

Western University

Scholarship@Western

Brain and Mind Institute Researchers'
Publications

Brain and Mind Institute

1-2014

An Invitation to the Study of Brain Networks, with Some Statistical Analysis of Thresholding Techniques

Mark Daley

The University of Western Ontario, London, ON, Canada, mdaley2@uwo.ca

Follow this and additional works at: <https://ir.lib.uwo.ca/brainpub>



Part of the [Neurosciences Commons](#), and the [Psychology Commons](#)

Citation of this paper:

Daley, Mark, "An Invitation to the Study of Brain Networks, with Some Statistical Analysis of Thresholding Techniques" (2014). *Brain and Mind Institute Researchers' Publications*. 148.

<https://ir.lib.uwo.ca/brainpub/148>

An Invitation to the Study of Brain Networks, with Some Statistical Analysis of Thresholding Techniques

Mark Daley

Abstract We provide a brief introduction to the nascent application of network theory to mesoscale networks in the human brain. Following an overview of the typical data-gathering, processing, and analysis methods employed in this field, we describe the process for inferring a graph from neural time series. A crucial step in the construction of a graph from time series is the thresholding of graph edges to ensure that the graphs represent physiological relationships rather than artifactual noise. We discuss the most popular currently employed methodologies and then introduce one of our own, based on the theory of random matrices. Finally, we provide a comparison of our random-matrix-theory thresholding approach with two dominant approaches on a data set of 1,000 real resting-state functional magnetic resonance imaging scans.

1 Introduction

In an age where humans are able, for the first time, to collect huge quantities of data on the relationships between objects, network/graph theory has become a rising star as a tool to help make sense of raw relational data. Nowhere is this more apparent than in the field of neuroimaging, which has begun to embrace graph theory as a tool for understanding relationships between brain areas at the mesoscale. In this field, one does not study individual neurons and traditional “neural networks”, as the imaging modalities available to us have a resolution which is far too coarse for this application; instead, one studies relationships between medium-scale (mesoscale) assemblies of tens of thousands, or more, neurons.

M. Daley (✉)
The University of Western Ontario, London, ON, Canada
e-mail: mdaley2@uwo.ca

Viewing this work from the standpoint of neuroscience, the motivations are immediate: if large, four-dimensional, data sets can be reduced to single static graphs – the topology of which varies according to the properties of interest – then the complexity of the analysis can be greatly reduced. For example, a developmental neuroscientist may hypothesize that global connectivity varies in the brain as we age from children to adults. Given a method for converting functional neuroimaging data into graphs, one now has an immediate way to test this hypothesis: scan the brains of children and adults, convert the scans to graphs, and compute a topological metric for each graph that captures the intuition of “global connectivity”; indeed, this precise study has been performed [9]. Likewise, one can imagine a clinician interested in whether or not there are metrics which can be applied to such graphs that would distinguish clinical conditions such as Alzheimer’s disease [7] or schizophrenia [17]. A broad introduction to the neuroscientific implications of this line of research can be found in [21].

Viewing this work as a mathematician interested in graph theory, one sees the opportunity to study very special classes of graphs and to investigate, for example, new metrics better suited to quantifying the topological differences between graphs in this class. As a computability theorist, one sees the opportunity to study snapshots of an intriguing computing machine (the brain) within the familiar framework of graph theory.

This chapter is meant to serve as a very brief introduction to the increasingly popular application of graph-theoretic techniques to data derived from neuroimaging experiments. In the background section, we will quickly review the fundamental concepts and procedures employed to move from neural activity in the brain, through several processing steps, to a static graph representation of the same activity. We note that the ability to threshold signal from noise in our final graph representation is a critical step in building physiologically credible graphs and, after reviewing the currently popular techniques for thresholding, we suggest a new approach based on the theory of random matrices. Finally, we conclude with a summary and discussion of the many classes of open problem in this field, both pragmatic and abstract.

2 Background

We now bring together the necessary background to understand how to map mesoscale neural time series, derived from neuroimaging, to static graphs. We begin with a short introduction to magnetic resonance imaging (MRI) and functional magnetic resonance imaging (fMRI) and clarify our use of graph-theoretic concepts and terminology. We discuss the specifics of the somewhat involved process of preprocessing such data to remove artifacts and make it suitable for analysis. With suitably preprocessed data, the next step towards constructing a graph involves the pairwise comparison of time series. We review three dominant approaches in current use: one linear time-domain measure (correlation), one linear spectral measure

(coherence), and one nonlinear measure (mutual information). Finally, we discuss the construction and analysis of a graph-theoretic representation of neuroimaging data.

2.1 MRI, fMRI, the Resting State, and Functional Connectivity

The data with which we will build our graphs is most commonly derived from resting-state functional magnetic resonance imaging (though similar approaches exist for other neuroimaging modalities). Magnetic resonance imaging leverages nuclear magnetic resonance in a carefully controlled fashion to generate volumes containing information about spin-relaxation properties in particular voxels.¹

We now give a very brief characterization of imaging via magnetic resonance. A subject is placed in a very strong magnetic field (typically 1.5–7 tesla), which causes the spin states of individual protons to become increasingly uniform, with each proton precessing around the direction of the magnetic field.² Electromagnetic radiation at the resonant frequency of the protons is then introduced, causing the protons to absorb the emitted photons and flip spin state; at the same time, this process causes a synchronization of the precession of individual protons. When the influx of photons from the electromagnetic field is removed, the spin states of the protons flip back to the lower-energy state, while the individual precessions drift increasingly out of synchronization. By measuring the photons emitted during this process, the magnetic resonance scanner can determine the length of time it takes for the spin vector of the bulk system to return to being parallel with the static magnetic field and the length of time for the individual precessions to become completely decoherent.

By superimposing a variable magnetic gradient on top of the static magnetic field, one can spatially localize this process and extract resonance contrast for a given spatial frequency in a particular location. After this k -space has been sampled densely, a simple Fourier transform yields a volume in Euclidean 3-space. By carefully choosing which properties (e.g., relaxation time, decoherence time, proton density, flow, and spectral shift) one records, and developing different pulse sequences,³ one can optimize the contrast for different tissue types. For a detailed

¹A *voxel* is simply the three-dimensional analog of a pixel: a (usually cube-shaped) volume assigned a homogenous scalar or vector value. In the context of fMRI, a typical voxel will represent a volume of size roughly 1–5 mm³, depending on several technical details of the hardware and software used for the imaging.

²Spin is a quantum property of elementary particles and, as such, may take on only a discrete number of states (or some quantum superposition thereof). In the case of fermions such as the proton, the two possible spin states are $\{\frac{1}{2}, -\frac{1}{2}\}$.

³Pulse sequences contain integrated information about the application of magnetic gradients, electromagnetic pulses, and the recording of electromagnetic radiation emitted from the subject. For more detail, see [4].

exposition of the physics of nuclear magnetic resonance imaging, we direct the reader to [11].

The contrast of interest for our purposes is the $T2^*$ contrast, which allows imaging of the BOLD (blood-oxygen-level-dependent) response, which, roughly, forms contrast from the oxygenation level of blood in a volume of tissue. Using earlier magnetic resonance techniques, it is difficult to visualize oxygenated blood versus deoxygenated blood without the use of an injectable contrast agent. The ability to image the BOLD signal without the use of exogenous contrast agents was first reported in [14, 18], with the suggestion that variation in blood flow to a particular brain region was correlated with activity in that region. Follow-up work, notably that in [16], demonstrated that this modulation of blood flow is strongly linked with underlying electrical activity, although the exact nature of the BOLD signal is still under active investigation. Interestingly, it has recently been linked to astrocytes [19], suggesting a greater than previously acknowledged functional role for glial cells. The BOLD signal is very slow, lagging behind electrical recordings of neural activation by several seconds and effectively acting as a temporal low-pass filter for the underlying neural activity. It is also spatially diffuse, as it is intrinsically linked to the vasculature surrounding the recently activated area rather than the area itself – one might view this as a spatial form of low-pass filtering. A thorough introduction to the methodologies of fMRI can be found in [12].

Traditional fMRI studies involve measuring the BOLD response in block-designed or event-related experiments and contrasting the BOLD signal between different experimental conditions. Beginning with [5], and to an increasingly extent recently, a different question has been addressed: “what can we learn about the brain by observing basal, ‘spontaneous’, fluctuations in the BOLD signal during rest?”. In this chapter, we shall describe work on resting-state data in which we record the BOLD signal from subjects lying in the scanner and given no particular task other than the instruction to relax and not think about anything in particular.

Resting-state fMRI data is typically used to build models of *functional connectivity*, in which two brain regions are said to be functionally connected if their BOLD time courses are highly similar according to some metric such as correlation. It is important to note that despite the use of the term “connectivity”, the information obtained in this type of analysis is purely correlational and cannot give any definitive information about actual anatomical connectivity. When we say that two brain regions are functionally connected, what we mean is simply that their BOLD signals “looked similar” during the period of time in which we recorded them. The study of functional connectivity inferred from resting-state data is referred to as resting-state functional-connectivity MRI (rs-fcMRI).

2.2 Graph Theory

Graph theory, at its simplest, is the study of objects which encode relationships. The two fundamental building blocks of graph theory are a *vertex* (also, equivalently,

called a *node*) and an *edge*. The common interpretation of a graph is very straightforward: the vertices represent entities of some sort and the edges represent relationships between those entities.

Here, we will limit ourselves to studying undirected graphs; while it is clear that on the lowest physiological level, the connectivity between individual neurons is directional, the issue becomes more complex at the macroscopic scale we wish to study. Connections between gross anatomical regions are complex and, in many cases, bidirectional, though both directions are likely not to be equally strong. We choose to study undirected graphs not because we believe that the true underlying system is undirected, but rather because of the limitations of our methods. The time series analysis methods used here are inherently symmetric. And, worse, even when one explicitly attempts to infer directionality, the best techniques are no more than about 60 % accurate,⁴ and some common techniques, such as Granger causality, perform no better than chance [20].

Graphs are often classified into broad categories according to their topological properties. For example, a graph where every vertex has the same number of edges is referred to as a regular graph. Clearly, regular graphs are very structured entities, and if spatial constraints regarding how the edges are connected are added (viz., an edge vertex is connected to its closest neighbors only), we reach the extreme of structure in a graph: a regular lattice. Imagine extending a small 4-regular graph to several hundred nodes, but keeping each vertex of degree 4. This graph would have high transitivity but a very long average path length.

At the far opposite extreme of structure, we have the Erdős–Renyi random graph, which one can imagine being generated thus: after picking a fixed set of vertices, we flip a coin for each pair and add an edge only if the coin comes up “heads”. With edges chosen randomly, this graph will, with high probability, have very low transitivity; however, because some “long-distance” edges will be chosen purely by luck, the average path length in the graph will typically be very small – the long-distance edges act as a shortcut through the graph.

Between these two extremes exists the class of small-world graphs. Small-world graphs are characterized by having both relatively high transitivity and relatively short path lengths; this is typically achieved with a structure consisting of densely connected local clusters (thus achieving high transitivity) with the occasional long-distance edge (thus shortening significantly the average path length). Small-world networks appear to be common in both natural systems (e.g., cellular metabolic networks, genetic transcription networks, food chains, and social networks) and artificial systems (e.g., road networks and power grids). Most significantly for the present chapter, the small-world structure has been proposed as a dominant feature of brain networks (see, e.g. [3] for a review).

⁴This figure comes from an analysis in [20] in which known “ground truth” networks we used to construct simulated fMRI data, to which various network inference techniques we then applied. The accuracy reflects how much of the topology of the ground truth network the inference method managed to capture.

2.3 Preprocessing

Raw fMRI data is noisy and contaminated by a host of artifacts, necessitating fairly aggressive preprocessing. Although the specific details of preprocessing pipelines vary somewhat between studies, we present here a “consensus” preprocessing pipeline based on the protocol optimization results in [23]. This is precisely the pipeline that we use for our own work on thresholding. For each subject, we acquire both a single high-resolution “anatomical” volume (3D) and a series of lower-resolution “functional” volumes (a time series of lower-spatial-resolution volumes that have been acquired sequentially).

The initial preprocessing steps consist of the following steps, applied to functional data:

1. Deletion of first four low-resolution volumes acquired to allow for stabilization of the T1 signal.
2. Brain extraction/skull stripping.
3. Slice-time correction – each “slice” in an fMRI volume is acquired at a slightly later time than the previous slice. For our time series to be comparable along the z -axis, it is necessary to compensate for this.
4. Motion correction – even when the subject is instructed to lie still, subject motion is unavoidable. Images can be corrected to a single spatial baseline using standard image registration algorithms.
5. Spatial smoothing (5 mm full width at half maximum Gaussian).
6. Prewhitening (removal of spurious/unwanted autocorrelations).
7. Removal of residual motion by linear regression.

This is followed by anatomically driven preprocessing on both the high-resolution anatomical data and the functional data:

1. Brain extraction/skull stripping of the anatomical image.
2. Tissue-type segmentation into gray matter, white matter, and cerebrospinal fluid. (white matter, and cerebrospinal fluid do not contain neuron soma, and any signal detected in these areas is noise, rather than functionally induced).
3. Transformation of tissue masks generated in higher-spatial-resolution imagery (viz., cerebrospinal fluid, white matter, or gray matter masks) into the corresponding lower-resolution coordinate space of the functional imagery.
4. Extraction of mean time series for the cerebrospinal fluid and white matter to use as regressors in the final processing step.
5. Registration of functional images to anatomical images according to the Montreal Neurological Institute’s MNI 152 T1 2 mm standard space. This is a purely linear, rigid, transformation.

In the penultimate step, we regress out the signals of the cerebrospinal fluid and white matter from our functional data set. We then perform temporal band-pass filtering ($0.009 \text{ Hz} < f < 0.08 \text{ Hz}$) and mask the output so that it includes only time series derived from voxels containing gray matter. Cortical gray matter is distributed

as a sheet on the surface of the brain which is approximately 2–4 mm thick; the gray matter contains the neuron cell bodies of the cortex and, consequently, is where we expect to find the BOLD signal when imaging.

2.4 Time Series Analysis

The first step towards building a graph of functional connectivity involves analyzing which voxels in the data under study have related time courses. We generate an $n \times n$ (where n is the number of gray-matter voxels) similarity matrix by analyzing every pair of gray-matter voxel time series. In principle, one may use any time series comparison method to evaluate the similarity of the time series obtained from a pair of voxels. A full description of every metric currently in use is beyond the scope of this review, so instead we explain here three commonly used, but significantly different, similarity measures: the Pearson correlation, which is a linear, time-domain measure; the band-averaged coherence, a linear frequency-domain (spectral) measure; and the mutual information, a nonlinear measure.

Given the time series of two voxels, x and y , from an fMRI data set preprocessed as described above (and, thus, hopefully consisting primarily of the low-frequency BOLD signal), we will now define our similarity measures.

Formally, we define the population Pearson correlation coefficient in the usual way:

$$\frac{\sum_{i=1}^m (x[i] - \bar{x})(y[i] - \bar{y})}{\sqrt{\sum_{i=1}^m (x[i] - \bar{x})^2} \sqrt{\sum_{i=1}^m (y[i] - \bar{y})^2}},$$

where $x[i]$ indicates the i th element in the time series x , and \bar{x} is the mean of that time series. Intuitively, one can think of the correlation coefficient in a very simple way: step through the two time series under consideration, in a parallel fashion, and construct an ordered pair (x, y) from the values of the two series at each point in time. Plot these pairs and fit a straight line to the plot. The closer the plotted points are to the best-fit line, the higher the correlation of the two series. Note carefully here the importance of the line. If two series are linearly correlated, then they will surely have a high correlation coefficient; imagine, though, that you were to perform such a plot and find a perfect monochrome rendering of the Sierra Nevada mountains. The Pearson correlation between these two series would be quite low (since there is no single line fitting a photograph of the Sierra Nevada mountains), but you might feel justified in supposing that, in truth, these two series do indeed have a very special relationship. In principle, this is an inherent weakness of linear methods – they can find only lines. In the case of low-frequency BOLD signals, however, this may pose

less of a problem than it seems as the bulk of the informative relationships do indeed appear to be linear [20].⁵

It is clear from the description of the correlation that it inherently provides a time-domain method answering the question “how similar is the variation in amplitude of these series over time?” An alternative question one might wish to ask is “how similar are the power spectra of these two time series?” That is to say, “how similar are the *frequency* components in each time series?” Instead of comparing the time series directly at the temporal level, we are suggesting transforming them into the frequency domain and comparing their spectra. We desire a metric that will be high if two series are composed of components having very similar frequencies and amplitudes, and low otherwise. The band-averaged spectral coherence is just such a metric.

In its most general form, the coherence C_{xy} between x and y can be computed as

$$C_{xy} = \frac{|G_{xy}|^2}{G_{xx}G_{yy}},$$

where G_{xx} is the autospectral density of x , and G_{xy} is the cross-spectral density between x and y . Note that C_{xy} is a function of frequency, and thus, to obtain a scalar value, we average C_{xy} over the frequency band of interest. In a typical implementation, one estimates power and cross spectra using Welch’s modified periodogram averaging methods, with a window length of 50 and an overlap of 25. Each windowed segment is then normalized and weighted by a Hanning window (also of length 50). The power spectral density is estimated as

$$G_{xx}(\lambda) = \frac{1}{N} \sum_{n=1}^N |X_n(\lambda)|^2,$$

where X_n is the discrete Fourier transform of the n th windowed segment of x , and λ is a variable depending on frequency. The cross-spectral density is estimated as

$$G_{xy}(\lambda) = \frac{1}{N} \sum_{n=1}^N X_n(\lambda)Y_n^*(\lambda).$$

Finally, we estimate the coherence, averaged over the band of interest ($\bar{\lambda}$), as

$$C_{xy}(\bar{\lambda}) = \frac{|\sum_{\lambda} G_{xy}(\bar{\lambda})|^2}{\sum_{\lambda} G_{xx}(\bar{\lambda}) \sum_{\lambda} G_{yy}(\bar{\lambda})}$$

⁵The careful reader may find, as does this author, this conclusion somewhat disturbing. The brain is a Turing-complete computational system, and accurate measures of its state should show useful statistics far transcending the first moment; yet, the smoothed, preprocessed BOLD signal does not, it would appear.

Although the coherence moves us from the time domain to the frequency domain – which enables the detection of temporal relationships regardless of phase shifts – it is still a linear measure, as it is only capable of detecting linear spectral relationships. To search for nonlinear relationships in time series, we turn to the mutual information.

Where as the correlation and coherence operate directly on the time series (or the Fourier transforms thereof) under investigation, the mutual information operates instead on the probability distributions of the hidden sources generating the series. Intuitively, if I am gathering time series data on two possibly related processes called A and B , the mutual information answers the following question: “How much does knowing the statistics of A allow me to infer about the statistics of B ?”

More formally; let A and B be the random variables modeling the process generating x and y , respectively. We compute the mutual information $I(A; B)$ between A and B as

$$I(A; B) = \sum_{b \in B} \sum_{a \in A} p(a, b) \log \left(\frac{p(a, b)}{p(a)p(b)} \right),$$

where $p(a, b)$ is the joint probability distribution function of A and B , and $p(a)$ and $p(b)$ are the marginal probability distribution functions of A and B , respectively. The quantity $I(A; B)$ measures how much knowing the distribution of A tells us about B , and vice versa. For example, if A and B are independent, then clearly $I(A; B) = 0$. If there is some statistical dependency between A and B , then $I(A; B)$ can quantify this dependency in units of bits.

In reality, of course, we do not have closed, analytical, expressions for A and B and must instead estimate them, ad hoc, from our data for example, with a k -nearest-neighbor estimation technique such as that of [13].

Even for relatively large data sets, the computational burden of these approaches is reasonably small. A Pearson correlation matrix for a data set with 15,000 gray-matter voxels can be computed in under 3 min on a typical quality laptop available in 2012 (viz., a hyperthreaded two-core (four effective cores), 2.2 GHz Intel i7-based 2011 MacBook Pro), and the calculation has only very modest memory requirements, as the results are streamed to disk. The band-averaged spectral coherence for the same data set, on the same machine, can be computed in approximately 1 h. Estimation of a mutual information matrix, however, requires more significant computational resources and would require approximately 1 day to complete.

2.5 Graph-Theoretical Analysis

Once a correlation (or coherence or mutual information) matrix has been generated, we may reinterpret it as the adjacency matrix of a graph thus:

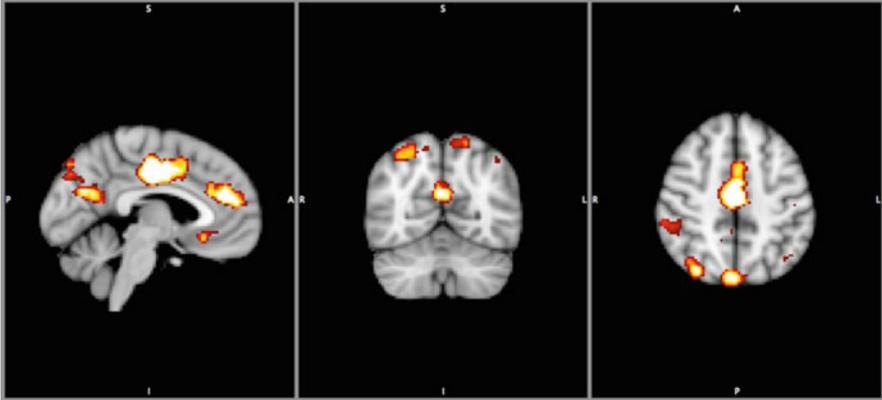


Fig. 1 Node degree projected back into voxel space and superimposed on an anatomical image. The degree is heatmapped so that lighter colors indicate higher degrees. The views from *left* to *right* are sagittal, coronal, and axial

- Each voxel is treated as a node in the graph.
- For every entry (i, j) in the correlation matrix:
 - If the entry falls below a user-specified threshold, do nothing.
 - If the entry exceeds the threshold, add an edge between nodes i and j , having a weight specified by this entry.

With the graph constructed, we can now begin to investigate its topological properties. A currently popular approach is to characterize such graphs in terms of both per-vertex and whole-graph metrics; for example, one might calculate the degree of each vertex in the hope of identifying regions of the brain which are “more connected” than others. In Fig. 1, we show the result of computing the degree of each vertex for a graph derived from a resting-state fMRI scan for a single subject. Once we have computed the degrees on the graph, we invoke our trivial mapping from graph vertices to voxels to visualize the degree in “brain space” rather than “graph space”; here, the degree is heatmapped so that brighter colors represent vertices (voxels) with higher degrees. The degree is a very coarse metric, and putatively neuroscientifically interesting results have been obtained using more sophisticated combinatorial and spectral methods such as those based on vertex betweenness, which can be informally characterized as measuring the fraction of shortest paths which travel through a given vertex.

An interesting consequence of using more sophisticated metrics such as betweenness is that although the topological interpretation of these metrics on a graph is as straightforward as for the degree, the same cannot be said of the neuroscientific interpretation of the results in “brain space”. A vertex with high degree represents a “more connected” brain region, but what does a vertex with high “betweenness” represent? Recall that the relationships in our graph are *not* defined by anatomical connectivity, but by *functional* connectivity, which is simply to say that two regions

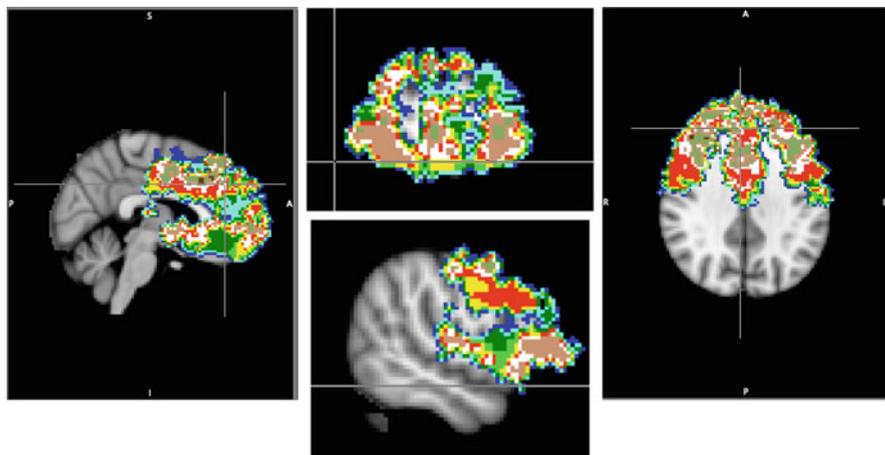


Fig. 2 Node modularity (computed using Markov clustering) projected into voxel space and superimposed on an anatomical image. Nodes with the same color belong to the same module

are considered related if their patterns of activity are similar enough. In such a framework the intuitive definition of betweenness, and very many other more complex graph metrics, breaks down; however, even without clear physiological semantics, such metrics have already proven useful in characterizing, for example, developmental trajectories in the brain [9].

Work in this area is still very exploratory in nature. It is clear that different metrics can yield different insights, although the correct physiological interpretation of these results is often less clear. It is equally apparent that many metrics recapitulate the same results; sometimes this is a consequence of a rigorously provable relationship between metrics, and other times it is a consequence of the particular topologies of the graphs themselves. We have recently provided a comparison of the repeatable intra-metric similarity (and robustness to parameters) for several commonly used metrics [8], but it seems likely that a more interesting direction for future work will be to develop metrics which are specialized for the restricted class of graphs derived from neuroimaging data.

An alternative approach to studying the topology of neuroimaging-derived graphs is to attempt a modular decomposition. In Fig. 2 we show an example of the result of applying Markov clustering [24] to the same resting-state scan used for Fig. 1, but restricted here to the prefrontal cortex. Following convergence of the clustering algorithm, the vertices were again mapped back into brain space and colored according to module membership. The modularization here yields strong support for a dorsal/ventral⁶ distinction in the prefrontal cortex. Support for a

⁶“Dorsal” refers here to the top of the brain – the part above the eyeline in a human standing upright – and “ventral” refers to the bottom part.

rostral-caudal⁷ gradient is less clear; instead, we find a complex pattern of interacting, nesting, and embedded networks along the rostral-caudal axis. Intriguingly, this challenges the model of two parallel strict hierarchies in dorsolateral prefrontal cortex and dorsomedial prefrontal cortex recently suggested in [22].

3 On Thresholding Graphs

Determining a reasonable threshold to use in the generation of the connectivity graph is a crucial issue, and we must be aware that, by design, any similarity matrix is composed of multiple comparisons (quadratic in the number of time series). We begin this section with a brief theoretical discussion of an overly conservative thresholding approach that applies only to correlation matrices. This is followed by a description of two thresholding procedures which are dominant in the current literature: thresholding by fixing topological expectations (the so-called *S*-value method) and thresholding by using the false discovery rate (FDR) to control the error. Noting the limitations of these techniques, we extend them with a new approach which computes thresholds based on the expectation of modular structure in the underlying graph; we implement this approach using the tools of random matrix theory (RMT).

3.1 Current Approaches

In the case of the Pearson correlation, we can rely upon theory to guide our thresholding. Given a desired global *p*-value, we apply a Bonferroni correction to obtain a corrected *p*-value, which we may then convert into a minimum *r*-score for the observed correlation.

For example, our preprocessing pipeline typically produces approximately 15,000 gray matter voxel time series, which results in

$$\frac{1}{2} \cdot (15,000^2 + 15,000) = 112,507,500$$

comparisons during correlation analysis. Assuming that our acceptable global threshold for significance is $p = 0.05$, we can correct to $p \approx 4.4 \times 10^{-10}$.

⁷The rostral-caudal axis follows a curved path through the head, beginning roughly in the region of the nose, and proceeding straight back towards the midbrain, at which point it bends 90° downwards to follow the spine.

Although this looks bleak, for time series of length 300, this yields a threshold of $r \approx 0.34$.⁸

This is, of course, overly conservative. Following this train of thought, a more statistically appropriate approach might be bootstrapping the correlation matrix. Unfortunately, the computational demands of computing the full correlation matrix make not only permutation analysis, but also a Monte Carlo approach, relatively infeasible. Worse still, a resampling approach to the band-averaged coherence and mutual information may be out of reach owing to significant computational requirements. Fortunately, in the discussion below, we find that these problems may be irrelevant.

Although the Bonferroni-corrected approach *should be* overly conservative, when we apply it to real data we find that it is, in fact, *overly liberal*. The Bonferroni-corrected r -value introduces edges in the graph which are well below the physiological noise floor in our data (viz., they are the product of spurious correlations). This is a consequence of the fact that our simple analysis treats the time series of each voxel as if it were generated by a unique, statistically independent source. The reality of the situation could scarcely be more different: voxels are highly temporally and spatially correlated as a consequence of both the underlying physiological processes being measured and the physical properties of the machines performing the measurement. Without appealing to detailed statistical models of these processes (in both spatial and temporal dimensions), a completely analytical approach to this problem is untenable.

An alternative approach is to define the cutoff threshold post hoc, attempting to match the properties of the induced graph to those expected from other analyses presented in the literature. A straightforward approach is to demand a constant relationship, S , between the number of nodes in the graph and the average node degree. The relationship $S = \log |V| / \log K$ (where $|V|$ is the number of vertices in the graph and K is the average degree) has been suggested in the literature. If one chooses a fixed S value at which to compare graphs, the only value that then needs to be computed is K , since the number of voxels, $|V|$, in each scan is fixed. With K computed from the formula, finding a threshold value for each particular graph is simply a matter of determining what threshold retains E edges in the graph, where $K = E / |V|$.

Note also that our definition of S is an approximation to the average path length in a large Erdős–Rényi graph (see the scaling relation of [1]), so we are, in some sense, fixing a target minimum path length. Of course, the graphs we are dealing with in practice are most certainly not random graphs (in the sense of an Erdős–Rényi graph), so this characterization of the metric S is at best heuristic, and at worst completely misleading.

That said, this approach has one overwhelmingly positive attribute: it facilitates easier cross-modality, cross-analysis, cross-subject comparison. By setting the

⁸We derived its r -value threshold by converting the corrected p threshold to a t -score with 298 degrees of freedom, giving $t = -6.33$, and computing $6.33 / \sqrt{298 + 6.33^2}$.

threshold based on structural properties of the graph, rather than statistical properties of the underlying analysis (such as r -values), we are more likely to end up “comparing apples with apples”. At the same time, we also run the risk of missing significant structural differences by forcing the structures to be similar.

Recently, a thresholding approach based upon controlling the false discovery rate (FDR) has become popular. When performing multiple comparisons, we refer to the rate of *false positives* (for us this would be ascribing a functional relationship to a pair of time series which are, in fact, unrelated) as the *false discovery rate*. We control the FDR by specifying a rate of false discovery, between 0 and 1, which we are willing to accept, and applying a procedure to ensure that, on average, this rate is maintained. Specifically, let E be the set of measurements for potential edges in our graph (equivalently, the set of entries in our similarity matrix). We then do the following:

1. Select an acceptable FDR bound α (a common choice is $\alpha = 0.05$).
2. Sort E into an ordered set, from smallest to largest: $E' = \{e_0, e_1, \dots, e_n\}, e_0 \leq e_1 \leq \dots \leq e_n$.
3. Find the largest $i \in \mathbb{N}$ such that $e_i < \alpha \cdot i / |E|$.
4. Select e_i as the threshold.

One potential drawback to the FDR approach is that it necessarily treats all values in E as exchangeable, in the sense that one can reasonably expect the semantics of comparing a pair of values to be constant across any pair in E . This conflicts with our understanding of real neuroimaging results, in that some brain areas (e.g., the primary visual cortex) are dominated by extremely strong correlations in activity while other areas (e.g., the parietal cortex) show correlations which are still very much physiologically “real” and scientifically interesting, but up to an order of magnitude weaker. A possible solution is to consider an approach in which one allows *local structure* to influence the threshold, rather than a simple globally compared magnitude.

3.2 Thresholds from Random Matrix Theory

We suggest here a mathematically elegant approach to choosing graph thresholds starting from a statistical basis rather than a post hoc heuristic. In particular, we note that a failing of the currently popular thresholding techniques is that they are conceived at the level of abstraction of the *graph* rather than that of the *correlation matrix*. Although a correlation matrix can certainly be profitably viewed as the adjacency matrix of a graph, it is an error to jump to this view prior to examining the statistics of the matrix. While all correlation matrices can be viewed as graphs, *most* graphs do not have adjacency matrices that are correlation matrices, and by analyzing thresholds at the graph level, we are thus throwing away our knowledge of the special structure of correlation matrices.

The theory of random matrices was proposed originally by Wigner and Dyson, in the context of studying the spectra of complex nuclei (see, e.g., [25]), and is specifically suited to studying phase transitions between disordered and ordered, modular, systems defined by correlation matrices. Consider the matrices studied here of correlations between the (filtered) time series of gray-matter voxels. A simple model for the value of the correlation between voxels i and j might look like this:

$$r_{i,j} = r_{i,j}^{\dagger} + \varepsilon,$$

where $r_{i,j}^{\dagger}$ represents the *true correlation* between the physiological processes underlying the signals observed at voxels i and j and ε represents the sum of the many noise sources (e.g., physiological noise and scanner noise). For a single correlation, we have no way of separating r^{\dagger} from ε , but if we have a priori a model for the global structure of the matrix of true correlations r^{\dagger} , then we can attempt to extract a reasonable estimate of r^{\dagger} from $r_{i,j}^{\dagger} + \varepsilon$.

In essence, this is exactly what the currently popular thresholding methods attempt to do, albeit in a nonrigorous way: remove ε from the matrix by setting a threshold that yields a matrix corresponding to a graph with an expected structure.

In random matrix theory, one studies – amongst many other things – the spectra of real, symmetric matrices representing systems composed of a sum of signal and noise. In particular, we are interested here in looking at the statistical properties of the eigenvalue spacing of our correlation matrices. RMT tells us that when we observe the spacing of eigenvalues of a correlation matrix, we should expect to find a distribution of spacings conforming to one of two possibilities: the Gaussian orthogonal ensemble (GOE), where there are strong correlations everywhere, and, at the other extreme, Poisson statistics, where there exist strong correlations only along the (block) diagonal of the matrix.

In the context of our correlation matrices, the former case – in which the eigenvalue spacings follow a GOE distribution – is indicative of a matrix which is dominated by noise and spurious correlations; the latter case, where the eigenvalues follow Poisson statistics, is indicative of a matrix describing a highly modular system.

If one is prepared to accept the hypothesis that the physiological networks generating the observed BOLD signal in our neuroimaging data sets are indeed modular, then we now have exactly the statistical tools we need to separate these modular networks from noise. We must simply find the threshold at which the distribution of eigenvalue spacings for our matrix completes the transition from GOE to Poisson statistics. Once we have identified a threshold which generates Poisson eigenvalue spacings, we can be mathematically sure that we have extracted, to the best of our ability given the signal-to-noise ratio of the data, the modular true correlations present in the data while having sacrificed the minimum number of true correlations during the elimination of spurious, noisy correlations.

More formally, for a correlation matrix of order n , let E_i for $i \in \{1, \dots, n\}$ denote the magnitude-ordered list of eigenvalues of the matrix. We perform a spectral unfolding procedure to obtain a distribution with the eigenvalue spacing represented in units of the local mean eigenvalue spacing. In particular, we unfold our eigenspectrum by estimating the integrated density of the spectrum, which we then fit to a cubic spline.⁹ Individual eigenvalues are then projected into the unfolded representation by evaluation on the spline; we denote these transformed eigenvalues by e_i . We then simply compute the pairwise difference between adjacent transformed eigenvalues (i.e., $d = e_{i+1} - e_i$), and from this generate the probability density $P(d)$ of the unfolded eigenvalue spacing. Formalizing the relationships noted above, we consider two extreme distributions for $P(d)$: the Wigner–Dyson distribution (for the GOE case), where

$$P(d) \approx \frac{1}{2} \pi d e^{-\pi d^2/4},$$

and the Poisson distribution

$$P(d) \approx e^{-d}.$$

Indeed, we can see in Fig. 3 that the nearest-neighbor spacing distribution (NNSD) for the eigenvalues of a raw, unthresholded, fMRI-derived correlation matrix pictured in the upper left very closely follows GOE statistics. If the same matrix is subjected to strict thresholding, the empirical NNSD looks a great deal “more Poisson”.

For a candidate threshold value, we can do the following: threshold the matrix at the candidate value, unfold the eigenspectrum and compute $P(d)$, and compare $P(d)$ with the GOE and Poisson distributions. If $P(d)$ follows the GOE, it is dominated by noise and our threshold is too low. If $P(d)$ is a Poisson distribution, it represents a modular network, and we assume it is thus dominated by signal rather than noise. Maximizing the signal-to-noise ratio then simply becomes a game of identifying the threshold at which the statistics of the unfolded eigenvalue difference (UED) change from Wigner–Dyson to Poisson statistics. The transition between GOE and Poisson distributions cannot occur instantaneously between thresholds, but rather resembles a phase transition, with intermediate threshold values having some degree of “Poissonness” and some degree of “Wignerness”.

A very liberal estimate of the threshold, certain to still include a great deal of noise, can be obtained by finding the first point at which the UED distribution begins to differ significantly from the Wigner–Dyson distribution. Likewise, a conservative

⁹One may, of course, fit the curve to an arbitrarily sophisticated function; we chose cubic splines here, as they have been demonstrated to work well in applications ranging from neutron scattering to quantitative finance.

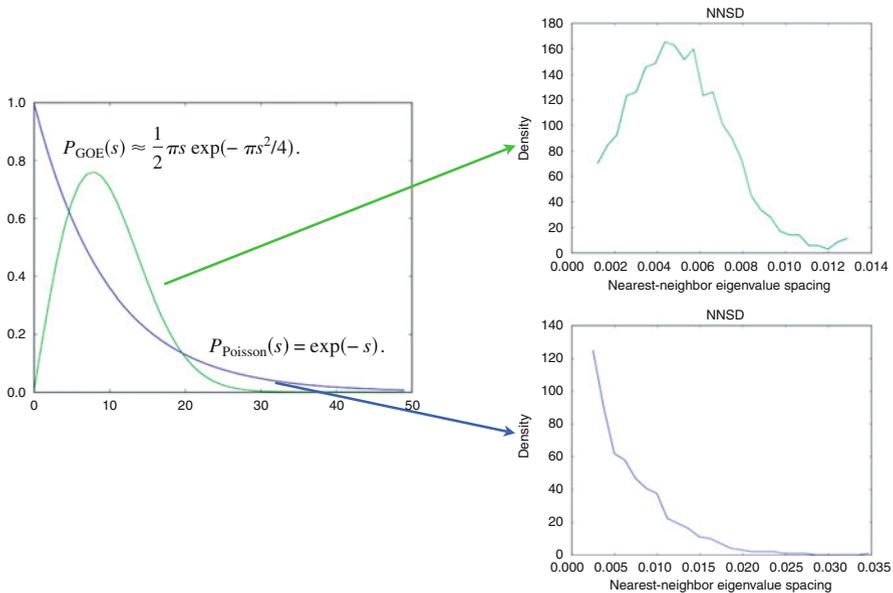


Fig. 3 Left: a plot of the GOE (green) and Poisson (blue) distributions. Right: above, the NNSD of eigenvalues from an unthresholded fMRI correlation matrix; below, the NNSD for the same matrix following strict thresholding

estimate may be obtained by identifying the threshold at which the UED distribution becomes significantly Poisson.

We now propose a very simple algorithm for finding such a threshold for a given correlation matrix M . Begin with a very low threshold $t \approx 0$. Compute $M' \leftarrow M$, where all matrix entries $m_{i,j} < t$ are set to zero, followed by $E \leftarrow$ eigenvalues of M' and $e \leftarrow U(E)$, where $U(E)$ is the smoothed, integrated eigenvalue density. Then, compute the distribution P of nearest-neighbor spacings in $U(E)$. If an Anderson–Darling goodness-of-fit test of P against a Poisson distribution yields $p < \alpha$ (for, say, $\alpha = 0.05$), terminate and report the threshold t ; otherwise, increase t by some small increment δ and repeat until this process terminates. The choice of value for δ controls a trade-off between computation time and threshold precision.

The result of applying this approach to matrices derived from ten resting-state fMRI scans can be seen in Fig. 4. The unthresholded matrices have very high Anderson–Darling scores (when tested against an exponential distribution), but these scores rapidly drop as we increase the threshold to an r -value of approximately 0.15. Beyond this value, the scores remain relatively stable, modulo some apparently stochastic factor. Thus, for the matrices considered in this diagram, thresholding in the neighborhood of $r > 0.15$ would be appropriate.

The computational complexity of this approach is tractable, even for large matrices, in a high-performance computing environment given the wide availability

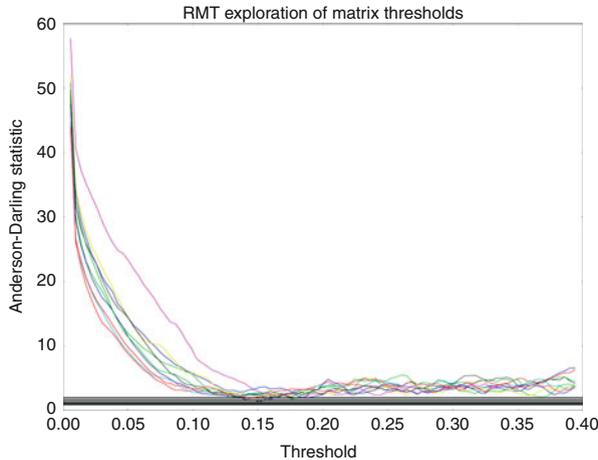


Fig. 4 Plots of Anderson–Darling test statistic, testing the fit of the empirical NNSD to an exponential distribution, for increasing thresholds (values below the threshold in the correlation matrix are set to zero) for ten resting-state fMRI scans. *Horizontal black lines* denote critical values for the Anderson–Darling score

of excellent parallel eigenvalue finding libraries¹⁰ (see, e.g., LAPACK [2]); finding the eigenvalues is the dominant computational step in our algorithm.

3.3 Evaluation of RMT-Based Thresholding

We now compare the real-world effectiveness of our RMT-derived thresholding approach with two approaches commonly used in the neuroimaging literature, based on the S -value and the false discovery rate. Our test data set consists of 1,000 resting-state fMRI scans from the 1,000 Functional Connectomes Project [6]. Each scan was preprocessed and then converted into a correlation matrix according to the protocol we outlined above. We focus solely on correlation matrices here, as these are presently by far the mostly commonly used basis for building graphs in the neuroimaging literature.

For each matrix, we computed an optimal threshold (below which all entries in the matrix should be set to zero) using the S -value approach (with $S = 2.0$), The FDR, and our RMT approach. The resulting thresholds are plotted in Fig. 5. Three trends are immediately visible:

¹⁰It is also worth noting that significant progress has been made in spectral decomposition on commodity GPU hardware; see, e.g., [15].

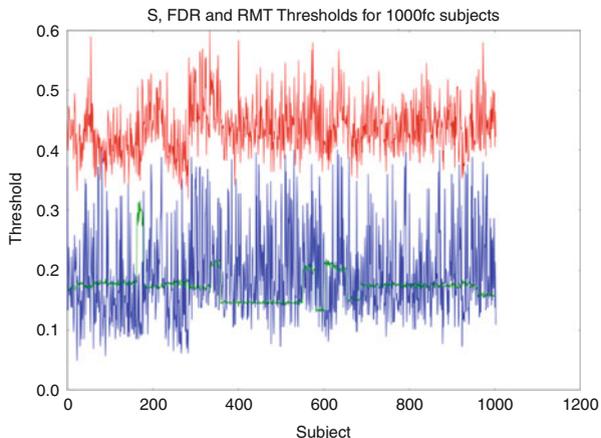


Fig. 5 Optimal thresholds for 1,000 resting-state fMRI scans as determined by the *S*-value approach (red), the FDR (green), and the proposed RMT approach (blue)

1. The variance in thresholds for the FDR approach is significantly lower than for either the *S*-value or the RMT-based approach.
2. The RMT and FDR thresholds are almost always lower than the *S*-value-based thresholds.
3. There is no consistent greater/less than relationship between the FDR- and RMT-based approaches, though it appears that were the variance of the RMT thresholds smaller, the two values would track each other relatively well.

It is not possible to draw a clear conclusion about which method is “best” from the data at hand; we can, however, offer some observations. If one expects thresholds to vary very little between subjects, then the significantly lower variance in the FDR scores would make this approach quite attractive. Conversely, if one expects some nontrivial physiological differences between subjects, then one might interpret the higher variance in the results from the RMT and *S*-value approaches as higher intrasubject sensitivity. We can investigate this hypothesis by looking at the distributions of thresholds while limiting ourselves to considering subjects from a single site (to avoid possible confounds from the use of different equipment at different sites contributing to the 1000 Functional Connectomes Project (http://fcon_1000.projects.nitrc.org/) data set).

In Fig. 6 we plot the distribution of thresholds for the *S*-value, FDR, and RMT approach for 198 scans from the Harvard site only. As expected, the width of the histogram is significantly lower for the FDR approach, but it is the shape of the histograms which is more interesting. While the FDR and *S*-value histograms have a roughly Gaussian-appearing shape, the RMT histogram appears to have a long right tail. It is, of course, impossible to make any definitive inferences from such a crude analysis, but this does suggest that, at the very least, the source of variance in the RMT thresholds is different from that in the FDR and *S*-value thresholds.

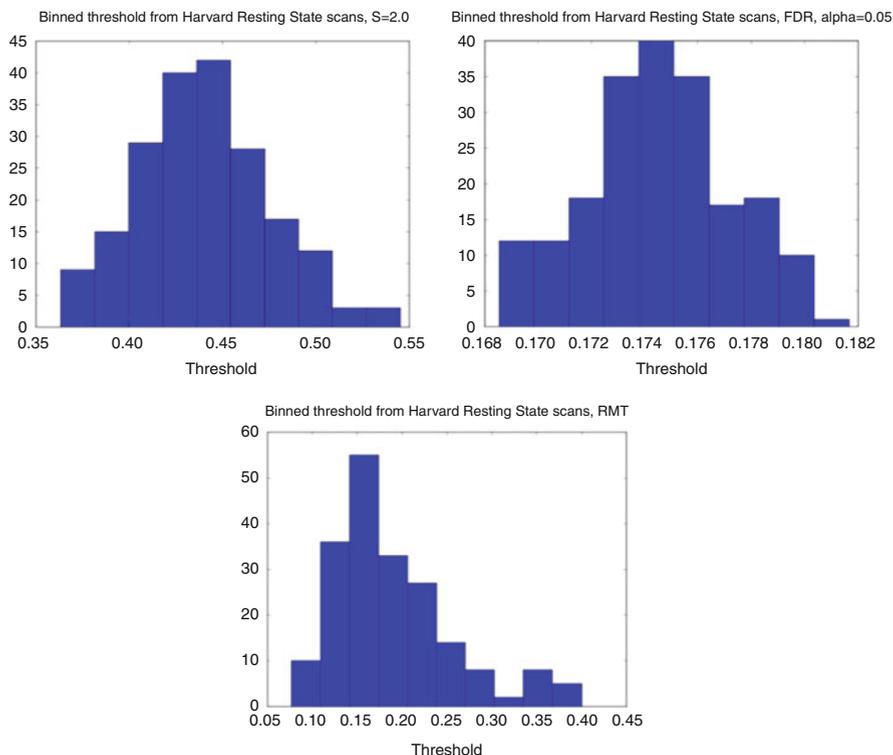


Fig. 6 Histograms of threshold values obtained for 198 resting-state fMRI scans from a single site for the S -value, FDR, and RMT approaches

Although providing food for thought, these results cannot conclusively identify a “best” approach to thresholding. Future work in which physiologically and physically accurate models generate simulated fMRI data based on precisely specified networks will allow qualitative comparison of thresholding techniques. If we have a contrived “ground truth” mesoscale brain network to hand, and a reasonable model for the physiological and physical processes that would produce a BOLD signal for this network (e.g., inverse dynamic casual modeling [10]), we can directly test the results of our thresholding against a 100% known ground truth – a concept not available to us in real data sets.

4 Closing Thoughts

We have provided a brief overview of the growing interest in applying tools from graph theory to the analysis of graphs derived from neuroimaging data. We introduced fMRI – currently the most popular functional neuroimaging technique – and touched on the details of what is required to preprocess fMRI data. Moving

from preprocessed neuroimaging-derived time series, we discussed methods for comparing time series and thus inferring a static graph describing relationships between these time series. We then focused in on the critical issue of choosing a threshold to separate “real” physiologically relevant relationships from noise in our graphs, including the description of a new thresholding technique based on random matrix theory.

The bulk of the work done in this area to date has been very pragmatic and applied in nature; the field is driven by neuroscientists who seek tools which can help them further understand the brain, with relatively less interest in the nature of the tools themselves. Several clear, pragmatic problems remain open, including two quite significant ones. First, the correct interpretation of complex graph metrics (e.g., betweenness or PageRank) on functionally derived graphs is still unclear. Second, the metrics that are typically employed are classical metrics optimized for studying abstract graph topologies, and networks which arise in physics and the social sciences. It would be surprising if there did not exist better metrics, tailored specifically to the questions most interesting to neuroscientists.

Moving to a more abstract, theoretical level, there has been little investigation of the deeper topological properties of neuroimaging-derived graphs. Is there a theory of descriptive complexity for these graphs? Can we define hierarchies of structure which correspond to, for example, clinical metrics of impaired neurological function? This has been done, for example, for simple measures such as the clustering coefficient, but is there a deeper unifying theory underlying these observations? Questions of computability arise naturally as well: if we imagine these graphs to represent a snapshot of synchronizations between the nodes of a computing machine, what can we say about the nature of the machine from such a snapshot? Is it even possible to reason about questions of computability and complexity given such a coarse projection of the dynamics of a computing system?

For one who studies the theory of computation and its application to natural systems, it is difficult to ignore the attraction of the dynamic computational system that is the human brain. With the recent, and rising, interest in graph theory from neuroscientists, it is now possible to find collaborators in the mathematical sciences and neurosciences who speak a common language and have a common interest in the dynamics of complex computational systems. Many problems – both straightforward and pragmatic and deeper, more philosophical – remain open and ready for investigation; and, as always, with each solution comes even more open problems.

References

1. R. Albert, A.-L. Barabási, Statistical mechanics of complex networks. *Rev. Modern Phys.* **74**(1), 47–97 (2002)
2. E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, *LAPACK Users' Guide*, 3rd edn. (Society for Industrial and Applied Mathematics, Philadelphia, 1999)

3. D.S. Bassett, E. Bullmore, Small-world brain networks. *Neuroscientist* **12**(6), 512–523 (2006)
4. M. Bernstein, K. King, X. Zhou, *Handbook of MRI Pulse Sequences* (Elsevier/Academic Press, Burlington, 2004)
5. B. Biswal, F.Z. Yetkin, V.M. Haughton, J.S. Hyde, Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* **34**(4), 537–541 (1995)
6. B.B. Biswal, M. Mennes, X.N. Zuo, S. Gohel, C. Kelly, S.M. Smith, C.F. Beckmann, J.S. Adelstein, R.L. Buckner, S. Colcombe et al., Toward discovery science of human brain function. *Proc. Natl. Acad. Sci.* **107**(10), 4734–4739 (2010)
7. R.L. Buckner, J. Sepulcre, T. Talukdar, F.M. Krienen, H. Liu, T. Hedden, J.R. Andrews-Hanna, R.A. Sperling, K.A. Johnson, Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer’s disease. *J. Neurosci.* **29**(6), 1860–1873 (2009)
8. M. Daley, Optimizing voxel scale graph theoretical analysis of fMRI-derived resting state functional connectivity. Master’s thesis, The University of Western Ontario, 2012
9. N.U.F. Dosenbach, B. Nardos, A.L. Cohen, D.A. Fair, J.D. Power, J.A. Church, S.M. Nelson, G.S. Wig, A.C. Vogel, C.N. Lessov-Schlaggar, K.A. Barnes, J.W. Dubis, E. Feczko, R.S. Coalson, J.R. Pruett, D.M. Barch, S.E. Petersen, B.L. Schlaggar, Prediction of individual brain maturity using fMRI. *Science* **329**(5997), 1358–1361 (2010)
10. K.J. Friston, L. Harrison, W. Penny, Dynamic causal modelling. *Neuroimage* **19**(4), 1273–1302 (2003)
11. E.M. Haacke, R.W. Brown, M.R. Thompson, R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design* (Wiley, New York, 1999)
12. S.A. Huettel, A.W. Song, G. McCarthy, *Functional Magnetic Resonance Imaging* (Sinauer Associates, Sunderland, 2004)
13. A. Kraskov, H. Stögbauer, P. Grassberger, Estimating mutual information. *Phys. Rev. E* **69**(6), 066138 (2004)
14. K.K. Kwong, J.W. Belliveau, D.A. Chesler, I.E. Goldberg, R.M. Weisskoff, B.P. Poncelet, D.N. Kennedy, B.E. Hoppel, M.S. Cohen, R. Turner, Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci.* **89**(12), 5675–5679 (1992)
15. S. Lahabar, P.J. Narayanan, Singular value decomposition on GPU using CUDA, in *IEEE International Symposium on Parallel & Distributed Processing, 2009. IPDPS 2009*. IEEE (2009)
16. N.K. Logothetis, J. Pauls, M. Augath, T. Trinath, A. Oeltermann, Neurophysiological investigation of the basis of the fMRI signal. *Nature* **412**(6843), 150–157 (2001)
17. M.E. Lynall, D.S. Bassett, R. Kerwin, P.J. McKenna, M. Kitzbichler, U. Muller, E. Bullmore, Functional connectivity and brain networks in schizophrenia. *J. Neurosci.* **30**(28), 9477–9487 (2010)
18. S. Ogawa, D.W. Tank, R. Menon, J.M. Ellermann, S.G. Kim, H. Merkle, K. Ugurbil, Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci.* **89**(13), 5951–5955 (1992)
19. J. Schummers, H. Yu, M. Sur, Tuned responses of astrocytes and their influence on hemodynamic signals in the visual cortex. *Science* **320**(5883), 1638–1643 (2008)
20. S.M. Smith, K.L. Miller, G. Salimi-Khorshidi, M. Webster, C.F. Beckmann, T.E. Nichols, J.D. Ramsey, M.W. Woolrich, Network modelling methods for fMRI. *NeuroImage* **54**(2), 875–891 (2011)
21. O. Sporns, *Networks of the Brain* (The MIT Press, Cambridge, MA 2010)
22. A.A. Taren, V. Venkataraman, S.A. Huettel, A parallel functional topography between medial and lateral prefrontal cortex: evidence and implications for cognitive control. *J. Neurosci.* **31**(13), 5026–5031 (2011)
23. K.R.A. Van Dijk, T. Hedden, A. Venkataraman, K.C. Evans, S.W. Lazar, R.L. Buckner, Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. *J. Neurophys.* **103**(1), 297–321 (2010)

24. S. Van Dongen, Graph clustering via a discrete uncoupling process. *SIAM J. Matrix Anal. Appl.* **30**(1), 121–141 (2008)
25. E.P. Wigner, Random matrices in physics. *SIAM Rev.* **9**(1), 1–23 (1967)