NeuroImage 135 (2016) 92–106

Contents lists available at ScienceDirect

NeuroImage

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

Overlapping communities reveal rich structure in large-scale brain networks during rest and task conditions

Mahshid Najafi [a,b,*], Brenton W. McMenamin [b], Jonathan Z. Simon [a,c,d], Luiz Pessoa [b,**]

[a] Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742, USA
[b] Department of Psychology, University of Maryland, College Park, MD 20742, USA
[c] Department of Biology, University of Maryland, College Park, MD 20742, USA
[d] Institute for Systems Research, University of Maryland, College Park, MD 20742, USA

ARTICLE INFO

Article history:
Received 1 October 2015
Revised 3 April 2016
Accepted 24 April 2016
Available online 26 April 2016

Keywords:
Clustering
Overlapping communities
Networks
Bayesian methods
Functional MRI

ABSTRACT

Large-scale analysis of functional MRI data has revealed that brain regions can be grouped into stable "networks" or *communities*. In many instances, the communities are characterized as relatively disjoint. Although recent work indicates that brain regions may participate in multiple communities (for example, hub regions), the extent of community overlap is poorly understood. To address these issues, here we investigated large-scale brain networks based on "rest" and task human functional MRI data by employing a mixed-membership Bayesian model that allows each brain region to belong to all communities *simultaneously* with varying membership strengths. The approach allowed us to 1) compare the structure of disjoint and overlapping communities; 2) determine the relationship between functional diversity (how diverse is a region's functional activation repertoire) and membership diversity (how diverse is a region's affiliation to communities); 3) characterize overlapping community structure; 4) characterize the degree of *non*-modularity in brain networks; 5) study the distribution of "bridges", including bottleneck and hub bridges. Our findings revealed the existence of *dense* community overlap that was not limited to "special" hubs. Furthermore, the findings revealed important differences between community organization during rest and during specific task states. Overall, we suggest that dense overlapping communities are well suited to capture the flexible and task dependent mapping between brain regions and their functions.

1. Introduction

Network analysis of human neuroimaging data has contributed to a view of brain function in which groups of brain regions participate in functions rather than brain function relying on just regions operating in isolation (for influential early ideas, see Mesulam, 1981; Rakic et al., 1986). Functional MRI data during the so-called "resting state" has been extensively investigated in order to characterize *network structure*. A central finding is that, at rest, brain regions can be grouped into a relatively small number of stable *communities*, also called clusters or subnetworks. For example, Yeo et al. (2011) described a seven-community parcellation of cortical areas based on a large sample of participants. Based on anatomical and functional considerations, the communities were labeled as "visual", "frontoparietal", "default", and so on.

Much of the work employing modern network methods to study brain community structure and other network measures makes the assumption that each node (that is, brain region) belongs to a single community — thus, the overall network is partitioned into disjoint sets of clusters. However, the importance of understanding and characterizing overlapping structure has been discussed for some time in many disciplines, including sociology (Wasserman and Faust, 1994) and biology (Gavin et al., 2002); for example, biologists exploring protein interactions have found that a substantial fraction of proteins interact with several communities at the same time (Gavin et al., 2002). As nicely summarized by Palla et al. (2005, p. 814): "actual networks are made of interwoven sets of overlapping communities".

Large-scale analysis of functional MRI data indicates that brain networks are also overlapping. This is indicated, for example, by Independent Component Analysis (ICA) and other methods that allow community overlap (Smith et al., 2012). Indeed, there is increasing realization that brain regions may belong to several brain communities simultaneously (Cocchi et al., 2013; Cole et al., 2013; Pessoa, 2014; see also Hilgetag et al., 1996; Mesulam, 1998). It is still unclear, however, whether network overlap is sparse or dense (Fig. 1). In the former case, some key regions work as hubs that participate flexibly in multiple communities. In this scenario, network overlap is relatively limited and may be a property mainly of specific parts of the brain (say, prefrontal cortex; see Miller and Cohen, 2001). In the latter case, network overlap

---

* Correspondence to: M. Najafi, Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742, USA.
** Corresponding author.
E-mail addresses: mahshid@umd.edu (M. Najafi), pessoa@umd.edu (L. Pessoa).

http://dx.doi.org/10.1016/j.neuroimage.2016.04.054

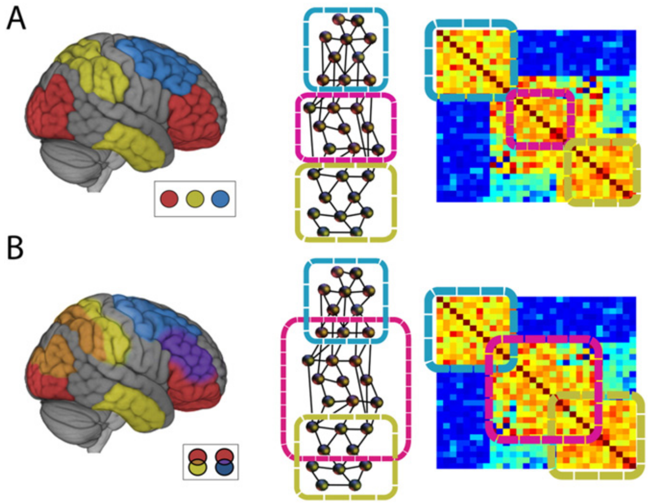

Fig. 1. Community organization in the brain. (A) Standard communities are disjoint (inset: colors indicate separate communities), as illustrated via three types of representation. The representation on the right corresponds to a schematic correlation matrix. (B) Overlapping communities are interdigitated, such that brain regions belong to multiple communities simultaneously (inset: community overlap shown on the brain indicated by intersection colors).

is extensive and many (possibly most) brain regions participate in multiple communities. For example, Yeo et al. (2014) reported that 44% of the vertices studied (away from clustering boundaries) participated in multiple networks.

One reason regions may belong to (or affiliate with) multiple communities is that they may participate in different "region assemblies" depending on task demands. Whereas evidence for regions with "flexible functional connectivity patterns" in frontal and parietal cortex has been recently described (Cole et al., 2013), the extent of such multiple-community participation, and its spatial distribution in the brain remains poorly understood. Moreover, the relationship of flexible affiliation and functional diversity (that is, the spectrum of tasks a region may participate in; see Anderson et al., 2013) is not understood.

At least since the work by Guimera and Amaral (2005a) and Guimera and Amaral, (2005b) there has been increased appreciation that particularly well-connected nodes, often called *hubs*, can be grouped into several distinct sub-types: *provincial* hubs (well-connected nodes with almost all of their links within a single community), *connector* hubs (well-connected nodes with at least half of their links within a community), and *kinless* hubs (hubs with fewer than half of their links within a community). The different hub sub-types are useful for understanding the general functional organization of networks because each of the defining connectivity patterns lends itself to a "universal role" that does not depend on the type of network being studied (social, technological, or biological). By using overlapping communities, hub sub-types can be naturally defined by characterizing each node's *bridgeness* (Nepusz et al., 2008), namely, the ability to participate with multiple communities simultaneously and "bridge" them together. Bridges are important because they have the potential to spread signals across multiple communities, thereby performing important roles in distributed processing.

The goals of the present study were several-fold. First, we sought to characterize the overlapping structure of brain communities during rest by using a state-of-the-art *mixed-membership* algorithm (Gopalan and Blei, 2013; see also Airoldi et al., 2014, 2008; Lancichinetti et al., 2009). In particular, one of our goals was to quantify how much "information" is lost when large-scale networks are treated as disjoint compared to when overlapping structure is characterized. Standard, disjoint clustering assigns membership values of 0 or 1 (all or none), and in so doing dichotomizes measures that may be informative. In contrast, when overlapping structure is determined, a node's participation is assigned across all communities, though with varying strengths; the strengths are summarized by the *membership values* (Fig. 2). Specifically, in the framework adopted, each node has a probability-like membership value associated with each of the existing communities. This community *membership vector* specifies a node's affiliation to all communities considered, with membership values between 0 and 1 (and summing to 1), with entries close to 1 indicating membership to essentially one community, and intermediate values indicating membership to multiple communities.

Second, we sought to investigate the relationship between *functional diversity* and community membership properties. Brain regions differ in terms of their functional diversity, namely, the repertoire of functions that engage them (Passingham et al., 2002; Anderson et al., 2013). Some regions are engaged by a wide variety of tasks (they have high functional diversity), whereas other regions are more narrowly tuned and are engaged by a limited range of paradigms (they have low functional diversity). Here, we asked the following question: Is functional diversity related to how brain regions affiliate with other regions in the absence of a task? In other words, is a region's functional diversity related to its membership values? To estimate a region's functional diversity, we employed the BrainMap database (Laird et al., 2005), which collates activation results across thousands of publications in the literature and organizes them in terms of a task taxonomy spanning perception, action, cognition, and emotion.

Third, we wished to characterize how network structure during rest is potentially altered by task execution. This is important because whereas the large-scale structure of brain networks at rest have been studied extensively, less is known about the organization during task performance. Critically, it is at times assumed that functional connectivity at rest is affected in minor ways by tasks (Cole et al., 2014). In this view, the activity covariation at rest is only mildly influenced by task execution. While some studies have indeed provided evidence in favor of this idea, an alternative proposal is that tasks alter patterns of functional connectivity more substantially (Buckner et al., 2013).

Fourth, our objective was to use the mixed-membership model to measure the "bridgeness" of each region and combine it with other network measures to extend our understanding of "universal roles" and identify key information processing nodes in the brain. Furthermore, we wished to determine how both hub and bridge properties changed during task execution relative to rest.

2. Materials and methods

2.1. Dataset

The structural and functional MRI data for this study were obtained from the Human Connectome Project (HCP; Van Essen et al., 2013) dataset (N = 100) as accessed in June 2014. For completeness, we briefly describe the main aspects of the HCP data (for details, see Glasser et al., 2013; Smith et al., 2013; Van Essen et al., 2013).

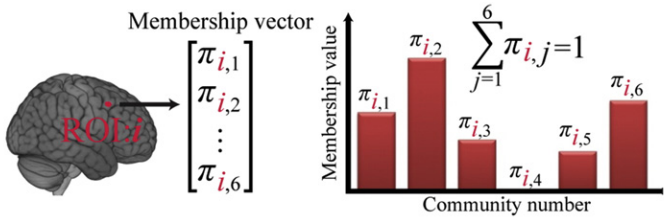

Fig. 2. Overlapping communities and membership values. Each brain region affiliates with each community with varying strengths that are captured by the membership value. These probability-like values are between 0 and 1 and sum to 1 (for each region of interest, or ROI).