NeuroImage xxx (2014) xxx–xxx

Contents lists available at ScienceDirect

NeuroImage

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

To cut or not to cut? Assessing the modular structure of brain networks

Yu-Teng Chang [a,*], Dimitrios Pantazis [a], Richard M. Leahy [b]

[a] McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[b] Signal and Image Processing Institute, University of Southern California, Los Angeles, CA 90089, USA

ARTICLE INFO

Article history:
Accepted 8 January 2014
Available online xxxx

Keywords:
Functional connectivity
Community structure
Graph partitioning methods
Modularity
Random graphs

ABSTRACT

A wealth of methods has been developed to identify natural divisions of brain networks into groups or modules, with one of the most prominent being modularity. Compared with the popularity of methods to detect community structure, only a few methods exist to statistically control for spurious modules, relying almost exclusively on resampling techniques. It is well known that even random networks can exhibit high modularity because of incidental concentration of edges, even though they have no underlying organizational structure. Consequently, interpretation of community structure is confounded by the lack of principled and computationally tractable approaches to statistically control for spurious modules. In this paper we show that the modularity of random networks follows a transformed version of the Tracy–Widom distribution, providing for the first time a link between module detection and random matrix theory. We compute parametric formulas for the distribution of modularity for random networks as a function of network size and edge variance, and show that we can efficiently control for false positives in brain and other real-world networks.

1. Introduction

The complexity in the macroscopic behavior of brain networks has been highlighted and quantified in a number of neuroscience studies in recent years (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010). A multitude of topological features has been reported in the literature, including modular structures (Bullmore and Bassett, 2011; Chang et al., 2012), hierarchical patterns (Meunier et al., 2009; Salvador et al., 2005), distribution of hubs (Dimitriadis et al., 2010; Sporns et al., 2007; Tomasi and Volkow, 2011), and core extraction (Hagmann et al., 2008a). It has also been shown that brain networks follow a small-world property both from a functional (Bassett and Bullmore, 2006; Van den Heuvel et al., 2008) and a structural perspective (Vaessen et al., 2010; Wang et al., 2012). Recent findings have revealed alterations of brain network topology with aging (Chen et al., 2011), brain development (Fan et al., 2011), and pathologies of schizophrenia, autism, and epilepsy (Alexander-Bloch et al., 2010; Chavez et al., 2010; Rudie et al., 2013), underscoring the importance of networks as biomarkers of the normal and diseased brain.

Fundamental to identification of the architecture and organization of brain networks is the detection of modules, also called communities or clusters. In the context of graph theory, modules are groups of interconnected nodes, typically regions of parcellated cerebral cortex, that share common properties or have similar function within the network. Identification of modules can facilitate the prediction and discovery of previously unknown connections and components, and show how the network constitutes a collective and integrative system. Individual nodes can be classified according to their structural position in the modules; nodes with central position are essential for the stability and robustness of their corresponding modules, and nodes lying at the boundaries contribute to interactions across communities. Studies of network topology can reveal important properties of brain organization, for exampling revealing potential vulnerabilities, or in the case of hierarchical networks, possibly encoding clues to the evolution of the brain (Meunier et al., 2010).

Underscoring the central role of module detection, numerous methods have been proposed to identify community structure in brain networks. Perhaps the most popular is modularity (Newman, 2006), which compares the network against a null model and favors within-module connections when edges are stronger than their expected values. Divisions that increase modularity are preferred because they lead to modules with high community structure. We recently proposed a new method to compute network null models based on conditional expected probabilities and provided exact analytical solutions for specific parametric distributions (Chang et al., 2012). Our models enhance module detection, provide a principled approach to deal with networks with negative connections, and accurately represent the topology of networks without necessitating self-loops.

Despite the popularity of modularity methods, the identification of stopping criteria for graph division and the evaluation of the statistical significance of modules remain largely unaddressed. Given that random networks can demonstrate spurious modules due to incidental concentration of edges, even though they have no underlying organizational structure, controlling for false positives in community detection is of

* Corresponding author at: McGovern Institute for Brain Research, Massachusetts Institute of Technology, 77 Mass Ave, Building 46-5147, Cambridge, MA 02139, USA.
E-mail addresses: yutengchang@gmail.com, ytchang@mit.edu (Y.-T. Chang).

1053-8119/$ – see front matter © 2014 Published by Elsevier Inc.
http://dx.doi.org/10.1016/j.neuroimage.2014.01.010

paramount importance. This is even more evident in large networks where the number of possible divisions increases rapidly with the network size (Guimerá et al., 2004; Karrer et al., 2008). Therefore, confirming the statistical significance of any identified modules is essential before discussing other findings related to those structures.

Existing methods that control for spurious modular structures fall into three categories. The first category relies on creating comparable random networks in order to compute an empirical null distribution of modularity and establish a threshold that controls error rate at a nominal level, typically 5%. For example, Alexander-Bloch et al. (2010) estimated the distribution of modularity using two types of random networks, the Erdős and Rényi (1960) random graphs or randomly rewired networks (Maslov and Sneppen, 2002). Meunier et al. (2009) created random networks by randomizing either the elements of the adjacency matrix, or the time points of the time series whose pairwise correlation defined the edges of a graph. He et al. (2009) generated a set of node- and degree-matched random graphs for comparison. Reichardt and Bornholdt (2006) computed the z-values of modularity after estimating its empirical distribution through multiple random network realizations. Mirshahvalad et al. (2013) studied how different resampling schemes influence significance analysis.

The second category of methods also relies on resampled networks. Here the aim is to measure the robustness of modular structures on network perturbations. For instance, Karrer et al. (2008) proposed a method to perturb network connections and measure the resulting change in community structure using mutual information. Hu et al. (2010) offered a generalization to this approach by incorporating together the number of clusters, content of the clusters, and random perturbation parameters. Mirshahvalad et al. (2012) studied the robustness of large sparse networks by randomly adding extra links based on local information. Lancichinetti et al. (2010) evaluated the importance of single communities using combinatorics and a modified null model. Seifi et al. (2012) measured the significance of modules based on the stability of structures from either randomly perturbed networks or different initialization of non-deterministic community detection techniques.

Since all the above methods depend on edge rewiring and random network realizations, they are network-specific and do not generalize. The computational cost of generating multiple realizations of random networks can be significant, and even prohibitive for very large networks in the order of thousands of nodes. The third category of methods offers analytical closed-form solutions for the distribution of modularity in random networks. To the best of our knowledge, due to the complex form of the modularity function, there exists only one closed form solution for a specific case. Reichardt and Bornholdt (2007) used the Potts spin-glass model to get a theoretical prediction for modularity value in binary random graphs, either Erdős–Rényi type or scale-free random networks. However, their formula is restricted to binary sparse networks, which prevents its use with most real-world networks.

Given the lack of a principled analytical approach or computationally efficient algorithms to control for false positives in network module detection, much of the literature overlooks statistical inference in networks. To address this problem we provide a new analytical approach for statistical inference in module detection. Modularity belongs to the wide class of spectral clustering algorithms (Von Luxburg, 2007), which use the extreme eigenvalues and corresponding eigenvectors of a spectral decomposition to partition data into groups with similar properties. To evaluate the statistical significance of spectral clustering results, we need to compare the spectral decomposition of a given network against those from random networks. Since networks are represented by their adjacency matrix, a connection between random networks and random matrix theory is natural. The eigenvalue distribution of a specific type of matrices, Gaussian random ensembles, has been thoroughly studied in random matrix theory (Tao, 2012; Tracy and Widom, 2000). In this paper, we provide for the first time a link between module detection and random matrix theory by showing that the Tracy–Widom mapping of the largest eigenvalue of Gaussian random ensembles can be modified to predict the distribution of the largest eigenvalue of matrices used for modularity-based spectral clustering. Using this finding, we derive an accurate parametric form of the distribution of modularity in random networks and compute formulas that control the type I error rate at a 5% level on modularity-based partitions of weighted graphs. Our modeling is valid for a wide range of network sizes and the utility of the method is all the more important for larger networks, given that resampling methods can be computationally infeasible in such networks. We demonstrate our method in the brain and other real-world networks.

2. Methods

In this section, we first describe the modularity partitioning problem and its solution using spectral decomposition. We then motivate the use of the maximum eigenvalue of a difference matrix (adjacency matrix minus the null model) as a surrogate of modularity. Using a transformed Tracy–Widom distribution, we derive empirical parametric formulas that accurately predict the distribution of the maximum eigenvalue. We estimate our model parameters through Monte Carlo simulations of weighted Gaussian random networks.

2.1. Overview to modularity

Large-scale brain networks are typically constructed by assigning nodes to represent regions of parcellated cerebral cortex and edges to represent the pairwise interactions or connection across these regions. These connections could be based on structural data, for example white matter fiber-tracts derived from diffusion data, or functional coupling measured between time series of brain activation. Assume a brain network of N nodes with weighted undirected connections and an underlying modular structure, as exemplified in Fig. 1. The network is represented with an adjacency matrix $\mathbf{A}$ with elements A_{ij} indicating the connection strength across nodes i and j. The degree vector $\mathbf{k}$ has elements $k_i = \Sigma_j A_{ij}$, equal to the sum of all edge strengths associated with node i. The total sum of edge weights of the network is denoted as $m = \frac{1}{2}\Sigma_i k_i$.

Modularity was originally introduced as a measure of the quality of a particular division of a network (Newman and Girvan, 2004), but later became a key graph clustering algorithm, after recognizing its direct maximization using spectral graph partitioning (Newman, 2006). According to modularity, the community structure of the network is compared against a null network, i.e. a randomized network with the same number of nodes and node degrees but otherwise no underlying structure. If a natural division of a network exists, we should expect within-module connections A_{ij} to be stronger than their expected values E_{ij}, and the opposite should hold true for between-module connections

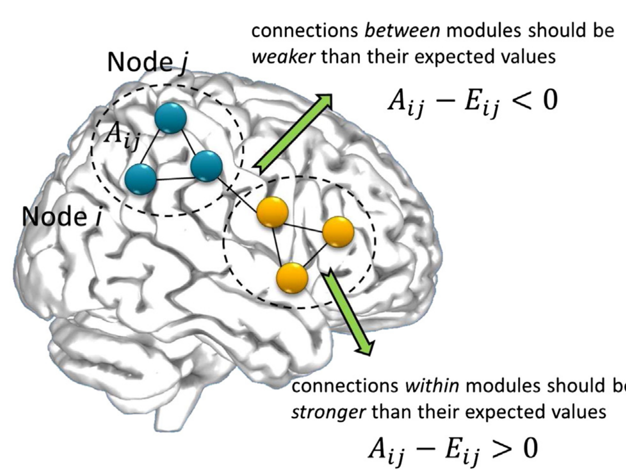

Fig. 1. Brain networks and modularity.

Please cite this article as: Chang, Y.-T., et al., To cut or not to cut? Assessing the modular structure of brain networks, NeuroImage (2014), http://dx.doi.org/10.1016/j.neuroimage.2014.01.010