NeuroImage 49 (2010) 3099–3109

Contents lists available at ScienceDirect

NeuroImage

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

Review

Ten simple rules for dynamic causal modeling

K.E. Stephan [a,b,*], W.D. Penny [b], R.J. Moran [b], H.E.M. den Ouden [c], J. Daunizeau [a,b], K.J. Friston [b]

[a] Laboratory for Social and Neural Systems Research, Institute for Empirical Research in Economics, University of Zurich, Blümlisalpstr. 10, 8006 Zurich, Switzerland
[b] Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London, WC1N 3BG, UK
[c] Donders Institute for Brain, Cognition and Behaviour, Centre for Cognitive Neuroimaging, P.O. Box 9101, 6500 HB Nijmegen, The Netherlands

ARTICLE INFO

Article history:
Received 3 September 2009
Revised 1 November 2009
Accepted 9 November 2009
Available online 12 November 2009

Keywords:
Effective connectivity
DCM
Bayesian model selection
BMS
Model evidence
Model comparison
Bayes factor
Nonlinear dynamics
fMRI
EEG
MEG
Synaptic plasticity

ABSTRACT

Dynamic causal modeling (DCM) is a generic Bayesian framework for inferring hidden neuronal states from measurements of brain activity. It provides posterior estimates of neurobiologically interpretable quantities such as the effective strength of synaptic connections among neuronal populations and their context-dependent modulation. DCM is increasingly used in the analysis of a wide range of neuroimaging and electrophysiological data. Given the relative complexity of DCM, compared to conventional analysis techniques, a good knowledge of its theoretical foundations is needed to avoid pitfalls in its application and interpretation of results. By providing good practice recommendations for DCM, in the form of ten simple rules, we hope that this article serves as a helpful tutorial for the growing community of DCM users.

Contents

Introduction . . . 3099
Know what is "causal" about dynamic causal models . . . 3100
Know your hypothesis and how to test it . . . 3101
Use Bayesian model selection as a first step . . . 3103
Motivate model space carefully . . . 3103
Choose an appropriate method for group-level inference on model structure . . . 3104
Know what you can and cannot do with Bayesian model selection . . . 3104
Choose an appropriate method for group-level inference on parameters . . . 3105
Optimize experimental design and data acquisition . . . 3106
Use anatomical information and computational models to refine your DCMs . . . 3107
Report the modeling approach and results in detail . . . 3108
Summary . . . 3108
Acknowledgments . . . 3108
References . . . 3108

Introduction

Over the last two decades, neuroimaging analyses have become progressively refined and sophisticated. For example, there has been a trend away from the analysis of manually defined regions of interest to whole-brain analyses; from classical frequentist statistics to Bayesian hypothesis testing; and, most recently, efforts to construct

* Corresponding author. Laboratory for Social and Neural Systems Research, Institute for Empirical Research in Economics, University of Zurich, Switzerland. Fax: +41 44 6344907.
E-mail address: k.stephan@iew.uzh.ch (K.E. Stephan).

1053-8119/$ – see front matter
doi:10.1016/j.neuroimage.2009.11.015

mechanistic models of brain function. A representative of the latter is dynamic causal modeling (DCM), a generic approach for inferring hidden (unobserved) neuronal states from measured brain activity. DCM was introduced in 2003 for fMRI data (Friston et al., 2003) and made available as open-source software within the Statistical Parametric Mapping (SPM) software. The mathematical basis and implementation of DCM for fMRI have since been refined and extended repeatedly (Friston et al., 2007; Kiebel et al., 2007; Marreiros et al., 2008; Stephan et al., 2008, 2007c). Dynamic causal models (DCMs)[1] have also been implemented for a range of measurement techniques other than fMRI, including electroencephalography (EEG), magnetoencephalography (MEG), and local field potentials (LFPs) obtained from invasive recordings in humans or animals, both in the time domain (Daunizeau et al., 2009b; David et al., 2006; Kiebel et al., 2006) and frequency domain (Chen et al., 2008; Moran et al., 2007, 2008, 2009; Penny et al., 2009).

DCMs are generative models of brain responses, which provide posterior estimates of neurobiologically interpretable quantities such as the effective strength of synaptic connections among neuronal populations and their context-dependent modulation. They are defined by five key features. First, DCMs are dynamic, using (linear or nonlinear) differential equations for describing (hidden) neuronal dynamics. Second, they are causal in the sense of control theory, that is, they describe how dynamics in one neuronal population cause dynamics in another and how these interactions are modulated by experimental manipulations or endogenous brain activity. Third, DCMs strive for neurophysiological interpretability. Fourth, they use a biophysically motivated and parameterized forward model to link the modeled neuronal dynamics to specific features of measured data (for example, regional hemodynamic time series in fMRI or spectral densities of electrophysiological data). Fifth, DCMs are Bayesian in all aspects. Each parameter is constrained by a prior distribution, which reflects empirical knowledge about the range of possible parameter values, principled considerations (e.g., certain parameters cannot have negative values) or a conservative attitude (e.g., "shrinkage" priors that express the assumption that coupling parameters are zero). Furthermore, Bayesian inversion not only provides posterior densities for each model parameter but also yields an approximation to the log model evidence, which is used to compare alternative DCMs of the same data.

Since their introduction in 2003, DCMs have gradually become part of mainstream neuroimaging analysis techniques. At the time of submitting this article (September 2009), the database PubMed listed more than 100 published papers on DCM. Its applications have concerned a wide range of domains in cognitive neuroscience, including language (Allen et al., 2008; Bitan et al., 2005; Leff et al., 2008; Noppeney et al., 2008; Schofield et al., 2009), motor processes (Eickhoff et al., 2005; Grefkes et al., 2008; Grol et al., 2007), vision and visual attention (Fairhall and Ishai, 2007; Haynes et al., 2005; Mechelli et al., 2003; Sonty et al., 2007), memory (Smith et al., 2006), perceptual decision making (Stephan et al., 2007b; Summerfield et al., 2006; Summerfield and Koechlin, 2008), and learning (den Ouden et al., 2009; Garrido et al., 2008, 2009). Given the relative complexity of DCM, compared to conventional analyses, many colleagues in the neuroimaging community have expressed an interest in a tutorial-like guide that addresses some of the most common questions about the theoretical foundations and empirical applications of DCM. This article represents an attempt to provide such a tutorial. It follows a recent tradition in the neuroimaging literature, inspired by the popular "10 simple rules" series in *PLoS Computational Biology* (Bourne, 2005), which has led to tutorial papers on, for example, voxel-based morphometry (Ridgway et al., 2008) and on reporting results from mass-univariate analyses (Poldrack et al., 2008).

In this article, we provide some generic "good practice" recommendations that address key conceptual and methodological issues in applying DCM to fMRI, EEG, MEG, or LFP measurements. Omitting any equations, we have tried to keep these recommendations as straightforward as possible. The suggestions made in this article should not be mistaken as dogmatic rules; instead, they are meant to provide guidelines for those users who are new to dynamic system theory, Bayesian statistics, and model selection procedures. Furthermore, some of the points below, such as the section on causality, are not concrete rules but outline the conceptual foundations of DCM. We anticipate that some of these guidelines and their underlying concepts may change over the forthcoming years, as both the theoretical foundations as well as the implementation of DCMs are progressively refined.

Know what is "causal" about dynamic causal models

Causality in DCM is based on control theory (Friston, 2009): causal interactions among hidden state variables[2] (e.g., specific aspects of neuronal population activity) are expressed by differential equations, which describe (i) how the present state of one neuronal population causes dynamics (i.e., rate of change) in another via synaptic connections and (ii) how these interactions change under the influence of external perturbations (i.e., experimental manipulations) or endogenous brain activity. The differential equations endow the system with memory such that future states are influenced by current states; the coupling parameters (rate constants) determine the speed of these influences. The ensuing coupling is influenced by where and when the system is subject to external perturbations; i.e., sensory inputs driving activity in specific neuronal populations or modulatory inputs that render the strength of coupling context-sensitive. In other words, causality in DCM does not only rely on temporal precedence but also takes into account when and where the system is perturbed by external influences.

An equivalent perspective is to interpret the state equation of a given DCM as encoding a particular causal structure–function relationship (Stephan, 2004). This is because the state equation of a given DCM prescribes explicitly how system dynamics arises from system structure: it specifies formally how neuronal state changes, induced by external inputs, propagate both in space (i.e., according to the system's connectivity structure) and in time (i.e., how current states influence future states). Therefore, changing the pattern of external inputs or the connectivity structure in a given DCM leads to different predictions about the spatiotemporal pattern of measured system responses. By simulating data from models with specified causal mechanisms, it is straightforward to assess whether, for a given level of observation noise, DCM is capable of correctly inferring these mechanisms. This has been done using both the same neuronal equations as in DCM (e.g., Stephan et al., 2008) and using independently designed large-scale biophysical models of spiking neurons (Lee et al., 2006). Perhaps even more convincingly, several animal studies using independent techniques such as invasive recordings and microdialysis demonstrated that DCM can successfully infer neuronal processes from BOLD responses and field potentials, respectively (David et al., 2008; Moran et al., 2008).

Critically, the hidden neuronal states give rise to noisy observations through a forward mapping (e.g., neurovascular coupling in fMRI). This transform is crucial for inferring causal interactions, particularly when it is nonlinear and may differ across brain regions, as is the case in fMRI (David et al., 2008; Stephan et al., 2004). Therefore, in contrast to Granger causality (Granger, 1969), causality in DCM does not describe interactions among the observations themselves. Instead, DCM aims to infer interactions among hidden neuronal states that cause noisy observations through a (possibly nonlinear and spatially variable) mapping.

[1] We use the acronym DCM both to refer to the general approach (dynamic causal modeling) and to refer to the instantiation of a specific dynamic causal model.

[2] The term "hidden state variables" refers to time-varying properties of systems that cannot be observed directly.