NeuroImage 60 (2012) 1055–1062

Contents lists available at SciVerse ScienceDirect

NeuroImage

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

Connectivity differences in brain networks

Andrew Zalesky [a,*], Luca Cocchi [b], Alex Fornito [a,f,g], Micah M. Murray [c,d], Ed Bullmore [e,h]

[a] Melbourne Neuropsychiatry Centre, The University of Melbourne and Melbourne Health, Melbourne, Australia
[b] Queensland Brain Institute, The University of Queensland, Brisbane, Australia
[c] Electroencephalography Brain Mapping Core, Center for Biomedical Imaging, Vaudois University Hospital Center and University of Lausanne Switzerland, Lausanne, Switzerland
[d] Department of Clinical Neurosciences and Department of Radiology, Vaudois University Hospital Center and University of Lausanne, Lausanne, Switzerland
[e] Behavioural and Clinical Neuroscience Institute, The University of Cambridge, Cambridge, UK
[f] NICTA Victorian Research Laboratory, The University of Melbourne, Melbourne, Australia
[g] Centre for Neural Engineering, The University of Melbourne, Melbourne, Australia
[h] Clinical Unit Cambridge, GlaxoSmithKline Cambridgeshire & Peterborough NHS Foundation Trust, Cambridge, UK

ARTICLE INFO

Article history:
Received 3 October 2011
Revised 4 January 2012
Accepted 8 January 2012
Available online 16 January 2012

Keywords:
Brain connectivity
Brain network
Connectome
Network-based statistic
Spatial pairwise clustering
Family-wise errors

ABSTRACT

The scenario considered here is one where brain connectivity is represented as a network and an experimenter wishes to assess the evidence for an experimental effect at each of the typically thousands of connections comprising the network. To do this, a univariate model is independently fitted to each connection. It would be unwise to declare significance based on an uncorrected threshold of $\alpha = 0.05$, since the expected number of false positives for a network comprising $N = 90$ nodes and $N(N-1)/2 = 4005$ connections would be 200. Control of Type I errors over all connections is therefore necessary. The network-based statistic (NBS) and spatial pairwise clustering (SPC) are two distinct methods that have been used to control family-wise errors when assessing the evidence for an experimental effect with mass univariate testing. The basic principle of the NBS and SPC is the same as supra-threshold voxel clustering. Unlike voxel clustering, where the definition of a voxel cluster is unambiguous, 'clusters' formed among supra-threshold connections can be defined in different ways. The NBS defines clusters using the graph theoretical concept of connected components. SPC on the other hand uses a more stringent pairwise clustering concept. The purpose of this article is to compare the pros and cons of the NBS and SPC, provide some guidelines on their practical use and demonstrate their utility using a case study involving neuroimaging data.

Introduction

There has been a shift in imaging neuroscience from brain activation to brain connectivity (Friston, 2009; Sporns, in press). Central to this shift in focus has been an emphasis on studying large-scale brain networks composed of nodes and connections (Bullmore and Sporns, 2009; Habeck and Moeller, 2011; He and Evans, 2010; Kaiser, 2011; Sporns, 2011; Wig et al., 2011). Nodes represent brain regions and the connections formed between pairs of nodes represent some measure of interaction between them, as inferred from neuroimaging data (Rubinov and Sporns, 2011).

Brain networks have been found to exhibit various nontrivial topological features, such as small-world organization, modular structure and highly connected hubs (Achard et al., 2006; Bassett and Bullmore, 2006; Hagmann et al., 2008; van den Heuvel et al., 2008). The goal of numerous studies has been to elucidate differences in these topological properties over developmental stages (Fair et al., 2009), in clinical conditions (e.g. Lynall et al., 2011; He et al., 2008; van den Heuvel et al., 2010) and in relation to different experimental conditions and cognitive states (e.g. Bassett et al., 2011; Fornito et al., 2011a; Kitzbichler et al., 2011) as well as genetic influences (Fornito et al., 2011b).

The interpretation of topological differences found in brain networks is not always straightforward, however. Topological properties derived from the characteristic path length in functional brain networks (Wang et al., 2010) are particularly difficult to interpret because functional networks are intrinsically fully connected. Therefore, the "path length" between a pair of regions is already explicitly captured by the strength of the direct connection (Rubinov and Sporns, 2010). Negative functional connections complicate the interpretation of path length as well (Chen et al., 2011).

Furthermore, topological differences can in some circumstances be a complex manifestation of simple differences in connectivity strength. Steps taken to disambiguate topological differences from simple differences in connectivity strength are equivocal and typically require the selection of arbitrary thresholds to transform connectivity strength from a continuous to a binary scale (Ginestet et al., 2011; van Wijk et al., 2010) (see Fig. 1).

* Corresponding author at: Melbourne Neuropsychiatry Centre, c/o National Neuroscience Facility, The University of Melbourne, Level 3, Alan Gilbert Building, 161 Barry Street, Carlton South, 3053, Australia. Fax: +61 3 9348 0469.
E-mail address: azalesky@unimelb.edu.au (A. Zalesky).

1053-8119/$ – see front matter © 2012 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2012.01.068

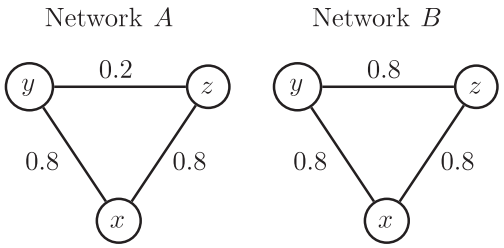

Fig. 1. Is there a topological difference between networks *A* and *B*, or do the networks simply differ in connectivity strength? If the networks are analyzed as weighted graphs, then they are topologically different. For example, $c_A = 0.5$ whereas $c_B = 0.8$, where c_x denotes the average clustering coefficient (Onnela et al., 2005) for network *x*. On the other hand, the networks are topologically indistinct if they are analyzed as binary graphs that are matched in terms of connection density.

Elucidating differences in the strength of connectivity is therefore an important undertaking in and of itself. Differences in connectivity strength are more basic than topological differences and as such are more straightforward to interpret.

This article considers exploratory methods for assessing the evidence of an experimental effect at each of the typically thousands of connections comprising a brain network. The experimental effect may be, for example, an association between connectivity strength and diagnostic status in a case-control study or contextual changes during performance of a cognitive task (Bressler and Menon, 2010).

The scenario considered is one in which connectivity is measured between every pair of many distinct brain regions. Connectivity includes anatomical connectivity inferred from fiber tracking methods (Bassett et al., 2010; Li et al., in press) and cortical thickness/volume estimates (Bassett et al., 2008; He et al., 2007) as well as functional connectivity inferred from functional imaging (van den Heuvel and Hulshoff Pol, 2010) or electromagnetic tomography (Schoffelen and Gross, 2011). To enable inferential statistics, the same connectivity measurements are repeated for each subject comprising a case-control study or in the same subject during different experimental conditions. A univariate model is then independently fitted to each connection to assess the evidence of an experimental effect. This involves computing a test statistic and corresponding *p*-value for the contrast of interest.

The total number of connections is typically in the thousands. Control of Type I errors among all connections is therefore essential. The network-based statistic (NBS) (Zalesky et al., 2010a) and spatial pairwise clustering (SPC) (Hipp et al., 2011; Zalesky et al., in press) are two distinct methods that have been used to control family-wise errors when assessing the evidence for an experimental effect with mass univariate testing. The family-wise error rate refers to the likelihood of committing one or more Type I errors among all connections (Nichols and Hayasaka, 2003).

The basic principle of the NBS and SPC is the same as supra-threshold voxel clustering in traditional task-based functional-MRI activation studies (Bullmore et al., 1999; Nichols and Holmes, 2001). Whereas voxel clustering pertains to mass univariate testing of brain activation, the NBS and SPC pertain to mass univariate testing of brain connectivity. Unlike voxel clustering, where the definition of a voxel cluster is unambiguous, 'clusters' formed among supra-threshold connections can be defined in different ways. The only obvious way to form voxel clusters is to cluster supra-threshold voxels that share at least one common face, edge or corner. Supra-threshold voxels refers to voxels having a test statistic that exceeds a chosen cluster-forming threshold.

In contrast, there is not one obvious way to form 'clusters' among supra-threshold connections. The NBS defines clusters using the graph theoretical concept of connected components. SPC on the other hand uses a more stringent pairwise clustering concept. With the development of these two complimentary methods, experimenters face a choice: NBS or SPC?

This choice is addressed by comparing the pros and cons of the NBS and SPC, providing some guidelines on their practical use and demonstrating their utility using a case study involving connectivity measurements inferred from electroencephalography data. Note that the NBS is freely available as part of the Brain Connectivity Toolbox (http://www.brain-connectivity-toolbox.net/) and Connectome Mapping Toolkit (http://www.connectomics.org).

Methods

Assume brain connectivity is measured between every pair of N distinct brain regions. To assess the evidence for an experimental effect, a univariate model is independently fitted to each connection and a test statistic is computed for the contrast of interest (e.g. *F*-statistic, *t*-statistic). Each connection is therefore endowed with a test statistic and corresponding uncorrected *p*-value.

A typical value of N is 90 (Zalesky et al., 2010b), in which case the total number of connections is $N(N-1)/2 = 4005$, since each node can form a connection with every other node, apart from itself. It would be unwise to declare significance based on an uncorrected threshold of $\alpha = 0.05$ because the expected number of false positives would be 200. Control of Type I errors among all connections is therefore essential.

The most conservative approach is to control the family-wise error rate with the Bonferroni correction or the Holm–Bonferroni method (Holm, 1979). However, the Bonferroni correction is generally too conservative, particularly given that the number of multiple tests, $N(N-1)/2$, grows quadratically with the number of nodes, N. For $N = 90$, a true positive must have a *p*-value that is less than approximately $\alpha/4005 \approx 0.00001$ if it is to survive the Bonferroni correction using a family-wise error rate of $\alpha = 0.05$. A less conservative approach is to control the false discovery rate (Genovese et al., 2002), but it too can be underpowered when N is large (Zalesky et al., 2010a).

To increase statistical power, a simple strategy is to reduce the total number of connections, and thereby reduce the number of multiple tests. The number of multiple tests can be reduced by excluding certain nodes or connections *a priori* based on case-specific assumptions. For example, it may be desirable to explicitly exclude short-range connections to alleviate spurious connectivity induced by volume conduction effects in the case of electroencephalography (Nolte et al., 2004; Schoffelen and Gross, 2009), or smoothing in the case of functional-MRI.

Another strategy is to use a coarser parcellation template to subdivide the brain into distinct nodes (Hayasaka and Laurienti, 2010; Wang et al., 2009). Fewer nodes mean fewer connections. However, using a coarser parcellation template is detrimental to the spatial resolution at which focal effects can be localized. A focal effect is likely to encompass a proportionally small volume of the larger node within which it is encapsulated. Therefore, a focal effect may be obscured once averaging is performed over the entire volume of the larger node.

These limitations have led to the development of two new methods for controlling Type I errors. The two methods are called the network-based statistic (NBS) and spatial pairwise clustering (SPC). Both methods exploit the tendency for experimental effects involving brain connectivity to exhibit specific spatial/topological characteristics that would not be expected as a matter of chance alone in the absence of an effect.

Connections with a test statistic exceeding a chosen threshold are first admitted to a set of supra-threshold connections. Supra-threshold connections essentially represent uncorrected effects that survive the application of a primary, uncorrected cluster-forming threshold. The NBS and SPC search for distinct spatial/topological patterns, referred to as clusters, in the set of supra-threshold connections that provide evidence of an experimental effect. Permutation testing is then used to ascribe a family-wise error corrected *p*-value to each cluster.

The NBS and SPC assess the evidence for an experimental effect at the cluster level, rather than at the level of each connection. In this

way, they control the family-wise error rate in the weak sense (Maris and Oostenveld, 2007; Nichols and Holmes, 2001). This means that the null hypothesis cannot be rejected individually for particular connections comprising a cluster, but only for the cluster in totality.

What has led to the development of two distinct methods? Unlike voxel clustering, where the definition of a voxel cluster is unambiguous, 'clusters' formed among supra-threshold connections can be defined in different ways, which has given rise to distinct methods.

Spatial pairwise clustering

SPC is described here for situations where nodes can be represented as discrete points in three-dimensional space. Let x_i denote the coordinate of node i. Two nodes x_i and x_j are considered neighbors if the distance by which they are separated is less than a predefined threshold, d. Furthermore, if x_i is a neighbor of x_j and y_i is a neighbor of y_j, then the two connections (x_i, y_i) and (x_j, y_j) are said to be pairwise neighbors and form a pairwise cluster. Pairwise clusters can comprise more than two connections. If x_k is a neighbor of x_i or x_j and y_k is a neighbor of y_i or y_j, then the three connections (x_i, y_i), (x_j, y_j) and (x_k, y_k) form a pairwise cluster.

Pairwise clusters provide evidence of an experimental effect. In the absence of a localized effect, supra-threshold connections are likely to be randomly distributed (see Fig. 2). SPC operates by systematically searching over all supra-threshold connections for pairwise clusters.

This search is performed by initializing an $N(N-1)/2 \times N(N-1)/2$ adjacency matrix, denoted $\mathbf{A}$, where each row/column corresponds to a unique connection, and thus each element of $\mathbf{A}$ corresponds to a unique pair of connections. Assume an ordering where row/column i corresponds to connection (x_i, y_i) and row/column j corresponds to connection (x_j, y_j). Set $\mathbf{A}(i, j) = \mathbf{A}(j, i) = 1$ if and only if (x_i, y_i) and (x_j, y_j) are pairwise neighbors and both reside in the set of supra-threshold connections; otherwise set $\mathbf{A}(i, j) = \mathbf{A}(j, i) = 0$. Specifically, $\mathbf{A}(i, j) = \mathbf{A}(j, i) = 1$ if $||x_i - x_j|| < d$ and $||y_i - y_j|| < d$, or equivalently, $||x_i - y_j|| < d$ and $||y_i - x_j|| < d$. These two conditions are equivalent due to symmetry. In this way, each unity element in $\mathbf{A}$ represents pairwise neighbors. A breadth first search can then be used to identify any connected components in $\mathbf{A}$, where each connected component defines a unique pairwise cluster. An efficient algorithm for computing connected components is provided as part of the Boost Graph Library (Gleich, 2009).

The size of a pairwise cluster can be defined as the number of pairwise relations it comprises, or the sum of the test statistic over all connections comprising the pairwise cluster. Drawing an analogy with voxel clustering, the number of pairwise relations is loosely equivalent to the notion of cluster extent, while the sum of the test statistic over all connections corresponds to cluster intensity. A weighted combination of these two measures is also possible (Hayasaka and Nichols, 2004).

The data is permuted K times using appropriate exchangeability criteria. The size of the largest pairwise cluster, denoted $s(k)$, is recorded for each permutation k. This results in an empirical null distribution for the largest pairwise cluster. A family-wise error corrected p-value for an observed pairwise cluster of size sz is then given by the total number of permutations for which the largest pairwise cluster is greater than or equal to sz divided by the total number of permutations considered.

SPC is described in pseudo-code in Algorithm 1 for the case of a one-sided test. For a two-sided test, $t(i, j)$ at line 4 is replaced with its absolute value, $|t(i, j)|$.

SPC has been used to elucidate between-group connectivity differences in a case-control study of schizophrenia (Zalesky et al., in press) as well as to localize connections modulated by the performance of an audiovisual task (Hipp et al., 2011). In the latter study, a space-, time- and frequency-resolved measure of connectivity was used (source-level coherence inferred from electroencephalography), which enabled localization of connectivity differences not only in space, but also in the dimensions of time and frequency (Maris and Oostenveld, 2007).

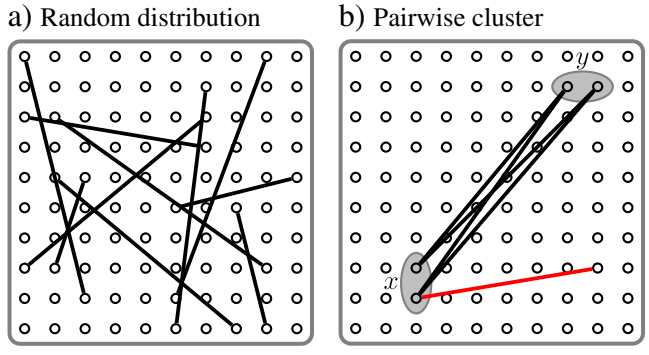

Fig. 2. Each line represents a supra-threshold connection. Circles represent nodes. Supra-threshold connections can be randomly distributed (a) or arranged into a pairwise cluster (b). The pairwise cluster provides evidence for an experimental effect between the pair of distant regions shaded gray. The node pair colored red is not part of the pairwise cluster because only one of its nodes is a neighbor.

Algorithm 1: Spatial pairwise clustering

Inputs: x_i = coordinate of node i; T = cluster-forming threshold; d = neighbor distance threshold; K = number of permutations

1: Measure connectivity between every pair of nodes (i, j)
2: Fit univariate model to each connection (i, j) and compute test statistic $t(i, j)$ to assess evidence for the contrast of interest
3: **for** all possible pairs of node pairs $\{(x_i, y_i), (x_j, y_j)\}$ **do**
4: **if** $||x_i - x_j|| < d$ **and** $||y_i - y_j|| < d$ **and** $t(i, j) > T$ **then**
5: $\mathbf{A}_{i,j} := 1, \mathbf{A}_{j,i} := 1$
6: **else**
7: $\mathbf{A}_{i,j} := 0, \mathbf{A}_{j,i} := 0$
8: **end if**
9: **end for**
10: Conduct breadth first search for connected components in $\mathbf{A}$ and let sz_n be the size of component $n = 1, \ldots, M$
11: Generate K samples, $s(k)$, from the null distribution of the largest pairwise cluster via permutation testing
12: Compute p-value for component sz_n as $p_n = \#\{s(k) \geq \text{sz}_n\}/K$
13: **return** p_n, $n = 1, \ldots, M$

Network-based statistic

The NBS is a more established algorithm. It has been used to map functional (e.g. Fornito et al., 2011a; Zhang et al., 2011) and anatomical (e.g. Verstraete et al., 2011; Zalesky et al., 2011) connectivity disturbances in psychiatric and neurological disorders.

The NBS differs from SPC in the criteria used to define clusters. With SPC, if (x_i, y_i) and (x_j, y_j) are two supra-threshold connections, they are clustered together if and only if x_i is a neighbor of x_j **and** y_i is a neighbor y_j. With the NBS on the other hand, they are clustered together if x_i is a neighbor of x_j **or** y_i is a neighbor y_j.

While this modification might seem rather inconsequential, it can lead to a substantial improvement or worsening in statistical power (see Fig. 3). Whether or not the modification leads to an improvement depends on many factors (see Discussion). However, in general, the NBS is suited to an effect spanning multiple interconnected regions, whereas SPC is suited to an effect between an isolated pair of regions.

The following three modifications to Algorithm 1 give rise to the NBS: the adjacency matrix $\mathbf{A}$ is now of size $N \times N$; at line 3, the for-loop now only needs to traverse all possible connections, not all possible

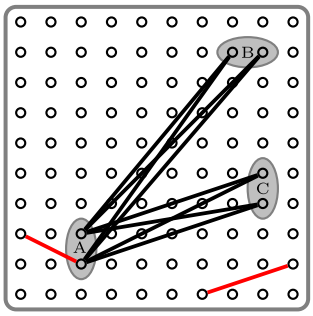

Three interconnected regions:
In this case, SPC identifies two separate pairwise clusters, A-B and A-C, whereas the NBS identifies one cluster corresponding to the entire network A-B-C. As such, the NBS only enables rejection of the null hypothesis at the level of the entire network A-B-C, whereas SPC enables rejection of the null hypothesis for A-B and A-C individually. The drawback of the NBS is that it detects the connection originating from region A, colored red, which can represent a false positive.

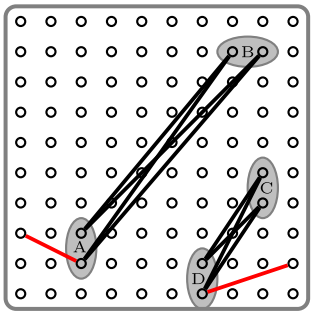

Two pairs of interconnected regions:
In this case, the NBS and SPC both identify two separate clusters corresponding to A-B and C-D. The null hypothesis can be rejected individually for A-B and A-C with both methods. The drawback of the NBS is that it detects the connections originating from regions A and D, colored red, which could be false positives. These false positives are disregarded by SPC because neither of them form *pairwise* neighbors.

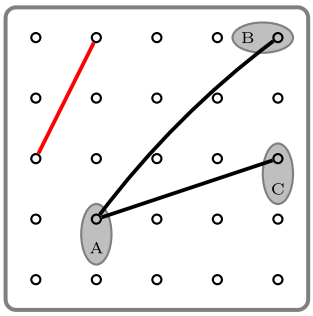

Coarser sampling resolution:
In this case, SPC does not identify any effect due to the lack of any pairwise relations; however, the NBS correctly identifies the cluster corresponding to the network A-B-C.

Fig. 3. Three examples explicating key differences between the NBS and SPC. Each line represents a supra-threshold connection. Circles represent nodes. Connections colored black correspond to an experimental effect (true positives), while those colored red correspond to false positives that have survived the cluster-forming threshold.

pairs of connections; and, at line 4, the if-loop condition now becomes $||x_i - x_j|| < d$ **or** $||y_i - y_j|| < d$ **and** $t(i,j) > T$.

Application

Electroencephalographic activity was recorded at 160 scalp electrodes in 11 healthy volunteers. Each volunteer performed a visual proximity task under different working memory load conditions: no load, low load and high load (Cocchi et al., 2011). Data were band-pass filtered from 0.1 to 40 Hz. Sources were estimated using a distributed linear inverse solution (ELECTRA) applying the local autoregressive average (LAURA) regularization approach to address the non-uniqueness of the inverse problem (Grave de Peralta Menendez et al., 2005). The space of the inverse solution was sampled with $N = 280$ discrete nodes uniformly distributed across the cortical surface.

Using Pearson's correlation coefficient, connectivity (zero lag amplitude covariation) was measured between every pair of nodes separated by a distance of at least 40 mm, resulting in a total of 10,861 connections. Node pairs separated by less than 40 mm were excluded to avoid spurious short-range correlations arising from volume conduction effects.

For each connection, a repeated measures analysis of variance was computed to test for a change in the strength of connectivity across the three conditions in the 160–190 ms post-stimulus period. This specific post-stimulus period was chosen based on a previous analysis (Cocchi et al., 2011). Each of the 10,861 connections was therefore endowed with an *F*-statistic and corresponding *p*-value.

The most significant of the 10,861 *p*-values was 0.00026. Because $\alpha/10861 \approx 10^{-6}$, $\alpha = 0.05$, is two orders of magnitude smaller than 0.00026, controlling the family-wise error rate with the Holm–Bonferroni method yielded no significant findings. Control of the false discovery rate also yielded no significant findings.

The NBS and SPC were separately applied for each of three different *F*-statistic cluster-forming thresholds: $T = 6.5, 7$ and 8. For the purposes of SPC, any two nodes were defined as neighbors if they were separated by no more than $d = 7$ mm. Based on this distance threshold, each node possessed 5.8 neighbors on average. The size of a pairwise cluster was defined as the number of pairwise relations it comprised.

A total of $K = 5000$ permutations were computed for each threshold. For each permutation, the three conditions were randomized using appropriate exchangeability criteria for a repeated measures analysis of variance (Suckling and Bullmore, 2004).

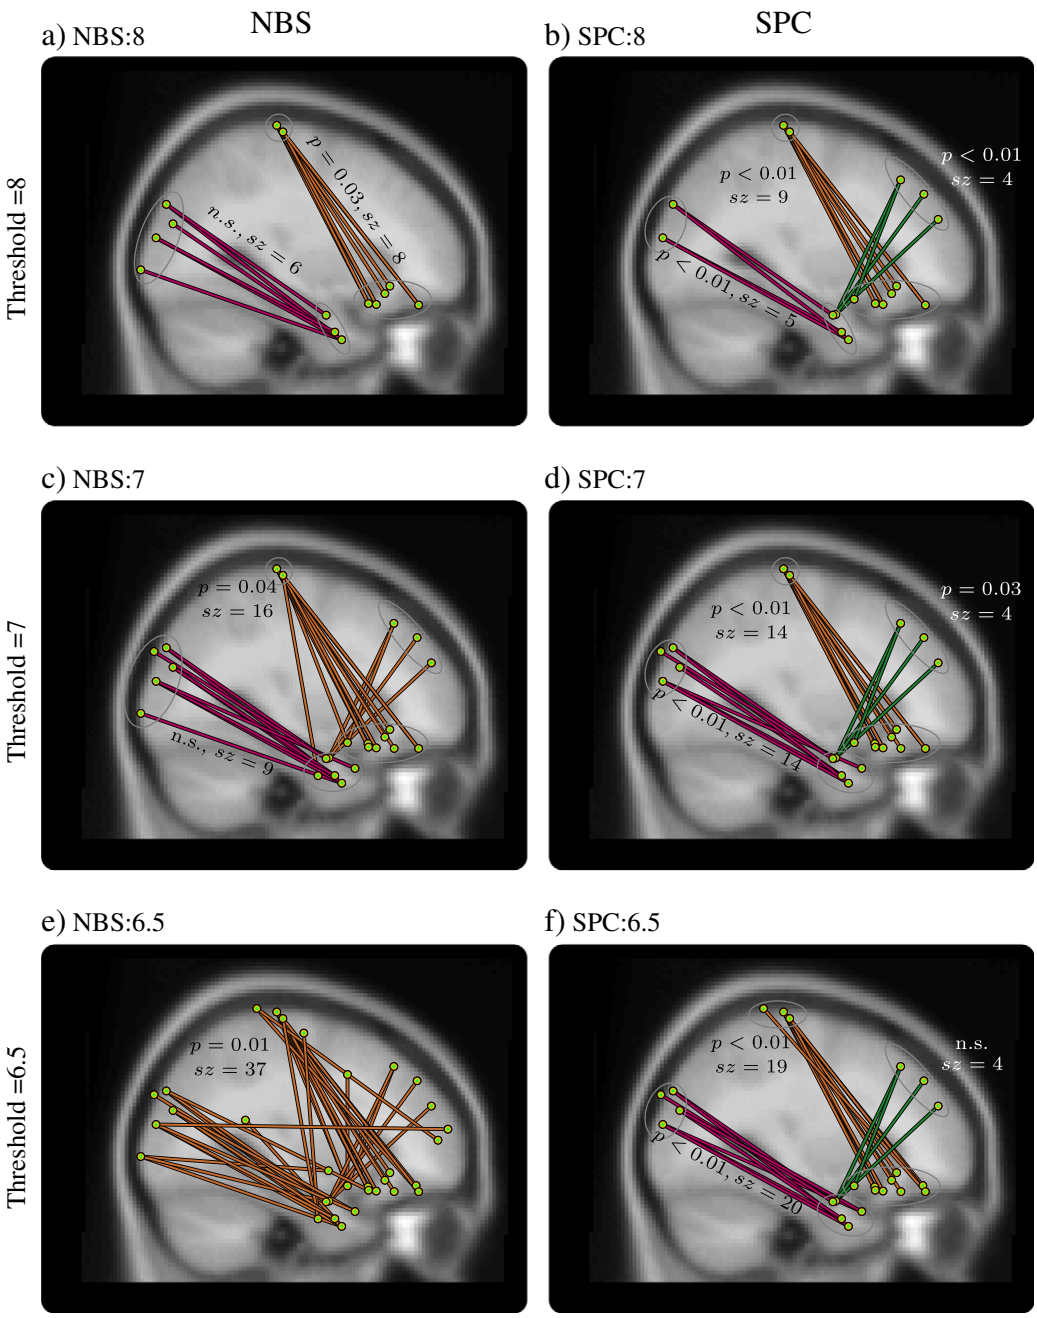

Fig. 4. The NBS and SPC were used to identify regions between which connectivity changed as a function of the working memory load. Each slice is a planar (sagittal) representation. The threshold refers to the *F*-statistic cluster-forming threshold. The connections comprising each distinct cluster are colored uniquely. The size (sz) and family-wise error corrected *p*-value of each cluster is indicated, where n.s. indicates the cluster is not significant ($p > 0.05$). The size of a pairwise cluster is the number of pairwise neighbors it comprises.

Fig. 4 shows the regions identified by the NBS and SPC showing evidence of an experimental effect. SPC identified the same three pairwise clusters for each of the three thresholds evaluated. The pairwise clusters are colored orange, green and red in Fig. 4. Each pairwise cluster represents a pair of brain regions between which connectivity changed as a function of the working memory load. The null hypothesis was rejected for each pairwise cluster individually ($p<0.05$), except for the threshold of 6.5, where the green pairwise cluster was no longer significant. When the threshold was reduced to 6.5, the size of the green pairwise cluster remained unchanged ($sz=4$), but the likelihood of a larger pairwise cluster emerging as a matter of chance in the randomized data naturally increased with a lower threshold; hence the loss of significance.

The NBS is better suited to an effect spanning multiple interconnected regions. This is exemplified in Fig. 4(c), where the green and orange pairwise clusters have been interconnected to form a single network, as well as in Fig. 4(e), where the green, orange and red pairwise clusters have been interconnected to form a single all-encompassing cluster. This is to the detriment of localizing power because the null hypothesis can only be rejected for the cluster as a whole in Fig. 4(e), but not individually for the three constituent pairwise clusters. In contrast, SPC preserves the distinction between the three pairwise clusters.

It is important to remark that in some cases, but not here, the pairwise clusters comprising a network might not be statistically significant individually, but the network they comprise may be significant as a whole. The NBS is clearly advantageous in these cases.

Discussion

Whereas supra-threshold voxel clustering (Bullmore et al., 1999; Nichols and Holmes, 2001) pertains to mass univariate testing of brain activation, the network-based statistic (NBS) and spatial pairwise clustering (SPC) are analogous methods enabling mass univariate testing of brain connectivity. More specifically, the NBS and SPC are clustering strategies to control family-wise errors, in the weak sense, when assessing the evidence for an experimental effect at each of the typically thousands of connections comprising a brain network.

A key advantage of the NBS and SPC is that they enable testing of an experimental effect at every connection comprising a brain network, not just those connections between a few select regions-of-interest. The region-of-interest approach is sensible if an experimenter wishes to test a single hypothesis relating to a specific connection, but in the case of exploratory testing, it may introduce biases stemming from the region-of-interest selection process (Hipp et al., 2011).

Another key advantage of the NBS and SPC is that they are both nonparametric. This means assumptions are not required on the probability distribution of connectivity measurements under the null hypothesis, or the distribution of the measure used to quantify the extent of an experimental effect. For the case study considered in this article, the *F*-statistic was used to quantify the extent to which connectivity changed as a function of the working memory load. However, any measure of variation would have been just as valid, even if it was nonstatistical in origin. In contrast, with parametric testing, an experimenter is limited to tests that conform to the sampling distribution (Maris and Oostenveld, 2007).

Threshold selection

Both the NBS and SPC require selection of a cluster-forming threshold to define a set of supra-threshold connections. The choice of threshold is rather arbitrary. While the arbitrariness of threshold selection is a shortcoming of supra-threshold clustering methods in general, it is important to remark that the choice of threshold does not affect specificity. In particular, weak control of the family-wise error rate is ensured irrespective of the choice of threshold.

The concern is that the choice of threshold can have a bearing on sensitivity. For example, consider an experimental effect comprising relatively many connections, but each of which shows a relatively weak effect. In this case, if the threshold chosen exceeds the cutoff value corresponding to this weak effect, both the NBS and SPC will fail to detect the effect, resulting in a false negative, even though the effect is 'significant' on the basis of its large spatial extent (i.e. it comprises relatively many connections).

This is analogous to the arbitrariness of threshold selection in voxel clustering, which has been alleviated to a certain extent with the threshold-free cluster enhancement approach (Smith and Nichols, 2009).

Fig. 4 shows that relatively small variations in the *F*-statistic threshold ($F=6.5$, 7 or 8) can have a bearing on the results. Most notably, the cluster colored green was no longer detected with the NBS when the threshold was increased from $F=7$ to $F=8$, while with SPC, the cluster colored green was no longer significant ($p>0.05$) when the threshold was reduced from $F=7$ to $F=6.5$. The latter case demonstrates an important point—reducing the cluster-forming threshold does not necessarily result in improved sensitivity.

NBS versus SPC

In general, the NBS is suited to an experimental effect spanning multiple interconnected regions (i.e. a network), whereas SPC is suited to an effect between an isolated pair of regions (i.e. a connection). This distinction between the two methods stems for a subtle, yet consequential difference in the criteria used to define clusters among the set of supra-threshold connections.

With the NBS, a supra-threshold connection can join a cluster if at least one of its nodes is a neighbor of one of the nodes already comprising the cluster. This is the minimum requirement necessary to form a cluster of interconnected connections, which is referred to as a connected component in the parlance of graphs. In contrast, SPC uses a more stringent pairwise clustering concept. With SPC, a supra-threshold connection can join a cluster if and only if both of its nodes are neighbors of nodes already comprising the cluster.

An advantage of SPC is that effects can be declared at a finer resolution. In particular, the null hypothesis can be rejected separately at the level of individual pairs of brain regions, rather than at the level of a network of many brain regions. This is most clearly demonstrated in Fig. 4 for the case of $F=6.5$. In this case, SPC identified three separate clusters, each corresponding to a different pair of brain regions between which connectivity changed as a function of the working memory load (Fig. 4f). The null hypothesis could therefore be rejected/accepted separately for each of the three clusters. In contrast, the NBS identified a single cluster representing a network that encompassed the three smaller clusters identified with SPC (Fig. 4e).

It is important to note that the finer localizing resolution of SPC can however be to the detriment of sensitivity. Specifically, a pairwise cluster in itself may be too small to reach statistical significance with SPC. However if that pairwise cluster forms a network with other pairwise clusters, the network may be large enough to declare significance at the network level with the NBS. The case of $F=6.5$ in Fig. 4 provides an example of such a scenario, where the pairwise cluster colored green failed to reach significance with SPC (Fig. 4e), but was significant at the network level with the NBS (Fig. 4f).

Lower computational complexity is an advantage of the NBS. With the NBS, the adjacency matrix used to perform the cluster search is of size $N\times N$, whereas with SPC, the size grows substantially to $N(N-1)/2\times N(N-1)/2$. Furthermore, the number of neighborhood relations that must be enumerated is much greater when clusters are defined in a pairwise manner. For example, a voxel has 26 immediate neighboring voxels, but a pair of voxels has $(26+1)\times(26+1)=729$ neighboring pairs of voxels. The $+1$ here represents the original voxel pair.

The NBS may be advantageous if a coarse parcellation template is used subdivide the brain into distinct nodes. In this case, it is more likely that regions corresponding to an experimental effect are sampled by only one node, and thus pairwise clusters cannot be formed (e.g. see third case, Fig. 3). In contrast, the NBS may be able to declare significance based on the presence of a network, since network structure does not require each region to be sampled by multiple nodes.

In many scenarios, an experimenter may not have a strong hypothesis about the spatial characteristics of an experimental effect. Therefore, the choice between the NBS and SPC is likely to be

Table 1
Key features distinguishing the NBS and SPC.

	NBS	SPC
Complexity	Lower	Higher
Experimental effect	Network (Coarser)	Connection (Finer)
Specificity	Lower	Higher
Sensitivity	Higher	Lower

guided by a tradeoff between factors such as specificity and sensitivity, computational complexity and the desired resolution at which an effect is to be localized. In particular, as summarized in Table 1, SPC usually provides greater specificity, whereas the NBS usually provides greater sensitivity; the NBS is always less computationally burdensome than SPC; and SPC enables localization at the higher resolution of individual connections, whereas the NBS is suited to localizing networks.

Cognitive interpretation

The case study considered in this paper was designed to investigate brain dynamics involved in concurrent execution of a visual–perceptual task while maintaining an unrelated visual–spatial working memory (VSWM) load. Full details of the task paradigm have been presented elsewhere (Cocchi et al., 2011). Concurrent execution of two unrelated tasks has been found to result in either enhanced or decreased overall performance compared to the case of serial execution (Kim et al., 2005; Rissman et al., 2009). Elucidating the dynamics of functional connectivity associated with this novel effect provides ideal motivation for the application of methods such as the NBS and SPC.

Three distinct connections were modulated by VSWM load: occipital–temporal, frontal–temporal and dorsofrontal–orbito-frontal (Fig. 4). Our findings support the hypothesis that neural networks encompassing frontal, temporal and occipitoparietal regions are central in managing dual-task demands. Specifically, our current results are in line with previous studies suggesting that frontal regions play a central role in linking sensory-perceptual and motor operations in dual task contexts (Dux et al., 2006; Sigman and Dehaene, 2008). Indeed, both the superior- and inferolateral–frontal regions identified were connected with regions involved in visual–perceptual and motor processes. Our findings are also in line with theoretical models implicating frontal/prefrontal regions in cognitive control mechanisms integrating sensory and motor processes (Badre, 2008).

There is significant scope to further analyze our findings in the context of dual-task performance. As part of a post hoc analysis, it might be useful to determine whether connectivity increased/decreased as a function of VSWM load for each of the three connections identified. Due to the unsigned nature of the *F*-statistic, a pairwise cluster can potentially comprise a mixture of nodes pairs—some between which connectivity increases as function of VSWM load, while others between which it decreases. Therefore, it might also be advantageous to replace the *F*-statistic with a signed measure of variability that is sensitive to only an increase or only a decrease in connectivity. Another line of investigation might be to use a time- and frequency-resolved measure of connectivity, such as source-localized coherence.

Acknowledgments

We are very grateful to Professor Olaf Sporns and Dr Stephan Gerhard for adding the NBS to the Brain Connectivity Toolbox and the Connectome Mapping Toolkit, respectively.

A.Z. is grateful for the support provided by Professor Trevor Kilpatrick as part of the inaugural Melbourne Neuroscience Institute Fellowship. This work was also supported by the Australian Research Council [DP0986320 to A.Z.], the Melbourne Neuroscience Institute, the National Health and Medical Research Council of Australia [C.J. Martin Fellowship to A.F.], and the Swiss National Science Foundation [grant 310030B_133136 to M.M.].

References

Achard, S., Salvador, R., Whitcher, B., Suckling, J., Bullmore, E., 2006. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. J. Neurosci. 26 (1), 63–72.

Badre, D., 2008. Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. Trends Cogn. Sci. 12, 193–200.

Bassett, D.S., Bullmore, E., 2006. Small-world brain networks. Neuroscientist 12 (6), 512–523.

Bassett, D.S., Bullmore, E., Verchinski, B.A., Mattay, V.S., Weinberger, D.R., Meyer-Lindenberg, A., 2008. Hierarchical organization of human cortical networks in health and schizophrenia. J. Neurosci. 28 (37), 9239–9248.

Bassett, D.S., Brown, J.A., Deshpande, V., Carlson, J.M., Grafton, S.T., 2010. Conserved and variable architecture of human white matter connectivity. Neuroimage 54 (2), 1262–1279.

Bassett, D.S., Wymbs, N.F., Porter, M.A., Mucha, P.J., Carlson, J.M., Grafton, S.T., 2011. Dynamic reconfiguration of human brain networks during learning. Proc. Natl. Acad. Sci. U. S. A. 108, 7641–7646.

Bressler, S.L., Menon, V., 2010. Large-scale brain networks in cognition: emerging methods and principles. Trends Cogn. Sci. 14 (6), 277–290.

Bullmore, E.T., Sporns, O., 2009. Complex brain networks: graph-theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. 10, 186–198.

Bullmore, E.T., Suckling, J., Overmeyer, S., Rabe-Hesketh, S., Taylor, E., Brammer, M.J., 1999. Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural MR images of the brain. IEEE Trans. Med. Imaging 18 (1), 32–42.

Chen, G., Chen, G., Xie, C., Shi-Jiang, L., 2011. Negative functional connectivity and its dependence on the shortest path length of positive network in the resting-state human brain. Brain Connect. 1 (3), 195–206.

Cocchi, L., Toepel, U., De Lucia, M., Martuzzi, R., Wood, S.J., Carter, O., Murray, M.M., 2011. Working memory load improves early stages of independent visual processing. Neuropsychologia 49 (1), 92–102.

Dux, P.E., Ivanoff, J., Asplund, C.L., Marois, R., 2006. Isolation of a central bottleneck of information processing with time-resolved FMRI. Neuron 52, 1109–1120.

Fair, D.A., Cohen, A.L., Power, J.D., Dosenbach, N.U.F., Church, J.A., Miezin, F.M., Schlagger, B.L., Petersen, S.E., 2009. Functional brain networks develop from a "local to distributed" organization. PLoS Comput. Biol. 5 (5), e1000381.

Fornito, A., Yoon, J., Zalesky, A., Bullmore, E.T., Carter, C.S., 2011a. General and specific functional connectivity disturbances in first episode schizophrenia during cognitive control performance. Biol. Psychiatry 70 (1), 64–72.

Fornito, A., Zalesky, A., Bassett, D., Meunier, D., Ellison-Wright, I., Yucel, M., Wood, S.J., Shaw, K., O'Connor, J., Nertney, D., Mowry, B., Pantelis, C., Bullmore, E., 2011b. Genetic influences on cost-efficient organization of human cortical functional networks. J. Neurosci. 31 (9), 3261–3270.

Friston, K.J., 2009. Modalities, modes, and models in functional neuroimaging. Science 326 (5951), 399–403.

Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage 15, 870–878.

Ginestet, C.E., Nichols, T.E., Bullmore, E.T., Simmons, A., 2011. Brain network analysis: separating cost from topology using cost-integration. PLoS One 6 (7), e21570.

Gleich, D.F., 2009. Models and algorithms for page rank sensitivity. PhD Thesis. Stanford University.

Grave de Peralta Menendez, R., Murray, M.M., Michel, C.M., Martuzzi, R., Gonzalez Andino, S.L., 2005. Electrical neuroimaging based on biophysical constraints. Neuroimage 21, 527–539.

Habeck, C., Moeller, J.R., 2011. Intrinsic functional-connectivity networks for diagnosis: just beatiful pictures? Brain Connect. 1 (2), 99–103.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C.J., Wedeen, V.J., Sporns, O., 2008. Mapping the structural core of human cerebral cortex. PLoS Biol. 6 (7), e159.

Hayasaka, S., Laurienti, P.J., 2010. Comparison of characteristics between region- and voxel-based network analyses in resting-state fMRI data. Neuroimage 50 (2), 499–508.

Hayasaka, S., Nichols, T.E., 2004. Combining voxel intensity and cluster extent with permutation test framework. Neuroimage 23, 54–63.

He, Y., Evans, A., 2010. Graph theoretical modeling of brain connectivity. Curr. Opin. Neurol. 23 (4), 341–350.

He, Y., Chen, Z.J., Evans, A.C., 2007. Small-world anatomical networks in the human brain revealed by cortical thickness from MRI. Cereb. Cortex 17 (10), 2407–2419.

He, Y., Chen, Z., Evans, A., 2008. Structural insights into aberrant topological patterns of large-scale cortical networks in Alzheimer's disease. J. Neurosci. 28, 4756–4766.

Hipp, J.F., Engel, A.K., Siegel, M., 2011. Oscillatory synchronization in large-scale cortical networks predicts perception. Neuron 69, 387–396.

Holm, S., 1979. A simple sequentially rejective multiple test procedure. Scand. J. Stat. 6 (2), 65–70.

Kaiser, M., 2011. A tutorial in connectome analysis: topological and spatial features of brain networks. Neuroimage 57, 892–907.

Kim, S.Y., Kim, M.S., Chun, M.M., 2005. Concurrent working memory load can reduce distraction. Proc. Natl. Acad. Sci. U. S. A. 102, 16524–16529.

Kitzbichler, M.G., Henson, R.N., Smith, M.L., Nathan, P.J., Bullmore, E.T., 2011. Cognitive effort drives workspace configuration of human brain functional networks. J. Neurosci. 31 (22), 8259–8270.

Li, L., Rilling, J.K., Preuss, T.M., Glasser, M.F., Hu, X., (in press). The effects of connection reconstruction method on the interregional connectivity of brain networks via diffusion tractography. Hum Brain Mapp. doi:10.1002/hbm.21332.

Lynall, M.-E., Bassett, D.S., Kerwin, R., McKenna, P.J., Kitzbichler, M., Muller, U., Bullmore, E., 2011. Functional connectivity and brain networks in schizophrenia. J. Neurosci. 30 (28), 9477–9487.

Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing EEG- and MEG-data. J. Neurosci. Methods 164, 177–190.

Nichols, T.E., Hayasaka, S., 2003. Controlling the familywise error rate in functional neuroimaging: a comparative review. Stat. Methods Med. Res. 12 (5), 419–446.

Nichols, T.E., Holmes, A.P., 2001. Nonparametric permutation tests for functional neuroimaging: a primer with examples. Hum. Brain Mapp. 15, 1–25.

Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., Hallett, M., 2004. Identifying true brain interaction from EEG data using the imaginary part of coherency. Clin. Neurophysiol. 115, 2292–2307.
Onnela, J.P., Saramaki, J., Kertesz, J., Kaski, K., 2005. Intensity and coherence of motifs in weighted complex networks. Phys. Rev. E 71, 065103.
Rissman, J., Gazzaley, A., D'Esposito, M., 2009. The effect of non-visual working memory load on top-down modulation of visual processing. Neuropsychologia 47 (7), 1637–1646.
Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretation. Neuroimage 52 (3), 1059–1069.
Rubinov, M., Sporns, O., 2011. Weight-conserving characterization of complex functional brain networks. Neuroimage 56 (4), 2068–2079.
Schoffelen, J.M., Gross, J., 2009. Source connectivity analysis with MEG and EEG. Hum. Brain Mapp. 30, 1857–1865.
Schoffelen, J.M., Gross, J., 2011. Improving the interpretability of all-to-all pairwise source connectivity analysis in MEG with nonhomogeneous smoothing. Hum. Brain Mapp. 2 (3), 426–437.
Sigman, M., Dehaene, S., 2008. Brain mechanisms of serial and parallel processing during dual-task performance. J. Neurosci. 28, 7585–7598.
Smith, S.M., Nichols, T.E., 2009. Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. Neuroimage 44 (1), 83–98.
Sporns, O., 2011. The human connectome: a complex network. Ann. N. Y. Acad. Sci. 1224, 109–125.
Sporns, O., (in press). From simple graphs to the connectome: networks in neuroimaging. Neuroimage. doi:10.1016/j.neuroimage.2011.08.085.
Suckling, J., Bullmore, E., 2004. Permutation tests for factorially designed neuroimaging experiments. Hum. Brain Mapp. 22, 193–205.
van den Heuvel, M.P., Hulshoff Pol, H.E., 2010. Exploring the brain network: a review on resting-state fMRI functional connectivity. Eur. Neuropsychopharmacol. 20 (8), 519–534.
van den Heuvel, M.P., Stam, C.J., Boersma, M., Hulshoff Pol, H.E., 2008. Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. Neuroimage 3 (3), 528–539.
van den Heuvel, M.P., Mandl, R.C., Stam, C.J., Kahn, R.S., Hulshoff Pol, H.E., 2010. Aberrant frontal and temporal complex network structure in schizophrenia: a graph theoretical analysis. J. Neurosci. 30 (47), 15915–15926.
van Wijk, B.C.M., Stam, C.J., Daffertshofer, A., 2010. Comparing brain networks of different size and connectivity density using graph theory. PLoS One 5 (10), e13701.
Verstraete, E., Veldink, J.H., Mandl, R.C.W., van den Berg, L.H., van den Heuvel, M.P., 2011. Impaired structural motor connectome in amyotrophic lateral sclerosis. PLoS One 6, e24239.
Wang, J., Wang, L., Zang, Y., Yang, H., Tang, H., Gong, Q., Chen, Z., Zhu, C., He, Y., 2009. Parcellation-dependent small-world brain functional networks: a resting-state fMRI study. Hum. Brain Mapp. 30, 1511–1523.
Wang, J., Zuo, X., He, Y., 2010. Graph-based network analysis of resting-state functional MRI. Front. Syst. Neurosci. 4, 16.
Wig, G.S., Schlaggar, B.L., Petersen, S.E., 2011. Concepts and principles in the analysis of brain networks. Ann. N. Y. Acad. Sci. 1224 (1), 126–146.
Zalesky, A., Fornito, A., Bullmore, E.T., 2010a. Network-based statistic: identifying differences in brain networks. Neuroimage 53 (4), 1197–1207.
Zalesky, A., Fornito, A., Harding, I.H., Cocchi, L., Yucel, M., Pantelis, C., Bullmore, E.T., 2010b. Whole-brain anatomical networks: does the choice of nodes matter? Neuroimage 50 (3), 970–983.
Zalesky, A., Fornito, A., Seal, M.L., Cocchi, L., Westin, C.-F., Bullmore, E.T., Egan, G.F., Pantelis, C., 2011. Disrupted axonal fiber connectivity in schizophrenia. Biol. Psychiatry 69, 80–89.
Zalesky, A., Fornito, A., Egan, G.F., Pantelis, C., Bullmore, E.T., (in press). The relationship between regional and inter-regional functional connectivity deficits in schizophrenia. Hum Brain Mapp. doi:10.1002/hbm.21379.
Zhang, J., Wang, J., Wu, Q., Kuang, W., Huang, X., He, Y., Gong, Q., 2011. Disrupted brain connectivity networks in drug-naive, first-episode major depressive disorder. Biol. Psychiatry 70, 334–342.