Bayesian modelling of fMRI time series

Pedro A. d. F. R. Højen-Sørensen, Lars K. Hansen and **Carl Edward Rasmussen**
Department of Mathematical Modelling, Building 321
Technical University of Denmark
DK-2800 Lyngby, Denmark
phs,lkhansen,carl@imm.dtu.dk

Abstract

We present a Hidden Markov Model (HMM) for inferring the hidden psychological state (or neural activity) during single trail fMRI activation experiments with blocked task paradigms. Inference is based on Bayesian methodology, using a combination of analytical and a variety of Markov Chain Monte Carlo (MCMC) sampling techniques. The advantage of this method is that detection of short time learning effects between repeated trails is possible since inference is based only on single trail experiments.

1 Introduction

Functional magnetic resonance imaging (fMRI) is a non-invasive technique that enable indirect measures of neuronal activity in the working human brain. The most common fMRI technique is based on an image contrast induced by temporal shifts in the relative concentration of oxyhemoglobin and deoxyhemoglobin (BOLD contrast). Since neuronal activation leads to an increased blood flow, the so-called *hemodynamic response*, the measured fMRI signal reflects neuronal activity. Hence, when analyzing the BOLD signal there are two unknown factors to consider, the task dependent neuronal activation and the hemodynamic response. Bandettini et al. [1993] analyzed the correlation between a binary reference function (representing the stimulus/task sequence) and the BOLD signal. In the following we will also make reference to the binary representation of the task as the *paradigm*. Lange and Zeger [1997] discuss a parameterized hemodynamic response adapted by a least squares procedure. Multivariate strategies have been pursued in [Worsley et al. 1997, Hansen et al. 1999]. Several explorative strategies have been proposed for finding spatio-temporal activation patterns without explicit reference to the activation paradigm. McKeown et al. [1998] used independent component analysis and found several types of activations including components with "transient task related" response, i.e., responses that could not simply be accounted for by the paradigm. The model presented in this paper draws on the experimental observation that the basic coupling between the net neural activity and hemodynamic response is roughly linear while the relation between neuronal response and stimulus/task parameters is often nonlinear [Dale 1997]. We will represent the neuronal activity (integrated over the voxel and sampling time interval) by a smooth binary on/off signal while we will represent the hemodynamic response as a linear filter of unknown form and temporal extent.

2 A Bayesian model of fMRI time series

Let $s = \{s_t : t = 1 \dots T\}$ be a hidden sequence of T binary state variables $s_t \in \{0, 1\}$, representing the state of a single voxel over time; the time variable, t, indexes the sequence of fMRI scans. Hence, s_t is a binary representation of the neural state. The model assumes that the observable $z = \{z_t : t = 1 \dots T\}$ is the result of filtering the hidden state sequence by a linear system of order M_f plus a constant and trend, corrupted by independent additive Gaussian noise with variance σ_{n}^2. The hidden sequence is governed by a symmetric first order Hidden Markov Model (HMM) with transition probability α. The generative model considered is therefore given by:

$$\begin{array}{rcl} p(s_i|s_{i-1}, \alpha) & = & \alpha\delta_{s_i,s_{i-1}} + (1-\alpha)(1-\delta_{s_i,s_{i-1}}), \\ p(z|s, \sigma_{\mathrm{n}}, \theta, M) & \sim & \mathcal{N}(y, \sigma_{\mathrm{n}}^2 I), \ \text{ where } \ y = H_{s,M}\theta, \end{array} \tag{1}$$

where $\delta_{s_i,s_{i-1}}$ is the usual Kronecker delta and θ is a real valued vector of size $M = M_f + 2$, containing an off-set, linear trend and the filter impulse response. The *predicted signal*, y is the convolution of θ and s and is equal to $H_{s,M}\theta$, where the $T \times M$ matrix $H_{s,M}$ is given by:

$$H_{s,M} = \begin{pmatrix} 1 & 1/T & s_1 & 0 & & 0 \\ 1 & 2/T & s_2 & s_1 & & \vdots \\ 1 & 3/T & s_3 & s_2 & \ddots & 0 \\ \vdots & \vdots & \vdots & \vdots & & s_1 \\ 1 & (T-1)/T & s_{T-1} & s_{T-2} & & \vdots \\ 1 & 1 & s_T & s_{T-1} & & s_{T-M_f+1} \end{pmatrix}. \tag{2}$$

The first two columns have been included to account for the off-set and linear trend which are typically seen in fMRI time series.

3 Analytic integration and Monte Carlo sampling

In this section we introduce priors over the model parameters and show how inference may be performed. The filter coefficients and noise parameters may be handled analytically, whereas the remaining parameters are treated using sampling procedures (a combination of Gibbs and Metropolis sampling).

The filter coefficients for the off-set and trend are given (improper) uniform priors. The remaining filter coefficients are given priors that are uniform on an interval of length β independently for each coefficient:

$$p(\theta|M) = \begin{cases} \beta^{-M_f} & \text{for} \quad |\theta_i| < \frac{\beta}{2}, \quad 2 < i \leq M \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Assuming that all the values of θ for which the associated likelihood has non-vanishing contributions lie inside the box where the prior has support, we may integrate out the filter coefficients via a Gaussian integral:

$$\begin{array}{rcll} p(z|\sigma_{\mathrm{n}}, s, M) & = & \displaystyle\int p(z|\theta, \sigma_{\mathrm{n}}, s, M)p(\theta|M)d\theta & (4) \\ & = & \beta^{-M_f}(2\pi\sigma_{\mathrm{n}}^2)^{\frac{M-T}{2}}|H'_{s,M}H_{s,M}|^{-1/2}\exp\big(-\dfrac{z'z - \hat{y}'_{s,M}\hat{y}_{s,M}}{2\sigma_{\mathrm{n}}^2}\big), & \\ & \text{where} & \hat{y}_{s,M} = H_{s,M}\hat{\theta}_{s,M}, \quad \text{and} \quad \hat{\theta}_{s,M} = (H'_{s,M}H_{s,M})^{-1}H_{s,M}z. & \end{array}$$

We have here defined the *mean filter*, $\hat{\theta}_{s,M}$ and *mean predicted signal*, $\hat{y}_{s,M}$, for given state and filter length. We set the interval-length, β to be 4 times the standard deviation of the observed signal z. This is done, since the response from the filter should be able to model the signal, for which it is thought to need an interval of plus/minus two standard deviations.

We now proceed to integrate over the noise parameter; using the (improper) non-informative Jeffreys prior, $p(\sigma_{\mathrm{n}}) \propto \sigma^{-1}$, we get a Gamma integral:

$$\begin{aligned} p(z|s,M) &= \int p(z|\sigma_{\mathrm{n}},s,M)p(\sigma_{\mathrm{n}})d\sigma_{\mathrm{n}} \\ &= \frac{1}{2}\Gamma\big(\frac{T-M}{2}\big)\left(\pi(z'z-\hat{y}'_{s,M}\hat{y}_{s,M})\right)^{(M-T)/2}\beta^{-M_f}|H'_{s,M}H_{s,M}|^{-1/2}. \end{aligned} \tag{5}$$

The remaining variables cannot be handled analytically, and will be treated using various forms of sampling as described in the following sections.

3.1 Gibbs and Metropolis updates of the state sequence

We use a flat prior on the states, $p(s_i = 0) = p(s_i = 1) = 1/2$, together with the first order Markov property for the hidden states and Bayes' rule to get the conditional posterior for the individual states:

$$p(s_i = j|s_{-i},\alpha,M) \quad \propto \quad p(s_i = j|s_{i-1},\alpha)p(s_{i+1}|s_i = j,\alpha)p(z|s,M). \tag{6}$$

These probabilities may (in normalized form) be used to implement Gibbs updates for the hidden state variables, updating one variable at a time and sweeping through all variables. However, it turns out that there are significant correlations between states which makes it difficult for the Markov Chain to move around in the hidden state-space using only Gibbs sampling (where a single state is updated at a time). To improve the situation we also perform global state updates, consisting of proposing to move the entire state sequence one step forwards or backwards (the direction being chosen at random) and accepting the proposed state using the Metropolis acceptance procedure. The proposed movements are made using periodic boundary conditions. The Gibbs sweep is computationally involved, since it requires computation of several matrix expressions for every state-variable.

3.2 Adaptive Rejection Sampling for the transition probability

The likelihood for the transition probability α is derived from the Hidden Markov Model:

$$p(s|\alpha) = Z^{-1}(\alpha)\prod_{t=1}^{T} p(s_t|s_{t-1},\alpha) = Z^{-1}(\alpha)\alpha^{E(s)}(1-\alpha)^{T-E(s)}\ , \tag{7}$$

where $E(s) = \sum_{i=1}^{T}\delta_{s_i,s_{i-1}}$ is the number of neighboring states in s with identical values. Here we have introduced a periodic boundary condition by setting $s_0 = s_T$. Using the periodic boundary condition, it may be shown that the normalizing constant is $Z(\alpha) = 1-(2\alpha-1)^T$, see eg. [Reichl, 1990]. We note that for $\alpha \in]0,1[$ the product of the conditional state probabilities becomes the normalized joint distribution when T becomes large. The prior on the transition probabilities is uniform, but restricted to be larger than one half, since we expect the activation to mimic the blocked structure of the experimental paradigm. It is readily shown that $p(\alpha|s) \propto p(s|\alpha)/Z(\alpha)$, $\alpha \in [\frac{1}{2},1]$ is log-concave (for $T > 1$). Hence, we may use the Adaptive Rejection Sampling algorithm [Gilks and Wild, 1992] to sample from the distribution for the transition probability.

3.3 Metropolis updates for the filter length

In practical applications using real fMRI data, we do typically not know the necessary length of the filter. The problem of finding the "right" model order is difficult and has received a lot of attention. Here, we let the Markov Chain sample over different filter lengths,

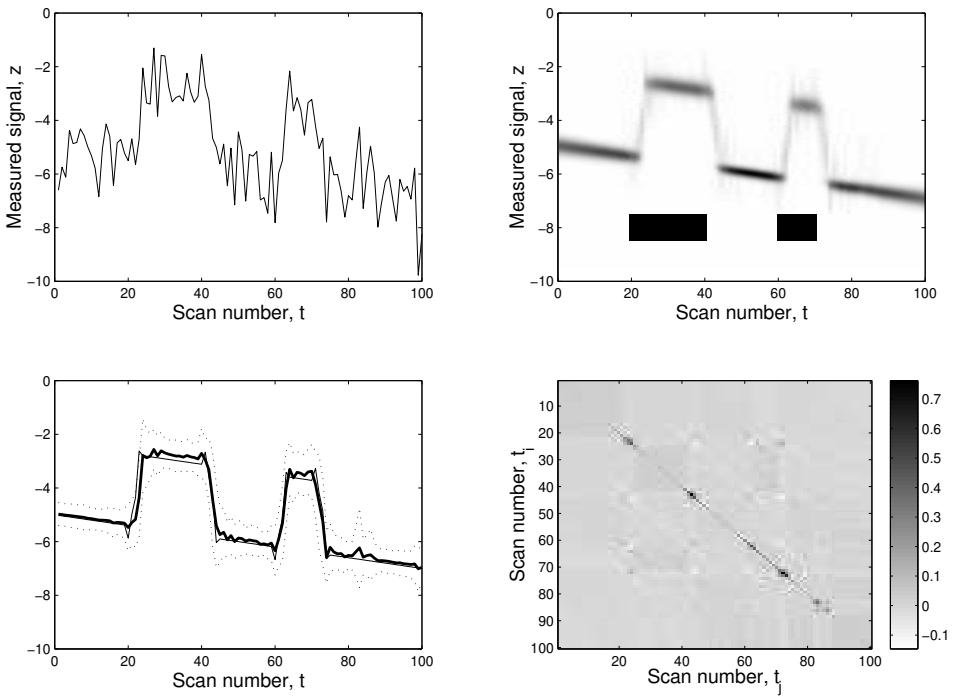

Figure 1: Experiments with synthetic data. Top left, the measured response from a voxel is plotted for 100 consecutive scans. In the bottom left, the underlying signal is seen in thin, together with the posterior mean, $\hat{y}$ (thick), and two std. dev. error-bars in dotted. Top right, the posterior probabilities are shown as a grey level, for each scan. Activated instances are indicated in black. Bottom right shows the posterior uncertainty, eq. (10)

effectively integrating out the filter-length rather than trying to optimize it. Although the value of M_f determines the dimensionality of the parameter space, we do not need to use specialized sampling methodology (such as Reversible Jump MCMC [Green, 1995]), since those parameters are handled analytically in our model. We put a flat prior on M_f and propose new filter lengths using a Gaussian proposal centered on the current value, with a standard deviation of 3 (non-positive proposed orders are rejected). This choice of the standard deviation only effects the mixing rate of the Markov chain and does not have any influence on the stationary distribution. The proposed values are accepted using the Metropolis algorithm, using $p(M|s, y) \propto p(y|s, M)$.

4 Example: synthetic data

In order to test the model, we first present results on a synthetic data set. A set of 100 values are generated using a $M_f = 5$ order filter, and a hidden state sequence consisting of two activation bursts (indicated by dark bars in figure 1 top right). The model is run for 10000 iterations, which is sufficient to generate 100 approximately independent samples from the posterior; figure 2 (right) shows the auto-covariance for M_f as a function of the iteration lag. It is thought that changes in M_f are indicative of correlation time of the overall system.

In order to verify that the model can reasonably recover the parameters used to generate the data, posterior samples from some of the model variables are shown in figure 3. For all these parameters the posterior density is large around the correct values. Notice, that the

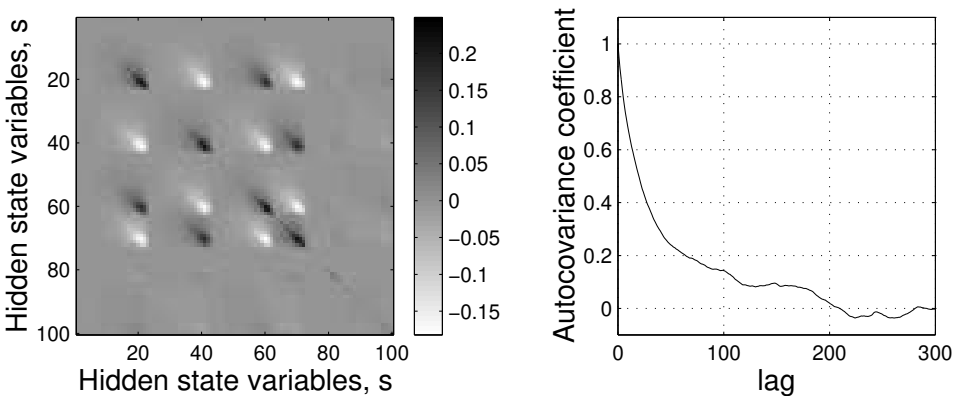

Figure 2: To the left, the covariance of the hidden states based on a long run of the model are shown to the left. Notice, that the states around the front (back) of the activity "bumps" are highly (anti-) correlated. Right: auto-covariance for the filter length M_f as a function of the lag time in iterations. The correlation length is about 100, computed as the sum of autocovariance coefficients from lag -500 to 500.

distribution of M_f is fairly broad, indicating that it is difficult to infer the model order from the data. We have not attempted to plot the inferred filter coefficients since the simultaneous inference of state sequence and filter parameters leads to an indeterminacy; for example, the same signal is predicted by shifting the states backward and introducing a leading zero filter coefficient. However, of course the is no indeterminacy in the predicted signal.

The correlation plot for the hidden states (figure 2, left) shows that the state activation onset correlates strongly with the second onset and negatively with the end of the activation. This indicates that the Metropolis updates described in section 3.1 may indeed be effective. Notice also that the very strong correlation among state variables does not strongly carry over to the predicted signal (figure 1, bottom right).

Since θ has a flat prior the conditional probability for the filter coefficients is proportional to the likelihood $p(z|\theta, \cdot)$ and by eq. (1) we get:

$$p(\theta|z, s, \sigma_{\mathrm{n}}, M) \sim \mathcal{N}(Dz, \sigma_{\mathrm{n}}^2 D_{s,M} D'_{s,M}), \quad D_{s,M} = (H'_{s,M} H_{s,M})^{-1} H'_{s,M} \; . \tag{8}$$

The posterior mean of the predicted signal, $\hat{y}$, is then readily computed as:

$$\hat{y} = \langle y_{\theta,\sigma_{\mathrm{n}},s,M} \rangle_{\theta,\sigma_{\mathrm{n}},s,M} = \langle \hat{y}_{s,M} \rangle_{s,M} = \langle H_{s,M} \hat{\theta}_{s,M} \rangle_{s,M} = \langle F_{s,M} \rangle z \; , \tag{9}$$

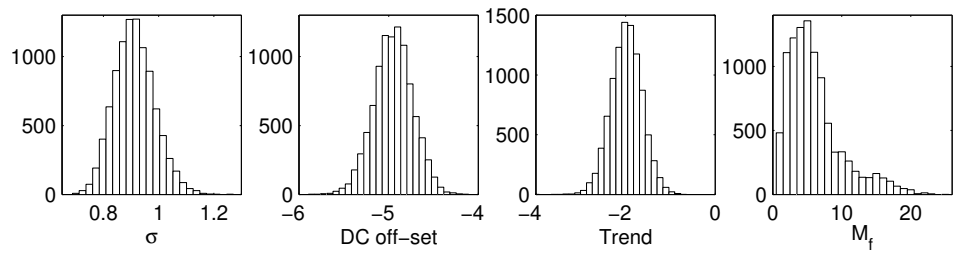

Figure 3: Posterior distributions of various model parameters. The parameters used for generation were: $\sigma = 0.9$, off-set $= -5$, trend $= -2$ and filter order $M_f = 5$.

where $F_{s,M} = H_{s,M} D_{s,M}$. Here, the average over θ and σ_{n} is done analytically, and the average over the state and filter length are done using Monte Carlo. The uncertainty in the posterior, can also be estimated partly by analytical averaging, and partly Monte Carlo:

$$\begin{aligned}\sigma_y^2 &= \langle (y_{\theta,\sigma_{\mathrm{n}},s,M} - \hat{y})(y_{\theta,\sigma_{\mathrm{n}},s,M} - \hat{y})' \rangle_{\theta,\sigma_{\mathrm{n}},s,M} \\ &= \frac{1}{T-M} \left\langle (z'z - \hat{y}'_{s,M}\hat{y}_{s,M}) F_{s,M} F'_{s,M} \right\rangle_{s,M} + \left\langle F_{s,M} z z' F'_{s,M} \right\rangle_{s,M} - \hat{y}'\hat{y},\end{aligned} \tag{10}$$

The posterior mean and the two standard deviations are plotted in figure 1 bottom left. Notice, however, that the distribution of y is not Gaussian, but rather a mixture of Gaussians, and is not necessarily well characterized by mean and variance alone. In figure 1, top left, the distribution of y_t is visualized using grey-scale to represent density.

5 Simulations on real fMRI data and discussion

In figure 4 the model has been applied to two measurements in the same voxel in visual cortex. The fMRI scans were acquired every 330 ms. The experimental paradigm consisted of 30 scans of rest followed by 30 scans of activation and 40 rest. Visual activation consisted of a flashing (8 Hz) annular checkerboard pattern. For each run 10000 samples were generated which in a Matlab implementation takes around 3 hours on a 400 MHz PC. The model readily identifies the activation burst of somewhat longer duration than the visual stimulus and delayed around 2 seconds. The delay is in part caused by the delay in the hemodynamic response.

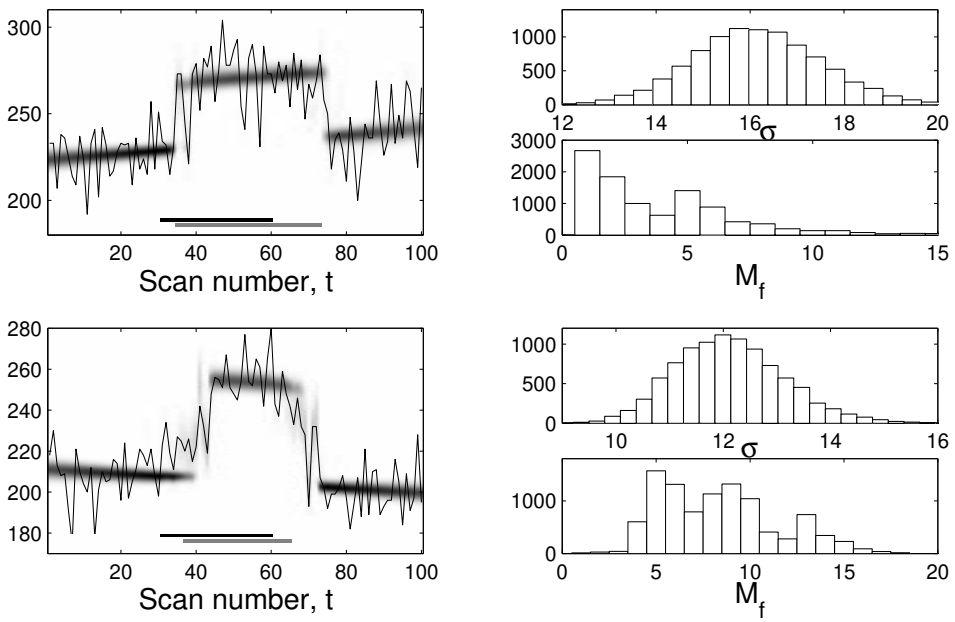

Figure 4: Analysis of two experimental runs in the same voxel in visual cortex. The left hand plot show the posterior inferred signal distribution superimposed by the measured signal. The colored bars show the experimental paradigm (top) and the hidden state variables with mean greater than 0.5. To the right the posterior noise level and inferred filter length are displayed. The noise is fairly concentrated whereas filter length has large uncertainty.

These results show that the integration procedure works in spite of the very limited data at hand. Because of the indeterminacy in the system we are not trying to make any statements

about the from of the specific physical filter (hemodynamic response) used in generating the measured fMRI signal. We are only interested in the predicted signal and the posterior of the hidden states. This also warrants caution when interpreting the inferred states because of the indeterminacy discussed earlier.

Since the posterior distribution of the filter length is very broad it is questionable whether an optimization based procedure such as maximum likelihood estimation would be able to make useful inference in this case were data is very limited. Also it is not obvious how one may use cross-validation in this setting. One might expect such optimization based strategies to get trapped in suboptimal solutions. This, of course, remains to be investigated.

6 Conclusion

We have presented a model for voxel based explorative data analysis of single trail fMRI signals during blocked task activation studies. The model is founded on the experimental observation that the basic coupling between the net neural activity and hemodynamic response is roughly linear. The preliminary investigation reported here are encouraging in that the model reliably detects reasonable hidden states from the very noisy fMRI data.

One drawback of this method is that it is computational expensive. However, implementing the algorithm in C would improve the speed substantially. Current work is in progress for generalizing the model to multiple voxels, including spatial correlation such as spill-over effects.

Acknowledgments

Thanks to Egill Rostrup for providing the fMRI data. This work is funded by the Danish Research Councils through the Computational Neural Network Center (CONNECT) and the THOR Center for Neuroinformatics.

References

Bandettini, P. A. (1993). Processing strategies for time-course data sets in functional MRI of the human brain *Magnetic Resonance in Medicine 30*, 161–173.

Dale, A. M. and R. L. Buckner (1997). Selective Averaging of Individual Trials Using fMRI. *NeuroImage 5*, Abstract S47.

Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika 82*, 711–732.

Gilks, W. R. and P. Wild (1992). Adaptive rejection sampling for Gibbs sampling. *Applied Statistics 41*, 337–348.

Hansen, L. K. et al. (1999). Generalizable Patterns in Neuroimaging: How Many Principal Components? *NeuroImage*, to appear.

Lange, N. and S. L. Zeger (1997). Non-linear Fourier time series analysis for human brain mapping by functional magnetic resonance imaging. *Journal of the Royal Statistical Society - Series C Applied Statistics 46*, 1–30.

McKeown, M. J. et al. (1998). Spatially independent activity patterns in functional magnetic resonance imaging data during the stroop color-naming task. *Proc. Natl. Acad. Sci. USA. 95*, 803–810.

Reichl, L. E. (1990). *A modern course in statistical physics*, Edward Arnold.

Worsley, K. J. et al. (1997). Characterizing the Response of PET and fMRI Data Using Multivariate Linear Models (MLM). *NeuroImage 6*, 305–319.