NeuroImage 58 (2011) 442–457

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Generalised filtering and stochastic DCM for fMRI

Baojuan Li [a,c], Jean Daunizeau [a,b], Klaas E. Stephan [a,b], Will Penny [a], Dewen Hu [c], Karl Friston [a,*]

[a] The Wellcome Trust Centre for Neuroimaging, University College London, Queen Square, London WC1N 3BG, UK
[b] Laboratory for Social and Neural Systems Research, Institute of Empirical Research in Economics, University of Zurich, Zurich, Switzerland
[c] College of Mechatronic Engineering and Automation, National University of Defense Technology, Changsha, Hunan 410073, PR China

ARTICLE INFO

Article history:
Received 17 September 2010
Revised 10 December 2010
Accepted 31 January 2011
Available online 17 February 2011

Keywords:
Bayesian
Filtering
Dynamic causal modelling
fMRI
Free energy
Dynamic expectation maximisation
Random differential equations
Neuronal

ABSTRACT

This paper is about the fitting or inversion of dynamic causal models (DCMs) of fMRI time series. It tries to establish the validity of stochastic DCMs that accommodate random fluctuations in hidden neuronal and physiological states. We compare and contrast deterministic and stochastic DCMs, which do and do not ignore random fluctuations or noise on hidden states. We then compare stochastic DCMs, which do and do not ignore conditional dependence between hidden states and model parameters (generalised filtering and dynamic expectation maximisation, respectively). We first characterise state-noise by comparing the log evidence of models with different a priori assumptions about its amplitude, form and smoothness. Face validity of the inversion scheme is then established using data simulated with and without state-noise to ensure that DCM can identify the parameters and model that generated the data. Finally, we address construct validity using real data from an fMRI study of internet addiction. Our analyses suggest the following. (i) The inversion of stochastic causal models is feasible, given typical fMRI data. (ii) State-noise has nontrivial amplitude and smoothness. (iii) Stochastic DCM has face validity, in the sense that Bayesian model comparison can distinguish between data that have been generated with high and low levels of physiological noise and model inversion provides veridical estimates of effective connectivity. (iv) Relaxing conditional independence assumptions can have greater construct validity, in terms of revealing group differences not disclosed by variational schemes. Finally, we note that the ability to model endogenous or random fluctuations on hidden neuronal (and physiological) states provides a new and possibly more plausible perspective on how regionally specific signals in fMRI are generated.

Introduction

This paper is about stochastic dynamic causal modelling of fMRI time series. Stochastic DCMs differ from conventional deterministic DCMs by allowing for endogenous or random fluctuations in unobserved (hidden) neuronal and physiological states, known technically as system or state-noise (Riera et al., 2004; Penny et al., 2005; Daunizeau et al., 2009). In this paper, we look more closely at the different ways in which stochastic DCMs can be treated. Deterministic DCMs provide probabilistic forward or generative models that explain observed data in terms of a deterministic response of the brain to known exogenous or experimental input. This response is a generalised convolution of the exogenous input (e.g. the stimulus functions used for defining design matrices in conventional fMRI analyses). In contrast, stochastic DCMs allow for fluctuations in the hidden states, such as neuronal activity or hemodynamic states like local perfusion and deoxyhemoglobin content. These fluctuations can be regarded as a result of (endogenous) autonomous dynamics that are not explained by (exogenous) experimental inputs. This state-noise can propagate around the system and, potentially, can have a profound effect on the correlations among observed fMRI signals from different parts of the brain. In this work, we ask whether it is possible to model endogenous or random fluctuations and still recover veridical estimates of the effective connectivity that mediates distributed responses. In particular, we compare and contrast DCMs with and without stochastic or random fluctuations in hidden states and explore variants of stochastic DCMs that make different assumptions about the conditional dependence between unknown (hidden) states and parameters.

Dynamic causal modelling (DCM) refers to the inversion of state-space models formulated with differential equations. Crucially, this inversion or fitting allows for uncertainty about both the states and parameters of the model. To date, DCMs for neuroimaging time series have been limited largely to deterministic DCMs, where uncertainty about the states is ignored (e.g., Friston et al., 2003). These are based on ordinary differential equations and assume that there are no random variations in the hidden neuronal and physiological states that mediate the effects of known experimental inputs on observed fMRI responses. In other words, the only uncertainty arises at the point of observation, through measurement noise. However, many

* Corresponding author at: The Wellcome Trust Centre for Neuroimaging, Institute of Neurology, Queen Square, London WC1N 3BG, UK. Fax: +44 207 813 1445.
E-mail address: k.friston@fil.ion.ucl.ac.uk (K. Friston).

1053-8119/$ – see front matter © 2011 Published by Elsevier Inc.
doi:10.1016/j.neuroimage.2011.01.085

studies suggest that physiological noise due to stochastic fluctuations in neuronal and vascular responses need to be taken into account (Biswal et al., 1995; Krüger and Glover, 2001; Riera et al., 2004). Recently, there has been a corresponding interest in estimating both the parameters and hidden states of DCMs based upon differential equations that include state-noise. Examples of these have been in the DCM literature for a while (e.g., Friston, 2008; Daunizeau et al., 2009). Early pioneering work in this area focussed on multivariate autoregression and state-space models formulated as difference equations (Riera et al., 2004; Valdes-Sosa, 2004; Penny et al., 2005; Valdés-Sosa et al., 2005). Riera et al. (2004) used stochastic differential equations to model hemodynamic responses in fMRI data, and estimated the underlying states and parameters from BOLD responses using a local linearisation innovation method. Penny et al. (2005) used difference equations to furnish a bilinear state-space model for fMRI time series and estimated its parameters and states using expectation maximisation (EM). This work was extended by Makni et al. (2008), who used a Variational Bayes inversion scheme that allowed for priors over model parameters and enabled model comparison (Penny et al., 2004). More recently, Daunizeau et al. (2009) introduced a general variational Bayesian approach for approximate inference on nonlinear models based on stochastic differential equations. In their recent work, Sotero et al. (2009) used the innovation method to invert a biophysical generative model of fMRI, which included both physiological and observation noise.

This paper deals with models based on random differential equations rather than stochastic differential or difference equations. This affords a model of state-noise that is not restricted to Wiener processes or Markovian assumptions. Furthermore, we will consider DCMs that comprise a network of regions (see also Valdés-Sosa et al., 2005), instead of the single regions considered previously (Penny et al., 2005; Makni et al., 2008). Our work in this area has focused on schemes that simplify the inversion problem, using various assumptions about the posterior or conditional density on unknown quantities in the model. Usually this density is assumed to have a Gaussian form. This is known as the Laplace approximation. In addition to this assumption, schemes based upon variational Bayes assume that the states and parameters (and any hyperparameters governing the amplitude of random noise) are conditionally independent. This is known as the mean-field approximation. Each set of conditionally independent quantities induces a separate optimisation step in the variational inversion scheme. For deterministic DCMs there are only two unknown quantities, the parameters and the hyperparameters. These are optimised by maximising a variational (free-energy) bound on the model log evidence in two steps. These are usually described as expectation and maximisation steps in variational EM schemes (Friston et al., 2003). Stochastic DCMs include a new set of unknown variables, namely, the hidden states. This introduces a third (dynamic) step, leading to schemes like dynamic expectation maximisation (DEM; Friston et al., 2008). Recently, we have developed a simpler and more general scheme called generalised filtering (GF; Friston et al., 2010) that dispenses with the (mean-field) conditional independence assumption. In this paper, we examine the utility and validity of modelling uncertainty about hidden states and the impact of conditional independence assumptions implicit in the difference between DEM and GF. We will show that estimates of effective connectivity (parameter estimates) from fMRI data are relatively robust to these fluctuations. Furthermore we demonstrate the potential usefulness of generalised filtering over its mean-field variant (DEM), when making inferences about differences in coupling among brain regions.

This paper comprises four sections. In the first, we present an illustrative application of generalised filtering to the same fMRI data set (attention to motion) that we have used previously to demonstrate DCM using EM (Friston et al., 2003; Stephan et al., 2008) and DEM (Friston et al., 2008). This section serves to illustrate the nature of the GF scheme and the results it produces. Our focus here will be on estimates of hidden neuronal and physiological states causing data and how their estimation affects inference on the parameters we are interested in (effective connectivity). Having established that it is possible to recover estimates of both parameters and states, the second section turns to the nature of noise or fluctuations in the hidden states. This section uses model comparison to search over models with different hyperpriors on the amplitude, form and smoothness of noise. In the third section, we turn to face validity and ensure that the accuracy of parameter estimates is robust to the introduction of state-noise. We generated data with and without state-noise (using the conditional parameter estimates from the first section) and fitted stochastic (GF) and deterministic (EM) DCMs. Using the conditional density on parameters and models, we then assessed the ability of each DCM to distinguish between data that were generated with and without state-noise and the impact of false assumptions about state-noise on parameter estimates. In the final section, we turn to construct validity and apply DCM to empirical data from an fMRI study of (clinical) group differences. Our focus here was on the conditional estimates of effective connectivity from EM, DEM and generalised filtering. Our objective in these analyses was to see if the deterministic and mean-field assumptions (implicit in EM and DEM) improved or subverted the ability of the estimators to distinguish between groups (under the assumption that group differences exist), in terms of their functional architectures (i.e. effective connectivity). We discuss the implications of our findings in the discussion, paying special attention to endogenous brain activity in dynamic causal modelling.

Stochastic DCM

In this section, we reanalyse an old data set that has been used extensively in demonstrating connectivity analyses over the years. These data were acquired during an attention to visual motion paradigm and have been used to illustrate psychophysiological interactions, structural equation modelling, multivariate autoregressive models, Kalman filtering, variational filtering, EM and DEM (Friston et al., 1997; Büchel and Friston, 1997, 1998; Friston et al., 2003, 2008; Harrison et al., 2003; Stephan et al., 2008). Here, we revisit questions about the generation of distributed responses by analysing the data using conventional deterministic DCMs (EM), stochastic DCMs under the mean-field approximation (DEM) and generalised filtering (GF). The mathematical details of these schemes are described in a series of technical papers (e.g., EM: Friston et al., 2007; DEM: Friston et al., 2008; GF: Friston et al., 2010). In this paper, we focus on the products of these schemes and how they differ from each other. One interesting thing that we will see is that modelling endogenous fluctuations allows one to infer neuronal and physiological states explicitly. This provides a different perspective on how to model brain dynamics, which we will return to in the discussion. We will first describe the data and then review comparative analyses, under the three different schemes.

Empirical data

Data were acquired from a normal subject at 2 T using a Magnetom VISION (Siemens, Erlangen) whole-body MRI system, during a visual attention study. Contiguous multi-slice images were obtained with a gradient echo-planar sequence (TE = 40 ms; TR = 3.22 s; matrix size = 64 × 64 × 32, voxel size 3 × 3 × 3 mm). Four consecutive 100 scan sessions were acquired, comprising a sequence of 10 scan blocks of five conditions. The first was a dummy condition to allow for magnetic saturation effects. In the second, Fixation, subjects viewed a fixation point at the centre of a screen. In an Attention condition, subjects viewed 250 dots moving radially from the centre at 4.7 ° per second and were asked to detect changes in radial velocity. In No