



Characterising group-level brain connectivity: A framework using Bayesian exponential random graph models

B.C.L. Lehmann^{a,b,c,*}, R.N. Henson^d, L. Geerligs^e, Cam-CAN^{f,#}, S.R. White^{a,g}

^a MRC Biostatistics Unit, University of Cambridge, UK

^b Big Data Institute, University of Oxford, UK

^c Department of Statistics, University of Oxford, UK

^d MRC Cognition and Brain Sciences Unit, University of Cambridge, UK

^e Donders Institute for Brain, Cognition and Behaviour, Radboud University, UK

^f Cambridge Centre for Ageing and Neuroscience (Cam-CAN), University of Cambridge and MRC Cognition and Brain Sciences Unit, Cambridge, UK

^g Department of Psychiatry, University of Cambridge, UK

ARTICLE INFO

Keywords:

Exponential Random Graph Model (ERGM)

Bayesian ERGM

Group studies

Network neuroscience

Fmri

ABSTRACT

The brain can be modelled as a network with nodes and edges derived from a range of imaging modalities: the nodes correspond to spatially distinct regions and the edges to the interactions between them. Whole-brain connectivity studies typically seek to determine how network properties change with a given categorical phenotype such as age-group, disease condition or mental state. To do so reliably, it is necessary to determine the features of the connectivity structure that are common across a group of brain scans. Given the complex interdependencies inherent in network data, this is not a straightforward task. Some studies construct a group-representative network (GRN), ignoring individual differences, while other studies analyse networks for each individual independently, ignoring information that is shared across individuals. We propose a Bayesian framework based on exponential random graph models (ERGM) extended to multiple networks to characterise the distribution of an entire population of networks. Using resting-state fMRI data from the Cam-CAN project, a study on healthy ageing, we demonstrate how our method can be used to characterise and compare the brain's functional connectivity structure across a group of young individuals and a group of old individuals.

1. Introduction

Brain connectivity analysis aims to understand how spatially distinct regions of the brain interact with each other. The use of networks to model whole-brain connectivity has become increasingly popular in recent years (Bullmore and Sporns, 2009): by treating distinct regions as nodes and the connections between them as edges, researchers have gained new insights into both the structure and function of the brain. To determine the salient features of the brain's connectivity structure, it is necessary to identify which of those are common across individuals. Given the complex interdependencies inherent in network data, however, it is not a trivial task to consider connectivity structure across multiple individuals; how best to combine connectivity information across participants remains a key challenge (Simpson et al., 2013b).

Many existing methods aim to construct a (single) group-representative network (GRN) across (multiple) individuals. There have been several methods proposed for constructing GRNs. For example, Achard et al. (2006) constructed a group-representative network using

mean functional connectivity (mean-GRN) by first taking the mean of the individuals' functional connectivity matrices and then thresholding the resulting matrix. Song et al. (2009) took a similar approach using the median of the individuals' functional connectivity matrices (i.e. median-GRN). Sinke et al. (2016) constructed a group-representative structural connectivity network from Diffusion Tensor Imaging (DTI) data by keeping those edges which are present in at least 35% of the individuals' networks (i.e. minimal-GRN). These edge-based GRN methods are computationally convenient, since each individual is processed separately and the subsequent analysis is based on a single network. However, these methods ignore higher-order topological properties present in each individual's network. Rather than construct a completely new network as a combination of each individual's, other methods represent a whole group by selecting a 'best' individual network (i.e. best-GRN) in terms of mutual information (Meunier et al., 2009) or Jaccard index (Joyce et al., 2010). While this approach preserves the topological properties of a single network, in subsequent analyses it may lead to undue weight being placed on features that are not present in many of the remaining networks.

* Corresponding author.

E-mail address: brieuc.lehmann@bdi.ox.ac.uk (B.C.L. Lehmann).

www.cam-can.org

The above issues are compounded when comparing brain networks across groups of individuals, identifying group differences in connectivity via differences in the GRNs from each group. In particular, network metrics are influenced by the overall density of the network (Ginestet et al., 2011). As a result, it is difficult to disentangle differences in more complex topological properties such as clustering from differences in the network density; especially in the context of the various GRN methods. This is particularly important in the context of ageing since mean functional connectivity is known to decrease with age (Geertjigs et al., 2017). Note that it is not trivial to simply control for network density when comparing network metrics due to their highly non-linear relationship.

A key consideration is whether to view the network as a single object or a single realisation of a random process. The former implies that the network is a fixed construct, the latter defines a random process (with inherent noise) that gives rise to a distribution of brain networks. An exponential random graph model (ERGM) defines a parametric statistical distribution across all possible networks with a given set of nodes (see Robins et al. (2007b) for a review). The aim of the model is to characterise the distribution of a network in terms of a set of *summary statistics*. These summary statistics are typically comprised of topological features of the network, such as the number of edges and subgraph counts. The summary statistics enter the likelihood via a weighted sum; the weights are (unknown) model parameters that quantify the relative influence of the corresponding summary statistic on the overall network structure and must be inferred from the data. ERGMs are thus a flexible way in which to describe the global network structure as a function of network summary statistics.

ERGMs have been applied successfully to characterise both functional connectivity (Dell'italia et al., 2018; Obando and Fallani, 2017; Simpson et al., 2011, 2012) and structural connectivity (Sinke et al., 2016). To date, there have been two proposed approaches for using ERGMs in group studies. The first approach constructs an edge-based GRN for each group and then fits the same ERGM (i.e. identical summary statistics) separately to each group's network, obtaining GRN-ERGM parameter estimates (Sinke et al., 2016). As described above, this edge-based approach may average over informative topological structure present in each of the individuals' networks. In contrast, the second approach fits an ERGM to each individuals' network and then takes the mean or median of the fitted parameters from individuals within a group to represent the group-level connectivity structure (Obando and Fallani, 2017; Simpson et al., 2012; i.e. mean-ERGM or median-ERGM). Note that it is then possible to generate a GRN from the resulting mean-ERGM and median-ERGM parameter estimates (Simpson et al., 2012). While this is preferable to the first approach, taking a simple mean or median may obscure important information.

While ERGMs are designed to represent network data in terms of network summary statistics, they are limited in their ability to study specific edges. A complementary approach for analysing brain network data is the mixed effects framework. This models the probability of a connection in terms of both network metrics and covariates (Simpson and Laurienti, 2015). Although both approaches incorporate network metrics into the models, the associated parameters have subtly different interpretations. The outcome variable in an ERGM is the entire network and the parameters quantify the effect of the network metrics on the global network structure. In contrast, the outcome variables in the mixed effects model are the individual network dyads and the parameters quantify the relationship between the network metrics and the probability of a connection. Note that, by including group membership as a covariate, it is straightforward to use the mixed effects approach for group studies. The aim of this work is to similarly expand the use of ERGMs for group studies.

In this article, we introduce a novel framework to model a population of individual brain networks which we call multi-BERGM (the 'B' denotes Bayesian; the framework could theoretically be analysed in a frequentist paradigm but we do not consider that further). Our

framework uses a Bayesian formulation of the ERGM within a multi-level (i.e. hierarchical) model. Bayesian ERGMs have previously been used in the context of a group study to characterise the distribution of a group-representative network (Sinke et al., 2016). In contrast, our method fits a Bayesian ERGM to each individual network and pools information across multiple individuals nested within multiple groups via the multilevel structure of the overarching model. This provides an approach to characterise connectivity at a group-level, and thus compare the brain's connectivity structure between groups. We demonstrate our method on functional connectivity networks derived from resting-state fMRI scans of participants from the Cam-CAN project, a study on healthy ageing (Shafto et al., 2014).

2. Methods

2.1. Data

The data were collected as part of Phase II of the CamCAN project (www.cam-can.org; Shafto et al., 2014). The MRI data were acquired on a 3T Siemens TIM Trio at the MRC Cognition & Brain Sciences Unit, with a 32 channel head-coil. Structural images were acquired using a 1mm³ isotropic, T1-weighted Magnetization Prepared Rapid Gradient Echo (MPRAGE) sequence and a 1mm³ isotropic, T2-weighted Sampling Perfection with Application optimized Contrasts using different flip angle Evolution (SPACE) sequence. The fMRI data for the eyes-closed, resting-state run were acquired using a Gradient-Echo (GE) Echo-Planar Imaging (EPI) sequence, and consisted of 261 volumes (lasting 8 min and 40 s). Each volume contained 32 axial slices (acquired in descending order), with slice thickness of 3.7 mm and interslice gap of 20% (for whole brain coverage including cerebellum; TR 1,970 ms; TE 30ms; voxel-size 3 mm 3 mm 4.44 mm). EPI fieldmaps with two TEs (5.19 ms and 7.65 ms) were also acquired. All the raw data from CamCAN Phase II, together with more acquisition details, are available on: <http://www.cam-can.org/index.php?content=dataset>.

The data were processed using r7219 of the SPM12 software (<http://www.fil.ion.ucl.ac.uk/spm>), automated with release r5.4 of the Automatic Analysis (AA) pipeline system¹ (Cusack et al., 2015; see Taylor et al. (2017) for an overview of the pipelines) in r2015 of MATLAB (The MathWorks). To obtain a good starting-point for image normalisation, the T1 image was coregistered to the Montreal Neurological Institute (MNI) template using rigid body transformation, and then the T2 image was coregistered to the T1. Both T1 and T2 images were bias corrected, and then combined in a multimodal segmentation to estimate images of each of six tissue classes, including grey-matter (GM), white-matter (WM) and cerebrospinal fluid (CSF). Diffeomorphic registration (DARTEL) was then applied to the GM and WM segments to create group templates, which were in turn transformed to MNI space using a 12-parameter affine transform.

The fMRI images were unwarped using distortion fields estimated from the fieldmaps, and corrected for motion using rigid-body realignment to the mean fMRI image across runs. The different slice acquisition times were corrected by interpolating to the middle slice. The images were rigid-body coregistered to the T1 image and the spatial transformations from that T1 to MNI space (diffeomorphic and affine) were applied to every fMRI image. Residual effects of abrupt motion were reduced by applying wavelet despiking (Patel et al., 2014). The mean time series for all voxels within the thresholded WM and CSF segments were calculated to use as later covariates of no interest.

The fMRI time series were then extracted from 90 cortical and subcortical regions of interest (ROIs) from the AAL atlas (Tzourio-Mazoyer et al., 2002). ROIs in the cerebellum were not included. The time series for each voxel in each ROI were adjusted for various con-

¹ Available at <http://github.com/automaticanalysis/automaticanalysis/releases/v5.4.0>.

finds by taking the residuals from a general linear model (GLM) that included: (1) the time series in WM and CSF segments, (2) the 6 rigid-body motion parameters from the realignment stage, (3) the first-order difference in those motion parameters across successive TRs and (4) a cosine basis set up to cut-off frequency of 0.008 Hz (implementing a form of high-pass filter). Second-order (squared) terms for (1), (2) and (3) were also included in the GLM. The autocorrelation in the GLM error was modelled by a family of 8 exponentials with half-lives from 0.5 to 64 TRs (estimated by pooling across voxels within each ROI), and used to prewhiten the time series. The average voxels of the residual time series was then used as a summary measure for each ROI. This approach was based on the optimised pipeline proposed by Geerligs et al. (2017).

To illustrate our method, we analysed two groups (denoted j , i.e. $j = 1, 2$) from the Cam-CAN study: the 100 youngest individuals, aged 18–33, and the 100 oldest individuals, aged 74–89 (denoted i within each group, i.e. $i = 1, \dots, n_j$, where $n_1 = n_2 = 100$).

2.2. Network construction

The preprocessed data consists of $N = 90$ ROI time series for each individual. To apply a standard (binary) ERGM, it is necessary to extract a network for each individual. For a given individual i , we first computed the pairwise Pearson correlation between each of the time series, yielding a $N \times N$ correlation matrix $\mathbf{C}^{(i)}$. In contrast with the pipeline proposed by Geerligs et al. (2017), we did not apply mean regression to the resulting functional connectivity matrices because we accounted for differences in mean connectivity via the thresholding procedure or through the specification of the ERGM (we return to this point in the discussion).

To each correlation matrix, we then applied a threshold $r^{(i)}$ to produce an $N \times N$ adjacency matrix, $\mathbf{A}^{(i)}$, with entries:

$$\mathbf{A}_{kl}^{(i)} = \begin{cases} 1 & \text{if } \mathbf{C}_{kl}^{(i)} \geq r^{(i)} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The adjacency matrix defines an individual's network, $\mathbf{y}^{(i)}$, with an edge between nodes k and l if and only if $\mathbf{A}_{kl}^{(i)} = 1$. The choice of threshold $r^{(i)}$ can have a significant impact on the resulting connectivity structure (van den Heuvel et al., 2017; van Wijk et al., 2010). The two most common strategies are “absolute” thresholding and “proportional” thresholding. Absolute thresholding picks a common threshold $r^{(i)} = r$ for each individual. In contrast, proportional thresholding ensures that the same number of edges are present in each network by allowing the threshold to vary by individual.

Both thresholding strategies have their drawbacks. Since the number of edges in a network inherently affects the overall network (Ginestet et al., 2014; van Wijk et al., 2010), differences in other network metrics may simply be attributable to variations in the overall connectivity. A proportional threshold may thus be preferred in order to keep the number of edges constant across individuals (in both groups) and facilitate comparison of other metrics of interest. On the other hand, a lower correlation value may be less reliable in indicating a functional connection between brain regions. Therefore, by including lower correlations as edges for individuals with lower overall connectivity, a proportional threshold may induce more randomness in the resulting network (van den Heuvel et al., 2017).

Given these issues, we analysed networks constructed using the two different thresholding procedures: absolute and proportional. Further, for both procedures we considered two distinct threshold values (i.e. different $r^{(i)}$), chosen to yield average node degrees of $K = 3$ and $K = 5$. The value of $K = 3$ corresponds to the efficiency-cost optimization (ECO) criterion (Fallani et al., 2017); this aims to optimise the trade-off between the efficiency of a network and its density, or wiring cost. The value of $K = 5$ has been used in previous applications of ERGMs to resting-state functional connectivity (Simpson et al., 2011, 2012).

2.2.1. Group-representative network construction

In order to contrast our framework with existing methods, we also constructed group-representative networks (GRN) following Achard et al. (2006) (mean-GRN) and Song et al. (2009) (median-GRN). The mean-GRN and median-GRN take the mean and median, respectively, of the correlation matrices across all individuals within the group (young and old) and then threshold these group mean and median matrices to yield networks with average node degrees of $K = 3$ and $K = 5$. Explicitly, we have $\mathbf{C}_{kl}^{\text{mean}} = \frac{1}{n} \sum_i \mathbf{C}_{kl}^{(i)}$ and $\mathbf{C}_{kl}^{\text{median}} = \text{median}(\{\mathbf{C}_{kl}^{(i)} : \forall i\})$, with associated adjacency matrices. These approaches produce a single network for each group, which can be modelled as an ERGM and we can assess the respective GRN-ERGM parameters.

2.3. Exponential random graph model specification

The probability mass function of a network \mathbf{Y} under an ERGM is given by $\pi(\mathbf{y}|\theta)$ where

$$\pi(\mathbf{y}|\theta) = \frac{\exp\{\theta^T s(\mathbf{y})\}}{Z(\theta)}. \quad (2)$$

Here, $s(\mathbf{y})$ is a vector of p network summary statistics, $\theta \in \Theta \subseteq \mathbb{R}^p$ is a vector of p corresponding model parameters that must be estimated from the data and $Z(\theta) = \sum_{\mathbf{y} \in \mathcal{Y}} \exp\{\theta^T s(\mathbf{y})\}$ is the normalising constant ensuring the probability mass function sums to one. Given data, that is, an observation \mathbf{y} of \mathbf{Y} , the goal is to infer which values of θ best correspond to the data under this distribution. The summary statistics included in a given ERGM represent those configurations expected to appear more frequently or less frequently than in a random graph. In other words, the choice of summary statistics is a modelling decision; it reflects our belief of how the global network structure may be summarised and is driven by the context of the network. The flexibility of ERGMs derives from the range and number of possible summary statistics that can be used (Robins et al., 2007a).

In what follows, we work in the Bayesian paradigm, treating the model parameters θ as random variables. The Bayesian formulation of ERGMs was first suggested by Koskinen (2004) and then expanded upon by Caimo and Friel (2011). Through the machinery of Bayesian hierarchical modelling, this provides a natural framework for describing the joint distribution of a group of networks. To fully specify a Bayesian ERGM, we need only augment the definition in (2) with a prior distribution $\pi(\theta)$ for the model parameters. Given an observation \mathbf{y} of the network, we can then perform inference by analysing the posterior distribution $\pi(\theta|\mathbf{y})$. Through Bayes' rule, we may write the posterior as

$$\begin{aligned} \pi(\theta|\mathbf{y}) &= \frac{\pi(\mathbf{y}|\theta)\pi(\theta)}{\pi(\mathbf{y})} \\ &= \frac{\exp\{\theta^T s(\mathbf{y})\}\pi(\theta)}{Z(\theta)\pi(\mathbf{y})} \end{aligned} \quad (3)$$

where $\pi(\mathbf{y}) = \int_{\Theta} \pi(\mathbf{y}|\theta)\pi(\theta)d\theta$ is the model evidence. Although the posterior distribution is generally not available in a closed-form expression, one may obtain samples from the posterior through a Markov chain Monte Carlo method known as the exchange algorithm (Murray et al., 2006).

Under absolute thresholding, our exponential random graph model (which is common to all individuals and groups) uses three summary statistics that have previously been applied to single brain functional connectivity networks (Simpson et al., 2013a, 2011). The three summary statistics used are: number of edges (E); geometrically weighted edgewise shared partners (GWESP); and geometrically weighted non-edgewise shared partners (GWNSP). For the proportional thresholding, we fixed the number of edges and modelled the network using two summary statistics: GWESP and GWNSP.

The number of edges $E = \sum_{k < l} Y_{kl}$ characterises the sparsity of the network. When using an absolute threshold, it is important to include the number of edges in the model in order to account for differences in the network density across individuals. When using a proportional threshold, however, each individual network is constructed to have a

predetermined density. As a result, the number of edges E is constrained to be the same across each network. In this case, we therefore opted not to include the summary statistic in the model, constraining the space of networks to those with the same number of edges as those observed. In other words, the model describes a probability distribution over networks with a given average node degree K . Note that this in contrast to previous applications of ERGMs in neuroimaging, which include the number of edges in the model despite using a proportional threshold (Simpson et al., 2011, 2012).

The geometrically weighted edgewise shared partner (GWESP) statistic of a network \mathbf{y} is a measure of network transitivity (or clustering) and is given by:

$$GWESP(\mathbf{y}) = e^{\tau_1} \sum_{w=1}^{V-2} \{1 - (1 - e^{-\tau_1})^w\} EP_w(\mathbf{y}) \quad (4)$$

where $EP_w(\mathbf{y})$ is the number of *connected* node pairs having exactly w shared partners and $\tau_1 > 0$ is a decay parameter. The decay parameter τ_1 serves to diminish the effect of the network having more higher-order edgewise shared partners relative to lower-order edgewise shared partners. In other words, the increase in GWESP from adding a single edge is smaller if the edge adds a shared partner to connected nodes that already share many partners.

The geometrically weighted non-edgewise shared partner (GWNSP) statistic is similarly defined as

$$GWNSP(\mathbf{y}) = e^{\tau_2} \sum_{w=1}^{V-2} \{1 - (1 - e^{-\tau_2})^w\} NP_w(\mathbf{y}) \quad (5)$$

where $NP_w(\mathbf{y})$ is the number of *non-connected* node pairs having exactly w shared partners and $\tau_2 > 0$ is a decay parameter. GWNSP is related to global network efficiency; a higher value of GWNSP indicates that non-connected nodes are more likely to have a shared partner.

While it is possible to treat the decay parameters for both geometrically weighted statistics as extra model parameters (leading to *curved* ERGMs; see Hunter, 2007), this increases the computational burden substantially. The decay parameters were therefore fixed at $\tau_1 = \tau_2 = 0.75$, as these values have previously been found to result in better fitting models for networks constructed from the AAL atlas (Simpson et al., 2011).

2.4. Multi-level (hierarchical) framework

The Bayesian exponential random graph model described above provides a flexible family of distributions for a *single* network. Our goal is to extend this to a model for a *population* of networks. Our proposed approach is simple: represent each network as a separate ERGM within a Bayesian multilevel (or hierarchical) model. By pooling information across individual networks, this approach allows us to characterise the distribution of the whole population.

We model each individual network $\mathbf{Y}^{(i)}$ as an exponential random graph with model parameter $\theta^{(i)}$. Importantly, each individual ERGM must consist of the same set of summary statistics $s(\cdot)$. The probability mass function of each network can then be written

$$\pi(\mathbf{y}^{(i)}|\theta^{(i)}) = \frac{\exp\{\theta^{(i)T} s(\mathbf{y}^{(i)})\}}{Z(\theta^{(i)})}, \quad i = 1, \dots, n. \quad (6)$$

This specifies the data-generating process for each individual network. To obtain a joint distribution for the set of networks, we assume that, conditional on their respective individual-level parameters, the $\mathbf{Y}^{(i)}$ are independent. Thus, the sampling distribution for the set of networks \mathbf{Y} is simply the product of the individual probability mass functions:

$$\pi(\mathbf{y}|\boldsymbol{\theta}) = \frac{\prod_{i=1}^n \pi(\mathbf{y}^{(i)}|\theta^{(i)})}{\exp\{\sum_{i=1}^n \theta^{(i)T} s(\mathbf{y}^{(i)})\}} = \frac{\prod_{i=1}^n \pi(\mathbf{y}^{(i)}|\theta^{(i)})}{\prod_{i=1}^n Z(\theta^{(i)})}. \quad (7)$$




Fig. 1. A diagrammatic representation of the hierarchical framework. Each network $\mathbf{Y}^{(i)}$ is modelled as an exponential random graph with individual-level parameter $\theta^{(i)}$. In turn, each θ_i is assumed to come from a common population-level distribution with hyperparameter ϕ .

2.4.1. First level of the multi-level model (within group)

To model a group of networks we need to specify the prior distribution of the individual-level ERGM parameters $\theta^{(1)}, \dots, \theta^{(n_j)}$ (for individuals $1, \dots, n_j$ in group j). To this end, we propose a multilevel model such that $\theta^{(1)}, \dots, \theta^{(n)}$ are drawn from a common population-level Normal distribution with parameters $\phi = (\mu_{\text{multi}}, \Sigma_{\theta})$, which are also treated as random variables. We write

$$\theta^{(i)} \sim \pi(\cdot|\phi), \quad i = 1, \dots, n_j \quad (8)$$

for the population-level distribution. Assuming that, conditional on ϕ , the $\theta^{(i)}$ are independent, we have

$$\pi(\boldsymbol{\theta}|\phi) = \prod_{i=1}^{n_j} \pi(\theta^{(i)}|\phi). \quad (9)$$

Finally, write $\pi(\phi)$ for the (hyper)prior distribution of ϕ . The joint distribution of $(\mathbf{Y}, \boldsymbol{\theta}, \phi)$ can be written as $\pi(\mathbf{y}, \boldsymbol{\theta}, \phi) = \pi(\mathbf{y}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|\phi)\pi(\phi)$. See Fig. 1 for a diagrammatic representation of the full hierarchical framework. Full details of the prior specifications can be found in the Supplementary Material.

2.4.2. Second level of the multi-level model (between group)

To extend the model for multiple groups of within a population, we add another level to the hierarchy. Denoting $\theta^{(i,j)}$ to be the (vector-valued) ERGM parameter for the i^{th} individual in the j^{th} group, we assume $\theta^{(i,j)} \sim \mathcal{N}(\mu_{\text{multi}}^{(j)}, \Sigma_{\theta})$. Thus, we assume different group-level means but the same covariance structure across groups (our framework allows per-group covariance, but the sample size of our illustrative example leads us to a common covariance). Writing $\phi^{(j)} = (\mu_{\text{multi}}^{(j)}, \Sigma_{\theta})$, we then complete the model by specifying the group-level prior $\phi^{(j)} \sim \pi(\cdot|\phi^{\text{pop}})$ and population-level hyperprior $\pi(\phi^{\text{pop}})$. Full details of the prior specifications can be found in the Supplementary Material.

2.5. Model fitting and comparison to alternative approaches

To generate samples from the joint posterior distribution, we devised a novel MCMC algorithm, the details of which can be found in Lehmann (2019). We used our algorithm to generate posterior samples

for two sets of networks. The first set consisted of a single group of networks from the 100 youngest individuals in the Cam-CAN study (a one-level model), while the second set also included the networks from the 100 oldest individuals (a two-level model).

2.5.1. One-level model for young group

For the single-group dataset, we derived estimates for the posterior mean and 95% credible regions for each pair of components of the group-level parameter μ_{multi} using our multi-BERGM model with one-level (i.e. individuals nested within one group).

In order to compare our method with existing approaches, we also generated posterior samples for the same (using three summary statistics defined earlier) Bayesian ERGM fit to the mean-GRN, $\mu_{\text{mean-GRN}}$, and median-GRN, $\mu_{\text{median-GRN}}$.

We also compare our method, which generates a full posterior distribution for μ and Σ , to a mean-BERGM approach. Specifically, we fit a BERGM to each individual within the group (i.e. we fit $n_1 = 100$ separate single-BERGMs and obtain 100 separate sets of posterior samples $\{\theta_t^{(i)}\}_{t=1}^T$ for $i = 1, \dots, 100$). The original proposal of this approach used frequentist ERGMs, giving a point estimate per individual (Simpson et al., 2011). The Bayesian approach we use yields a separate posterior per individual, which can then be combined into a group-level estimate $\mu_{\text{mean-BERGM}} \sim \mathcal{N}(\bar{\mu}, \hat{\Sigma})$ with mean

$$\bar{\mu} = \frac{1}{n} \sum_{i=1}^n \bar{\theta}^{(i)}$$

and covariance

$$\hat{\Sigma} = \frac{1}{T-1} \sum_{t=1}^T (\mu_t - \bar{\mu})^2 + \frac{1}{T} \sum_{t=1}^T \left[\frac{1}{n(n-1)} \sum_{i=1}^n (\theta_t^{(i)} - \mu_t)^2 \right].$$

where $\bar{\theta}^{(i)} = \frac{1}{T} \sum_{t=1}^T \theta_t^{(i)}$ and $\mu_t = \frac{1}{n} \sum_{i=1}^n \theta_t^{(i)}$.

2.5.2. Two-level model for young and old groups

For the two-group dataset, we used the two-level multi-BERGM model and derived posterior density estimates for the group-level mean parameters $\mu_{\text{multi}}^{(1)}, \mu_{\text{multi}}^{(2)}$, thus providing a way to compare the connectivity structure across the two groups. We also applied the mean-BERGM approach to the group of old individuals, i.e. we fit a single-BERGM to each individual within the old group and used the resulting posterior samples to construct a group-level estimate (see previous section).

2.6. Goodness-of-fit assessment

For the single-group dataset, we assessed the goodness-of-fit of the posterior distribution to the data by simulating networks from the poste-

rior predictive distribution and comparing the network metrics of these simulated networks to those of the observed networks. Specifically, we chose uniformly at random $S = 1000$ samples μ_1, \dots, μ_S generated from the posterior distribution of the group-level parameter μ_{multi} . For each sample, s , we also randomly selected (uniformly, with replacement) an individual i_s and then simulated a network $Y^{(s)}$ from the ERGM with parameter $\mu_s + \theta_s^{(i_s)}$. We compared the simulated networks to the observed networks based on three network metrics: degree distribution, geodesic distance distribution (length of shortest paths) and edge-wise shared partners distribution (note these are different to, but highly correlated with, the summary statistics used to define the ERGM).

It is possible to assess goodness-of-fit on any network metrics by comparing simulated networks to observed networks. Since age-related differences in local efficiency (Achard and Bullmore, 2007; Geertligs et al., 2015; Sala-Llonch et al., 2014; Song et al., 2014) and global efficiency (Achard and Bullmore, 2007; Sala-Llonch et al., 2014) have previously been observed, we assessed the goodness-of-fit for the two-group dataset in terms of these two metrics. For both groups, we chose uniformly at random $S = 1000$ samples $\mu_1^{(j)}, \dots, \mu_S^{(j)}$ generated from the posterior distribution of the group-level parameter $\mu_{\text{multi}}^{(j)}$. For each sample, we simulated a network $Y^{(j,s)}$ from the ERGM with parameter $\mu_s^{(j)}$, and compared the local and global efficiency of the simulated networks to the observed networks.

2.7. Availability of data and code

Access to the Cam-CAN dataset for the purpose of scientific investigation, teaching or the planning of clinical research studies can be requested at <https://camcan-archive.mrc-cbu.cam.ac.uk/dataaccess/datarequest.php>. The code to fit the Bayesian multilevel framework with exponential random graph models is available to download as an R package at <https://github.com/briuclehmann/multibergm/releases/v0.1>. The correlation matrices and scripts used to generate the results for this paper are available at <https://doi.org/10.17605/osf.io/5nh94>.

3. Results

We performed analyses for two populations of networks, the first with a single group (young only) and the second with two groups (young and old). Both the one and two group settings were considered under two different thresholding procedures (absolute and proportional) and two different threshold values (chosen to yield average node degrees of $K = 3$ and $K = 5$); resulting in eight sets of networks that we analysed




Fig. 2. Posterior mean estimates of models fit to networks constructed under absolute thresholding with group-wide average node degree $K = 3$. The histogram (grey bars) corresponds to the single-BERGM posterior mean values of the model parameters fitted to each of the $n = 100$ individual networks. The vertical lines correspond to the posterior mean values of group-level mean parameter under multi-BERGM (red line), mean-BERGM (blue dotted line), mean-GRN BERGM (green dashed line) and median-GRN BERGM (purple dashed line). For each of the model components, the posterior mean under multi-BERGM lay close to the mean-BERGM estimate, i.e. the centre of the single-BERGM mean values. In contrast, for the edges and GWESP components, posterior means for the mean-GRN BERGM and median-GRN BERGM lay in the tails of the single-BERGM mean values.

using our multi-BERGM framework. In this section we present the results of the analyses performed on the sets of networks constructed under absolute thresholding with a population-wide average node degree of $K = 3$. The other results can be found in the Supplementary Material. In the context of this paper the supplementary results are in line with those presented here, except where explicitly stated.

For the single-group setting we compare our method, specifically the estimated ERGM parameters, to several alternative approaches from the literature.

3.1. Single-group mean ERGM parameter estimation

We first consider the point estimate of the ERGM parameters, μ , for the young only (i.e. single-group setting) under various approaches: our one-level multi-BERGM framework; the mean-BERGM, the mean-GRN BERGM; the median-GRN BERGM; and the distribution of point estimates (posterior sample mean) from a single-BERGM fitted to each individual's network. Briefly, multi-BERGM fits all the model parameters simultaneously in a multilevel framework; mean-BERGM combines estimates from the separate single-BERGM fits; mean-GRN and median-GRN are single-BERGM fits on the group-representative mean and median networks respectively (see Section 2.5 for details).

Fig. 2 presents the three components of μ (corresponding to the three summary statistics) showing the mean-BERGM estimate (blue dotted line); the posterior sample means of the one-level multi-BERGM (red line), the mean-GRN BERGM (green dashed line), and the median-GRN

BERGM (purple dashed line); and the distribution of posterior sample means for single-BERGM fitted to each individual (grey histogram).

The multi-BERGM estimates are very close to the mean-BERGM estimates, i.e. the centre of the posterior mean values across the individual network fits (single-BERGMs). In contrast, the posterior means of the edges and GWESP components for the mean-GRN and median-GRN are in the tails of single-BERGM distribution. Similar results were observed under the proportional thresholding procedures while the estimates under absolute thresholding with average node degree $K = 5$ were similar under each of the approaches (see Supplementary material Fig. A.1). This suggests that the mean-GRN and median-GRN do not always accurately capture the topological structure of the majority of the individual networks.

3.2. Individual-level covariance under multi-BERGM and single-BERGM

We consider the covariance of the posterior distribution of the individual-level parameters for three randomly selected participants under our multi-BERGM approach compared to the distribution from the single-BERGM fits. For each approach, we calculated 95% credible regions for each pair of model components based on the respective posterior samples.

While the mean values were similar between the multi-BERGM and single-BERGM distributions, the multi-BERGM yielded tighter credible regions (Fig. 3). This illustrates one of the benefits of modelling the networks simultaneously: the borrowing of information across individuals




Fig. 3. Posterior means and 95% credible regions of the individual-level parameters for three subjects and for each pair of model components under the multi-BERGM (dot and solid line) and under the single-BERGM (cross and dashed line) approaches. The correlation between parameter estimates is shown by the elliptical shape of the regions.




Fig. 4. Posterior means and 95% credible regions of the group-level mean parameters for each pair of model components under the multi-BERGM (dot and solid line) and under the mean-BERGM (cross and dashed line) approaches.

leads to more precise estimates of the model parameters. The credible regions derived from the other thresholding procedures were similar or tighter under multi-BERGM compared to single-BERGM (see Supplementary material Fig. A.2).

3.3. Single-group covariance under multi-BERGM and mean-BERGM

We consider the covariance of the posterior distribution under our multi-BERGM approach compared to the mean-BERGM distribution constructed from the posterior samples of the single-BERGMs for each individual. From the posterior samples for the multi-BERGM, we calculated 95% credible regions for each pair of model components. For the mean-BERGM approach, we calculate a 95% region from a normal distribution with mean and covariance based on the single-BERGM fits for each individual.

While the mean values were similar between the multi-BERGM and mean-BERGM distributions, the multi-BERGM yielded moderately tighter credible regions (Fig. 4). This again illustrates how the borrowing of information across individuals can lead to more precise estimates of the model parameters. The credible regions derived from the other thresholding procedures were similar or tighter under multi-BERGM compared to the mean-BERGM approach (see Supplementary material Fig. A.3).

3.4. Single-group goodness-of-fit under multi-BERGM

To assess the goodness-of-fit under multi-BERGM, we simulated networks from the ERGM with parameters taken uniformly at random from the posterior samples. The simulated networks appeared broadly similar to the observed networks in terms of the network metrics we considered: degree distribution, geodesic distance distribution and edgewise-shared partner distribution (Fig. 5). This indicates that the group-level posterior distribution adequately captured the important network characteristics across the entire group of individuals. Similar results were observed under the remaining three thresholding procedures (see Supplementary material Figs A.4–A.6).

3.5. Two-group multi-BERGM: group differences

Using a two-level multi-BERGM framework we obtained posteriors for the group-level ERGM parameters. These posterior distributions reveal that network differences between the young group and the old group are driven by differences in the GWNSP parameter (Fig. 6). In particular, the posteriors for the edges and GWESP parameters are almost

identical for the two groups, while the GWNSP parameter is markedly larger for the young group. This difference in the group-level GWNSP parameters was also observed under the remaining three thresholding procedures (Supplementary Figs A.8–A.10).

By considering the posterior distribution of difference in the group-level parameters, $\mu_1 - \mu_2$, we also calculated 95% credible regions to assess the degree of certainty in the group differences under both the multi-BERGM and the mean-BERGM approaches (Fig. 7). The left panel shows that difference in the (edges-GWESP) pair of parameters is centered on zero, while the middle and right panels demonstrate that, under the multi-BERGM approach, the credible regions for the (GWESP-GWNSP) and (GWNSP-edges) pair do not contain zero. This provides evidence that the group-level differences in network structure are driven by the differences in GWNSP. We note that a permutation testing approach based on single-BERGM and single-ERGM fits also found the strongest group difference between the GWNSP parameter estimates (Supplementary Table A.1).

Under the mean-BERGM approach, the credible regions are slightly wider and contain zero for each pair of parameters. This again highlights the benefit of more precise estimates of the model parameters via the multi-BERGM approach by modelling the networks simultaneously. Similar results were observed under absolute thresholding with average node degree $K = 5$ but no significant differences were found under proportional thresholding (Fig. A.11)

Age-related differences in both local efficiency (Achard and Bullmore, 2007; Geertligs et al., 2015; Sala-Llonch et al., 2014; Song et al., 2014) and global efficiency (Achard and Bullmore, 2007; Sala-Llonch et al., 2014) have previously been observed. To check whether the multi-BERGM approach could identify similar age-related group differences, we also simulated $S = 1000$ networks for each group from the ERGM with parameters taken uniformly at random from the group-level posterior samples. Overall, the simulated networks for the young group exhibited higher global and local efficiency, in correspondence with the observed data (Fig. 8). Similar results were observed for networks constructed via absolute thresholding with average node degree $K = 5$ (Supplementary Fig. A.12). Under proportional thresholding, however, the posterior predictive distribution of local and global efficiency did not correspond well with the observed networks, indicating a poor model fit (Supplementary Figs A.13 and A.14).

Based on the results from previous sections, we have shown that the various GRN methods do not properly reflect the distribution of individual ERGMs, and hence we do not consider group differences under these alternative approaches.




Fig. 5. Goodness-of-fit assessment of the multi-BERGM for a single group. $S = 1000$ networks were generated from the ERGM with model parameters sampled from the respective group-level distributions. The simulated networks are compared to the observed data (shown in the box plots) across three network metrics: degree distribution, geodesic distance distribution, edge-wise shared partner distribution. The lines correspond to the respective means across the simulated networks, while the ribbons correspond to 95% credible intervals.

4. Discussion

The characterisation of connectivity structure at a group level is a necessary step towards understanding how the brain changes with a given phenotype such as age, disease condition or mental state. We have proposed multi-BERGM, a multi-level (hierarchical) Bayesian framework based on exponential random graph models (ERGMs), as a coherent method for understanding group level connectivity. By pooling information across the individual networks, this framework provides a principled (and extendable) approach to model the relational structure for groups of networks.

Our work builds on previous approaches attempting to investigate differences between groups of networks, either by constructing a group-representative network (GRN; Achard et al., 2006; Joyce et al., 2010; Meunier et al., 2009; Sinke et al., 2016; Song et al., 2009) or by combining separate individual-level estimates, e.g. mean-ERGM (Obando and Fallani, 2017; Simpson et al., 2012). There are three main advantages of our framework: using the Bayesian paradigm allows us to formally

quantify uncertainty in the group-level model parameters; the group-level posterior densities provide a way to directly compare connectivity structure across groups; and the multi-level (hierarchical) framework allows information to be borrowed across individuals thus, in principle, improving the accuracy of the parameter estimates.

In line with previous findings (Ginestet et al., 2011; Simpson et al., 2012), our results illustrate that the mean-GRN and median-GRN approaches for constructing group-representative networks do not accurately characterise the topological structure of the individual networks. This is at least partially due to the *quasilinearity* of the thresholding operation used to construct the networks: the expectation of thresholded correlation matrices will not in general be equal to the network constructed by thresholding the expectation of the correlation matrices (Ginestet et al., 2011). As a result, the network metrics from such a GRN will not accurately summarise the topological information across the individual networks.

Our proposed method takes as input any group of unweighted networks. While there are a variety of methods to construct networks from




Fig. 6. Posterior density estimates under multi-BERGM for the group-level mean parameters for networks constructed via absolute thresholding with average node degree $K = 3$. The young group displayed markedly larger values for the GWNSP parameter while the edges and GWESP posteriors were almost identical between the two groups.




Fig. 7. Posterior means and 95% credible regions of the difference in group-level mean parameters for each pair of model components under the multi-BERGM (dot and solid line) and under the mean-BERGM (cross and dashed line) approaches.

functional neuroimaging data, each with their own benefits and disadvantages, we took the most common and simple approach of applying a threshold to the correlation matrices, keeping edges between nodes that exhibited a sufficiently strong positive correlation. The choice of threshold can have a significant impact on the resulting connectivity structure (van den Heuvel et al., 2017; van Wijk et al., 2010) and we found that different threshold values did indeed lead to small differences in the posterior distribution of the model parameter. Regardless of the threshold used, our approach provides a framework to characterise the group-level brain connectivity structure, though further investigation into the effect of thresholding procedure on ERGM parameter estimates is warranted.

From a modelling perspective, the appropriate selection of summary statistics is crucial. This is true of fitting ERGMs to single networks, let alone populations of networks. Moreover, the ‘correct’ choice likely depends on the network construction method. We used three summary statistics: the number of edges, the geometrically-weighted edgewise shared partners (GWESP) statistic and the geometrically-weighted non-edgewise shared partners (GWNSP) statistic. The choice of summary statistics follows Simpson et al. (2011), in which this particular combination was selected via a graphical goodness-of-fit approach, with the ‘best’ metrics chosen from a prespecified set of potential metrics. While it is possible to perform Bayesian model selection for an ERGM on a




Fig. 8. Local and global efficiency in the observed networks (bar plots) compared to $S = 1000$ networks simulated from the posterior predictive distribution (density plots) for the young group (red) and the old group (blue).

single network (Bouranis et al., 2018; Caimo and Friel, 2013), further work is needed to extend these approaches to perform model selection for a group of networks.

Although ERGMs provide a concise representation of network data via a set of summary statistics, the biological interpretation of these summary statistics requires some care. The number of edges corresponds to overall network density and is a natural measure of connectedness. The GWESP/GWNSP statistics are weighted sums of the number of connected/non-connected node pairs having exactly w shared partners. The GWESP statistic quantifies the propensity to form triangles and is thus a measure of local clustering (functional segregation). The GWNSP statistic can be seen as an indirect measure of global network efficiency (functional integration); a higher value of GWNSP indicates that non-connected nodes are more likely to have a shared partner. Both GWNSP and GWESP were highly correlated with local as well as global network efficiency (Supplementary Table A.2) for the 200 networks constructed under absolute thresholding with average node degree $K = 3$. This is in line with previous work (Obando and Fallani, 2017) which concluded that it may not be possible to uniquely relate GWNSP and GWESP to global and local efficiency respectively. Note however that GWESP and GWNSP were also highly correlated with each other. In fact, when Spearman's *partial* correlation is used, the correlation between GWNSP and local efficiency (controlling for GWESP) is close to zero while the correlation between GWNSP and global efficiency is still strong (see Supplementary Tables A.3–A.4). Similarly, the correlation between GWESP and local efficiency (controlling for GWNSP) is close to zero while the correlation between GWESP and global efficiency is moderately large. These results indicate that GWNSP and GWESP are indeed useful surrogates for global and local efficiency respectively.

The parameter estimates obtained from the fitting procedure quantify the relative influence of the corresponding summary statistics on the overall network structure. To be precise, the parameter estimates represent the estimated change in log-odds of an edge existing for each unit increase of the corresponding summary statistic. The negative estimates for the edges parameter corresponds to a lower expected edge density relative to a random graph in which each edge exists with probability $1/2$. Note that this relates strongly to the overall sparsity of the network and is thus highly dependent on the threshold level used. The positive estimates for the GWESP parameter indicate that many connections contribute to the formation of clusters in the brain. This corresponds to the well-established local efficiency and functional segregation of the brain (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010). The negative estimates for the GWNSP parameter suggest that pairs of non-connected nodes sharing one or more neighbours are relatively uncommon. The lower estimates of this parameter for the old group indicate lower levels of functional integration. Note that the parameter estimates, as well

as the corresponding summary statistics, are strongly correlated, so caution should be exercised when interpreting the estimates independently.

A graphical assessment of the model indicated a reasonable goodness-of-fit under constant thresholding. In particular, the model was able to recover differences in observed global and local efficiency, which were both higher in the young group than the old group. The observed age-related decrease in local efficiency is in line with previous studies (Achard and Bullmore, 2007; Geertligs et al., 2015; Sala-Llonch et al., 2014; Song et al., 2014). A decline in global efficiency has also been observed previously (Achard and Bullmore, 2007; Sala-Llonch et al., 2014), though not in other studies (Geertligs et al., 2015; Song et al., 2014). In our study, the thresholding method had a pronounced effect on the observed age-differences. Under constant thresholding, we observed a clear difference in global efficiency between the two age groups and a relatively small difference in local efficiency. Under proportional thresholding, however, both local and global efficiency were slightly lower for the young group. The disparities between thresholding procedures may be associated with alterations in mean connectivity due to age-related differences in vascular health (Geertligs et al., 2017), while the discrepancies with previous studies are likely due to differences in network construction including the parcellation and definition of functional connectivity.

Future research will be needed to extend our approach to alternative settings. For example, the flexibility of Bayesian hierarchical modelling could be exploited to deal with continuous covariates or more complex group structures such as factorial designs. Given multiple networks per individual, one could extend the framework by adding another layer to the model. Specifically, one could again posit an ERGM for each network observation, with individual-level parameters to pool information for networks on the same individual and group-level parameters to pool information across several individuals. It would then be possible to apply our framework to dynamic functional connectivity or task-based scans in order to query how connectivity structure changes with time or across states. Although our framework is currently limited to binary networks, more work is needed to handle weighted networks (thus bypassing the thresholding issue). A promising avenue for this is the development of weighted ERGMs (Desmarais and Cranmer, 2012; Wilson et al., 2017), though the associated computational cost is likely to be prohibitive.

A downside to our framework, especially compared to the single-(B)ERGM and GRN approaches, is in terms of computation cost. Using a 40-core computing cluster (Intel Xeon E7-8860 v3, 2.2 GHz), the algorithm took between 8 and 12 h to produce the posterior samples for each multi-BERGM (with additional time to tune the Bayesian procedure). However, Krivitsky and Handcock (2014) provide empirical evidence suggesting the computational cost may grow on the order of $p(N + E) \log(E)$ where p is the number of summary statistics,

N is the number of nodes, and E is the (typical) number of edges (as well as increase roughly linearly in the number of networks, n). In practical terms this limits the size of networks that our framework can currently handle to relatively coarse parcellations of the brain (approximately 100 regions). Note that the computational cost also increases with the density of the networks. The efficiency cost optimisation (ECO) criterion, corresponding to an average node degree of $K = 3$, offers a principled threshold that yields sparse networks while maintaining their essential structure (Fallani et al., 2017). Improving the efficiency of fitting (B)ERGMs is an ongoing area of research (Bouranis et al., 2017; Cerqueira et al., 2020; Tan and Friel, 2020) and our framework will benefit from ongoing developments and may be applicable to more spatially resolved networks in the future. The code used to apply our framework is available as an R package at <https://github.com/briuclehmann/multiBERGM/releases/v0.1>.

Declaration of Competing Interest

The authors declare that they do not have any financial or nonfinancial conflict of interests

Credit authorship contribution statement

B.C.L. Lehmann: Conceptualization, Formal analysis, Methodology, Software, Writing - original draft. **R.N. Henson:** Supervision, Writing - review & editing. **L. Geerligs:** Data curation, Supervision, Writing - review & editing. **S.R. White:** Conceptualization, Methodology, Supervision, Writing - review & editing.

Acknowledgments

BCLL was supported by the UK Engineering and Physical Sciences Research Council [Programme number EP/R018561/1] and Jesus College, Oxford. BCLL, SRW and RNH were supported by the UK Medical Research Council [Programme numbers U105292687 and SUAG/046 G101400]. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) was supported by the Biotechnology and Biological Sciences Research Council (grant number BB/H008217/1). LG was supported by a Veni grant [451-16-013] from the Netherlands Organization for Scientific Research. This research was supported by the NIHR Cambridge Biomedical Research Centre. The views expressed are those of the author(s) and not necessarily those of the NIHR or the Department of Health and Social Care The Cam-CAN corporate author consists of the project principal personnel: Lorraine K Tyler, Carol Brayne, Edward T Bullmore, Andrew C Calder, Rhodri Cusack, Tim Dalgleish, John Duncan, Richard N Henson, Fiona E Matthews, William D Marslen-Wilson, James B Rowe, Meredith A Shafto; Research Associates: Karen Campbell, Teresa Cheung, Simon Davis, Linda Geerligs, Rogier Kievit, Anna McCarrey, Abdur Mustafa, Darren Price, David Samu, Jason R Taylor, Matthias Treder, Kamen Tsvetanov, Janna van Belle, Nitin Williams; Research Assistants: Lauren Bates, Tina Emery, Sharon Erziniolgu, Andrew Gadie, Sofia Gerbase, Stanimira Georgieva, Claire Hanley, Beth Parkin, David Troy; Affiliated Personnel: Tibor Auer, Marta Correia, Lu Gao, Emma Green, Rafael Henriques; Research Interviewers: Jodie Allen, Gillian Amery, Liana Amunts, Anne Barcroft, Amanda Castle, Cheryl Dias, Jonathan Dowrick, Melissa Fair, Hayley Fisher, Anna Goulding, Adarsh Grewal, Geoff Hale, Andrew Hilton, Frances Johnson, Patricia Johnston, Thea Kavanagh-Williamson, Magdalena Kwasniewska, Alison McMinn, Kim Norman, Jessica Penrose, Fiona Roby, Diane Rowland, John Sargeant, Maggie Squire, Beth Stevens, Aldabra Stoddart, Cheryl Stone, Tracy Thompson, Ozlem Yazlik; and administrative staff: Dan Barnes, Marie Dixon, Jaya Hillman, Joanne Mitchell, Laura Willis.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.neuroimage.2020.117480](https://doi.org/10.1016/j.neuroimage.2020.117480)

References

- Achard, S., Bullmore, E., 2007. Efficiency and cost of economical brain functional networks. *PLOS Comput. Biol.* 3 (2), 1–10. doi:[10.1371/journal.pcbi.0030017](https://doi.org/10.1371/journal.pcbi.0030017).
- Achard, S., Salvador, R., Whitcher, B., Suckling, J., Bullmore, E., 2006. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *J. Neurosci.* 26 (1), 63–72. doi:[10.1523/JNEUROSCI.3874-05.2006](https://doi.org/10.1523/JNEUROSCI.3874-05.2006).
- Bouranis, L., Friel, N., Maire, F., 2017. Efficient Bayesian inference for exponential random graph models by correcting the pseudo-posterior distribution. *Soc. Netw.* 50 (Supplement C), 98–108. doi:[10.1016/j.socnet.2017.03.013](https://doi.org/10.1016/j.socnet.2017.03.013).
- Bouranis, L., Friel, N., Maire, F., 2018. Bayesian model selection for exponential random graph models via adjusted pseudolikelihoods. *J. Comput. Graph. Stat.* 27 (3), 516–528. doi:[10.1080/10618600.2018.1448832](https://doi.org/10.1080/10618600.2018.1448832).
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10 (3), 186–198. doi:[10.1038/nrn2575](https://doi.org/10.1038/nrn2575).
- Caimo, A., Friel, N., 2011. Bayesian inference for exponential random graph models. *Soc. Netw.* 33 (1), 41–55. doi:[10.1016/j.socnet.2010.09.004](https://doi.org/10.1016/j.socnet.2010.09.004).
- Caimo, A., Friel, N., 2013. Bayesian model selection for exponential random graph models. *Soc. Netw.* 35 (1), 11–24. doi:[10.1016/j.socnet.2012.10.003](https://doi.org/10.1016/j.socnet.2012.10.003).
- Cerqueira, A., Garivier, A., Leonardi, F., 2020. A note on perfect simulation for Exponential Random Graph Models. *ESAIM: Probability and Statistics* 24, 138–147. doi:[10.1051/ps/2019024](https://doi.org/10.1051/ps/2019024).
- Cusack, R., Vicente-Grabovetsky, A., Mitchell, D.J., Wild, C.J., Auer, T., Linke, A.C., Peelle, J.E., 2015. Automatic analysis (aa): efficient neuroimaging workflows and parallel processing using matlab and xml. *Frontiers in Neuroinformatics* 8, 90. doi:[10.3389/fninf.2014.00090](https://doi.org/10.3389/fninf.2014.00090).
- Dell'Italia, J., Johnson, M.A., Vespa, P.M., Monti, M.M., 2018. Network analysis in disorders of consciousness: Four problems and one proposed solution (exponential random graph models). *Front. Neurol.* 9, 439. doi:[10.3389/fneur.2018.00439](https://doi.org/10.3389/fneur.2018.00439).
- Desmarais, B.A., Cranmer, S.J., 2012. Statistical inference for valued-edge networks: the generalized exponential random graph model. *PLoS One* 7 (1), e30136.
- Fallani, F.D.V., Latora, V., Chavez, M., 2017. A topological criterion for filtering information in complex brain networks. *PLOS Comput. Biol.* 13 (1), e1005305. doi:[10.1371/journal.pcbi.1005305](https://doi.org/10.1371/journal.pcbi.1005305).
- Geerligs, L., Renken, R.J., Saliassi, E., Maurits, N.M., Lorist, M.M., 2015. A brain-wide study of age-related changes in functional connectivity. *Cereb. Cortex* 25 (7), 1987–1999. doi:[10.1093/cercor/bhu012](https://doi.org/10.1093/cercor/bhu012).
- Geerligs, L., Tsvetanov, K.A., Cam-CAN, Henson, R.N., 2017. Challenges in measuring individual differences in functional connectivity using Fmri: the case of healthy aging. *Hum. Brain Mapp.* 38 (8), 4125–4156. doi:[10.1002/hbm.23653](https://doi.org/10.1002/hbm.23653).
- Ginestet, C.E., Fournel, A.P., Simmons, A., 2014. Statistical network analysis for functional MRI: summary networks and group comparisons. *Front. Comput. Neurosci.* 8, 51. doi:[10.3389/fncom.2014.00051](https://doi.org/10.3389/fncom.2014.00051).
- Ginestet, C.E., Nichols, T.E., Bullmore, E.T., Simmons, A., 2011. Brain network analysis: separating cost from topology using cost-integration. *PLOS ONE* 6 (7), 1–17. doi:[10.1371/journal.pone.0021570](https://doi.org/10.1371/journal.pone.0021570).
- van den Heuvel, M.P., de Lange, S.C., Zalesky, A., Seguin, C., Yeo, B.T., Schmidt, R., 2017. Proportional thresholding in resting-state fMRI functional connectivity networks and consequences for patient-control connectome studies: issues and recommendations. *NeuroImage* 152, 437–449. doi:[10.1016/j.neuroimage.2017.02.005](https://doi.org/10.1016/j.neuroimage.2017.02.005).
- Hunter, D.R., 2007. Curved exponential family models for social networks. *Soc. Netw.* 29 (2), 216–230. doi:[10.1016/j.socnet.2006.08.005](https://doi.org/10.1016/j.socnet.2006.08.005). Special Section: Advances in Exponential Random Graph (p^*) Models
- Joyce, K.E., Laurienti, P.J., Burdette, J.H., Hayasaka, S., 2010. A new measure of centrality for brain networks. *PLOS ONE* 5 (8), 1–13. doi:[10.1371/journal.pone.0012200](https://doi.org/10.1371/journal.pone.0012200).
- Koskinen, J., 2004. Bayesian analysis of exponential random graphs-estimation of parameters and model selection. Technical Report. Stockholm University.
- Krivitsky, P.N., Handcock, M.S., 2014. Supplementary material: a separable model for dynamic networks. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* 76 (1), 29–46. doi:[10.1111/rssb.12014](https://doi.org/10.1111/rssb.12014).
- Lehmann, B.C.L., 2019. Inferring differences between networks using Bayesian exponential random graph models. University of Cambridge.
- Meunier, D., Lambiotte, R., Fornito, A., Ersche, K., Bullmore, E., 2009. Hierarchical modularity in human brain functional networks. *Front. Neuroinf.* 3, 37. doi:[10.3389/fninf.2009.11.037](https://doi.org/10.3389/fninf.2009.11.037).
- Murray, I., Ghahramani, Z., MacKay, D.J.C., 2006. MCMC for doubly-intractable distributions. In: *Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*. AUAI Press, pp. 359–366.
- Obando, C., Fallani, F.D.V., 2017. A statistical model for brain networks inferred from large-scale electrophysiological signals. *J. R. Soc. Interface* 14 (128), 20160940. doi:[10.1098/rsif.2016.0940](https://doi.org/10.1098/rsif.2016.0940).
- Patel, A.X., Kundu, P., Rubinov, M., Jones, P.S., Vrtes, P.E., Ersche, K.D., Suckling, J., Bullmore, E.T., 2014. A wavelet method for modeling and despiking motion artifacts from resting-state fMRI time series. *NeuroImage* 95, 287–304. doi:[10.1016/j.neuroimage.2014.03.012](https://doi.org/10.1016/j.neuroimage.2014.03.012).
- Robins, G., Pattison, P., Kalish, Y., Lusher, D., 2007. An introduction to exponential random graph (p^*) models for social networks. *Soc. Netw.* 29 (2), 173–191. doi:[10.1016/j.socnet.2006.08.002](https://doi.org/10.1016/j.socnet.2006.08.002). Special Section: Advances in Exponential Random Graph (p^*) Models
- Robins, G., Snijders, T., Wang, P., Handcock, M., Pattison, P., 2007. Recent developments in exponential random graph (p^*) models for social networks. *Soc. Netw.* 29 (2), 192–215. doi:[10.1016/j.socnet.2006.08.003](https://doi.org/10.1016/j.socnet.2006.08.003). Special Section: Advances in Exponential Random Graph (p^*) Models

- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* 52 (3), 1059–1069. doi:[10.1016/j.neuroimage.2009.10.003](https://doi.org/10.1016/j.neuroimage.2009.10.003). Computational Models of the Brain
- Sala-Llonch, R., Junqu, C., Arenaza-Urquijo, E.M., Vidal-Pieiro, D., Valls-Pedret, C., Palacios, E.M., Domnech, S., Salv, A., Bargall, N., Bartró-Faz, D., 2014. Changes in whole-brain functional networks and memory performance in aging. *Neurobiol. Aging* 35 (10), 2193–2202. doi:[10.1016/j.neurobiolaging.2014.04.007](https://doi.org/10.1016/j.neurobiolaging.2014.04.007).
- Shafto, M.A., Tyler, L.K., Dixon, M., Taylor, J.R., Rowe, J.B., Cusack, R., Calder, A.J., Marslen-Wilson, W.D., Duncan, J., Dalgleish, T., Henson, R.N., Brayne, C., Matthews, F.E., 2014. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC Neurol.* 14 (1), 1–25. doi:[10.1186/s12883-014-0204-1](https://doi.org/10.1186/s12883-014-0204-1).
- Simpson, S., Lyday, R., Hayasaka, S., Marsh, A., Laurienti, P., 2013. A permutation testing framework to compare groups of brain networks. *Front. Comput. Neurosci.* 7, 171. doi:[10.3389/fncom.2013.00171](https://doi.org/10.3389/fncom.2013.00171).
- Simpson, S.L., Bowman, F.D., Laurienti, P.J., 2013. Analyzing complex functional brain networks: fusing statistics and network science to understand the brain. *Statist. Surv.* 7, 1–36. doi:[10.1214/13-SS103](https://doi.org/10.1214/13-SS103).
- Simpson, S.L., Hayasaka, S., Laurienti, P.J., 2011. Exponential random graph modeling for complex brain networks. *PLoS One* 6 (5), e20039.
- Simpson, S.L., Laurienti, P.J., 2015. A two-part mixed-effects modeling framework for analyzing whole-brain network data. *NeuroImage* 113, 310–319. doi:[10.1016/j.neuroimage.2015.03.021](https://doi.org/10.1016/j.neuroimage.2015.03.021).
- Simpson, S.L., Moussa, M.N., Laurienti, P.J., 2012. An exponential random graph modeling approach to creating group-based representative whole-brain connectivity networks. *NeuroImage* 60 (2), 1117–1126. <https://doi.org/10.1016/j.neuroimage.2012.01.071>.
- Sinke, M.R., Dijkhuizen, R.M., Caimo, A., Stam, C.J., Otte, W.M., 2016. Bayesian exponential random graph modeling of whole-brain structural networks across lifespan. *NeuroImage* 135 (Supplement C), 79–91. doi:[10.1016/j.neuroimage.2016.04.066](https://doi.org/10.1016/j.neuroimage.2016.04.066).
- Song, J., Birn, R.M., Boly, M., Meier, T.B., Nair, V.A., Meyerand, M.E., Prabhakaran, V., 2014. Age-related reorganizational changes in modularity and functional connectivity of human brain networks. *Brain Connect.* 4 (9), 662–676. doi:[10.1089/brain.2014.0286](https://doi.org/10.1089/brain.2014.0286). PMID: 25183440
- Song, M., Liu, Y., Zhou, Y., Wang, K., Yu, C., Jiang, T., 2009. Default network and intelligence difference. In: *Proceedings of the 2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 2212–2215. doi:[10.1109/IEMBS.2009.5334874](https://doi.org/10.1109/IEMBS.2009.5334874).
- Tan, L.S., Friel, N., 2020. Bayesian variational inference for exponential random graph models. *J. Comput. Graph. Stat.* 1–19. doi:[10.1080/10618600.2020.1740714](https://doi.org/10.1080/10618600.2020.1740714).
- Taylor, J.R., Williams, N., Cusack, R., Auer, T., Shafto, M.A., Dixon, M., Tyler, L.K., Cam-CAN, Henson, R.N., 2017. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage* 144, 262–269. doi:[10.1016/j.neuroimage.2015.09.018](https://doi.org/10.1016/j.neuroimage.2015.09.018). Data Sharing Part II
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., 2002. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15 (1), 273–289. doi:[10.1006/nimg.2001.0978](https://doi.org/10.1006/nimg.2001.0978).
- van Wijk, B.C.M., Stam, C.J., Daffertshofer, A., 2010. Comparing brain networks of different size and connectivity density using graph theory. *PLOS ONE* 5 (10), 1–13. doi:[10.1371/journal.pone.0013701](https://doi.org/10.1371/journal.pone.0013701).
- Wilson, J.D., Denny, M.J., Bhamidi, S., Cranmer, S.J., Desmarais, B.A., 2017. Stochastic weighted graphs: flexible model specification and simulation. *Soc. Netw.* 49 (Supplement C), 37–47. doi:[10.1016/j.socnet.2016.11.002](https://doi.org/10.1016/j.socnet.2016.11.002).