New Approaches for Exploring Anatomical and Functional Connectivity in the Human Brain

Narender Ramnani, Timothy E.J. Behrens, Will Penny, and Paul M. Matthews

Information processing in the primate brain is based on the complementary principles of modular and distributed information processing. The former emphasizes the specialization of functions within different brain areas. The latter emphasizes the massively parallel nature of brain networks and the fact that function also emerges from the flow of information between brain areas. The localization of function to specific brain areas ("functional segregation") is the commonest approach to investigating function; however, an emerging, complementary approach ("functional integration") describes function in terms of the information flow across networks of areas. Here, we highlight recent advances in neuroimaging methodology that have made it possible to investigate the anatomical architecture of networks in the living human brain with diffusion tensor imaging (DTI). We also highlight recent thinking on the ways in which functional imaging can be used to characterize information transmission across networks in the human brain (functional and effective connectivity).

Key Words: Diffusion, magnetic resonance imaging, functional connectivity, human

A major challenge for neuroscience is to understand brain function in terms of connectional anatomy and the dynamic flow of information across neuronal networks. In nonhuman primates, synaptic connectivity between brain regions can be established by the injection of tracers into target brain areas and observation of the patterns of transport of tracers in the brain postmortem (see Ramnani and Miall [2001] for a description of recent advances; Kobbert et al 2000). Such methods can even identify connectivity between individual synapses, but their invasiveness makes them unsuitable for use in humans. Magnetic resonance imaging (MRI) now offers an entirely noninvasive, alternative approach. In white matter, water diffusion is highly directional ("anisotropic"), with preferential diffusion along the long axis of fibre tracts. With the application of large magnetic field gradients during image acquisition, MR images can be sensitized to the diffusion of water molecules within the voxel, and from these images we can compute the local direction of greatest diffusion. With these principal diffusion directions (PDD), the organization of major fibre tracts can be mapped. The resolution of these methods is still limited by the inherently low signal/noise ratio of MRI, and the methods cannot achieve the levels of spatial resolution of conventional anatomical tracer methods that can establish synaptic connectivity. For example, typical diffusion-weighted images used for tractography might have a voxel resolution on the order of $2.5 \times 2.5 \times 2.5$ mm^3, whereas conventional tract-tracing methods can track the projections of single axons (with spatial resolution measured in micrometers). Furthermore, a major limitation of these methods is that they do not distinguish between efferent and afferent projections. Nonetheless, data are easily acquired from individual subjects, and the analysis of tracts across the brain can proceed relatively quickly. We will describe selected recent advances in this field and give examples of how they can be used to determine aspects of the organization of the human brain.

A second, complementary approach is concerned with establishing the ways in which information is transmitted and integrated across brain networks. These are dynamic, context-dependent processes, in which variations in task demands lead to the preferential recruitment of some networks over others. Methods for analysis of these processes are based on the premise that functionally interacting regions will show correlated patterns of activity. Thus, simultaneously recording the activities of two groups of neurons in an animal preparation allows testing for conditions under which they become functionally coupled (Scannell et al 1995; Young et al 1994). The advantage of functional neuroimaging methods is that they can they can be used to detect activity not just in a limited set of areas but across the entire brain simultaneously. This makes it possible to examine the statistical relationships between the activities of not just two but of several areas across the brain. We will describe exciting new strategies for use of functional MRI (fMRI) data in the analysis of functional connectivity in the human brain. The review of diffusion tractography and functional mapping together highlights the possibility that future strategies for understanding interactions between regions of the human brain will benefit from integrating anatomically informed models of functional interactions.

Diffusion Tractography: Exploring the Connectional Architecture of the Human Brain

Recent advances in diffusion-weighted imaging and its derivative, diffusion tensor imaging (DTI), have brought to light the possibility of in vivo explorations of anatomical connectivity in the human brain. Magnetic resonance diffusion-weighted imaging sensitizes the nuclear MR signal to the random motion of water molecules along a single diffusion-encoding direction (Le Bihan 2003; Stejskal and Tanner 1965). By taking measurements along many such directions, it is possible to characterize the mean diffusion properties within a voxel. Diffusion tensor imaging then makes the assumption that this local diffusion might be explained by a three-dimensional Gaussian process and fits the diffusion tensor (Basser et al 1994) as its covariance matrix at each voxel. This tensor might be represented by a diffusion ellipsoid and, if the assumption of Gaussian diffusion holds true, the principal axis of this ellipsoid corresponds to the direction of greatest diffusion, or principal diffusion direction (PDD), and its

From the Centre for Functional Magnetic Resonance Imaging of the Brain (NR, TEJB, PMM), Department of Clinical Neurology, University of Oxford, Oxford; and Wellcome Department of Imaging Neuroscience (WP), Institute of Neurology, University College London, London, United Kingdom.

Address reprint requests to Dr. Narender Ramnani, Centre for Functional Magnetic Resonance Imaging of the Brain, Department of Clinical Neurology, University of Oxford, Headley Way, Oxford OX3 9DU, United Kingdom.

Received August 13, 2003; revised January 26, 2004; accepted February 3, 2004.

0006-3223/04/$30.00
doi:10.1016/j.biopsych.2004.02.004

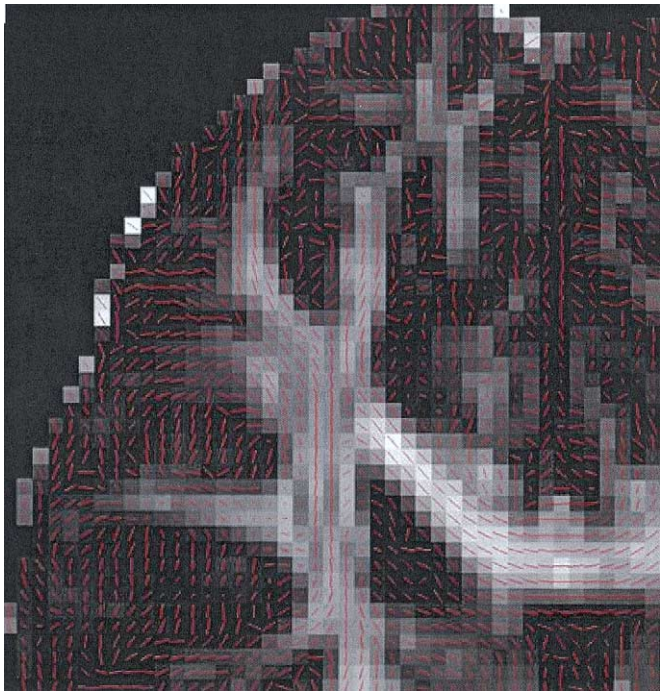

Figure 1. Local principal diffusion directions (discrete red lines) overlaid on fractional anisotropy in a typical diffusion-weighted image. Section shown is an axial slice through the splenium of the corpus callosum.

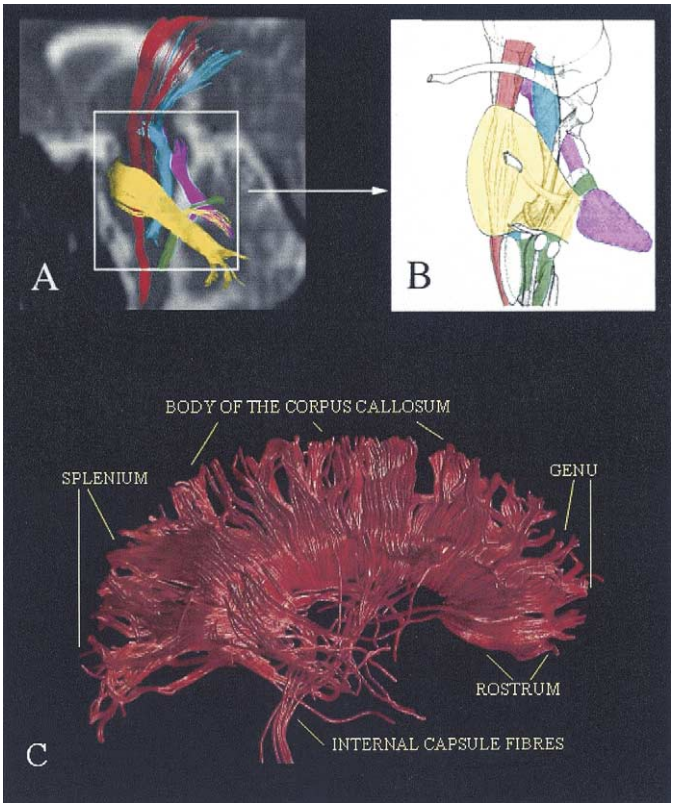

Figure 2. (A) In vivo diffusion tensor imaging tractography in the fibres around the human brainstem and cerebellar peduncles. **(B)** Drawing from postmortem dissection of the same fibre systems. In both **A** and **B**, the cortico-spinal tract is shown in red, the medial lemniscus in blue, the inferior cerebellar peduncle in green, and the superior cerebellar peduncle in pink. **A** and **B** are adapted from Stieltjes et al 2001 with permission from Elsevier. **(C)** In vivo reconstruction of the human callosal system (adapted from Catani et al 2002 with permission from Elsevier).

prolateness or anisotropy corresponds to the degree to which diffusion is preferred along this direction over other directions.

In tissue with a high degree of directional organization, diffusion is more hindered in some directions than others (see Le Bihan 2003 for a recent review). For example, in white matter, the PDD corresponds well with the dominant orientation of fibres within the voxel (Beaulieu and Allen 1994). Therefore, by visualizing the field of PDDs measured by DTI we can estimate the local orientation of major white matter tracts at each voxel (Figure 1). In fact, the directional patterns of vector fields displayed graphically are so compelling that it is tempting to make specific inferences of the spatial trajectories of white matter tracts from them.

This process is based on two implicit assumptions. First, we assume that the underlying vector field along the pathways is "smoother" than the resolution of the diffusion-weighted image (i.e., that the fibre architecture in every voxel is well-represented by a single vector, and if we followed the true PDD, we would confidently follow the tract). Second, we assume the measured PDD field to be a faithful representation of the true PDD field (i.e., that noise has a negligible effect on the measured PDD).

In the case of a voxel in which fibre tracts fork into two pathways, a single vector might show the mean direction instead of describing the true direction of either fibre. Nonetheless, so-called "streamlining algorithms" (Basser et al 2000; Conturo et al 1999; Mori et al 1999) have been highly successful in reconstructing major fibre systems in the deep white matter (Figure 2; Catani et al 2002; Stieltjes et al 2001). There is a close correspondence between in vivo DTI-based reconstructions and information from postmortem studies; however, these approaches also have limitations. For example, they are only able to define paths when diffusion anisotropy is high. So, whereas large deep white matter paths are well-defined, pathways toward the neocortex, where there can be considerable fibre divergence, and hence low diffusion anisotropy, are not. Another limitation is that DTI-based reconstructions express results qualitatively and do not quantitatively measure the strength or confidence in the pathways described, which makes between-subject comparisons reliant on qualitative analyses.

Recent work has started to tackle the assumptions discussed above. The first assumption is a relatively difficult one to address. The voxel resolution in diffusion-weighted images is orders of magnitude greater than the diameter of an axon. As we seek to identify increasingly finer pathways, it will inevitably be the case that some areas are poorly described by Gaussian diffusion, because such areas might be poorly described in terms of a single dominant fibre direction. This topic is the focus of much current research. For example, Tuch et al (2003) propose techniques for recovering arbitrarily complex distributions on diffusion within each voxel. This allows the investigators to resolve areas of complex fibre structure, such as crossing fibres. Figure 3 is a typical example of the result of q-ball diffusion imaging in areas with complex fibre architecture, showing the potential for resolving more complex fibre structure within a voxel. Note the multiple lobes on the distributions, for example at the crossing of corona radiata and superior longitudinal fasciculus.

The second assumption is easier to tackle. By taking repeated measurements along the same diffusion-encoding directions and "bootstrapping" these data to create many DTI data sets, Jones (2003) was able to quantify the uncertainty in the measured PDD

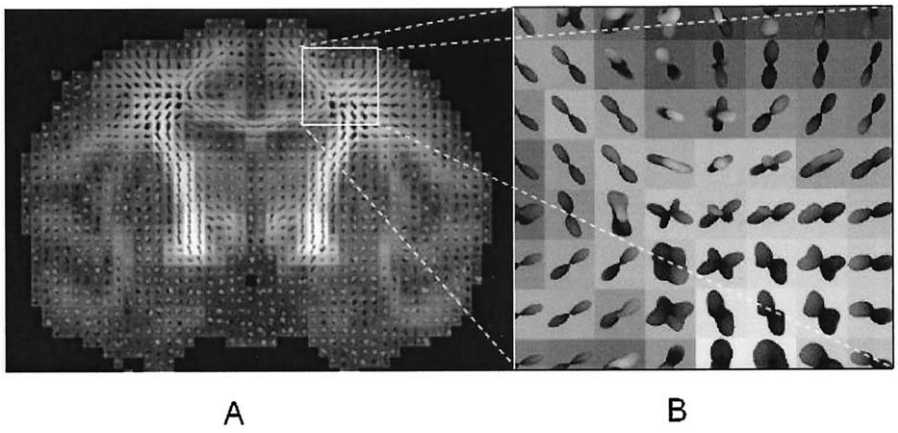

Figure 3. Q-ball diffusion imaging in the macaque monkey brain reveals fibre complexity within a voxel. Normalized orientational density functions are shown at each voxel. **(A)** A single coronal slice from macaque brain. **(B)** Detail of an area of crossing fibres in **A** (intersection of the superior longitudinal fasciculus, the corona radiata, and the corpus callosum). Image courtesy of D. Tuch.

field. Uncertainty is typically low in deep white matter fibres, thus explaining the reproducibility of streamlining results in these areas; however, uncertainty is high in areas with geometrically complex structure (such as crossing fibres). Behrens et al (2003a) propose a method for estimating this uncertainty from a single data set and propagating this local uncertainty on PDD through to a global probability density function (PDF) on the recovered connecting streamlines. Thus, pathways seeded in a given location might encounter regions of high uncertainty as they approach their targets but be able to progress into target areas, with the PDF spreading spatially to account for the uncertainty earlier in the pathway. This is illustrated in Figure 4A, which depicts thalamo-cortical pathways (note the broadening of PDFs in the approach to the cortex; see figure legend for details).

An important advantage in computing a PDF on the location of the pathway is that it is possible to express a level of confidence in the resulting projection. Instead of discretizing the PDF on a voxel-by-voxel basis, Behrens et al (2003b) computed the probability of projection between seed points in the thalamus and each of seven anatomically defined cortical masks (Figure 4B). Seeds were classified according to the masked cortical region with which they had the greatest probability of connection (Figure 4C). Figure 4C (inset) shows a schematic diagram of the thalamus subdivided into histologically defined nuclei. Color

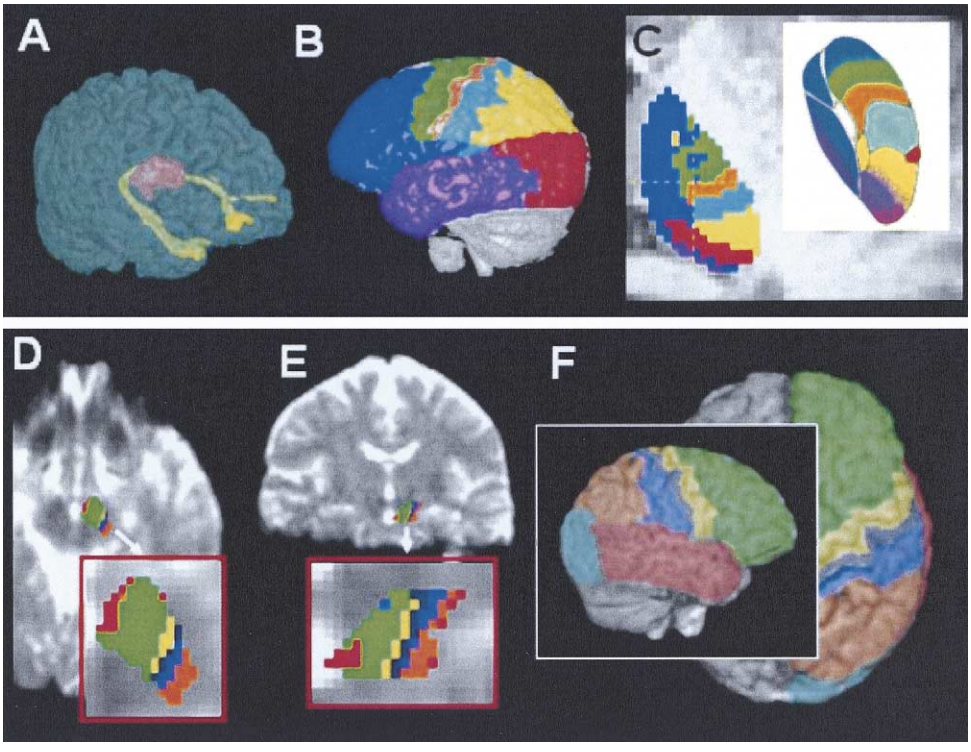

Figure 4. (A–C) Probabilistic tractography between thalamus and cortex. **(A)** Probabilistic connectivity distribution from medio-dorsal thalamus to the prefrontal cortex and temporal lobe. **(B)** Anatomically defined cortical masks. **(C)** Axial section through thalamus showing result of connectivity analysis. Probabilistic tractography (see text) was run from each seed voxel in thalamus, and the seed voxel was labeled according to the cortical zone in **B** with the highest probability of connection. Inset: schematic of thalamus with (overlaid in color) predictions from the monkey literature of the dominant cortical connections within each thalamic nuclear cluster. **(D–F)** Probabilistic tractography between cerebral peduncle and cortex. Axial **(D)** and coronal **(E)** section through the cerebral peduncle, parcellated according to the highest probability of connection from cortical zones defined in **F**. **A–C** adapted from Behrens et al 2003a.

overlays represent predictions from the monkey literature of the dominant cortical connection of each nuclear cluster. Figure 4C (background image) shows the result of classifying seed voxels in the thalamus to the cortical area with the highest probability of connection. The authors examined the probability values of connections between seed voxels and target areas and found areas in the thalamus with a high probability of a connection to more than one cortical mask and areas in which connection probability is small with all cortical masks.

Establishing the validity of diffusion-based fibre tracking methods is an important challenge, which can be accomplished by evaluating both between-subject reproducibility and by comparing results with those from conventional tract-tracing methods in the same brain. The former confers the advantage that between-subject analyses can be quantitative, but it might be complicated by the complexities associated with the spatial warping of imaging data from several cases into a common anatomical frame of reference (Xu et al 2003); deep white matter tractography often relies on qualitative comparisons between subjects. The latter is a particularly powerful approach to validation: nonhuman primate models might be used to demonstrate that DTI and conventional tract-tracing methods reveal common anatomical architecture when both methods are applied to the same brains. It was recently reported that the uptake of manganese through fibre tracts can be detected with MRI (Pautler et al 1998) and can therefore potentially be used as an anatomical tracer. Saleem et al (2002) reported that this method yields projection patterns that are comparable with those that result from the use of conventional tracers and histologic methods. The reliability of manganese-based methods to establish projections suggests that it should be used to validate DTI.

With increasing confidence in the methods, results can be interpreted in ways that provide novel anatomical information. In preliminary work, Ramnani and colleagues (unpublished data) applied the same methods to study the topography of the cortico-cerebellar system in the cerebral peduncles, which is part of the controversial debate on the cognitive functions of the cerebellum. This work confirmed that rostro-caudal gradients in the cerebral cortex are represented by medio-laterally organized cortico-cerebellar fibre tracts in the ipsilateral cerebral peduncles (See Figure 4D–F). The representation of fibres from the prefrontal cortex was very much larger than expected, which suggests an important role for the human cerebellum in processing information from the prefrontal cortex. The large differences in the neuroanatomical organization of this pathway in macaques and humans suggest that inferences about human neuroanatomy that are drawn on the basis of neuroanatomical studies in nonhuman primates might often be misleading. Diffusion tractography promises to play an increasingly important role in our efforts to understand the anatomical architecture of the human brain and its relation to that of other species.

Functional and Effective Connectivity: Exploring Models of Information Flow Through Networks

The functional mapping of different brain regions is the primary approach to the analysis of functional imaging data. Classic examples include the use of positron emission tomography (PET) by Zeki et al (1991) to localize color and motion centers of the human visual cortex (V4 and V5, respectively). More recently, these analyses have been augmented by functional integration analyses that describe how functionally specialized regions interact with each other. This can be thought of as the functional mapping of different brain pathways or networks. A recent example is the study by Buchel et al (1999), who found that the success with which a subject learned an object–location association task was correlated with the coupling between regions in the dorsal and ventral visual streams (Ungerleider and Mishkin 1982).

Functional Connectivity

Early analyses of functional integration used principal component analysis (PCA) to decompose neuroimaging data into a set of modes that are mutually uncorrelated, both spatially and temporally. The modes are also ordered according to the amount of variance they explain. By comparing the temporal expression of the first few modes with the variation in experimental task, a distributed functional system associated with that task can be identified (Friston et al 1993). A more sophisticated use of PCA occurs in the context of generalized eigenimage analysis (Friston 1997), in which the principal component is found that is maximally expressed in one experimental condition/population and minimally expressed in another (e.g., control vs. patient groups). More recently, independent component analysis (ICA) has been used to identify modes describing activity in a sparsely distributed network (McKeown et al 1998). Such PCA/ICA-based methods are referred to as analyses of functional connectivity, because they are data-driven transform methods that make no assumptions about the underlying biology. They are therefore of greatest practical benefit when it is not known which regions are involved in a given task and/or what is the underlying structural connectivity. In contrast, analyses of "effective connectivity" (described below) are based on statistical models that make anatomically motivated assumptions (e.g., knowledge of structural connectivity) and restrict their inferences to networks comprising a number of preselected regions. These analyses are hypothesis-driven rather than data-driven and are most applicable when one has knowledge of the relevant functional areas (e.g., from analyses of functional specialization). Detailed discussions of both approaches are found in Frackowiak et al (2003).

Structural Equation Modeling

Structural equation models (SEMs) were developed in the field of econometrics and were first applied to neuroimaging data by McIntosh and Gonzalez-Lima (1994). They comprise a set of regions and a set of directed connections. Importantly, a causal semantics is ascribed to these connections, whereby an arrow from A to B means that A causes B. Causal relationships are thus not inferred from the data but are assumed a priori (see Figure 5).

An SEM with particular connection strengths implies a particular set of instantaneous correlations between regions. One can therefore set the connection strengths so as to minimize the discrepancy between the observed and implied correlations and thereby fit a model to data.

If, for example, one partitions a given fMRI data set into those scans obtained under two different levels of an experimental factor, then one can attribute differences in connectivity to that factor and so conclude that a pathway has been activated. Structural equation models have, to date, been the most widely used model for connectivity analyses in neuroimaging (see Goncalves and Hull 2003), and we envisage that this will remain the case for experiments in which PET data are used.

There are, however, two major drawbacks to SEM. First, because SEM only makes use of information in correlation matrices, it is only possible to specify networks with a limited

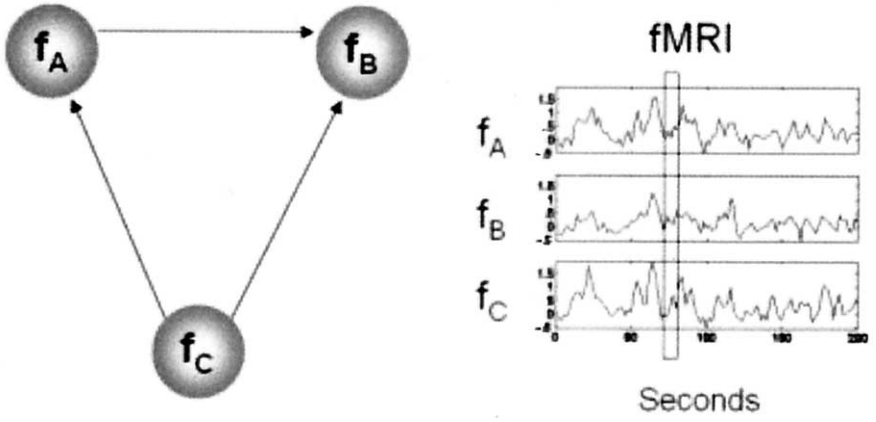

Figure 5. Structural equation models (SEMs) posit a set of causal relationships between variables. These can be shown graphically, for example, by the network in the left panel. The right panel shows a set of functional magnetic resonance imaging (fMRI) time series in which the superimposed narrow rectangle indicates that SEMs model the instantaneous correlations, that is, the correlation between regions at the same time point. Instantaneous activity is assumed to be the result of random fluctuations (i.e., activity that cannot be directly attributed to known experimental manipulation) and connections between regions. Changes in connectivity can be attributed to experimental manipulation by partitioning the data set.

number of connections. Such sparse structures that neglect, for example, reciprocal connections between regions are often biologically implausible and might result in poor model fits. Second, SEMs do not make use of temporal information—if the time indexes of the data were randomly permuted, SEM would give the same results.

Multivariate Autoregressive Modeling

To overcome these difficulties, Harrison et al (2003) have proposed the use of multivariate autoregressive (MAR) models for the analysis of fMRI data. An autoregressive approach is used to characterize structure in a time series, whereby the current value of a time series is modeled as a weighted linear sum of previous values. Multivariate autoregressive models extend this approach to multiple time series, so that the vector of current values of all regions is modeled as a linear sum of previous vector values. The optimal number of preceding time points can be found with Bayesian model order selection (see Figure 6).

In a MAR model, the dependencies between time points and between regions are characterized by a matrix of weighting values. Estimation of these weights is a noniterative linear fitting process. Thus, estimation is fast, which opens up the possibility of readily comparing connectivity models comprising different regions and connectivity patterns.

The parameters of a MAR model can be used to compute a number of further quantities, each of which can be used to describe network connectivity. These include coherences (correlation at particular frequencies), partial coherences (the coherence between two time series after the effects of others have been taken into account), phase relationships (the lag between two signals at a given frequency), and Granger causality (the dependence of region A on region B, as assessed by comparing two MAR models, one with the A-to-B connection and one without).

By partitioning an fMRI data set into different levels of a factor, one can then infer that pathways have been activated or that, for example, Granger causality between regions has changed. Multivariate autoregressive models are only beginning to be applied in fMRI but have a history of application in electroencephalography (EEG)/magnetoencephalography (Bressler and Scott Kelso 2001).

Dynamic Causal Modeling

Whereas SEM and MAR models were developed in other areas of science, dynamic causal modeling (DCM) (Friston et al 2003) has been specifically designed for the analysis of functional imaging data. Dynamic causal modeling posits a causal model, whereby neuronal activity in a given region causes changes in neuronal activity in other regions, via interregional connections, and in its own activity, via self-connections (see Figure 7). The neuronal activity in each region then gives rise to changes in blood volume, flow, and deoxyhemoglobin content. These then determine the blood oxygen level–dependent signal that is measured with fMRI. In DCM, these hemodynamic relationships are quantified by the Balloon model (Friston et al 2003).

Thus, DCM models neuronal connectivity, whereas SEM and MAR model correlations at the level of observed fMRI time series. Dynamic causal models are able to work at the neuronal level because they use a "forward model" (with hemodynamic parameters), relating neuronal activity to fMRI activity, and this model is

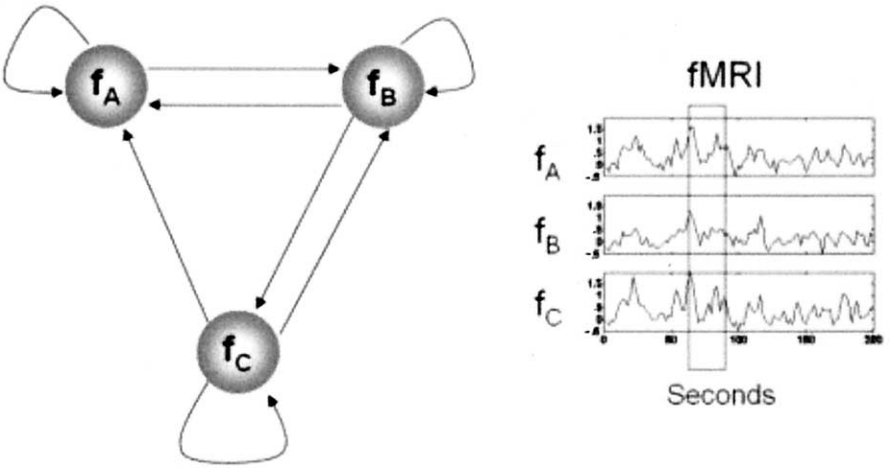

Figure 6. Multivariate autoregressive (MAR) models posit a set of causal relationships between variables as shown, for example, in the left panel. The self-connections highlight the fact that activity in each region is modeled as an autoregressive process. The right panel shows a set of functional magnetic resonance imaging (fMRI) time series in which the superimposed wide rectangle indicates that MAR models take into account the ongoing correlations, that is, the correlation between regions at the same and neighboring time points. Instantaneous activity is the result of random fluctuations, local dynamics, and connections between regions. Changes in connectivity can be attributed to experimental manipulation by partitioning the data set.

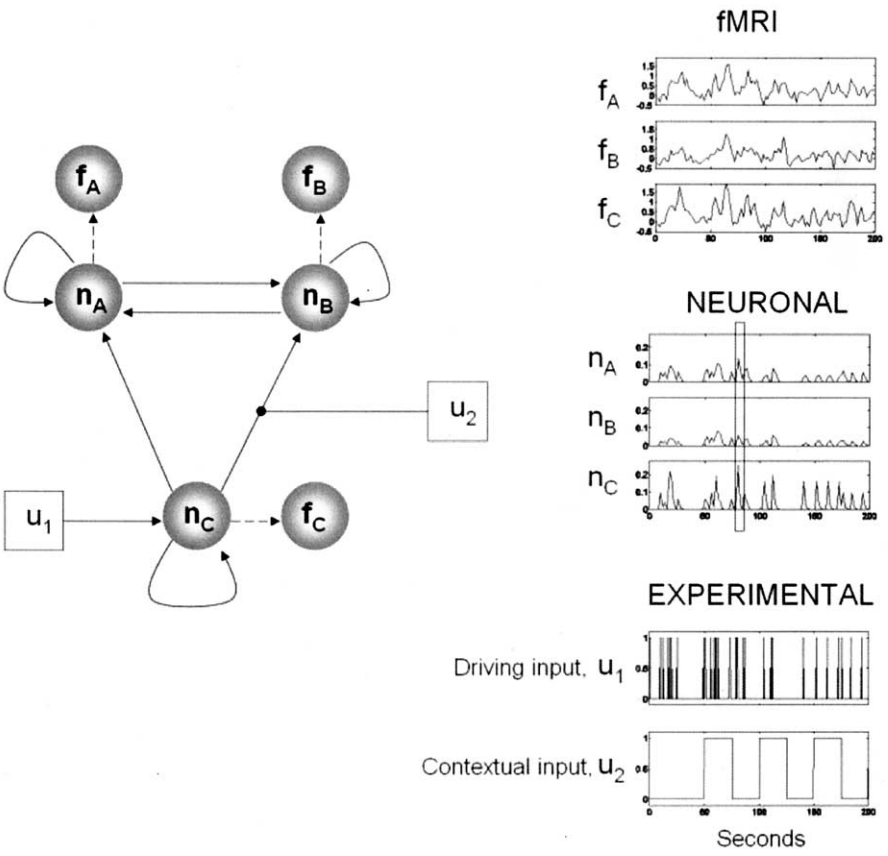

Figure 7. Dynamic causal modeling (DCM) models transient dependence in neuronal signals. Neuronal activity is the result of driving experimental input and neuronal dynamics, and changes in connectivity are directly attributable to experimental manipulation. This figure shows that input 1, a driving input, causes activity in region C, which in turn causes activity in regions A and B. This activity gradually decays according to a neurodynamic model that can be estimated from functional magnetic resonance imaging (fMRI) data. Input 2, a contextual input, changes the connectivity from the neuronal ensembles in region C to those in region B. This changes network activity and results in different observed hemodynamics (note, for example, that fMRI activity in region B is stronger and more correlated to activity in region A when input 2 is "high").

inverted during the model-fitting process. Another important distinction is that DCM explicitly models the effect of experimental manipulation on network dynamics. Mathematically, neuronal activity is described by a bilinear differential equation, whereby transient responses are initiated via driving external inputs, and the time constants of these transients can be altered via modulatory inputs. The strength of these driving and modulatory input connections (or "neurodynamic parameters") can be estimated from data.

A DCM is fitted to data by tuning the neurodynamic and hemodynamic parameters so as to minimize the discrepancy between predicted and observed fMRI time series. This takes place via an iterative nonlinear fitting process. A current limitation of DCM is that, because this model fitting is computationally demanding, one must restrict analyses to a small number of regions.

An example of an analysis with DCM is a study of whether category specificity effects in infero-temporal cortex are mediated by top-down or bottom-up activity (Mechelli et al 2003). We anticipate that the DCM approach rapidly will become widely used because it both 1) explicitly models how experimental manipulations cause network activity; and 2) models this activity at a neuronal rather than hemodynamic level, a level that is most appropriate for understanding information flow.

A second current limitation of DCM is that neurodynamics in each region are characterized by a single state variable ("neuronal activity"). This prohibits inferences that can be meaningfully linked to specific neurotransmitter systems, because these would require multiple state variables in each region that describe activity in excitatory and inhibitory subpopulations. The parameters of such models could only be identified with DCMs that use high temporal resolution data, such as from EEG. The development of such models therefore requires integration of information from fMRI (to determine where activity occurs) and from EEG (to determine when it occurs) and is an exciting area for future research that would significantly strengthen the bridge between data from imaging neuroscience and our understanding of the neurobiology underlying cognitive processing.

A Multidisciplinary Approach to Understanding Connectivity

This review has provided an overview of recently developed methods that permit investigations of anatomical and functional connectivity in the human brain. Although these methods have yet to reach the peak of their sophistication, it is clear that they have already made significant contributions to our understanding of how the human brain operates as a collection of networks. The same methods also promise to transform the ways in which we think about the underlying causes of neuropsychiatric conditions. For example, DTI has been useful in the identification of connectional abnormalities in fronto-parietal and fronto-temporal circuitry in schizophrenia (Burns et al 2003; see Lim and Helpern 2002 for a review). Investigations of functional connectivity have been useful in studies of neurotransmitter systems closely linked to schizophrenia. As a recent example, Honey et al (2003) demonstrated that dopaminergic drugs alter the functional connectivity between areas of the prefrontal cortex and interconnected regions of the striatum and thalamus.

In future work, the combined use of DTI and functional connectivity analyses will also serve to overcome important limitations. Methods for examining effective connectivity (e.g., DCM, as described above) often require the a priori specification of anatomical connectivity models in the system of interest, but

these are inevitably inaccurate because they are derived from nonhuman primate studies, and the connectivity between areas in the human brain is almost always unknown. The methods described here offer the prospect of using DTI to specify the anatomical model to inform functional connectivity analyses, not only in the same species but also in the same subjects.

NR and PMM are supported by the Medical Research Council (United Kingdom), TEJ is supported by the Engineering and Physical Sciences Research Council (United Kingdom), and WP is supported by the Wellcome Trust.

Basser PJ, Mattiello J, Le Bihan D (1994): Estimation of the effective self-diffusion tensor from the NMR spin echo. *J Magn Reson B* 103:247–254.

Basser PJ, Pajevic S, Pierpaoli C, Duda J, Aldroubi A (2000): In vivo fiber tractography using DT-MRI data. *Magn Reson Med* 44:625–632.

Beaulieu C, Allen PS (1994): Determinants of anisotropic water diffusion in nerves. *Magn Reson Med* 31:394–400.

Behrens TE, Johansen-Berg H, Woolrich MW, Smith SM, Wheeler-Kingshott CA, Boulby PA, et al (2003b): Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat Neurosci* 6:750–757.

Behrens TEJ, Woolrich MW, Jenkinson M, Johansen-Berg H, Nunes RG, Clare S, et al (2003a): Characterisation and propagation of uncertainty in diffusion weighted MR imaging. *Magn Reson Med* 50:1077–1088.

Bressler S, Scott Kelso JA (2001): Cortical coordination dynamics and cognition. *Trends Cogn Sci* 5:26–36.

Buchel C, Coull J, Friston KJ (1999): The predictive value of changes in effective connectivity for human learning. *Science* 283:1538–1541.

Burns J, Job D, Bastin ME, Whalley H, Macgillivray T, Johnstone EC, Lawrie SM (2003): Structural disconnectivity in schizophrenia: A diffusion tensor magnetic resonance imaging study. *Br J Psychiatry* 182:439–443.

Catani M, Howard RJ, Pajevic S, Jones DK (2002): Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *Neuroimage* 17:77–94.

Conturo TE, Lori NF, Cull TS, Akbudak E, Snyder AZ, Shimony JS, et al (1999): Tracking neuronal fiber pathways in the living human brain. *Proc Natl Acad Sci U S A* 96:10422–10427.

Frackowiak R, Friston KJ, Frith C, Dolan R, Price C, Zeki S, et al (2003): *Human Brain Function*. 2nd ed. Amsterdam: Elsevier Academic Press.

Friston KJ (1997): Characterizing distributed functional systems. In: Frackowiak R, Friston KJ, Frith C, Dolan R, Mazziota J, editors. *Human Brain Function*. Amersterdam: Elsevier Academic Press, 107–126.

Friston KJ, Frith C, Liddle FP, Frackowiak R (1993): The principal component analysis of large (PET) data sets. *J Cereb Blood Flow Metab* 13:5–14.

Friston KJ, Harrison L, Penny W (2003): Dynamic causal modeling. *Neuroimage* 19:1273–1302.

Goncalves MS, Hull DA (2003): Connectivity analysis with structural equation modeling: An example of the effects of voxel selection. *Neuroimage* 20:1455–1467.

Harrison L, Penny W, Friston KJ (2003): Multivariate autoregressive modeling of fMRI time series. *Neuroimage* 19:1477–1491.

Honey GD, Suckling J, Zelaya F, Long C, Routledge C, Jackson S, et al (2003): Dopaminergic drug effects on physiological connectivity in a human cortico-striato-thalamic system. *Brain* 126:1767–1781.

Jones DK (2003): Determining and visualising uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magn Reson Med* 49:7–12.

Kobbert C, Apps R, Bechmann I, Lanciego JL, Mey J, Thanos S (2000): Current concepts in neuroanatomical tracing. *Prog Neurobiol* 62:327–351.

Le Bihan D (2003): Looking into the functional architecture of the brain with diffusion MRI. *Nat Rev Neurosci* 4:469–480.

Lim KO, Helpern JA (2002): Neuropsychiatric applications of DTI—a review. *NMR Biomed* 15:587–593.

McIntosh AR, Gonzalez-Lima F (1994): Structural equation modeling and its application to network analysis in functional brain imaging. *Hum Brain Mapp* 2:2–22.

McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ (1998): Analysis of fMRI data by blind separation into independent spatial components. *Hum Brain Mapp* 6:160–188.

Mechelli A, Price CJ, Noppeney U, Friston KJ (2003): A dynamic causal modelling study on category effects: Bottom-up or top-down mediation? *J Cogn Neurosci* 15:925–934.

Mori S, Crain BJ, Chacko VP, van Zijl PC (1999): Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol* 45:265–269.

Pautler RG, Silva AC, Koretsky AP (1998): In vivo neuronal tract tracing using manganese-enhanced magnetic resonance imaging. *Magn Reson Med* 40:740–748.

Ramnani N, Miall C (2001): Expanding cerebellar horizons. *Trends Cogn Sci* 5:135–136.

Saleem KS, Pauls JM, Augath M, Trinath T, Prause BA, Hashikawa T, Logothetis NK (2002): Magnetic resonance imaging of neuronal connections in the macaque monkey. *Neuron* 34:685–700.

Scannell JW, Blakemore C, Young MP (1995): Analysis of connectivity in the cat cerebral cortex. *J Neurosci* 15:1463–1483.

Stejskal EO, Tanner JE (1965): Spin diffusion measurements: Spin echoes in the presence of time-dependent field gradient. *J Chem Phys* 42:288–292.

Stieltjes B, Kauffman WE, van Zijl PC, Frederickson K, Pearlson GD, Solaiyappan M, Mori S (2001): Diffusion tensor imaging and axonal tracking in the human brainstem. *Neuroimage* 14:723–725.

Tuch DS, Reese TG, Wiegell WR, Wedeen VJ (2003): Diffusion MRI of complex neural architecture. *Neuron* 40:885–895.

Ungerleider LG, Mishkin M (1982): Ingle DJ, Goodale MA, Mansfield RJW, editors. *Analysis of Visual Behavior*. Cambridge, MA: MIT Press, 549–586.

Xu D, Mori S, Shen D, van Zijl PC, Davatzikos C (2003): Spatial normalization of diffusion tensor fields. *Magn Reson Med* 50:175–182.

Young MP, Scannell JW, Burns GA, Blakemore C (1994): Analysis of connectivity: Neural systems in the cerebral cortex. *Rev Neurosci* 5:227–250.

Zeki S, Watson JD, Lueck CJ, Friston KJ, Kennard C, Frackowiak RS (1991): A direct demonstration of functional specialization in human visual cortex. *J Neurosci* 11:641–649.