Review

A short history of causal modeling of fMRI data

Klaas Enno Stephan [a,b,*], Alard Roebroeck [c]

[a] Laboratory for Social and Neural Systems Research, Dept. of Economics, University of Zurich, Switzerland
[b] Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, UK
[c] Dept. of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University, The Netherlands

ARTICLE INFO

ABSTRACT

Twenty years ago, the discovery of the blood oxygen level dependent (BOLD) contrast and invention of functional magnetic resonance imaging (MRI) not only allowed for enhanced analyses of regional brain activity, but also laid the foundation for novel approaches to studying effective connectivity, which is essential for mechanistically interpretable accounts of neuronal systems. Dynamic causal modeling (DCM) and Granger causality (G-causality) modeling have since become the most frequently used techniques for inferring effective connectivity from fMRI data. In this paper, we provide a short historical overview of these approaches, describing milestones of their development from our subjective perspectives.

© 2012 Elsevier Inc. All rights reserved.

Contents

Introduction

The advent of fMRI in the early 1990s revolutionized neuroimaging of the human brain. Beyond enhanced analyses of functional specialization, it enabled novel analyses of effective connectivity, i.e. the causal influences that neural units exert over another, opening new avenues to mechanistic accounts of neuronal systems. Prior to fMRI, such analyses could only be performed at the group level (with positron emission tomography, PET, data) or were restricted to parts of the brain near the scalp at limited spatial resolution (as in MEEG/EEG studies). The first analysis of effective connectivity from fMRI data (which the authors are aware of) relied on a regression model describing modulatory interactions between human visual areas V1 and V2 (Friston et al., 1995). This approach was subsequently extended and became known as "psychophysiological interaction" analysis (Friston et al., 1997). At the same time, many fMRI researchers started using structural equation modeling (SEM; Büchel and Friston, 1997; Bullmore et al., 2000; Horwitz et al., 1999), whose utility for network analysis had previously been established for 2-deoxyglucose and PET data by Randy McIntosh (McIntosh and Gonzales-Lima, 1994; McIntosh and Gonzalez-Lima, 1991). While

PPI and SEM have played an important role for establishing effective connectivity analyses, they either assume no (SEM) or fixed (PPI) temporal dependencies across data points (or their underlying causes) but do not estimate them from the data. Consequently, over the years, there was a growing sense that more powerful techniques were needed to exploit temporally resolved information contained in fMRI signals. This widespread notion was expressed at a landmark meeting that took place in May 2002: the first "Brain Connectivity Workshop" (BCW), organized by Rolf Kötter and Karl Friston in Düsseldorf (Lee et al., 2003; Stone and Kotter, 2002). At this meeting, a community was born that has since been shaping research on brain connectivity in neuroimaging and beyond (see www.brain-connectivity-workshop.org). It was at this meeting that two emerging techniques, dynamic causal modeling (DCM) and multivariate/vector autoregressive models (MAR/VAR[1]), with and without reference to Granger causality (or G-causality), were first discussed in a wider forum, one year prior to the first publications (Friston et al., 2003; Goebel et al., 2003; Harrison et al., 2003). They have since become the most frequently used techniques for inferring effective connectivity from fMRI data.

It was not a coincidence that the first encounter of the authors, both of whom were PhD students at the time, took place at this meeting. One of us (KES) was a student of Rolf Kötter, developing a database of anatomical connectivity in primates (CoCoMac; Stephan et al., 2001) that Rolf would later turn into a freely accessible web repository (Kötter 2004). While the original goal of my CoCoMac work had been to provide anatomical constraints for large-scale biophysical models of brain function, I had become a little frustrated with the difficulties of verifying the goodness of such large-scale models and was now keen to learn more about the alternative approach, where simpler models are used to infer effective connectivity in circumscribed systems from empirical measurements. The other one of us (AR) had just gained his first experience with fMRI and diffusion-weighted MRI data at a time when 3T machines started becoming a new research standard and had started investigating G-causality for fMRI.

For both of us, the initial BCW meeting had an important influence on our scientific trajectories, setting up the stage for subsequent methodological and empirical work within the frameworks of DCM and G-causality, respectively. In this paper, we provide a short historical overview, written from our subjective perspectives, describing the history of developments in DCM and G-causality. A graphical summary of what we (subjectively) perceive as the most important milestones over the past decade is provided by Fig. 1. Due to the topic of this Special Issue of *NeuroImage*, this paper focuses on fMRI only, and we ask our colleagues to forgive us for not discussing in depth the contributions that DCM and GCM have made to the analyses of electrophysiological data. Similarly, we regret not being able to cover other approaches to causal modeling of fMRI data, such as graphical causal models (e.g., Ramsey et al. 2011) or extensions of SEM and its combinations with VAR (e.g., Gates et al. 2010, 2011).

Dynamic causal modeling

DCM for fMRI data was introduced in a seminal paper by Karl Friston, Lee Harrison and Will Penny in 2003 (Friston et al., 2003). Combining concepts from control theory and probability theory, DCM represents a Bayesian framework for specifying and comparing generative models of measured brain responses. These models provide estimates of neurophysiologically interpretable quantities, including (but not limited to) the effective connectivity among neuronal populations. While numerous types of DCMs have been implemented, all share five key characteristics[2] (cf. Stephan et al., 2010): First, hidden (unobserved) neuronal dynamics are described by (potentially nonlinear and stochastic) differential equations. Second, DCMs are hierarchical models, where a forward model links the neuronal state equations to measured data. Third, DCM is based on the control theory concept of causality, describing how dynamics in one neuronal population causes dynamics in other populations, and how these interactions are modulated by experimentally controlled perturbations. Fourth, model inversion (fitting) rests on Bayesian principles, providing both posterior estimates of the parameters and an estimate of the model evidence. Fifth, the central goal of DCM is not to decide whether an experimental condition elicited an effect; rather it serves to compare the relative plausibility of alternative neurophysiological mechanisms that may have caused an established effect (i.e., model selection).

One of us (KES) joined Karl Friston's group as a post-doc shortly before the initial DCM paper (Friston et al., 2003) was published. While attending the 2003 Human Brain Mapping meeting at New York, I read a preprint of the paper that Karl had given me. I was immediately attracted to the approach, for three reasons. First, I noted the close mathematical relation of DCM to General System Theory (Von Bertalanffy, 1969), which had been a source of inspiration for me since my undergraduate studies. Secondly, the paper introduced me to Bayesian inference techniques for dynamic systems, a field that I was unfamiliar with but found fascinating. I would spend the next years familiarizing myself with (variational) Bayesian techniques, a learning process that was greatly aided not only by Karl's thoughtful supervision, but also by the help from my colleagues and friends in the FIL Methods Group at London, most notably Lee Harrison, Will Penny and Jean Daunizeau. And perhaps most importantly, I felt that the ambition of DCM to provide probabilistic estimates of "hidden" neuronal processes could provide one key methodology for my long-term research goal which started emerging at that time: developing model-based approaches for inferring pathophysiological processes and predicting optimal treatment in individual patients (Stephan, 2004; Stephan et al., 2006). Eight years later, I feel encouraged by proof-of-concept studies (e.g., Brodersen et al., 2011b; Moran et al., 2011b) that support the feasibility of establishing model-based diagnostics of individual patients in the future.

The initial DCM paper (Friston et al., 2003) drew on a number of important previous developments, including Bayesian techniques for identification of dynamic systems (Friston, 2002) and the hemodynamic "Balloon" model for BOLD signals (Buxton et al., 1998; Friston et al., 2000). With this biophysically motivated forward model, DCM was the first generative model of fMRI data which allowed for inference on effective connectivity at a neuronal level, not the BOLD level, and which could deal with potential confounds due to inter-regional variations in BOLD responses. Several refinements of this hemodynamic model have taken place since (e.g., Kiebel et al., 2007; Riera et al., 2004; Stephan et al., 2007b), and further future improvements are expected concerning, for example, the effects of different field strengths and acquisition techniques (Uludag et al., 2009) or the relative contributions of glutamatergic vs. GABAergic transmission.

The original DCM for fMRI used bilinear differential equations as a low-order approximation to any nonlinear system, describing how neuronal population dynamics arises from effective connectivity and its context-dependent modulation. Three major extensions have been suggested subsequently: (i) *nonlinear DCMs* which account for synaptic interactions and activity-dependent gating of connections

---

[1] Both notations are common and interchangeable; for simplicity, we will use VAR throughout this paper.

[2] We use the acronym DCM to refer both to the modeling approach and to a specific dynamic causal model.

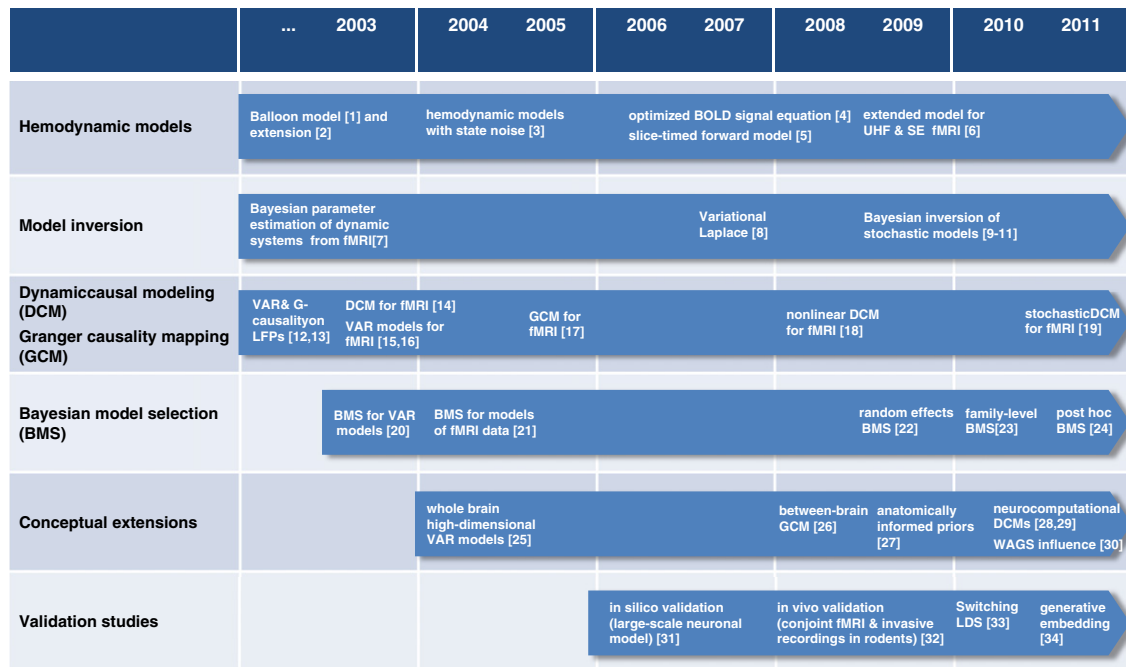|	...	2003	2004	2005	2006	2007	2008	2009	2010	2011
Hemodynamic models		Balloon model [1] and extension [2]	hemodynamic models with state noise [3]		optimized BOLD signal equation [4] slice-timed forward model [5]		extended model for UHF & SE fMRI [6]			
Model inversion		Bayesian parameter estimation of dynamic systems from fMRI[7]				Variational Laplace [8]	Bayesian inversion of stochastic models [9-11]			
Dynamiccausal modeling (DCM) **Granger causality mapping (GCM)**	VAR& G-causalityon LFPs [12,13]	DCM for fMRI [14] VAR models for fMRI [15,16]		GCM for fMRI [17]			nonlinear DCM for fMRI [18]			stochasticDCM for fMRI [19]
Bayesian model selection (BMS)		BMS for VAR models [20]	BMS for models of fMRI data [21]					random effects BMS [22]	family-level BMS[23]	post hoc BMS [24]
Conceptual extensions			whole brain high-dimensional VAR models [25]				between-brain GCM [26]	anatomically informed priors [27]	neurocomputational DCMs [28,29] WAGS influence [30]	
Validation studies					in silico validation (large-scale neuronal model) [31]		in vivo validation (conjoint fMRI & invasive recordings in rodents) [32]		Switching LDS [33]	generative embedding [34]

Fig. 1. A graphical summary of what the authors subjectively perceive as the most important milestones in developing VAR models and DCM over the past decade. [1] Buxton et al. 1998; [2] Friston et al. 2000; [3] Riera et al. 2004; [4] Stephan et al. 2007b; [5] Kiebel et al. 2007; [6] Uludag et al. 2009; [7] Friston 2002; [8] Friston et al. 2007; [9] Friston et al. 2008; [10] Daunizeau et al. 2009; [11] Friston et al. 2010; [12] Freiwald et al. 1999; [13] Bernasconi & Konig, 1999; [14] Friston et al. 2003; [15] Goebel et al. 2003; [16] Harrison et al. 2003; [17] Roebroeck et al. 2005; [18] Stephan et al. 2008; [19] Friston et al. 2011; [20] Penny & Roberts 2002;.[21] Penny et al. 2004; [22] Stephan et al. 2009a; [23] Penny et al. 2010; [24] Friston & Penny 2011; [25] Valdes-Sosa 2004; [26] Schippers et al. 2010; [27] Stephan et al. 2009b; [28] den Ouden et al. 2009; [29] den Ouden et al. 2010; [30] Valdes-Sosa et al. 2011; [31] Lee et al. 2006; [32] David et al. 2008; [33] Smith et al. 2010; [34] Brodersen et al. 2011b.

(Stephan et al., 2008), (ii) *two-state DCMs* which represent separate excitatory and inhibitory populations within each region (Marreiros et al., 2008), and (iii) *stochastic DCMs* which account for endogenous fluctuations in fMRI signals and can also be applied in the absence of experimental control, e.g., resting-state fMRI (Daunizeau et al., 2009; Friston et al., 2011; Li et al., 2011). The latter DCMs were made possible by novel inversion techniques, such as Dynamic Expectation Maximization (DEM; (Friston et al., 2008)) or Generalised Filtering (Friston et al., 2010), which address the triple estimation problem of identifying states, parameters and hyperparameters in stochastic hierarchical system models.

The first application paper of DCM followed the original report (Friston et al., 2003) by just a few months. This was a paper by Andrea Mechelli et al. (2003), examining bottom-up and top-down mechanisms of object category processing in visual cortex. DCM for fMRI has since been used in approximately 150 empirical and methodological studies, addressing a broad range of neurophysiological and cognitive questions. Over the years, the application papers followed the methodological developments closely. A good example is Bayesian model selection (BMS), a generic procedure from probability theory for assessing the fit/complexity trade-off (and thus generalizability) when comparing models. While model comparison had been a central theme to DCM from the beginning (Penny et al., 2004), early DCM studies were typically more concerned with parameter estimates (given a particular model) than model structure. Subsequently, however, DCM studies have been focusing more and more on model selection. This was driven by a series of methodological papers, starting with Will Penny's initial work on order selection for VAR models (Penny and Roberts, 2002) and AIC/BIC based comparison of DCMs (Penny et al., 2004), followed by papers on the utility of the free energy approximation to the model evidence (Friston et al., 2007; Penny, 2012; Stephan et al., 2007b), by novel methods for random effects group BMS and family-level inference (Penny et al., 2010; Stephan et al., 2009a), and, most recently, rapid post-hoc scoring

schemes for models (Friston and Penny, 2011). Additionally, the benefits of systematically structured model spaces for testing competing hypotheses about neuronal mechanisms using fMRI data were demonstrated by several empirical applications of BMS (e.g., Leff et al., 2008; Rowe et al., 2010; Stephan et al., 2007a; van Leeuwen et al., 2011).

In our interactions with colleagues and students over the years, however, we have learned that not everyone finds the statistical principles of DCM and BMS easy to understand. This reflects the diversity of the neuroimaging community and the fact that Bayesian inference and dynamic systems theory are not part of the curriculum in most disciplines. Therefore, "good practice" rules have been derived that provide guidance for the application of DCM and BMS (Stephan et al., 2010). The first step generally is to define the model space, motivating a set of competing hypotheses (alternative models) whose relative plausibility is tested against the measured data.[3] One can then proceed in one of three ways. Most studies to date have searched for a single best model (among those considered) in order to examine its structure and/or parameter estimates. An alternative approach of increasing popularity is to address more broadly defined questions by partitioning model space and comparing families of models (Penny et al., 2010; Stephan et al., 2009a), e.g., models with and without particular types of connections. This approach is particularly useful for large model spaces to prevent overfitting (at the level of model inference) and dilution. A third option is to use Bayesian model averaging (BMA) where parameter estimates are averaged across all models considered, weighted by the posterior model probabilities (Penny et al. 2010). Random effects BMS and BMA provide a principled way to

---

[3] Formally, this corresponds to specifying a prior density on models such that those models considered implausible are assigned a prior probability of zero.

deal with situations where the optimal model may differ across subjects and groups, as in clinical studies (e.g., Banyai et al., 2011).

Personally, I find that one of the most important aspects of DCM is that it forces one to formalize and make explicit one's set of hypotheses (i.e., the model space) and how exactly they are compared. Neuroimaging analyses have often been conducted in an exploratory way, and this tendency has been further reinforced with the popularity of "resting-state" studies. While exploratory analyses of fMRI data are certainly useful (and indeed mandatory) as long as systems are not well understood, the richness and complexity of their results can invite unconstrained interpretations and *post-hoc* injection of meaning. Instead, exploration should be used to generate initial hypotheses; these hypotheses then need to be represented as models whose relative plausibility is tested by subsequent studies. This is precisely the typical sequence of a DCM study, where the results of initial mass-univariate analyses motivate a set of hypotheses which are transcribed into competing DCMs that are then compared using one of the BMS variants described above.

G-causality modeling

The application of G-causality to fMRI data was introduced in 2003 in a paper by Rainer Goebel et al. (2003) and then further extended and validated in a following publication (Roebroeck et al., 2005). Thus, one of us (AR) was heavily involved in its development. Already before we me met, Rainer and I were both influenced by the same animal studies that introduced the use of autoregressive modeling for spectral analysis (Ding et al., 2000) and G-causality modeling (Bernasconi and Konig, 1999; Freiwald et al., 1999) of electrophysiological recordings. When subsequently starting my PhD studies on the analysis of structural and functional brain connectivity using MRI techniques with Rainer, we asked whether G-causality could be applied to fMRI data despite the sluggishness of the BOLD response. We were further inspired by the application of Volterra kernel expansions to fMRI connectivity analysis (Friston et al., 2000) and decided to evaluate the usefulness of G-causality in fMRI effective connectivity modeling. The fundamental motivation for exploring this was the principle, also expressed above, that complex data modeling efforts must be a symbiosis between exploration of the model space and prior assumption constrained hypothesis testing. A delicate balance must be struck between the two: too much exploration might lead to meaning injected into small parts of a large body of results; too many constraints on a model space without rigid justification might lead to bias at the level of model inference. Therefore our aim was to add exploration oriented techniques for connectivity to complement the hypothesis driven ones that were already available (or becoming available) at the time, such as SEM and DCM. In parallel to our efforts, Pedro Valdes-Sosa also explored the use of whole brain autoregressive modeling, aided by sparse regression methods, interpreted in the framework of G-causality (Valdes-Sosa, 2004; Valdes-Sosa et al., 2005), and Lee Harrison explored bilinear extensions of VAR models for fMRI effective connectivity (Harrison et al., 2003; albeit outside the context of G-causality).

The implementation of connectivity exploration in these initial papers was led by two guiding ideas. First and foremost, effective connectivity models can be constructed that consider many or all regions in the brain as potential network nodes. This exploration of the structural model (the set of assumptions that determine what—which set of regions—can interact) can avoid the missing region problem: spurious inference on connectivity due to regions left out of the model, for instance a source of common input (Roebroeck et al., 2011a). Thus the initial publications focused on mapping G-causality over the entire brain (Granger Causality Mapping, GCM; Goebel et al., 2003; Roebroeck et al., 2005) and whole brain G-causality analysis aided by high-dimensional regression approaches (Valdes-Sosa, 2004; Valdes-Sosa et al., 2005; see also the recent

work by Garg et al. 2011). Second, for the structural model to be permissive enough to allow for this exploration, the dynamical model (the signal equations that determine how regions interact) must contain rather strongly directive assumptions. Here, the concept of G-causality dictates that past information of a causing time series must predict present or future values of a caused time series, given that other relevant information is taken into account.

In the initial phases of development of fMRI G-causality mapping a lot of thought went into what is clearly its greatest challenge: the presence and variability over the brain of the hemodynamics that intervene between neuronal population activity and measured BOLD signals. In our thinking we separated this into two sequential validation steps, the results of which were eventually reported in Roebroeck et al. (2005). In the first, we asked whether the combined aggregating effect of the sluggish hemodynamics and the relatively sparse temporal sampling of the BOLD signal (even if assumed to be the same in all regions) makes application of G-causality to fMRI possible at all. Extensive ground-truth software simulations led to some hesitation in the case of fully multivariate models but, encouragingly, bi-variate G-causality results could be used with high sensitivity and specificity. Given this result, in the second step, we asked ourselves whether the variability of the hemodynamics within the brain could lead to loss of sensitivity or bias in estimates of the strength and direction of causality if that variability was simply ignored. To us, the obvious but uninformative answer was: yes. When the variability of hemodynamic delays in two regions around an equal average delay approaches the size of neuronal delays, sensitivity is lost. When a systematic difference in mean delay between two regions exceeds neuronal delays, bias in careless directionality estimates will ensue. To our minds, the limited informativeness of this answer lay in the uncertainty, spanning orders of magnitude, about the relevant neuronal delays and hemodynamic variability. In the human brain, neuronal conduction delays are 10 or 20 ms at most, but event related potentials (ERP) related to perceptual or cognitive processes are measured in hundreds of milliseconds and a few challenging tasks (such as the so-called clock-task; Formisano et al., 2002), have measureable fMRI responses seconds apart. Hemodynamic variability between brain areas has been reported to be on the order of hundreds of milliseconds extending to over a second, although this variability has sometimes included task-related neuronal delays (which are thus tacitly assumed to be negligible). Our internal consensus was that ignorance of the intervening role of hemodynamics in this context is careless, but a downright dismissal of all relative temporal structure of fMRI signals could be equally wasteful of potentially useful information. Rather than pursuing either of these extreme strategies we felt—and still do—that a careful interrogation of temporal order structure in fMRI data should start (but not end) with looking for experimentally induced changes in the detected G-causality. Much of these thoughts and discussion were expressed in writing in Roebroeck et al. (2005). The initial publications led to an increasing wave of studies investigating directed interactions in the brain by fMRI G-Causality analysis. Although (with very few exceptions) these studies have all shown awareness of the hemodynamic variability confound, the ways and means by which it is accounted for are almost as diverse as the applications they investigate, which is perhaps why it has been a source of some discussion in more recent years (see below).

It is important to stress that the concept of G-causality is historically not at all bound to discrete-time autoregressive modeling (cf. Granger, 1980), although it is often (wrongly) equated with it. Conversely, VAR models of effective connectivity do not necessarily refer to the concept of G-causality, e.g., (Harrison et al., 2003; Ryali et al., 2011). Early on in the neuroscience context, G-causality was instantiated in nonlinear models (Freiwald et al., 1999) and time-varying models for non-stationary data (Havlicek et al., 2010; Hesse et al., 2003), and it has been framed in terms of non-parametric spectral factorization (Dhamala et al., 2008). In fact, Valdes-Sosa et al.

(2011) have coined the term "Wiener–Akaike–Granger–Schweder (WAGS) influence" to acknowledge the contribution of all these scientists to generalizing predictability-based causality in time series to a wealth of data types—continuous time as well as discrete time and real valued as well as spiking processes—and to a multitude of prediction models. In addition, this general concept helps in unifying time series approaches to causality in a discrete-time tradition and dynamic systems and control theory approaches in a continuous time tradition (Roebroeck et al., 2011c).

A particularly powerful use of fMRI G-Causality has recently been reported in Schippers et al. (2010), where between-brain (rather than within-brain) influence is investigated in the context of human social communication by gestures in the game of charades. Results show that a guesser's brain activity in the human mirror system regions influences the gesturer's brain activity. This illustrates the use of the G-causality concept to answer interesting neuroscience questions that (i) are not easily formulated by a generative model, for instance incorporating the entire interpersonal communication chain and (ii) are concerned with the identification of the unknown areas that interact to perform a perceptual, cognitive or motor task. These types of questions have characterized many of the applications of fMRI G-causality, such as top-down control of visual spatial attention (Bressler et al., 2008), switching between executive control and default-mode networks (Sridharan et al., 2008), fatigue (Deshpande et al., 2009) and the resting state (Uddin et al., 2009).

Model validation

Model selection is not about finding the "true" model. Generally, models are never "true" or "false" in an absolute sense (with the exception of synthetic data where the underlying generative model was constructed and is thus known). Instead, models are deliberately simplistic caricatures of the real world, aiming to unmask a particular mechanism that is not visible from the high-dimensional data. As the famous saying by Box and Draper (1987) puts it: "Essentially, all models are wrong, but some are useful" (p. 424). Determining whether or not a model is "useful" for a particular application domain requires systematic studies of the model's reliability and its face, construct and predictive validity. Simply put, we need to know what can and what cannot be safely inferred from our models. For example, even when a model is well motivated by theory and prior knowledge, it is possible that it allows for inference on certain (neuronal) mechanisms, but not on others, due to conditional dependencies among parameters that are typical for biological systems (Gutenkunst et al., 2007; Stephan et al., 2007b). In technical terms this is the question of identifiability of parameters and state trajectories. For these aspects of face validity, numerical analyses and simulation studies with known "ground truth" play an important role.

These *in silico* approaches have played a major role during the history of both DCM (e.g., Friston et al., 2003; Lee et al., 2006; Stephan et al., 2008; Stephan et al., 2009a) and G-causality models (e.g., Roebroeck et al., 2005; Ryali et al., 2011; Schippers et al., 2011; Smith et al., 2011a; Smith et al., 2011b). While the merits and plausibility of each simulation study need to be examined carefully case-by-case, it is clear that they have greatly contributed (and continue to do so) to our understanding under what conditions certain aspects of model-based inference may fail, and which other aspects may remain robust.

The most difficult challenge, however, is to establish the predictive validity of models. This requires one to test how a model performs in relation to external criteria that are independent of the data to which the model is fitted. In the history of DCM and G-causality models, such tests of predictive validity follow three major strategies. One way is to assess how well "unseen" (test) data can be predicted based on model parameter estimates obtained from known (training) data (for example applications,

see Smith et al., 2010, and Brodersen et al., 2011a). This 'training-set/test-set' logic has the additional advantage that can be used for *structural model selection* for fMRI (cf. Roebroeck et al., 2011b). A second approach is to validate models against external (independent) labels or facts, for example, known diagnostic states of individual patients or their individual treatment response (e.g., generative embedding; Brodersen et al., 2011b). The third and most widely used approach to date for addressing the predictive validity of DCMs and G-causality models tests whether a given model can infer known or measurable consequences of a controlled experimental intervention. This approach requires carefully planned invasive studies in animals (e.g., neurochemistry, electrical stimulation, genetic manipulations or lesions) or humans (e.g., neuropharmacology). While an increasing number of such studies have been performed in recent years, they have mainly addressed the validity of DCMs for electrophysiological data (e.g., Moran et al., 2008; Moran et al., 2011a; Moran et al., 2011b); in contrast they have been relatively rare for models of fMRI data. One important exception is the study by David et al. (2008) who obtained simultaneous fMRI and invasive electrophysiological measurements from rats with a genetically defined type of epilepsy. Using the electrophysiological data for establishing a "ground truth", they asked whether models were capable of inferring, from fMRI data alone, in which region the seizure originated. They found that DCM was capable of doing this, despite profound regional variations in the shape and latency of the BOLD signal; these confounding effects, however, were accounted for by the hemodynamic forward model in DCM. A VAR model of G-causality was also capable of identifying the seizure origin, but only once this model was augmented with a hemodynamic deconvolution procedure (which, however, did not operate on the BOLD data alone but was informed by the simultaneous electrophysiological recordings, in the same way as the DCM analysis, using it as driving input to the model). Overall, this study demonstrated empirically that a hemodynamic forward model can be critical for unconfounded inference on effective connectivity from fMRI data, a notion that has been discussed extensively during the history of effective connectivity models for fMRI data (Friston et al., 2003; Gitelman et al., 2003; Roebroeck et al., 2011a; Stephan et al., 2004).

Controversies and points of convergence

The paper by David et al. (2008) and an associated commentary (Friston, 2009) sparked a lively debate on the pros and cons of DCM and G-causality models for inferring effective connectivity from fMRI data. This led to a series of articles in NeuroImage (e.g., Daunizeau et al., 2011; David, 2011; Friston, 2011a; Roebroeck et al., 2011a, b; Valdes-Sosa et al., 2011) and other journals (Friston, 2011b; Roebroeck et al., 2011c) that discussed differences between the approaches and potential problems, such as neurophysiological interpretability and unknown uncertainty of model parameters in GCM, or robustness of variational Bayesian procedures and issues of structural model selection in DCM. However, there are also points of convergence that were discussed in these papers. For instance, hemodynamic forward models, once the unique hallmark of DCM, have recently been integrated into state space models of an extended VAR type (Ryali et al., 2011; Smith et al., 2010). In turn, DCM has been augmented to include random innovations, albeit assuming temporally smooth as opposed to Markovian noise in VAR models (Friston et al., 2010; Li et al., 2011). A final summary paper written by authors representing both approaches (Valdes-Sosa et al., 2011) reviewed these constructive discussions and future possibilities of integration, concluding that "We are not saying that DCM and GCM are equivalent, but rather that an integration is possible within a Bayesian state space modeling framework and the use of model comparison methods."

Where do we go from here?

It is probably fair to say that G-causality models and DCM for fMRI have considerably expanded the tool-kit for fMRI data analysis and have already provided potentially important results for various domains of cognitive neuroscience. However, it is also fair to say that much methodological work remains to be done. As discussed above, this concerns model validation in particular. For this, we believe that close collaborations between modelers and neurophysiologists are crucial and should be sought more actively in the computational neuroimaging community.

Another challenge for the future concerns novel application concepts that go beyond "classical" effective connectivity analyses. Here, several innovative approaches have surfaced over the past few years are likely to play a major role in future. For example, one popular theme has been the analysis of structure–function relationships in neuronal systems by juxtaposing estimates of functional or effective connectivity to estimates of anatomical connectivity obtained by diffusion-weighted imaging (e.g., Koch et al., 2002; Upadhyay et al., 2008). Although useful, this approach has remained largely descriptive, and a more formal way is to incorporate the anatomical information directly in the definition of the prior densities that specify a model of effective connectivity. For example, such "anatomically informed priors" can be used to tune the prior variance of coupling parameters in a DCM and establish a mathematical relationship of how anatomical connectivity constrains effective connectivity (Stephan et al., 2009b). Future extensions of this approach will have to take account of the bias and variance that can exist in tractography results (probabilistic or deterministic; e.g., due to low spatial resolution; Roebroeck et al., 2008).

A second domain of innovation concerns models for "multimodal fusion", e.g., formally integrating models of fMRI and EEG data. The ultimate goal would be to specify a generative model accounting for all simultaneously acquired data; this would consists of a single neuronal model that is linked by separate forward models to the different measurement types (Daunizeau et al., 2010; Deneux and Faugeras, 2010; Riera et al., 2007; Rosa et al., 2010; Valdes-Sosa et al., 2009).

A third important recent development is the integration of computational models of learning into DCMs, thus enabling one to test for the expression of prediction error dependent plasticity in specific neuronal circuits (den Ouden et al., 2009; den Ouden et al., 2010). Such neurocomputational models have considerable potential as non-invasive assays of synaptic function and neuromodulatory regulation, for example in model-based classification of psychiatric spectrum diseases (Stephan et al., 2006). Clearly, careful and systematic validation studies will be required to test the robustness of model-based inference for clinical applications. Ideally, validation of such models in pharmacological and animal studies will go hand-in-hand with their application to real-world problems, such as predicting treatment responses or inferring pathophysiological states (e.g., Brodersen et al., 2011b).

Finally, the advent of ultra-high field fMRI has greatly increased the level of spatial detail that is accessible with this technique. For instance, fMRI at 7T provides sufficient spatial resolution to resolve orientation column maps in human primary visual cortex (Yacoub et al., 2008) and axis-of-motion maps in the human middle temporal motion area MT (Zimmermann et al., 2011). Potentially, this ongoing development can shift the level of causal and computational modeling for cognitive and translational neuroscience down to cortical columns and layers, arguably the fundamental mesoscale at which computational units of the brain operate. Importantly, having more fine-grained data will not obviate the need for careful modeling. On the contrary, as our methods for data acquisition become progressively refined, models of the neuronal causes underlying our measurements become ever more

necessary to avoid drowning in complexity. Hopefully, by the time *NeuroImage* publishes its next anniversary issue of fMRI in 10 years or so, we will see some models that have passed rigorous validation studies and have entered the practical application domain, solving important real-world problems, e.g., in clinical diagnostics.

References

Banyai, M., Diwadkar, V.A., Erdi, P., 2011. Model-based dynamical analysis of functional disconnection in schizophrenia. Neuroimage 58, 870–877.

Bernasconi, C., Konig, P., 1999. On the directionality of cortical interactions studied by structural analysis of electrophysiological recordings. Biol. Cybern. 81, 199–210.

Box, G.E.P., Draper, N.R., 1987. Empirical Model-Building and Response Surfaces. John Wiley and Sons, New York.

Bressler, S.L., Tang, W., Sylvester, C.M., Shulman, G.L., Corbetta, M., 2008. Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention. J. Neurosci. 28, 10056–10061.

Brodersen, K.H., Haiss, F., Ong, C.S., Jung, F., Tittgemeyer, M., Buhmann, J.M., Weber, B., Stephan, K.E., 2011a. Model-based feature construction for multivariate decoding. NeuroImage 56, 601–615.

Brodersen, K.H., Schofield, T.M., Leff, A.P., Ong, C.S., Lomakina, E.I., Buhmann, J.M., Stephan, K.E., 2011b. Generative embedding for model-based classification of fMRI data. PLoS Comput. Biol. 7, e1002079.

Büchel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI. Cereb. Cortex 7, 768–778.

Bullmore, E., Horwitz, B., Honey, G., Brammer, M., Williams, S., Sharma, T., 2000. How good is good enough in path analysis of fMRI data? NeuroImage 11, 289–301.

Buxton, R.B., Wong, E.C., Frank, L.R., 1998. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magn. Reson. Med. 39, 855–864.

Daunizeau, J., Friston, K.J., Kiebel, S.J., 2009. Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. Phys. D 238, 2089–2118.

Daunizeau, J., Vaudano, A.E., Lemieux, L., 2010. Bayesian multi-modal model comparison: a case study on the generators of the spike and the wave in generalized spike-wave complexes. NeuroImage 49, 656–667.

Daunizeau, J., David, O., Stephan, K.E., 2011. Dynamic causal modelling: a critical review of the biophysical and statistical foundations. NeuroImage 58, 312–322.

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. 6, 2683–2697.

David, O., 2011. fMRI connectivity, meaning and empiricism Comments on: Roebroeck et al. The identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. NeuroImage 58, 306–309.

den Ouden, H.E., Friston, K.J., Daw, N.D., McIntosh, A.R., Stephan, K.E., 2009. A dual role for prediction error in associative learning. Cereb. Cortex 19, 1175–1185.

den Ouden, H.E., Daunizeau, J., Roiser, J., Friston, K.J., Stephan, K.E., 2010. Striatal prediction error modulates cortical coupling. J. Neurosci. 30, 3210–3219.

Deneux, T., Faugeras, O., 2010. EEG-fMRI fusion of paradigm-free activity using Kalman filtering. Neural Comput. 22, 906–948.

Deshpande, G., LaConte, S., James, G.A., Peltier, S., Hu, X., 2009. Multivariate Granger causality analysis of fMRI data. Hum. Brain Mapp. 30, 1361–1373.

Dhamala, M., Rangarajan, G., Ding, M., 2008. Analyzing information flow in brain networks with nonparametric Granger causality. NeuroImage 41, 354–362.

Ding, M., Bressler, S.L., Yang, W., Liang, H., 2000. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. Biol. Cybern. 83, 35–45.

Formisano, E., Linden, D.E., Di Salle, F., Trojano, L., Esposito, F., Sack, A.T., Grossi, D., Zanella, F.E., Goebel, R., 2002. Tracking the mind's image in the brain I: time-resolved fMRI during visuospatial mental imagery. Neuron 35, 185–194.

Freiwald, W.A., Valdes, P., Bosch, J., Biscay, R., Jimenez, J.C., Rodriguez, L.M., Rodriguez, V., Kreiter, A.K., Singer, W., 1999. Testing non-linearity and directedness of interactions between neural groups in the macaque inferotemporal cortex. J. Neurosci. Methods 94, 105–119.

Friston, K.J., Ungerleider, L.G., Jezzard, P., Turner, R., 1995. Characterizing modulatory interactions between areas V1 and V2 in human cortex: a new treatment of functional MRI data. Hum. Brain Mapp. 2, 211–224.

Friston, K.J., Buechel, C., Fink, G.R., Morris, J., Rolls, E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. NeuroImage 6, 218–229.

Friston, K.J., Mechelli, A., Turner, R., Price, C.J., 2000. Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. NeuroImage 12, 466–477.

Friston, K.J., 2002. Bayesian estimation of dynamical systems: an application to fMRI. NeuroImage 16, 513–530.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. NeuroImage 19, 1273–1302.

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W., 2007. Variational free energy and the Laplace approximation. NeuroImage 34, 220–234.

Friston, K.J., Trujillo-Barreto, N., Daunizeau, J., 2008. DEM: a variational treatment of dynamic systems. NeuroImage 41, 849–885.

Friston, K., 2009. Causal modelling and brain connectivity in functional magnetic resonance imaging. PLoS Biol. 7, e33.

Friston, K., Stephan, K.E., Li, B., Daunizeau, J., 2010. Generalised filtering. Math. Probl. Eng. 2010, 621670.

Friston, K., 2011a. Dynamic causal modeling and Granger causality Comments on: the identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. NeuroImage 58, 303–305.

Friston, K.J., 2011b. Functional and effective connectivity: a review. Brain Connectivity 1, 13–36.

Friston, K., Penny, W., 2011. Post hoc Bayesian model selection. NeuroImage 56, 2089–2099.

Friston, K.J., Li, B., Daunizeau, J., Stephan, K.E., 2011. Network discovery with DCM. NeuroImage 56, 1202–1221.

Garg, R., Cecchi, G.A., Rao, A.R., 2011. Full-brain auto-regressive modeling (FARM) using fMRI. NeuroImage 58, 416–441.

Gates, K.M., Molenaar, P.C.M., Hillary, F.G., Ram, N., Rovine, M.J., 2010. Automatic search for fMRI connectivity mapping: an alternative to Granger causality testing using formal equivalences among SEM path modeling, VAR, and unified SEM. NeuroImage 50, 1118–1125.

Gates, K.M., Molenaar, P.C.M., Hillary, F.G., Slobounov, S., 2011. Extended unified SEM approach for modeling event-related fMRI data. NeuroImage 54, 1151–1158.

Gitelman, D.R., Penny, W.D., Ashburner, J., Friston, K.J., 2003. Modeling regional and psychophysiologic interactions in fMRI: the importance of hemodynamic deconvolution. NeuroImage 19, 200–207.

Goebel, R., Roebroeck, A., Kim, D.S., Formisano, E., 2003. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. Magn. Reson. Imaging 21, 1251–1261.

Granger, C.W.J., 1980. Testing for causality: a personal viewpoint. J. Econ. Dyn. Control. 2, 329–352.

Gutenkunst, R.N., Waterfall, J.J., Casey, F.P., Brown, K.S., Myers, C.R., Sethna, J.P., 2007. Universally sloppy parameter sensitivities in systems biology models. PLoS Comput. Biol. 3, 1871–1878.

Harrison, L., Penny, W.D., Friston, K., 2003. Multivariate autoregressive modeling of fMRI time series. NeuroImage 19, 1477–1491.

Havlicek, M., Jan, J., Brazdil, M., Calhoun, V.D., 2010. Dynamic Granger causality based on Kalman filter for evaluation of functional network connectivity in fMRI data. NeuroImage 53, 65–77.

Hesse, W., Moller, E., Arnold, M., Schack, B., 2003. The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies. J. Neurosci. Methods 124, 27–44.

Horwitz, B., Tagamets, M.A., McIntosh, A.R., 1999. Neural modeling, functional brain imaging, and cognition. Trends Cogn. Sci. 3, 91–98.

Kiebel, S.J., Kloppel, S., Weiskopf, N., Friston, K.J., 2007. Dynamic causal modeling: a generative model of slice timing in fMRI. NeuroImage 34, 1487–1496.

Koch, M.A., Norris, D.G., Hund-Georgiadis, M., 2002. An investigation of functional and anatomical connectivity using magnetic resonance imaging. NeuroImage 16, 241–250.

Kötter, R., 2004. Online retrieval, processing, and visualization of primate connectivity data from the CoCoMac database. Neuroinformatics 2, 127–144.

Lee, L., Harrison, L.M., Mechelli, A., 2003. A report of the functional connectivity workshop, Dusseldorf 2002. NeuroImage 19, 457–465.

Lee, L., Friston, K., Horwitz, B., 2006. Large-scale neural models and dynamic causal modelling. NeuroImage 30, 1243–1254.

Leff, A.P., Schofield, T.M., Stephan, K.E., Crinion, J.T., Friston, K.J., Price, C.J., 2008. The cortical dynamics of intelligible speech. J. Neurosci. 28, 13209–13215.

Li, B., Daunizeau, J., Stephan, K.E., Penny, W., Hu, D., Friston, K., 2011. Generalised filtering and stochastic DCM for fMRI. NeuroImage 58, 442–457.

Marreiros, A.C., Kiebel, S.J., Friston, K.J., 2008. Dynamic causal modelling for fMRI: a two-state model. NeuroImage 39, 269–278.

McIntosh, A.R., Gonzalez-Lima, F., 1991. Structural modeling of functional neural pathways mapped with 2-deoxyglucose: effects of acoustic startle habituation on the auditory system. Brain Res. 547, 295–302.

McIntosh, A.R., Gonzales-Lima, F., 1994. Structural equation modelling and its application to network analysis in functional brain imaging. Hum. Brain Mapp. 2, 2–22.

Mechelli, A., Price, C.J., Noppeney, U., Friston, K.J., 2003. A dynamic causal modeling study on category effects: bottom-up or top-down mediation? J. Cogn. Neurosci. 15, 925–934.

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 42, 272–284.

Moran, R.J., Jung, F., Kumagai, T., Endepols, H., Graf, R., Dolan, R.J., Friston, K.J., Stephan, K.E., Tittgemeyer, M., 2011a. Dynamic causal models and physiological inference: a validation study using isoflurane anaesthesia in rodents. PLoS One 6, e22790.

Moran, R.J., Symmonds, M., Stephan, K.E., Friston, K.J., Dolan, R.J., 2011b. An in vivo assay of synaptic function mediating human cognition. Curr. Biol. 21, 1320–1325.

Penny, W., Roberts, S., 2002. Bayesian multivariate autoregressive models with structured priors. IEE Proc. Visi. Image Signal. Process. 149, 33–41.

Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004. Comparing dynamic causal models. NeuroImage 22, 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. 6, e1000709.

Penny, W.D., 2012. Comparing dynamic causal models using AIC, BIC and free energy. NeuroImage 59, 319–330.

Ramsey, J.D., Hanson, S.J., Glymour, C., 2011. Multi-subject search correctly identifies causal connections and most causal directions in the DCM models of the Smith et al. simulation study. NeuroImage 58, 838–848.

Riera, J.J., Watanabe, J., Kazuki, I., Naoki, M., Aubert, E., Ozaki, T., Kawashima, R., 2004. A state-space model of the hemodynamic approach: nonlinear filtering of BOLD signals. NeuroImage 21, 547–567.

Riera, J.J., Jimenez, J.C., Wan, X., Kawashima, R., Ozaki, T., 2007. Nonlinear local electrovascular coupling. II: from data to neuronal masses. Hum. Brain Mapp. 28, 335–354.

Roebroeck, A., Formisano, E., Goebel, R., 2005. Mapping directed influence over the brain using Granger causality and fMRI. NeuroImage 25, 230–242.

Roebroeck, A., Galuske, R., Formisano, E., Chiry, O., Bratzke, H., Ronen, I., Kim, D.S., Goebel, R., 2008. High-resolution diffusion tensor imaging and tractography of the human optic chiasm at 9.4 T. NeuroImage 39, 157–168.

Roebroeck, A., Formisano, E., Goebel, R., 2011a. The identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. NeuroImage 58, 296–302.

Roebroeck, A., Formisano, E., Goebel, R., 2011b. Reply to Friston and David after comments on: the identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. NeuroImage 58, 310–311.

Roebroeck, A., Seth, A.K., Valdes-Sosa, P., 2011c. Causal time series analysis of functional magnetic resonance imaging data. J. Mach. Learn. Res. Workshop Conf. Proc. 12, 65–94.

Rosa, M.J., Daunizeau, J., Friston, K.J., 2010. EEG-fMRI integration: a critical review of biophysical modeling and data analysis approaches. J. Integr. Neurosci. 9, 453–476.

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 52, 1015–1026.

Ryali, S., Supekar, K., Chen, T., Menon, V., 2011. Multivariate dynamical systems models for estimating causal interactions in fMRI. NeuroImage 54, 807–823.

Schippers, M.B., Roebroeck, A., Renken, R., Nanetti, L., Keysers, C., 2010. Mapping the information flow from one brain to another during gestural communication. Proc. Natl. Acad. Sci. U. S. A. 107, 9388–9393.

Schippers, M.B., Renken, R., Keysers, C., 2011. The effect of intra- and inter-subject variability of hemodynamic responses on group level Granger causality analyses. NeuroImage 57, 22–36.

Smith, J.F., Pillai, A., Chen, K., Horwitz, B., 2010. Identification and validation of effective connectivity networks in functional magnetic resonance imaging using switching linear dynamic systems. NeuroImage 52, 1027–1040.

Smith, S.M., Bandettini, P.A., Miller, K.L., Behrens, T.E., Friston, K.J., David, O., Liu, T., Woolrich, M.W., Nichols, T.E., 2011a. The danger of systematic bias in group-level FMRI-lag-based causality estimation. NeuroImage.

Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011b. Network modelling methods for FMRI. NeuroImage 54, 875–891.

Sridharan, D., Levitin, D.J., Menon, V., 2008. A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. Proc. Natl. Acad. Sci. U. S. A. 105, 12569–12574.

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. 356, 1159–1186.

Stephan, K.E., 2004. On the role of general system theory for functional neuroimaging. J. Anat. 205, 443–470.

Stephan, K.E., Harrison, L.M., Penny, W.D., Friston, K.J., 2004. Biophysical models of fMRI responses. Curr. Opin. Neurobiol. 14, 629–635.

Stephan, K.E., Baldeweg, T., Friston, K.J., 2006. Synaptic plasticity and dysconnection in schizophrenia. Biol. Psychiatry 59, 929–939.

Stephan, K.E., Marshall, J.C., Penny, W.D., Friston, K.J., Fink, G.R., 2007a. Interhemispheric integration of visual processing during task-driven lateralization. J. Neurosci. 27, 3512–3522.

Stephan, K.E., Weiskopf, N., Drysdale, P.M., Robinson, P.A., Friston, K.J., 2007b. Comparing hemodynamic models with DCM. NeuroImage 38, 387–401.

Stephan, K.E., Kasper, L., Harrison, L.M., Daunizeau, J., den Ouden, H.E., Breakspear, M., Friston, K.J., 2008. Nonlinear dynamic causal models for fMRI. NeuroImage 42, 649–662.

Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., Friston, K.J., 2009a. Bayesian model selection for group studies. NeuroImage 46, 1004–1017.

Stephan, K.E., Tittgemeyer, M., Knosche, T.R., Moran, R.J., Friston, K.J., 2009b. Tractography-based priors for dynamic causal models. NeuroImage 47, 1628–1638.

Stephan, K.E., Penny, W.D., Moran, R.J., den Ouden, H.E., Daunizeau, J., Friston, K.J., 2010. Ten simple rules for dynamic causal modeling. NeuroImage 49, 3099–3109.

Stone, J.V., Kotter, R., 2002. Making connections about brain connectivity. Trends Cogn. Sci. 6, 327–328.

Uddin, L.Q., Kelly, A.M., Biswal, B.B., Xavier Castellanos, F., Milham, M.P., 2009. Functional connectivity of default mode network components: correlation, anticorrelation, and causality. Hum. Brain Mapp. 30, 625–637.

Uludag, K., Muller-Bierl, B., Ugurbil, K., 2009. An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. NeuroImage 48, 150–165.

Upadhyay, J., Silver, A., Knaus, T.A., Lindgren, K.A., Ducros, M., Kim, D.S., Tager-Flusberg, H., 2008. Effective and structural connectivity in the human auditory cortex. J. Neurosci. 28, 3341–3349.

Valdes-Sosa, P.A., 2004. Spatio-temporal autoregressive models defined over brain manifolds. Neuroinformatics 2, 239–250.

Valdes-Sosa, P.A., Sanchez-Bornot, J.M., Lage-Castellanos, A., Vega-Hernandez, M., Bosch-Bayard, J., Melie-Garcia, L., Canales-Rodriguez, E., 2005. Estimating brain functional connectivity with sparse multivariate autoregression. Philos Trans R Soc Lond B Biol Sci 360, 969–981.

Valdes-Sosa, P.A., Sanchez-Bornot, J.M., Sotero, R.C., Iturria-Medina, Y., Aleman-Gomez, Y., Bosch-Bayard, J., Carbonell, F., Ozaki, T., 2009. Model driven EEG/fMRI fusion of brain oscillations. Hum. Brain Mapp. 30, 2701–2721.

Valdes-Sosa, P.A., Roebroeck, A., Daunizeau, J., Friston, K., 2011. Effective connectivity: influence, causality and biophysical modeling. NeuroImage 58, 339–361.

van Leeuwen, T.M., den Ouden, H.E., Hagoort, P., 2011. Effective connectivity determines the nature of subjective experience in grapheme-color synesthesia. J. Neurosci. 31, 9879–9884.

Von Bertalanffy, L., 1969. General System Theory. George Braziller, New York.

Yacoub, E., Harel, N., Ugurbil, K., 2008. High-field fMRI unveils orientation columns in humans. Proc. Natl. Acad. Sci. U. S. A. 105, 10607–10612.

Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P.F., Feinberg, D., Adriany, G., Chaimow, D., Shmuel, A., Uğurbil, K., Yacoub, E., 2011. Mapping the organization of axis of motion selective features in human area MT using high-field fMRI. PLoS One 6, e28716.