NeuroImage 80 (2013) 330–338

Contents lists available at SciVerse ScienceDirect

NeuroImage

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

Biophysical network models and the human connectome

Mark W. Woolrich [a,b,*], Klaas E. Stephan [c,d,e]

[a] Oxford Centre for Human Brain Activity (OHBA), Warneford Hospital, Oxford, UK
[b] FMRIB Centre, University of Oxford, Oxford, UK
[c] Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich & ETH Zurich, Switzerland
[d] Wellcome Trust Centre for Neuroimaging, University College London, UK
[e] Laboratory for Social and Neural Systems Research (SNS), University of Zurich, Switzerland

ARTICLE INFO

Article history:
Accepted 24 March 2013
Available online 6 April 2013

Keywords:
Biophysical model
Networks
Connectivity
Connectome
Bayes
DCM
Generative embedding
Bottom-up model
MEG
EEG
FMRI
Diffusion
Multi-modal

ABSTRACT

A core goal of human connectomics is to characterise the neural pathways that underlie brain function. This can be largely achieved noninvasively by inferring white matter connectivity using diffusion MRI data. However, there are challenges. First, diffusion tractography is blind to directed connections, or whether a connection is expressed functionally. Second, we need to be able to go beyond the characterization of anatomical pathways, to understand distributed brain function that results *from* them. In particular, we need to characterise effective connectivity using functional imaging modalities, such as FMRI and M/EEG, to understand its context-sensitivity (e.g., modulation by task), and how it changes with synaptic plasticity. Here, we consider the critical role that biophysical network models have to play in meeting these challenges, by providing a principled way to conciliate information from anatomical and functional data. They also provide biophysically meaningful parameters, through which we can better understand brain function. In a translational setting, well-validated models may shed light on the mechanisms of individual disease processes.

Introduction

The Human Connectome Project (HCP) is collecting a wealth of state-of-the-art data across a range of imaging modalities; in particular, functional MRI, magnetoencephalography (MEG) and diffusion MRI (Van Essen et al., 2012). Arguably, each of these modalities could be used to obtain a different connectome (Behrens and Sporns, 2012; Friston, 2011). For example, FMRI can provide us with a map of functional/effective connectivity (Biswal et al., 1995; Friston, 2011), and diffusion MRI with a map of anatomical white matter connectivity (Basser et al., 1994, 2000; Behrens et al., 2007). But it is not immediately clear how these different modalities can be related. Or indeed, what governing principles we should use to resolve differences among these connectomes.

A useful unifying principle is that the *anatomical connectome* underlies (is necessary for) the functional connectome. This idea has previously been expressed in terms of the concept of "connectional fingerprints" (Passingham et al., 2002), which postulates that the functional profile of any given cortical area depends on the structural pattern of its incoming and outgoing connections. More recently, it has been demonstrated that models of effective connectivity are improved when formally integrating quantitative anatomical information (Stephan et al., 2009b). This is *not* to say that structural connectivity is a sufficient or complete description of connectivity, but rather that function depends on structural connectivity. Armed with this perspective, we can start to consider how we might characterise the architecture of a multi-modal human connectome.

Firstly, diffusion MRI data is not a panacea even for inferring vanilla anatomical connectivity — it has blind spots. Comparison with invasive studies in non-human primates reveal that current tractography approaches can suffer from both false positive and false negative results (Behrens and Sporns, 2012). Functional connectomics can help inform the anatomical connectome when structural information is missing, or is inaccurate. Arguably, the best way to do this is through *network models*; because these can embody both the structural and functional architecture, and allow information from the different modalities to be fused in a mathematically principled way.

We also want to go be able to go beyond the characterization of the anatomical connectome, to understand the brain function that rests upon it. Patterns of functional network connectivity emerge as the result of neuronal interactions taking place on this anatomical skeleton (Deco et al., 2011; Honey et al., 2007). The best way we can understand these patterns is by using biophysical network models that combine models of the anatomy with dynamic models of neuronal interactions. In principle, if we had sufficient knowledge of the system, including

* Corresponding author at: Oxford Centre for Human Brain Activity (OHBA), Warneford Hospital, Oxford, UK.

1053-8119/$ – see front matter
http://dx.doi.org/10.1016/j.neuroimage.2013.03.059

information at the level of synapses, we could predict context sensitive coupling or the dynamics one might expect to see, i.e. how functional pathways are modulated depending on task or cognitive set. This requires biophysical models that can be informed by experimentally controlled context, and can represent connection strengths as a function of activity (non-linear postsynaptic effects) or time (synaptic plasticity).

FMRI has a key role to play, and has already been successful in mapping functional networks, for example, using resting state data (Beckmann et al., 2005). However, evidence from MEG suggests that these network interactions are likely underpinned by oscillatory activity in particular frequency bands (Brookes et al., 2011a; Hipp et al., 2012). Understanding these oscillations and the biophysical models that underpin them will provide unique and important insights into the function of the brain. For example, one possibility is that long-range connectivity may be mediated by synchronisation of oscillatory activity (Fries, 2005). To illuminate these models, direct measures of neural activity at high temporal resolution are needed, such as can provided by the increasingly relevant modality of MEG.

Eventually, the understanding we can gain about the physiology of network dynamics can be used to elucidate the mechanisms of aging and disease in a clinical setting. Through approaches like generative embedding (Brodersen et al., 2011), we can investigate disease mechanisms; e.g. by looking at the population variability in certain biophysical model parameters. By moving closer to the actual mechanisms of brain function, this approach should ultimately be more sensitive and more interpretable than descriptive or normative approaches.

This paper will focus primarily on systems-level biophysical network models of non-invasive neuroimaging data at the macroscopic, whole brain, system level. There is a particular focus on biophysical network models of function. However, we will also consider models of anatomical connectivity, particularly with regards to informing models of network dynamics using the anatomical connectome.

Functional biophysical network models

In recent years there has been a noticeable move away from the spatial mapping of task related activity towards inferring brain connectivity. This is motivated by the idea that connectivity brings us closer to the distributed mechanisms of brain function.

A popular approach to looking at connectivity in functional data has been to look at measures of statistical dependency, otherwise known as *functional connectivity*. This includes approaches such as partial or full correlation (Smith et al., 2011). In FMRI, these correlations are typically computed on the raw BOLD time series, whereas in MEG (due to the non-zero lag correlations) correlations are typically computed over band-limited power timeseries, particularly in the alpha and beta bands (Brookes et al., 2011a; Hipp et al., 2012).

Full correlation simply corresponds to:

$$C = yy^T/(N-1) \quad (1)$$

where y is the $N \times T$ matrix of functional neuroimaging data for N brain regions and T time points, and y is normalised to have zero mean and unit variance for each brain region.[1] Notably, full correlation cannot distinguish between direct and indirect connections, whereas *partial correlation* can — at least to some extent.

Partial correlation refers to the correlation between two timeseries, after each has been adjusted by regressing out other variables (e.g., activity in other brain regions/network nodes). An efficient way to estimate the full set of partial correlations is via the inverse of the covariance matrix (Marrelec et al., 2006). Under the constraint that this matrix is expected to be sparse, regularisation is often applied, for example, using the Lasso method (Friedman et al., 2008). Partial correlation has been advocated (Marrelec et al., 2006) as a good surrogate for structural equation modelling (SEM) (see the Appendix A for a description of SEM, and for a mathematical perspective relating partial correlation to SEM). While partial correlation does seem to improve the distinction between direct and indirect connections (Smith et al., 2011), it also introduces Berkson's paradox, where there can be artifactual negative correlations between brain regions.[2]

Functional connectivity measures such as full and partial correlation are popular measures because they are easy to compute, and have been shown to perform relatively well in FMRI network discovery (Smith et al., 2011). The most straightforward functional human connectomes are likely to be based on the partial or full correlation matrix from FMRI (or MEG) data using an appropriate parcellation.

The problem with functional connectivity

However, there are fundamental problems with the use of functional connectivity. For example, we have already commented on how existing methods for distinguishing direct from indirect connections can only be achieved at the expense of artifactual negative correlations (Berkson's paradox). Another important issue is the increasingly popular approach of using functional connectivity (often in the form of correlation) as a feature to predict or classify the group from which a particular subject was sampled (Craddock et al., 2009). The problem is that changes in functional connectivity; e.g., between conditions or between two population groups, can occur simply due to changes in signal-to-noise ratio, or due to changes in other parts of a wider network; even when there is no change in the effective connectivity between the two nodes. This issue has been demonstrated and documented elsewhere (Friston, 2011).

False positive connectivity can also be inferred if correlations caused by the measurement process itself are not accounted for. This includes erroneously inferring neuronal causality in FMRI data when the hemodynamic blurring is ignored, or inferring artifactual connectivity due to volume conductance (zero-lag spatial correlations) in source reconstructed M/EEG data (Schoffelen and Gross, 2009).

Functional connectivity also provides limited insight into the mechanisms of the dynamics that underlie brain activity, and does not directly provide biologically relevant information. The best way to overcome these limitations is to turn to *effective connectivity*.

Essentially, effective connectivity is an estimate of directed influence, inferred using a generative model, which, to at least to some extent, is grounded in bio-physiology. In other words, we need a biophysical model of the network.

Biophysical functional network models

Here, we consider a biophysical functional network model as a mathematical description of how we can generate measurements of brain activity (e.g. using MRI). This modelling assumes that we know the underlying dynamic interactions between, and within, different brain areas; and the stochastic or deterministic properties of the exogenous or endogenous fluctuations that drive the network. As such "biophysics" includes the neurophysiology of the brain, and the physics of the measurement device.

We can consider the complete biophysical network model of functional neuroimaging data as being decomposed into two main components. First, there is the neurophysiological model of the network's neuronal interactions (neuronal dynamics or state equations), which predicts the neuronal activity. Second, there is the forward model of the imaging measurement, which predicts the imaging data given the neuronal activity.

[1] Without normalizing to unit variance, C would correspond to the covariance matrix.

[2] Consider 3 nodes, where B depends on A and C, and where A and C are uncorrelated ($A \rightarrow B \leftarrow C$). By regressing B out of A and C (as in partial correlation), we induce an artifactual negative correlation between A and C.

Neuronal interaction models

For a plausible biophysical model of neuronal interaction, we consider that an entry-level requirement is that the model should be cyclic and connections should be able to be reciprocal (bi-directional). This is based on evidence from the analyses of large scale anatomical databases, where the majority of brain regions appear to be connected reciprocally (Markov et al., 2011). Sadly, this makes inference on these models much trickier, as it precludes the use of Directed Acyclic Graphs (DAGs), as they do not allow for reciprocal (bi-directional) connections.

We also know that the effect of one brain area on another cannot occur instantaneously, due to conduction delays. In other words, any biophysical network model needs to be causal. This leads us to the general framework of Dynamic Causal Models (DCM) pioneered by Friston and colleagues (Friston et al., 2003). These models express the interactions between brain regions using differential equations. These equations also allow for known external inputs (experimentally controlled perturbations) and can therefore model both "resting" brain activity and task- or stimulus-related responses:

$$\dot{x} = f(x, \theta, u) + e_x \tag{2}$$

where x is $P \times T_x$ matrix of P hidden neuronal states, $\dot{x}$ is its temporal derivative, at T_x time points, θ are the biophysical neuronal model parameters, u, are the known external inputs, and e_x is a stochastic neuronal noise or fluctuation. Note that the presence of e_x makes this equation a stochastic differential equation; a special case is when $e_x = 0$, which would correspond to a deterministic (ordinary) differential equation.

The inclusion of stochastic noise is necessary if these DCMs are to be used with "resting state" data. This is because models of resting data are not equipped with external inputs, u, and so neuronal activity can only be explained in terms of spontaneous (endogenous) neuronal fluctuations that can be modelled by the stochastic term (Daunizeau et al., 2009; Li et al., 2011).

Dynamic causal models can be thought of as comprising individual brain areas (or network nodes), whose directed connections are described by a connectivity matrix (normally designated as A). The dynamics within a brain region, or network node, often appeal to the notion of a *cortical micro-circuit (CMC)*. CMCs have been described as "functional modules that act as elementary processing units bridging single cells and systems levels" (Grillner and Graybiel, 2006). Each CMC can consist of one or more interacting sub-populations that also interact with the sub-populations in CMCs of other brain areas; see Fig. 1 for an illustrative example.

Examples

There is a continuum of models that could be considered at varying temporal and spatial scales, and with different amounts of biophysical realism and complexity. These models take us from abstract representations of (lumped) neuronal activity, through to full blown biophysical models of individual neuron dynamics and inter-neuron interactions. Some typical examples will now be described.

"Classic" fMRI DCM

Here, we consider the special case of the form of $f(x,\theta)$ used for modelling "resting" state fMRI data using DCM. This can easily be extended to include external inputs, u, (Friston et al. 2011). In this case, $f(x,\theta,u)$ corresponds to a linear state model:

$$\dot{x} = Ax + e_x \tag{3}$$

where A is the $P \times P$ matrix that encodes directed connectivity between P brain regions. In these sorts of models, neuronal dynamics within a node are described by a single neuronal state with self-inhibition (by having negative values on the diagonal of A). It is worth noting an observation from (Penny et al. 2004) that if we assume $\dot{x} = 0$, then this equation corresponds to a structural equation model $x = \overline{A}x + e_x$, where $A = \overline{A} - I$.

FMRI DCM has also been extended to a nonlinear model, which allows second order or multiplicative interactions between states to produce changes in target activity. This can be thought of in terms of the modulation of coupling between two regions by a third (Stephan et al. 2008).

Mean-field and neural mass models

Electrophysiological data (e.g. M/EEG) is a more direct neuronal measure, with high temporal resolution. In this case, the linear state model in Eq. (3) is far too simple; instead models at the timescale of neuronal dynamic interactions need to be considered. However, non-invasive electromagnetic data are not sufficiently resolved for models of individual neurons and so approximations to population dynamics are used. Thankfully, MEG and EEG measure the combined effect of populations of large numbers of neurons. These neuronal populations can be modelled using mean-field or neural mass approximations, in which the population behaviour is captured using probability distributions over the neuronal state variables (Deco et al., 2008).

The "neuronal populations" we are referring to typically correspond to sub-populations of similar neurons within a brain area, source or node. For example, a Wilson-Cowan node contains two populations, one inhibitory and one excitatory (Deco et al., 2009; Wilson and Cowan, 1972). These within-brain source populations, or sub-populations, interact with each other, and are coupled to sub-populations in other sources; see Fig. 1 for an illustrative example. Another example is M/EEG DCM, in which there are three sub-populations within each source, corresponding to excitatory pyramidal neurons, excitatory spiny stellate neurons, and inhibitory inter neurons (David et al., 2006).

But how do we model the population distribution of neuronal dynamics? Typically the population is represented using a single neural mass at the mean of this distribution, so-called *neural mass models*. However, one can also consider a mean-field model, which accounts for higher order statistics via stationary solutions of the Fokker–Planck equation (Deco et al., 2008). For example, (Marreiros et al., 2010) used this approach to model dynamics of covariances by using a Laplace approximate around the mean in DCM. These full mean field models only become relevant when the local dynamics are nonlinear — and the dispersion or [co]variances affect the mean and vice versa (Marreiros et al., 2010).

Conductance-based models

A particularly useful class of model is the conductance-based model (Morris and Lecar, 1981). This sort of model corresponds to a system of coupled nonlinear first order differential equations with state variables for the transmembrane potentials and for different channel conductances (Marreiros et al., 2010). Note that these conductance-based models are particularly attractive biophysical models to work with, because they are directly related to specific synaptic processes, and can be used to investigate experimental effects of altering specific neurotransmitters (Moran et al., 2011). These models are inherently non-linear (second-order) because they consider the interaction between voltage differences and conductances. While conductance-based models describe the behaviour of individual neurons, the Fokker–Planck approach can be used to translate them into equivalent neural mass, or mean field, models at the population level, as discussed in the last section (Marreiros et al., 2010).

Kuramoto oscillators

Coupled mean-field models of neurons tend to produce dynamics that are oscillatory. So an alternative strategy has been to circumvent the complexity of the neuronal model (albeit at the expense of biophysical interpretability of the parameters), and to model each brain area's population dynamics as a Kuramoto oscillator (Breakspear et al., 2010;

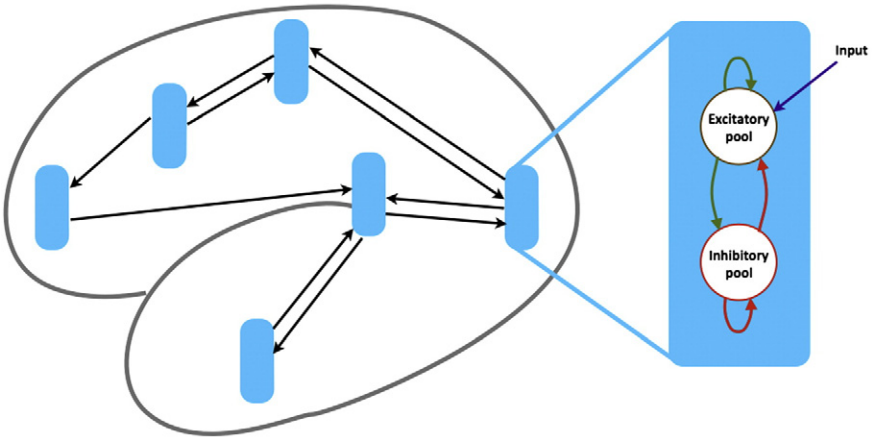

Fig. 1. Example of a biophysical network model. The model comprises individual brain areas or sources (or network nodes), whose connections are described by a connectivity matrix that encodes directed connectivity (black lines) between brain regions (blue patches). The local model that describes dynamics within a brain region – sometimes referred to as a *cortical micro-circuit* – can consist of one or more sub-populations. Here the local model is a Wilson–Cowan model, which contains two sub-populations: one inhibitory pool and one excitatory pool of neurons. These pools interact with each other (red, green arrows). The networks are driven by inputs into each of the nodes (blue arrow), which can be a combination of external stimuli and endogenous fluctuations.

Cabral et al., 2011; Shanahan, 2010). As in the full biophysical models, these phenomenological models can incorporate endogenous noise, and biophysical parameters such as conduction delays; and have been shown to simulate emergent dynamics, e.g. multistability, similar to those found in "resting" state neuroimaging data (Cabral et al., 2011).

The observation model

The observation model captures the biophysics of the imaging modality. This predicts the imaging data given the neuronal activity, x:

$$y = g(x, \alpha) + e_y \tag{4}$$

where $g(x,\alpha)$ describes the mapping between the neuronal activity and predicted imaging measurement (or modality) in question. This is a function of the hidden neuronal states, x, and some observation model parameters, α. Here, e_y is the observation noise. In the case of FMRI, $g(x,\alpha)$ could correspond to a model from MRI physics of how hidden physiological states (e.g. blood flow, volume) produce measurement signals (e.g. T2* effects). For example, in DCM for fMRI, this observation model is based upon the hidden states entailed by the so-called "Balloon model" (Buxton et al., 2004). This consists of coupled differential equations describing changes in vasodilatory signals, blood flow, blood volume and deoxyhemoglobin content (Friston, 2000; Stephan et al., 2007b). This part of the model for fMRI model is crucial for dealing with hemodynamic confounds (David et al., 2008) when inferring hidden neuronal states and parameters. In MEG, $g(x,\alpha)$ would typically correspond to the classical electromagnetic forward model — incorporating solutions to Maxwell's equation given the head and sensor array geometry (i.e. the lead field matrix) (Mosher et al., 1999).

Which model to use?

It is all very well emphasising the need for biophysical network models, but the benefits over functional connectivity will only be properly obtained with an appropriate choice of model — both in terms of the form of the model (e.g., neural mass versus Kuramoto oscillator) and its graphical structure (e.g., fully connected versus sparse). For example, a model that has missing nodes may be susceptible to exhibiting apparent changes due to changes in the wider network (full graph).

The "best" model to use in any given context will be an inevitable balancing act between inter-related considerations:

- Ease of inference (e.g. simpler models have less parameters and are generally easier to invert, in terms of numerics and convergence).
- Model evidence: a model with higher evidence (the probability of observing the measured data given the model) has a better trade-off between accuracy (fit) and complexity and is thus less prone to overfitting (Penny et al., 2004). A complementary notion is that of model generalisability as assessed by cross-validation (Strother et al., 2002).
- What are the data modalities being considered? (e.g., FMRI models do not necessarily require modelling of populations of neurons at fine temporal scales).
- What will the results be used for? What are the requirements on the output with regards to:
  - o Level of biophysical interpretability (e.g., conductance-based neural mass models allow modelling of neurotransmitter specific channels).
  - o Temporal and spatial scale.
  - o Diagnostic utility.

Modelling data transforms

There is also a choice as to the form of the data that are modelled. It is not necessary to always have a full generative biophysical model of the data as it comes off the imaging device. In other words, the data features one tries to model can sometimes be as important as the model itself. For example, should one be modelling the entire timeseries or the covariance matrix based on timeseries? One example, from electrophysiology is to separate the source reconstruction model of M/EEG data from the neuronal modelling. However, while this can allow for the use of source reconstruction approaches that cannot be readily expressed as a generative model (e.g. beamforming (Vrba and Robinson, 2000; Woolrich et al., 2011)), this can introduce problems such as volume conduction and spatial leakage (i.e. the existence of zero-lag correlations between brain areas as an artifact of the source reconstruction) that need to be accounted for in other ways.

Another example is that we do not necessarily need to model the full time-series in the time-domain. Instead, we can look to model the frequency (spectral) characteristics. This is done in DCM for steady state responses (Moran et al., 2009), where the temporally stationary frequency response of M/EEG or LFP data is modelled directly. Similarly in (Cabral et al., submitted for publication), the frequency profile of functional connectivity is matched between real MEG data, and a biophysical model of resting state networks using Kuramoto oscillators.

Models of "resting" state brain activity based on anatomical skeletons

Functional neuroimaging studies have found that the brain's "resting" state activity patterns can be decomposed into networks of brain areas with known task-related rules, such as the default mode, attention, and sensorimotor networks. These have not only been identified in FMRI data (Beckmann et al., 2005), but also now independently in MEG data (Brookes et al., 2011b).

There has been increasing interest in models of "resting state" data that make use of the information about white matter anatomical connectivity from diffusion MRI data {Cabral, 2011 #91;Deco, 2012 #252; Honey, 2009 #418;Ghosh, 2008 #435}. These so-called "bottom-up" models have used data from the Macaque connectivity database CoCoMac (Stephan et al., 2001), or tractography on human diffusion MRI data, to calculate an anatomical network matrix between all pairs of nodes. This provides information about the existence of an anatomical connection, and sometimes about the length of the connections (by virtue of the length of the tracts estimated from diffusion tractography). A dynamic model of the neuronal interactions among nodes is then added, which is constrained to interactions allowed by the anatomical skeleton.

Activity patterns simulated from these models have been shown to reproduce the spatiotemporal characteristics of real functional neuroimaging data. For example, they can simulate anti-correlated functional networks models using neural mass models at each node ({Ghosh, 2008 #435}, Deco et al., 2009, Honey et al., 2007). They can also simulate slow oscillatory spontaneous activity using, for example, dynamic models of Wilson–Cowan units (Deco et al., 2009), Kuramoto oscillators (Cabral et al., 2011), neural mass models (Honey et al., 2007), or spiking neuron models (Deco and Jirsa, 2012). Kuramoto oscillator models, which only oscillate locally in the gamma band, have also been shown to produce functional connectivity specifically in the alpha and beta bands (Cabral et al., submitted for publication), matching the observations made in real MEG data (Brookes et al., 2011b; de Pasquale et al., 2010; Hipp et al., 2012).

In general, these network dynamics arise as emergent properties that result from the nature of the anatomical connectivity, which underpins the different models. In this context, evidence is emerging about the relative importance of biophysical assumptions about the node dynamics, endogenous noise, or indeed conduction delays (Deco et al., 2011).

The HCP offers an exciting new resource for approaches that investigate emergent dynamics from structure. The state-of-the art diffusion MRI data (enhanced by new developments in tractography methods) will allow for more accurate anatomical skeletons to be used. Furthermore, the HCP will provide an extensive database of resting state and task data, both FMRI and MEG, along with the corresponding anatomical skeletons for the same individuals.

While providing an interesting frontier, these approaches are not without their potential problems. For example, one has to assume that the anatomical connectivity is known precisely; however, this cannot be known with great certainty for all node-pairs from diffusion MRI data. Missing anatomical links could cause the model to make spurious predictions. Arguably, the broad characteristics being investigated so far should be robust to moderate variations in the anatomical skeleton (and this robustness can be demonstrated numerically). However, this problem may become more acute as more detailed predictions are pursued. As discussed elsewhere in this paper, the ultimate solution is to have a generative model of both the functional and anatomical imaging data that can capture the inherent uncertainties.

Models of anatomical imaging data

As discussed earlier, one might hope to infer the underlying anatomical connectome, and then characterise the functional interactions that arise from these anatomical constraints. However, for this to be a complete framework, we need to augment the functional models discussed so far, with models of the anatomical imaging. Through probabilistic inference, this should then allow information to be appropriately garnered and integrated from all modalities.

While important progress has been made in local biophysical models of diffusion data, relatively little progress has been made in modelling biophysical tracts at the global, systems level. One exception to this is the approach of global tractography (Jbabdi et al., 2007; Yendiki et al., 2011), which models tracts using splines — and ties these (using a hierarchical model) to the local biophysical models of diffusion data (Behrens et al., 2007). This allows for top-down regularisation, which provides more robust tracking through low SNR, or regions of the white matter with crossing fibres. Furthermore, this model also provides parameters for the existence of tracts, for the end location of the tracts, and for the length of tracts. These global tract parameters are important, in as they are the kind of parameters that also appear in generative models of functional imaging data and functional networks. As illustrated in Fig. 2 this could provide a unified generative model, linked by common biophysical characteristics, which can predict both functional and anatomical imaging data.

Choosing the parcellation

A core goal of the HCP – to describe the connections between different brain areas – requires the specification of a set of brain areas between which to characterise the connectivity. This raises the question as to what makes for an appropriate parcellation for a given model and modality (or modalities). While this is a topic for other articles in this special issue, we consider here some of the key issues in the context of generative network modelling.

Anatomical atlases

A common approach is the use of atlases. For example, in (Deco et al., 2009; Honey et al., 2009) the authors use a pre-defined AAL atlas. However, this can be problematic. In (Smith et al., 2011), network simulations based on DCM for FMRI were used to demonstrate how node leakage can cause substantial errors in network discovery. This work illustrates the need for data-driven parcellations.

Data-driven parcellations

The most straightforward approach to data-driven parcellations is to use functional localisers from traditional brain mapping methods; e.g., general linear modelling of task FMRI data (Friston et al., 2003). However, this approach tailors the parcels (or regions of interest, ROIs) to those relevant in brain mapping, and it may not always be obvious which to choose for network modelling. Furthermore, when considering network modelling of resting state data, it is non-trivial to ensure that you have available the relevant task-FMRI data to provide sufficient brain coverage.

Emerging data-driven approaches use unsupervised learning approaches based on clustering. In particular, independent component analysis (ICA) has been shown at low dimensions (~25 components) to reliably identify the spatial maps of classic resting state networks, such as the default mode, attention, and sensory-motor networks. These have not only been identified using ICA with fMRI data, but also now independently with MEG data (Brookes et al., 2011b; Luckhoo et al., 2012) and in task data (Luckhoo et al., 2012). However, each of these components contains multiple brain areas. What we really need are nodes that correspond to individual brain areas (or sources), and then subsequently allow the network modelling to characterise the interactions among those nodes. For this, we will need to turn to higher dimensional parcellations (>50 components); at this scale the components will tend to include only single brain areas. This was used, for example, in (Smith et al., 2012) to find individual brain nodes from spatial (~150 dimensional) ICA.

Modelling of node parcellations

The parcellation approaches discussed thus far pre-calculate nodes (ROIs) and then impose them as 'knowns' in the subsequent network analysis. However, an attractive alternative is to parameterise the position and shape of the ROIs as part of the generative model. For example, in (Woolrich et al., 2009) multivariate normal distributions were used to model the nodes in fMRI DCM. The network model and the data itself will then drive the inferred position and shape of the ROIs.

Crucially, the position and shape of the ROIs for each brain region in the network can then be inferred, alongside connectivity parameters in the network model. Eventually, these position and shape parameters can be tied to the end points of global diffusion tractography models developed in (Jbabdi et al., 2007), in a symmetric analysis of fMRI and diffusion data, see Fig. 2.

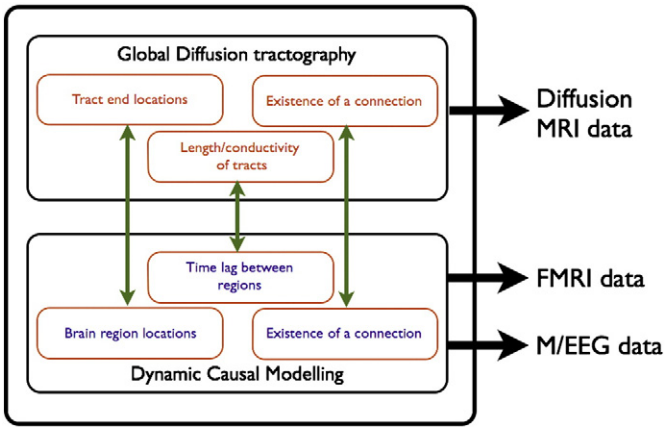

Fig. 2. Schematic of a combined biophysical model that predicts both anatomical and functional imaging data. This generative model can regarded as separate generative models for anatomical and functional modalities, linked probabilistically by common parameters (green arrows).

Estimation/inference on biophysical models

Biophysical models are typically used in one of two modes:

1) To predict emergent spatiotemporal characteristics observed in neuroimaging data.
2) As generative models that are inverted using (fitted to) neuroimaging data.

Although the boundaries between these two options can often get blurred, we will discuss them in turn.

Predicting emergent spatiotemporal characteristics

As discussed in the "Resting state activity" section this approach has been particularly prevalent in predicting the emergent spatiotemporal characteristics of spontaneous brain activity, through the use of underlying anatomical skeletons (Cabral et al., 2011; Deco and Jirsa, 2012; Honey et al., 2009). These approaches typically control broad network behaviour with a limited number of parameters. Using this approach, we can see qualitatively if the predictions made by the biophysical models match empirical phenomena. This comparison can be made more quantitative; e.g., by using correlation in time and/or space between the model predictions of, data-derived estimates of, the spatiotemporal characteristics.

However, there is often the nagging concern that any sufficiently complex model can be tweaked to match any empirical observation. While the use of biophysical modelling helps to restrict the tweaking to that which remains biophysically plausible, this can still remain an issue. Explicit consideration should be given to limiting the parameter space where possible. For example, in (Cabral et al., 2011), only two parameters were varied (a global scaling of the connectivity strength and of the time delays), and yet good prediction of real imaging phenomena was still obtained. These approaches are not restricted to "resting state" data. For example, in (Hunt et al., 2012) the temporal dynamics of neural activity was predicted in a decision making task, and shown to match real MEG data (albeit in a single brain area).

Full inference: Using Bayes

Biophysical generative models are a natural way for us to incorporate our understanding of the brain and of the neuroimaging measures to make predictions about what the data will look like. However, in practice we want to do the opposite. We want to be able to take acquired data and extract pertinent information about the brain (i.e., "infer" the model parameters). Bayesian statistics offers a complete framework to solve this problem, and also provides a framework in which we can do much more besides (Woolrich, 2012).

Within the Bayesian setting, the biophysical network model provides the data likelihood; i.e., the probability of the data given the model parameter values. Biophysical constraints (or regularisation) can be encoded using probabilistic priors, which keep parameters within biologically plausible ranges. Bayes rule can then be used to calculate the posterior distribution over the model parameters, *given* the neuroimaging data. What is more, Bayes provides a measure of the extent to which different biophysical models are supported by the evidence in the data, namely the Bayesian *model evidence* — this is the probability of the data given the model (averaging over uncertainty about the model parameters).

Bayesian inference in DCM

The best example of Bayesian inference on biophysical network models is DCM (Friston et al., 2003). DCM uses approximate Bayesian inference approaches to invert and compare DCMs. Approximate Bayesian inference usually uses the variational free energy as an approximation to the model evidence to provide a measure for model comparison (Penny et al., 2004; Stephan et al., 2007b). When combined with the Laplace assumption this is known as Variational Laplace. The combined use of Bayesian model comparison and DCM provides a framework for evaluating competing hypotheses about the architecture of networks. Recent developments also allow for random effects inference in heterogeneous groups of subjects (Stephan et al., 2009a) and for model comparisons across families of models to overcome problems with "model dilution" — when there are a large number of models to compare (Penny et al., 2010).

DCM has been used successfully in a wide variety of settings, including task and rest, fMRI and M/EEG data (Stephan et al., 2007a). However, we should be cautious about offering up Bayesian inference as a solution to all problems: We will now look at some of the outstanding problems in more detail.

Inaccuracy of inference

The most immediate problem is that the solution to Bayes rule to get posterior probability distributions is not analytically tractable for all practical biophysical network models. Either time-consuming numerical methods (e.g. Markov Chain Monte Carlo (MCMC) sampling) or posterior approximate approaches (e.g. Variational Bayes or Variational Laplace) are needed. These approaches are only as good as any of the inherent approximations allow, and do *not* guarantee that they explore the most probable areas of parameter space.

Network discovery

A particular challenge facing the community is the problem of network discovery. Network discovery is needed when we do not have a strong a priori hypothesis about the configuration of effective connections in the network, and instead want to "discover", or infer, the existence of connections from the data. This is a particular problem (although not exclusively) in "resting state" data, due to the large number of nodes that are likely to be engaged in the resting state. Each possible configuration of connections between these nodes can be thought of a candidate model, and so a huge set of models must be considered.

A pragmatic approach is the recent idea of post hoc Bayesian model comparison, in which the Bayesian model evidence for nested sub-models can be approximated from the inference on the full model — with all parameters being non-zero (Friston and Penny, 2011; Friston et al., 2011). While fast, this solution will not allow for exhaustive exploration of networks with large numbers of nodes, and heuristic search strategies will be required (e.g., greedy search Friston and Penny, 2011). See (Seghier and Friston, 2012) for recent developments along these lines that use the prior constraint that nodes (or modes) with independent dynamics (no functional connectivity) are unlikely to share effective connections.

What can we do with biophysical network models?

Since they are generative models (see above), a general trait of biophysical network models is their ability to allow for inference on hidden variables. So far, we have mainly referred to inference on the hidden parameters describing the effective connectivity between brain areas. However, this idea extends to other hidden variables; for example, those relating to network nodes for which we have little, or no, sensitivity in our measurements. In other words, even when we are not able to measure the activity in a network node directly, the inversion of a biophysical network model may enable us to infer it from activity changes in the rest of the network. Important examples include inferring the activity of deep structures from scalp electrophysiology, or the activity in brain areas not directly measured using intracranial electrodes (Moran et al., 2011; Marreiros et al., 2010; Garrido et al., 2012).

The potential of biophysical modelling to provide compact, mechanistically interpretable and quantitative summaries of neuronal systems is of great interest for clinical applications. This is particularly relevant for psychiatry, where we presently lack any objective diagnostic procedures — and where diagnostic classifications purely rest on symptoms, as opposed to pathophysiological mechanisms. If biophysically plausible models of non-invasive single-subject data can be established and validated, this would provide exciting opportunities for redefining psychiatric disease classifications; enable diagnostic procedures that rest on pathophysiology not symptoms, and perhaps even allow for precise treatment recommendations. While the idea of model-based inference on individual disease mechanisms is not new (Stephan et al., 2004), it is only over the last few years that advances in biophysical modelling (in particular with DCM) and machine learning have begun to explore the practical utility of this idea.

There are two main challenges that need to be addressed in order to unlock the potential of model-based diagnostics.

1) *Systematic validation*
   Systematic validation studies are needed that can evaluate a model's capacity to recover "ground truth" processes (face validity). This requires experimental studies, in which a characterised physiological process is induced (e.g., by drugs or stimulation techniques) and then we ask whether model inversion can recognise the known states or parameters of the system, from measured fMRI or electrophysiological data. Several studies of this sort have already been undertaken; mainly in the context of DCM. These include the demonstration that DCM can recover known changes in synaptic transmission following neurochemical modulation in rodents (Moran et al., 2008) or humans (Moran et al., 2011) and that it can track dose-dependent changes in excitation and inhibition, under varying levels of anaesthesia in rodents (Moran et al., 2011). However, we also require validation studies in patients. This is needed because even if we have a model that can capture known physiological processes, it still remains to be established whether quantitative measures of this process are useful (have predictive validity) for clinical decision-making. In other words, we need longitudinal patient studies, which examine whether model-based estimates can predict clinical trajectories and treatment responses.
2) *Clinical prediction on individuals*
   The second challenge is that we need methods for generating clinical predictions from model parameter estimates or structure for the individual. There are two ways of accomplishing this: the first is to use Bayesian model selection (BMS). This is possible whenever we can represent a particular disease state or mechanism by a particular model structure. In this case, we can simply fit all of these models to the individual patient's data and compute their relative model evidence. A nice example was provided by (Rowe et al., 2010) who showed that BMS, applied to fMRI data during a motor task, could distinguish between the presence and absence of dopaminergic medication in patients with Parkinson's disease. However, this approach is not always feasible because we often lack knowledge about the mapping between a particular clinical state and a specific model structure. In this case, a useful alternative is "*generative embedding*".

Generative embedding

Generative embedding uses estimates of model parameters, obtained from inverting a generative model, to inform the kernel of a discriminant classifier, such as a support vector machine. In neuroimaging, generative embedding has been pioneered by (Brodersen et al., 2011). Using groups of aphasic patients (due to stroke) and healthy controls, both measured with fMRI while listening to passive speech, they extracted fMRI signals from the early auditory system, which was not lesioned. They then tried to predict, subject by subject, whether or not there was a remote lesion (beyond the field of view). Using generative embedding (based on a six-area DCM of the early auditory system), almost perfect predictive accuracy (98%) was achieved. By comparison, conventional classification approaches operating on the BOLD signals themselves (or on functional connectivity estimates) were significantly less sensitive — with the best accuracies around 80%.

In addition to superior predictive accuracy, generative embedding also allows for mechanistic insights. This is because one can go back to model parameter space and examine which parameter combination drove the discrimination. In the case of (Brodersen et al., 2011), the predictive power of the model mainly rested on parameters representing transfer from the right to the left hemisphere. One can intuit that if generative embedding is based on models with more physiological interpretability than DCM for fMRI (e.g., conductance-based DCMs for electrophysiological data as described above), much more fine-grained insights into disease mechanisms may be in reach.

Finally, generative embedding can also be employed in an unsupervised way; e.g., by combining generative models of biophysical processes with clustering methods. This is particularly useful when dealing with spectrum disorders, such as schizophrenia or depression, where disease mechanisms are likely to vary considerably across patients and an unknown number of subgroups may exist. In ongoing work, Brodersen et al. (2011) have used this approach to examine a group of 40 schizophrenic subjects with DCM for fMRI (during a visual working memory task) and found evidence of three distinct subgroups, characterised by different prefrontal–parietal–visual network architectures. Crucially, the subgroups identified by these physiological differences also displayed significant differences in

clinical symptom profiles, information that was not available to the clustering procedure.

Summary

We have considered the key role that biophysical functional network modelling has to play in characterising the human connectome using multi-modal non-invasive neuroimaging data. Furthermore, network models provide the means by which we can understand the relationship between structural and functional connections; and can provide us with more sensitive and interpretable parameters — through which we can better understand normal and diseased brain function. The suggestion here is that there is great potential for future clinical applications, although there is much work to be done with regard to validation studies in animals, human pharmacological and patient studies. Given successful validation, techniques such as generative embedding could provide exciting future opportunities for predicting individual clinical outcomes and for detecting hitherto unknown subgroups in spectrum diseases.

Acknowledgments

The authors would like to thank Prof. Karl Friston for his extensive and invaluable comments on the paper, and also thank Faysal Kahn for useful discussions. MWW is funded by the Wellcome Trust, and supported by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre based at Oxford University Hospitals Trust Oxford University (the views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health). KES is supported by the René and Susanne Braginsky Foundation.

Appendix A

The use of the partial correlation matrix is normally motivated by noting that full correlations pick up indirect connections as well as direct connections, whereas the partial correlation for a pair of brain areas partials out all other brain areas, other than the pair in question. Here we offer an alternative perspective by relating it to a structural equation model (SEM). Like DCM, SEM's comprise a set of regions and a set of directed connections. However, there are limitations on structural model complexity and only instantaneous effects are modelled {Penny, 2004}:

$$y_t = \overline{A} y_t + e_t$$

where $\overline{A}$ is the $N \times N$ matrix of connections (or path coefficients) between N brain regions, and e_t is Gaussian noise at time t, $e_t \sim N(0,\sigma^2)$. Then:

$$y_t = \left(I - \overline{A}\right)^{-1} e_t$$

and so the $N \times N$ full covariance matrix of y, C_y, is given by

$$C_y \propto \left(I - \overline{A}\right)^{-1} \left(I - \overline{A}\right)^{-T}.$$

If we rearrange this and assume that $\overline{A}$ is symmetric, then we can see that the partial correlation matrix, $(C_y)^{-1}$, is related to $\overline{A}$ by:

$$\overline{A} = I - \sqrt{\left(C_y^{-1}\right)}$$

where we are using the matrix square root. This means that the off-diagonal elements of the (matrix square root of the) partial correlation matrix correspond to the negative values in the matrix of connections, $\overline{A}$, when that matrix is symmetric.

References

Basser, P.J., Mattiello, J., LeBihan, D., 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. J. Magn. Reson. B 247–254.
Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using DT-MRI data. Magn. Reson. Med. 625–632.
Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. Philos. Trans. R. Soc. Lond. B Biol. Sci. 1001–1013.
Behrens, T.E.J., Sporns, O., 2012. Human connectomics. Curr. Opin. Neurobiol. 144–153.
Behrens, T.E.J., Berg, H.J., Jbabdi, S., Rushworth, M.F.S., Woolrich, M.W., 2007. Probabilistic diffusion tractography with multiple fibre orientations: what can we gain? Neuroimage 144–155.
Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 537–541.
Breakspear, M., Heitmann, S., Daffertshofer, A., 2010. Generative models of cortical oscillations: neurobiological implications of the Kuramoto model. Front. Hum. Neurosci. 14.
Brodersen, K.H., Haiss, F., Ong, C.S., Jung, F., Tittgemeyer, M., Buhmann, J.M., Weber, B., Stephan, K.E., 2011. Model-based feature construction for multivariate decoding. Neuroimage 56 (2), 601–615.
Brookes, M.J., Hale, J.R., Zumer, J.M., Stevenson, C.M., Francis, S.T., Barnes, G.R., Owen, J.P., Morris, P.G., Nagarajan, S.S., 2011a. Measuring functional connectivity using MEG: Methodology and comparison with fcMRI. Neuroimage 1082–1104.
Brookes, M.J., Woolrich, M., Luckhoo, H., Price, D., Hale, J.R., Stephenson, M.C., Barnes, G.R., Smith, S.M., Morris, P.G., 2011b. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. Proc. Natl. Acad. Sci. U. S. A. 16783–16788.
Buxton, R., Uludag, K., Dubowitz, D., Liu, T., 2004. Modeling the hemodynamic response to brain activation. Neuroimage S220–S233.
Cabral, J., Hugues, E., Sporns, O., Deco, G., 2011. Role of local network oscillations in resting-state functional connectivity. Neuroimage 57 (1), 130–139.
Cabral, J., Luckhoo, H., Woolrich, M., Joensson, M., Mohseni, H., Baker, A., Kringelbach, M.L., Deco, G., 2013. Revealing underlying neural mechanisms of spontaneous MEG functional connectivity: oscillatory network interactions can lead to structured band-limited power fluctuations (submitted for publication).
Craddock, R.C., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S., 2009. Disease state prediction from resting state functional connectivity. Magn. Reson. Med. 1619–1628.
Daunizeau, J., Friston, K.J., Kiebel, S.J., 2009. Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. Physica D 2089–2118.
David, O., Kiebel, S.J., Harrison, L.M., Mattout, J., Kilner, J.M., Friston, K.J., 2006. Dynamic causal modeling of evoked responses in EEG and MEG. Neuroimage 1255–1272.
David, O., Guillemain, I., Saillet, S., Reyt, S., Deransart, C., Segebarth, C., Depaulis, A., 2008. Identifying neural drivers with functional MRI: an electrophysiological validation. PLoS Biol. 2683–2697.
de Pasquale, F., Della Penna, S., Snyder, A.Z., Lewis, C., Mantini, D., Marzetti, L., Belardinelli, P., Ciancetta, L., Pizzella, V., Romani, G.L., Corbetta, M., 2010. Temporal dynamics of spontaneous MEG activity in brain networks. Proc. Natl. Acad. Sci. U. S. A. 107 (13), 6040–6045.
Deco, G., Jirsa, V.K., 2012. Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. J. Neurosci. 3366–3375.
Deco, G., Jirsa, V.K., Robinson, P.A., Breakspear, M., Friston, K., 2008. The dynamic brain: from spiking neurons to neural masses and cortical fields. PLoS Comput. Biol. e1000092.
Deco, G., Jirsa, V., McIntosh, A.R., Sporns, O., Kötter, R., 2009. Key role of coupling, delay, and noise in resting brain fluctuations. Proc. Natl. Acad. Sci. U. S. A. 10302–10307.
Deco, G., Jirsa, V.K., Mcintosh, A.R., 2011. Emerging concepts for the dynamical organization of resting-state activity in the brain. Nat. Rev. Neurosci. 43–56.
Friedman, J., Hastie, T., Tibshirani, R., 2008. Sparse inverse covariance estimation with the graphical lasso. Biostatistics 432–441.
Fries, P., 2005. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends Cogn. Sci. 9 (10), 474–480.
Friston, K., 2000. Nonlinear responses in fMRI: the balloon model, Volterra Kernels, and other hemodynamics. Neuroimage 466–477.
Friston, K.J., 2011. Functional and effective connectivity: a review. Brain Connect. 13–36.
Friston, K., Penny, W., 2011. Post hoc Bayesian model selection. Neuroimage 2089–2099.
Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 1273–1302.
Friston, K.J., Li, B., Daunizeau, J., Stephan, K.E., 2011. Network discovery with DCM. Neuroimage 56 (3), 1202–1221.
Garrido, M.I., Barnes, G.R., et al., 2012. Functional evidence for a dual route to amygdala. Curr. Biol. 22 (2), 129–134.
Grillner, S., Graybiel, A.A.M., 2006. Microcircuits: The Interface Between Neurons and Global Brain Function. MIT Press.
Hipp, J.F., Hawellek, D.J., Corbetta, M., Siegel, M., Engel, A.K., 2012. Large-scale cortical correlation structure of spontaneous oscillatory activity. Nat. Neurosci. 15 (6), 884–890.
Honey, C.J., Kötter, R., Breakspear, M., Sporns, O., 2007. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. Proc. Natl. Acad. Sci. U. S. A. 10240–10245.
Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. Proc. Natl. Acad. Sci. U. S. A. 2035–2040.
Hunt, L.T., Kolling, N., Soltani, A., Woolrich, M.W., Rushworth, M.F.S., Behrens, T.E.J., 2012. Mechanisms underlying cortical activity during value-guided choice. Nat. Neurosci. 470–476.

Jbabdi, S., Woolrich, M.W., Andersson, J.L.R., Behrens, T.E.J., 2007. A Bayesian framework for global tractography. Neuroimage 116–129.
Li, B., Daunizeau, J., Stephan, K.E., Penny, W., Hu, D., Friston, K., 2011. Generalised filtering and stochastic DCM for fMRI. Neuroimage 58 (2), 442–457.
Luckhoo, H., Hale, J.R., Stokes, M.G., Nobre, A.C., Morris, P.G., Brookes, M.J., Woolrich, M.W., 2012. Inferring task-related networks using independent component analysis in magnetoencephalography. Neuroimage 530–541.
Markov, N.T., Misery, P., Falchier, A., Lamy, C., Vezoli, J., Quilodran, R., Gariel, M.A., Giroud, P., Ercsey-Ravasz, M., Pilaz, L.J., Huissoud, C., Barone, P., Dehay, C., Toroczkai, Z., Van Essen, D.C., Kennedy, H., Knoblauch, K., 2011. Weight consistency specifies regularities of macaque cortical networks. Cereb. Cortex 1254–1272.
Marreiros, A.C., Kiebel, S.J., Friston, K.J., 2010. A dynamic causal model study of neuronal population dynamics. Neuroimage 91–101.
Marrelec, G., Krainik, A., Duffau, H., Pélégrini-Issac, M., Lehéricy, S., Doyon, J., Benali, H., 2006. Partial correlation for functional brain interactivity investigation in functional MRI. Neuroimage 228–237.
Moran, R.J., Stephan, K.E., Kiebel, S.J., Rombach, N., O'Connor, W.T., Murphy, K.J., Reilly, R.B., Friston, K.J., 2008. Bayesian estimation of synaptic physiology from the spectral responses of neural masses. Neuroimage 272–284.
Moran, R.J., Stephan, K.E., Seidenbecher, T., Pape, H.-C., Dolan, R.J., Friston, K.J., 2009. Dynamic causal models of steady-state responses. Neuroimage 796–811.
Moran, R.J., Stephan, K.E., Dolan, R.J., Friston, K.J., 2011. Consistent spectral predictors for dynamic causal models of steady-state responses. Neuroimage 1694–1708.
Morris, C., Lecar, H., 1981. Voltage oscillations in the barnacle giant muscle fiber. Biophys. J. 193–213.
Mosher, J.C., Leahy, R.M., Lewis, P.S., 1999. EEG and MEG: forward solutions for inverse methods. IEEE Trans. Biomed. Eng. 245–259.
Passingham, R.E., Stephan, K.E., Kötter, R., 2002. The anatomical basis of functional localization in the cortex. Nat. Rev. Neurosci. 606–616.
Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004. Comparing dynamic causal models. Neuroimage 1157–1172.
Penny, W.D., Stephan, K.E., Daunizeau, J., Rosa, M.J., Friston, K.J., Schofield, T.M., Leff, A.P., 2010. Comparing families of dynamic causal models. PLoS Comput. Biol. e1000709.
Rowe, J.B., Hughes, L.E., Barker, R.A., Owen, A.M., 2010. Dynamic causal modelling of effective connectivity from fMRI: are results reproducible and sensitive to Parkinson's disease and its treatment? Neuroimage 1015–1026.
Schoffelen, J.-M., Gross, J., 2009. Source connectivity analysis with MEG and EEG. Hum. Brain Mapp. 1857–1865.
Seghier, M.L., Friston, K.J., 2013. Network discovery with large DCMs. Neuroimage 68, 181–191.
Shanahan, M., 2010. Metastable chimera states in community-structured oscillator networks. Chaos 013108.
Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011. Network modelling methods for FMRI. Neuroimage 875–891.
Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K., 2012. Temporally-independent functional modes of spontaneous brain activity. Proc. Natl. Acad. Sci. U. S. A. 3131–3136.
Stephan, K.E., Kamper, L., Bozkurt, A., Burns, G.A., Young, M.P., Kötter, R., 2001. Advanced database methodology for the Collation of Connectivity data on the Macaque brain (CoCoMac). Philos. Trans. R. Soc. Lond. B Biol. Sci. 1159–1186.
Stephan, K.E., Harrison, L.M., Penny, W.D., Friston, K.J., 2004. Biophysical models of fMRI responses. Curr. Opin. Neurobiol. 14 (5), 629–635 (Oct., Review).
Stephan, K.E., Weiskopf, N., Drysdale, P.M., Robinson, P.A., Friston, K.J., 2007b. Comparing hemodynamic models with DCM. Neuroimage 387–401.
Stephan, K.E., Harrison, L.M., Kiebel, S.J., David, O., Penny, W.D., Friston, K.J., 2007a. Dynamic causal models of neural system dynamics:current state and future extensions. J. Biosci. 129–144.
Stephan, K., Kasper, L., Harrison, L., Daunizeau, J., den Ouden, H., Breakspear, M., 2008. Nonlinear dynamic causal models for FMRI. Neuroimage 42 (2), 649–662.
Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., Friston, K.J., 2009a. Bayesian model selection for group studies. Neuroimage 1004–1017.
Stephan, K.E., Tittgemeyer, M., Knösche, T.R., Moran, R.J., Friston, K.J., 2009b. Tractography-based priors for dynamic causal models. Neuroimage 1628–1638.
Strother, S.C., Anderson, J., Hansen, L.K., Kjems, U., Kustra, R., Sidtis, J., Frutiger, S., Muley, S., LaConte, S., Rottenberg, D., 2002. The quantitative evaluation of functional neuroimaging experiments: the NPAIRS data analysis framework. Neuroimage 747–771.
Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E.J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., Della Penna, S., Feinberg, D., Glasser, M.F., Harel, N., Heath, A.C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E., Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A.Z., Xu, J., Yacoub, E., Consortium, W.-M.H., 2012. The Human Connectome Project: a data acquisition perspective. Neuroimage 2222–2231.
Vrba, J., Robinson, S.E., 2000. Linearly constrained minimum variance beamformers, synthetic aperture magnetometry, and MUSIC in MEG applications. Signals, Systems and Computers, 2000. Conference Record of the Thirty-Fourth Asilomar Conference on, vol. 311, pp. 313–317.
Wilson, H.R., Cowan, J.D., 1972. Excitatory and inhibitory interactions in localized populations of model neurons. Biophys. J. 1–24.
Woolrich, M.W., 2012. Bayesian inference in FMRI. Neuroimage 801–810.
Woolrich, M., Jbabdi, S., Behrens, T., 2009. fMRI dynamic causal modelling with inferred regions of interest. Abstract presented at the annual meeting of the Organsation for Human Brain Mapping, San Francisco.
Woolrich, M., Hunt, L., Groves, A., Barnes, G., 2011. MEG beamforming using Bayesian PCA for adaptive data covariance matrix regularization. Neuroimage 1466–1479.
Yendiki, A., Panneck, P., Srinivasan, P., Stevens, A., Zöllei, L., Augustinack, J., Wang, R., Salat, D., Ehrlich, S., Behrens, T., Jbabdi, S., Gollub, R., Fischl, B., 2011. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. Front. Neuroinform. 23.