

Special Issue: The Connectome – Feature Review

Functional connectomics from resting-state fMRI

Stephen M. Smith¹, Diego Vidaurre², Christian F. Beckmann^{1,3,4},
Matthew F. Glasser⁵, Mark Jenkinson¹, Karla L. Miller¹, Thomas E. Nichols⁶,
Emma C. Robinson¹, Gholamreza Salimi-Khorshidi¹, Mark W. Woolrich^{1,2},
Deanna M. Barch⁵, Kamil Uğurbil⁷, and David C. Van Essen⁵

¹ Oxford Centre for Functional MRI of the Brain (FMRIB), University of Oxford, Oxford, UK

² Oxford Centre for Human Brain Activity (OHBA), University of Oxford, Oxford, UK

³ Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The Netherlands

⁴ MIRA Institute for Biomedical Technology and Technical Medicine, University of Twente, Enschede, The Netherlands

⁵ Washington University School of Medicine, Washington University, St. Louis, MO, USA

⁶ Department of Statistics & WMG, University of Warwick, Coventry, UK

⁷ Center for Magnetic Resonance Research, University of Minnesota Medical School, Minneapolis, MN, USA

Spontaneous fluctuations in activity in different parts of the brain can be used to study functional brain networks. We review the use of resting-state functional MRI (rfMRI) for the purpose of mapping the macroscopic functional connectome. After describing MRI acquisition and image-processing methods commonly used to generate data in a form amenable to connectomics network analysis, we discuss different approaches for estimating network structure from that data. Finally, we describe new possibilities resulting from the high-quality rfMRI data being generated by the Human Connectome Project and highlight some upcoming challenges in functional connectomics.

Resting-state fMRI and the connectome: what's the connection?

rfMRI has been used to study spontaneous fluctuations in brain activity since it was first noted that the rfMRI time series from one part of the motor cortex were temporally correlated with other parts of the same functional network, even with the subject at rest [1]. Many other large-scale networks of correlated temporal patterns in the 'resting brain' have subsequently been identified. These patterns can be distinguished from each other because, although each has a relatively consistent time course across its set of involved regions, the different networks have distinct time courses [2]. These 'resting-state networks' (RSNs) persist even during sleep and under anaesthesia [3] and are consistent across subjects [4] and, to some extent, across species [5]. RSNs have also been investigated with other

modalities such as magnetoencephalography (MEG) [6,7], but most RSN research to date has used rfMRI. In addition to providing new information about the structure and function of the healthy brain, the study of RSNs has already been shown to have potential clinical value, providing rich and sensitive markers of disease [8]. Although there has been concern that some patterns of spatially-extended spontaneous signals may be of non-neural physiological origin, these concerns are increasingly being addressed [9]. It is now generally accepted not only that RSNs do reflect networks of brain function [10], but also that the extensive set of functional networks identified in the task fMRI literature (e.g., as encoded in the BrainMap meta-analysis database [11]) can be found in rfMRI data [12].

The biological significance of such a rich and continuously present set of spontaneous, correlated activities in the resting brain remains poorly understood. Presumably, the brain is continuously engaged in undirected cognitive activities (both conscious thought processes and subconscious activity such as learning/unlearning) and also responds to uncontrolled external stimuli. However, the high level of overall spontaneous activity measured in rfMRI and the correspondingly large energy expenditure [13] have been surprising to many. Linking rfMRI investigations of the macroscopic-scale functional connectome to other modalities should further our understanding of resting-state activity and functional connectomics. Other types of connectome-related information include:

- macroscopic structural connectomics from diffusion MRI data [14–19];
- the 'mesoscopic' connectome of long-distance connections studied at the cellular level [20–22];
- the 'microscopic' connectome of all neurons and synapses [23];
- electrophysiological measures [24–26];
- *ex vivo* histological mapping [27];
- covariance of anatomical measures such as cortical thickness [28];

Corresponding author: Smith, S.M. (steve@fmrib.ox.ac.uk).

Keywords: connectomics; resting-state fMRI; network modelling.

1364-6613/\$ – see front matter

© 2013 Elsevier Ltd. All rights reserved. <http://dx.doi.org/10.1016/j.tics.2013.09.016>

Glossary

Connectomics: the mapping of the brain's structural and/or functional connections.

Diffusion MRI: an MRI modality that allows for the tracking of the major white matter fibre bundles and hence the mapping of the macroscopic structural connectome.

Dynamic causal modelling (DCM): the estimation of effective connectivity in the brain through the application of biophysical models, applied to neuroimaging data such as fMRI and inferred on using Bayesian techniques.

Effective connectivity: the estimated direct connection strengths and the causal connection directionality between distinct macroscopic-scale functional regions.

Functional connectivity: (i) as opposed to effective connectivity, functional connectivity estimates whether there is any functional connection between functional regions, even if indirect (polysynaptic); or (ii) can also be used more generally to refer to any approach to connectivity modelling based on functional data, in contradistinction to structural connectivity.

Graph theory: the study of network characteristics. Once a functional or structural connectivity network matrix (or graph) has been estimated (and often thresholded to identify a binarised set of connections), graph theory can be applied to study characteristics of the network. For example, graph theory may be used to study subclusters in the network, to identify highly connected network nodes, or to investigate overall network efficiency in terms of information passing around the network.

Independent component analysis (ICA): a data-driven method for identifying structured processes in the data (e.g., where multiple voxels behave with the same time course). Spatial ICA separates distinct components from each other that are spatially independent, whereas temporal ICA separates distinct components that are temporally independent of each other.

Network matrix: the representation of functionally distinct brain regions as network nodes and the connectivities between these as network edges. The strengths of the estimated network edges are placed as elements in an $N_{\text{nodes}} \times N_{\text{nodes}}$ network matrix that is a compact representation of the entire network.

Nonstationarity: variability in any given statistic (e.g., signal variance changing over time). In rfMRI, nonstationarity is most commonly used to refer to the changing of correlation strength over time, which may inform about functional network dynamics.

Parcellated connectome: another term for the network matrix, assuming that the network nodes are defined as the parcels from a parcellation of the brain's grey matter.

Partial correlation: a variant of correlation that attempts to infer the direct connections between nodes from analysis of correlations between multiple nodes' time series. Before correlating any two nodes' time series, all other time series are first regressed out of the two in question. Partial correlation is an approximation to solving a structural equation model, but cannot estimate causal directions.

Resting-state functional MRI (rfMRI): an MRI modality that measures spontaneous temporal fluctuations in brain activity (i.e., with the subject 'at rest'). rfMRI is primarily used to estimate connectivity in the brain, given that functionally connected areas have related spontaneous time series.

Resting-state networks (RSNs): functional networks in the brain that are most commonly estimated from rfMRI data.

Sparsity: reference to the existence of a large number of zero (or close to zero) values in a set of parameters. For example, a network matrix (of edge strengths) might be sparse, which would indicate that many node pairs are not directly connected.

Structural equation modelling (SEM): the estimation of effective connectivity through statistical modelling of, for example, neuroimaging time series data, as opposed to explicit biological and/or physical modelling (for example, as applied in DCM). In this context, SEM is most commonly applied to functional time series data and the term 'structural' here refers to the type of modelling and not the use of structural connectivity data (such as from diffusion MRI).

- task fMRI, behavioural measurements, and genotyping [10,11,29,30].

Early rfMRI studies typically characterised functional connectivity via a small number of large-scale spatial maps [1,2]. By contrast, the nascent field of 'connectomics' [31] generally attempts to study brain connectivity in a different way, first identifying a number of network nodes (functionally distinct brain regions) and then estimating the functional connections (network edges) between these nodes (Figure 1). To generate nodes, parcellation of the brain is often conducted by

clustering together neighbouring voxels (3D pixels) on the basis of similarity of their time series. This typically yields a large number of non-overlapping parcels, with a single contiguous group of voxels in each parcel or node, and is then generally referred to as a 'hard parcellation' [32,33]. Another approach to generating nodes involves high-dimensional independent-component analysis (ICA) [34]. Using ICA, each node is described by a spatial map of varying weights; each map may overlap with other nodes' maps and may span more than one group of contiguously neighbouring points. Network edges (connections between nodes) are estimated by comparing the fMRI time series associated with the nodes (e.g., the average time series of all voxels in a parcel). In some approaches, the directionality of these connections is estimated in an attempt to infer the direction of information flow through the network (see detailed discussion and references in [35]). As a result, brain connectivity can be represented as a 'parcellated connectome', which can be visualized simply as an $N_{\text{nodes}} \times N_{\text{nodes}}$ network matrix or a graph (explicitly showing nodes and the strongest edges) or using more sophisticated visualization approaches that embed nodes and edges into spatial representations of the brain [36].

rfMRI acquisition and image-processing overview

fMRI data (both task based and resting state) is acquired as a series of volumetric images over time, with each image generally taking 2–3 s to acquire. rfMRI data is typically acquired for 5–15 min, with the subject asked to 'lie still, think of nothing in particular, and not fall asleep'. The fMRI acquisition is tuned such that the image intensity reflects local blood flow and oxygenation changes resulting from variations in local neural activity [37]. To achieve this sensitivity, and to acquire the fMRI data rapidly, it is common to utilise echo planar imaging (EPI) [38], which acquires the data one 2D slice at a time. Standard acquisitions working at a magnetic field strength of 3 T can achieve a temporal resolution of 2–3 s with a spatial resolution of 3–5 mm. More recently, faster acquisitions have emerged. For example, multiband accelerated EPI acquires multiple slices simultaneously [39,40]. Such approaches enable major improvements in spatial and/or temporal resolution; for example, acquiring data with 2 mm spatial resolution in less than 1 s. Higher temporal resolution of the fMRI data can improve overall statistical sensitivity and also increase the information content of the data (e.g., in terms of reflecting the richness of the neural dynamics) [41,42], although the sluggish response of the brain's haemodynamics (to neural activity) will ultimately place a limit on the usefulness of further improvements in temporal resolution.

A four-dimensional rfMRI dataset requires extensive preprocessing before resting-state network analyses can be conducted. The preprocessing reduces the effects of artefacts (such as subject head motion and non-neural physiological signals), spatially aligns the functional data to the subject's high-resolution structural scan, and may subsequently align the data into a 'standard-space' reference coordinate system; for example, based on a population-average brain image. A standard sequence of processing steps [43,44] is as follows.

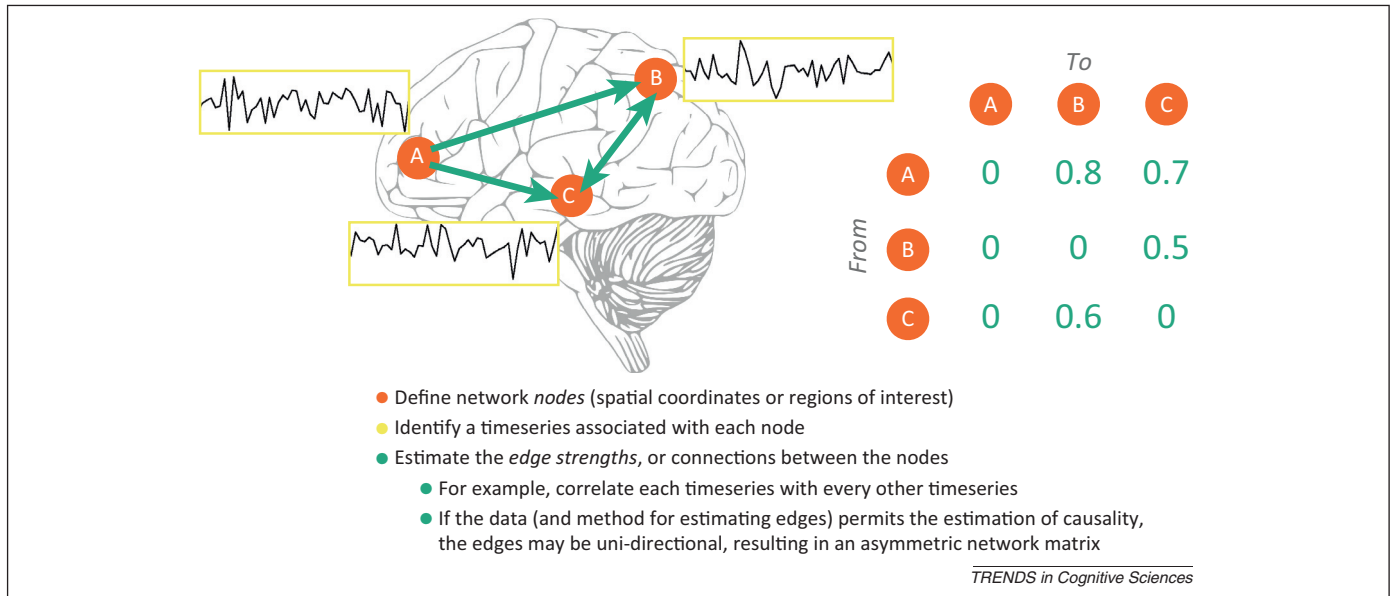


Figure 1. The main steps in taking resting-state functional MRI (rfMRI) data (with an activity time series at every point in the brain), identifying network nodes, and then estimating network edges.

- Realign each time point's image to a reference image, reducing the effects of subject head motion over the duration of the rfMRI acquisition.
- Correct the data for MRI spatial distortions.
- Remove non-brain parts of the image.
- Estimate the alignment transformations between the rfMRI data and the same subject's high-resolution structural image and between the structural image and a population-average brain image.
- Optionally, map the cortical data from the 3D voxel matrix ('volume based') onto the vertices of a cortical surface representation ('surface based'), in which a surface mesh follows the intricate convolutions of the cortical sheet. This aids visualization and enables better intersubject alignment (see below).
- Optionally, apply spatial smoothing to improve signal-to-noise ratio and ameliorate the effects of functional misalignments across subjects. In the best datasets and when using the most advanced methods for intersubject alignment, smoothing can be minimised. Unless smoothing is constrained to act within the cortical sheet, it will cause undesirable mixing of signal across tissue compartments and across sulcal banks between functionally distinct regions.
- Apply filtering to remove the slowest temporal drifts in the data.
- Remove other artefacts.
- regression of confound time series out of the data; for example, those derived from
 - average white matter and/or ventricle time series
 - head motion parameters (to further remove residual motion-related artefacts)
 - separately measured cardiac and breathing signals
 - global-average time series (although many researchers consider this to be a blunt tool that makes the interpretation of the final correlations difficult [44]);
- removal of corrupted time points [45];
- data-driven structured noise removal, using ICA with automated component classification to remove remaining artefacts [46];
- filtering out the highest temporal frequencies (this is commonly applied when more targeted artefact removal approaches are not available, because the balance between signal and noise is expected to be worst at the highest frequencies; related to this, it is widely presumed that resting-state signals of interest are fundamentally low frequency, but, as discussed below, there is increasing evidence that there is useful signal at relatively high frequencies [47]).

This last stage – the removal of other artefacts in the data – includes a diversity of commonly used approaches. Effective artefact removal is particularly important for resting-state analyses, which rely fundamentally on correlations between different voxels' time series, because these will be corrupted by artefacts that span multiple voxels. (By contrast, task fMRI has the advantage of fitting a prespecified temporal model, which provides greater robustness against artefactual influences.) Artefact clean up [9] can involve one or more of:

Once preprocessing is complete, the data is ready for some form of connectivity analysis. In early rfMRI studies, connectivity was often summarised by one (or a few) spatial maps spanning the whole brain. For example, in seed-based correlation (Figure 2), a single seed voxel or region of interest is selected, such as a 5-mm radius sphere centred in the precuneus. The average time course from this seed region is extracted and every other voxel's time series is correlated against it, creating a correlation-strength map spanning all of the brain. Such an approach contains fine spatial detail, but only informs about average correlation with the selected seed. More information is obtained by low-dimensional ICA decomposition, for example, reducing the data to ten to 30 independent spatial maps, each of which is analogous to a distinct seed-based correlation map. This therefore

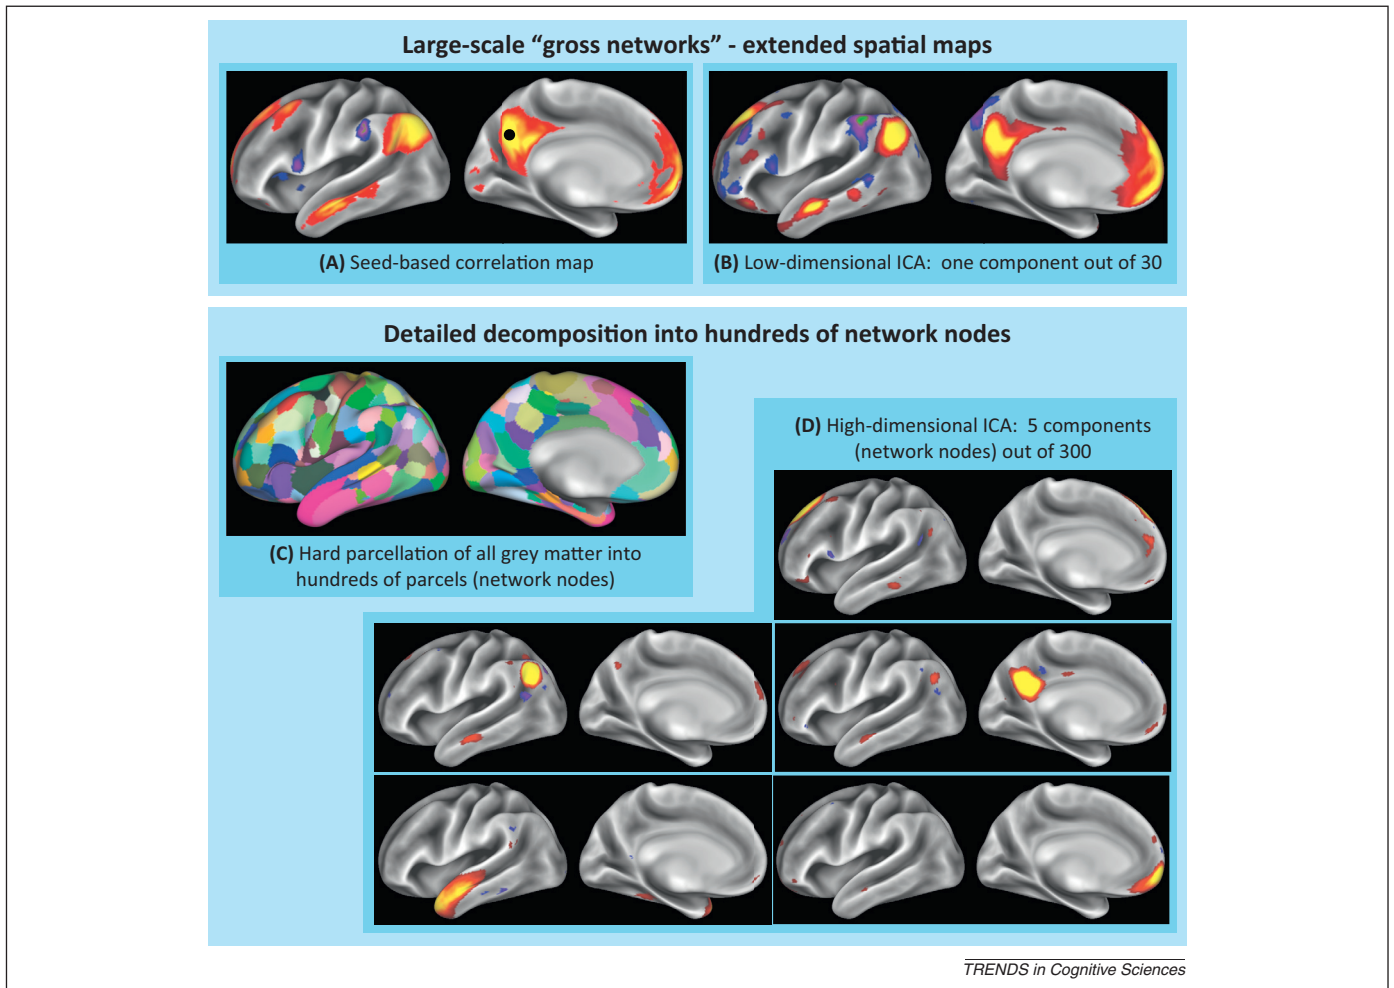


Figure 2. Examples of different approaches to connectivity analysis. All results are derived from group-average connectivity from the first 131 subjects' resting-state functional MRI (rsfMRI) data released publicly by the Human Connectome Project (HCP) (available from <http://www.humanconnectome.org>) and displayed on lateral and medial views of the inflated left hemisphere. (A) and (B) show large-scale networks represented by extended spatial maps; (A) shows a single seed-based correlation map, with the seed (marked with a dot) placed in the precuneus, whereas (B) shows the spatial map from a single component from a low-dimensional (30) independent-component analysis (ICA) decomposition. (C) and (D) show high-dimensional decompositions of the data into hundreds of network nodes; (C) shows an exemplar 'hard' parcellation of the grey matter into hundreds of non-overlapping parcels, whereas (D) shows several ICA components from a high-dimensional (300) ICA decomposition. Seed-correlation maps may contain negative values (seen here in blue/purple), which indicate cortical regions that are anticorrelated over time with the seed. Likewise, ICA components' spatial maps can contain negative values, which indicate anticorrelation with that component's associated time series. In the case of high-dimensional parcellations, anticorrelations would probably be seen not in the spatial maps, but as negative correlations between the time series associated with different parcels. The cortical surface renderings were generated using the Connectome Workbench display tool (<http://www.humanconnectome.org/connectome/connectome-workbench.html>).

generates a richer description of multiple networks in the brain, but at the expense of no longer dictating in advance the regions with which the connectivity maps are related.

In contrast to seed-based correlation and low-dimensional ICA, high-dimensional parcellation into many nodes (potentially hundreds) allows a richer analysis of the network connections; by shifting the emphasis from large-scale maps with fine spatial detail into a network description of nodes and edges, new information can be obtained. For example, whereas two large-scale networks might have some functional interaction (seen as non-zero correlation between their associated time series), a detailed nodes+edges network model may reveal which nodes (subparts of the large-scale networks) are responsible for the correlations seen between the larger-scale networks. Put another way, detailed network modelling facilitates analysis of both functional specialisation (investigating the functional connectivities of each node separately) and

functional integration (investigating how nodes interact with each other and form communities of functionally related clusters of nodes).

'Mapping the functional connectome' may be regarded as taking a nodes+edges approach to connectivity modelling. To combine or compare connectomes across subjects, it is important to strive for 'the same' parcellation in each subject; comparisons between two network models are inherently flawed if they are derived from non-corresponding sets of nodes. One simple approach to this problem is to generate a group-level parcellation and then impose this parcellation onto each subject. If every subject has been transformed into the same space as part of the preprocessing, this is conceptually straightforward. In reality, accurate alignment of functionally corresponding cortical parcels is an exceedingly challenging problem, largely owing to individual variability of cortical folding patterns along with variability of functional parcel locations relative to these folds (see below).

Once a parcellation has been applied (e.g., as a set of parcel masks) to a given subject's dataset, each parcel can then be assigned a representative time series based on that subject's rfMRI data, for example, by averaging the time series from all voxels within a parcel. From the resulting $N_{\text{timepoints}} \times N_{\text{nodes}}$ data matrix, one can then compute the subject-specific $N_{\text{nodes}} \times N_{\text{nodes}}$ parcellated connectome matrix; for example, by correlating each time series with every other. However, correlation is just one approach (albeit the simplest and most commonly used) for inferring the network edges (elements in the network matrix). The strengths and weaknesses of more sophisticated approaches are considered in the next section.

Functional network modelling

We now discuss in more detail the estimation of network edges, given a set of nodes' time series. The simplest method, correlation between the time courses of any two brain regions, allows one to infer whether the regions are functionally connected, although many factors other than the direct anatomical node-to-node connection 'true strength' can affect correlation coefficients, including variations in signal amplitude and noise level [48]. Furthermore, correlation cannot reveal anything about causality or even whether connectivity is direct versus indirect [49]. A common implicit assumption – that correlation is

unambiguously indicative of a direct connection – creates a major problem for network modelling (and graph theory applied to networks estimated from rfMRI) [35]. The distinction between simple correlation and trying to estimate the underlying, direct, causal connections (sometimes referred to as the distinction between functional and effective connectivity [50]) is important for deciphering the underlying biological networks. For example, in a three-node network $A \rightarrow B \rightarrow C$, all three nodes' time series will be correlated, so correlation will incorrectly estimate a fully-connected network (including an $A-C$ connection). However, another simple estimation method, partial correlation, aims to more accurately estimate the direct connections (though not their directionalities). In the three-node network example, this works by taking each pair of time series in turn and regressing out the third from each of the two time series in question before estimating the correlation between the two. (For more than three nodes, all of the other $N_{\text{nodes}} - 2$ nodes are regressed out of the two under consideration.) In this case, regression of B out of A and C removes the correlation between A and C and hence the spurious third edge ($A-C$).

A wide range of different network modelling approaches can be placed along a spectrum (Figure 3). This starts with neural-level brain simulations at one end [51], includes network modelling methods in the middle that can be

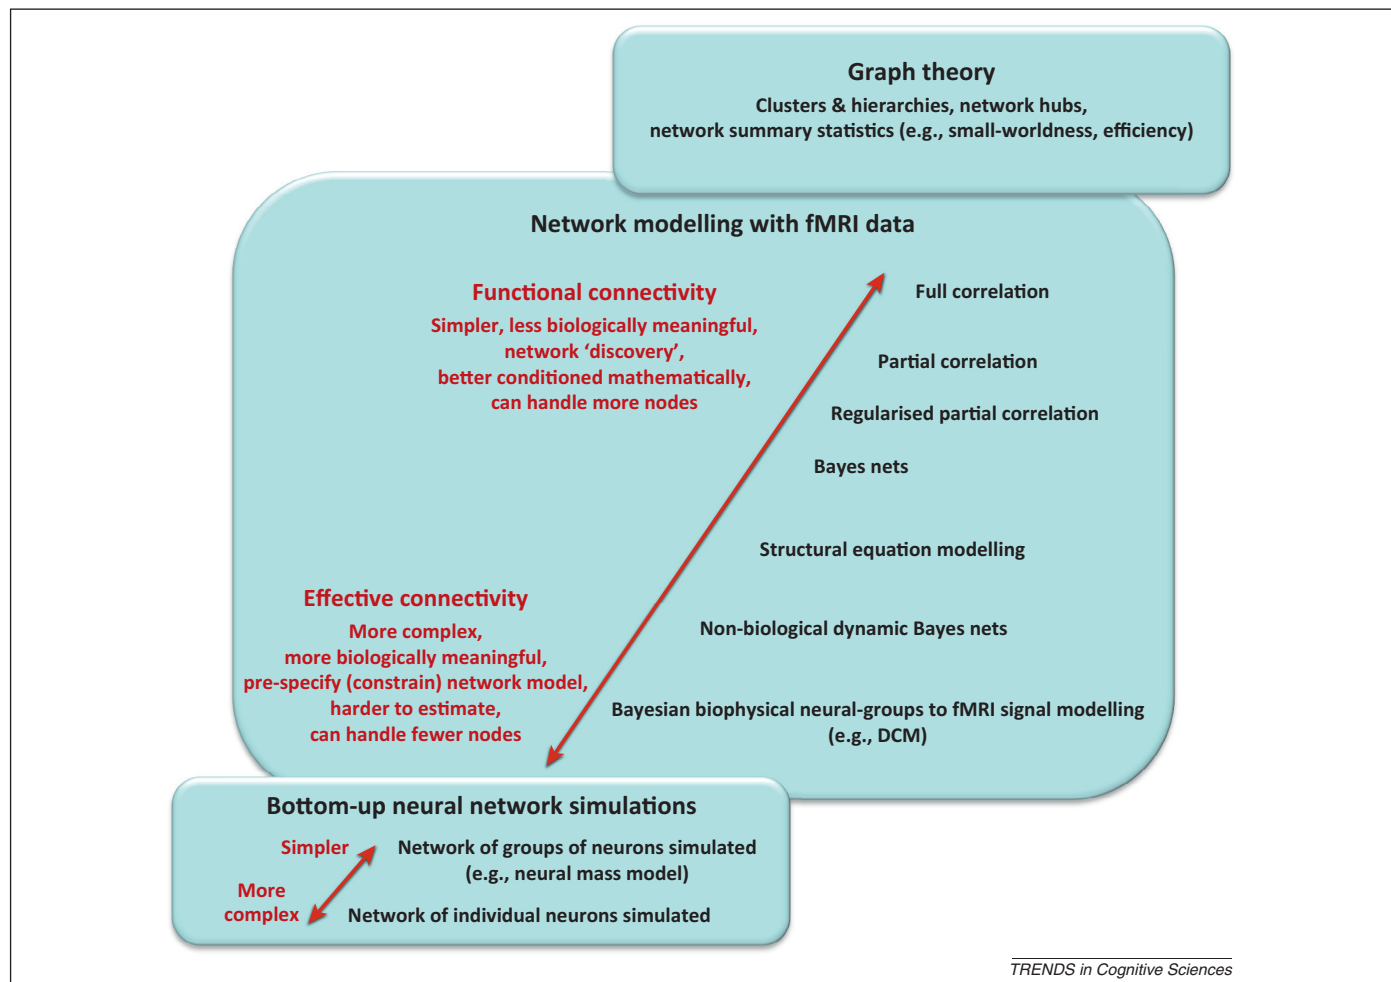


Figure 3. Schematic of relationships between various network modelling analyses for/from functional MRI (fMRI). Adapted from [12].

applied to real rfMRI data (with full and partial correlation sitting at the simple extreme), and ends with abstract graph-theoretical network analyses that require the network to have already been estimated.

Of the methods for connectivity modelling that have been applied to rfMRI data, at one extreme are complex models of effective connectivity with many parameters, each representing a biological or physical measure such as average neuronal activity and (separately) the haemodynamic response to neural activity; this model is ideally fit to data using probabilistic (e.g., Bayesian) methods. One such approach is dynamic causal modelling (DCM) [52]. Not only is the model complex, but so is the Bayesian inference method, which is computationally more sophisticated than simple ‘point estimate’ model fitting. Estimating quantitative, biologically meaningful parameters is clearly of great value if we want to find and interpret changes in functional networks, for example, as a result of disease. Moving towards the simpler end of the modelling spectrum, with methods such as structural equation modelling [53], model parameters refer to statistical relationships between data variables (e.g., causations or correlations between node time series) as opposed to underlying biological or physical quantities. At the simplest extreme are mathematically straightforward methods such as correlation. The simpler methods are in general more robust (with respect to fitting the model to the data) and faster to compute. Related to this, and the fact that they have many fewer parameters to estimate, the simpler methods can handle a much larger number of network nodes than the more complex methods; they do not require the scope of possible network models to be pre-constrained, making them computationally feasible for attempting network discovery. However, simpler methods provide descriptions of the data rather than being directly tied to underlying, more biologically interpretable network parameters.

These considerations motivate a desire to work at the more complex, biophysically interpretable level, but they also restrict the practicality of the analyses that are feasible. For example, although DCM was originally developed for task data, it has recently been extended [54] to allow modelling of resting-state data and to search across a wide range of possible network matrix models (rather than requiring the prespecification of just a few). However, this is currently practical (in terms of both computational expense and mathematical robustness) only for networks with a small number of nodes (<10). A major challenge for functional connectomics will be to enable application of biologically interpretable models using large numbers of nodes in a robust and practical way.

Currently, if one wishes to conduct connectomics network modelling using a reasonably large number of nodes (50–500), one pragmatic compromise is to use partial correlation. The estimation of partial correlation effectively involves inverting the full correlation matrix, a process that is often quite unstable depending on the quality of the data and the number of original time points. Hence, improved estimation can often be achieved by ‘regularising’ the estimated inverse covariance matrix. For example, L1-norm regularisation of the inverse covariance matrix [55]

shrinks the estimated values in the partial correlation matrix so that very small, noisy values are forced to zero and all other values are better estimated. In recent work involving networks of simulated fMRI time series with up to 50 nodes, such methods performed the most accurately [56]. These methods also scale well, handling hundreds of nodes given sufficient data (primarily, a large number of time points). Additionally, it may be of value to apply regularisation across different subjects’ network matrix estimation, utilising the similarities in network structure across the population to improve within-subject estimation [57].

The directionality of edges (causality) is often investigated, but in general it is harder to estimate the direction of a connection from fMRI data than to detect whether an edge exists [56]. Many methods, such as correlation, give no directionality information at all. Those that do attempt to estimate directionality fall into several classes, as follows.

- Temporal lag between pairs of time series [58] is often used, although it is arguably too confounded by the slow and variable haemodynamics to be useful with fMRI data [56] and is better suited to electrophysiological modalities.
- Bayes nets model the full set of covariances (or conditional dependencies) between all nodes’ time series [59] and show some promise.
- Non-Gaussianities in the time series [60,61] provide a distinct, potentially valuable source of causality information.

Because most brain connections are bidirectional [62], simplistic views of causality may not be appropriate when considering rfMRI-based connectomics. Furthermore, causality estimation when the brain is ‘at rest’ may involve many different functional processes mixed together and would be estimated as the average ‘relative causality’ over all possible spontaneous fluctuations. This may be so different from what occurs during individual focussed tasks that dominant average causalities found from rfMRI may not relate meaningfully to the route by which information flows around the ‘brain network’ when triggered by distinct external events. Finally, observational data (such as rfMRI) is in general not a robust and safe tool for inferring causality, compared with interventional studies (such as task fMRI experiments) [63]. It may therefore be profitable to utilise task fMRI in a confirmatory role after an exploratory study using rfMRI, when attempting to study causalities in the brain.

At the highest network modelling level, many efforts fall in the domain of graph theory [64,65] (see also van den Heuvel and Sporns, this issue). This includes: the study of node clustering and hierarchies; the study of hubs (nodes, or clusters of nodes, that are particularly highly connected to other parts of the network); and deriving global network summary metrics such as small worldness (looking at how the communication and clustering acts over multiple scales) or measures of general network efficiency. Naturally, these techniques are dependent on accurate network modelling at the lower level (carefully constructed network matrices); the abstraction of these metrics from the

underlying data makes it challenging to evaluate results, regardless of how advanced or conceptually elegant a given graph-theoretical measure may be. One specific risk is the use of inappropriate node definition [56,66], such as a gross structural atlas-based parcellation that corresponds poorly to functional boundaries and results in network matrices whose neurobiological interpretability is limited (as is therefore any graph-theory analysis fed by such network matrices). Another concern is that it is frequently the (thresholded) full correlation matrix that is input into graph-theoretical analyses rather than an estimate of direct network connections; interpretations of graph theory-based measures (such as communication path length) are often, as a result, somewhat questionable. A third concern is that graph theory can abstract the network matrix to a very high degree (e.g., summarising an entire study down to a single measure representing overall network efficiency, or small worldness) and any apparent change in this measure (e.g., between patients and controls) might not reflect a genuine change in the brain connectivity but rather any of a myriad of potential confounds (e.g., factors as basic as systematic group differences in head motion or heart rate).

Despite such concerns, graph-theory investigations of network topological properties should become more meaningful as improved data and better network modelling methods yield more accurately estimated sets of network matrices. One interesting challenge for network modelling will be to ascertain whether information flowing between multiple nodes really is being passed around the network; even if a method such as partial correlation correctly shows a direct functional connection between node A and node B, and also a connection between node B and node C, it will still be important to know whether information originating in node A does reach C. For example, a potential confound would be if two such direct connections were never active at the same points in time; it cannot be assumed that the brain's connections are static. A further complication arises if brain activity in one region modulates the connection strengths between other nodes. This renders analyses based on assumptions of linearity inaccurate, but should add extra richness to the data that more sophisticated analysis methodologies may be able to take advantage of. Other important issues for future network modelling research, such as dynamic network estimation and the estimation of overlapping clusters of nodes, are discussed below.

The Human Connectome Project and future directions for macroscopic functional connectomics

In 2009, the National Institutes of Health (NIH) announced a program targeted at characterising the human connectome and its variability using cutting-edge neuroimaging methods, seeking applications that would accelerate advances in imaging technologies and apply these advances to a large population of healthy adults. In 2010, the NIH awarded Human Connectome Project grants to two consortia, one led by Washington University, the University of Minnesota and the University of Oxford (referred to below as the HCP) [67] and the other led by Massachusetts General Hospital (MGH) and the

Box 1. Macroscopic *in vivo* connectome projects

Here we list the major macroscopic-level connectome projects that are disseminating data publicly.

- As described in the main text, the Washington University–University of Minnesota–Oxford (WU–Minn) HCP is generating high-quality and high-resolution 3T rfMRI and diffusion MRI data from over 1000 healthy adults, as well as task fMRI, genotyping, and behavioural data. MEG and 7T MRI data will be acquired on a subset of subjects.
- The MGH–UCLA HCP is using extremely powerful MRI gradients to generate leading-edge 3T diffusion MRI data.
- The Developing Human Connectome Project (dHCP), led by King's College London, Imperial College London and Oxford, will map functional and structural connectomes in 1000 babies (*in utero* and after birth, from 20 to 44 weeks' post-conceptual age) using rfMRI and diffusion MRI. It will also conduct 25 *post mortem* diffusion MRI scans at 7 T.
- The Thousand Functional Connectomes (and related) projects [81,82], led by Mike Milham, are seeking to generate even larger numbers of subjects than the HCP, bringing together data from a wide range of imaging studies and scanners and covering a wider range of subject groups, including a range of pathologies. This growing dataset comes with the caveat of having greater heterogeneity of scanning parameters (and lower data quality) than the HCP, although some of the most recent additions are utilising EPI accelerations and should approach the quality of HCP data. Hopefully, such larger, heterogeneous databases of connectivity data, although not supporting the most sophisticated analysis techniques, will complement studies such as the HCP by being able to find gross imaging phenotypes and conducting very large-*N* subject–pathology correlations.

University of California, Los Angeles (UCLA) [16]. The latter is concentrating on leading-edge structural connectomics (via diffusion MRI) and so is not discussed further in this review (see also Box 1). The former is generating and sharing both functional and structural connectome data. In this section we give an overview of the HCP, present examples of connectomic investigations using publicly released HCP rfMRI data, and discuss how these developments raise exciting possibilities for functional connectomics.

The HCP is generating a detailed *in vivo* mapping of functional connectivity in a large cohort (over 1000 subjects) and is making these datasets freely available for use by the neuroimaging community (over 200 subjects' datasets are already acquired and publicly released, available via <http://www.humanconnectome.org>). Subjects are drawn from a population of healthy adult twins and their non-twin siblings (in the age range 22–35 years), thus aiding the study of the influence of heritability and environmental factors on the connectome. From each subject, 1 h of whole-brain rfMRI data is acquired at 3 T (in two pairs of 15-min runs on separate days). A spatial resolution of $2 \times 2 \times 2$ mm and temporal resolution of 0.7 s are achieved utilising an EPI acceleration factor of $\times 8$. Subsets of the cohort will additionally be scanned at higher field strength (and resolution) and using MEG [26]. The rfMRI acquisitions (including the use of leading-edge, customized MRI hardware and acquisition software) and image processing are covered in detail in [43,44,68]. The HCP is also acquiring data from diffusion MRI, task fMRI with a broad range of behavioural paradigms, extensive behavioural phenotyping outside the scanner [29], and genotyping, to allow future researchers to relate the HCP functional and structural connectomes to behaviour and genetics.

In addition to achieving unusually high spatial and temporal resolution and acquisition duration, significant effort has been put into minimizing MRI spatial distortions and signal loss and achieving accurate alignment of the functional data to the high-resolution (0.7 mm) structural data acquired for each subject. This allows for the transformation of the cortical rfMRI signal from the originally acquired 3D voxel matrix onto a grey matter surface mesh. Surface-based analysis is advantageous, because it: (i) restricts data storage and analysis to just the grey matter domains of interest (bypassing the storage of white matter and non-brain data); (ii) represents grey matter in a way that respects its natural geometry – 2D mesh-surface vertices for cerebral cortex, plus 3D image-matrix voxels for subcortical grey matter (when all combined together, referred to as ‘grayordinates’); and (iii) therefore allows for better functional alignment across subjects, because the variability in cortical folding patterns in different subjects is greatly ameliorated by modelling on the cortical surface.

The 3T rfMRI data is currently represented using about 90,000 grayordinates with ~ 2 mm spacing.

As discussed above, if intersubject alignment of rfMRI data is driven only by information from high-resolution structural images, it is expected that alignment of cortical functional areas across subjects will be imperfect in many regions, given the variable relationship between cortical areas and structural features such as tissue intensities and cortical folding patterns [69]. The smaller the individual parcels (node maps), the more crucial it becomes to address this problem. As a result, the HCP has developed a method for multimodal intersubject alignment, capable of using many sources of spatial localisation information, including geometric structural features, myelin-weighted images, task fMRI maps, and connectivity information (from rfMRI and diffusion MRI) [70]. By appropriate alignment of multiple sources of information across subjects, we can hope to achieve much more accurate alignment of functional regions (see below).

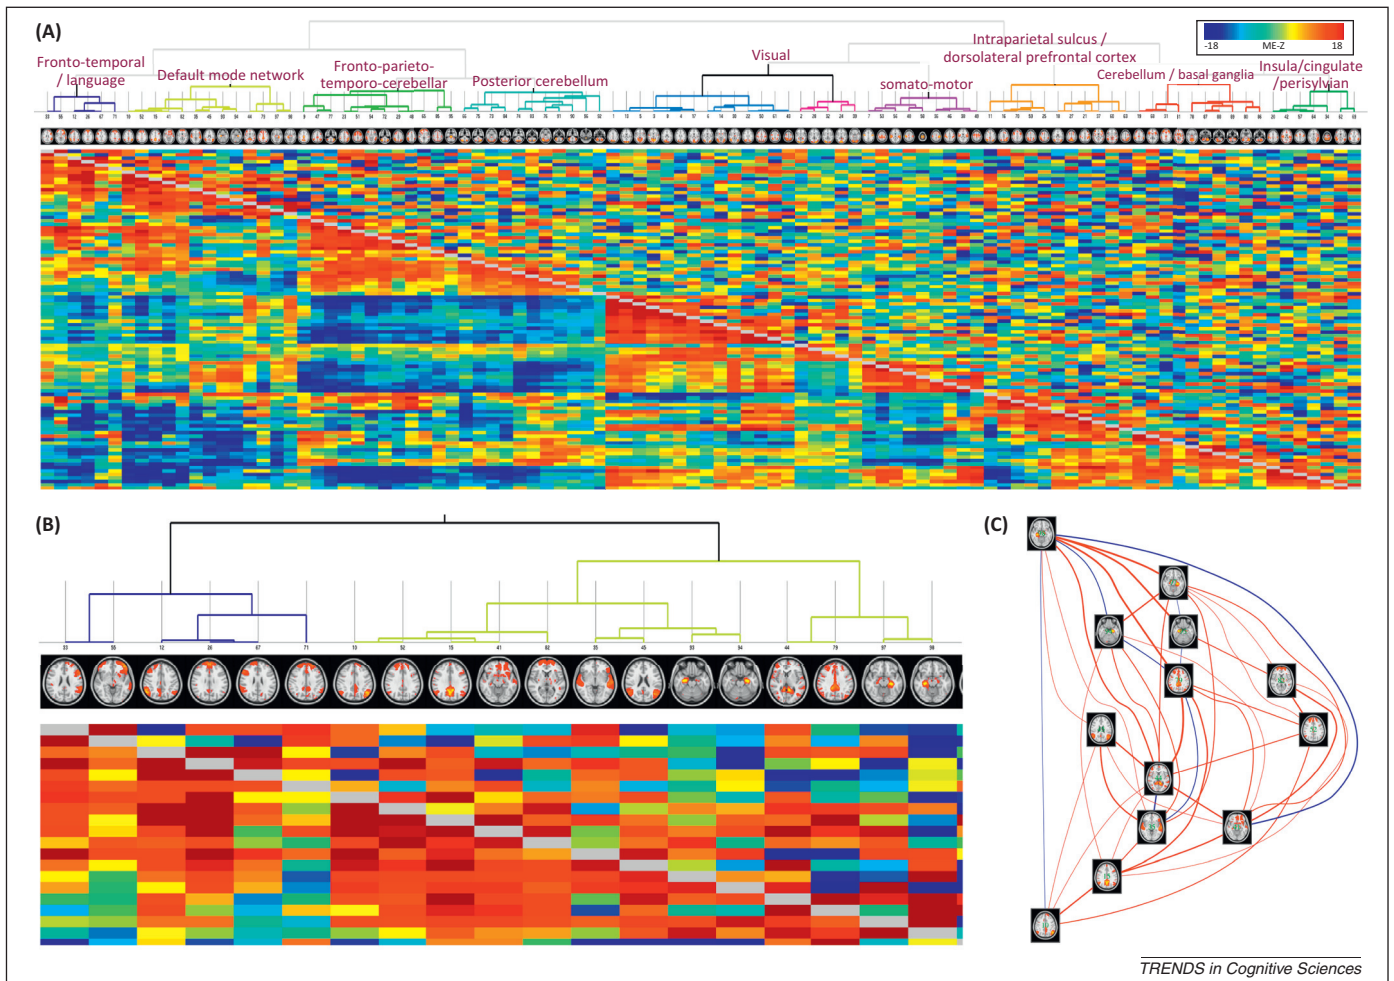


Figure 4. Functional connectomes estimated for 98 nodes derived from 131 Human Connectome Project (HCP) subjects’ resting-state functional MRI (rfMRI) data using group independent-component analysis (ICA) (via the FMRIB Software Library [FSL] MELODIC tool). (A) Full and partial correlation matrices. Full correlation is shown below the diagonal and partial correlation above. Intersubject alignment was conducted using the Multimodal Surface Matching (MSM) method (Figure 5). Small images at the top of each column summarise each node’s spatial map. The nodes were reordered according to a hierarchical clustering of the full correlation matrix (using Ward’s method as implemented in MATLAB), visualized at the top. (B) An expanded view of the top-left part of (A). The individual nodes’ spatial maps can now be more clearly seen; these two clusters involve default mode [13] and language areas. (C) The second of these clusters – covering the default mode network – is shown in an alternative representation of nodes and edges derived from the thresholded ($|Z| > 10$) group-level partial correlation matrix. The partial correlation values are displayed in terms of their strength (connection line thickness) and sign (colour; positive connections shown in red). Network analysis and visualization are conducted using the FSLNets package (<http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets>), with (C) created using Graphviz (<http://www.graphviz.org>).

Following image preprocessing (primarily using the FMRIB Software Library [FSL] [71], FreeSurfer [72], and Connectome Workbench [73] software packages), the HCP applies data-driven artefact removal to rfMRI time series. First, ICA is applied to each 15-min rfMRI dataset separately, to identify both structured artefact and non-artefact components. These components are then automatically classified into artefact versus non-artefact, using machine learning – a hand-trained hierarchical classifier (FMRIB’s ICA-based X-noisifier [FIX]) [44]. FIX achieves classification accuracy on HCP data of over 99%. Removal of the artefact components from the data is then conducted, improving data quality. Further additional clean-up approaches may be of value and are under continuing investigation.

The HCP is already disseminating both raw and pre-processed rfMRI time series data, with group-level parcellations to follow, from which the associated subject-wise parcellated connectome matrices can be estimated. At this point, no single parcellation technique has emerged as the best approach and investigations are ongoing. In the following sets of results, parcellations were generated using high-dimensional ICA. We now present several analyses designed to illustrate the potential of the HCP rfMRI data and some relevant methodological issues.

Figure 4 shows full and partial correlation matrices derived from the first 131 subjects’ rfMRI data released by the HCP – a dataset aggregating over 600,000 time points. Group ICA was run at a dimensionality of 100 and two components were discarded as artefactual, leaving 98 network nodes. (Although FIX clean up is applied to each separate time series dataset, some artefactual components

can still emerge at the group level. For example, low-level artefactual processes that are too weak to be identified by single-session ICA may be consistent across subjects and hence appear more strongly at the group level.) Each row and column represents one of the 98 nodes; small images at the top of each column summarize each node’s spatial map. Because correlation matrices are symmetrical, both full and partial correlation can be shown on the same matrix, with full correlation below the diagonal and partial above. We estimate the correlation matrices separately for each subject (converting from correlation coefficients to z -statistics), conduct a one-group t -test across all 131 subjects (separately for each element in the matrix), and display the results as z -statistics. As a result, the values in the matrices shown reflect both the group-average correlation strength and the consistency across subjects. The colour scale is the same for both full and partial correlations. Notably, the overall range of values is similar for the two matrices. For the partial correlation matrices to be as strong and consistent across subjects as the full correlation is unusual and is indicative of the high quality of the data and large number of time points in each subject’s rfMRI dataset. The nodes have been reordered according to a hierarchical clustering driven by the full correlation matrix, which brings groups of correlated nodes together into clusters (seen as blocks on the diagonal of the full correlation matrix); the hierarchical clustering is visualized at the top, where distinct clusters and subclusters can be seen (and are annotated by the dominant brain regions).

Different methods for aligning data across subjects can be partially evaluated by estimating how similar the resulting network matrices are to each other. The

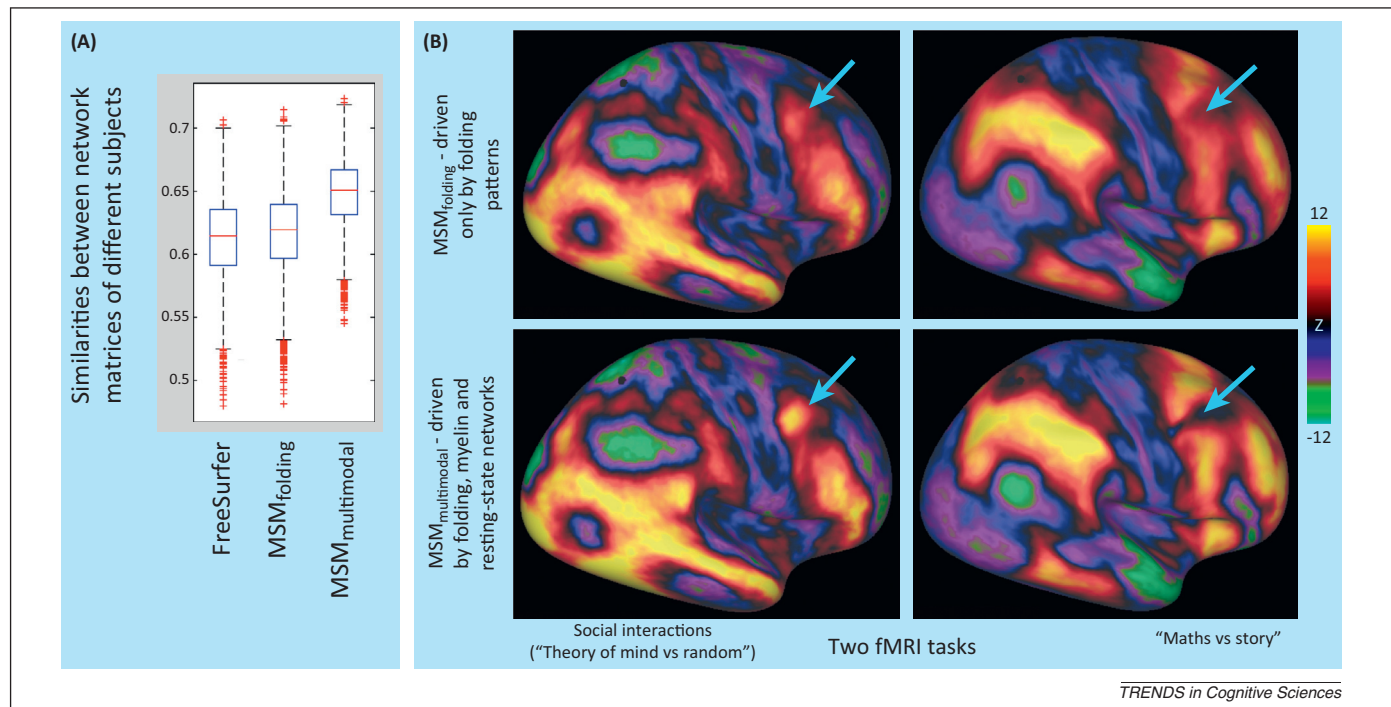


Figure 5. Example analyses showing the effects of improved surface alignment across subjects. (A) Similarities between pairs of subjects’ network matrices. All 8417 pairings of nonrelated subjects in the HCP data were used, after application of three different cross-subject surface-alignment methods. All three group means are statistically significantly different from each other. (B) Group-level analyses of two tasks from 120 Human Connectome Project (HCP) subjects’ task functional MRI (fMRI) datasets. Above is shown the group maps when subjects are aligned with each other using just the folding information. Below is shown the group maps when Multimodal Surface Matching (MSM) utilizes multimodal information, including resting-state networks. The arrows mark example regions showing improved spatial localization of task activation.

cross-subject alignments are conducted before creating a group-level parcellation, propagating that parcellation to individual subjects, and estimating subject-wise network matrices. Therefore, better alignments should yield more accurate network matrices, because the functional parcels defined at the group level should be more accurately aligned to the true subject-specific versions of those same parcels. Once each subject's cortical surface has been estimated from the structural data using FreeSurfer, FreeSurfer also provides one way to improve intersubject alignment of features on the surface. It uses folding (shape) features to move mesh points around the surface, aiming to better align multiple subjects' functional data on the basis of this structural information. We have recently developed an alternative surface-warping approach, Multimodal Surface Matching (MSM) [70]. When MSM is given only folding information to drive the alignment (MSM_{folding}), the overall degree of alignment is similar to that achieved by FreeSurfer. However, if MSM is also fed cortical myelin-weighted maps [74], as well as low-dimensional resting-state network maps, it can utilise this richer structural and

connectional information and further improve the inter-subject alignment ($MSM_{\text{multimodal}}$) [70].

By correlating pairs of subjects' network matrices with each other, we quantified between-subject network similarity. We found the following average correlation values for the three alignment methods: FreeSurfer = 0.613, MSM_{folding} = 0.618, and $MSM_{\text{multimodal}}$ = 0.649 (Figure 5A). However, although this is encouraging, an increase in cross-subject network similarity is not conclusive proof of a better alignment method; for example, it might indicate over-strong application of prior information (e.g., in the extreme, setting all subjects to a constant network), which would reduce sensitivity when attempting to relate the network matrices to other measures. Therefore, we evaluated the quality of alignment for task fMRI data as an independent modality (each subject's alignment warps were estimated from their resting-state fMRI data and applied to task fMRI data). Significant improvements in the alignment of task activation regions were found, both quantitatively and qualitatively (E. Robinson, unpublished); for example, Figure 5B shows stronger and more

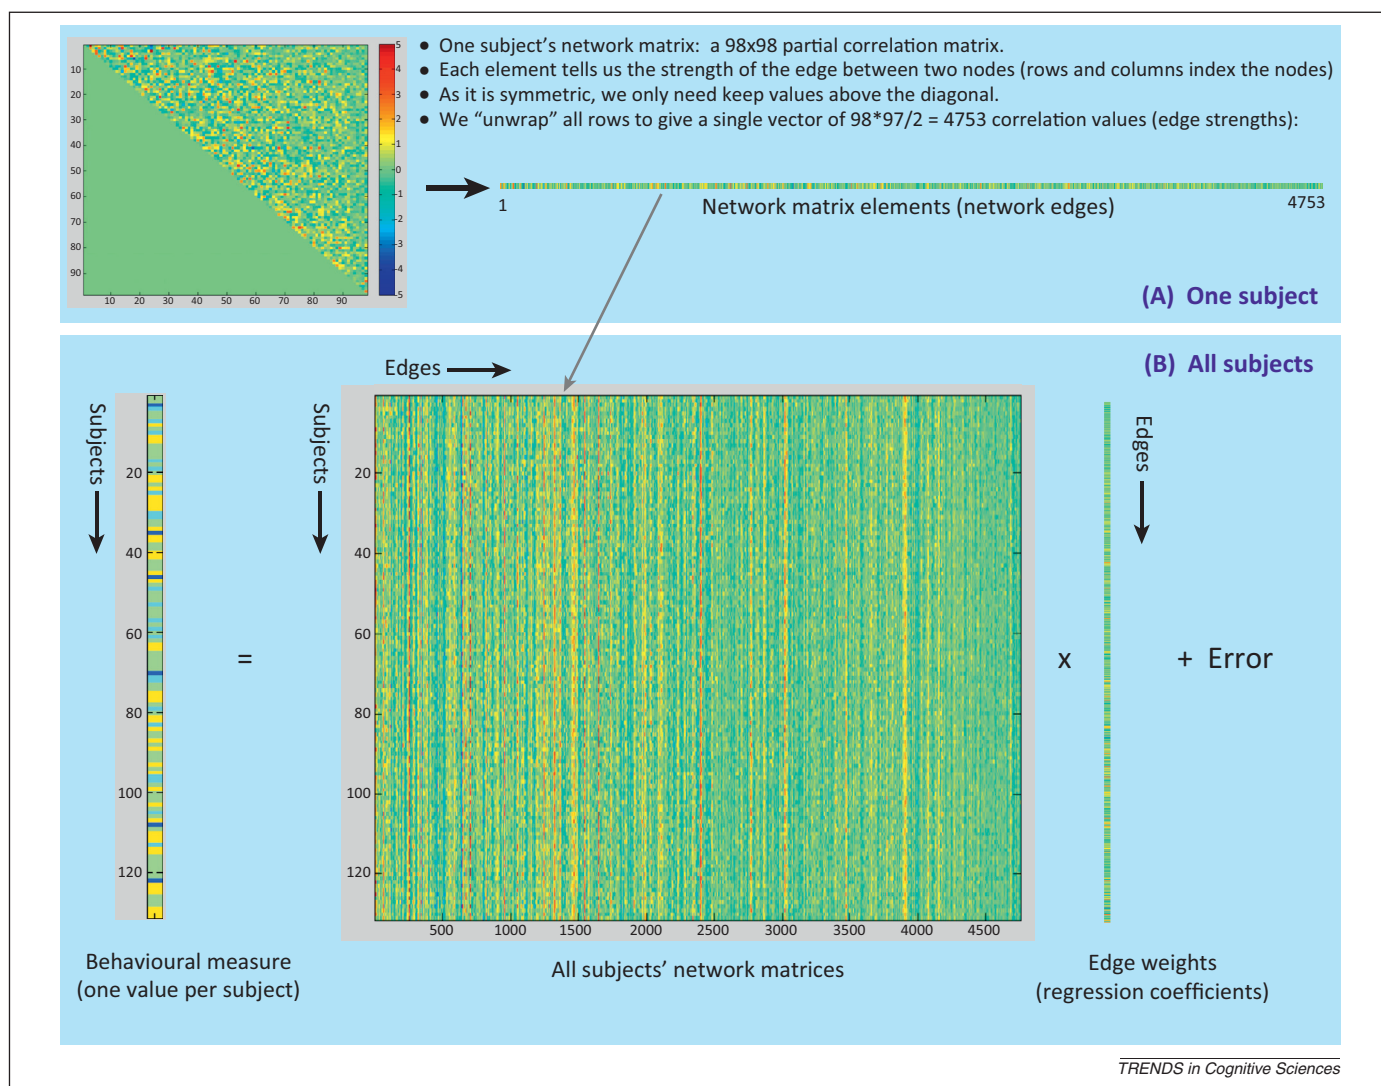


Figure 6. Illustration of how multiple subjects' network matrices can be combined and modelled against other subject variables. (A) The elements of the partial correlation matrix from each subject are reordered into a single row to prepare for combining all network matrix estimates across subjects. (B) The resultant $N_{\text{subjects}} \times N_{\text{edges}}$ matrix represents all parcellated connectomes from all subjects, which can be used to relate the parcellated connectome to subjects' behaviour, genotype, and other personal measures.

sharply defined activation patterns from a group-level analysis of two different tasks ('maths versus story' and a social-interaction task) when $MSM_{multimodal}$ is utilised (lower panels) compared with $MSM_{folding}$ (upper panels), especially in parts of the prefrontal cortex (blue arrows).

Figure 6 illustrates how the elements of the partial correlation matrix from each subject can be reordered into a single row and then combined with network matrix estimates from all subjects to form a matrix that represents all connectomes from all subjects. We can use this as a set of regressors in a large multiple regression that seeks to model (across subjects) a given behavioural variable

(e.g., intelligence quotient [IQ]). This approach represents a reversal of the regression analysis that is typically performed in task fMRI, in which one independently analyses each voxel to determine how much a given temporal regressor explains that voxel's time course, resulting in a spatial map of explanatory power. In that type of analysis, one aims to explain the data (fMRI time series) in terms of a set of model regressors that are derived from the study design (task time course). Our goal here is different; we wish to find a set of edge weights that is able to explain the behaviour of interest (e.g., IQ). Hence, the network edges are the model regressors in this analysis. If the number of

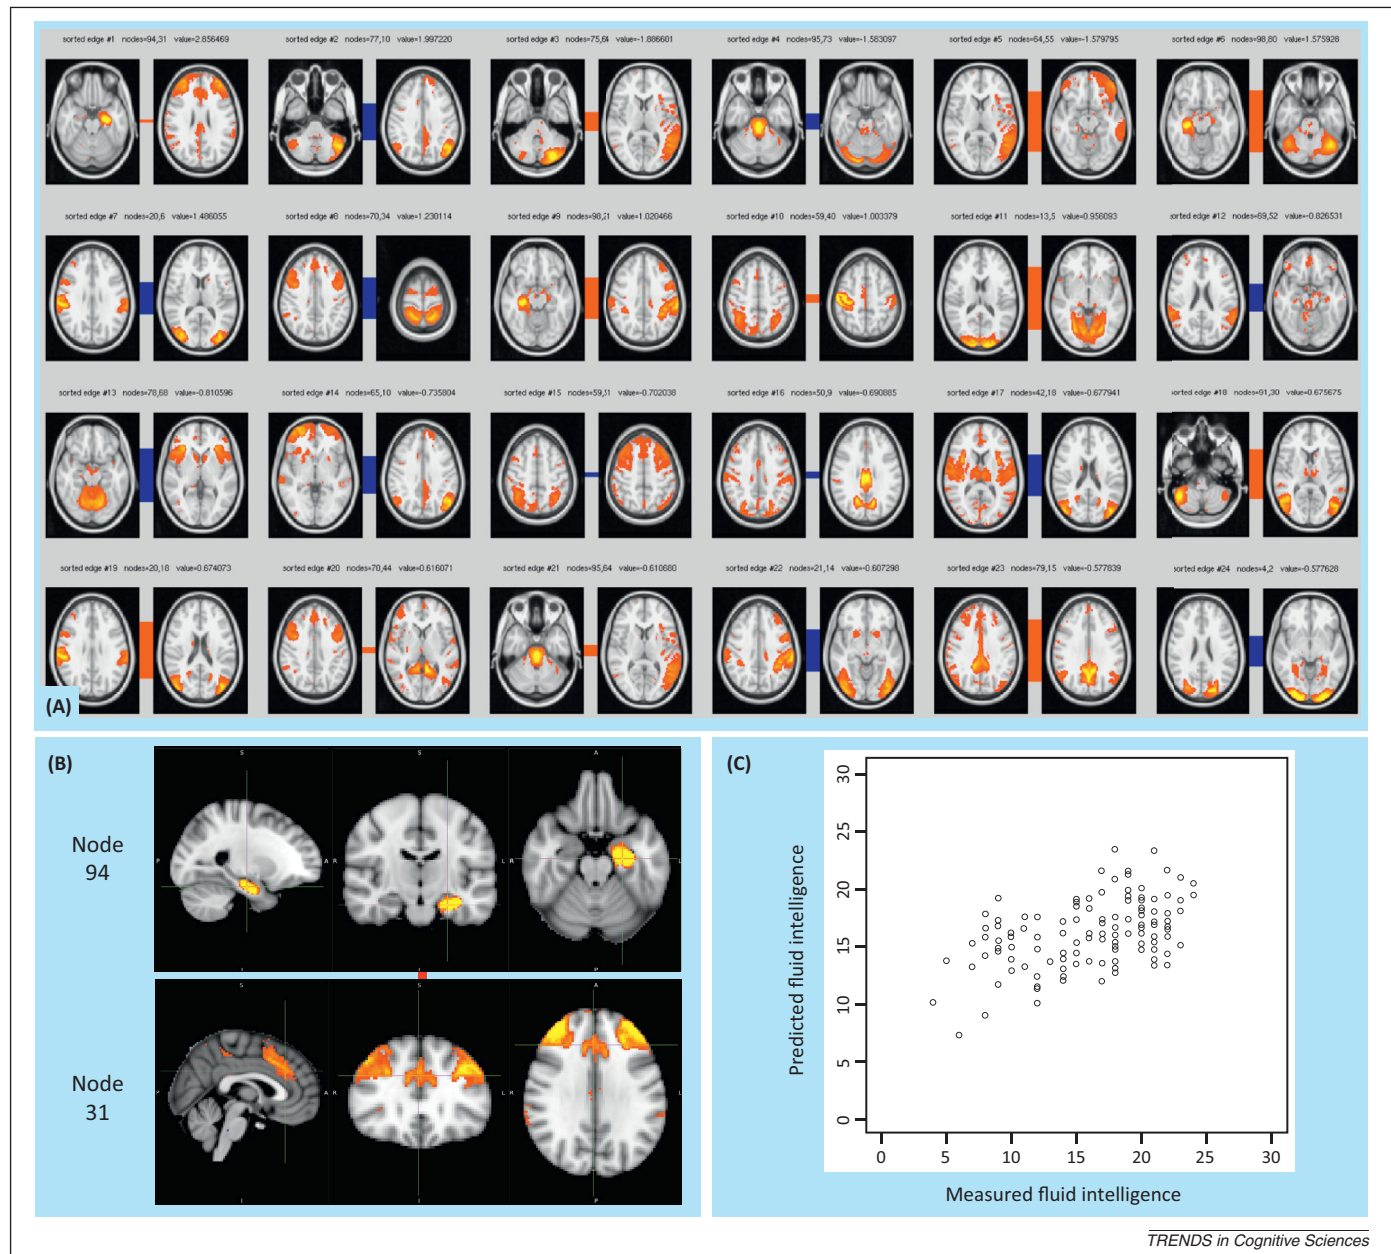


Figure 7. A significant association was found between network matrices and fluid intelligence (FI) ($p < 0.05$, corrected for multiple comparisons across all 91 behavioural variables tested). The 4753 edges were reduced using a Bayesian feature-selection method [83] that kept only the 98 edges most strongly predictive of FI and applied ridge regression (L2) shrinkage on the kept features (Box 2); over-fitting was avoided by conducting the feature selection and shrinkage inside a leave-one-out loop. Statistical significance was estimated using subject-wise permutation testing to derive p -values on the model fitting, taking into account the family structure in the data, such that cross-subject correlations were correctly handled. (A) The 24 edges (node pairs) that (on average) had the strongest weights in the regression are shown. The coloured bars connecting the two nodes in each pair reflect the overall group-average connection strength. Each edge's weight in the multiple regression is noted as the 'value'. (B) The node pair that contributes most strongly to the regression against FI is shown in more detail; left hippocampus (generally associated with memory/recall) and medial/lateral frontal regions (generally associated with cognitive control). (C) Predicted versus measured FI, with one data point per subject; each subject's FI was predicted using his or her network matrix, where the linear regression model was trained excluding the data from that subject (and all of his or her family members).

subjects is small relative to the number of edges, the dimensionality of this problem is such that the network matrix cannot be used directly, but rather must first be reduced (e.g., see Figure 7 legend). If the overall model fit is found to be significant, it provides evidence for a link between the functional connectome and a specific behavioural characteristic. More specifically, the regression coefficients (one weight per network edge) indicate the relative contributions from the edges to the overall behavioural modelling. This kind of analysis has previously been used (albeit from full correlation network matrices) to predict subjects' psychiatric condition [75] and cognitive state [76]; for additional discussions on cross-subject network modelling, see also [77,78].

In Figure 7, we utilise the above approach to show a significant association between individual subjects' network matrices and fluid intelligence (a component of IQ). The

24 edges (node pairs) having the strongest weights in the regression are shown; the coloured bars connecting the two nodes in each pair reflect group-average connection strength. There is not necessarily a relationship between the strength (or even sign) of the population average connection between two nodes and the extent to which the cross-subject variability in the connection strength correlates with a given behavioural measure. Indeed, the node pair that contributes most strongly to the regression against fluid intelligence (shown in more detail in Figure 7B) has only a weak connection on average in the population, but it varies across subjects in a way that correlates with fluid intelligence.

In Figure 8, we show results relating subjects' network matrices to their sex. In the first analysis we applied a two-group *t*-test (between the two sexes) separately for every edge in the partial correlation network matrix. It is

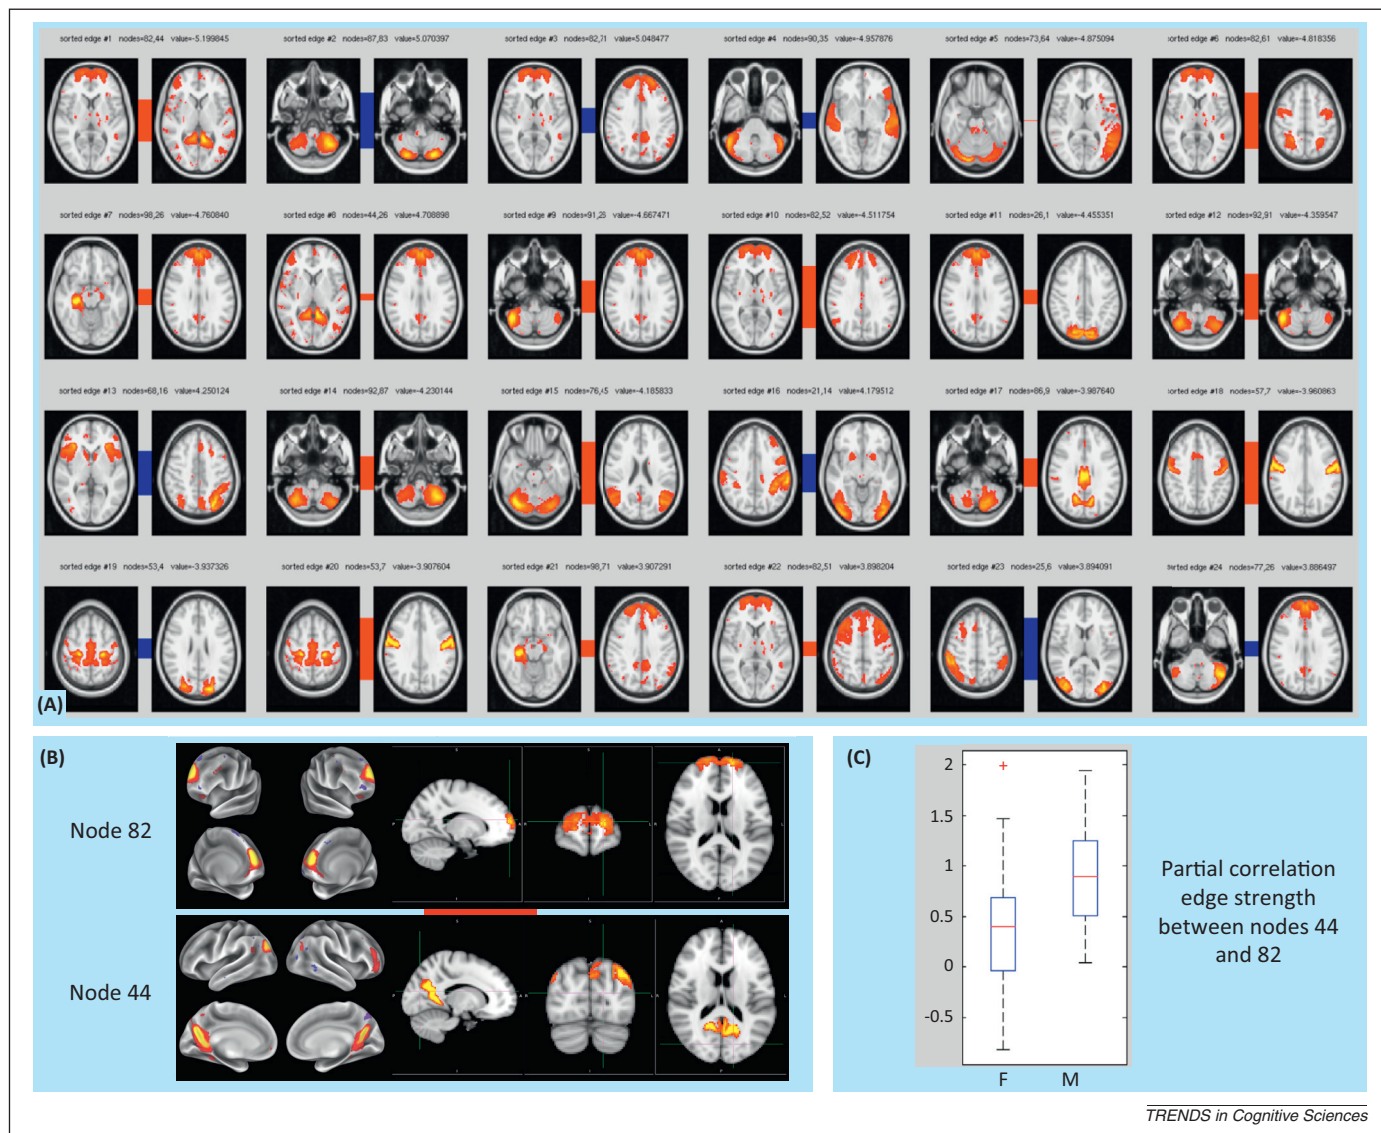


Figure 8. Relationship between edge strength and sex. We discarded the second subject of any pair of twins (leaving 104 subjects) to avoid the danger of over-fitting (in the case of sex prediction) or confounding the statistical inference (when conducting the univariate *t*-tests), given that twins are likely to have similar network matrices. (A) The 24 edges that are most different in connection strength between males and females are shown; the first nine are significantly different ($p < 0.05$, two-tailed, corrected for multiple comparisons using the FSL randomise tool). The coloured bars connecting the two nodes in each pair reflect the overall connection strength, averaged across all subjects of both sexes. (B) The two nodes whose connection (edge) is most different between the sexes are shown in greater detail: posterior cingulate/precuneus and frontal pole. (C) For this edge, the connection strengths in the two groups are shown as separate cross-subject distributions.

important to correct the resulting p -values for the multiple comparisons across all of the network edges; however, Bonferroni correction would probably be conservative, because we do not expect the subject-wise variations of all edges to be fully independent of each other. A preferable way to correct for multiple comparisons is to utilise permutation testing and build up a null distribution (over thousands of random permutations of the subject ordering) of the maximum t -value across all edges. Comparing all edges' group-difference t -values against the 97.5th percentile of this null distribution provides accurate control of the family-wise error rate at 5% (for a two-sided test) and enables identification of edges that are significantly different between the groups. This identified nine edges that were significantly different in strength between the two sexes.

We then conducted multivariate analyses, attempting to predict subjects' sex, utilising all network edges simultaneously. The $N_{\text{subjects}} \times N_{\text{edges}}$ connectomes matrix and the sex of all subjects were fed into a linear classifier (simple linear discriminant analysis, assuming equal group covariances). Using leave-one-out training and testing, sex was successfully predicted in 87% of subjects (see [Box 2](#) for further discussion of multivariate modelling and statistics). The three alignment methods described above gave classification accuracies of: FreeSurfer = 78%; $\text{MSM}_{\text{folding}}$ = 82%; and $\text{MSM}_{\text{multimodal}}$ = 87%. This suggests that multimodal MSM improves the alignment of function across subjects (because it improves the accuracy with which the derived network matrices can predict sex) and also supports the hypothesis that this sex prediction is indeed primarily being driven by network connection strength and not (for example) by systematic functional misalignments across (groups of) subjects.

The fact that some connections are stronger in females than in males, and some are weaker, suggests that the discrimination is not driven primarily by uninteresting gross effects such as group biases in head size or within-scan head motion. Indeed, the histogram of edge-strength mean group differences is centred at zero, with no gross asymmetry. When we used the full correlation network matrices instead of partial correlation, the sex classification accuracy was reduced from 87% to 70%. This supports the hypothesis that partial correlation matrices provide a better estimate of the true macroscopic functional connectome than full correlation.

Another area of high current interest is the investigation of connectivity changes over time ('nonstationarities'). For example, nonstationarities in correlation have been studied utilising either wavelet decompositions or sliding-window correlation [79]. In the latter approach, a windowed subset of time points (e.g., taking all images from the first minute of the data) is used to estimate the full correlation between two or more nodes, and the time window is shifted to different temporal positions to yield a 'time series' of correlation values (or matrices). One important issue is to distinguish between changes in correlation due to some nodes being part of multiple overlapping functional networks versus the internal connections within any given network being non-constant. Both possibilities are neurobiologically interesting and it is

Box 2. Advances in multivariate modelling and statistics

Recent advances in multivariate modelling and statistics are likely to be useful in relating connectomes to behavioural and genetic data.

Feature selection may aid the sensitivity of multivariate regressions or classifications. For example, if network edges that do not contribute to a prediction of a specific behaviour are rejected from consideration, multivariate modelling may be better conditioned and hence more sensitive [84]. To validate hypotheses such as the predictability of a certain behavioural measure from the networks, care must be taken to avoid bias when feature selection is involved. For example, in a standard permutation analysis we shuffle the data (subject ordering) with respect to the behavioural measures (respecting the structural characteristics of the data, such as the presence of twins) and compute some statistic for each permutation. A natural choice for this test statistic is the leave-one-out prediction error, the null distribution of which will be compared with the corresponding leave-one-out error from the non-permuted data. The feature-selection process must be rerun for each permutation.

Sparse estimation techniques, where feature selection is a consequence of regularisation in the estimation process, are a popular alternative to 'hard' feature selection. Well-known examples are L1 regularisation and the elastic net (which combines an L1-norm penalty with an L2 penalty over the regression coefficients). Whereas pure L1 regularisation discards the redundant edges, the elastic net selects all of the relevant features in a balanced fashion and is particularly convenient for interpretation purposes [85]. If cross-validation is used to optimise the regularisation parameters, this must be done separately for each permutation. Bayesian inference, which naturally estimates the distribution of the regularisation parameters within the learning process, is a valuable alternative to classical selection of the regularisation parameters [86,87].

Finally, methods that co-model two (or more) multivariate datasets are of great potential value. For example, instead of relating each of the many behavioural variables to the $N_{\text{subjects}} \times N_{\text{edges}}$ population network matrix one at a time, we can form the $N_{\text{subjects}} \times N_{\text{behavioural-variables}}$ population behavioural matrix and directly attempt to co-model this with the connectomes matrix. Approaches such as canonical correlation analysis (CCA) [88] and partial least squares (PLS) [89] use a set of latent variables to model the relation between the two matrices. Such approaches can provide great gains in sensitivity by increasing the signal-to-noise ratio and reducing problems of multiple comparisons.

not trivial to determine which of these factors dominates in practice for the apparent nonstationarities seen in typical rfMRI data. In addition, a problem can arise when applying a window length that does not encompass (at least) several cycles of the resting fluctuations, because the rfMRI signal is dominated by low frequencies (~ 0.015 Hz), and a sliding window containing a fraction of a cycle is expected to show correlations appearing to fluctuate wildly over time, even if the underlying network structure is stationary. Recent evidence [47] indicates that resting-state network signals occur up to at least 0.2 Hz in good-quality data; hence, aggressive high-pass filtering might ameliorate this problem (although the data signal-to-noise ratio may impose a practical limitation on this, implying a minimum practical window length).

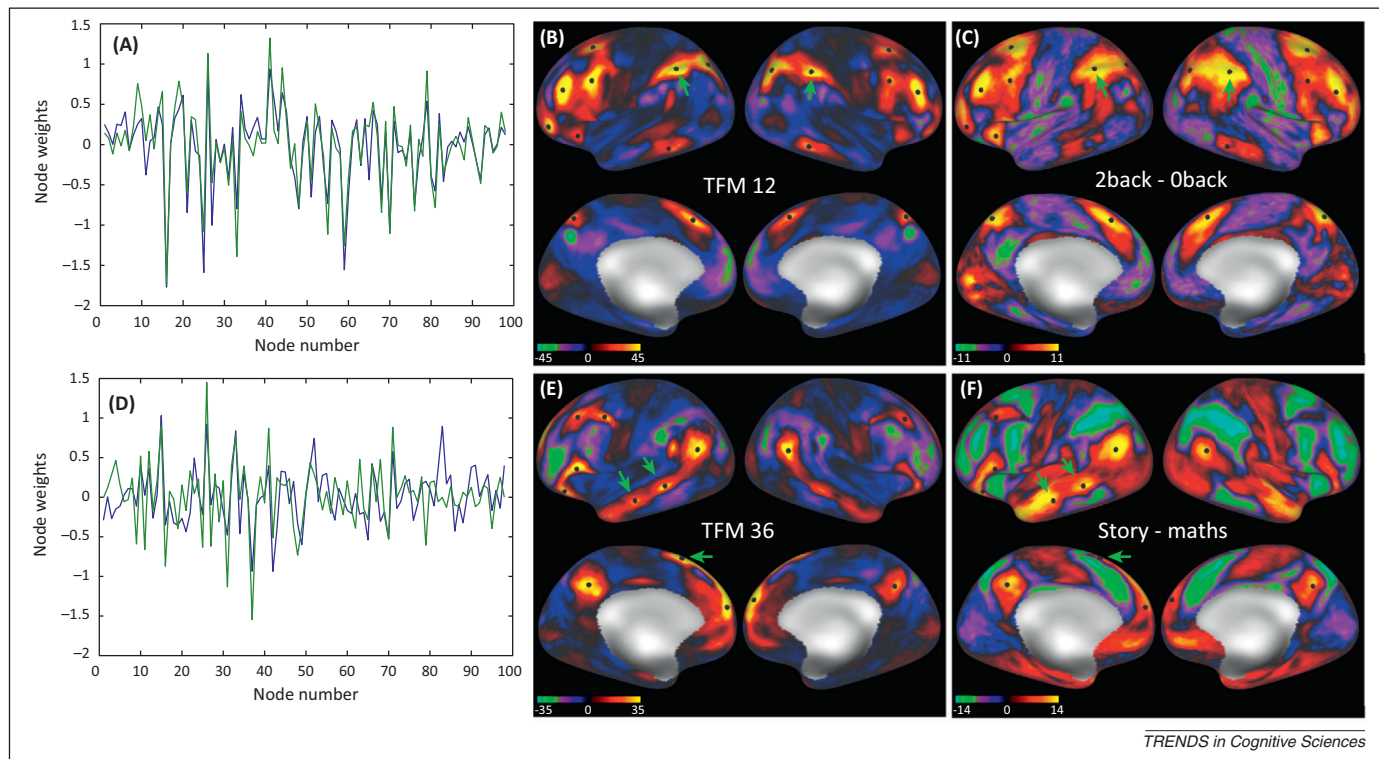
Once a 'time series of correlation matrices' has been generated through sliding-window correlation, clustering analysis can be used to identify distinct correlation patterns that repeatedly reappear over an extended period [79]. This has the potential to be a useful exploratory tool for identifying multiple modes of correlation. A limitation, however, is that multiple distinct functional networks might overlap each other in space and also vary over time

in how they interact with each other. If that is the case, windowed correlation patterns may not be very informative, because distinct overlapping and coactivating networks' varying amplitudes will be a large factor in the apparent correlation pattern, confounding estimation of their true distinct internal connectivities. An alternative model for multiple, spatially overlapping networks is to feed the node time series themselves (as opposed to correlations derived from them) into a clustering approach that separates distinct networks from each other on the basis of being temporally distinct, while allowing for spatial overlap (allowing nodes to be members of multiple functional networks). For example, feeding node time series into temporal ICA (as opposed to spatial ICA that may be used to generate the parcellation), the resulting ICA components can be considered as distinct 'temporal functional modes', where there is very little restriction on the spatial overlap between distinct modes; [42] showed several such modes, with plausible functional interpretability and large amounts of spatial overlap – for example, the patterns of average correlation and anticorrelation associated with the default mode network can be decomposed into several modes of distinct but overlapping spatial layout (and hence presumably distinct function). The hypothesis is that

without such decomposition, the patterns seen through time-averaged correlations tell a cruder story that only reflects multiple distinct processes all averaged together.

A major challenge for temporal ICA, however, is the need for a large number of time points to perform robustly. This technique thus benefits greatly from long rfMRI sessions and accelerated acquisition, such as that obtained by the HCP. Figure 9 shows two example results from feeding 98 nodes' time series into temporal ICA, where we analysed the high-quality HCP rfMRI data containing more time points (>600,000, concatenating data from 131 subjects) than spatial points (~90,000 grayordinates). In both cases we show strong spatial correspondence to specific task-activation maps from the HCP task fMRI datasets, thus helping to interpret the functional modes identified from the resting-state data.

One limitation of temporal ICA is that all modes are forced to be fully independent of each other over time; this is unfortunate insofar as we would prefer to identify functionally distinct modes that may have some temporal dependence. For example, two distinct functional modes might in reality be mutually exclusive, if they share one node that strongly participates in only one function at any given time; this would imply a negative correlation



TRENDS in Cognitive Sciences

Figure 9. Two Temporal Functional Modes (TFMs) [42] estimated from 131 subjects' data from the Human Connectome Project (HCP), with cross-subject alignment using multimodal Multimodal Surface Matching ($MSM_{multimodal}$). Ninety-eight nodes' time series were fed into temporal independent-component analysis (ICA) and 98 ICA components were estimated. Split-half reproducibility testing [90] (running temporal ICA separately on two halves of the complete dataset, with subjects randomly assigned to one group or the other) indicated that 29 of the estimated TFMs were significantly spatially reproducible. Each TFM spatial map was generated by multiplying each node's voxel-wise spatial map (obtained by the original 100-dimensional group-level spatial ICA) by the weight identified by the temporal ICA for that node for that TFM and summing across all such weighted node maps. (A) Split-half reproducibility for TFM 12. The two plots (differently coloured for the different subgroups of subjects) show the set of node weights for this TFM; this is the temporal ICA 'mixing matrix' and determines the TFM's spatial map. (B) Spatial map of TFM 12 on inflated cortical surfaces. (C) An activation contrast map from the HCP task functional MRI (fMRI) datasets, with a specific task chosen that matches spatially the TFM found from the resting data. This TFM closely matches the 2-back versus 0-back contrast in the HCP working memory task. Black dots in (B) and (C) are in the same anatomical location in the two cases and placed to be centred on TFM 12 patches. (D) Split-half reproducibility for TFM 36. (E) Spatial map of TFM 36. (F) Task fMRI activation contrast map for the 'story versus maths' contrast, comparing listening to stories with answering spoken arithmetic questions; in this case the TFM analysis has identified a functional network showing great similarity to language function. Black dots in (E) and (F) are centred on TFM 36 regions. Green arrows indicate TFM patches that are spatially more focal than the corresponding task fMRI activations. It is clear that, although sharing several nodes, these two TFMs have differing overall spatial make ups and are reproducible across the two sets of distinct subjects. See also a movie that shows TFM brain dynamics from 10 subjects, one after another (see supplementary movie). For each subject, 1 min of dynamics is shown, at $\times 4$ real time. Different colours are different TFMs; for a given TFM, darker clusters are anticorrelated with brighter clusters.

between the modes' time series, a situation that would not be well identified through the application of temporal ICA. A promising area of future research is thus to develop more sophisticated, dynamic network modelling approaches that are able to identify such scenarios and improve our understanding of dynamic and overlapping functional modes in the brain. Another limitation with temporal ICA is the assumption that the connections within a given mode are not changing in strength over time; this is surely an over-simplistic view of brain networks. A related challenge is raised by evidence that rfMRI signal includes complex spatiotemporal patterns of spontaneous activity that propagate across the brain [80] and are therefore not at all well modelled by spatially fixed network modelling approaches. There clearly remains much to be done in the spatial and temporal modelling of functional brain networks.

Concluding remarks

The macroscopic functional connectome as elucidated through rfMRI provides just one view on the 'complete' brain connectome. It does not tell us directly about the microscopic, neuronal-level structural connectome or even about the macroscopic-level structural connectome such as is inferred from diffusion MRI. It does not provide the more direct view of neuronal activity and information flow that can be obtained from electrophysiological techniques or inform us regarding the biochemistry of brain connectivity. However, rfMRI does provide us with something valuable and unique – the ability to create a richly detailed mapping of the functional connectome, at the millimetre scale,

covering the whole brain, with *in vivo*, non-invasive imaging across thousands of subjects. In coming years we can expect to learn much regarding the topologies and functions of the brain's networks and how they relate to behaviour, genotype, and pathologies, as acquisition and modelling methods for rfMRI connectomics continue to make exciting advances (Box 3).

Acknowledgements

The authors are grateful for funding via the following grants: 1U54MH091657-01 (NIH Blueprint for Neuroscience Research), P30-NS057091, P41-RR08079/EB015894, F30-MH097312 (NIH), and 098369/Z/12/Z (Wellcome Trust). They thank their many colleagues within the WU-Minn HCP Consortium for their invaluable contributions in generating the publicly available HCP data and in implementing the many procedures needed to acquire, analyse, visualise, and share these datasets.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.tics.2013.09.016>.

References

- 1 Biswal, B. *et al.* (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34, 537–541
- 2 Beckmann, C.F. *et al.* (2005) Investigations into resting-state connectivity using independent component analysis. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 360, 1001–1013
- 3 Picchioni, D. *et al.* (2013) Sleep and the functional connectome. *Neuroimage* 80, 387–396
- 4 Yeo, B.T. *et al.* (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106, 1125–1165
- 5 Li, L. *et al.* (2013) Mapping putative hubs in human, chimpanzee and rhesus macaque connectomes via diffusion tractography. *Neuroimage* 80, 462–474
- 6 Brookes, M.J. *et al.* (2011) Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proc. Natl. Acad. Sci. U.S.A.* 108, 16783–16788
- 7 de Pasquale, F. *et al.* (2010) Temporal dynamics of spontaneous MEG activity in brain networks. *Proc. Natl. Acad. Sci. U.S.A.* 107, 6040–6045
- 8 Castellanos, F.X. *et al.* (2013) Clinical applications of the functional connectome. *Neuroimage* 80, 527–540
- 9 Murphy, K. *et al.* (2013) Resting-state fMRI confounds and cleanup. *Neuroimage* 80, 349–359
- 10 Sadaghiani, S. and Kleinschmidt, A. (2013) Functional interactions between intrinsic brain activity and behavior. *Neuroimage* 80, 379–386
- 11 Laird, A.R. *et al.* (2013) Networks of task co-activations. *Neuroimage* 80, 505–514
- 12 Smith, S.M. *et al.* (2009) Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A.* 106, 13040–13045
- 13 Raichle, M.E. (2010) The brain's dark energy. *Sci. Am.* 302, 44–49
- 14 Passingham, R.E. (2013) What we can and cannot tell about the wiring of the human brain. *Neuroimage* 80, 14–17
- 15 Sotiropoulos, S.N. *et al.* (2013) Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *Neuroimage* 80, 125–143
- 16 Setsompop, K. *et al.* (2013) Pushing the limits of *in vivo* diffusion MRI for the Human Connectome Project. *Neuroimage* 80, 220–233
- 17 McNab, J.A. *et al.* (2013) The Human Connectome Project and beyond: initial applications of 300 mT/m gradients. *Neuroimage* 80, 234–245
- 18 O'Donnell, L.J. *et al.* (2013) Fiber clustering versus the parcellation-based connectome. *Neuroimage* 80, 283–289
- 19 Mangin, J.F. *et al.* (2013) Toward global tractography. *Neuroimage* 80, 290–296
- 20 Kennedy, H. *et al.* (2013) Why data coherence and quality is critical for understanding interareal cortical networks. *Neuroimage* 80, 37–45

Box 3. Outstanding questions

- How can we better understand and model the relationship between true biological brain connections and what is measured by functional and structural neuroimaging? Put another way, how can we bridge the gap between large-scale macroscopic network modelling from connectomics neuroimaging data and low-level biophysical measurements and modelling of neural connectivity?
- In the shorter term, how can we optimally generate functional and structural connectome matrices from connectomics neuroimaging data (such as that being generated by the HCP)?
- Tied in with this, what are the most useful methods for generating functional parcellations of the brain, in individuals as well as in group averages, and how should these be supplemented by richer models that account for the variation of connectivity and function within parcels?
- More generally, how can we combine the information from detailed spatial mapping (e.g., consideration of the full voxel-wise or 'dense' connectome or from low-dimensional spatial ICA) with the richer temporal network modelling that can be conducted after parcellation?
- How should we identify and model subjects' datasets where the connectome is fundamentally different from population averages, for example, where the parcellation or connectivity is topologically incompatible or otherwise fundamentally different from the common group norm?
- What are the best models for identifying functional subnetworks and studying how they dynamically evolve over time and interact with each other?
- Ultimately, what are the structural connections in the brain, how do these connections and their electrochemical modulation of each other give rise to functional dynamics, behaviour, and consciousness, and how are these altered in early development, disease, and ageing?

- 21 da Costa, N.M. and Martin, K.A. (2013) Sparse reconstruction of brain circuits: or, how to survive without a microscopic connectome. *Neuroimage* 80, 27–36
- 22 Stephan, K.E. (2013) The history of CoCoMac. *Neuroimage* 80, 46–52
- 23 White, J.G. *et al.* (1986) The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 314, 1–340
- 24 Scholvinck, M.L. *et al.* (2013) The contribution of electrophysiology to functional connectivity mapping. *Neuroimage* 80, 297–306
- 25 David, O. *et al.* (2013) Probabilistic functional tractography of the human cortex. *Neuroimage* 80, 307–317
- 26 Larson-Prior, L. *et al.* (2013) Adding dynamics to the Human Connectome Project with MEG. *Neuroimage* 80, 190–201
- 27 Caspers, S. *et al.* (2013) Microstructural grey matter parcellation and its relevance for connectome analyses. *Neuroimage* 80, 18–26
- 28 Evans, A.C. (2013) Networks of anatomical covariance. *Neuroimage* 80, 489–504
- 29 Barch, D.M. *et al.* (2013) Function in the Human Connectome: task-fMRI and individual differences in behavior. *Neuroimage* 80, 169–189
- 30 Thompson, P.M. *et al.* (2013) Genetics of the connectome. *Neuroimage* 80, 475–488
- 31 Sporns, O. (2013) The human connectome: origins and challenges. *Neuroimage* 80, 53–61
- 32 Blumensath, T. *et al.* (2013) Spatially constrained hierarchical parcellation of the brain with resting-state FMRI. *Neuroimage* 76, 313–324
- 33 de Reus, M.A. and van den Heuvel, M.P. (2013) The parcellation-based connectome: limitations and extensions. *Neuroimage* 80, 397–404
- 34 Beckmann, C.F. (2012) Modelling with independent components. *Neuroimage* 62, 891–901
- 35 Smith, S.M. (2012) The future of FMRI connectivity. *Neuroimage* 62, 1257–1266
- 36 Margulies, D.S. *et al.* (2013) Visualizing the human connectome. *Neuroimage* 80, 445–461
- 37 Ogawa, S. *et al.* (1992) Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci. U.S.A.* 89, 5951–5955
- 38 Mansfield, P. (1977) Multi-planar image formation using NMR spin echoes. *J. Phys. C: Solid State Phys.* 10, L55–L58
- 39 Moeller, S. *et al.* (2010) Multiband multislice GE–EPI at 7 Tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magn. Reson. Med.* 63, 1144–1153
- 40 Setsompop, K. *et al.* (2012) Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn. Reson. Med.* 67, 1210–1224
- 41 Feinberg, D.A. *et al.* (2010) Multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging. *PLoS ONE* 5, e15710
- 42 Smith, S.M. *et al.* (2012) Temporally-independent functional modes of spontaneous brain activity. *Proc. Natl. Acad. Sci. U.S.A.* 109, 3131–3136
- 43 Glasser, M.F. *et al.* (2013) The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* 80, 105–124
- 44 Smith, S.M. *et al.* (2013) Resting-state fMRI in the Human Connectome Project. *Neuroimage* 80, 144–168
- 45 Power, J.D. *et al.* (2011) Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* 59, 2142–2154
- 46 De Martino, F. *et al.* (2007) Classification of fMRI independent components using IC-fingerprints and support vector machine classifiers. *Neuroimage* 34, 177–194
- 47 Niazy, R.K. *et al.* (2011) Spectral characteristics of resting state networks. *Progress Brain Res.* 193, 259–276
- 48 Friston, K.J. (2011) Functional and effective connectivity: a review. *Brain Connect.* 1, 13–36
- 49 Marrelec, G. *et al.* (2006) Partial correlation for functional brain interactivity investigation in functional MRI. *Neuroimage* 32, 228–237
- 50 Friston, K.J. (1994) functional and effective connectivity in neuroimaging: a synthesis. *Hum. Brain Mapp.* 2, 56–78
- 51 Nakagawa, T.T. *et al.* (2013) Bottom up modeling of the connectome: linking structure and function in the resting brain and their changes in aging. *Neuroimage* 80, 318–329
- 52 Woolrich, M.W. and Stephan, K.E. (2013) Biophysical network models and the human connectome. *Neuroimage* 80, 330–338
- 53 McIntosh, A.R. and Gonzales-Lima, F. (1994) Structural equation modeling and its application to network analysis in functional brain imaging. *Hum. Brain Mapp.* 2, 2–22
- 54 Friston, K.J. *et al.* (2011) Network discovery with DCM. *Neuroimage* 56, 1202–1221
- 55 Friedman, J. *et al.* (2008) Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 9, 432–441
- 56 Smith, S.M. *et al.* (2011) Network modelling methods for fMRI. *Neuroimage* 54, 875–891
- 57 Varoquaux, G. *et al.* (2010) Brain covariance selection: better individual functional connectivity models using population prior. *Adv. Neural Inf. Process. Syst.* 23, 2334–2342
- 58 Granger, C.W.J. (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37, 424–438
- 59 Ramsey, J.D. *et al.* (2010) Six problems for causal inference from fMRI. *Neuroimage* 49, 1545–1558
- 60 Patel, R.S. *et al.* (2006) A Bayesian approach to determining connectivity of the human brain. *Hum. Brain Mapp.* 27, 267–276
- 61 Hyvarinen, A. and Smith, S.M. (2013) Pairwise likelihood ratios for estimation of non-Gaussian structural equation models. *JMLR* 14, 111–152
- 62 Markov, N. *et al.* (2013) A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cereb. Cortex* <http://dx.doi.org/10.1093/cercor/bhs270>
- 63 Pearl, J. (2009) Causal inference in statistics: an overview. *Stat. Surv.* 3, 96–146
- 64 Rubinov, M. and Sporns, O. (2010) Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52, 1059–1069
- 65 Fornito, A. *et al.* (2013) Graph analysis of the human connectome: promise, progress, and pitfalls. *Neuroimage* 80C, 426–444
- 66 Craddock, R.C. *et al.* (2011) A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum. Brain Mapp.* 33, 1914–1928
- 67 Van Essen, D.C. *et al.* (2013) The WU–Minn Human Connectome Project: an overview. *Neuroimage* 80, 62–79
- 68 Ugurbil, K. *et al.* (2013) Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *Neuroimage* 80, 80–104
- 69 Van Essen, D.C. *et al.* (2012) Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cereb. Cortex* 22, 2241–2262
- 70 Robinson, E. *et al.*, eds (2013) *Multimodal Surface Matching: Fast and Generalisable Cortical Registration using Discrete Optimisation*, Springer
- 71 Jenkinson, M. *et al.* (2012) FSL. *Neuroimage* 62, 782–790
- 72 Fischl, B. (2012) FreeSurfer. *Neuroimage* 62, 774–781
- 73 Marcus, D.S. *et al.* (2013) Human Connectome Project informatics: quality control, database services, and data visualization. *Neuroimage* 80, 202–219
- 74 Glasser, M.F. *et al.* (2011) Comparison of surface gradients derived from myelin maps and functional connectivity analysis. In *Proceedings of the 17th International Conference on Functional Mapping of the Human Brain*. Organization for Human Brain Mapping
- 75 Craddock, R.C. *et al.* (2009) Disease state prediction from resting state functional connectivity. *Magn. Reson. Med.* 62, 1619–1628
- 76 Shirer, W.R. *et al.* (2012) Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb. Cortex* 22, 158–165
- 77 Varoquaux, G. and Craddock, R.C. (2013) Learning and comparing functional connectomes across subjects. *Neuroimage* 80, 405–415
- 78 Meskaldji, D.E. *et al.* (2013) Comparing connectomes across subjects and populations at different scales. *Neuroimage* 80, 416–425
- 79 Hutchison, R.M. *et al.* (2013) Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* 80, 360–378
- 80 Majeed, W. *et al.* (2010) Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *Neuroimage* 54, 1140–1150
- 81 Biswal, B.B. *et al.* (2010) Towards discovery science of human brain function. *Proc. Natl. Acad. Sci. U.S.A.* 107, 4734–4739
- 82 Yan, C.G. *et al.* (2013) Standardizing the intrinsic brain: towards robust measurement of inter-individual variation in 1000 functional connectomes. *Neuroimage* 80, 246–262

- 83 Carbonetto, P. and Stephens, M. (2012) Scalable variational inference for Bayesian variable selection in regression, and its accuracy in genetic association studies. *Bayesian Anal.* 7, 73–108
- 84 Pereira, F. *et al.* (2009) Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage* 45, S199–S209
- 85 Vidaurre, D. *et al.* (2013) A survey on L1 regression. *Int. Stat. Rev.* (in press)
- 86 Park, T. and Casella, G. (2008) The Bayesian lasso. *J. Am. Stat. Assoc.* 103, 681–686
- 87 Li, Q. and Lin, N. (2010) The Bayesian elastic net. *Bayesian Anal.* 5, 151–170
- 88 Virtanen, S. *et al.* (2011) Bayesian CCA via group sparsity. In *Proceedings of the 28th International Conference on Machine Learning*. International Machine Learning Society
- 89 Vidaurre, D. *et al.* (2013) Bayesian sparse partial least squares. *Neural Comput.* http://dx.doi.org/10.1162/NECO_a_00524
- 90 Hyvarinen, A. (2011) Testing the ICA mixing matrix based on inter-subject or inter-session consistency. *Neuroimage* 58, 122–136