Published in final edited form as:
Trends Cogn Sci. 2013 December ; 17(12): 666–682. doi:10.1016/j.tics.2013.09.016.

Functional connectomics from resting-state fMRI

Stephen M Smith[1], **Diego Vidaurre**[2], **Christian F Beckmann**[3,1], **Matthew F Glasser**[4], **Mark Jenkinson**[1], **Karla L Miller**[1], **Thomas E Nichols**[5], **Emma Robinson**[1], **Gholamreza Salimi-Khorshidi**[1], **Mark W Woolrich**[2,1], **Deanna M Barch**[4], **Kamil Uğurbil**[6], and **David C Van Essen**[4]

[1]FMRIB (Oxford Centre for Functional MRI of the Brain), Oxford University, Oxford, UK [2]OHBA (Oxford Centre for Human Brain Activity), Oxford University, Oxford, UK [3]Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The Netherlands & MIRA Institute for Biomedical Technology and Technical Medicine, University of Twente, Enschede, The Netherlands [4]Washington University School of Medicine, Washington University, St. Louis, Missouri, USA [5]Statistics Department, Warwick University, Warwick, UK [6]Center for Magnetic Resonance Research, University of Minnesota Medical School, Minneapolis, Minnesota, USA

Abstract

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 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.

Keywords

connectomics; resting-state fMRI; network modelling

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

Resting-state functional magnetic resonance imaging (rfMRI) has been used to study spontaneous fluctuations in brain activity since it was first noted that the rfMRI timeseries 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, while each has a relatively consistent timecourse across its set of involved regions, the different networks have distinct timecourses [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

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

[5]. RSNs have also been investigated in other modalities such as magnetoencephalography (MEG [6, 7]), but the majority of 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 into 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 sub-conscious 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 corresponding 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].
- 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]. In 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 carried out by clustering together neighbouring voxels (3D pixels) on the basis of similarity of their timeseries. 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 timeseries associated with the nodes (e.g., the average timeseries 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 $N_{nodes} \times N_{nodes}$ network matrix, as 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

Functional MRI data (both task-based and resting-state) is acquired as a series of volumetric images over time, with each image generally taking 2–3s to acquire. rfMRI data is typically acquired for 5–15 minutes, 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 Tesla can achieve a temporal resolution of 2–3s with a spatial resolution of 3–5mm. 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 2mm spatial resolution in less than a second. 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 4-dimensional rfMRI dataset requires extensive pre-processing before resting-state network analyses can be carried out. The pre-processing 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 co-ordinate system, for example, based on a population-average brain image. A standard sequence of processing steps [43, 44] is:

- Realign each timepoint'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 3-dimensional 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 in visualization and enables better inter-subject 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 inter-subject alignment, smoothing can be minimised. Unless smoothing is constrained to acting 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 very slowest temporal drifts in the data.
- Remove other artefacts.

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' timeseries, as these will be corrupted by artefacts that span multiple voxels. (In contrast, task-fMRI has the advantage of fitting a pre-specified temporal model, which provides greater robustness against artefactual influences.) Artefact cleanup [9] can involve one or more of:

- Regression of confound timeseries out of the data, for example, derived from:
  - Average white-matter and/or ventricle timeseries.
  - Head motion parameters (to further remove residual motion-related artefacts).
  - Separately measured cardiac and breathing signals.
  - Global-average timeseries (although many researchers consider this to be a blunt tool that makes the interpretation of the final correlations difficult [44]).
- Removal of corrupted timepoints [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, as 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].

Once pre-processing 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 5mm radius sphere centred in the precuneus. The average timecourse from this seed region is extracted, and every other voxel's timeseries 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(s)*. More information is obtained by a low-dimensional ICA decomposition, for example, reducing the data to 10–30 independent spatial maps, each of which is analogous to a distinct seed-based correlation map. This therefore 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 map(s) are related.

In contrast to seed-based correlation and low-dimensional ICA, a *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 timeseries), a detailed nodes+edges network model may reveal which nodes (sub-parts of the large-scale networks) are responsible for the correlations seen between the larger-scale networks. Put another way, a 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. In order 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 pre-processing, 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 timeseries based on that subject's rfMRI data, for example by averaging the timeseries from all voxels within a parcel. From the resulting $N_{timepoints} \times N_{nodes}$ data matrix, one can then compute the subject-specific $N_{nodes} \times N_{nodes}$ parcellated connectome matrix, for example by correlating each timeseries with each 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' timeseries. The simplest method, correlation between the timecourses 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 vs 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 very important for deciphering the underlying biological networks. For example, in a 3-node network A→B→C, all three nodes' timeseries 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 3-node network example, this works by taking each pair of timeseries in turn, and regressing out the third from each of the two timeseries in question, before estimating the correlation between the two. (For more than 3 nodes, all the other N_{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 applied to real rfMRI data (with full and partial correlation sitting at the "simple" extreme), and ends with abstract graph-theoretic 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 timeseries), 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, while DCM, originally developed for task data, 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 pre-specification 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 carry out 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 which is often quite unstable, depending on the quality of the data and the number of original timepoints. 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 timeseries 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 timepoints). 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:

- *Temporal lag* between pairs of timeseries [58] is often used, though 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' timeseries [59] and show some promise.
- *Non-Gaussianities* in the timeseries [60, 61] provide a distinct, potentially valuable source of causality information.

Because the majority of brain connections are bi-directional [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] [*ref 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 theoretic 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, rather than 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 actually 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 connections 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 NIH announced a program targeted at characterizing 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, NIH awarded Human Connectome Project grants to two consortia, one led by Washington University, the University of Minnesota and Oxford University (referred to below simply as "HCP") [67], and the other led by MGH and UCLA [16]. The latter is concentrating on leading-edge structural connectomics (via diffusion MRI), and so is not discussed further in this paper (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.

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 **WashU-UMinn-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 Kings 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-conceptional age), using rfMRI and diffusion MRI. It will also carry out 25 post mortem diffusion MRI scans at 7T.
- The **Thousand Functional Connectomes** (and related) projects [74, 75], 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, while not supporting the most sophisticated analysis techniques, will complement studies such as HCP by being able to find gross

imaging phenotypes and carrying out very-large-*N* subject-pathology correlations.

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 humanconnectome.org). Subjects are drawn from a population of healthy adult twins and their non-twin siblings (in the age range of 22–35 years), thus aiding the study of the influence of heritability and environmental factors on the connectome. From each subject 1h of whole-brain rfMRI data is acquired at 3T (in two pairs of 15-minute runs on separate days). A spatial resolution of 2×2×2mm and temporal resolution of 0.7s are achieved utilising an EPI acceleration factor of x8. 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 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 (0.7mm) structural data acquired for each subject. This allows for the transformation of the cortical rfMRI signal from the originally-acquired 3-dimensional voxel matrix onto a grey-matter surface mesh. Surface-based analysis is advantageous, as 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 - 2-dimensional mesh-surface vertices for cerebral cortex, plus 3-dimensional image-matrix voxels for subcortical grey matter (when all combined together, referred to as "grayordinates"); (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 ~2mm spacing.

As discussed above, if inter-subject alignment of rfMRI data is only driven 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).

Following image pre-processing (primarily using the FSL [71], FreeSurfer [72] and Connectome Workbench [73] software packages), the HCP applies data-driven artefact removal to rfMRI timeseries. First, ICA is applied to each 15-minute rfMRI dataset separately, to identify both structured artefact and non-artefact components. These components are then automatically classified into artefact vs non-artefact, using machine learning – a hand-trained hierarchical classifier (FIX - FMRIB's ICA-based X-noisifier (Salimi-Khorshidi, unpublished manuscript)). FIX achieves classification accuracy on HCP data of over 99%. Removal of the artefact components from the data is then carried out, improving data quality. Further additional cleanup approaches may be of value, and are under continuing investigation.

The HCP is already disseminating both raw and pre-processed rfMRI timeseries 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.

In Figure 4 we show full and partial correlation matrices, derived from the first 131 subjects' rfMRI data released by the HCP – a dataset aggregating over 600,000 timepoints. Group-ICA was run at a dimensionality of 100, and 2 components were discarded as artefactual, leaving 98 network nodes. (Although FIX cleanup is applied to each separate timeseries 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 symmetric, 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), carry out 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 both 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 timepoints 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 sub-clusters 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 cross-subject alignments are carried out prior to 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 only given 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 [76], as well as low-dimensional resting-state network maps, it can utilise richer structural and connectional information, and further improve the intersubject alignment ("$MSM_{multimodal}$") [43].

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 3 alignment methods: FreeSurfer=0.613, $MSM_{folding}$=0.618 and $MSM_{multimodal}$=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 [Robinson, in submission]; for example, Figure 5B shows stronger and more sharply defined activation patterns from a group-level analysis of two different tasks ("maths vs story" and a social interaction task) when $MSM_{multimodal}$ is utilised (lower panels) compared to $MSM_{folding}$ (upper panels), especially in parts of prefrontal cortex (blue arrows).

Figure 6 illustrates how the elements of the partial correlation matrix from each subject can be re-ordered 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., 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 timeseries) 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 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 caption). 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 [77] and cognitive state [78]; for additional discussions on cross-subject network modelling, see also [79, 80].

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 which contributes most strongly to the regression against fluid intelligence (shown in more detail in B) 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 important to correct the resulting p-values for the multiple comparisons across all of the network edges; however, Bonferroni correction would likely 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 9 edges that were significantly different in strength between the two sexes.

We then carried out *multivariate* analyses, attempting to predict subjects' sex, utilising *all* network edges simultaneously. The $N_{subjects} \times N_{edges}$ 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 3 alignment methods described above gave classification accuracies of: FreeSurfer=78%, $MSM_{folding}$=82% and $MSM_{multimodal}$=87%. This suggests that multimodal MSM improves the alignment of function across subjects (as 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.

Box 2

Advances in multivariate modelling and statistics

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

Feature selection may well 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 [82]. In order 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 re-run 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 a L2-penalty over the regression coefficients). Whereas pure L1-regularisation discards the redundant edges, the elastic net selects all the relevant features in a balanced fashion, and is particularly convenient for interpretation purposes [83]. 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 [84, 85].

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 $N_{subjects} \times N_{edges}$ population network matrix one at a time, we can form the $N_{subjects} \times N_{behavioural\text{-}variables}$ population behavioural matrix, and directly attempt to co-model this with the connectomes matrix. Approaches such as canonical correlation analysis (CCA) [86] and partial least squares (PLS) [87] use a set of latent variables to model the relation between the two matrices. Such approaches can provide great gains in sensitivity by increasing signal-to-noise and reducing problems of multiple comparisons.

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 [88]. In the latter approach, a windowed subset of timepoints (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 "timeseries" 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*, vs. *the internal connections within any given network being non-constant*. Both possibilities are neurobiologically interesting, and it is 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, as the rfMRI signal is dominated by low frequencies (~0.015Hz), 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 0.2Hz in good quality data; hence, aggressive highpass 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 "timeseries of correlation matrices" has been generated through sliding-window correlation, clustering analysis can be used to identify distinct correlation patterns that repeatedly re-appear over an extended period [88]. 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 co-activating networks' varying amplitudes will play 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 timeseries 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 timeseries 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 anti-correlation 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 a 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 timepoints 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 2 example results from feeding 98 nodes' timeseries 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 provide interpretation for 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 well 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 between the modes' timeseries, and this situation will 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 oversimplistic 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 [92], 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.

Conclusions

The macroscopic functional connectome as elucidated through resting-state functional MRI 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.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We are grateful for funding via the following grants: 1U54MH091657-01 (NIH Blueprint for Neuroscience Research), P30-NS057091, P41-RR08079/EB015894, F30-MH097312 (NIH), 098369/Z/12/Z (Wellcome Trust). We thank our 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, analyze, visualize, and share these datasets.

Glossary

Connectomics	the mapping of the brain's structural and/or functional connections
Diffusion Magnetic Resonance Imaging (dMRI)	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	A) As opposed to effective connectivity, functional connectivity estimates whether there is *any* functional connection between functional regions, even if indirect (polysynaptic). B) 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 sub-clusters 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 timecourse). 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 $N_{nodes} \times N_{nodes}$ 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 a 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, which attempts to infer the *direct* connections between nodes from analysis of correlations between multiple nodes' timeseries. Before correlating any two nodes' timeseries, all other timeseries 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 Magnetic Resonance Imaging (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 timeseries
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 (e.g.) neuroimaging timeseries 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 timeseries data, and the word "structural" here refers to the type of modelling, and not that it is structural connectivity data (such as dMRI) being used

References

1. Biswal B, et al. Functional Connectivity in the Motor Cortex of Resting Human Brain Using Echo-Planar MRI. Magnetic Resonance in Medicine. 1995; 34:537–541. [PubMed: 8524021]

2. Beckmann CF, et al. Investigations into Resting-State Connectivity using Independent Component Analysis. Philosophical Transactions of the Royal Society. 2005; 360:1001–1013.
3. Picchioni D, et al. Sleep and the functional connectome. Neuroimage. 2013; 80:387–396. [PubMed: 23707592]
4. Yeo BT, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol. 2011; 106:1125–1165. [PubMed: 21653723]
5. Li L, et al. Mapping putative hubs in human, chimpanzee and rhesus macaque connectomes via diffusion tractography. Neuroimage. 2013; 80:462–474. [PubMed: 23603286]
6. Brookes MJ, et al. Investigating the Electrophysiological Basis of Resting State Networks using Magnetoencephalography. Proc Natl Acad Sci USA (PNAS). 2011; 108:16783–16788.
7. de Pasquale F, et al. Temporal dynamics of spontaneous MEG activity in brain networks. Proceedings of the National Academy of Sciences. 2010; 107:6040–6045.
8. Castellanos FX, et al. Clinical applications of the functional connectome. Neuroimage. 2013; 80:527–540. [PubMed: 23631991]
9. Murphy K, et al. Resting-state fMRI confounds and cleanup. Neuroimage. 2013; 80:349–359. [PubMed: 23571418]
10. Sadaghiani S, Kleinschmidt A. Functional interactions between intrinsic brain activity and behavior. Neuroimage. 2013; 80:379–386. [PubMed: 23643921]
11. Laird AR, et al. Networks of task co-activations. Neuroimage. 2013; 80:505–514. [PubMed: 23631994]
12. Smith SM, et al. Correspondence of the brain's functional architecture during activation and rest. Proc Natl Acad Sci USA (PNAS). 2009; 106:13040–13045.
13. Raichle ME. The brain's dark energy. Scientific American. 2010; 302:44–49. [PubMed: 20184182]
14. Passingham RE. What we can and cannot tell about the wiring of the human brain. Neuroimage. 2013; 80:14–17. [PubMed: 23321152]
15. Sotiropoulos SN, et al. Advances in diffusion MRI acquisition and processing in the Human Connectome Project. NeuroImage. 2013; 80:125–143. [PubMed: 23702418]
16. Setsompop K, et al. Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. Neuroimage. 2013; 80:220–233. [PubMed: 23707579]
17. McNab JA, et al. The Human Connectome Project and beyond: Initial applications of 300mT/m gradients. Neuroimage. 2013; 80:234–245. [PubMed: 23711537]
18. O'Donnell LJ, et al. Fiber clustering versus the parcellation-based connectome. Neuroimage. 2013; 80:283–289. [PubMed: 23631987]
19. Mangin JF, et al. Toward global tractography. Neuroimage. 2013; 80:290–296. [PubMed: 23587688]
20. Kennedy H, et al. Why data coherence and quality is critical for understanding interareal cortical networks. Neuroimage. 2013; 80:37–45. [PubMed: 23603347]
21. da Costa NM, Martin KA. Sparse reconstruction of brain circuits: Or, how to survive without a microscopic connectome. Neuroimage. 2013; 80:27–36. [PubMed: 23624494]
22. Stephan KE. The history of CoCoMac. Neuroimage. 2013; 80:46–52. [PubMed: 23523808]
23. White JG, et al. The structure of the nervous system of the nematode Caenorhabditis elegans. Philos Trans Roy Soc Lond B Biol Sci. 1986; 314:1–340. [PubMed: 22462104]
24. Scholvinck ML, et al. The contribution of electrophysiology to functional connectivity mapping. Neuroimage. 2013; 80:297–306. [PubMed: 23587686]
25. David O, et al. Probabilistic functional tractography of the human cortex. Neuroimage. 2013; 80:307–317. [PubMed: 23707583]
26. Larson-Prior L, et al. Adding dynamics to the Human Connectome Project with MEG. NeuroImage. 2013; 80:190–201. [PubMed: 23702419]
27. Caspers S, et al. Microstructural grey matter parcellation and its relevance for connectome analyses. Neuroimage. 2013; 80:18–26. [PubMed: 23571419]
28. Evans AC. Networks of anatomical covariance. Neuroimage. 2013; 80:489–504. [PubMed: 23711536]

29. Barch DM, et al. Function in the Human Connectome: Task-fMRI and Individual Differences in Behavior. NeuroImage. 2013; 80:169–189. [PubMed: 23684877]
30. Thompson PM, et al. Genetics of the connectome. Neuroimage. 2013; 80:475–488. [PubMed: 23707675]
31. Sporns O. The human connectome: Origins and challenges. Neuroimage. 2013; 80:53–61. [PubMed: 23528922]
32. Blumensath T, et al. Spatially constrained hierarchical parcellation of the brain with resting-state FMRI. NeuroImage. 2013; 76:313–324. [PubMed: 23523803]
33. de Reus MA, van den Heuvel MP. The parcellation-based connectome: Limitations and extensions. Neuroimage. 2013; 80:397–404. [PubMed: 23558097]
34. Beckmann CF. Modelling with independent components. NeuroImage. 2012; 62:891–901. [PubMed: 22369997]
35. Smith SM. The Future of FMRI Connectivity. NeuroImage. 2012; 62:1257–1266. [PubMed: 22248579]
36. Margulies DS, et al. Visualizing the human connectome. Neuroimage. 2013; 80:445–461. [PubMed: 23660027]
37. Ogawa S, et al. Intrinsic Signal Changes Accompanying Sensory Stimulation: Functional Brain Mapping with Magnetic Resonance Imaging. Proc Natl Acad Sci USA. 1992; 89(13):5951–5955. [PubMed: 1631079]
38. Mansfield P. Multi-planar image formation using NMR spin echoes. Journal of Physics C: Solid State Physics. 1977; 10:L55–L58.
39. Moeller S, et al. 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. 2010; 63:1144–1153. [PubMed: 20432285]
40. Setsompop K, et al. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planer imaging with reduced g-factor penalty. Magnetic Resonance in Medicine. 2012; 67:1210–1224. [PubMed: 21858868]
41. Feinberg DA, et al. Multiplexed Echo Planar Imaging for Sub-second Whole Brain FMRI and Fast Diffusion Imaging. PLoS ONE. 2010; 5:e15710. [PubMed: 21187930]
42. Smith SM, et al. Temporally-independent functional modes of spontaneous brain activity. Proc Natl Acad Sci USA (PNAS). 2012; 109:3131–3136.
43. Glasser MF, et al. The Minimal Preprocessing Pipelines for the Human Connectome Project. NeuroImage. 2013; 80:105–124. [PubMed: 23668970]
44. Smith SM, et al. Resting-state fMRI in the Human Connectome Project. NeuroImage. 2013; 80:144–168. [PubMed: 23702415]
45. Power JD, et al. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage. 2011; 59:2142–2154. [PubMed: 22019881]
46. De Martino F, et al. Classification of fMRI independent components using IC-fingerprints and support vector machine classifiers. Neuroimage. 2007; 34:177–194. [PubMed: 17070708]
47. Niazy, RK., et al., editors. Spectral characteristics of resting state networks. Elsevier; 2011.
48. Friston KJ. Functional and effective connectivity: A review. Brain Connectivity. 2011; 1:13–36. [PubMed: 22432952]
49. Marrelec G, et al. Partial correlation for functional brain interactivity investigation in functional MRI. NeuroImage. 2006; 32:228–237. [PubMed: 16777436]
50. Friston KJ. Functional and Effective Connectivity in Neuroimaging: A Synthesis. Human Brain Mapping. 1994; 2:56–78.
51. Nakagawa TT, et al. Bottom up modeling of the connectome: Linking structure and function in the resting brain and their changes in aging. Neuroimage. 2013; 80:318–329. [PubMed: 23629050]
52. Woolrich MW, Stephan KE. Biophysical network models and the human connectome. Neuroimage. 2013; 80:330–338. [PubMed: 23571421]
53. McIntosh AR, Gonzales-Lima F. Structural Equation Modeling and Its Application to network Analysis in Functional Brain Imaging. Human Brain Mapping. 1994; 2:2–22.

54. Friston KJ, et al. Network discovery with DCM. NeuroImage. 2011; 56:1202–1221. [PubMed: 21182971]
55. Friedman J, et al. Sparse inverse covariance estimation with the Graphical Lasso. Biostat. 2008; 9:432–441.
56. Smith SM, et al. Network Modelling Methods for FMRI. NeuroImage. 2011; 54:875–891. [PubMed: 20817103]
57. Varoquaux G, et al. Brain covariance selection: Better individual functional connectivity models using population prior. Advances in Neural Information Processing Systems. 2010
58. Granger CWJ. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. Econometrica. 1969; 37:424–438.
59. Ramsey JD, et al. Six problems for causal inference from fMRI. NeuroImage. 2010; 49:1545–1558. [PubMed: 19747552]
60. Patel RS, et al. A Bayesian Approach to Determining Connectivity of the Human Brain. Human Brain Mapping. 2006; 27:267–276. [PubMed: 16092131]
61. Hyvarinen A, Smith SM. Pairwise Likelihood Ratios for Estimation of Non-Gaussian Structural Equation Models. Journal of Machine Learning Research. 2013; 14:111–152.
62. Markov N, et al. 2013A Weighted and Directed Interareal Connectivity Matrix for Macaque Cerebral Cortex. Cerebral Cortex. 10.1093/cercor/bhs270
63. Pearl J. Causal inference in statistics: An overview. Statistics Surveys. 2009; 3:96–146.
64. Rubinov M, Sporns O. Complex network measures of brain connectivity: Uses and interpretations. Neuroimage. 2010; 52:1059–1069. [PubMed: 19819337]
65. Fornito A, et al. Graph analysis of the human connectome: Promise, progress, and pitfalls. Neuroimage. 2013; 80C:426–444. [PubMed: 23643999]
66. Craddock RC, et al. A whole brain fMRI atlas generated via spatially constrained spectral clustering. Human Brain Mapping. 2011; 33:1914–1928. [PubMed: 21769991]
67. Van Essen DC, et al. The WU-Minn Human Connectome Project: An Overview. NeuroImage. 2013; 80:62–79. [PubMed: 23684880]
68. Ugurbil K, et al. Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. NeuroImage. 2013; 80:80–104. [PubMed: 23702417]
69. Van Essen DC, et al. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atalses. Cerebral Cortex. 2012; 22:2241–2262. [PubMed: 22047963]
70. Robinson, E., et al., editors. Multimodal Surface Matching: Fast and Generalisable Cortical Registration using Discrete Optimisation. Springer; Berlin Heidelberg: 2013.
71. Jenkinson M, et al. FSL. NeuroImage. 2012; 62:782–790. [PubMed: 21979382]
72. Fischl B. FreeSurfer. Neuroimage. 2012; 62:774–781. [PubMed: 22248573]
73. Marcus DS, et al. Human Connectome Project Informatics: Quality control, database services, and data visualization. NeuroImage. 2013; 80:202–219. [PubMed: 23707591]
74. Biswal BB, et al. Towards discovery science of human brain function. Proc Natl Acad Sci USA (PNAS). 2010; 107:4734–4739.
75. Yan CG, et al. Standardizing the intrinsic brain: Towards robust measurement of inter-individual variation in 1000 functional connectomes. Neuroimage. 2013; 80:246–262. [PubMed: 23631983]
76. Glasser, MF., et al. Comparison of Surface Gradients Derived from Myelin Maps and Functional Connectivity Analysis. Int. Conf. on Functional Mapping of the Human Brain; 2011.
77. Craddock RC, et al. Disease state prediction from resting state functional connectivity. Magnetic Resonance in Medicine. 2009; 62:1619–1628. [PubMed: 19859933]
78. Shirer WR, et al. Decoding subject-driven cognitive states with whole-brain connectivity patterns. Cereb Cortex. 2012; 22:158–165. [PubMed: 21616982]
79. Varoquaux G, Craddock RC. Learning and comparing functional connectomes across subjects. Neuroimage. 2013; 80:405–415. [PubMed: 23583357]
80. Meskaldji DE, et al. Comparing connectomes across subjects and populations at different scales. Neuroimage. 2013; 80:416–425. [PubMed: 23631992]

81. Carbonetto P, Stephens M. Scalable variational inference for Bayesian variable selection in regression, and its accuracy in genetic association studies. Bayesian Analysis. 2012:73–108.
82. Pereira F, et al. Machine learning classifiers and fMRI: A tutorial overview. Neuroimage. 2009:S199–209. [PubMed: 19070668]
83. Vidaurre D, et al. A survey on L1 regression. International Statistical Review. 2013 In press.
84. Park T, Casella G. The Bayesian Lasso. Journal of the American Statistical Association. 2008; 103(482):681–686.
85. Li Q, Lin N. The Bayesian elastic net. Bayesian Analysis. 2010:151–170.
86. Virtanen, S., et al. Bayesian CCA via group sparsity. 28th International Conference on Machine Learning; 2011.
87. Vidaurre D, et al. Bayesian Sparse Partial Least Squares. Neural Computation. 201310.1162/NECO_a_00524
88. Hutchison RM, et al. Dynamic functional connectivity: Promise, issues, and interpretations. Neuroimage. 2013; 80:360–378. [PubMed: 23707587]
89. Hyvarinen A. Testing the ICA mixing matrix based on inter-subject or inter-session consistency. Neuroimage. 2011; 58:122–136. [PubMed: 21704714]
90. Binder JR, et al. Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. Cerebral Cortex. 2009; 19:2767–2796. [PubMed: 19329570]
91. Watkins KE, et al. Language networks in anophthalmia: maintained hierarchy of processing in ‘visual’ cortex. Brain. 2012; 135:1566–1577. [PubMed: 22427328]
92. Majeed W, et al. Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. Neuroimage. 2010; 54:1140–1150. [PubMed: 20728554]

Box 3

Questions for future research

- 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 measurement 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 voxelwise or "dense" connectome, or from low-dimensional spatial-ICA) with the richer "*temporal*" network modelling that can be carried out 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 sub-networks, 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 electro-chemical modulation of each other give rise to functional dynamics, behaviour and consciousness, and how are these altered in early development, disease and aging?

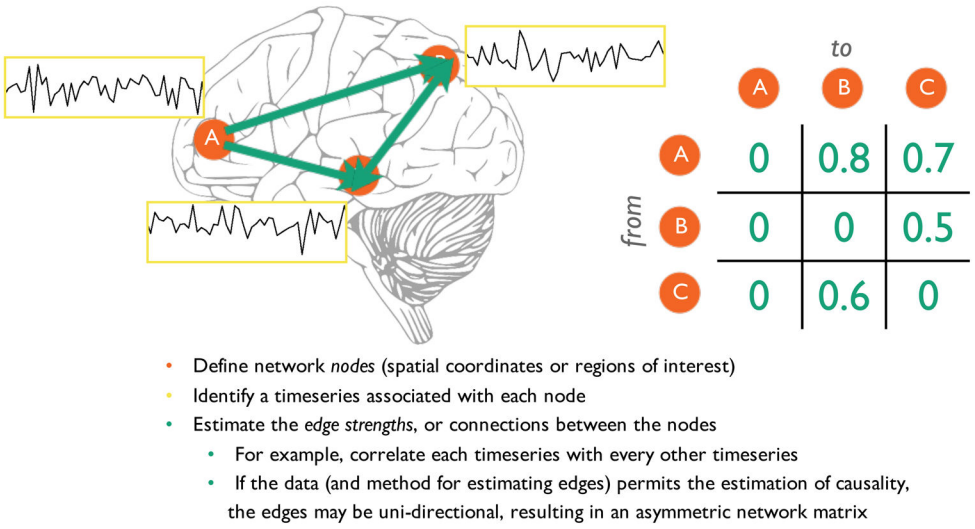

Figure 1.

Illustration of the main steps that take rfMRI data (with an activity timeseries at every point in the brain), identify network nodes, and then estimate network edges.

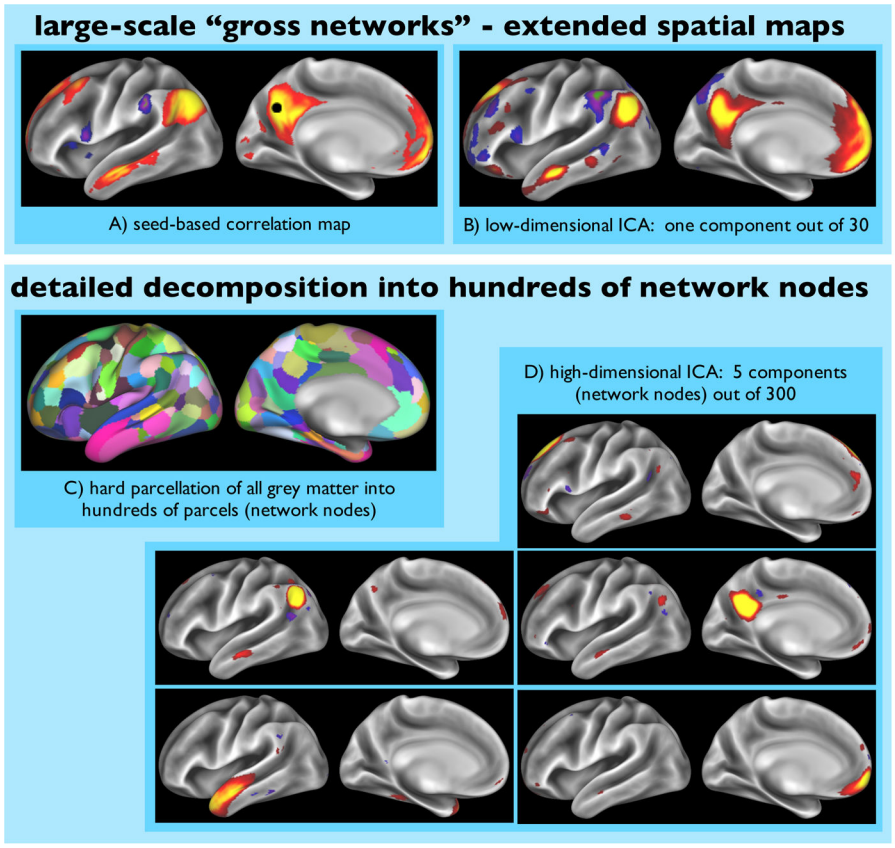

Figure 2.

Examples of different approaches for connectivity analysis. All results are derived from group-average connectivity from the first 131 subjects’ rfMRI data released publicly by the HCP (available from 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) 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 indicates anticorrelation with that component’s associated timeseries. In the case of high-dimensional parcellations, anticorrelations would likely be seen not in the spatial maps, but as negative correlations between the timeseries associated with different parcels. The cortical surface renderings were generated using the Connectome Workbench display tool (humanconnectome.org/connectome/connectome-workbench.html).

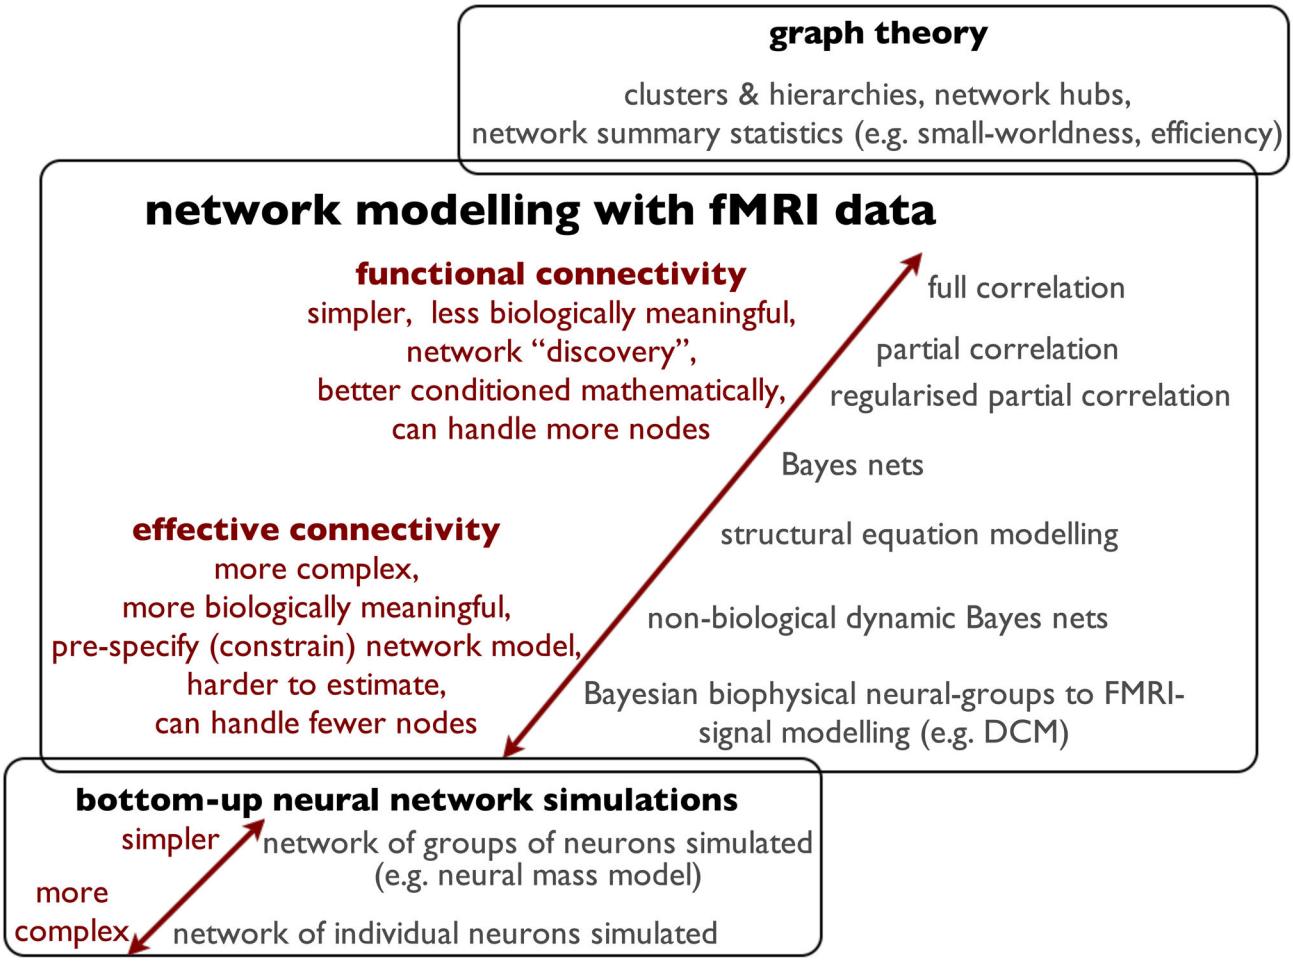

Figure 3.

Schematic of relationships between various network modelling analyses for/from fMRI, adapted from Smith12b

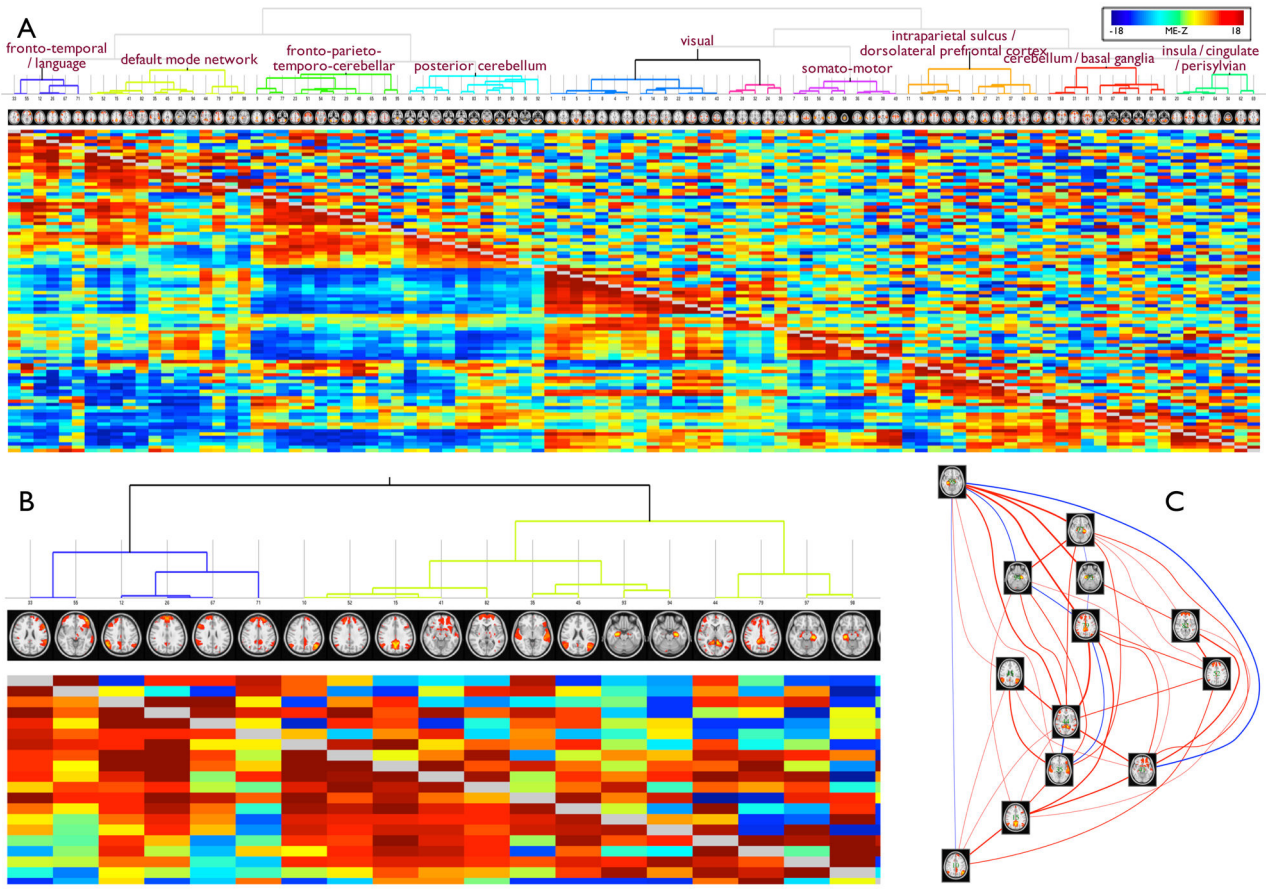

Figure 4.

Functional connectomes, estimated for 98 nodes derived from 131 HCP subjects' rfMRI data using group-ICA (with FSL's MELODIC tool). A) Full and partial correlation matrices. Full correlation is shown below the diagonal, and partial correlation above. Intersubject alignment was carried out using the MSM method (see Fig. 5). Small images at the top of each column summarize 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 is carried out using the FSLNets package (fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets), with (C) created using Graphviz (www.graphviz.org).

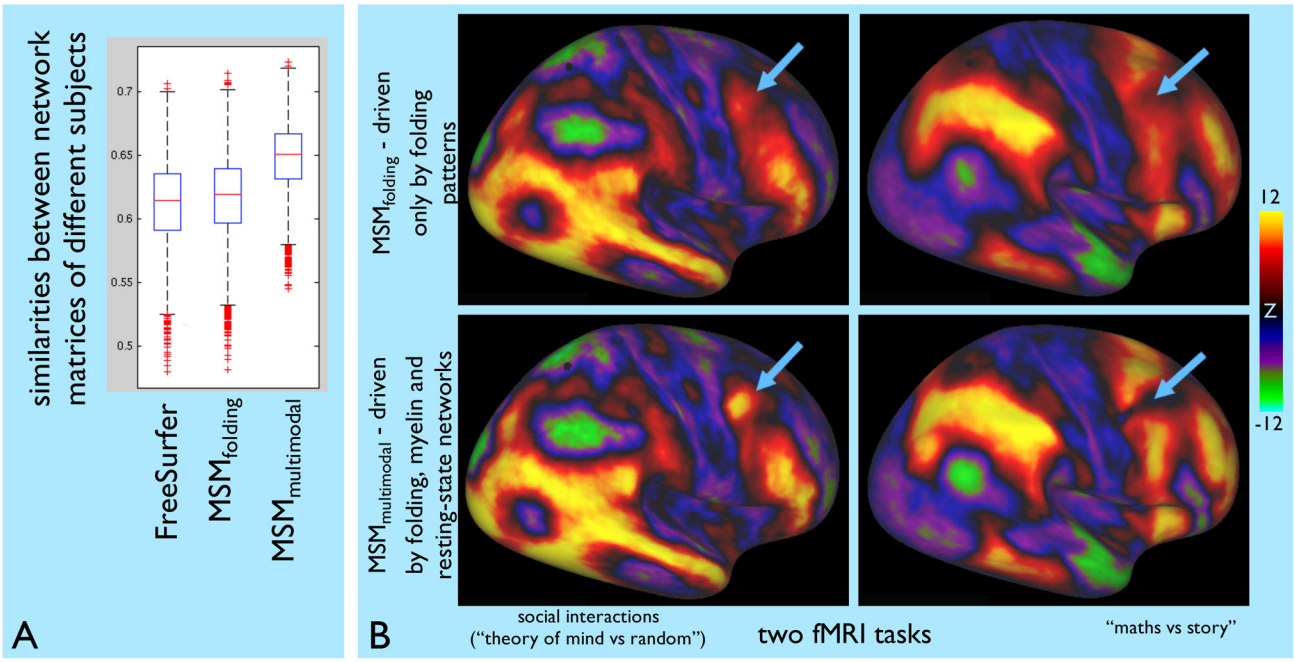

Figure 5.

Example analyses showing the effects of improved surface alignment across subjects. A) Similarities between pairs of subjects' network matrices. All 8417 pairings of non-related subjects in the HCP data were used, after application of 3 different cross-subject surface alignment methods. All 3 group means are statistically significantly different from each other. B) Group-level analyses of two tasks from 120 HCP subjects' task-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 MSM utilizes multimodal information, including resting-state networks. The arrows mark example regions showing improved spatial localization of activation.

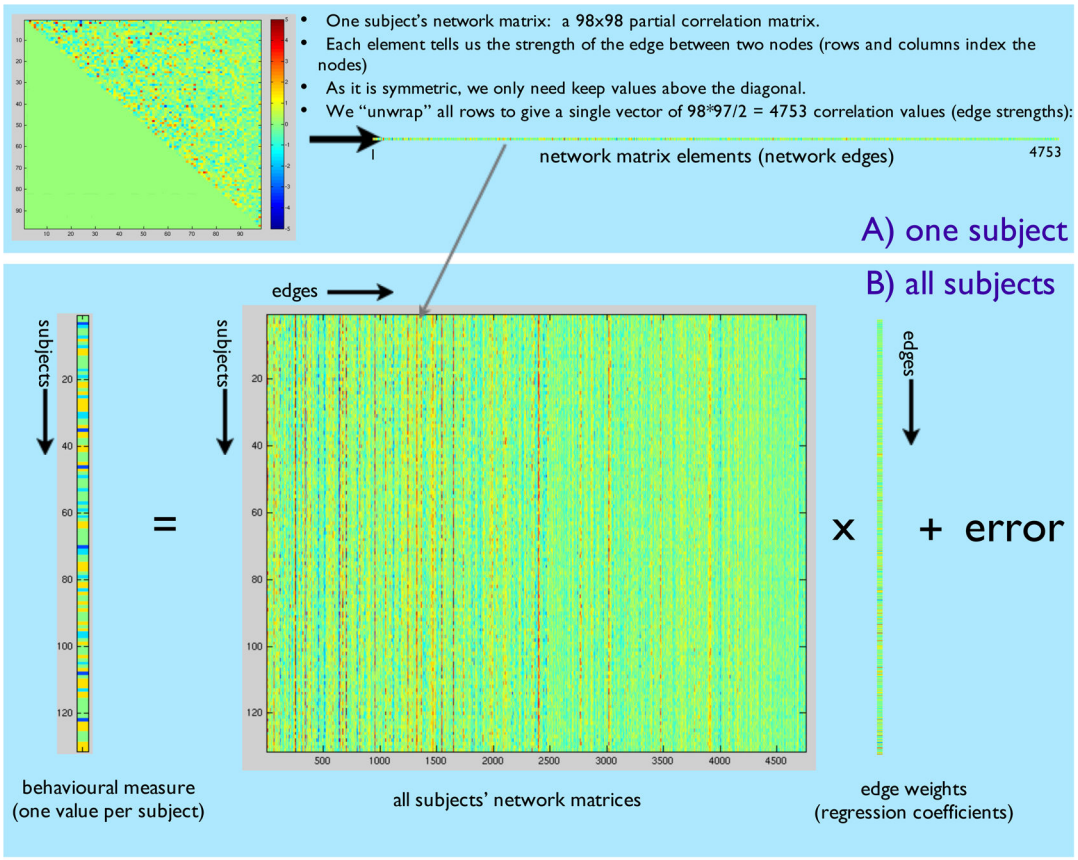

Figure 6.

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 re-ordered into a single row, in order to prepare for combining all network matrix estimates across subjects. B) The resultant $N_{subjects} \times N_{edges}$ 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.

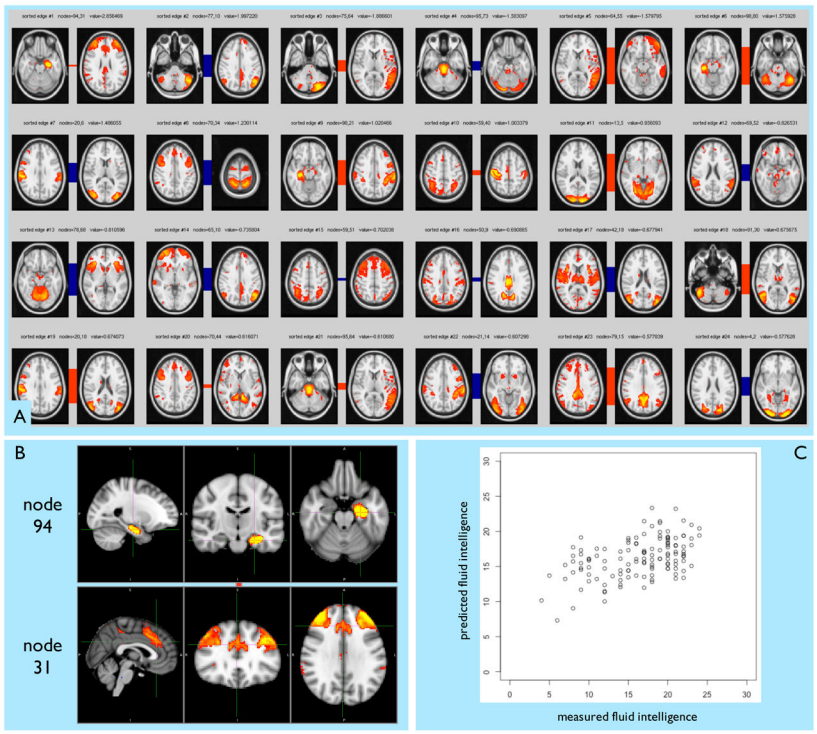

Figure 7.

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 [81] that only kept the 98 edges most strongly predictive of FI, and applied ridge regression (L2) shrinkage on those kept features (see Box 2); over-fitting was avoided by carrying out 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 which 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 vs. measured FI, with one data point per subject; each subject's FI was predicted using their network matrix, where the linear regression model was trained excluding the data from that subject (and all of their family members).

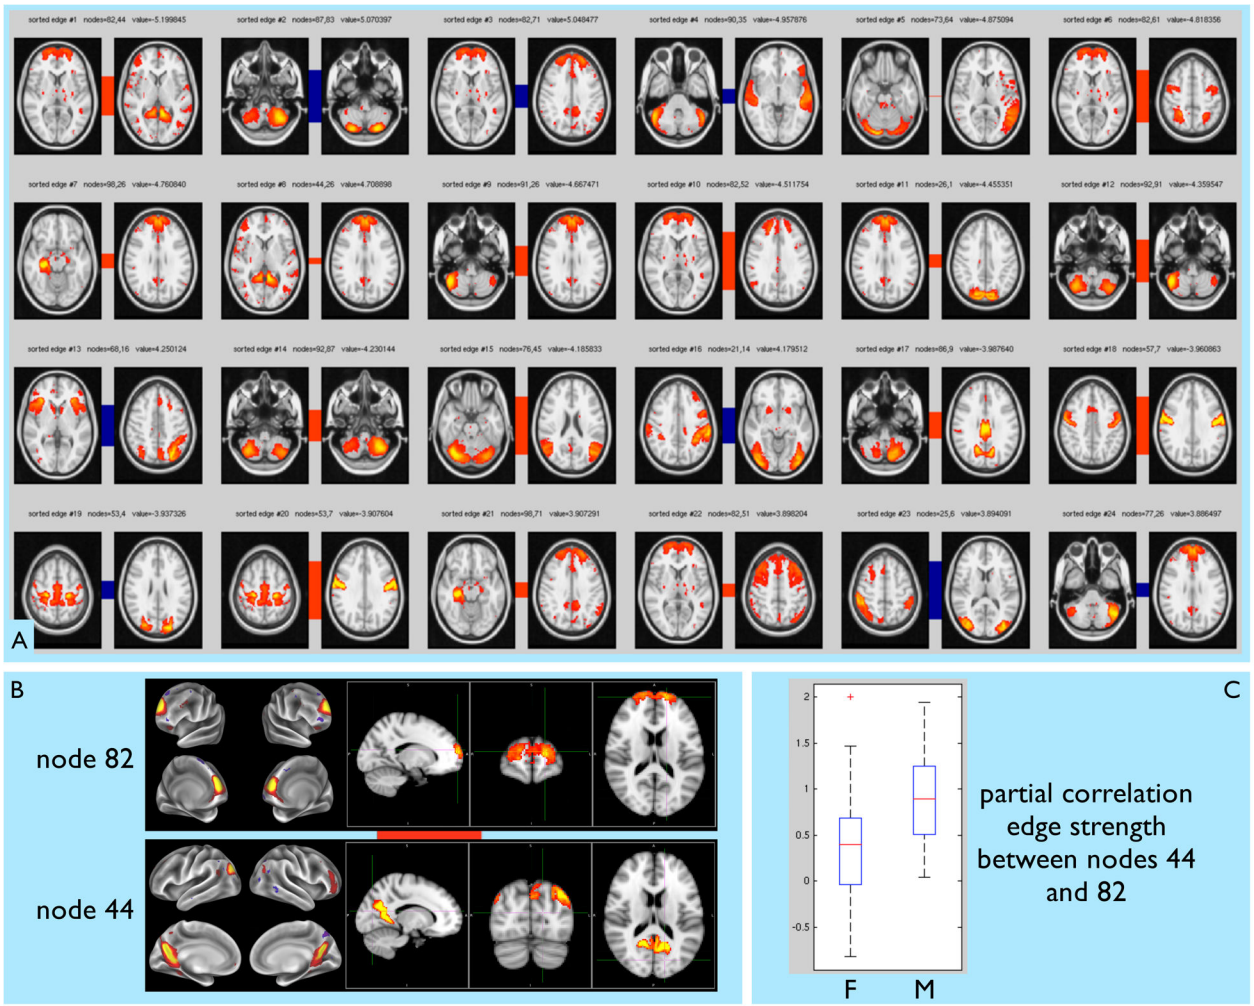

Figure 8.

Relationship between edge strength and sex. We discarded the second subject of any pair of twins (leaving 104 subjects), to avoid the danger of overfitting (in the case of sex prediction) or confounding the statistical inference (when carrying out the univariate t-tests), given that twins are likely to have similar network matrices to each other. A) The 24 edges which are most different in connection strength between males and females are shown; the first 9 are significantly different, $p<0.05$, two-tailed, corrected for multiple comparisons using FSL's *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.

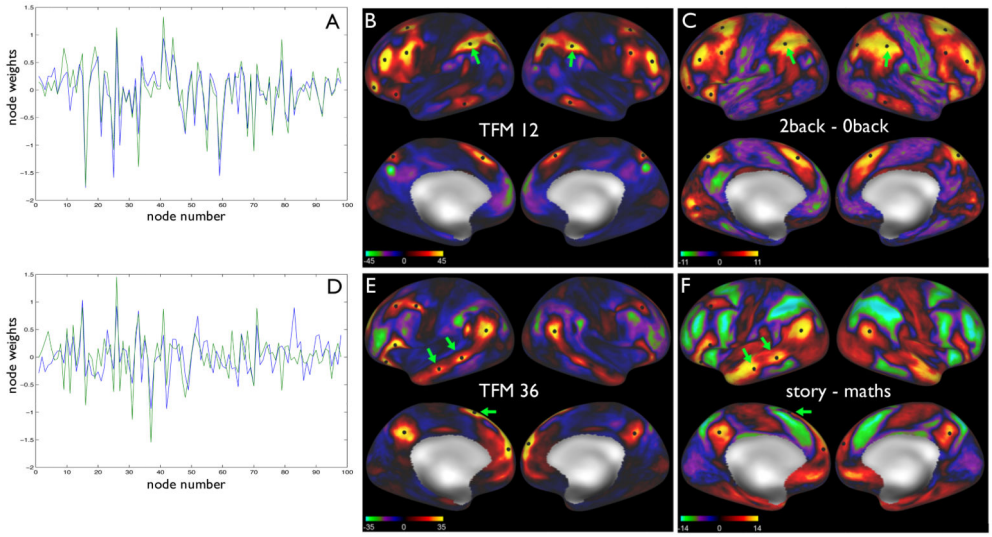

Figure 9.

Two “Temporal Functional Modes” (TFMs) [42], estimated from 131 subjects’ data from the HCP, with cross-subject alignment carried out using $MSM_{multimodal}$. 98 nodes’ timeseries were fed into temporal-ICA and 98 ICA components were estimated. Split-half reproducibility testing ([89], 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 voxelwise 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 sub-groups 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-fMRI datasets, with a specific task chosen that matches spatially the TFM found from the resting data. This TFM closely matches the 2-back vs. 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 vs. maths’ contrast, comparing listening to stories vs. 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 while sharing several nodes, these two TFMs have a different overall spatial makeup than each other, and are reproducible across the two sets of distinct subjects. See also the [Supplemental_TFMs_Movie], which shows TFM brain dynamics from 10 subjects, one after another. For each subject, 1 minute of dynamics is shown, at X4 real time. Different colours are different TFMs; for a given TFM, darker clusters are anticorrelated with brighter clusters (www.fmrib.ox.ac.uk/~steve/ftp/TFMs.mov).