



Spatially Adaptive Mixture Modeling for Analysis of fMRI Time Series.

Thomas Vincent, Laurent Risser, Philippe Ciuciu

► To cite this version:

Thomas Vincent, Laurent Risser, Philippe Ciuciu. Spatially Adaptive Mixture Modeling for Analysis of fMRI Time Series.. IEEE Transactions on Medical Imaging, 2010, 29 (4), pp.1059-1074. 10.1109/TMI.2010.2042064 . cea-00470594

HAL Id: cea-00470594

<https://cea.hal.science/cea-00470594>

Submitted on 7 Apr 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Spatially adaptive mixture modeling for analysis of fMRI time series

Thomas Vincent, Laurent Risser and Philippe Ciuciu, *Member, IEEE*

Abstract—Within-subject analysis in fMRI essentially addresses two problems, the *detection* of brain regions eliciting evoked activity and the *estimation* of the underlying dynamics. In [1, 2], a detection-estimation framework has been proposed to tackle these problems jointly, since they are connected to one another. In the Bayesian formalism, detection is achieved by modeling activating and non-activating voxels through independent mixture models (IMM) within each region while hemodynamic response estimation is performed at a regional scale in a nonparametric way. Instead of IMMs, in this paper we take advantage of spatial mixture models (SMM) for their non-linear spatial regularizing properties. The proposed method is *unsupervised* and *spatially adaptive* in the sense that the amount of spatial correlation is automatically tuned from the data and this setting automatically varies across brain regions. In addition, the level of regularization is specific to each experimental condition since both the signal-to-noise ratio and the activation pattern may vary across stimulus types in a given brain region. These aspects require the precise estimation of multiple partition functions of underlying Ising fields. This is addressed efficiently using first *path sampling* for a small subset of fields and then using a recently developed fast extrapolation technique for the large remaining set. Simulation results emphasize that detection relying on supervised SMM outperforms its IMM counterpart and that unsupervised spatial mixture models achieve similar results without any hand-tuning of the correlation parameter. On real datasets, the gain is illustrated in a localizer fMRI experiment: brain activations appear more spatially resolved using SMM in comparison with classical General Linear Model (GLM)-based approaches, while estimating a specific parcel-based HRF shape. Our approach therefore validates the treatment of unsmoothed fMRI data without fixed GLM definition at the subject level and makes also the classical strategy of spatial Gaussian filtering deprecated.

Index Terms—fMRI; estimation of partition function; Bayesian analysis; MCMC; Spatial mixture models; joint detection estimation; MRF; Ising field.

I. INTRODUCTION

Since the first report of the BOLD effect in human [3], functional MRI (fMRI) has represented a powerful tool to non-invasively study the relation between cognitive task and the hemodynamic (BOLD) response, which indirectly reflects evoked neuronal activity. 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 given stimulus types or to behavioral tasks, while the

second one concerns the *estimation* of the underlying temporal dynamics, usually referenced as the characterization of the Hemodynamic Response Function (HRF). The first question has been extensively treated in the context of General Linear Models (GLM) in the last decade; see for instance [4, 5]. The second question has been investigated in a parametric [6–10], semi-parametric [11–13] or nonparametric [14–18] setting. The main differences between these approaches lie in the underlying assumptions regarding the evoked responses: parametric methods treat the HRF as a deterministic function whose parameters have to be identified while nonparametric approaches define the hemodynamic filter as a stochastic object of finite length. In the same vein, a stochastic modeling of hemo- and neurodynamics has been conducted in a state space formulation either relying on physiological models [19–24] or on less informed evolution schemes (*e.g.*, ARX, ARMAX models) [25–27].

Physiological Balloon-like models [19–24] are attractive in their ability to jointly estimate the HRF properties and the neuronal events eliciting evoked responses that can be seen under a delayed and filtered form in fMRI BOLD signals. Nonetheless, their highly nonlinear structure induces identifiability problems especially when fitting their parameters from non-activating or evoked-free fMRI time series. Hence, they are not able to provide a reliable answer to the joint question of brain activity detection and hemodynamics estimation from every fMRI time series. In [1, 2], a combined approach has been proposed to reconcile both issues in a region-based analysis *i.e.*, on a set of prespecified *parcels*, which are supposed to be functionally homogeneous. In [28–30], such parcels were identified using parcellation algorithms that resort either to statistical or agglomerative clustering under connectivity constraints. Given the parcellation result, a key idea is to assign to each parcel a single but unknown nonparametric HRF shape, making the definition of spatially varying general linear models feasible. The main motivation here was to adapt the GLM definition to the local microvascular network with a restricted number of regressors. Within a given parcel, the evoked response is supposed to be modulated in magnitude across experimental conditions (or stimulus types) and voxels to accommodate voxel-dependent and stimulus-specific signal fluctuations. Following [16, 17], the HRF was modeled as a random quantity and physiological prior information is introduced to recover a smooth time course. Recently, several works segregating neuro- and hemodynamics events from fMRI time series have been proposed in the context of linear dynamic systems [31, 32]. Regarding the dynamic aspect, these contributions extend previous works since they are able to

Thomas Vincent, Laurent Risser and Philippe Ciuciu are with NeuroSpin (CEA), 91191 Gif-sur-Yvette, cedex, France and with IFR 49, Paris, France. Tel./Fax: 33 16908 7785/7980. E-mails: firstname.lastname@cea.fr.

Copyright (c) 2009 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

recover the Neuronal Response (NR) or both the HRF and NR. However, from a spatial point of view, these works remain massively univariate and proceed either voxel by voxel or on averaged profiles (*e.g.*, profile of independent components in ICA [33]).

To enhance the estimation of regression coefficients in the GLM context, global stationary and more recently local non-stationary spatial regularization models have been introduced [13, 34–36]. These Bayesian spatial priors encode similarity between neighboring voxels and estimate one or several degrees of smoothness of the parameter images throughout the brain. In this regard, they do not require data to be smoothed prior to entering a statistical model. However, since these models take place in the family of edge-preserving regularization models [37, 38], posterior inference¹ is not numerically tractable at low computational cost without considering a mean-field approximation [39] and the Variational Bayesian (VB) formalism [40]. This kind of approximations breaks down numerical complexity by assuming conditional posterior independence, which means independence across voxels in the fMRI context. It also provides closed form updating rules for iterative estimation of parameters of interest and hyper-parameters provided that conjugate priors have been retained. Nonetheless, the literature in spatial regularization offers alternative choices that make the definition of non-stationary spatial models and their global optimization [41, 42] or simulation feasible [43–49], even if VB approximation can also be derived [50]. These models enter in the class of Spatial Mixture Models (SMM) in which the spatial structure is embedded on *hidden* allocation variables that specify the voxel states in the fMRI context (activating or non-activating) through a *discrete* Markov Random Field (MRF) [46, 48, 51, 52]. In this regard, the detection problem becomes a segmentation issue that can be addressed using a region-based method [43, 53] instead of an edge-preserving restoration one [37, 38, 54].

In [1, 2], the authors have investigated the choice of the best prior mixture model that serves as prior distribution on the evoked response magnitude, called the Neural Response Level (NRL) hereafter. More precisely, it has been shown that the use of inhomogeneous IMM (*e.g.*, gamma-Gaussian) permits to better disentangle activating from non-activating voxels at the expense of higher computational cost. Nonetheless, in case of very few activations in a given parcel, the IMM model overestimates the false positive rate (specificity). Both on synthetic and real fMRI datasets, it has been demonstrated in [55] that the homogeneous *supervised* SMM (SSMM) outperforms its IMM counterpart in terms of false positive control and activation cluster recovery. In this regard, our work [55] is close to previous contributions [46, 48, 50, 52, 56, 57] but more general in the sense that our purpose is not limited to the segmentation of a statistical map.

Akin to [46, 51], in the present paper, we jointly estimate HRFs and perform activation detection. However, our HRF estimation procedure proceeds parcelwise and not voxelwise

and does not rely on any parametric assumption. We also generalize contributions [46, 51] in the following directions:

- We do not impose HRF parameters (delay, peak magnitude, width, ...) to be sampled from parametric distributions (*eg.* Gaussian or gamma densities).
- Whatever the type of event-related fMRI paradigm (*ie* slow or fast), our approach can handle multiple experimental conditions through the estimation of stimulus varying NRLs.
- The trial-by-trial variability or the non-stationarity of the BOLD response captured in [51] in the context of slow event-related designs can also be efficiently addressed in the proposed JDE framework irrespective of the experimental design. For the sake of simplicity, this paper assumes a stationary model of NRLs (constant across trials) but non-stationary extensions already exist to model the so-called *repetition suppression* effect in cognitive neuroscience [58, 59].
- Following [44, 48] but in contrast to [46, 51], we develop an *unsupervised* extension that makes the segmentation of brain activations fully automatic. The unsupervised aspect means that the MRF coupling parameter is automatically tuned from the data. Our unsupervised regularization is also spatially adaptive meaning that the amount of regularization varies across parcels. This requires a precise estimation of the partition function of the MRF defined over each parcel. A small subset of these partition functions are estimated by path sampling [60–62] while the large remaining part is accurately approximated using fast and precise extrapolation techniques [63, 64].
- Our spatial regularization varies across *experimental conditions* in every parcel since both the Signal-to-Noise Ratio (SNR) and the activation pattern may fluctuate from one condition to another according to the involvement of the given parcel in the experimental paradigm.
- Our approach is also able to account for deactivations or negative BOLD responses [65] occurring for instance in epilepsy, by introducing a 3-state Potts MRF as spatial model instead of an Ising one; see [66] for details.

Hence, the proposed methodology introduces spatially adaptive levels of regularization at a reasonable computational cost in adopting a fully *exact* Bayesian inference of brain activity. Parcel-based parameters of interest (HRF shape, NRLs) as well as hyper-parameters are estimated in the Posterior Mean (PM) sense from unsmoothed fMRI time series after convergence of a hybrid Metropolis within Gibbs sampling procedure [67].

The rest of this paper is structured as follows: Section II summarizes our parcel-based approach for whole brain analysis of fMRI time series by extending the Joint Detection-Estimation (JDE) framework proposed in [1, 2] to spatial mixture models. The extension of this methodology to the whole brain analysis is addressed in Section III where the specific concerns of unsupervised and spatially adaptive regularization (USMM) are dealt with. As shown, they respectively require specific numerical schemes for estimating single and then multiple partition functions (PF) of Ising fields. The construction of the brain parcellation is then discussed and

¹*e.g.*, the computation of the maximum a posteriori or the posterior mean estimates do not admit a closed form expression.

illustrated in Section IV. Section V then shows the interest of USMM in comparison with IMM and SSMM on artificial datasets. Simulations results also demonstrate the importance of accounting for spatial fluctuations of the HRF shape. Then, our approach is illustrated on real fMRI datasets in Section VI in order to demonstrate that there exist relevant alternative strategies to spatial filtering of fMRI time series, specifically for contrasts of interest associated to low SNR. The last section finally summarizes our contributions and opens the discussion about the pros and cons of our method and its future improvements and extensions.

II. THE SPATIALLY REGULARIZED JDE APPROACH

A. Notations

We denote vectors and matrices with bold lower and upper case letters, respectively (e.g., $\mathbf{y}$ and $\mathbf{P}$). A vector is by convention a column vector. Scalars are denoted with non-bold lower case letters. The transpose is denoted by t . Unless stated otherwise, subscripts i, j, m and n are respectively indexes over mixture components, voxels, stimulus types and time points. The probability distribution functions (pdf) are denoted using calligraphic letters (eg, $\mathcal{N}$ and $\mathcal{G}$ for the Gaussian and gamma distributions).

The JDE framework proposed in [2, 55] relies on a prior parcellation of the brain into $\mathcal{P} = (\mathcal{P}_\gamma)_{\gamma=1:\Gamma}$ functionally homogeneous and connected parcels [28], where typically $\Gamma \approx 500$ to cover the whole brain (see Subsection IV-A and [68, 69] for their computation and the assessment of Γ^2); see also Fig. 2 for illustrations. Every parcel $\mathcal{P}_\gamma$ comprising voxels $(V_j)_{j=1:J_\gamma}$ is characterized by a specific model of the BOLD signal, which consists in estimating a *single* HRF shape $\mathbf{h}_\gamma$ whatever the number of experimental conditions involved in the paradigm. In what follows, we first summarize the parcel-based modeling of the BOLD signal before specifying the priors involved over each parcel. In other words, our Bayesian parcel-based inference of brain activity is independently repeated over Γ different parcels with different sets of parameters on the corresponding fMRI time series. For notational simplicity, in what follows we drop off the γ index that makes reference to the parcel except for $\mathbf{h}_\gamma$ and J_γ .

B. Forward parcel-based model of the BOLD signal

The forward bilinear model of the BOLD signal introduced in [1] and extended in [2] to account for serial correlation of fMRI time series is a time-invariant model that characterizes each and every parcel $\mathcal{P}_\gamma$ by a single neurovascular impulse response $\mathbf{h}_\gamma$ and a NRL for each voxel and stimulus type. As illustrated in Fig. 1, this means that the HRF shape $\mathbf{h}_\gamma$ is assumed constant within $\mathcal{P}_\gamma$, while its magnitude a_j^m can vary in space ($j = 1 : J_\gamma$) and across experimental conditions ($m = 1 : M$), where M is the total number of stimulus types. Then, the generative BOLD model reads:

$$\forall j = 1 : J_\gamma, \quad \mathbf{y}_j = \sum_{m=1}^M a_j^m \mathbf{X}^m \mathbf{h}_\gamma + \mathbf{P} \ell_j + \mathbf{b}_j, \quad (1)$$

²The algorithm is available to the community in the fMRI Toolbox of the Brainvisa software at <http://brainvisa.info>.

where:

- $\mathbf{y}_j = (y_{j,n})_{n=1:N}$ denotes the fMRI signal measured in voxel V_j at times $n = 1:N$ (N is the number of scans).
- $\mathbf{X}^m = (x_{n-d\Delta t}^m)_{n=1:N, d=0:D}$ is a $N \times (D+1)$ binary matrix coding for the occurrences of the m stimulus type, with Δt is the sampling period of the unknown HRF $\mathbf{h}_\gamma = (h_{d\Delta t, \gamma})_{d=0:D}$ in $\mathcal{P}_\gamma$.
- a_j^m stands for the NRL in voxel V_j for condition m . Hence, the activation time course associated to the m th stimulus type in voxel V_j is given by $\mathbf{h}_\gamma \times a_j^m$. Let also $\mathbf{A} = [\mathbf{a}^1 | \dots | \mathbf{a}^M]$ be the whole NRL matrix in $\mathcal{P}_\gamma$ where $\mathbf{a}^m = (a_j^m)_{j=1:J_\gamma}$.
- $\mathbf{P}$ is a low frequency orthogonal matrix of size $N \times Q$. To each voxel is attached an unknown weighting vector ℓ_j to estimate the trend in V_j . We denote $\mathbf{L} = [\ell_1 | \dots | \ell_{J_\gamma}]$ the set of low frequency drifts involved in $\mathcal{P}_\gamma$.
- $\mathbf{b}_j \in \mathbb{R}^N$ is the noise in V_j and follows a first-order autoregressive process: $\mathbf{b}_j \sim \mathcal{N}(\mathbf{0}, \sigma_j^2 \mathbf{\Lambda}_j^{-1})$ where $\mathbf{\Lambda}_j$ is tridiagonal and depends on the AR parameter ρ_j [2].

Although the noise structure is correlated in space [70, 71], we neglect such spatial dependency and consider the fMRI time series $\mathbf{Y} = [\mathbf{y}_1 | \dots | \mathbf{y}_{J_\gamma}]$ independent in space but *not* identically distributed. The reason is twofold: first, neglecting the spatial dependencies of noise is tenable when the BOLD signal model itself is flexible enough to account for HRF shape fluctuations. Indeed, part of the usually observed spatial correlation of the noise is due to a misspecification of the BOLD signal model. Second, the noise correlation is much lower than that of the evoked BOLD response. While its modeling introduces additional computational complexity, there is no evidence in the literature that ignoring this correlation induces strong bias on the sought parameters. Hence, our likelihood reads:

$$p(\mathbf{Y} | \mathbf{h}_\gamma, \mathbf{A}, \mathbf{L}, \boldsymbol{\theta}_0) \propto \prod_{j=1}^{J_\gamma} |\mathbf{\Lambda}_j|^{1/2} \sigma_j^{-N} \exp\left(-\frac{\tilde{\mathbf{y}}_j^t \mathbf{\Lambda}_j \tilde{\mathbf{y}}_j}{2\sigma_j^2}\right) \quad (2)$$

where $\theta_{0,j} = (\rho_j, \sigma_j^2)$, $\boldsymbol{\theta}_0 = (\theta_{0,j})_{j=1:J_\gamma}$ and $\tilde{\mathbf{y}}_j = \mathbf{y}_j - \sum_m a_j^m \mathbf{X}^m \mathbf{h}_\gamma - \mathbf{P} \ell_j$.

Note that model (1) is bilinear in the sense that Eq. (1) linearly depends on $\mathbf{h}_\gamma$ when $\mathbf{A}$ is fixed and *vice-versa*. This means that the Maximum Likelihood (ML) solution $(\mathbf{h}_\gamma^*, \mathbf{A}^*)$ cannot be distinguished from any other pair $(\mathbf{h}_\gamma^*/s, \mathbf{A}^* \times s)$ whatever the scale parameter $s > 0$. The Bayesian formalism is helpful to get rid of such identifiability problems and define a reference scale. In the next subsection, we introduce priors on $\mathbf{h}_\gamma$ and $\mathbf{A}$, which are helpful to fix this scale to an arbitrary value c . However, this value is not necessarily optimal for exploring the posterior distribution. Hence, instead of normalizing $\mathbf{h}_\gamma$ deterministically ($|\mathbf{h}_\gamma| = c$), it is shown in Appendix C how this reference scale s can be selected to speed up the exploration of the posterior distribution.

C. Bayesian inference

The Bayesian approach developed in [2] introduces proper priors on the unknown parameters $(\mathbf{h}_\gamma, \mathbf{A})$ in order to recover




Fig. 1. **Parcel-based regional BOLD model.** The size of each parcel $\mathcal{P}_\gamma$ varies typically between a few tens and a few hundreds of voxels: $80 \leq J_\gamma \leq 350$. The number M of experimental conditions involved in the model usually varies from 1 to 5. In our example, $M = 2$, the NRLs (a_j^1, a_j^2) corresponding to the first and the second conditions are surrounded by circles and squares, respectively. Note that our model accounts for asynchronous paradigms in which the onsets do not necessarily match acquisition time points. As illustrated, the NRLs take different values from one voxel to another. The HRF h_γ can be sampled at a period of 1s and estimated on a range of 20 to 25s (e.g., $D = 25$). Most often, the LFD coefficients ℓ_j are estimated on a few components ($Q = 4$).

a robust estimate of brain activity (localization and activation profile). Akin to [2, 17], the prior density for the HRF remains Gaussian, $h_\gamma \sim \mathcal{N}(\mathbf{0}, v_h \mathbf{R})$ with $\mathbf{R} = (D_2^T D_2)^{-1}$, which allows us to estimate a smooth HRF shape since D_2 is the second-order finite difference matrix penalizing therefore abrupt changes. Moreover, the extreme time points of the HRF can be constrained to zero if necessary [17].

1) *Neural Response Levels:* Regarding NRLs, according to the maximum entropy principle we assume that different types of stimulus induce statistically independent NRLs i.e., $p(\mathbf{A} | \boldsymbol{\theta}_a) = \prod_m p(\mathbf{a}^m | \boldsymbol{\theta}^m)$ with $\boldsymbol{\theta}_a = (\boldsymbol{\theta}^m)_{m=1:M}$. Vector $\boldsymbol{\theta}^m$ denotes the set of unknown hyperparameters related to the m th stimulus type. However, for certain class of paradigms (e.g., attention and motor dual tasks; priming effects, ...) this assumption may potentially be unlikely. In such cases, a between-condition prior covariance matrix could be introduced, the difficulty lying in the choice of a relevant and know correlation model to limit the computational complexity.

Mixture models are introduced to segregate activating voxels from non-activating ones. To this end, let q_j^m be the allocation variable that states whether voxel V_j is activating ($q_j^m = 1$) or not ($q_j^m = 0$) in response to stimulus m . The NRLs still remain independent conditionally upon $\mathbf{q}^m$. This means that $p(\mathbf{a}^m | \mathbf{q}^m, \boldsymbol{\theta}^m) = \prod_j p(a_j^m | q_j^m, \boldsymbol{\theta}^m)$ for every condition m . In the case of an IMM [1, 2], the marginal density of the NRLs factorizes over voxels and reads:

$$p(\mathbf{a}^m | \boldsymbol{\theta}^m) = \prod_{j=1}^{J_\gamma} \sum_{i=0}^1 p(a_j^m | q_j^m, \boldsymbol{\theta}^m) \Pr(q_j^m = i | \boldsymbol{\theta}^m). \quad (3)$$

An important feature of IMM lies in the definition of the mixing probability (or weight) $\Pr(q_j^m = i)$ which is constant over voxels and so independent of j . This means that IMM explicitly estimates the proportions of voxels in the activating and non-activating classes.

Instead, *spatial* mixture models are introduced here to favor clustered activations and the mixing probabilities become

space-varying. Hereafter, the marginal density no longer factorizes over voxels and reads:

$$p(\mathbf{a}^m | \boldsymbol{\theta}^m) = \sum_{\mathbf{q}^m} \left[\prod_{j=1}^{J_\gamma} p(a_j^m | q_j^m, \boldsymbol{\theta}^m) \right] \Pr(\mathbf{q}^m | \boldsymbol{\theta}^m). \quad (4)$$

Spatial correlation is directly incorporated in the probabilities of activation through a *hidden* Ising MRF on the allocation variables $\mathbf{q}^m$, as already done in image analysis [44, 45] or in neuroimaging [46, 57]. Here, as opposed to IMM, the proportions of voxels for the different classes is not explicit. The prior density on the allocation variables reads:

$$\Pr(\mathbf{q}^m | \beta_m) = Z(\beta_m)^{-1} \exp(\beta_m U(\mathbf{q}^m)) \quad (5)$$

with $U(\mathbf{q}^m) = \sum_{j \sim k} I(q_j^m = q_k^m)$

and $I(A) = 1$ if A is true and $I(A) = 0$ otherwise. The notation $j \sim k$ means that the sum extends over all pairs (j, k) of neighbouring sites. The neighbouring system can be defined either in 3D in the brain volume intersecting region $\mathcal{P}_\gamma$ or in 2D along the cortical surface. In this paper, we only consider the 3D case using 6-connexity. Extensions to 18 or 26 neighborhood system are straightforward. The Ising MRF in (5) has no *external field*. Previous works have shown that anatomical prior information can be embedded in an external field to increase activation probability in the grey matter [46]. The parameter $\beta_m > 0$ in (5) controls the amount of spatial regularization: large values of β_m associate higher probabilities to configurations containing clusters of like-valued neighboring binary variables. Since activation patterns within parcel $\mathcal{P}_\gamma$ should be different from one stimulus type to another, different parameters β_m are considered across m stimulus types.

The normalization constant of the MRF, also called the

partition function $Z(\cdot)$ reads:

$$Z(\beta_m) = \sum_{\mathbf{q}^m \in \{0,1\}^J} \exp(\beta_m U(\mathbf{q}^m)) \quad (6)$$

and guarantees that the MRF defines a pdf.

In what follows, we assume that $(a_j^m | q_j^m = i) \sim \mathcal{N}(\mu_{i,m}, v_{i,m})$, with $i = 0, 1$. We impose $\mu_{0,m} = 0$ for the mean of the NRLs in non-activating voxels, leading to $\boldsymbol{\theta}^m = [v_{0,m}, \mu_{1,m}, v_{1,m}, \beta_m]$. Note that a Bernoulli-Gaussian formulation has also been tested in fMRI in [46]. This modeling corresponds to a degenerated mixture ($v_{0,m} = 0$).

2) *Noise and drift parameters*: To complete the Bayesian model, priors are required for all the remaining parameters. The noise and drift parameters, $\boldsymbol{\theta}_0$ and $\mathbf{L}$ respectively, are assumed independent in space: $p(\boldsymbol{\theta}_0, \mathbf{L} | v_\ell) = \prod_j p(\theta_{0,j}) p(\ell_j | v_\ell)$ and without informative prior knowledge, the following priors are chosen: $\ell_j \sim \mathcal{N}(\mathbf{0}, v_\ell \mathbf{I}_Q)$ and $p(\rho_j, \sigma_j^2) = \sigma_j^{-1} I(|\rho_j| < 1)$ to ensure stability of the AR(1) noise process [72]. Non-informative Jeffrey priors are retained for hyper-parameters such as the drift and HRF shape variances: $p(v_h, v_\ell) = (v_h v_\ell)^{-1/2}$.

3) *Mixture parameters*: Similarly, the prior considered for $v_{0,m}$ is $p(v_{0,m}) = v_{0,m}^{-1/2}$ because we *do* expect non-activating voxels in any parcel. Hence, class 0 should never be empty a priori. If this assumption is not tenable, we could introduce a conjugate prior (an inverse Gamma law) denoted as $\mathcal{IG}(a_{v_0}, b_{v_0})$, as already done for the class of activating voxels. We would then avoid degeneracy problem that could prevent its sampling. In the same way, a proper prior $\mathcal{N}(a_{\mu_1}, b_{\mu_1})$ is chosen for $\mu_{1,m}$ ($a_{\mu_1} = x$ and $b_{\mu_1} = y$). Finally, the prior on $\boldsymbol{\beta} = (\beta_m)_{m=1:M}$ is independent and identically distributed (iid) across conditions and follows a uniform pdf over fixed range: $p(\beta_m) = \mathcal{U}([0, \beta_{\max}])$ with $\beta_{\max} = 1.6$. Interestingly, such proper prior defined over $[0, \beta_{\max}]$ allows us to easily compute Ising field PF at discrete β -values defined over the same range; see Section III.

4) *Posterior distribution*: Considering the constructed model and assuming no further prior dependence between parameters, Bayes' rule gives us:

$$\begin{aligned} p(\mathbf{h}_\gamma, \mathbf{A}, \mathbf{L}, \boldsymbol{\Theta} | \mathbf{Y}) &\propto p(\mathbf{Y} | \mathbf{h}_\gamma, \mathbf{A}, \mathbf{L}, \boldsymbol{\theta}_0) p(\mathbf{A} | \boldsymbol{\theta}_\mathbf{A}) p(\mathbf{h}_\gamma | v_h) \\ &\quad p(\mathbf{L} | v_\ell) p(\boldsymbol{\theta}_0) p(\boldsymbol{\theta}_\mathbf{A}) p(v_h, v_\ell) \quad (7) \\ &\propto v_h^{-\frac{D}{2}} v_\ell^{-\frac{J_\gamma Q}{2}} \prod_{j=1}^{J_\gamma} \frac{(1 - \rho_j^2)^{1/2}}{\sigma_j^{N+1}} I(|\rho_j| < 1) \\ &\quad \exp\left(-\sum_{j=1}^{J_\gamma} \left[\frac{1}{2\sigma_j^2} \tilde{\mathbf{y}}_j^t \boldsymbol{\Lambda}_j \tilde{\mathbf{y}}_j + \frac{1}{2v_\ell} \|\ell_j\|^2\right]\right) \\ &\quad \exp\left(-\frac{\mathbf{h}_\gamma^t \mathbf{R}^{-1} \mathbf{h}_\gamma}{2v_h}\right) \prod_{m=1}^M p(\boldsymbol{\theta}^m) p(\mathbf{a}^m | \boldsymbol{\theta}^m). \end{aligned}$$

Our Bayesian model is too complex to be amenable to analytical calculations. Hence, we resort to Gibbs sampling to sample the posterior distribution (7). To facilitate sampling and avoid manipulations of the marginal pdf $p(\mathbf{a}^m | \boldsymbol{\theta}^m)$, the allocation variables $\mathbf{q}$ are introduced into the Gibbs sampler. However, the generated Markov chain admits the same target distribution

(7). Posterior mean (PM) estimates are then computed from these samples according to the following rule: $\hat{\mathbf{x}}^{\text{PM}} = (T_c - T_0)^{-1} \sum_{t=T_0+1}^{T_c} \mathbf{x}^{(t)}$, $\forall \mathbf{x} \in \{\mathbf{h}_\gamma, \mathbf{A}, \boldsymbol{\Theta}\}$ where T_0 stands for the length of the burn-in period. Also, for classification or detection purpose, the marginal maximum a posteriori criterion is employed: $(\hat{q}_j^m)^{\text{MAP}} = \arg \max_i \Pr(q_j^m = i | \mathbf{y}_j)$.

The sampling scheme for the posterior mixtures $(\mathbf{A}, \mathbf{Q})$ is detailed in Appendix B while for other quantities of interest $(\mathbf{h}_\gamma, \boldsymbol{\Theta})$, the reader may refer to [2, Appendix B] since their sampling remains unchanged. Table I summarizes the complete procedure for a given parcel.

III. WITHIN-PARCEL ANALYSIS: UNSUPERVISED SMMS

Unsupervised spatial regularization consists in automatically tuning parameter vector $\boldsymbol{\beta}$ from the dataset $\mathbf{Y}$ in a *given* parcel $\mathcal{P}_\gamma$. In the proposed hybrid Gibbs sampler (see Table I), this is implemented by making $\boldsymbol{\beta}$ a random vector and then adding a sampling block involving $\boldsymbol{\beta}$ within the sampling loop. As shown below, a Metropolis-Hastings (MH) algorithm is designed to draw candidates (see subsection III-A) but importantly this requires the knowledge of the partition function $Z(\cdot)$, which will be addressed in subsection III-B.

A. Metropolis-Hastings (MH) $\boldsymbol{\beta}$ sampling step

The full conditional posterior distribution $p(\boldsymbol{\beta} | \mathbf{Q})$ only depends on the hidden MRFs $p(\mathbf{q}^m | \beta_m)$ and the prior $p(\boldsymbol{\beta})$:

$$\begin{aligned} p(\boldsymbol{\beta} | \mathbf{Q}) &= \prod_{m=1}^M p(\beta_m | \mathbf{q}^m) \propto \prod_{m=1}^M \Pr(\mathbf{q}^m | \beta_m) p(\beta_m) \\ &\propto \prod_{m=1}^M Z(\beta_m)^{-1} \exp(\beta_m U(\mathbf{q}^m)) \mathcal{U}_{[0, \beta_{\max}]}(\beta_m). \quad (8) \end{aligned}$$

This distribution depends on the partition function $Z(\cdot)$, which is *independent* of m . Therefore, its precise estimation is a prerequisite to any attempt for sampling from $p(\beta_m | \mathbf{q}^m)$. Note that here, we consider a MH sampling step while in other contributions like [73] a Gibbs sampler on a discrete β -grid was proposed. As shown below, our approach is computationally more expensive but is less dependent on the grid definition and can provide a more accurate estimate of $\boldsymbol{\beta}$. For every condition m , let us consider a constrained random walk to generate a candidate $\beta_m^{(c)} \sim \mathcal{N}_{[0, \beta_{\max}]}(\beta_m^{(t)}, \xi^2)$ from the current value $\beta_m^{(t)}$ at iteration t . The acceptance probability of this candidate reads $\alpha(\beta_m^{(t)} \rightarrow \beta_m^{(c)}) = \min(1, A_{t \rightarrow c}^m)$ where the acceptance ratio $A_{t \rightarrow c}^m$ is given by:

$$A_{t \rightarrow c}^m = \frac{p(\beta_m^{(c)} | \mathbf{q}_m^{(t)}) g(\beta_m^{(t)} | \beta_m^{(c)})}{p(\beta_m^{(t)} | \mathbf{q}_m^{(t)}) g(\beta_m^{(c)} | \beta_m^{(t)})}.$$

For a truncated Gaussian instrumental law $g(\cdot | x) \sim \mathcal{N}_{[0, \beta_{\max}]}(x, \xi^2)$, this ratio reads:

$$A_{t \rightarrow c}^m = \frac{Z(\beta_m^{(t)})}{Z(\beta_m^{(c)})} \exp((\beta_m^{(c)} - \beta_m^{(t)}) U((\mathbf{q}^m)^{(t)})) B_{t \rightarrow c}^m, \quad (9)$$

$$\text{with } B_{t \rightarrow c}^m = \frac{\text{erf}(-\xi^{-2} \beta_m^{(c)}) \text{erf}(\xi^{-2} (\beta_{\max} - \beta_m^{(t)}))}{\text{erf}(-\xi^{-2} \beta_m^{(t)}) \text{erf}(\xi^{-2} (\beta_{\max} - \beta_m^{(c)}))}.$$

Note that the $B_{t \rightarrow c}^m$ term handles the Gaussian truncation since the transition kernel is not symmetric.

If $\beta_m^{(c)}$ is not accepted, then $\beta_m^{(t+1)} = \beta_m^{(t)}$. As shown in Eq. (9), an accurate evaluation of $A_{t \rightarrow c}^m$ relies on a precise estimation of the ratio of partition functions $Z(\beta)$. Exact evaluations of $Z(\beta)$ in a reasonable time are actually impossible except for very small grids since it requires the evaluation of $U(\mathbf{q}^m)$ for all possible combinations of $\mathbf{q}^m$ whose number is exponentially related to the grid dimensions. Fast and robust estimation of $Z(\cdot)$ is thus a key issue to get an unsupervised approach. Robustness is achieved by avoiding numerical overflow, thus considering Eq. (9) under its logarithmic form. Akin to [44, 57], efficiency is obtained here adopting path sampling to perform this estimation on a discrete grid of β -values, as detailed in what follows. Combined with interpolation of $Z(\cdot)$ outside the grid points, this strategy makes the MH sampler of β feasible.

B. Estimation of single partition function

It is worth noting that in the special case where $\beta_m = 0$, the elements of $\mathbf{q}^m$ are spatially independent, and Z has a closed form expression $Z(0) = 2^{J\gamma}$. This property is fundamental for path sampling since this algorithm essentially relies on relative evaluations of $Z(\beta_m)$ for distinct values of β_m [60–62, 74]. Path-sampling is actually an extension of importance sampling for estimating ratios of normalizing constants, by considering series of easy-to-sample unnormalized intermediate densities. Such a strategy was proven efficient to tabulate the PF for the Ising case; see [73] for details.

Let $f_1(\mathbf{q})$ and $f_2(\mathbf{q})$ be two densities with respective support Ω_1 and Ω_2 such that $\Omega_1 \subseteq \Omega_2$. Each density is known up to a normalizing constant, $f_s(\mathbf{q}) = p_s(\mathbf{q})/Z_s$. The importance sampling identity is meaningful in order to estimate ratios of normalizing constants [60]:

$$\frac{Z_1}{Z_2} = \mathbb{E}_2 \left[\frac{p_1(\mathbf{q})}{p_2(\mathbf{q})} \right], \quad (10)$$

where $\mathbb{E}_s$ denotes the expectation with respect to f_s . Applying Eq. (10) to the context of Ising fields and considering a Monte Carlo approximation, we get:

$$\forall \beta_m, \hat{Z}(\beta_m) \simeq \frac{Z(\beta_m^0)}{L} \sum_{l=1}^L \exp((\beta_m - \beta_m^0) U(\mathbf{q}_l^m)). \quad (11)$$

where $Z(\beta_m^0)$ is assumed to be known and $\{\mathbf{q}_1^m, \dots, \mathbf{q}_L^m\}$ are distributed according to $\text{Pr}(\mathbf{q}^m | \beta_m^0)$. To avoid numerical overflow, Eq. (11) is usually handled in logarithmic scale:

$$\log \hat{Z}(\beta_m) \simeq \log \frac{Z(\beta_m^0)}{L} + \log \sum_{l=1}^L \exp((\beta_m - \beta_m^0) U(\mathbf{q}_l^m)).$$

To estimate Z around β_m^0 , L realizations $\mathbf{q}_l^m$ are thus generated at $\beta_m = \beta_m^0$ before computing $U(\mathbf{q}_l^m)$, $\forall l = 1 : L$. Starting with $\beta_m^0 = 0$, the overall scheme consists in computing $\log \hat{Z}(\cdot)$ on a relatively dense β -grid (i.e., $\beta_m^g = g\beta_{\max}/G$, $g \in \{0, 1, \dots, G\}$ with e.g., $G = 10^3$) using Eq. (11) and then in interpolating $\log \hat{Z}(\cdot)$ for β_m -values outside the grid points. This incremental procedure makes the estimation of $Z(\beta_m^{g+1})$

relatively robust since it only relies on $\hat{Z}(\beta_m^g)$. Interpolation between grid points is robust even if it is linearly performed since $Z(\beta_m)$ is an *increasing* function of β_m and does not present strong irregularities³. Furthermore, results strongly depend on our ability to generate draws $\mathbf{q}^m$ that cover the whole support of the field.

While the Gibbs sampler is practically efficient at small β -values i.e., in little correlated configurations, it becomes impractical in highly correlated fields (large β -values) because it moves too slowly from one configuration to another. In such cases, a more efficient algorithm is the iterative Swendsen-Wang (SW) sampling scheme [75] provided that the Ising MRF to be sampled is symmetric. For non-symmetric models embedding external fields that could be useful to account for anatomical-based priors in the neuroimaging context [46, 52, 57, 76], several generalizations have been proposed [44, 49, 77, 78]. However in this paper, we only focus on the symmetric case.

IV. WHOLE BRAIN ANALYSIS: SPATIALLY ADAPTIVE USMM

A. Derivation of brain parcellation from fMRI data

As outlined in Section II, our BOLD signal modeling is parcel-dependent and spatial regularization differs over each parcel $\mathcal{P}_\gamma$. The methodology proposed in Section III to make this regularization unsupervised is therefore applicable to all parcels separately making our whole brain analysis of fMRI data fully spatially adaptive. This makes sense given that the stimulus-specific SNRs also vary in space. Of course, the critical issue is to exhibit such functionally homogeneous parcellation of brain. To this end, several algorithms have been proposed [28, 29, 68, 69]: they segregate the brain into connected and functionally homogeneous regions by minimizing a criterion reflecting both the spatial and functional structures of the dataset. The functional part of this criterion can be computed either from the raw fMRI time series or from voxel-based hemodynamic features (time-to-peak and time-to-undershoot, peak and undershoot magnitudes, ...) [79], which can be extracted from nonparametric HRF estimates [17, 18].

Since all these parcels have a specific geometry, we proceed to the PF estimation of 3D Ising MRF for each out of them. Still, the computation cost increase is reasonable as several hundreds of parcels are typically necessary to cover the entire brain. For an averaged-size parcel of 250 voxels, the PF estimation requires around 10 seconds which results in a increase of about 30 minutes for a whole brain analysis, compared to a time of 1.5 hour for the complete analysis, it yields a large increase of 33%. Hence, in the next part, we introduce fast numerical alternatives based on efficient approximations of 3D Ising partition functions, so that the overall cost would be quite negligible.

B. Multiple 3D Ising MRF partition function estimation

Algorithms with polynomial time complexity [39, 80] provide efficient alternatives to a single PF estimation. However,

³except for large grids, around the critical value β_m^* where the phase transition occurs

TABLE I

STEPS OF THE UNSUPERVISED JDE ALGORITHM IN A GIVEN PARCEL $\mathcal{P}_\gamma$.

- Estimate $Z(\cdot)$ using Eq. (11) for $\beta^g = g\beta_{\max}/G, g = 0 : G$
- Interpolate linearly $Z(\cdot)$ outside the grid points
- Setting up: choose $\mathbf{h}_\gamma^0, \mathbf{A}^0, (\boldsymbol{\sigma}, \boldsymbol{\rho}^2)^0, \boldsymbol{\theta}_a^0$.
- Iteration t : draw samples $\mathbf{h}_\gamma^{(t)}, \mathbf{A}^{(t)}, \mathbf{Q}^{(t)}, (\boldsymbol{\sigma}^2, \boldsymbol{\rho})^{(t)}, \boldsymbol{\theta}_a^{(t)}$ from the full conditional posterior pdfs:
 - HRF: $\mathbf{h}_\gamma^{(t)} \sim p(\mathbf{h}_\gamma | \mathbf{y}, (\ell_j^{(t-1)}, \mathbf{A}^{(t-1)}, \boldsymbol{\theta}_0^{(t-1)}), \text{cf. [2, Eq. (B.1)]}$.
 - HRF variance: $v_h^{(t)} \sim p(v_h | \mathbf{h}_\gamma^{(t)}), \text{cf. [2, Eq. (B.2)]}$.
 - Drift: $\forall j = 1 : J_\gamma, \ell_j^{(t)} \sim p(\ell_j^t | \mathbf{y}, \mathbf{h}_\gamma^{(t)}, \mathbf{A}^{(t-1)}, (\boldsymbol{\theta}_{0,j})^{(t-1)}), \text{cf. [2, Eq. (B.1)]}$.
 - Drift variance: $v_\ell^{(t)} \sim p(v_\ell | \ell^{(t)}), \text{cf. [2, Eq. (B.2)]}$.
 - NRLs: External loop over stimuli m ; Internal loop over voxels j :
 - * Identify $\sum_{i=0}^1 \lambda_{i,j}^m \mathcal{N}(\mu_{i,j}^m, v_{i,j}^m)$ using Eqs. (14)-(16).
 - * $(u_j^m)^{(t)} \sim \mathcal{U}[0, 1]$; if $(u_j^m)^{(t)} \leq (\lambda_{1,j}^m)^{(t)}$, then $(q_j^m)^{(t)} = 1$ otherwise $(q_j^m)^{(t)} = 0$.
 - * $(a_j^m)^{(t)} | (q_j^m)^{(t)} = i \sim \mathcal{N}((\mu_{i,j}^m)^{(t)}, (v_{i,j}^m)^{(t)}),$
 - * Update $\bar{\mathbf{y}}_j$ accordingly.
 - Noise parameters: $\forall j = 1 : J_\gamma, (\sigma_j^2, \rho_j)^{(t)} \sim p((\sigma_j^2, \rho_j) | \mathbf{y}_j, \mathbf{h}_\gamma^{(t)}, \mathbf{a}_j^{(t)}), \text{cf. [2, Eqs. (B.15)-(B.16)]}$.
 - Mixture parameters: for every condition m ,
 - * Variance of NRLs for nonactivated voxels: $(v_{0,m})^{(t)} \sim p(v_{0,m} | \eta_{0,m}^{(t)}, \nu_{0,m}^{(t)}), \text{cf. [1, Eq. (12)]}$.
 - * Variance of NRLs for activated voxels: $(v_{1,m})^{(t)} \sim p(v_{1,m} | \eta_{1,m}^{(t)}, \nu_{1,m}^{(t)}), \text{cf. [1, Eq. (13)]}$.
 - * Mean of NRLs for activated voxels: $(\mu_{1,m})^{(t)} \sim p(\mu_{1,m} | v_{1,m}^{(t)}, \eta_{1,m}^{(t)}, \nu_{1,m}^{(t)}), \text{cf. [1, Eq. (14)]}$.
 - * Regularization level β_m for every m : cf. Eq. (9).
- Iterate until convergence is achieved and compute posterior mean estimates (Discard $x^1, \dots, x^{T_0}$ as burn-in period): $\hat{x}^{\text{PM}} = (T_c - T_0)^{-1} \sum_{t=T_0+1}^{T_c} x^{(t)}, \forall x \in \{\mathbf{h}_\gamma, \mathbf{A}, \mathbf{Q}, \boldsymbol{\sigma}^2, \boldsymbol{\rho}^2, \boldsymbol{\theta}_a\}$




Fig. 2. **Parcellations of the brain:** (a) using the Voronoi based method - (b) as obtained with an optimal anatomo-functional parcellation [28].

none of them is able to perform multiple PF estimation at the same time, a problem that arises in our whole brain parcellation-based analysis. Since several hundreds of grids of variable size and shape are manipulated in our fMRI application as illustrated in Fig. 2, fast estimation of multiple PF is necessary. To this end, in [63, 64] we have proposed several hybrid schemes which consist in resorting to path sampling to get

log-scale estimates $(\log \hat{Z}_{\mathcal{G}_p}(\beta))_{p=1:P}$ in a small subset of *reference* graphs $(\mathcal{G}_p)_{p=1:P}$ and then in using extrapolation formulas to obtain $\log \tilde{Z}_T(\beta)$ for the large remaining set of brain regions to be analyzed, referenced here by a *test* graph $\mathcal{T}$. In what follows, we summarize the main ideas and underlying assumptions of this extrapolation techniques.

1) *Linear/Bilinear extrapolation schemes:* In [73], the authors have proposed a linear regression procedure to estimate $(\log \hat{Z}_{\mathcal{G}_p}(\beta))_{p=1:P}$ as a function of the number of cliques in the grids $(\mathcal{G}_p)_{p=1:P}$. Estimates of $\log \tilde{Z}_T(\beta)$ are then linearly computed using the estimated regression coefficients and the number of cliques in $\mathcal{T}$: $\log \tilde{Z}_T(\beta_k) = \hat{a}_k c_T$ at each β_k regularization level ($\beta_k = k\Delta\beta$). A bilinear extension of this technique, which also takes the number of sites in the grid into account, has been developed in [63]. This procedure was shown to be efficient to estimate log-PFs in *small and irregular grids*⁴ such as those appearing in our fMRI application. However, the accuracy of linear/bilinear PF extrapolations strongly depends on the homogeneity and the number of *reference* grids: the less homogeneous the reference set, the larger the approximation error. These reasons motivate the derivation of a more reliable and versatile approach to grid inhomogeneities as requested by our fMRI application.

2) *Min-max extrapolation scheme:* We have therefore proposed a more robust 2-step technique in [64]:

- a) Akin to [73], reference log-PFs $\log \hat{Z}_{\mathcal{G}_p}(\beta_k)$ are estimated using path sampling on inhomogeneous reference grids $(\mathcal{G}_p)_{p=1:P}$.
- b) For any test grid $\mathcal{T}$, the quantity $\log Z_T$ is approximated from a *single* reference log-PF estimate out of $(\log \hat{Z}_{\mathcal{G}_p}(\beta))_{p=1:P}$ selected by a grid homogeneity criterion $\mathcal{L}_T(\mathcal{G}_p)$ as well as an approximation error criterion $\mathcal{A}_T(\beta, \mathcal{G}_p)$ defined by:

$$\mathcal{A}_T(\beta, \mathcal{G}_p) = \|\log Z_T(\beta) - \log \tilde{Z}_T(\beta, \mathcal{G}_p)\|^2 / \|\log Z_T(\beta)\|^2$$

$$\text{with } \log \tilde{Z}_T(\beta, \mathcal{G}_p) = \left(\frac{c_T}{c_{\mathcal{G}_p}} (\log \hat{Z}_{\mathcal{G}_p}(\beta) - \log 2) + \log 2 \right), \quad (12)$$

and where $(c_T, c_{\mathcal{G}_p})$ and $(n_T, n_{\mathcal{G}_p})$ are the number of cliques and sites of the Ising MRFs defined over $\mathcal{T}$ and $\mathcal{G}_p$, respectively. The theoretical justification of this construction can be found in [64].

In short, the extrapolation formula (12) is built up according to two principles: *i.*) an unbiased asymptotic approximation error⁵ and *ii.*) an exact approximation of $(\log Z_T(\beta))'$ for $\beta \rightarrow 0^+$. The reference grid $\mathcal{G}_{\text{ref}}$ is exhibited using a min-max principle, which consists in minimizing wrt all reference grids $(\mathcal{G}_p)_{p=1:P}$ the maximal approximation error $\mathcal{A}(\beta, \mathcal{G}_p)$ given by (see [64]):

$$\mathcal{A}_T(0, \mathcal{G}_p) \triangleq \|(n_T - 1) - c_T(n_{\mathcal{G}_p} - 1)/c_{\mathcal{G}_p}\|^2 / n_T^2. \quad (13)$$

Once $\mathcal{G}_{\text{ref}}$ has been identified, the log-PF estimate in $\mathcal{T}$ is thus given by $\log \tilde{Z}_T(\beta, \mathcal{G}_{\text{ref}})$ according to Eq. (12).

⁴Here, by irregular grids we make reference to regular lattices combined with non-straight boundaries.

⁵ $\lim_{\beta \rightarrow +\infty} \mathcal{A}_T(\beta, \mathcal{G}_p) = 0$.




Fig. 3. From top to bottom rows: simulated labels Q^{true} , NRLs A^{true} and examples of artificial time series ($M = 2$). On the bottom row, the solid line corresponds to the stimulus induced signal, the dashed line to the artificial data and vertical blue and green lines correspond to stimulus onsets for the two conditions. Simulated signals on the left side depict an active voxel for both conditions with $a^1 = 2.25$ and $a^2 = 1.9$. Signals on the right side depict an inactive voxel with $a^1 = -0.1$ and $a^2 = -0.3$.

V. SIMULATION RESULTS

A. Generation of Artificial fMRI datasets

Artificial fMRI datasets were simulated according to the forward BOLD model defined by 1 and illustrated in Fig. 1, considering two experimental conditions ($M = 2$). Simulation setup at low Signal-to-Noise Ratio (SNR) is reported in Table II which is the general context to examine JDE behaviour on artificial data either deriving from the true Ising prior for the labels (subsection V.D) or from hand drawn labels maps (subsections V. B,C,E). The considered map size is 400 voxels, which is close to the size a parcel obtained from a brain parcellation on real data. Convergence to simulated values is theoretically guaranteed for spatial variables when considering an asymptotic situation, *i.e.*, very large size. Here, we rather focus on a small problem in order to grasp our model behaviour on a more realistic situation.

TABLE II
SETUP FOR THE GENERATION OF ARTIFICIAL FMRI DATASETS.

Paradigm:	
30 trials, $\forall m$, and fast event-related: $\mu_{ISI} = 5s$, $\sigma_{ISI} = 2.9s$	labels: 20x20 maps Drift: $\ell_j \sim \mathcal{N}(0, 3I_4)$ Noise: $\epsilon_j^2 = 2$
NRLs:	$\forall m$ HRF:
$(a_j^m q_j^m = 0) \sim \mathcal{N}(0, 0.3)$ $(a_j^m q_j^m = 1) \sim \mathcal{N}(1.8, 0.3)$	canonical shape duration of 25 seconds <i>i.e.</i> , $D = 49$, $\delta t = 0.5s$




Fig. 4. Comparison of different supervised settings of β (SSMM) versus USMM with respect to the area under ROC (AUROC) for conditions $m = 1$ (left) and $m = 2$ (right). AUROC for IMM is not displayed, but is almost the same as SSMM- $\beta = 0$. The solid line shows the evolution of the AUROC index for supervised JDE runs. The diamond point indicates the AUROC index for the unsupervised setting.

B. IMM, SSMM, USMM comparison

We have compared the efficiency of different models: IMM (independent MM, which do not model any spatial correlation), SSMM (supervised SMM, where β is set to a fixed value) and USMM (unsupervised SMM, where β is estimated). SSMM was iterated over a grid of β values ranging in $[0, 2]$ to manually search for optimal detection results. The USMM version relied on a path sampling estimate of $\log \hat{Z}(\beta_m)$, over the considered 20x20 map. Our first assessment step dealt with the detection performance when label maps are composed of realistic activation clusters and so do not follow the Ising prior, therefore there is no obvious β value associated to these maps. Moreover, in order to illustrate the need for adaptive β estimation, cluster configuration differs from one condition to another. As shown in Fig. 3, for $m = 1$, clusters are numerous with variable sizes, whereas for $m = 2$ there is only one cluster with a stretched shape.

A good synthetic criterion for detection assessment is the Area Under the Receiver Operating Characteristic curve (AUROC) which represents the sensitivity-specificity compromise. Results for this criterion are shown in Fig. 4. In terms of detection quality, IMM is comparable to using $\beta = 0$ with SSMM, whatever the condition of interest. When varying β with SSMM, too low β yields not enough spatial regularization whereas too high β values yield too much correlation and poor detection results, so that there is an “optimal” range of β values for a given label map. This optimal range varies across conditions: $[0.4, 1]$ for $m = 1$ and $[0.8, 1.5]$ for $m = 2$, because of different activation patterns. This illustrates the need for an adaptive estimation of β . USMM fulfills this role by providing β estimates within the good ranges, *i.e.* $\hat{\beta}_1^{PM} = 0.86$ and $\hat{\beta}_2^{PM} = 0.99$. USMM is able to adapt across condition and adjust $\hat{\beta}^{PM}$ to the actual level of correlation within the field.

To go one step further in the mixture model behaviour, we then present results on label maps estimates, NRLs estimates and prior mixture components estimates. Regarding SSMM, we present these results for relevant β values extracted from the AUROC curves in Fig. 4: $\beta = 0.2$ - few spatial regularization, $\beta = 0.8$ - optimal setting for $m = 1$, $\beta = 1.1$ - optimal setting for $m = 2$ and $\beta = 1.8$ - very high spatial regularization level.

Label map estimates are depicted in Fig. 5. First, we can state that SSMM is not beneficial over IMM if β is not tuned to a right value. Indeed Fig. 5- $[\beta = 0.2]$ and Fig. 5-IMM show close results and Fig. 5- $[\beta = 1.8]$ illustrates a misspecification of β where a too high value yields too much spatial regularization and results worse than IMM.

The optimal β setting regarding SSMM for $m = 1$ is $\beta_1 = 0.8$ but this value is not suitable for $m = 2$. Conversely, the optimal setting for $m = 2$ is not suitable for $m = 1$. The unsupervised approach (see Fig. 5-USMM) is able to recover these nearly optimal β settings: $(\hat{q}^1)^{PM}$ maps for USMM and SSMM- $[\beta = 0.8]$ are quite identical and so are $(\hat{q}^2)^{PM}$ maps for USMM and SSMM- $[\beta = 1.1]$.

The same conclusions can be drawn from results on $\hat{\alpha}^{PM}$ depicted in Fig. 6, where we can state the impact of a right regularization with USMM, SSMM ($\beta_1 = 0.8$ and SSMM- $\beta_2 = 1.1$) whereas IMM and SSMM with misspecified β -values lead to less sensitive results.

Considering the estimation of the prior mixture model components $\hat{\theta}$, as shown in Fig. 7, IMM leads to very overlapped estimated components with a high variance for the activating class and so does SSMM- $[\beta = 0.2]$. A misspecification of β as shown with SSMM- $[\beta = 1.8]$ leads to degenerate mixtures where component variances are very high. Again, the optimal setting for one condition may not be suitable for the other condition, see Fig. 7- $[\text{SSMM-}\beta = 0.8]$. USMM yields consistent estimates of mixture components for all conditions.

C. Robustness to SNR decrease

We are interested here in the impact of signal degradation by noise. To this end, we varied the noise variance in simulations such as $\sigma_j^2 \in [0.5, 5]$ and we compared the efficiency of USMM versus IMM. Fig. 8 depicts the evolution of the AUROC index for the two conditions in terms of average SNR computed for each simulation, according to: $\overline{\text{SNR}} = \frac{1}{J} \sum_j 10 \log_{10} \frac{\|\sum_m a_j^m X^m \mathbf{h}\|^2}{\sigma_j^2}$. For $m = 1$, where activations are more scattered, the USMM asset only holds for $\overline{\text{SNR}} > 0.2$. Indeed, below 0.2 the degradation is too important and there is not enough regularity within the data to enable a proper recovery. For $m = 2$, where activations are more clustered, USMM clearly outperforms IMM whatever the noise level and still maintains a good AUROC index around 0.9 in the worst situation. Similarly, estimated $\hat{\beta}^{PM}$ is stable against SNR decrease when there is enough signal in the data and then strongly decreases when there is too much noise (for $\overline{\text{SNR}} < 0.2$).

D. Reproducibility of USMM results

Previous simulation results tested the USMM JDE procedure only on two particular configurations. We focus here on result reproducibility for the estimation of β . For doing so with a hidden field, we need to simulate multiple realizations of label maps and the most convenient way is to generate label maps following the Ising model, each of them being associated to a β value. Accordingly, each simulation data set was generated as explained in Table II, with label maps q^1 and




Fig. 8. Impact of degradation by noise. $\hat{\beta}$ (top row) and AUROC (bottom row) in terms of SNR for $m = 1, 2$. On bottom row, the dashed line corresponds to IMM, the solid line to USMM.

q^2 deriving here from the Ising prior at a given β value and we applied the JDE procedure to these data. This simulation step was repeated over 100 random Ising field realizations and over a β -grid defined as $\beta^{\text{true}} \in \{0, 0.1, \dots, 1\}$ leading to a Monte Carlo validation scheme where the recovery of true β -values can be assessed as well as variability across estimations. Fig. 9 shows the results of this procedure and some typical maps of Ising fields at increasing β -values with detection results in terms of false positives/negatives. Note that we only present results for $m = 1$ since they are the same for $m = 2$. The overall estimation of the regularization level β is accurate since the regression β_m^{true} vs. $\hat{\beta}_m^{PM}$ does not strongly deviate from the first bisector. Optimal recovering appears for $\beta_m^{\text{true}} > 0.2$ with a small bias, while highest errors occur for very uncorrelated situations. This quite imperfect adequation of $\hat{\beta}_m^{PM}$ estimates regarding simulated quantities is first due to the small grid dimensions: convergence of hyper-parameters are only guaranteed in the asymptotic situation, which means in the case of very large grids and second, due to a low SNR, which can favor misclassification and induce biased estimates $\hat{\beta}_{1,2}^{PM}$. suitable for real data analysis: salt-and-pepper situations are not realistic in the fMRI context where activations are expected to be clustered. We also emphasize the fact that Ising fields are *hidden* here so that hyperparameter estimation is a much more delicate issue than in the case of observed fields. We tested the latter situation using our PF extrapolation scheme and results were indeed very accurate (not shown).

Finally, note that dynamics estimation variability regarding this Monte Carlo validation is very little. For any simulated activation field, $\hat{\mathbf{h}}_\gamma^{PM}$ was very close to the simulated $\mathbf{h}_\gamma$ (results not shown). Indeed, the HRF temporal reproducibility was not affected by varying activation localizations.

E. Hemodynamic Response Function shape variability

Spatial modeling is performed at the parcel level through SMM as was illustrated in previous sections V-B to V-D. At




Fig. 5. Comparison of different supervised settings of β (SSMM) versus IMM and USMM with respect to the PM label maps $(\hat{q}^1)^{PM}$ (a) and $(\hat{q}^2)^{PM}$ (c), true activating clusters are rounded in green. Histograms on (b) and (d), for conditions $m = 1$ and $m = 2$ respectively, illustrate the segregation of true activations (red) and true inactivations (blue) in the corresponding maps. Y-axis shows densities. Note: purple regions stand for histogram overlaps.




Fig. 6. Comparison of different supervised settings of β (SSMM) versus IMM and USMM wrt NRL estimates $(\hat{a}^1)^{PM}$ (top) and $(\hat{a}^2)^{PM}$ (bottom).




Fig. 7. Comparison of different supervised settings of β (SSMM) versus IMM and USMM wrt *prior* mixture model component estimates $\hat{\theta}^m$ for $m = 1$ (top) and $m = 2$ (bottom). Dashed line densities correspond to the Gaussian pdf associated to inactivating voxels in the mixture while continuous line densities define the Gaussian pdf for activating voxels.




Fig. 9. Right: illustration of detection results for increasing values of β - blue and red positions are false positives and false negatives, respectively. Left: the mean values over Monte Carlo simulations $\text{mean}_k((\hat{\beta}_1^{\text{PM}})^k)$ (solid line) - error bars indicate the standard deviation across 100 simulations when generating the BOLD signal. The dashed line indicates the ideal target situation $\hat{\beta}_1^{\text{PM}} = \beta_1^{\text{true}}$.

a higher regional scale, the JDE approach enables the HRF shape to vary across parcels, thus also performing spatial modeling, which is the focus of the current section. The same setup presented in V-A was used and we also kept the label maps depicted in Fig. 3 but these data were equally splitted vertically into two ROIs. Two different HRFs that differ from the canonical shape were simulated, as depicted in Fig. 10(a). As we wanted to test the impact of the spatial relaxation of the HRF, we compared detection results for three situations:

- (i) fixing the HRF to its canonical version, which is analogous to the classical GLM situation.
- (ii) estimating one HRF over the concatenation of the two ROIs, which can be compared to GLM using a Finite Impulse Response (FIR) model.
- (iii) estimating one HRF for each ROI.

Apart from the HRF modeling, we applied the USMM approach in all cases.

Fig. 11 shows how HRF modeling impacts detection. The most inflexible case, in Fig. 11(a) (one single fixed canonical HRF) yields the least sensitive results. The simulated HRF for the left ROI is close to the canonical shape (Fig. 10(a)), therefore we get more sensitive results in this part. However, for the right ROI, simulation strongly departs from the canonical shape, especially for the time-to-peak value so that sensitivity is dramatically decreased.

A more flexible situation is estimating one HRF for the whole data, results of which are depicted in Fig. 11(b). Compared to the previous situation, sensitivity is decreased in the left ROI and increased in the right one. Indeed, the estimated HRF is a compromise (or average) between the two simulated ones, as shown in Fig. 10(b-purple line), so that the estimated HRF is a worse solution than the canonical version for the left ROI, but a better one for the right ROI. However, sensitivity is globally enhanced.

Finally, the most flexible situation, in Fig. 11(c), allows a proper recovery of response levels and HRF are well estimated, see Fig. 10(b-orange and blue lines). In another respect, the case (ii) (one HRF for two ROIs) is equivalent to using an inadequate parcellation whereas case (iii) resorts to the true parcellation. Hence, the worse detection results for case (ii) compared to case (iii) also bring up the question of the sensitivity to the input parcellation. This issue is tackled in

section VII.




Fig. 10. HRF simulation (a) and estimation (b). Dashed lines: canonical HRF. (a): HRFs used in simulation (orange is left, blue is right) (b): estimated HRFs when considering two HRFs for the two regions. The purple estimated HRF is obtained when modeling one single HRF across the two regions.




Fig. 11. Impact of HRF modeling on detection - maps of $\hat{q}^1$: (a) fixed canonical HRF (b) one HRF estimated for the 2 ROIs (c) HRF estimation is adaptive across ROIs.

VI. RESULTS ON REAL FMRI DATASETS

We applied the JDE procedure to real *unsmoothed* fMRI data recorded during an experiment designed to map auditory, visual and motor brain functions as well as higher cognitive tasks such as number processing and language comprehension. It consisted of a single session of $N = 125$ scans lasting $TR = 2.4$ s each, yielding 3-D volumes composed of $64 \times 64 \times 32$ voxels. The paradigm was a fast event-related design comprising sixty auditory, visual and motor stimuli, defined in ten experimental conditions (auditory and visual sentences, auditory and visual calculations, left/right auditory and visual clicks, horizontal and vertical checkerboards).

We still compare the three versions of the JDE procedure: IMM, SSMM ($\beta = 0.8$) and USMM, in order to assess the impact of the adaptive spatial correlation model. Posterior mean estimates $\hat{\mathbf{A}}$ have been computed over 5000 realizations of the Gibbs sampler after a burn-in period of 10^3 iterations. Fig. 12 shows normalized contrasts maps of auditory computation (AC) versus auditory sentence (AS), where the modeling of spatial correlation seems to lead more sensitive results as activations in the parietal cortex are highlighted with SSMM and USMM whereas they are not with IMM. Moreover, these activations are coherent with the anatomy since they seem to follow the posterior part of the cingulate sulcus, which implication in numbers processing has been identified [81]. In another respect, Fig. 13 shows normalized contrast maps of auditory induced right click (RAC) versus auditory induced left click (LAC). As expected, the activations lie in the contralateral left motor cortex. Here, only USMM is more sensitive and we illustrate the advantage of an *adaptive*

spatial correlation model. Indeed, estimated $\hat{\beta}^{\text{PM}}$ with USMM for the left auditory click was 0.56 so that the supervised setting of SSMM with $\beta = 0.8$ leads to too much correlation and less sensitive results.

Interestingly, Figs. 12-13 also depict the parcel-dependent maps of $\hat{\beta}^{\text{PM}}$ estimates for the RAC and LAC experimental conditions. The gain in sensitivity in the USMM contrast map ($\hat{a}^{\text{RAC}} - \hat{a}^{\text{LAC}}$) results from a difference in the amount of spatial regularization introduced between the two conditions involved in the contrast. In parcels located in the left motor cortex, the BOLD signal is known to be stronger for the RAC than for the LAC condition. This is a possible interpretation of the lower regularization level estimated for RAC compared to LAC ($\hat{\beta}_{\text{LAC}} \approx 0.5$ vs. $\hat{\beta}_{\text{RAC}} \approx 0.75$) in the activating region outlined in Fig. 13.

Finally, following an illustrative purpose, Figs. 12-13 show estimated HRF shapes in the most active regions. They do not strongly depart from the canonical shape, except for $\hat{h}_{26}$ which time-to-peak is shifted by around 2 seconds. The bumped tail of $\hat{h}_{232}$ can be explained by several hypothesis. There may be a sort of periodic scheme in the stimulus involved by the concerned activating regions. This may also be a particular behaviour of the local vascular system which responds to the undershoot. In any case, we would have to resort to a specific paradigm for a precise investigation. This concern goes beyond the scope of this paper and would be an interesting focus for future works centered on the neuro-vascular dynamics study, which is enabled by our approach.

VII. DISCUSSION AND CONCLUSION

In the present paper, unsupervised and spatially adaptive regularization has been integrated in a Bayesian approach that performs a conjoint detection-estimation of brain activity from fMRI time series. This extension has been achieved by estimating parcel-dependent regularization parameters β , which requires a precise estimation of the underlying 3D Ising field partition function. To this end, an extrapolation algorithm based on path-sampling has been systematically used over each parcel. This algorithm can be carried out prior to cycling over random samples drawn through Gibbs sampling since the partition functions only depend on the topological parcel configurations within the pre-computed parcellation.

Our approach first relies on the definition of an appropriate spatial scale which should be small enough so that the HRF shape invariance can be assumed and big enough so that we benefit from enough HRF reproductibility. This optimal spatial trade-off that impacts temporal modeling is fulfilled by the parcellation scheme we rely on. If we resort to such partitioning for HRF modeling, one could also think of taking the whole brain as the spatial support for the detection part, so that we would consider only one SMM to model response levels. However, this would obviously prevent us from making regularization spatially adaptive. Even if parcellation is well justified for HRF modeling from a physiological point of view, it is still questionable for response levels modeling. Indeed, two regions may be well suited to explain different vascular system properties so that we consider two different HRFs but

they may not be suited for underlying activations which may span these two regions for example. From a practical point of view, we did not observe any impact of parcel boundaries that would prevent activation clusters to span different regions. Moreover the parcellation optimality has been assessed by a sensitivity analysis in [82], where the parcellation method yields the same detection quality as a heavy Monte Carlo scheme which aims at marginalizing the parcellation. Ideally, the parcellation should be jointly exhibited so that parcel labels become random variables in our setting. In this respect, multivariate spatial Gaussian mixture modeling (MSGMM) was applied in [79] to perform the required parcellation. Even if it settled a proof of concept as to merge parcellation with JDE, MSGMM alone suffers from a heavy computation burden.

One could also produce a meaningful functional parcellation by estimating for each parcel a label associated to a class of HRF. After treating an over-segmented parcellation (more parcels than HRF classes), spatially connected parcels which share the same labels could then be merged together and then be subsequently used in another JDE procedure to yield final reliable results. In another respect, reckoning on the second level analysis, parcellation at the group level can also be performed to account for between-subject variability [83]. Indeed, linking together brain parcels rather than voxels allows us to overcome the drawbacks of normalization due to spatial variability.

As a validation purpose, the right estimation of β was assessed on simulations that followed prior modeling. Worst cases arised with highly correlated configurations, which are not expected to be interesting for dealing with real data. Simulations focusing on activation maps that do not follow prior Ising modeling confirmed the asset of SMM versus IMM but also highlighted that a misspecification of a fixed β value (SSMM) can lead to wrongly estimated activation label maps so that spatial correlation modeling does not bring any advantage compared to IMM. On the same simulated data, this limitation of SSMM was cleared by the USMM approach where the optimal setting of β was correctly recovered. On real fMRI datasets, the latter situation is even worse for the SSMM setting. Indeed, the optimal setting of β varies when considering different regions of the brain. In this respect, we identified regions where SSMM- $\beta = 0.8$ provides the same effect maps as IMM whereas USMM was more sensitive ($\hat{\beta} = 0.56$). Hence, our claim is that our approach is not only unsupervised but also an *adaptive* spatial regularization scheme. As a remark, the sensitivity gain compared to IMM was mainly observed for low contrasts between conditions. For contrast involving higher response levels (for example auditive vs. visual conditions), there were no noticeable gain in resorting to USMM. Indeed, in this case there is enough information within data and spatial regularization may be useless.

To summarize, our approach enables a finer recovering of subtle contrasts by adapting the spatial regularization to varying contrast-to-noise ratios as well as various underlying activation patterns.

From a perspective point of view, the current work can




Fig. 12. Comparison of the IMM, SSMM and USMM models wrt the Auditory Computation vs. Sentence (AC/S) normalized contrast maps $(\hat{\mathbf{a}}^{\text{AC}} - \hat{\mathbf{a}}^{\text{AS}}) / \text{std}(\hat{\mathbf{a}}^{\text{AC}} - \hat{\mathbf{a}}^{\text{AS}}) - \hat{\beta}$ maps for the corresponding conditions – $\hat{\mathbf{h}}$ for the active region outlined in purple, compared to a canonical shape. Neurological orientation: left is left.




Fig. 13. Comparison of the IMM, SSMM and USMM models wrt the Right vs. Left Auditory Click (R/LAC). From left to right: normalized contrast maps: $(\hat{\mathbf{a}}^{\text{RAC}} - \hat{\mathbf{a}}^{\text{LAC}}) / \text{std}(\hat{\mathbf{a}}^{\text{RAC}} - \hat{\mathbf{a}}^{\text{LAC}}) - \hat{\beta}$ maps for the corresponding conditions – $\hat{\mathbf{h}}$ for the active region outlined in purple, compared to a canonical shape. Neurological orientation: left is left.

be generalized to 3-class USMM, thus involving Potts fields, as the detection of an additional deactivation voxel state is relevant in several fMRI contexts, like in epilepsy. The path-sampling step extension to Potts model is straightforward and NRL sampling demands little adaptation. However, model selection strategies have to be developed to avoid overfitting or interpretation issues when resorting to 3-class modeling in a region where 2-class modeling is more adequate, *i.e.*, when there is no deactivation. Finally, we intend to validate the current method on group analysis to appraise the impact of our spatial regularization scheme to the sensitivity of group level effect maps.

APPENDIX

A. Generalized Inverse Gaussian distribution

The density of the $\mathcal{GIG}$ distribution reads [84]:

$$f_{\mathcal{GIG}}(z) = \left(\frac{\alpha}{\beta}\right)^{\frac{\lambda}{2}} \frac{z^{\lambda-1}}{2K_{\lambda}(\sqrt{\alpha\beta})} \exp\left(-\frac{1}{2}(\beta z^{-1} + \alpha z)\right),$$

$$z > 0, \quad \lambda \in \mathbf{R}, \quad \begin{cases} \alpha > 0, \beta \geq 0 \text{ pour } \lambda > 0 \\ \alpha > 0, \beta > 0 \text{ pour } \lambda = 0 \\ \alpha \geq 0, \beta > 0 \text{ pour } \lambda < 0 \end{cases}$$

where $K_{\lambda}(\cdot)$ is the third-type Bessel function. Efficient accept-reject simulation techniques do really exist for generating samples distributed according to this pdf.

B. The “neural” response levels

The prior on the NRLs ($\mathbf{A}$) being a Gaussian mixture and the likelihood being Gaussian when $\mathbf{h}_{\gamma}$ is fixed, the full conditional posterior density of $\mathbf{A}$ is also a Gaussian mixture. From Eq. (7), it can be shown that each $a_j^m \in \mathbf{A}$ is obtained by sampling a 2-class posterior spatial Gaussian mixture in

voxel V_j for the m th stimulus type:

$$p(a_j^m | \mathbf{y}_j, a_j^{m' \neq m}, q_{k \sim j}^m, \dots) = \sum_{i=0,1} \lambda_{i,j}^m \mathcal{N}(\mu_{i,j}^m, v_{i,j}^m),$$

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 $\mathcal{N}(\mu_{i,j}^m, v_{i,j}^m)$. As detailed in [2], we have for $i = 0, 1$:

$$v_{i,j}^m = \left(v_{i,m}^{-1} + \frac{\mathbf{g}_m^t \Lambda_j \mathbf{g}_m}{\sigma_j^2}\right)^{-1}, \quad (14)$$

$$\mu_{i,j}^m = v_{i,j}^m \left(\frac{\mathbf{g}_m^t \Lambda_j \mathbf{e}_{m,j}}{\sigma_j^2} + i \frac{\mu_{i,m}}{v_{i,m}}\right) \quad (15)$$

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

$$\lambda_{i,j}^m = \left(1 + \frac{r_{1-i,j}^m \pi_{1-i,j}^m}{r_{i,j}^m \pi_{i,j}^m}\right)^{-1} \quad (16)$$

with $r_{i,j}^m = (v_{i,j}^m / v_{i,m})^{1/2} \exp((\mu_{i,j}^m)^2 / v_{i,j}^m - i(\mu_{i,m}^2) / v_{i,m}^m)$ and $\pi_{i,j}^m = \Pr(q_j^m = i | q_{k \sim j}^m, \beta_m)$. To compute (16), we only need to evaluate

$$\pi_{1-i,j}^m / \pi_{i,j}^m = \exp\left(2\beta_m(2i-1) \sum_{k \sim j} (2q_k^m - 1)\right).$$

which only depends on the labels in the neighborhood of V_j .

C. Sampling of the reference scale s

This new step reads as follows: $[\mathbf{A}^{\text{new}}, \mu_1^{\text{new}}] = [\mathbf{A}^{\text{old}}, \mu_1^{\text{old}}] / s$ and $\mathbf{h}_{\gamma}^{\text{new}} = \mathbf{h}_{\gamma}^{\text{old}} \times s$, where s stands for the scale of the problem that will be drawn along each iteration of the Gibbs sampler. Straightforward application of

the principles exposed in [85, § 3] allows us to derive the full conditional distribution of s from the full posterior pdf:

$$\begin{aligned} f_S(s | \text{rest}) &= |s|^{P-2-M(J_\gamma+1)} f\left(\frac{\mathbf{A}^{\text{old}}}{s}, \frac{\boldsymbol{\mu}_1^{\text{old}}}{s}, s\mathbf{h}_\gamma^{\text{old}} | \mathbf{Q}, \boldsymbol{\Theta}\right) \\ &\propto |s|^{P-2-M(J_\gamma+1)} \exp\left(-\frac{s^2(\mathbf{h}_\gamma^{\text{old}})^t \mathbf{R}^{-1} \mathbf{h}_\gamma^{\text{old}}}{2v_h}\right) \\ &\quad p\left(\frac{\boldsymbol{\mu}_1^{\text{old}}}{s}\right) \prod_{j=1}^{J_\gamma} \mathcal{N}\left(\frac{\mathbf{a}_j^{\text{old}}}{s} \mid \frac{\boldsymbol{\mu}_{q_j}^{\text{old}}}{s}, v_{q_j}\right) \end{aligned} \quad (17)$$

where $(\boldsymbol{\mu}_{q_j}^{\text{old}}, v_{q_j})$ are mean and variance vectors for all stimulus types. Interestingly, we can note that the full conditional pdf (17) of s is independent of the data $\mathbf{Y}$. This is of course due to the scale ambiguity present in our bilinear observation model in Eq. (1). Then, from Eq. (17) it can be deduced that s^2 follows a *Generalized Inverse Gaussian* pdf (see Appendix A), denoted as $\mathcal{GIG}(\lambda, \alpha, \beta)$ where:

$$\lambda = (P - M(J_\gamma + 1))/2, \quad \alpha = (\mathbf{h}_\gamma^{\text{old}})^t \mathbf{R}^{-1} \mathbf{h}_\gamma^{\text{old}} / \sigma_h^2 \quad (18)$$

$$\beta = \frac{\|\boldsymbol{\mu}^{\text{old}}\|^2}{2\sigma_\mu^2} + \sum_{j=1}^{J_\gamma} (\mathbf{a}_j^{\text{old}})^t \Sigma_j^{-1} \mathbf{a}_j^{\text{old}} \quad (19)$$

with $\Sigma_j = \text{diag}_M[v_{q_j}^m]$.

The global sampling scheme of the full posterior distribution attached to a 2-class Gaussian SMM and AR(1) noise becomes:

- a) Sample first $\mathbf{Q}^{(t)}$ and then $(\mathbf{A}^{\text{old}} | \mathbf{Q}^{(t)})$ given that $a_j^m \sim \sum_{i=0}^1 \lambda_{i,j}^m \mathcal{N}(\mu_{i,j}^m; v_{i,j}^m)$, $\forall j, m$
- b) Sample $(\mathbf{h}_\gamma^{\text{old}})$ according to $\mathcal{N}(\boldsymbol{\mu}_h, \Sigma_h)$;
- c) Sample the drift $(\ell_j)_j$ and hyper-parameters $\boldsymbol{\Theta}$ especially $\boldsymbol{\mu}_1^{\text{old}}$;
- d) Simulate $(s^2)^{(t)} \sim \mathcal{GIG}(\lambda, \alpha, \beta)$ with parameters given in (18)-(19) then take its square root $s^{(t)}$.
- e) Apply the stochastic normalization: $[\mathbf{A}^{(t)}, \boldsymbol{\mu}_1^{(t)}] = [\mathbf{A}^{\text{old}}, \boldsymbol{\mu}_1^{\text{old}}] / s^{(t)}$ and $\mathbf{h}_\gamma^{(t)} = \mathbf{h}_\gamma^{\text{old}} \times s^{(t)}$.

As seen in Fig. 14(a-b), normalizing the HRF to constant values during Gibbs sampling may have a dramatic impact and disable the recovery of the HRF shape. In contrast, the proposed stochastic sampling step enables this recovery, up to a multiplicative constant which is inherent to the scale ambiguity. Still, for $\|\mathbf{h}_\gamma^n\| > 10^{-2}$ we do not observe any HRF estimation degradation with the deterministic approach, see Fig. 14(c-d). However, this optimal setting for the normalizing constant depends on the data so that we actually have to automatically adapt the HRF scale and rather consider the stochastic approach.

ACKNOWLEDGMENT

The authors thank the Région île de France and the ANR for financial support (grant number ANR-005-MMSA-0014-03, OPTIMED project). We are also grateful to Jérôme Idier and Jean-François Giovannelli for fruitful discussions and comments.




Fig. 14. **Impact of HRF normalization:** illustration on simulations (true HRF in red). (a), (b) and (c) show the impact of deterministic normalizations of the HRF estimate $\hat{\mathbf{h}}_\gamma^n$ using increasing normalization values versus a stochastic sampling of the scale parameter s associated to the HRF estimate $\hat{\mathbf{h}}_\gamma^s$. (d) show mean square errors (MSE) in terms of normalization value (blue line), versus the optimal MSE obtained with scale sampling (green line).

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. Makni, J. Idier, T. Vincent, B. Thirion, G. Dehaene-Lambertz, and P. Ciuciu, "A fully Bayesian approach to the parcel-based detection-estimation of brain activity in fMRI," *Neuroimage*, vol. 41, no. 3, pp. 941–969, July 2008. [Online]. Available: <http://dx.doi.org/10.1016/j.neuroimage.2008.02.017>
- [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] K. Friston, "Imaging neuroscience: Principles or maps?" *Proc. Natl. Acad. Sci. USA*, vol. 95, pp. 796–802, 1998.
- [5] R. Henson, C. Price, M. Rugg, R. Turner, and K. Friston, "Detecting latency differences in event-related BOLD responses: application to words versus nonwords and initial versus repeated face presentations," *Neuroimage*, vol. 15, no. 1, pp. 83–97, 2002.
- [6] K. Friston, "Statistical parametric mapping," in *Functional Neuroimaging: Technical Foundations*, R. Thatcher, M. Hallet, T. Zeffiro, E. John, and M. Huerta, Eds., 1994, pp. 79–93.
- [7] N. Lange, "Empirical and substantive models, the Bayesian paradigm, and meta-analysis in functional brain imaging," *Hum. Brain Mapp.*, vol. 5, pp. 259–263, 1997.
- [8] M. S. Cohen, "Parametric analysis of MRI data using linear systems methods," *Neuroimage*, vol. 6, pp. 93–103, 1997.
- [9] J. C. Rajapakse, F. Kruggel, J. M. Maisog, and D. Von Cramon, "Modeling hemodynamic response for analysis of functional MRI time-series," *Hum. Brain Mapp.*, vol. 6, pp. 283–300, 1998.
- [10] F. Kruggel and D. Y. Von Crammon, "Modeling the hemodynamic response in single-trial functional MRI experiments," *Magn. Reson. Med.*, vol. 42, pp. 787–797, 1999.
- [11] C. Genovese, "A Bayesian time-course model for functional magnetic resonance imaging data (with discussion)," *J. Amer. Statist. Assoc.*, vol. 95, pp. 691–719, 2000.
- [12] C. Gössl, D. P. Auer, and L. Fahrmeir, "Bayesian spatio-temporal modeling of the hemodynamic response function in BOLD fMRI," *Biometrics*, vol. 57, pp. 554–562, June 2001.
- [13] 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.
- [14] F. A. Nielsen, L. K. Hansen, P. Toft, C. Goutte, N. Lange, S. C. Stroher, N. Mørch, C. Svarer, R. Savoy, B. Rosen, E. Rostrup, and

- P. Born, "Comparison of two convolution models for fMRI time series," *Neuroimage*, vol. 5, p. S473, 1997.
- [15] C. Goutte, F. A. Nielsen, and L. K. Hansen, "Modeling the haemodynamic response in fMRI using smooth filters," *IEEE Trans. Med. Imag.*, vol. 19, no. 12, pp. 1188–1201, Dec. 2000.
 - [16] G. Marrelec, H. Benali, P. Ciuciu, M. Péligrini-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.
 - [17] P. Ciuciu, J.-B. Poline, G. Marrelec, J. Idier, C. Pallier, and H. Benali, "Unsupervised robust non-parametric estimation of the hemodynamic response function for any fMRI experiment," *IEEE Trans. Med. Imag.*, vol. 22, no. 10, pp. 1235–1251, Oct. 2003.
 - [18] G. Marrelec, P. Ciuciu, M. Péligrini-Issac, and H. Benali, "Estimation of the hemodynamic response function in event-related functional MRI: Bayesian networks as a framework for efficient Bayesian modeling and inference," *IEEE Trans. Med. Imag.*, vol. 23, no. 8, pp. 959–967, Aug. 2004.
 - [19] R. B. Buxton, E. C. Wong, and F. L. R., "Dynamics of blood flow and oxygenation changes during brain activation: the balloon model," *Magn. Reson. Med.*, vol. 39, pp. 855–864, June 1998.
 - [20] K. J. Friston, A. Mechelli, R. Turner, and C. J. Price, "Nonlinear responses in fMRI: the balloon model, Volterra kernels, and other hemodynamics," *Neuroimage*, vol. 12, pp. 466–477, June 2000.
 - [21] R. B. Buxton, K. U. g. D. J. Dubowitz, and T. T. Liu, "Modeling the hemodynamic response to brain activation," *Neuroimage*, vol. 23, Supplement 1, pp. S220–S233, 2004.
 - [22] J. Riera, J. Watanabe, I. Kazuki, M. Naoki, E. Aubert, T. Ozaki, and R. Kawashima, "A state-space model of the hemodynamic approach: nonlinear filtering of BOLD signal," *Neuroimage*, vol. 21, pp. 547–567, 2004.
 - [23] T. Deneux and O. Faugeras, "EEG-fMRI fusion of non-triggered data using kalman filtering," in *Proc. 3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro*, 6–9 April 2006, pp. 1068–1071.
 - [24] K. E. Stephan, N. Weiskopf, P. M. Drysdale, P. A. Robinson, and K. J. Friston, "Comparing hemodynamic models with dcm," *Neuroimage*, vol. 38, no. 3, pp. 387–401, Nov. 2007. [Online]. Available: <http://dx.doi.org/10.1016/j.neuroimage.2007.07.040>
 - [25] J. Riera, J. Bosch, O. Yamashita, R. Kawashima, N. Sadato, T. Okada, and T. Ozaki, "fMRI activation maps based on the NN-ARx model," *Neuroimage*, vol. 23, pp. 680–697, 2004.
 - [26] P. Baraldi, A. Manginelli, M. Maieron, D. Liberati, and A. Porro, "An ARX model-based approach to trial by trial identification of fMRI-BOLD responses," *Neuroimage*, vol. 37, pp. 189–201, 2007.
 - [27] L. A. Johnston, E. Duff, I. Mareels, and G. F. Egan, "Nonlinear estimation of the BOLD signal," *Magn. Reson. Med.*, vol. 40, no. 2, pp. 504–514, Apr. 2008.
 - [28] 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.
 - [29] G. Flandin, F. Kherif, X. Pennec, D. Rivière, N. Ayache, and J.-B. Poline, "A new representation of fMRI data using anatomo-functional constraints," in *Proc. 8th HBM*, Sendai, Japan, June 2002.
 - [30] D. Lashkari and P. Golland, "Exploratory fMRI analysis without spatial normalization," in *21st Proceedings of IPMI*, Williamsburg, VA., July 2009.
 - [31] W. Penny, Z. Ghahramani, and K. Friston, "Bilinear dynamical systems," *Philos Trans R Soc Lond B Biol Sci*, vol. 360, no. 1457, pp. 983–993, May 2005.
 - [32] S. Makni, C. Beckmann, S. Smith, and M. Woolrich, "Bayesian deconvolution fMRI data using bilinear dynamical systems," *Neuroimage*, vol. 42, no. 4, pp. 1381–1396, Oct. 2008.
 - [33] C. Beckmann and S. Smith, "Probabilistic independent component analysis for functional magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 23, no. 2, pp. 137–152, 2004.
 - [34] W. D. Penny, N. Trujillo-Barreto, and K. J. Friston, "Bayesian fMRI time series analysis with spatial priors," *Neuroimage*, vol. 23, no. 2, pp. 350–362, 2005.
 - [35] G. Flandin and W. D. Penny, "Bayesian fMRI data analysis with sparse spatial basis function priors," *Neuroimage*, vol. 34, no. 3, pp. 1108–1125, Feb. 2007.
 - [36] L. M. Harrison, W. Penny, J. Daunizeau, and K. J. Friston, "Diffusion-based spatial priors for functional magnetic resonance images," *Neuroimage*, vol. 41, no. 2, pp. 408–423, Jun. 2008. [Online]. Available: <http://dx.doi.org/10.1016/j.neuroimage.2008.02.005>
 - [37] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 3, pp. 367–383, Mar. 1992.
 - [38] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Processing*, vol. 6, no. 2, pp. 298–311, Feb. 1997.
 - [39] F. Forbes and N. Peyrard, "Hidden Markov Random Field model selection criteria based on mean field-like approximations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 9, pp. 1089–1101, Sep. 2003.
 - [40] M. Beal, "Variational algorithms for approximate Bayesian inference," PhD thesis, University College of London, London, United Kingdom, May 2003.
 - [41] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1222–1239, Nov. 2001.
 - [42] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 9, pp. 1124–1137, Sep. 2004.
 - [43] F. C. Jeng and J. W. Woods, "Compound Gauss-Markov random fields for image estimation," *IEEE Trans. Signal Processing*, vol. 39, no. 3, pp. 683–697, Mar. 1991.
 - [44] 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.
 - [45] P. J. Green and S. Richardson, "Hidden Markov models and disease mapping," *J. Amer. Statist. Assoc.*, vol. 97, no. 460, pp. 1–16, Dec. 2002.
 - [46] M. Smith, B. Pütz, D. Auer, and L. Fahrmeir, "Assessing brain activity through spatial Bayesian variable selection," *Neuroimage*, vol. 20, pp. 802–815, 2003.
 - [47] S. Fernández and P. J. Green, "Modelling spatially correlated data via mixtures: a Bayesian approach," *J. R. Statist. Soc. B*, vol. 64, no. 4, pp. 805–826, 2002.
 - [48] M. Woolrich, T. Behrens, C. 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.
 - [49] A. Barbu and S.-C. Zhu, "Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1239–1253, Aug. 2005.
 - [50] M. Woolrich and T. Behrens, "Variational Bayes inference of spatial mixture models for segmentation," *IEEE Trans. Med. Imag.*, vol. 25, no. 10, pp. 1380–1391, Oct. 2006.
 - [51] M. Svensen, F. Kruggel, and D. von Crammon, "Probabilistic modeling of single-trial fMRI data," *IEEE Trans. Med. Imag.*, vol. 19, pp. 19–35, Jan. 2000.
 - [52] W. Ou and P. Golland, "From spatial regularization to anatomical priors in fMRI analysis," in *IPMI, Glenwood Springs, Colorado*, July 2005.
 - [53] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, pp. 721–741, Nov. 1984.
 - [54] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," *IEEE Trans. Image Processing*, vol. 4, no. 7, pp. 932–946, July 1995.
 - [55] 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*, vol. I, Honolulu, Hawaii, Apr. 2007, pp. 325–328.
 - [56] N. Vaever Hartvig and J. Jensen, "Spatial mixture modeling of fMRI data," *Hum. Brain Mapp.*, vol. 11, no. 4, pp. 233–248, 2000.
 - [57] 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.
 - [58] S. Dehaene, L. Naccache, L. Cohen, D. Le Bihan, J.-F. Mangin, J.-B. Poline, and D. Rivière, "Cerebral mechanisms of word masking and unconscious repetition priming," *Nat. Neurosci.*, vol. 4, no. 7, pp. 752–758, 2001.
 - [59] P. Ciuciu, S. Sockeel, T. Vincent, and J. Idier, "Modelling the neurovascular habituation effect on fMRI time series," in *34th Proc. IEEE ICASSP*, Taipei, Taiwan, Apr. 2009, pp. 433–436.
 - [60] X. Meng and W. Wong, "Simulating ratios of normalizing constants via a simple identity: a theoretical exploration," *Statistica Sinica*, vol. 6, pp. 831–860, 1996.
 - [61] D. M. Higdon, J. E. Bowsher, V. E. Johnson, T. G. Turkington, D. R. Gilland, and R. J. Jaszczak, "Fully Bayesian estimation of Gibbs

- hyperparameters for emission computed tomography data," *IEEE Trans. Med. Imag.*, vol. 16, no. 5, pp. 516–526, Oct. 1997.
- [62] A. Gelman and X.-L. Meng, "Simulating normalizing constants: from importance sampling to bridge sampling to path sampling," *Statistical Science*, vol. 13, pp. 163–185, 1998.
- [63] L. Risser, J. Idier, and P. Ciuciu, "Bilinear extrapolation scheme for fast estimation of 3D Ising field partition function. Application to fMRI time course analysis," in *16th Proc. IEEE ICIP*, Cairo, Egypt, Nov. 2009.
- [64] L. Risser, T. Vincent, P. Ciuciu, and J. Idier, "Robust extrapolation scheme for fast estimation of 3D Ising field partition functions. application to within-subject fMRI data analysis," in *12th Proc. MICCAI'09*, ser. LNCS 5761, G.-Z. Yang, Ed. London, UK: Springer Verlag Berlin Heidelberg, Sep. 2009, pp. 975–983.
- [65] A. Shmuel, M. Augath, A. Oeltermann, and N. K. Logothetis, "Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1," *Nat Neurosci*, vol. 9, no. 4, pp. 569–577, Apr 2006.
- [66] L. Risser, T. Vincent, and P. Ciuciu, "Schéma d'extrapolation de fonctions de partition de champs de pots. application à l'analyse d'images en IRMf," in *Actes du 22^e colloque GRETSI*, Dijon, France, Sep. 2009.
- [67] J. Liu, *Monte Carlo strategies in scientific computing*, ser. Springer series in Statistics. New-York: Springer-Verlag, 2001.
- [68] B. Thyreau, B. Thirion, G. Flandin, and J.-B. Poline, "Anatomo-functional description of the brain: a probabilistic approach," in *Proc. 31th Proc. IEEE ICASSP*, vol. V, Toulouse, France, May 2006, pp. 1109–1112.
- [69] A. Tucholka, B. Thirion, M. Perrot, P. Pinel, J.-F. Mangin, and J.-B. Poline, "Probabilistic anatomo-functional parcellation of the cortex: how many regions?" in *11th Proc. MICCAI, LNCS Springer Verlag*, New-York, USA, 2008.
- [70] K. Worsley, C. Liao, J. Aston, V. Petre, G. Duncan, F. Morales, and A. Evans, "A general statistical analysis for fMRI data," *Neuroimage*, vol. 15, no. 1, pp. 1–15, Jan. 2002.
- [71] M. Woolrich, M. Jenkinson, J. M. Brady, and S. Smith, "Constrained linear basis set for HRF modelling using variational Bayes," *Neuroimage*, vol. 21, no. 4, pp. 1748–1761, 2004.
- [72] S. M. Kay, *Modern Spectral Estimation*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [73] A. Trillon, J. Idier, and P. Peureux, "Unsupervised Bayesian 3D reconstruction for non-destructive evaluation using gammagraphy," in *EUSIPCO*, Lausanne, Suisse, Aug. 2008.
- [74] C. Geyer and E. Thompson, "Constrained monte carlo maximum likelihood for dependent data," *J. Roy. Statist. Soc.*, vol. 54, pp. 657–699, 1992.
- [75] R. H. Swendsen and J. S. Wang, "Nonuniversal critical dynamics in Monte Carlo simulations," *Physical Review Letters*, vol. 58, pp. 86–88, 1987.
- [76] D. Smith and L. Fahrmeir, "Spatial Bayesian variable selection with application to Functional Magnetic Resonance Imaging," *J. Amer. Statist. Assoc.*, vol. 102, no. 478, pp. 417–431, 2007.
- [77] R. Edwards and A. Sokal, "Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm," *Phys. Rev. D*, vol. 38, no. 6, pp. 2009–2012, 1988.
- [78] A. Barbu and S. Zhu, "Generalizing swendsen-wang for image analysis," *Journal of Computational & Graphical Statistics*, vol. 16, no. 4, pp. 877–900, 2007.
- [79] A.-L. Fouque, P. Ciuciu, and L. Risser, "Multivariate spatial Gaussian mixture modeling for statistical clustering of hemodynamic parameters in functional MRI," in *34th Proc. IEEE ICASSP*, Taipei, Taiwan, Apr. 2009, pp. 445–448.
- [80] M. Jerrum and A. Sinclair, "Polynomial-time approximation algorithms for the Ising model," *SIAM J. Comput.*, vol. 22, pp. 1087–1116, 1993.
- [81] O. Gruber, P. Indefrey, H. Steinmetz, and A. Kleinschmidt, "Dissociating Neural Correlates of Cognitive Components in Mental Calculation," *Cereb. Cortex*, vol. 11, no. 4, pp. 350–359, 2001. [Online]. Available: <http://cercor.oxfordjournals.org/cgi/content/abstract/11/4/350>
- [82] T. Vincent, P. Ciuciu, and B. Thirion, "Sensitivity analysis of parcellation in the joint detection-estimation of brain activity in fMRI," in *5th Proc. IEEE ISBI*, Paris, France, May 2008, pp. 568–571.
- [83] P. Ciuciu, T. Vincent, A.-L. Fouque, and A. Roche, "Improved fMRI group studies based on spatially varying non-parametric BOLD signal modeling," in *5th Proc. IEEE ISBI*, Paris, France, May 2008, pp. 1263–1266.
- [84] H. Snoussi and J. Idier, "Bayesian blind separation of generalized hyperbolic processes in noisy and underdeterminate mixtures," *IEEE Trans. Signal Processing*, vol. 54, no. 9, pp. 3257–3269, Sep. 2006.
- [85] T. Veit, J. Idier, and S. Moussaoui, "Rééchantillonnage de l'échelle dans les algorithmes MCMC pour les problèmes inverses bilinéaires," *Traitement du Signal*, vol. 25, no. 4, pp. 329–343, 2008.