tau}^{m\ (k)} \in I^{(k)}, I^{(k)} = [a_j^m \pm k 20\%]_{95\%}$ with $k \in [0; 15]$. As shown in Fig. 4, the detection is robust against the inter-trial variability of the NRLs up to a deviation of 100% from the basal value, where both sensitivity and specificity fall under 0.8. When investigating the values of $(\widehat{a}^m)^{\mathrm{PM}}$, they also tend towards the average over all the time-varying NRLs values used in simulation (result not shown).
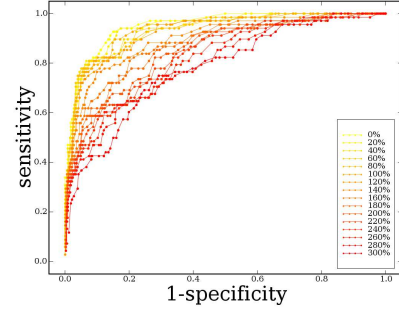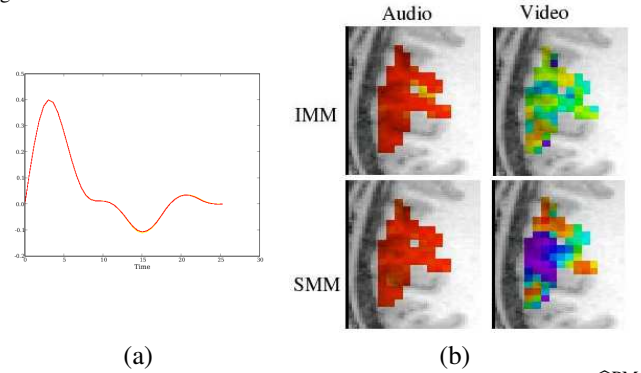

Fig. 5. Results with real fMRI Localizer experiment. (a) Normalised $\widehat{h}^{\mathrm{PM}}$ estimate. (b) Maps of PM label estimates - SMM vs IMM. Rainbow colors range from 0.0-blue to 1.0-red

V. REAL DATA RESULTS

Real fMRI data were recorded during an experiment designed to map auditive and visual brain functions, which consisted of a single session of $N = 125$ scans lasting $TR = 2.4$ s each. The chosen parcel $\mathcal{P}$ (632 voxels) is a cluster obtained with SPM from thresholded t maps at $p = 0.001$. This cluster results from the "*audio minus video*" comparison that elicits activation when the auditory stimulus leads to a stronger response than the visual one. Therefore, the voxels in $\mathcal{P}$ should not be involved in visual perception *a priori*.

We compare the SMM and IMM approaches on the same real data (*cf.* [1]). Fig. 5(a) depicts the HRF estimate obtained with the SMM extension, which is comparable to the one reported in [1] using IMM. Fig. 5(b) shows results on activation detection (PM estimations of labels), mapped on anatomical images. Maps of $\widehat{q}^{\mathrm{PM}}$ first appear to be more contrasted with SMM, yielding to more certainty in the estimation. Moreover, detecting less activations, SMM should lead to a higher specificity than IMM. Since this region is almost only dedicated to auditory processes, SMM is thus more consistent by detecting much less activations in response to the visual condition than IMM.

VI. CONCLUSION

In [2], we stated the gain of a SMM versus IMM approach in term of sensitivity and specificity on simulated data.

Such improvement is confirmed here on real fMRI data with more contrasted label estimates. In another respect, we assess the robustness and flexibility of this Bayesian MCMC estimation procedure: even if the activation delay is not taken into account, the PM HRF estimate still provides a sensible response which "averages" every time-to-peak values. Moreover, the introduction of an inter-trial NRL variability corresponding for instance to habituation effects is well supported by the model. We also observed that the detection seems to be strongly conditioned by the HRF model. Indeed, a fixed HRF approach as in SPM yields a loss in detection of whole clusters when performed on data generated with spatially variable HRFs. Whereas our adaptive HRF approach is able to make a better recovery of these clusters.

Tests were also performed by using different mixture models for NRLs simulations (results not shown). Distribution of posterior mean estimation of NRLs tended to match a simulated Gamma-Gaussian mixture, even if the prior and posterior formulation are bi-Gaussian. We highlight here the strong flexibility of Gibbs sampling, which is able to fit a target distribution which diverges by nature from prior knowledge.

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] T. Vincent, P. Ciuciu, and J. Idier, "Spatial mixture modelling for the joint detection-estimation of brain activity in fMRI," in *32th Proc. IEEE ICASSP*, Honolulu, Hawaii, Apr. 2007, vol. I, pp. 325–328.

[3] 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.

[4] 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*, Toulouse, France, May 2006, vol. V, pp. 1093–1096.

[5] K.J. Friston, "Imaging neuroscience: Principles or maps?," *Proc. Natl. Acad. Sci. USA*, vol. 95, pp. 796–802, 1998.

[6] M. Woolrich, M. Jenkinson, J. Brady, and S. Smith, "Fully Bayesian spatio-temporal modelling of fMRI data," *IEEE Trans. Med. Imag.*, vol. 23, no. 2, pp. 213–231, Feb. 2004.

[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] Shu-Chiuan Chang, Jesper Lykke Jacobsen, Jesus Salas, and Robert Shrock, "Exact potts model partition functions for strips of the triangular lattice," *J.STAT.PHYS.*, vol. 114, pp. 763, 2004.

[9] Tomasz Luczak and Eric Vigoda, "Torpid mixing of the wang-swendsen-koteckÃ$\frac{1}{2}$ algorithm for sampling colorings.," *J. Discrete Algorithms*, vol. 3, no. 1, pp. 92–100, 2005.