Dynamic causal modelling of lateral interactions in the visual cortex

D.A. Pinotsis *, D.S. Schwarzkopf, V. Litvak, G. Rees, G. Barnes, K.J. Friston

The Wellcome Trust Centre for Neuroimaging, University College London, Queen Square, London WC1N 3BG, UK

ABSTRACT

This paper presents a dynamic causal model based upon neural field models of the Amari type. We consider the application of these models to non-invasive data, with a special focus on the mapping from source activity on the cortical surface to a single channel. We introduce a neural field model based upon the canonical microcircuit (CMC), in which neuronal populations are assigned to different cortical layers. We show that DCM can disambiguate between alternative (neural mass and field) models of cortical activity. However, unlike neural mass models, DCM with neural fields can address questions about neuronal microcircuitry and lateral interactions. This is because they are equipped with interlaminar connections and horizontal intra-laminar connections that are patchy in nature. These horizontal or lateral connections can be regarded as connecting macrocolumns with similar feature selectivity. Crucially, the spatial parameters governing horizontal connectivity determine the separation (width) of cortical macrocolumns. Thus we can estimate the width of macro columns, using non-invasive electromagnetic signals. We illustrate this estimation using dynamic causal models of steady-state or ongoing spectral activity measured using magnetoencephalography (MEG) in human visual cortex. Specifically, we revisit the hypothesis that the size of a macrocolumn is a key determinant of neuronal dynamics, particularly the peak gamma frequency. We are able to show a correlation, over subjects, between columnar size and peak gamma frequency — that fits comfortably with established correlations between peak gamma frequency and the size of visual cortex defined retinotopically. We also considered cortical excitability and assessed its relative influence on observed gamma activity. This example highlights the potential utility of dynamic causal modelling and neural fields in providing quantitative characterisations of spatially extended dynamics on the cortical surface — that are parameterised in terms of horizontal connections, implicit in the cortical micro-architecture and its synaptic parameters.

© 2012 Elsevier Inc. All rights reserved.

Introduction

This work combines neural field models — that model the activity of layers of cells in cortical patches — with a Bayesian framework for optimising model parameters — known as Dynamic Causal Modelling (DCM). This combination allows one to address questions about lateral cortical interactions in terms of optimal models and model parameters. DCM has been applied extensively to fMRI and electrophysiological data (David et al., 2006; Friston et al., 2003; Penny et al., 2004) and has been used recently to model spatiotemporal dynamics on the cortical surface (Pinotsis et al., 2012). Combining generative models with Bayesian optimisation techniques enables one to characterise the functional architectures that generate empirical data. In this paper, we show that DCM can disambiguate between alternative spatiotemporal models of cortical activity — using Bayesian model comparison — and furnish quantitative explanations of observed responses, in terms of the biophysical properties of lateral cortical connections. This work focuses on two classes of biophysical models describing mesoscale brain

activity: neural mass and field models. Neural field models describe how hidden neuronal states (such as the average depolarisation of a neural population or layer) evolve over both space and time. In contrast, neural mass models only characterise dynamics over time, under the assumption that all the neurons of a particular population are located at (approximately) the same point.

In previous work (Pinotsis et al., 2012), we considered the relationship between neural mass and field models and showed that field models can be reduced to neural masses by applying shrinkage priors to spike propagation times — such that lateral interactions were effectively instantaneous. We introduced a DCM that provides an explicit model of spatially extended cortical activity that allows one to make inferences about key parameters controlling the topographic distribution of cortical activity using LFP data, like the extent of lateral cortical connections and the conduction velocity of spike propagation. We were able to show that including spatial parameters enables one to explain effects that other models — such as neural mass models — attribute to variations in temporal parameters, like synaptic rate constants. Dynamic causal models based on neural fields enable one to characterise the propagation of activity on the cortex and provide a formal understanding of the mechanisms generating spatiotemporal responses. In our previous work above, we considered

* Corresponding author at: The Wellcome Trust Centre for Neuroimaging, Institute of Neurology, Queen Square, London, UK WC1N 3BG, UK. Fax: +44 207 813 1445.
E-mail address: d.pinotsis@ucl.ac.uk (D.A. Pinotsis).

model optimisation under the assumption that one could measure local signals (local field potentials) that are sensitive to all spatial frequencies. In this paper, we consider neural field models of non-invasive (EEG or MEG) data. This development allows one to make inferences about the nature of lateral interactions in cortical sources, without spatially resolved measurements.

Spectral responses of neural fields

The characterisation of electrophysiological signals depends upon models of how they are generated in source space and how the resulting (hidden) neuronal states are detected by sensors. At the source level we consider a model based upon a canonical microcircuit that allows one to separate the sources of forward and backward connections in cortical hierarchies. In terms of the mapping from source to sensor space — we use a conventional lead field formulation that is expanded in terms of spatial basis functions. As in previous work, we focus on the modelling of power spectra. There is a long history of modelling steady-state (or ongoing) activity spectra, associated with neural fields, usually in models of the whole cortex, e.g., (Jirsa, 2009). Robinson (2006) has developed a neurophysiologically grounded model of corticothalamic activity, which reproduces many properties of empirical EEG signals; such as the spectral peaks seen in various sleep states and seizure activity. Technically, the spectra summarising the response of cortical sources can be defined in terms of transfer functions, mapping endogenous neuronal fluctuations to observed responses (Freeman, 1972; Nunez, 1995; Robinson et al., 2001; Robinson et al., 2003). In Pinotsis and Friston (2011), we derived the transfer function — and an expression for the spectral responses — for a source described by a classical neural field equation, while in Pinotsis et al. (2012) we extended our approach to a cortical source that comprises multiple layers. This allowed us to model the spectral activity of cortical fields as measured on the cortical surface. Here, we pursue a similar approach and generate the corresponding spectral responses, as measured by non-invasive sensors. These predictions are generated in an efficient manner that exploits the nature of the mapping from sources to sensors in EEG and MEG.

The resulting scheme can be regarded as inverting or fitting population models of the Amari type, using real data and Bayesian model inversion. Previous work in a similar vein includes the use of Kalman filters to develop estimation schemes for both neural mass (Riera et al., 2007; Valdes et al., 1999) and neural field models of a single population (Galka et al., 2008; Schiff and Sauer, 2008). In a related approach, Daunizeau et al. (2009)) replaced the standard dipole source — used in neural mass models — with the principal Fourier mode of a neural field, for the particular case of exponentially decaying synaptic density over the cortical surface. Finally (Markounikau et al. (2010) used a combination of linear and nonlinear optimisation methods to invert a two-layered neural field model of voltage-sensitive dye data, describing inhibitory and excitatory populations (without conduction delays). The neural field model considered here has four layers and is based on canonical cortical microcircuitry that accounts for several aspects of local cortical computations in theoretical neurobiology. This model provides an extension of the well known Jansen and Rit neural mass model and incorporates conduction delays associated with the propagation of neuronal spikes.

Lateral interactions and neural fields

Modelling lateral interactions with neural field models has a long history. Pioneering work was introduced in papers by Amari, Wilson and Cowan, Grossberg and colleagues (Amari and Arbib, 1977; Amari and Takeuchi, 1978; Grossberg and Levine, 1975; Wilson and Cowan, 1973). These developments can be traced back to the work of physicists in the 19th century — such as Helmholtz and Mach — on visual perception. The first neural field models considered spontaneous

pattern formation, by analysing the steady-state behaviour of underlying field equations: for example, Wilson and Cowan developed a treatment of Turing instabilities in the context of neural fields (Wilson and Cowan, 1972). In a similar vein, Grossberg initiated a line of work on shunting interactions — via nonlinearly coupled inputs — and considered limits as the solutions approached steady state (Grossberg, 1973). At about the same time, Amari proved that systems of neural fields typically approach steady-state, in which some parts remain active, thus providing a metaphor for short-term memory.

From an anatomical viewpoint, the functional specialisation of visual (and auditory) cortex is reflected in its patchy or modular organisation — in which local cortical structures share common response properties. This organisation may be mediated by a patchy distribution of horizontal intrinsic connections that can extend up to 8 mm, linking neurons with similar receptive fields: see, e.g., Angelucci and Bressloff (2006), Burkhalter and Bernardo, (1989), and Wallace and Bajwa (1991). The existence of patchy connections in different cortical areas (and species) has been established with tracer studies in man, macaque and cat: Burkhalter and Bernardo (1989), Stettler et al. (2002) and (Wallace and Bajwa (1991), respectively. It has been shown that such connections can have profound implications for neural field dynamics: see Baker and Cowan (2009). The precise form of such connections may be explained by self-organisation under functional and structural constraints; for example, minimising the length of myelinated axons to offset the cost of transmitting action potentials (Cherniak, 1994; Wen and Chklovskii, 2008). Generic constraints of this sort have been used to motivate general principles of connectivity; namely, that evolution attempts to optimise a trade-off between metabolic cost and topological complexity (Bassett et al., 2010). In short, visual and auditory cortices can be characterised by a patchy organisation that is conserved over the cortex and which allows for both convergence and divergence of cortical connections. Synaptic densities can then be approximated by isotropic distributions with an exponential decay over the cortical surface. In this work, we use a combination of patchy and isotropic distributions, using connectivity kernels with non-central peaks to model sparse intrinsic connections in cortical circuits that mediate both local and non-local interactions. In other words, we consider models in which neurons receive signals both from their immediate neighbours and remote populations that share the same functional selectivity (Pinotsis and Friston, 2011). We focus on the particular problem of identifying the parameters of lateral connections within a bounded cortical patch or manifold, as measured at a distance, with an array of non-invasive (MEG or EEG) sensors.

This paper comprises three sections. The first reviews a canonical neural mass model based on anatomical data and theoretical constraints from the theory of predictive coding. The second section describes a neural field model based upon this canonical microcircuitry. Our focus here is on equipping the resulting neuronal model with an electromagnetic forward model to predict responses in non-invasive sensors. In the third section, we use Bayesian model comparison to adjudicate between various formulations of the ensuing neural field model (DCM) and establish its construct validity by optimising parameters pertaining to GABAergic concentrations, that have been shown to correlate with the peak gamma frequency of steady-state activity (Muthukumaraswamy et al., 2009).

A canonical model of cortical activity

This section introduces the neuronal field model used in subsequent sections on dynamic causal modelling. The particular neuronal model of source activity used here is based upon a refinement of conventional (convolution-based) neuronal models that explicitly model the neuronal sources of forward and backward connections in cortical hierarchies — these are the superficial and deep pyramidal cell