

Estimation of the Hemodynamic Response Function in event-related functional MRI: directed acyclic graphs for a general Bayesian inference framework.

Guillaume Marrelec, Philippe Ciuciu, Mélanie Péligrini-Issac, Habib Benali

► To cite this version:

Guillaume Marrelec, Philippe Ciuciu, Mélanie Péligrini-Issac, Habib Benali. Estimation of the Hemodynamic Response Function in event-related functional MRI: directed acyclic graphs for a general Bayesian inference framework.. Information processing in medical imaging: proceedings of the .. conference., Springer-Verlag, 2003, 18, pp.635-46. cea-00333749

HAL Id: cea-00333749

<https://hal-cea.archives-ouvertes.fr/cea-00333749>

Submitted on 24 Oct 2008

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.

Estimation of the Hemodynamic Response Function in Event-Related Functional MRI: Directed Acyclic Graphs for a General Bayesian Inference Framework

Guillaume Marrelec^{1,4}, Mélanie Péligrini-Issac^{2,4}, Philippe Ciuciu^{3,4}, and Habib Benali^{1,4}

¹ INSERM U494

{marrelec, benali, u494}@imed.jussieu.fr

² INSERM U483

pelegrin@imed.jussieu.fr

³ CEA, SHFJ

ciuciu@cea.shfj.fr

⁴ IFR 49

Abstract. Functional MRI (fMRI) is a non-invasive technique allowing for the evolution of brain processes to be dynamically followed in various cognitive tasks. In BOLD fMRI, what is actually measured is only indirectly related to neuronal activity through a process that is still under investigation. A convenient way to analyze such data consists of considering the whole brain as a system characterized by a transfer response function, called the Hemodynamic Response Function (HRF). In this paper we develop the general linear model and show how it can be translated in terms of graphical models. This conceptual step has the advantage of making clear all structural and functional assumptions entailed by the model. Moreover, once embedded in a DAG, the preparatory step for Gibbs sampling –variable block definition for sequential updating and derivation of the corresponding conditional distributions– is readily performed from the model specification. From there, probabilistic numerical estimation of all parameters is straightforward, leading to robust inference.

1 Introduction

Discovered in the early 90s, functional MRI (fMRI) is a recent, non-invasive technique allowing for the evolution of brain processes to be dynamically followed in various cognitive and behavioral tasks. In the most common fMRI technique, based on the so-called Blood Oxygen Level Dependent (BOLD) contrast, what is actually measured is only indirectly related to neuronal activity through a process that is not well understood yet [1, 2]. For this reason, a convenient way to analyze BOLD fMRI data consists of considering the whole brain as a “black box” system characterized by its transfer response function, also called Hemodynamic Response Function (HRF) [3]. The HRF is the theoretical signal that

BOLD fMRI would measure in response to a single, very short stimulus of unit intensity. Under such hypotheses, the data are related to the HRFs through the so-called General Linear Model (GLM). This model assumes stationarity and linearity of the underlying physiological process. Such assumptions are a good approximation of the properties of the real system as long as the inter-stimulus interval does not decrease beyond about two seconds [4, 5]. Under such hypotheses,

Estimation of the HRF is of the greatest interest when analysing of fMRI data, since it can give a deep insight into the underlying dynamics of brain activation and the relationships between activated areas. HRFs are increasingly suspected to vary from region to region, from task to task, and from subject to subject [6–8].

Precise and robust estimation of the HRF still belongs to ongoing research, since the problem is badly conditioned, and various methods have been devised so far to estimate the HRF. Parametric methods include Gaussian [9] or spline-like fitting [10], spatial regularization [11], or consideration of physiological models [12]. However, assuming the shape of the hemodynamic response to be known *a priori* and invariant throughout the brain is a very strong constraint, since it fluctuates greatly. Non-parametric methods have hence been developed in an attempt to infer the HRF at each time sample. Methods include: averaging over regions [13], selective averaging [4], introduction of non-diagonal models for the temporal covariance of the noise [14], or temporal regularization [15].

In [16], we proposed a Bayesian, non-parametric estimation of the HRF for event-related designs. Basic yet relevant physiological information was introduced to temporally constrain the problem and calculate robust estimators of the parameters of interest. In [17], we quantified the performances of the method and showed great improvement compared to Maximum-Likelihood methods. Robustness with regard to the noise structure and level were also proven. [18] and [19] generalized the method to asynchronous event-related designs, taking different trial types into account, and making it possible to incorporate several fMRI sessions to further enhance the estimation. For calculation reasons, all variants proposed so far have the drawback of not integrating the hyperparameter uncertainty. Moreover, the general model cannot treat all parameters probabilistically (e.g., drift parameters).

In this paper, we propose to consider the GLM under a new light. We still place ourselves in a Bayesian framework, to permit integration of information originating from models, data, and prior information and lead efficient inference on the parameters of interest. A general model is set, as in [19], that allow to analyze most event-related fMRI data. In a common Bayesian approach, we would then calculate the joint posterior distribution of all parameters, which will be the pivotal quantity for all further inference. On the contrary, we advocate that, in our case, calculation of the posterior probability density function (pdf) is unnecessary. We propose a novel approach that focuses on modeling and that, once the model has been properly set, makes it possible to directly lead probabilistic inference about all parameters. More precisely, we utilize graph theory

to conveniently deal with the model. Graphs give a very simple and efficient representation of the model, however complex it may be. In this framework, we translate the model into a directed acyclic graph (DAG) and by functional relationships between the DAG variables. Using the Markov property of DAGs, solving the problem in a general manner becomes straightforward: application of Gibbs sampling scheme provides us with numerical approximation of the joint posterior probability density function.

In the first part of this paper, we develop the general framework of Bayesian estimation of the HRF, presenting an extended version of the General Linear Model (GLM) to enable analysis of most fMRI experiments. In the next section, the GLM is translated in terms of graphical models, and it is shown how inference can readily be performed from there. We quickly present simulations and finally apply our resolution model to real data, showing variability in HRFs.

2 HRF estimation in fMRI data analysis

2.1 Notations

In the following, x denotes a real number, \mathbf{x} a vector, and \mathbf{X} a matrix. (x_i) stands for $(x_i)_{1 \leq i \leq I}$. “ t ” is the regular matrix transposition. \mathbf{I}_N stands for the N -by- N identity matrix. “ \propto ” relates two expressions that are proportional. For two variables x and y , “ $x|y$ ” stands for “ x given y ”, and $p(x)$ for the probability of x . $\mathcal{N}(\mathbf{m}, \mathbf{V}; \mathbf{x})$ is the Gaussian density function with mean \mathbf{m} and covariance matrix \mathbf{V} calculated at point \mathbf{x} . $\text{Inv-}\chi^2(d, s^2; u)$ is the scaled inverse-chi-square density function with d degrees of freedom and scale parameter s^2 evaluated at point u .

2.2 General linear model (H)

Data. Let an fMRI experiment be composed of S sessions, each session involving I different stimulus types. Define $\mathbf{x}_{s,i} = (x_{s,i,t})_{t_{s,0} \leq t \leq t_{s,N}}$ as the time series of the i th stimulus for session s and $\mathbf{y}_s = (y_{s,t_{s,n}})_{1 \leq n \leq N}$ the corresponding BOLD fMRI time course of a voxel at (not-necessarily uniformly sampled) times $(t_{s,n})$. A discrete linear convolution model is assumed to hold between the stimulus vectors and the data:

$$\mathbf{y}_{s,t_{s,n}} = \sum_{i=1}^I \sum_{k=0}^{K_i} h_{i,k \Delta t} x_{s,i,t_{s,n}-k \Delta t} + \sum_{m=1}^{M_s} \lambda_{s,m} d_{m,t_{s,n}} + \ell_{t_n} \quad n = n_s + 1, \dots, N_s,$$

where n_s is the largest integer so, that $t_{s,n} - K_i \Delta t < t_{s,1}$ for all i . The $(K_i + 1)$ -dimensional vectors $\mathbf{h}_i = (h_{i,k \Delta t})^t$ represent the unknown HRFs to be estimated, sampled every Δt . They are assumed to be constant across sessions. $L_s = N_s - n_s$ is the actual amount of data used in the calculation for each session. $\mathbf{X}_{s,i} = (x_{s,i,t_{s,n}-k \Delta t})$ are the regular L_s -by- $(K + 1)$ design matrices, consisting of the lagged stimulus covariates. The L_s -by- M_s matrices $\mathbf{D}_s = (d_{m,t_{s,n}})$ are the values

at times $(t_{s,n})$ of a basis of M_s functions that takes a potential drift and any other nuisance effect into account, and the $\boldsymbol{\lambda}_s = (\lambda_{s,m})^t$ are the corresponding coefficients. For the sake of simplicity, the bases are assumed to be orthonormal, *i.e.* $\frac{1}{L_s} \mathbf{D}_s^t \mathbf{D}_s = \mathbf{I}_{L_s}$. Vectors $\mathbf{e}_s = (e_{s,t_{s,n}})^t$ account for noise and are supposed to consist of independent and identically distributed Gaussian variables of unknown variances (σ_s^2), assumed to be independent from the HRF. In matrix form, (H) boils down to

$$\mathbf{y}_s = \sum_{i=1}^I \mathbf{X}_{s,i} \mathbf{h}_i + \mathbf{D}_s \boldsymbol{\lambda}_s + \mathbf{e}_s, \quad s = 1, \dots, S,$$

also called General Linear Model (GLM). In this model, the likelihood of the data reads

$$p((\mathbf{y}_s)|H, (\mathbf{h}_i), (\sigma_s^2), (\boldsymbol{\lambda}_s)) = \prod_{s=1}^S p(\mathbf{y}_s|H, (\mathbf{h}_i), \sigma_s^2, \boldsymbol{\lambda}_s),$$

with each term in the product reading

$$p(\mathbf{y}_s|H, (\mathbf{h}_i), \sigma_s^2, \boldsymbol{\lambda}_s) = \mathcal{N} \left(\sum_{i=1}^I \mathbf{X}_{s,i} \mathbf{h}_i + \mathbf{D}_s \boldsymbol{\lambda}_s, \sigma_s^2 \mathbf{I}_{L_s}; \mathbf{y}_s \right).$$

HRFs and hyperparameters. The GLM being ill-conditioned, prior information must be incorporated in order to constrain the problem. Since the underlying physiological process of BOLD fMRI is as of yet only partially understood, we set the following soft constraints [16, 19]:

- (P1) the HRFs start and end at 0. This amounts to setting the first and last samples of each HRFs to 0, so that only $K_i - 1$ parameters (instead of $K_i + 1$) are now unknown.
- (P2) the HRFs are smooth. Quantification is achieved by setting Gaussian priors for the norm of the second derivative of the HRFs, whose variances are adjusted by hyperparameters ϵ_i 's:

$$p(\mathbf{h}_i|H, \epsilon_i^2) = \mathcal{N}(\mathbf{0}, \epsilon_i^2 \mathbf{R}_i^{-1}; \mathbf{h}_i) \quad i = 1, \dots, I$$

where \mathbf{R}_i is the following $(K_i - 1)$ -by- $(K_i - 1)$ matrix:

$$\mathbf{R} = \frac{1}{(\Delta t)^4} \begin{pmatrix} 5 & -4 & 1 & 0 & \dots & & 0 \\ -4 & 6 & -4 & 1 & 0 & & \\ 1 & -4 & 6 & -4 & 1 & 0 & \\ 0 & 1 & -4 & 6 & -4 & 1 & 0 & \vdots \\ & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ \vdots & & 0 & 1 & -4 & 6 & -4 & 1 & 0 \\ & & & 0 & 1 & -4 & 6 & -4 & 1 \\ & & & & 0 & 1 & -4 & 6 & -4 \\ 0 & & \dots & & 0 & 1 & -4 & 5 \end{pmatrix}.$$

(P3) No prior dependence is assumed between HRFs, so that

$$p((\mathbf{h}_i), (\epsilon_i^2)|H) = \prod_{i=1}^I p(\mathbf{h}_i|H, \epsilon_i^2) \cdot p(\epsilon_i^2|H)$$

The ϵ_i^2 's are assumed to be *a priori* i.i.d. with common pdf set to a scaled inverse- χ^2 with n_ϵ degrees of freedom and scale parameter s_ϵ^2 .

Drifts and noise variances. Unlike the HRFs, noise variances and drift parameters may vary across sessions. The σ_s^2 's are assumed to be i.i.d. with common pdf a scaled inverse- χ^2 distribution with n_σ degrees of freedom and scale parameter s_σ^2 . The $\boldsymbol{\lambda}_s$ are assumed to be i.i.d. with common pdf a Gaussian pdf of mean \mathbf{m}_0 and covariance matrix \mathbf{V}_0 .

Joint posterior distribution. Considering the model so constructed and assuming no further prior dependence between parameters, formal application of the chain rule yields

$$\begin{aligned} p((\mathbf{y}_s), (\boldsymbol{\lambda}_s), (\mathbf{h}_i), (\sigma_s^2), (\epsilon_i^2)|H) &= \prod_{s=1}^S p(\mathbf{y}_s|H, (\mathbf{h}_i), \sigma_s^2) \cdot p(\boldsymbol{\lambda}_s|H) \cdot p(\sigma_s^2|H) \\ &\times \prod_{i=1}^I p(\mathbf{h}_i|H, \epsilon_i^2) \cdot p(\epsilon_i^2|H) \end{aligned} \quad (1)$$

Replacing all distributions by their functional forms, this joint posterior pdf could be calculated in closed form, as is indeed done in most research papers applying Bayesian analysis. From there, conditional pdfs could be derived if Gibbs sampling is required. We propose to avoid this step and directly proceed to inference. In order to do so, we beforehand embed our model in a framework that allows for convenient representation and handling: directed acyclic graphs.

3 Graphical modeling

3.1 Directed acyclic graphs (DAGs)

A graph G is a mathematical object that consists of two sets, a set of vertices, K , and a set of edges, E , consisting of pairs of elements taken from K . There is a directed edge or arrow between vertices a and b in K if the set E contains the ordered pair (a, b) ; vertex a is a parent of vertex b , and vertex b is a child of vertex a . An oriented graph is a graph whose edges are all oriented.

The major feature of DAGs is that any probability density that is compatible with the graph must factorize according to the so-called *factorization property*. More precisely, let G be a DAG structure, and $(\mathbf{x}_n)_{1 \leq n \leq N}$ a set of variables. If

the variable set is compatible with G , then the pdf $p(\mathbf{x})$ that can be decomposed as:

$$p(\mathbf{x}) = \prod_{n=1}^N p(\mathbf{x}_n | \text{pa}(n)), \quad (2)$$

where $\text{pa}(n)$ is the set of parents of vertex \mathbf{x}_n . Defining a graph amounts to 1/ defining relevant variables (i.e., nodes) \mathbf{x}_n , 2/ defining structural relationships (i.e., edges) $\mathbf{x}_n \rightarrow \mathbf{x}_m$, and 3/ defining functional relationships $p(\mathbf{x}_n | \text{pa}(n))$. Pearl [20] showed a property that proves to be very efficient for numerical sampling, namely that nothing more is required to calculate the conditional probability of any DAG node: the probability distribution of any variable \mathbf{x}_i in the network, conditioned on the state of all other variables, is given by the product

$$p(\mathbf{x}_n | \text{rest}) \propto p(\mathbf{x}_n | \text{pa}(n)) \cdot \prod_{j \in \text{ch}(n)} p(\mathbf{x}_j | \text{pa}(j)), \quad (3)$$

where $\text{ch}(n)$ stands for the children nodes of \mathbf{x}_n . In other words, the conditional probabilities can be derived from local quantities that are part of the model specification.

3.2 DAG model for HRF estimation

The GLM can easily be expressed in terms of DAG. Indeed, consider the DAG proposed in Figure 1. Irrespective of the functional relationships between nodes, Relation (2) states that the joint posterior pdf for all DAG variables decomposes exactly like the posterior pdf in Equation (1). Identifying all functional relationships of the DAG to their counterparts for model (H) then makes the DAG a perfect representation of the GLM. However complicated (H) may be, it is still much simpler to conceptualize in graph form than as it was presented before. Whereas determination of structural relationships between two given variables in model (H) remains a tough problem to tackle, the corresponding DAG clearly and unambiguously represent all possible independence relationships, that can be read off the graph using Markov properties.

3.3 Numerical inference

To obtain a numerical approximation of the joint posterior pdf, we apply Gibbs sampling scheme. Gibbs sampling consists of starting with a seed vector and sequentially modifying one vector component at a time by sampling according to the conditional pdf of that component given all other components. Samples are composed of the set of all vectors whose components have been updated an equal amount of times.

An issue with Gibbs sampling is to break the vector of all parameters into components whose conditional sampling can easily be performed. Another one is derivation of the conditional pdfs corresponding to the chosen clustering. In




Fig. 1. DAG corresponding to the General Linear Model. The gray nodes represent available information.

our case, both questions are answered at once, thanks to the previous step of graph modeling. As a matter of fact, it first allows us to decompose the parameter vector onto its $2I + 2S$ canonical components: I ϵ_i 's and \mathbf{h}_i 's, S σ_s^2 's and $\boldsymbol{\lambda}_s$'s (all $\mathbf{y} - s$'s begin given, no sampling need to be done on these variables). The updating steps are performed on these variables, and we therefore need access to the following conditional pdfs: $p(\mathbf{h}_i|H, \text{rest})$, $p(\epsilon_i^2|H, \text{rest})$, $p(\sigma_s^2|H, \text{rest})$, and $p(\boldsymbol{\lambda}_s|H, \text{rest})$. But, according to Pearl's theorem, knowledge of the functional relationships is sufficient to infer these conditional pdfs. Application of Equation (3) yields

$$\begin{aligned}
 p(\epsilon_i^2|H, \text{rest}) &\propto p(\epsilon_i^2|H) \cdot p(\mathbf{h}_i|H, \epsilon_i^2) = \text{Inv-}\chi^2(\mu_i, \tau_i^2; \epsilon_i) \\
 p(\mathbf{h}_i|H, \text{rest}) &\propto p(\mathbf{h}_i|H, \epsilon_i^2) \cdot \prod_{s=1}^S p(\mathbf{y}_s|H, (\mathbf{h}_i), \sigma_s^2, \boldsymbol{\lambda}_s) = \mathcal{N}(\boldsymbol{\delta}_i, \boldsymbol{\Delta}_i; \mathbf{h}_i) \\
 p(\sigma_s^2|H, \text{rest}) &\propto p(\sigma_s^2|H) \cdot p(\mathbf{y}_s|H, (\mathbf{h}_i), \sigma_s^2, \boldsymbol{\lambda}_s) = \text{Inv-}\chi^2(\nu_s, \omega_s^2; \sigma_s^2) \\
 p(\boldsymbol{\lambda}_s|H, \text{rest}) &\propto p(\boldsymbol{\lambda}_s|H) \cdot p(\mathbf{y}_s|H, (\mathbf{h}_i), \sigma_s^2, \boldsymbol{\lambda}_s) = \mathcal{N}(\boldsymbol{\gamma}_s, \boldsymbol{\Gamma}_s; \boldsymbol{\lambda}_s),
 \end{aligned}$$

with

$$\begin{aligned}
\mu_i &= (K - 1) + n_\epsilon \\
\tau_i^2 &= \frac{n_\epsilon s_\epsilon^2 + \mathbf{h}_i^t \mathbf{R} \mathbf{h}_i}{(K - 1) + n_\epsilon} \\
\nu_s &= L_s + n_\sigma \\
\omega_s^2 &= \frac{n_\sigma s_\sigma^2 + \|\mathbf{y}_s - \sum_i \mathbf{X}_{s,i} \mathbf{h}_i - \mathbf{D}_s \boldsymbol{\lambda}_s\|^2}{L_s + n_\sigma} \\
\boldsymbol{\delta}_i &= \left(\frac{1}{\epsilon_i^2} \mathbf{R} + \sum_s \frac{1}{\sigma_s^2} \mathbf{X}_{s,i}^t \mathbf{X}_{s,i} \right)^{-1} \left(\sum_s \frac{1}{\sigma_s^2} \mathbf{X}_{s,i}^t \left(\mathbf{y}_s - \sum_{j \neq i} \mathbf{X}_{s,j} \mathbf{h}_j - \mathbf{D}_s \boldsymbol{\lambda}_s \right) \right) \\
\boldsymbol{\Delta}_i &= \left(\frac{1}{\epsilon_i^2} \mathbf{R} + \sum_s \frac{1}{\sigma_s^2} \mathbf{X}_{s,i}^t \mathbf{X}_{s,i} \right)^{-1} \\
\boldsymbol{\gamma}_s &= \left(\mathbf{V}_0^{-1} + \frac{1}{\sigma_s^2} \mathbf{D}_s^t \mathbf{D}_s \right)^{-1} \left(\mathbf{V}_0^{-1} \mathbf{m}_0 + \frac{1}{\sigma_s^2} \mathbf{D}_s^t \left(\mathbf{y}_s - \sum_i \mathbf{X}_{s,i} \mathbf{h}_i \right) \right) \\
\boldsymbol{\Gamma}_s &= \left(\mathbf{V}_0^{-1} + \frac{1}{\sigma_s^2} \mathbf{D}_s^t \mathbf{D}_s \right)^{-1}.
\end{aligned}$$

The sampling can then be performed. We sequentially update the ϵ_i 's, then the \mathbf{h}_i 's, the σ_s^2 's and finally the $\boldsymbol{\lambda}$'s. We are admittedly mostly interested in the HRFs, but knowledge of the values taken by the other parameters are relevant as well for our analysis and a better understanding of brain processing. Gibbs sampling gives us access to estimates for all parameters or any quantity of interest related to them.

4 Simulations

We simulated data with two HRFs, as depicted in Figure 2, and one session of $N = 100$ time samples. A quadratic drift and Gaussian white noise were also added. The noise variance was set to $\sigma^2 = 50$. This setting corresponds to a SNR⁵ of 13. For the analysis, we set $K = 20$, $M = 3$ with a quadratic drift, $\mathbf{m} = (1, 000 \ 0 \ 0)$, \mathbf{V}_0 diagonal, with $\text{diag}(\mathbf{V}_0) = (100^2 \ 1 \ 1)$, $n_\epsilon = n_\sigma = 1$, $s_\epsilon^2 = 1$, $s_\sigma^2 = 10^2$. Our unoptimized Matlab program took about 20 s on a Ultra Spark workstation to run 1,000 updates. We kept the 500 last samples to lead inference. The noise variance was estimated at $\sigma^2 \approx 53.5 \pm 10.6$, and the hyperparameters $\epsilon_1^2 \approx 0.84 \pm 0.64$ and $\epsilon_2^2 \approx 0.42 \pm 0.36$. As shown on Figure 2, HRF estimates are very accurate.

⁵ defined as $\text{SNR} = 20 \log_{10}(\sum_{s,i} \|\mathbf{X}_{s,i} \mathbf{h}_i\| / \sqrt{L_s \sigma_s^2})$




Fig. 2. Simulations. Estimated (dashed line) and true (solid line) HRFs. Top: Maximum-Likelihood estimate; bottom: estimate with temporal prior.

5 Real data

6 Discussion

Our approach made it possible to associate the well-known general linear model for HRF estimation in fMRI data analysis to a directed acyclic graph. This had the first advantage of making clear all modeling hypotheses. Moreover, in a Bayesian framework, the complex, yet central, step of calculating the joint posterior pdf was avoided. Instead, the graph provided us with a very convenient tool to first break down the set of all variables into coherent subsets, namely its nodes. Using the Markov properties, it is straightforward to derive all conditional pdfs that are required for Gibbs sampling as products of conditional pdfs that have been specified with the modeling. Fully probabilistic numerical inference is then straightforward at a reasonable time cost. For an increased effectiveness, the sampling procedure could be optimized. Convergence could be monitored through comparison of within- and between-variances of parallel chains as proposed in [21].

In this framework, increasing the model complexity can easily be performed. As a matter of fact, the proposed model can already be seen as an improvement of the simpler one-HRF, one-session DAG structure. Now, consideration of local spatial information, as in [11], could be achieved by gathering of all independent

voxel graphical models and addition of relationships between $h_{v,i}$'s. As more and more information is incorporated into the model, the corresponding graph will become more and more complex. However, tools have been developed to deal with them. A general procedure is proposed in [22] to sample from pdfs that have the structure implied by Equation (3) using rejection sampling. For huge graphs, [23] proposed an efficient variant of Gibbs sampling. Last, but not least, parallel processing of Gibbs sampling can be implemented. To avoid the problem of simultaneous updating of neighboring variables, one has to apply the so-called "edge reversal" control policy: Initially, the links of the graph are assigned arbitrary acyclic orientation of arrows (independently of the model graph). Each processor inspects the orientation of the arrows on its incident links and waits until all arrows point inward, in which case it becomes activated. When it completes the computation, it reverses the direction of all its incident arrows. This strategy has very interesting features, as detailed in [20].

We finally believe that this novel approach has a much broader application range than just fMRI data analysis. Indeed, we are confident in the fact that any Bayesian model can be embedded in a graphical framework, allowing for efficient and automated inference.

7 Conclusion

In this paper, we proposed a novel Bayesian inference framework for HRF estimation in fMRI data analysis, based on translating the existing Bayesian model into a DAG to combine the features of graphical modeling and Bayesian analysis. This approach makes extensive use of directed acyclic graphs to 1/ represent the model in a compact, yet efficient way, and 2/ lead probabilistic inference through Gibbs sampling. This technique takes advantage of Markov properties of DAGs. Models can easily be designed, and both structural (i.e., of independence) and functional relationships are clearly presented. Moreover, the local properties of structural relationships renders the model graph very simple to structurally or functionally modify, either because it does not correctly explain the phenomenon under interest, or because a model that is more complex is sought. Using Gibbs sampling on the DAG, fully probabilistic, numerical inference is straightforward. Further research includes integration of spatial constraint and of relevant physiological information.

Acknowledgements

The authors are grateful to Pr. J. Doyon for providing us with the data and to Carine Posé for her technical support. Guillaume Marrelec is supported by the Fondation pour la Recherche Médicale.

References

1. T.-Q. Li, T. N. Haefelin, B. Chan, A. Kastrup, T. Jonsson, G. H. Glover, and M. E. Moseley, "Assessment of hemodynamic response during focal neural activ-

- ity in human using bolus tracking, arterial spin labeling and BOLD techniques,” *NeuroImage*, vol. 12, pp. 442–451, 2000.
2. A. Aubert and R. Costalat, “A model of the coupling between brain electrical activity, metabolism, and hemodynamics: application to the interpretation of functional neuroimaging,” *NeuroImage*, vol. 17, no. 3, pp. 1162–1181, 2002.
 3. K. J. Friston, P. Jezzard, and R. Turner, “Analysis of functional MRI time-series,” *Human Brain Mapping*, vol. 1, pp. 153–171, 1994.
 4. A. M. Dale and R. L. Buckner, “Selective averaging of rapidly presented individual trials using fMRI,” *Human Brain Mapping*, vol. 5, pp. 329–340, 1997.
 5. R. L. Buckner, “Event-related fMRI and the hemodynamic response,” *Human Brain Mapping*, vol. 6, pp. 373–377, 1998.
 6. R. L. Buckner, W. Koutstaal, D. L. Schacter, A. D. Wagner, and B. R. Rosen, “Functional-anatomic study of episodic retrieval using fMRI (i),” *NeuroImage*, vol. 7, pp. 151–162, 1998.
 7. R. L. Buckner, W. Koutstaal, D. L. Schacter, A. M. Dale, M. Rotte, and B. R. Rosen, “Functional-anatomic study of episodic retrieval using fMRI (ii),” *NeuroImage*, vol. 7, pp. 163–175, 1998.
 8. G. K. Aguirre, E. Zarahn, and M. D’Esposito, “The variability of human BOLD hemodynamic responses,” *NeuroImage*, vol. 7, no. 4, p. S574, 1998.
 9. F. Kruggel and D. Y. von Cramon, “Modeling the hemodynamic response in single-trial functional MRI experiments,” *Magnetic Resonance in Medicine*, vol. 42, pp. 787–797, 1999.
 10. C. Gössl, L. Fahrmeir, and D. P. Auer, “Bayesian modeling of the hemodynamic response function in BOLD fMRI,” *NeuroImage*, vol. 14, pp. 140–148, 2001.
 11. C. Gössl, D. Auer, and L. Fahrmeir, “Bayesian spatiotemporal inference in functional magnetic resonance imaging,” *Biometrics*, vol. 57, no. 2, pp. 554–562, 2001.
 12. K. J. Friston, “Bayesian estimation of dynamical systems: an application to fMRI,” *NeuroImage*, vol. 16, pp. 513–530, 2002.
 13. J. Kershaw, S. Abe, K. Kashikura, X. Zhang, and I. Kanno, “A Bayesian approach to estimating the haemodynamic response function in event-related fMRI,” *NeuroImage*, vol. 11, no. 5, p. S474, 2000.
 14. M. A. Burock and A. M. Dale, “Estimation and detection of event-related fMRI signals with temporally correlated noise: a statistically efficient and unbiased approach,” *Human Brain Mapping*, vol. 11, pp. 249–260, 2000.
 15. C. Goutte, F. Årup Nielsen, and L. K. Hansen, “Modeling the haemodynamic response in fMRI using smooth FIR filters,” *IEEE Transactions in Medical Imaging*, vol. 19, no. 12, pp. 1188–1201, 2000.
 16. G. Marrelec, H. Benali, P. Ciuciu, and J.-B. Poline, “Bayesian estimation of the hemodynamic response function in functional MRI,” in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 21st International Workshop* (R. Fry, ed.), pp. 229–247, AIP, Melville, 2001.
 17. 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 fMRI using basic physiological information,” *Human Brain Mapping*, 2003. In press.
 18. P. Ciuciu, G. Marrelec, J.-B. Poline, J. Idier, and H. Benali, “Robust estimation of the hemodynamic response function in asynchronous multitasks multisessions event-related fMRI paradigms,” in *2002 IEEE International Symposium on Biomedical Imaging Proceedings*, pp. 847–850, IEEE, 2002.

19. 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 Transactions in Medical Imaging*, 2002. Submitted.
20. J. Pearl, "Evidential reasoning using stochastic simulation of causal models," *Artificial Intelligence*, vol. 32, pp. 245–257, 1987.
21. A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*. Texts in Statistical Science, Chapman & Hall, London, 1998.
22. G. Marrelec and H. Benali, "Automated rejection sampling from product of distributions," *Computational Statistics*, 2003. Accepted.
23. C. S. Jensen, A. Kong, and U. Kjærulff, "Blocking Gibbs sampling in very large probabilistic expert systems," *International Journal of Human Computer Studies. Special Issue on Real-World Applications of Uncertain Reasoning*, 1993.