

Technical Note

Dynamic causal modeling: A generative model of slice timing in fMRI

Stefan J. Kiebel,* Stefan Klöppel, Nikolaus Weiskopf, and Karl J. Friston

The Wellcome Department of Imaging Neuroscience, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3AR, UK

Received 21 June 2006; revised 16 October 2006; accepted 29 October 2006
Available online 11 December 2006

Dynamic causal modeling (DCM) of functional magnetic resonance imaging (fMRI) data allows one to make inferences about the architecture of distributed networks in the brain, in terms of effective connectivity. fMRI data are usually acquired using echo planar imaging (EPI). EPI sequences typically acquire slices at different times over a few seconds. DCM, in its original inception, was not informed about these slice timings and assumed that all slices were acquired simultaneously. It has been shown that DCM can cope with slice timing differences of up to 1 s. However, many fMRI studies employ a repetition time (TR) of 3 to 5 s, which precludes a straightforward DCM of these data.

We show that this limitation can be overcome easily by including slice timing in the DCM. Using synthetic data we show that the extended DCM furnishes veridical posterior means, even if there are large slice-timing differences. Model comparisons show that, in general, the extended DCM out-performs the original model. We contrast the modeling of slice timing, in the context of DCM, with the less effective approach of ‘slice-timing correction’, prior to modeling. We apply our procedure to real data and show that slice timings are important parameters. We conclude that, generally, one should use DCM with slice timing.

© 2006 Elsevier Inc. All rights reserved.

Keywords: Functional magnetic resonance imaging; Dynamic causal modeling; Sampling

Introduction

The aim of dynamic causal modeling (DCM) is to describe interactions among neuronal populations (Friston et al., 2003). DCM explains observed fMRI time-series by modeling interactions at a neuronal level and passing the ensuing neuronal states through a hemodynamic model to give region-specific blood oxygenation level-dependent (BOLD) signals (Buxton et al., 2004; Friston, 2002). The DCM is inverted using a variational Bayesian scheme to give posterior distributions of the model parameters. Inferences about a specific network are made using the posterior distributions of interesting parameters; e.g., the connectivity between areas or

the modulation of connectivity by experimental manipulations. This makes DCM useful for testing hypotheses about modulatory effects; i.e., how experimental conditions change directed connections among areas. For example, it has been shown, using DCM, that the superior posterior parietal cortex (SPC) exerts a modulatory role, that mediates attentional effects, on V5 responses (Friston et al., 2003; Friston and Buchel, 2000). Other examples of DCM can be found in Mechelli et al. (2004); Noppeney et al. (2006); and Stephan et al. (2005). DCM uses biophysically and physiologically inspired generative models.¹ This makes it relatively simple to extend existing DCMs. The usefulness of changes to the model can then be assessed using model comparison (Penny et al., 2004).

Currently, DCM for fMRI assumes that data from all areas are acquired at the same time. Clearly, this is an approximation, because fMRI data are usually acquired slice by slice. For a typical fMRI study, it is likely that data from remote areas (orthogonal to slice orientation) are acquired with a slice-timing difference in the order of seconds. Friston et al. (2003) showed, for block designs that the model can cope with slice-timing differences of up to 1 s. This robustness rests on changes in (non-interesting) hemodynamic parameters, which explain away slice-timing effects.

The assumption of simultaneous data acquisition restricts the application of DCM to data that are acquired with slice-timing differences of less than a second. The typical repetition time (TR), i.e., the time needed for acquiring one volume, lies between 3 and 5 s. Note that the TR itself is not the limiting factor; the key factor is the timing differences among modeled areas. With a TR of several seconds, it is quite likely that some areas have slice-timing differences of more than 1 s. This questions the validity of DCM for such data. Furthermore, for event-related designs, the robustness of DCM to slice-timing differences has never been assessed. This is important, because event-related designs are prevalent in fMRI.

In this paper, we propose an extension that accommodates slice-timing differences: instead of assuming that data were acquired at the same time, we incorporate information about temporal

* Corresponding author. Fax: +44 20 7813 1420.

E-mail address: skiebel@fil.ion.ucl.ac.uk (S.J. Kiebel).

Available online on ScienceDirect (www.sciencedirect.com).

¹ By ‘generative’, we mean that a DCM can be regarded as a prescription of how the observed data were generated (or how we can simulate synthetic data).

sampling into the DCM. This is implemented using an informed integration scheme, which computes the predicted output given the model parameters and sampling times. The resulting DCM can be regarded as a three-level model. The first level models (interesting) interactions at a neuronal level. The second level models the hemodynamics. The third (and now explicit) level is the discrete temporal sampling of continuous second-level output. At this third level, we specify when the data were acquired from each area. The contribution of the present work is to show that this temporal sampling level is a useful extension to the model.

We establish the validity of this extension using simulations. By comparing models using their log-evidences, we show that the extended model is superior. We also compare the extended model with the current practice of ‘slice-timing correction’, i.e., aligning the data temporally (before using DCM) in an attempt to discount timing differences. Finally, we demonstrate, using real data, the importance of modeling slice-timing differences.

Theory

Dynamic causal modeling for fMRI: levels one and two

In this section we review briefly dynamic causal models of fMRI data (Friston et al., 2003). In the next section, we extend this model to accommodate slice timing (see Fig. 1 for an overview).

A dynamic causal model is a multiple-input multiple-output system that comprises N_u inputs and N_r outputs with one output per region. The N_u inputs correspond to designed causes (e.g., boxcar or impulse stimulus functions used in conventional analyses). Each region produces a measured output that corresponds to the observed BOLD signal. These N_r time-series would normally be taken as the averages or first eigenvariables of N_r selected regions.

The first level models interactions among areas at the neuronal level. At this level, each region has a single state variable. This

state is a simple summary of neuronal ‘activity’ in an area. Friston et al. (2003) used a bilinear form to describe these neuronal dynamics:

$$F(z, u, \theta) = \dot{z} = \left(A + \sum_j^{N_u} u_j B^j \right) z + Cu$$

$$A = \frac{\partial F}{\partial z} = \frac{\partial \dot{z}}{\partial z} \Big|_{u=0}$$

$$B^j = \frac{\partial^2 F}{\partial z \partial u_j} = \frac{\partial}{\partial u_j} \frac{\partial \dot{z}}{\partial z}$$

$$C = \frac{\partial F}{\partial u} \Big|_{z=0} \quad (1)$$

The state vector z contains one scalar per region. The change in states is described by the sum of three effects. First, the matrix A describes directed connectivity between pairs of areas. The elements of this connectivity matrix are not a function of the input, and can be considered as a fixed connectivity that is intrinsic to the network of areas. Second, the elements of B^j represent the changes of connectivity induced by the inputs u_j (i.e., by condition-specific modulations). These bilinear terms B^j are usually the interesting parameters. Third, there is a direct influence of each input u_j on each area, encoded by the matrix C . The parameters of this system, at the neuronal level, are given by $\theta^n = (A, B^1, \dots, B^{N_u}, C)$. At this level, one can specify which connections, i.e., elements of the matrices, one wants to include in the model. Connections are removed by setting their prior mean and variance to zero. We will describe some examples below.

Put simply, a dynamic causal model is, like the general linear model, an equation which expresses predicted responses in terms of some parameters and explanatory variables. In our case the explanatory variables are the experimental inputs u , which correspond to stimulus functions in conventional models. The causal aspect comes from control theory, in which the response of causal models can be predicted from the current and past input. Critically, dynamic models do not predict the response *per se* but its rate of change. This rate of change can be a complicated nonlinear function of the models unknown parameters and known inputs. The form and parameterization of this function is entailed by the specific DCM used. In fMRI, people generally use a simple bilinear form with coupling parameter matrices A , B and C . Critically, the B matrix parameterizes interactions between the inputs and states; hence bilinear. This bilinear form can also be regarded as an approximation to any function $F(z, u, \theta)$ because one would get the bilinear equation if one took its Taylor expansion around $z=0$ and $u=0$; and kept only the terms that were first-order in the states or input. This is summarized in Eq. (1). Having chosen the form of the DCM, the objective is to estimate its unknown parameters, which in our case represent a useful summary of the coupling among brain regions.

At the second level, for each area, the neuronal state forms an input to a hemodynamic model that generates the BOLD signal. Area-specific hemodynamics are modeled by four hemodynamic state-variables z^h . The corresponding first-order ordinary differential equations are parameterized by five area-specific parameters θ^h . See Friston et al. (2003) for a complete description. Here, we




Fig. 1. The neuronal response is formed at the first level. The second level models the ensuing BOLD response. The third level emulates sampling by the MRI scanner.

summarize the integration of the neuronal and hemodynamic states by the generalized convolution

$$v(t) = h(u(t), \theta) \quad (2)$$

By integrating the differential equations of both levels we can compute the system's predicted response $v(t)$ as a continuous function of time given the neuronal and hemodynamic parameters $\theta = (\theta^n, \theta^h)$, input $u(t)$, and some initial states.

Dynamic causal modeling for fMRI: the third level

The first two levels are not a complete model for fMRI. The output $v(t)$ is continuous, whereas the measured data y are discrete. To bridge the gap between model and data, we have to apply some sampling function to the output; i.e., the sampled response of area i at scan j is

$$\begin{aligned} y_{ij} &= s_j(v, T_i) + \varepsilon_{ij} \\ &= v((j-1)TR + T_i) + \varepsilon_{ij} \end{aligned} \quad (3)$$

where $\varepsilon_i \sim \mathcal{N}(0, \Sigma_i)$ is a vector of serially correlated and normally distributed measurement errors ε_{ij} and $i=1, \dots, N_r$ and $j=1, \dots, N_T$. The vector T encodes the relative timing of region. Each element of T is the region-specific time from the onset (0 s) of a scan, until measurement of the slice in which the area is located² (Fig. 2). The sampling function simply samples, for each region, the response at the appropriate time, using the region-specific offset T_i that models slice timing. Note that Friston et al. (2003) assumed that the sampling times are the same for each area and, consequently, there was no need to make this sampling explicit. Here, we relax their assumption by letting T take different values for each region. We will describe below how T is specified operationally.

This completes the description of the model. The goal is to identify the parameters that maximize an objective function, which measures the discrepancy between the observed data y and their predicted values. We do this by approximating the posterior distribution on the parameters using standard variational techniques.

Estimation, inference and model comparison

For a given DCM, say model m , parameter estimation corresponds to approximating the moments of the posterior distribution given by Bayes rule

$$p(\theta|y, m) = \frac{p(y|\theta, m)p(\theta|m)}{p(y|m)} \quad (4)$$

The estimation procedure employed in DCM is described in Friston et al. (2003) and Kiebel et al. (2006). The posterior moments (conditional mean η and covariance Σ) are updated iteratively using variational Bayes under a fixed-form Laplace (i.e., Gaussian) approximation to the conditional density $q(\theta) = \mathcal{N}(\eta, \Sigma)$. This is formally equivalent to expectation–maximization (EM) that employs a local linear approximation of Eqs. (1), (2) and (3) about

² Note that the response of an area is typically computed over several, neighboring voxels. By the location of an area we mean some representative location, e.g., the center of mass.




Fig. 2. The difference between the original and the extended model: two-area example. In a hypothetical experiment, the scanner progresses through acquisition of a session. Each scan is measured in TR seconds. The arrows indicate the time points of the modeled MRI sampling. Top: the original model assumes that both areas are sampled from the last slice. Bottom: in the extended model, the MRI sampling of both areas is modeled at the (known) times at which they were actually acquired. In our example, the first area was measured after a quarter of TR (T_1), and the second area after three quarters of TR (T_2).

the current conditional expectation. The source code (Matlab, MathWorks) of these functions can be accessed in the file 'spm_nlsi_GN.m' in the software package <http://www.fil.ion.ucl.ac.uk/spm/software/spm5/>.

Bayesian inference proceeds using the conditional or posterior density. Usually this involves specifying a parameter or compound of parameters as a contrast $c^T \eta$. In the results section, we will focus on the posterior mean (corresponding to a contrast vector c with a single one and zeros elsewhere) for analysis. Inferences about this contrast are made using its conditional covariance $c^T \Sigma c$. For example, one can compute the probability that any contrast is greater than zero or some meaningful threshold, given the data. This inference is conditioned on the particular model specified.

Often, one wants to compare different models for a given data set. We use Bayesian model comparison using the model evidence (Penny et al., 2004), which is

$$p(y|m) = \int p(y|\theta, m)p(\theta|m)d\theta. \quad (5)$$

Note that the model evidence is simply the normalization constant in Eq. (4). The evidence can be decomposed into two components: an accuracy term, which quantifies the data fit, and a complexity term, which penalizes models with redundant parameters. Therefore, the evidence embodies the two conflicting requirements of a good model; that it explains the data and is as simple as possible. In the following, we approximate the model evidence for model m , under the Laplace approximation, with

$$\ln p(y|m) \approx \ln p(y|\lambda, m) \quad (6)$$

This is simply the maximum value of the objective function attained by EM. The most likely model is the one with the largest

log-evidence. This enables Bayesian model selection. Model comparison rests on the likelihood ratio of the evidence for two models. This ratio is the Bayes factor B_{ij} . For models i and j

$$\ln B_{ij} = \ln p(y|m=i) - \ln p(y|m=j) \quad (7)$$

Conventionally, strong evidence in favor of one model requires the difference in log-evidence to be about three or more (Penny et al., 2004). Under the assumption that all models are equally likely *a priori*, the marginal densities $p(y|m)$ can be converted into the probability of the model given the data $p(m|y)$ (by normalizing so that they sum to one over models). We will use this probability to quantify model comparisons below.

Slice timing in fMRI

Data for most fMRI studies are acquired using single-shot echo planar imaging (EPI) techniques (Schmitt et al., 1998). The advantage of EPI lies in its speed. After an initial radio-frequency excitation of a single slice, the plane is encoded and all data are acquired in about 20–100 ms. One slice is encoded as a matrix whose elements correspond to the measured MR signal within a voxel (volume element). Voxels typically cover a brain volume of roughly $3 \times 3 \times 3$ mm³. Depending on slice thickness, one requires around 30–50 slices to cover the whole brain. The acquisition time for a slice is about 50–150 ms, including the wait time that is required to achieve the long echo times (TE) necessary for BOLD contrast. For instance, a typical sequence, at 1.5 T, would have a TR of 4.32 s with 48 slices. We refer to the vector encoding the time of slice acquisitions as ‘slice timings’.

Slice-timing correction

A potential slice-timing difference of several hundred milliseconds should be explained by the model without altering inference on (interesting) model parameters, e.g., the induced connections B^i . For a block design with a short TR (1.7 s) Friston et al. (2003) showed that DCM is indeed robust to slice-timing differences of up to a second. A widely used procedure to remove slice-timing differences involves temporal realignment to a single acquisition time (Henson et al., 1999). This alignment, usually referred to as ‘slice-timing correction’, is sometimes used as a pre-processing step, prior to DCM (and other analyses). We assess the usefulness of slice-timing correction in comparison to the extended model below. In this paper we implemented the correction as a sinc interpolation with the Matlab command ‘interp1’ (Matlab 7.01, MathWorks).

The slice timing for a given location

Here we describe how the slice timings for a given DCM are retrieved. The data consist of one representative time-series per area. These are usually taken as the first eigenvariate of a voxel cluster, centered on the spatial maximum of an SPM. The rationale for summarizing the response over voxel clusters is that BOLD effects are dispersed over several millimeters.

The timings of all slices can be derived from the parameters of the EPI sequence. The next step is to identify the acquisition slices that encompass each area. For raw (reconstructed) data this is trivial, because one coordinate axis coincides with the acquisition order. However, spatial normalization changes the orientation of the image and can introduce a nonlinear transform (Ashburner and




Fig. 3. Three-area network used for generating the first synthetic data set. ‘Sensory’ input either conformed to a block design (shown) or an event-related design. See text for further details.

Friston, 1999, 2005). To accommodate this one needs the (forward) mapping from the original image to the normalized space. The original slice numbers can be computed by applying the inverse of this mapping to the areas coordinates in standard space. We refer the reader to instructive Matlab code which performs this inverse mapping (see ‘get_orig_coord.m’; <http://www.fil.ion.ucl.ac.uk/spm/software/spm5/>). Having determined the slice from which each area was sampled one can then assign the corresponding slice timing to that area in the vector $T=[T_1, \dots, T_N]$.

Simulations

Synthetic data

In this section, we validate our model in terms of conditional inference and model comparison, using simulated data; in which the true parameters and models are known. We used two synthetic data sets. We generated the first from a system that was used by Friston et al. (2003) to validate the original DCM. The second was generated by a three-area network used in a visual attention study (Buchel and Friston, 1997; Friston et al., 2003). To generate synthetic data we selected some model parameters (described below), and generated data using the generative model described in Eqs. (1), (2) and (3). The errors ϵ_{ij} (Eq. (3)) were identically and independently distributed (i.i.d.). For all our simulations, we used a signal-to-noise ratio (SNR) of five (see Discussion). The routine that generates synthetic data is available as ‘spm_dcm_create.m’ under <http://www.fil.ion.ucl.ac.uk/spm/software/spm5/>.

With the first data set we show the impact of slice-timing misspecification on parameter (i.e., connectivity) estimation. The second data set is a typical example of modulating connectivity between a sensory area and some higher area and was used to replicate and generalize the conclusions based on the analysis of




Fig. 5. First set of simulations: posterior means (solid lines) and confidence intervals (grey area) for the modulation of the intrinsic connection $A1 \rightarrow A2$. The true value was 0.4 (dashed line). The format is the same as in Fig. 4. However, we pooled over eight iterations of the same simulation to enhance important differences between models. The confidence intervals are computed as the average over the individual confidence intervals.

the first. In brief, we show the best results obtain when slice-timing information enters the model. This was corroborated by an analysis (see below) of the corresponding real data (Buchel and Friston, 1997).

Simulations: first study

The first system is a three-area network with two inputs (Fig. 3). Area $A1$ receives ‘sensory’ input and has forward connections to




Fig. 4. First simulations: posterior means (solid lines) and 90% confidence intervals (grey area) for the intrinsic connection $A1 \rightarrow A2$. The true value was 0.8 (dashed line). We show four cells, each analyzed with three methods: in each cell, the left-hand plot shows the result of the extended method, the middle the original, and the right-hand ‘slice-timing correction’. The x-axes refer to the slice timing of the 2nd area (measured from the onset of a scan), ranging from 0 to TR seconds. The other two areas were always acquired at 1/2 TR seconds. Note that the confidence intervals are very small and hardly visible.




Fig. 6. First set of simulations: model comparison. For each cell and slice timing we plot the probability of the extended model relative to the original model, under the assumption that all models are equally likely a priori (i.e., flat priors on the models). A high probability means that the extended model is very likely in comparison to the original model. We added a threshold of 0.95 to help visual interpretation. For the derivation of the probabilities see Eq. (7) and text thereafter.




Fig. 7. Three-area network used for generating the second set of synthetic data. The connections and their modulations were based on estimates from a real single subject data set (see main text).

A2 and **A3**. There are reciprocal connections between **A2** and **A3**. Additionally, the self-connectivity of **A3** is modulated by ‘sensory’ input. This ‘contextual’ input is a series of exponential decays modulating the forward connection from **A1** to **A2** (top left of Fig. 3).

With this model, we computed the posterior means of the parameters under a $3 \times 2 \times 2$ factorial design: the levels of the first factor corresponded to different inversion schemes, while the remaining two factors corresponded to different ways of generating synthetic data; in terms of experimental design³ and TR:

1. *Factor 1: method*
 - * original model (without slice timings)
 - * extended model (with slice timings)
 - * slice-timing correction (with slice timings)
2. *Factor 2: design*
 - * event-related design
 - * block design
3. *Factor 3: TR*
 - * short TR (1.7 s)
 - * long TR (4.0 s)

We chose these factors to test the hypothesis that event-related studies, with a long TR (i.e., potentially large slice-timing difference), would benefit most from the inclusion of slice timing

³ Note that we do not analyze this factorial design with a classical ‘analysis of variance’ (ANOVA), but use a Bayesian analysis following Eqs. (4), (5), (6) and (7).

during model inversion. Our reasoning was that event-related inputs generate rather transient effects, which make timing more important. For block designs, we expected less dependence on temporal sampling, because the responses are longer lasting. Furthermore, robustness with short TRs has been established already (Friston et al., 2003). We derived an event-related design by replacing the blocked ‘sensory’ input (Fig. 3) by an event-related input. The event-related input had an inter-stimulus interval of 3.37 s, using a random jitter of -0.5 to 0.5 s (uniformly distributed), and 40% null events. We chose 1.7 and 4 s as short and long TRs to emulate typical repetition times.

For each cell of our multifactorial simulations, we varied the temporal sampling of area **A2** from 0 to TR seconds, in 16 steps. The other two areas **A1** and **A3** were always acquired at the same time ($1/2$ TR). In other words, we simulated data sets, for which we sampled **A1** and **A3** in the middle of each scan, while **A2** was sampled at different times (ranging from the top to the bottom of the volume). The maximum slice time differences between **A2** and **A1/A3** were $-1/2$ and $1/2$ TR.

To show the effects of the slice timing on the posterior density of an interesting parameter, we selected one representative connection (**A1** \rightarrow **A2**) and looked at the posterior mean (i.e., conditional expectation, see Eq. (4) and thereafter) of its fixed connectivity and its modulation by experimental input. We observed qualitatively similar results for other connections.

Fig. 4 shows the posterior means for the intrinsic connection **A1** \rightarrow **A2** as a function of slice timing. The true value was 0.8. The most striking finding is that, for all TRs, designs and sampling differences, inversion of the slice-timing model delivers veridical estimates. In contrast, inverting the original model returns posterior

means with large deviations from the true value. This is most apparent for the event-related design (both short and long TR) for which the posterior means deviate substantially (>50%) from their true value. This is consistent with our conjecture that event-related designs are more susceptible to timing errors. For the block design and long TR, we observed deviations of around 25% from the true value for slice-timing differences as small as half a second. Note that Friston et al. (2003) looked exclusively at block designs with short TR. Under these models, we confirm the robustness of the original model, with respect to slice timing.

Slice-timing correction proved a reliable pre-processing method for short TRs. In Fig. 4, the estimates are mostly veridical and improved over the original method, for both event-related and block designs. However, with the long TR, there are large biases for the event-related design. For the block design, the estimates are much improved over the original model, but still worse than for the extended model. The difference in performance between short and long TR can be explained by the nature of the correction: this approach uses a simple interpolation device, which only works when the sampling frequency is sufficiently high. The difference in the improvement, between event-related and block designs can be explained by noting that experimental variance in block designs is deployed in lower frequencies. These components are interpolated more accurately in relation to the higher frequencies of event-related designs.

Fig. 5 shows the results for the modulation of the $A1 \rightarrow A2$ connection. Again, the extended model gives veridical results for all cells (true value: 0.4). Due to large variability, we plotted the posterior means, pooled over eight replications of the simulations. For the extended model, we obtain veridical estimates, except for some small positive bias. For the original model, the estimates mirror the findings for the $A1 \rightarrow A2$ connectivity (Fig. 4): For all designs, there is a mismatch between true values and estimates, with the least error in the short block design. The slice timing correction actually performs well, with the notable exception for the event-related design with long TR.

In Fig. 6, we show the results of a pair-wise model comparison between the original and extended model for each cell of our multifactorial design and each slice-timing difference.⁴ This shows the probability of the extended DCM under flat priors over both models. For the long TR, the relative probability of the extended model, is nearly always 100% except when the slice-timing differences are very small (<250 ms). For the short TR and the block design, model comparison does not find strong evidence for the extended model, except for large slice-timing differences of around 1/2 TR. For the short TR and event-related design, model comparison suggests that for differences exceeding 300 ms, the extended model supervenes.

Simulations: second study

The second simulation was based on a DCM of an fMRI study of attention to visual motion (Buchel and Friston, 1997; Friston et al., 2003). We consider this data set as representative for models of bottom-up and top-down modulations, from and to primary sensory areas (Fig. 7). A primary visual area (V1) receives ‘photic’ input and has reciprocal connections with the motion-sensitive area V5. ‘Motion’ input modulates the forward

connection from V1 to V5. Area V5 has reciprocal connections with superior parietal cortex (SPC). ‘Attention’ input modulates the top-down connection from SPC to V5. For a detailed description of the experiment, see Buchel and Friston (1997). The data were acquired on a 2-T Magnetom VISION (Siemens, Erlangen) whole-body MRI system equipped with a head volume coil. $T2^*$ -weighted fMRI images (TE=40 ms; 90 ms/image; 64×64 pixels [19.2×19.2 cm]) were obtained with echo planar imaging (EPI) using an axial slice orientation. The volume acquired covered the whole brain except for the lower half of the cerebellum and the most inferior part of the temporal lobes (32 slices; slice thickness 3 mm, giving 9.6 cm vertical field of view). The effective repetition time was 3.22 s. Slices were acquired starting from the top of the head in a contiguous fashion, such that the primary visual cortex was covered by the last slices of each volume. The pre-processing included spatial normalization and smoothing.

We simulated a single-subject data set. We set the slice timings to the last slice for all three areas. Taking the posterior means from the real data (see below) as true parameters, we generated a series of data sets, which had different slice timings for the second area V5. This was varied from 0 to 3.22 s in 16 steps. We estimated parameters using the original and the extended model. In Fig. 8A we show the posterior means and confidence intervals for a single




Fig. 8. Second set of simulations. Left: extended model, right: original model. (A) Posterior means (solid lines) and confidence intervals (grey area) for the intrinsic connection $V5 \rightarrow SPC$. The true value was 0.58 (dashed line). (B) Posterior means (solid lines) and confidence intervals (grey area) for the modulation of the intrinsic connection $SPC \rightarrow V5$. The true value was 0.32 (dashed line). Both plots: the x-axes refer to the slice timing of the V5 area, ranging from 0 to TR seconds. The other two areas were always acquired at TR seconds (i.e., with the last slice).

⁴ Note that we cannot include ‘slice-timing correction’ in the comparison, because the temporal alignment changes the data.

connection. The extended model leads to veridical estimates. As with the first simulation, the posterior means of this connection's modulation, for both the original and extended model, are close to the true value (Fig. 8B). This confirms that DCM is robust to slice-timing errors, with respect to induced or bilinear effects. In Fig. 9, we show that the probability of the extended model (relative to the original) as a function of simulated slice-timing differences. When the original model assumes an early (erroneous) slice timing for V5, the probability that the extended model is better is 100%. This is strong evidence in favor of the extended model.

Model inversion using real data

In this section we present an analysis of the single-subject data used to provide the parameters for the second simulation of the previous section. These data have also been used for DCM validation purposes (Friston et al., 2003) and are available from <http://www.fil.ion.ucl.ac.uk/spm/data/>.

We specified extended models for all combinations of 8 different slice timings for each of the three areas (ranging from 0 to 3.22 s in 8 steps of 0.40 s). This resulted in $8^3=512$ different models. Each of these models places the three areas V1, V5 and SPC into a different combination of slices. For each model, we computed its log-evidence. The true (known) slice timings of V1, V5, and SPC were (2.07, 1.89, 0.81 s). The slice timings of the first two areas V1 and V5 have by far the biggest impact on model evidence. In Fig. 10 we plot the model evidence as a function of the slice timing of V1 and V5, averaged over slice timings of area SPC. The average log-evidences range from -1731 to -1645 . The greatest log-evidence was achieved when the areas were sampled at 3.22, 1.84 and 1.38 s. For this model, the slice timing of V5 was identical to its true value, and the SPC slice timing was close to its true value. The mismatch for V1 is most likely due to a late onset of the V1 BOLD response, which is well captured by a slice timing of 3.22 s. Clearly, slice timing of area V1 is the most important determinant of model evidence. This is probably due to the relative magnitude of the response from each area. The activity in V1 varied between -10 to 10 (percent whole brain mean), while the other two areas lay between -2 and 2 (data not shown).




Fig. 9. Second set of simulations, model comparison. For each slice timing of area V5, ranging from 0 to TR seconds, we show the probability of the extended model relative to the original model.




Fig. 10. Real data set: log-evidence as a function of slice timing in the first two areas V1 and V5, averaged over slice timing of the third area SPC. The higher the model evidence, the better the model.

In Fig. 11, we show the posterior means of the two modulation parameters for motion and attention under different slice timings (again averaged over SPC timings). As expected, the parameters vary systematically with slice timing. Note that inversion with the original model corresponds to inversion of the extended model





Fig. 11. Real data set: posterior means of the two modulation parameters as a function of slice timing in the first two areas V1 and V5, averaged over slice timing of the third area SPC.

with slice timing fixed at the last slice (3.22 s) for both areas. Therefore, the variation in parameter estimates about this value reflects their quantitative dependence on the third level of the model and the slice timings actually used. In this case, the posterior probability that the bilinear parameters were greater than zero was greater than 99%; i.e., the inference does not change with different slice timings.

Because the objective function in Eq. (6) is exactly the same as our approximation to the log-evidence, the series of model inversions summarized above could also be construed as a search for the most likely slice-timing parameters (i.e., those that optimize the objective function). This speaks to the possibility of making slice timing a free parameter for each area, which can be estimated using observed data. Although this would be easy to implement, we have chosen to treat the slicing-timing parameters as fixed, because they are known quantities. In addition, our results stress that slice timings are, of course, correlated with the area-specific onset of the BOLD response. These are in turn correlated with connectivity parameters. To minimize conditional dependencies between non-interesting parameters (BOLD and slice timings) and coupling parameters, we recommend fixing known parameters like slice timings. In summary, we present evidence, using real data, that slice-timing information is an important model parameter.

Discussion

Our analyses indicate that event-related designs are the most susceptible to failures in modeling slice timing correctly. Our simulations suggest that even with a short TR of less than 2 s and a small difference in slice timing (between areas) of 0.5 s, one can get large deviations from the true connectivity parameters, for the fixed coupling and their modulations parameters (Figs. 4 and 5). Model comparisons show that including slice-timing results in better models (Fig. 6). For block designs, the original model is more robust to slice-timing differences, especially for short TRs (<2 s). The extended model always furnishes veridical estimates, for both design types, and both TRs.

From our simulations, we find that slice-timing correction (i.e., alignment of data in time) is quite useful. The notable exception is the event-related design with a long TR. For such a design, the error in the connectivity and modulation estimates is quite large, and we recommend using the extended model for these designs.

Our results were derived under a rather high SNR of five. We repeated all the simulations for a low SNR of two and found qualitatively equivalent results (data not shown).⁵ However, we did note an increased variability in the individual estimates and inflated posterior variances. Also, for some connections, we noticed a (slight) systematic deviation from the true value, for all approaches (due to the use of shrinkage priors). However, this bias was small in comparison to the variability of estimates over simulations.

A situation in which we do not recommend using this procedure is when the data were acquired using an interleaved EPI sequence. An interleaved sequence acquires slices, not in a contiguous fashion, but ‘jumps’ between slices such that neighboring slices are acquired at different times. The issue with this sequence is that all pre-processing steps (motion correction, normalization and especially spatial smoothing) average data from neighboring slices,

which are sampled at different times (usually separated by 1/2 TR). This renders the activity of each area a mixture of data from different slice timings. It is unlikely that the extended model is useful for these data. Instead, we suggest using slice-timing correction before normalization and smoothing. One can then invert the extended model with slice timing for all areas set to the reference slice selected during ‘correction’.

Conclusion

We have presented an extension to dynamic causal modeling for functional magnetic resonance imaging data. Our analyses of synthetic and real data suggest that slice-timing differences between areas should and can be modeled. The extended model gives veridical estimates and higher model evidences, as compared to the original formulation. Furthermore, we can drop the former limitation that slice-timing differences between areas must not exceed 1 s. This makes DCM applicable to a much wider range of fMRI data with large slice-timing differences and a range of TRs.

Software note

We have implemented the extended model in the most recent version of Statistical Parametric Mapping (SPM5), see <http://www.fil.ion.ucl.ac.uk/spm/>. In practice, to use the extended model, one performs the following steps:

1. Omit ‘slice-timing correction’
2. Retrieve the acquisition slice for each area and specify DCM
3. Estimate parameters (as usual)

Acknowledgments

The Wellcome Trust funded this work. We thank Felix Blankenburg and Klaas Stephan for their helpful discussions.

References

- Ashburner, J., Friston, K.J., 1999. Nonlinear spatial normalization using basis functions. *Hum. Brain Mapp.* 7, 254–266.
- Ashburner, J., Friston, K.J., 2005. Unified segmentation. *NeuroImage* 26, 839–851.
- Buchel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI. *Cereb. Cortex* 7, 768–778.
- Buxton, R.B., Uludag, K., Dubowitz, D.J., Liu, T.T., 2004. Modeling the hemodynamic response to brain activation. *NeuroImage* 23 (Suppl. 1), S220–S233.
- Friston, K.J., 2002. Bayesian estimation of dynamical systems: an application to fMRI. *NeuroImage* 16, 513–530.
- Friston, K.J., Buchel, C., 2000. Attentional modulation of effective connectivity from V2 to V5/MT in humans. *Proc. Natl. Acad. Sci. U. S. A.* 97, 7591–7596.
- Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. *NeuroImage* 19, 1273–1302.
- Henson, R., Buchel, C., Josephs, O., Friston, K.J., 1999. The slice-timing problem in event-related fMRI. *Human Brain Mapping Conference* 1999.
- Kiebel, S.J., David, O., Friston, K.J., 2006. Dynamic causal modelling of evoked responses in EEG/MEG with lead field parameterization. *NeuroImage* 30 (4), 1273–1284.

⁵ Note that the input data are averaged over a cluster of voxels, which increases their SNR.

- Mechelli, A., Price, C.J., Friston, K.J., Ishai, A., 2004. Where bottom-up meets top-down: neuronal interactions during perception and imagery. *Cereb. Cortex* 14, 1256–1265.
- Noppeney, U., Price, C.J., Penny, W.D., Friston, K.J., 2006. Two distinct neural mechanisms for category-selective responses. *Cereb. Cortex* 16 (3), 437–445.
- Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004. Comparing dynamic causal models. *NeuroImage* 22, 1157–1172.
- Schmitt, F., Stehling, M.K., Turner, R., 1998. *Echo-Planar Imaging*. Springer.
- Stephan, K.E., Penny, W.D., Marshall, J.C., Fink, G.R., Friston, K.J., 2005. Investigating the functional role of callosal connections with dynamic causal models. *Ann. N. Y. Acad. Sci.* 1064, 16–36.