Technical University of Denmark

DTU

Identification of non-linear models of neural activity in bold fmri

Jacobsen, Danjal Jakup; Madsen, Kristoffer Hougaard; Hansen, Lars Kai

DTU Library
Technical Information Center of Denmark

IDENTIFICATION OF NON-LINEAR MODELS OF NEURAL ACTIVITY IN BOLD FMRI

Daniel J. Jacobsen, Kristoffer Hougaard Madsen, Lars Kai Hansen

Intelligent Signal Processing
Technical University of Denmark

ABSTRACT

Non-linear hemodynamic models express the BOLD signal as a nonlinear, parametric functional of the temporal sequence of local neural activity. Several models have been proposed for this neural activity. We identify one such parametric model by estimating the distribution of its parameters. These distributions are themselves stochastic, therefore we estimate their variance by epoch based leave-one-out cross validation, using a Metropolis-Hastings algorithm for sampling of the posterior parameter distribution.

1. INTRODUCTION

Neuroimaging has made major contributions to our understanding of the relation between behavior and distribution of brain information processing. The richest neuroimaging modality is fMRI based on the so-called BOLD effect, involving the hemodynamic response which is rather sluggish, non-linear, and non-local. With the long term goal of increasing the spatio-temporal resolution of BOLD fMRI, we are interested in models linking subject behavior, neural activity, the so-called hemodynamic response, and fMRI BOLD observations, see e.g., [1, 2, 3, 4].

We will examine the model proposed by Friston et al. [2] which consists of a set of ordinary differential equations (ODE's) that model the evolution in time of four basic physiological state variables: The blood volume $v(t)$, blood inflow $f(t)$, amount of de-oxyhemoglobine $q(t)$ and a so-called 'flow inducing signal' $s(t)$, collected in the state vector, $\mathbf{x}(t) = [v(t)\ q(t)\ f(t)\ s(t)]^T$. The flow inducing signal is driven by an underlying neural activation function $\nu(t)$ - a time function describing the local neural activity.

The measured BOLD signal y_n is then modeled as a nonlinear function of 'snapshots' of the continuously evolving states, with additive white Gaussian noise w_n; subscript indices are used for these variables to emphasize the discrete 'sampled' nature.

$$\frac{\partial \mathbf{x}}{\partial t} = f(\mathbf{x}(t), \nu(t))$$
$$y_n = g(\mathbf{x}(t_n)) + w_n \tag{1}$$

The BOLD signal is measured with a sampling interval denoted *TR*. The model has seven parameters: σ_w^2, the vari-

ance of w_n, and six physiological parameters[1], combined in $\theta = \begin{bmatrix} \alpha\ \epsilon\ \tau_0\ \tau_s\ \tau_f\ E_0\ \sigma_w^2 \end{bmatrix}^T$. E_0 is the so-called 'resting net oxygen extraction fraction in the capillary bed'.

In addition, we assume that the states evolve from an initial known resting state $\mathbf{x}_0 = [0\ 1\ 1\ 1]^T$. The latter assumption is reasonable if a suitably long resting period precedes stimulation sequences. The dynamical model thus consists of the set of non-linear differential equations

$$\frac{\partial v(t)}{\partial t} = \frac{1}{\tau_0} \left(f(t) - v(t)^{1/\alpha} \right)$$

$$\frac{\partial q(t)}{\partial t} = \frac{1}{\tau_0} \left[f(t) \frac{1 - (1 - E_0)^{1/f(t)}}{E_0} - v(t)^{(1-\alpha)/\alpha} q_t \right]$$

$$\frac{\partial s(t)}{\partial t} = \epsilon \nu(t) - s(t)/\tau_s - (f(t) - 1)/\tau_f$$

$$\frac{\partial f(t)}{\partial t} = s(t)$$

Finally, the BOLD observation model involves the non-linearity,

$$\begin{aligned} g(\mathbf{x}(t)) = V_0[&(k_1 + k_2)(1 - q(t)) \\ &- (k_2 + k_3(1 - v(t)))] \end{aligned} \tag{2}$$

with a set of empirical constants taking values $V_0 = 0.02$, $v_0 = 40.3, TE = 0.03, r_0 = 25, \epsilon = 1.43$, $k_1 = 4.3 E_0 v_0 TE, k_2 = E_0 \epsilon r_0 TE, k_3 = \epsilon - 1$ for BOLD imaging at $1.5T$ [5]. The BOLD fMRI measurements are spatially sampled in volume elements (voxels). Experiments are typically ivided temporally in quasi-independent baseline-activation stimulus 'epochs'.

2. STATISTICAL MODELING

For given hemodynamic parameters and neural activity, the likelihood of an epoch is straightforward to set up. First, the hidden states will evolve deterministically according to (1) driven by the given neural activity $\nu(t)$. We use a variable step-size 4th/5th-order embedded Runge-Kutta method to solve these [6], with the starting condition $\mathbf{x}(t = 0) =$

---

[1]we assume the resting blood volume fraction, V_0, to be constant at 0.02 [2]

$\mathbf{x}_0$, the initial (relaxed) state, $\mathbf{x}_0 = [1\ 1\ 1\ 0]^T$ (all values relative to resting state). This gives a sequence of states, $\mathbf{x}_{1:N} \triangleq \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$, corresponding to the sampling times $\{t_1, t_1 + TR, \ldots, t_1 + N \cdot TR\}$, where t_1 is the starting time of the epoch. The mean BOLD signal is given by

$$y_n = g(\mathbf{x}_n; \theta)$$

with the observed output modelled as

$$y_n^* = y_n + w_n(\theta) \tag{3}$$

As the residuals are assumed normal i.i.d., $y_n^* \sim \mathcal{N}(y_n, \sigma_w^2)$, the likelihood becomes

$$p(y_{1:N}^* | \theta) = \prod_{n=0}^{N} p(y_n^* | \theta) = \prod_{n=0}^{N} \mathcal{N}_{y_n^*}(y_n, \sigma_w^2).$$

The neural activity function $\nu(t)$ is traditionally assumed to be a simple square wave signal representing signal 'on' during stimulus and signal 'off' during baseline [2]. Buxton et al. [3] have recently introduced an alternative dynamical model representing neural activity which we will use here. This model posits

$$\nu(t) = a(t) - I(t)$$
$$\frac{dI}{dt} = \frac{\kappa \nu(t) - I(t)}{\tau_1}, \tag{4}$$

where $a(t)$ is the square wave stimulus reference function and $I(t)$ is an inhibitory feedback signal. The values of the parameters are unknown a priori, although in [3], ranges are given as $\tau_1 \in [1; 3]$, $\kappa \in [0; 3]$. We refer to them jointly as $\phi = [\kappa, \tau_i]$.

Note that the square wave model obtains as a special case of this non-linear model as $\kappa/\tau_i \rightarrow 0$.

In this report we will apply this model to real fMRI data and investigate the posterior distribution of κ and τ_i. By way of leave-one-out cross-validation, the uncertainty on these distributions will further be described.

3. GENERALIZATION AND ESTIMATION PROCEDURES

The target for this investigation is the *posterior* distribution of the parameters of the non-linear neural model, ϕ:

$$p(\phi | D) = \frac{p(D|\phi)p(\phi)}{p(D)} = \frac{p(D|\phi)p(\phi)}{\int p(D|\phi)p(\phi)d\phi}$$

i.e. the distribution of the parameters conditioned on the observed BOLD data, D. This distribution cannot be obtained analytically; instead, we employ a Metropolis-Hastings Markov-chain Monte Carlo (MCMC) method to generate samples from the posterior.

3.1. Markov-chain Monte Carlo sampling

We use a Metropolis-Hastings (MH) algorithm [7] starting at an arbitrary state, ϕ_0, and at each step proposing small changes in ϕ from a *proposal distribution*, in our case a Gaussian centered on the 'current' state $p(\phi_{n+1}|\phi_n) = \mathcal{N}(\phi_n; \Sigma)$. The parameters of the hemodynamic model are sampled simultaneously, but as they are not the focus of this report, they are not discussed further and ignored for clarity (also the approximate marginal distribution of κ and τ_i is just the sampled values of these, so no further work is required to obtain the desired distribution).

The MH method produces samples from the true posterior distribution in the limit of large number of samples. We employ a set of heuristics to ensure convergence before averaging, e.g., inspection of saturation of the training set likelihood and stabilization of the actual parameter values.

A good proposal distribution is major determinant for succes of the Metropolis-Hastings algorithm. Again we invoke heuristics: Starting with a spherical normal distribution of dimension $\dim \phi$, we perform several (short) scout sampling runs. After each of these, the covariance of the generated samples is used to adapt the covariance of the proposal, Σ, scaled to give an acceptance rate of around 0.3. This procedure greatly speeds up the final sampling run (the samples of the initial runs are not used in averaging).

For most of the parameters, we use simple uniform priors ($p(\phi)$) over positive parameters; for the oxygen extraction fraction E_0 we use a uniform distribution over the interval $[0, 1]$.

These samples (ϕ_n, $n = 1..N$) can then be used as an empirical approximation of the target distribution, e.g. for prediction of a BOLD signal y:

$$p(y|D) = \int p(y|\phi)p(\phi|D)d\phi$$
$$\approx \frac{1}{N}\sum_{n=1}^{N} p(y|\phi_n), \quad \phi_n \sim p(\phi|D)$$

The resulting approximate distribution clearly depends on the particular data set D. To obtain information on the uncertainty of the posterior, we employ epoch-wise leave-one-out cross-validation. With K quasi-independent epochs available, each split of the data leaves out one epoch for a test set, which can be used to validate the ability of the model to prediction test data, while $K - 1$ are used for a training set to give an estimate of the posterior distribution.

To get an estimate of the uncertainty of the posterior distribution approximations, we fit a normal distribution to the samples for each cross-validation split. This yields a distribution of means and variances that can be used to illustrate the variance of the distribution. Each split yields an unbiased estimate of the true mean and variance of the distribution, and the mean over all splits is a convex combination thus also an unbiased estimate, but with a reduced variance.

4. EXPERIMENTAL EVALUATION

The described method was evaluated on a synthetically generated data set and on a real BOLD fMRI data set.

4.1. Synthetic data

The synthetic data was obtained by simulating the hemodynamic model with parameters set to the maximum likelihood values reported by Friston et al. [2],

$$\theta = \begin{bmatrix} 0.33 & 0.54 & 0.98 & 1.54 & 2.46 & 0.34 & \sigma_w^2 \end{bmatrix}^T,$$

with σ_w^2 set to produce a desired SNR (signal-to-noise ratio, measured as the ratio of the de-noised BOLD and observation noise signals) close to 5.0 dB, which is similar to real recording conditions. κ and τ_i were set to 2.0 and 1.6 respectively. The model is initialized in $\mathbf{x}_0$, the states are evolved using a Runge-Kutta solver, and observations are made, adding Gaussian white noise with the prescribed variance. Each epoch contains 100 samples with sampling interval $TR = 1.0s$.

To justify our assumption that the BOLD signal is independent between epochs, the stimulus for each epoch is set to zero for the last 30 seconds. This is helpful for two further reasons. Preprocessing (e.g. removing low-frequency noise), is aided in that such artifacts can be more accurately estimated using these 'resting' portions of data. Finally, it allows us to assume a known, resting, physiological state ($\mathbf{x}_0$) at the start of each epoch. The stimulus is designed to evoke non-linear behavior in the model; this is achieved by inter-stimulus intervals (ISI) and stimulus durations (SD) being sampled from a suitable gamma distribution.

Figure 1 shows parts of the reconstructed neural and BOLD signals. The posterior mean of κ and τ was 1.92 ± 0.51 and 1.32 ± 0.22 respectively, resulting in very close reconstructions.


(a) (b)

Fig. 1. (a): One pulse of the reconstructed and true neural activity. The shape of the signal is retrieved, although the slope is slightly off, corresponding to a bias in the τ_i estimate. (b): Reconstructed and true BOLD signal (first 3 epochs).

4.2. Real data

The data was acquired at Hvidovre Hospital, Denmark, using a 3T scanner (Magnetom Trio, Siemens). We obtained 1382 GRE EPI volumes each consisting of twelve 3mm slices oriented along the calcarine sulcus. Additional parameters where TR=725 ms, TE=30 ms FOV=192 mm, 64x64 acquisition matrix, FA = 82. The stimulus consisted in a circular black/white flickering checkerboard (24 degrees horizontal, 18 degrees vertical) on a grey background. The checkers reversed black/white at 8 Hz. The activation pattern ($a(t)$) used to determine on- and offset of this stimulus was the same as was used to generate the synthetic data.

Fifty significantly activated (as determined by SPM2 analysis [2] [8]) voxels in visual cortex were selected, and the mean of these was used as the BOLD signal.

The results are shown in figure 2. The resulting shape of the neural model is similar to the one found for the synthetic data, and is significantly different from square. The BOLD reconstruction (see (3.1)) on test data is satisfactory, validating the model and method. We found posterior mean values of 3.11 ± 0.76 for κ and 0.87 ± 0.19 ($\pm$ one standard deviation) for τ_i; again, not close to a square pulse.

Figure 3 shows the normal approximations to the posterior histograms. These indicate that although there is significant variation, the mean of all the posterior means obtained from the leave-one-out cross-validation is identifiable. We expect with longer sampling runs to bring down this variability - the important point is that this is a tool that gives guidance on the reliability of the estimated distribution.


(a) (b)

Fig. 2. (a): One pulse of the estimated neural activity. The posterior mean reconstruction is shown together with the reconstruction corresponding to $\pm$ one standard deviation of the distribution of the posterior mean. (b): Prediction on real test data epochs (first 3 epochs).

5. CONCLUSION

The method detailed here can be used obtain to obtain approximate posterior distributions of model parameters together with estimates on the reliability of these approximations. For the

---

[2]Software available from http://www.fil.ion.ucl.ac.uk/spm/

954

Fig. 3. (a): Normal approximations of posterior parameter histograms. (a) and (b) show the variation of the mean and variance respectively, for the posterior distribution of κ. (c) and (d) show the same for τ_i (real data).

parameters of the non-linear neural activity model, we found posterior mean values of 3.11 ± 0.76 for κ and 0.87 ± 0.19 ($\pm$ one standard deviation) for τ_i. Both of these are on the edge of the ranges given in [3], although not statistically significantly so. There is little other information available on statistical estimation of these parameters.

We find that both κ and τ are identifiable for real BOLD data, and that κ/τ_i for our data is significantly greater than zero. Thus the square model of neural activity which is widely used in BOLD analysis is not supported by our findings.

The present model may be inverted to produce estimates of neural activity as indicated in the work of Riera et al. [4]. In [4] a regularized radial basis function set is used, with parameters estimated using a likelihood based approach which leads to rather smooth activation estimates. Using our Bayesian sampling approach from an augmented posterior distribution including parameters of the neural activity time course (such as stimulus onset times etc.) may be a way to let data determine the level of regularization, hence, potentially lead to more crisp estimates of non-trivial neural activation sequences. This would be of particular interest in more complex activation designs involving different stimulus activation conditions within epochs. In the present model we have focused on the local hemodynamics in average data from a region. The BOLD hemodynamics is non-local and it is an important future task to produce a spatio-temporal hemodynamic model, which could also lead to improved spatial resolution.

6. REFERENCES

[1] R. B. Buxton, E. C. Wong, and L. R. Frank, "Dynamics of blood flow and oxygenation changes during brain activation: the balloon model," *MRM*, vol. 39, pp. 855–864, June 1998.

[2] 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, 2000.

[3] R. B. Buxton, K. Uludag, D. J. Dubowitz, and T. T. Liu, "Modeling the hemodynamic response to brain activation," *NeuroImage*, vol. 23, pp. 220–33, 2004.

[4] J.J. Riera, J. Watanabe, K. Iwata, N. Miura, E. Aubert, T. Ozaki, and R. Kawashima, "A state-space model of the hemodynamic approach: nonlinear filtering of bold signals," *Neuroimage*, vol. 21, pp. 547–567, 2004.

[5] T. Obata, T.T. Liu, K.L. Miller, W.M Luh, E.C. Wong, L.R. Frank, and R.B. Buxton, "Discrepancies between bold and flow dynamics in primary and supplementary motor areas: application of the balloon model to the interpretation of bold transients," *Neuroimage*, vol. 21, pp. 144–153, 2004.

[6] M. Abramowitz and I. A. Stegun, Eds., *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, 9th printing*, New York: Dover, 1972.

[7] D.J.C. MacKay, *Information theory, inference, and learning algorithms*, Cambridge University Press, 2003.

[8] K.J. Friston, *Introduction: Experimental design and statistical parametric mapping*, Academic Press, second edition, 2003.