NeuroImage 101 (2014) 473–484

Contents lists available at ScienceDirect

NeuroImage

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

Stochastic geometric network models for groups of functional and structural connectomes

Eric J. Friedman [a,b,*], Adam S. Landsberg [c], Julia P. Owen [d,e], Yi-Ou Li [d,e], Pratik Mukherjee [d,e]

[a] International Computer Science Institute, Berkeley, USA
[b] Department of Computer Science, University of California, Berkeley, USA
[c] W.M. Keck Science Department, Claremont McKenna College, Pitzer College, and Scripps College, Claremont,USA
[d] Department of Radiology & Biomedical Imaging, University of California, San Francisco, USA
[e] Department of Bioengineering & Therapeutic Sciences, University of California, San Francisco, USA

ARTICLE INFO

Article history:
Accepted 17 July 2014
Available online 25 July 2014

Keywords:
Network
Graph analysis
Connectome
fMRI
dMRI
Geometric networks

ABSTRACT

Structural and functional connectomes are emerging as important instruments in the study of normal brain function and in the development of new biomarkers for a variety of brain disorders. In contrast to single-network studies that presently dominate the (non-connectome) network literature, connectome analyses typically examine groups of empirical networks and then compare these against standard (stochastic) network models. The current practice in connectome studies is to employ stochastic network models derived from social science and engineering contexts as the basis for the comparison. However, these are not necessarily best suited for the analysis of connectomes, which often contain groups of very closely related networks, such as occurs with a set of controls or a set of patients with a specific disorder. This paper studies important extensions of standard stochastic models that make them better adapted for analysis of connectomes, and develops new statistical fitting methodologies that account for inter-subject variations. The extensions explicitly incorporate geometric information about a network based on distances and inter/intra hemispherical asymmetries (to supplement ordinary degree-distribution information), and utilize a stochastic choice of network density levels (for fixed threshold networks) to better capture the variance in average connectivity among subjects. The new statistical tools introduced here allow one to compare groups of networks by matching *both* their average characteristics and the variations among them. A notable finding is that connectomes have high "smallworldness" *beyond that* arising from geometric and degree considerations alone.

Introduction

The study of the empirical brain networks has taken great strides in recent years, allowing analysis of the brain "system" with its complex interconnections. The construction of the brain networks, or connectomes, from clinical MR data is becoming commonly available and is providing both deep insights into the functioning of the human brain and also into the differences between normal and abnormal (diseased or injured) brains (Bullmore and Sporns, 2009; Sporns, 2011).

The foundations of these approaches have been largely based on techniques developed in the social sciences and engineering, in particular for networks of people or computer networks (Albert et al., 1999; Jackson, 2010; Watts, 2004), as well as applications in biological and biochemical networks (Jeong et al., 2001). In these settings one typically has only a single, very large network (or several related but fundamentally different networks) to analyze, which has led to the development of very powerful approaches in those settings (Newman, 2003).

However, the study of groups of brain networks requires different tools. Here, one often has groups of closely related networks wherein although the exact edges may differ from subject to subject, nonetheless the number and basic attributes of nodes remain comparable between subjects. This comparability of nodes between different networks allows for a variety of new types of analyses and models, including the construction of detailed *geometric properties* of the network.

This consideration also allows one to view a group of networks from a distributional sense — for example, one can ask what is the distribution of networks for a population of subjects of a certain type, such as controls or those with a specific disorder or injury. In many instances, understanding the entire distribution is in fact crucial, as simple averages may sometimes conceal critical information (as noted but not formalized in Simpson et al., 2012). For instance, a recent paper analyzing structural connectomes in subjects with agenesis of the corpus callosum (AgCC) revealed that a key difference between the AgCC subjects and

* Corresponding author at: International Computer Science Institute, 1947 Center St. #600, Berkeley, CA 94704, USA.
E-mail address: ejf@icsi.berkeley.edu (E.J. Friedman).

http://dx.doi.org/10.1016/j.neuroimage.2014.07.039
1053-8119/

the controls was that the AgCC patients exhibited higher inter-subject variability in their networks (Owen et al., 2012).

In order to understand these distributions of networks, an underlying stochastic network model is commonly assumed in brain network studies. The choice of underlying model figures implicitly in the design of network measures. For example, computations of modularity and the clustering of nodes in a connectome typically employ a definition of "modularity" that is inherently based on the assumption of an underlying Degree-Distributed stochastic network, since it "weights" edges based on the degree of the nodes that it connects (e.g. If two nodes are both of high degree then an edge between them is not as "informative" as an edge between two low-degree nodes, which is in some sense less likely to arise by chance (Girvan and Newman, 2002)). The choice of underlying model also figures prominently in computing the significance of a network measure. For example, the "smallworldness" of a network is often compared to the smallworldness of a matched random network (Sporns and Zwi, 2004). The choice of such a comparison network can prove to be crucial. For example, for resting state fMRI networks, the smallworldness of the two most popular random network models – the Erdos–Renyi model and the Degree Distributed random model – typically differ by a factor of 2 on empirical brain networks (Newman et al. 2001,Newman, 2009a). Alternatives include choosing the average or median consensus network or a single representative one (Simpson et al., 2011). As well discussed in Simpson et al. (2012) there are many more examples exposing the importance of the underlying model network, ranging from their use as null networks as discussed above, to modularity analyses (Joyce et al., 2010; Meunier et al., 2009a,b; Valencia et al., 2009), to representing an individual's network based on several experimental runs (Zuo et al., 2011), to visualization tools (Song et al., 2009; Zuo et al., 2011), to their ability to assess a group of networks (Achard et al., 2006), to identifying hub/node types (Joyce et al., 2010), and to constructing representative networks for brain dynamics studies (Jirsa et al., 2010). Additional examples for modularity include Expert et al. (2011), Bassett et al. (2013) and Henderson and Robinson (2013).

The goal of modeling a group of networks, as in this paper, does affect our choices for analysis. The goal is to have a stochastic model that generates networks that "fits the entire group of networks" and is constructed to match basic network properties, such as degree distribution or geometry. This is in contrast to random network models that try to fit network measures, such as implemented by Vértes et al. (2012) or Simpson et al. (2011) which consider classes of random networks and then fit them to the empirical measures — while these can provide deep insights into the structure of the empirical networks, they do not provide simple intuitive models for comparison.

To see this point more clearly, consider the work by Vértes et al. (2012) which considers similar stochastic network models to those in this paper. While there are important differences in the models such as the use of preferential attachment terms in their models and inter/intra-hemispheric terms in ours, the differences in implementation are more significant. They choose important parameters in their models (such as those for the distance and preferential attachment terms) by maximizing an energy function that tries to match the *mean* of the subjects' global network measures (efficiency, clustering and modularity) to those of the stochastic networks. (Note that their models use the variability in network measures to scale the energy function but do not match the variability of these measures, as we do below.) This generates a network that fits the means of the data closely, but because of the complex nonlinear interactions between the parameters and the network measures can lead to networks with different parameters than would have been attained by directly fitting to the baseline network information. For example, the distribution of edge lengths in their models often differ significantly from the empirical distribution as seen clearly in their figures.

To see why this arises, consider a simple stochastic network model in which the probability of an edge between two nodes is given by a function of the distance between those two nodes. Clearly one could find a distance function that differs significantly from the true one that yields the same specified clustering coefficient. Similarly, one could likely fit the random networks to exactly match the small-worldness of the empirical networks, but then one cannot discuss the *excess* small-worldness (as we do later in this paper). In addition, standard random network modeling, such as that typically used for degree distributions (wherein each empirical network is individually matched to one or several random networks with the exact same degree distribution) differs significantly from our approach as this is in some sense over-fitting and only generates networks that have the exact degree distribution of one of the empirical networks (Newman et al. 2001, Newman, 2009a), while one would expect that a new subject would not exactly match the degree distribution of one of the existing networks. (One could see this statistically using standard cross-validation techniques, such as the well known leave-one-out cross-validation.)

Another key difference between connectomes and most traditional network models is that nodes in connectomes have a physical location. This is extremely important as connections between different areas of the brain definitively depend upon relative location, particularly the distance between various regions (Kaiser and Hilgetag, 2004a,b; Scannell et al., 1999; Sporns et al., 2004). As we will show, the use of such geometric information appears to be important in the development of good generative models, as was suggested by Expert et al. (2011), applied in Vértes et al. (2012) and motivated by the analysis in Alexander-Bloch et al. (2013). Fig. 1 (which will be explained more fully later), previews the various stochastic network models (both traditional ones and newer ones incorporating geometric information) that will be considered and compared in this paper.

An additional important aspect in the study of connectomes is the choice of threshold type and value, as both fMRI and dMRI generate continuous valued matrices that are then thresholded to create a binary matrix representing the network, where the network density is determined by the threshold value which can be chosen for fixed density (every network has exactly the same density) or variable density (every network uses the same threshold). While we do not directly analyze the optimal choice of threshold type and value (if there indeed is one; see van Wijk et al., 2010), we do consider the effects of such a threshold on the distribution of generated networks. Note that, as discussed in van Wijk et al. (2010), while fixed threshold networks may be superior in certain settings to fixed density networks, they are also more difficult to analyze due to the effects of the density variations on network measures; however we believe that the use of appropriate null networks can mitigate these difficulties. We also note that recent work has also considered using weights directly in the network analysis and not thresholding the data (e.g., Liu et al., 2013; Rubinov and Sporns, 2010). Alternatively, one can treat "multiple thresholds simultaneously" (Bassett et al., 2013; Ginestet et al., 2011).

One important but unrecognized consequence of applying a *fixed* threshold to all the empirical networks in the group is that it leads to wide variations in their densities, which, as we show later, well exceed those appearing in standard stochastic network models. Accordingly, we will demonstrate that in order to effectively capture the variability found in real connectome studies, one needs to allow density to vary in the underlying stochastic network models (see Fig. 2).

In this paper, we take a principled empirical approach towards these issues. We compare a variety of stochastic network models on both functional and structural brain networks to understand which of the standard network metrics are well captured and which are not and how to design models that better capture these properties of connectome data. In addition, we extend some traditional statistical methods so as to quantify and illuminate the *variation* in