Published in final edited form as:

J Comput Graph Stat. 2019 ; 28(1): 185–196. doi:10.1080/10618600.2018.1505633.

Multiresolution Network Models

Bailey K. Fosdick^a, Tyler H. McCormick^b, Thomas Brendan Murphy^c, Tin Lok James Ng^c, Ted Westling^d

^aDepartment of Statistics, Colorado State University, Fort Collins, CO; ^bDepartment of Statistics, Department of Sociology, University of Washington, Seattle, WA; ^cSchool of Mathematics and Statistics, University College Dublin, Belfield, Dublin, Ireland; ^dDepartment of Statistics, University of Washington, Seattle, WA

Abstract

Many existing statistical and machine learning tools for social network analysis focus on a single level of analysis. Methods designed for clustering optimize a global partition of the graph, whereas projection-based approaches (e.g., the latent space model in the statistics literature) represent in rich detail the roles of individuals. Many pertinent questions in sociology and economics, however, span multiple scales of analysis. Further, many questions involve comparisons across disconnected graphs that will, inevitably be of different sizes, either due to missing data or the inherent heterogeneity in real-world networks. We propose a class of network models that represent network structure on multiple scales and facilitate comparison across graphs with different numbers of individuals. These models differentially invest modeling effort within subgraphs of high density, often termed communities, while maintaining a parsimonious structure between said subgraphs. We show that our model class is projective, highlighting an ongoing discussion in the social network modeling literature on the dependence of inference paradigms on the size of the observed graph. We illustrate the utility of our method using data on household relations from Karnataka, India. Supplementary material for this article is available online.

Keywords

Latent space; Multiscale; Projectivity; Social network; Stochastic blockmodel

CONTACT Bailey K. Fosdick or Tyler H. McCormick, bailey.fosdick@colostate.edu; tvlermc@u.washington.edu, Colorado State University, Fort Collins, CO 80523; University of Washington, Seattle, WA 98195.

Supplementary Materials

The supplementary material includes: 1) a discussion of the interplay between exchangeability, projectivity, and asymptotic sparsity of random networks; 2) details on the prior distributions, MCMC sampling procedure, and post-processing steps used to perform Bayesian inference for the LS-SBM; 3) an algorithm for choosing the number the blocks in the model; 4) a fast estimation scheme for fitting the LS-SBM to large data sets; 5) details on the LS-SBM and LPCM estimates for the Karnataka data set; and 6) additional details on the simulation study.

Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JCGS.

1. Introduction

Social network data consist of a sample of actors and information on the presence/absence of pairwise relationships among them. These data are often represented as a graph where nodes correspond to actors and edges (ties) connect nodes with a relationship. A relationship may represent, for example, friendship between students, co-authorship between academics on a journal article, or a financial transaction between organizations. Understanding structure in social networks is essential to appreciating the nuances of human behavior and is an active area of research in the social sciences (Borgatti et al. 2009). Existing statistical models for social networks typically focus on either (i) carefully representing structure among actors that have a relatively high likelihood of interaction, or (ii) clearly differentiating between groups of actors, that is, communities, within the graph that have high within-group connectivity and low between-group connectivity (see, e.g., Salter-Townshend et al. 2012, for a review). Unfortunately, neither of these approaches fully characterize the complexities displayed in many real-world social networks.

Observed graphs frequently exhibit a mixture structure that manifests through a combination of global sparsity and local density. Global sparsity implies that the propensity for a tie between any two randomly selected actors is incredibly small. Yet, massive heterogeneity in the propensity for actors to connect often creates local graph structure concentrated in dense subgraphs, frequently termed *communities*. This structure is typically particularly pronounced in very large graphs. For example, in the context of online communication networks, Ugander et al. (2011) described the Facebook graph as containing pockets of “surprisingly dense” structure, though overall the graph is immensely sparse.

In this article, we propose a multiresolution model for capturing heterogeneous, complex structure in social networks that exhibit strong community structure. Our modeling framework decomposes network structure into a component that describes between-community relations, that is, relations between actors belonging to different communities, and another component describing within-community relations. The proposed framework has two distinct advantages over existing methods. First, our framework is able to accommodate a wide variety of models for between- and within-community relations. This feature allows the model to be tailored to reflect different scientific questions that arise when exploring the behavior within and across these communities. The second advantage of our model is that it balances parsimony with model richness by selectively directing modeling efforts toward representing interesting, relevant network structure. Typically, this structure is found within actors’ local communities. In such cases, we can exert the most modeling effort (i.e., model complexity and computational effort) within dense pockets, where we expect the most complex dependence structure, and use a parsimonious model to capture between community patterns. A similar approach has been adopted in spatial statistics where locations are partitioned into disjoint dependence neighborhoods (Page and Quintana 2016). Compared to popular network models that capture global structure, our approach can provide increased resolution on intricate structure within communities. Furthermore, our model is able to apportion little effort to modeling simple structure, resulting in a model that is substantially less complex than existing models focused on local structure for networks, even with only a few hundred actors.

After defining our model framework, we discuss its statistical properties. In exploring these properties, we take a traditional sampling perspective and consider our observed network as that pertaining to a collection of actors sampled from an infinite population of actors. Our goal is to learn features of the infinite population from the observed graph. In our model, these features include the distribution of within-community structure across the network. Communities are defined by their structure and may not have a consistent size. Therefore, inference about the population-level parameters requires that we be able to coherently compare and summarize parameters associated with subgraphs of different sizes. We may also desire to compare network-level parameters to those from another network of a different size and a different number of communities. In order for these properties to hold and comparisons to be meaningful, the model class must be a *projective family*, in the “consistency under sampling” sense of Shalizi and Rinaldo (2013). We introduce this concept and show the class of multiresolution network models proposed has this property. We also discuss the important implications of this for population inference.

In the remainder of this section, we explore the two existing approaches to multiresolution modeling of networks, highlighting their strengths and weaknesses. In Section 2, we introduce the general form of our multiresolution modeling framework, and in Section 3 we present one of many possible model instantiations, called the Latent Space Stochastic Blockmodel (LS-SBM). Section 4 describes the projectivity properties of this framework and provides context through comparison with other available methods. Finally, we conclude with a discussion in Section 5.

1.1. Related Models

In this section, we describe two existing models, the Latent Position Cluster Model (Handcock, Raftery, and Tantrum 2007) and the Locally Dependent Exponential Random Graph Model (Schweinberger and Handcock 2014), which capture aspects of our multiresolution approach.

The Latent Position Cluster Model (LPCM) of Handcock, Raftery, and Tantrum (2007) is an extension of the latent geometry framework introduced in Hoff, Raftery, and Handcock (2002), where the probability of network ties is a function of the distance between actor positions in a latent space. The LPCM performs model-based clustering (Fraley and Raftery 2002) on the positions in the unobserved social space. Cluster memberships then capture group structure and within-group analysis is performed by examining the actor-specific latent positions within each cluster. The likelihood for the LPCM, like that for the original latent space models, requires estimating a distance between every two pairs of actors in the unobserved social space. In even moderately large graphs, these distance calculations are computationally expensive and the propensity for actors in different groups to interact is often very small. In addition, since both ties and nonties are weighted equally by the LPCM, the latent position for each node is heavily influenced by the numerous other nodes with which it has no relation.

An attractive property of the LPCM is that it parsimoniously encodes patterns among ties in the network using a low-dimensional structure. As a consequence of this, the model tie probabilities are constrained by the latent geometry. Often these constraints are seen as

model features, for example, the triangle inequality encourages transitivity, which is known to be prevalent in empirical social networks. However, these constraints also restrict the types of network structure that can be represented. Consider a two-dimensional Euclidean latent space and suppose there are four groups of actors such that each pair of actors which belong to different groups interact with the same probability. To model this type of between group structure in the latent space, all groups need to be positioned equidistant from one another. Unfortunately, it is impossible to place four points in $\mathbb{R}^2$ equidistant from one another. The dimension of the latent space could be increased to accommodate this structure. However, ultimately a $K - 1$ -dimensional space is required to model all possible relationships among K groups and continually increasing the latent dimension greatly complicates the model.

The second recent and related model is the locally dependent Exponential Random Graph Model proposed in Schweinberger and Handcock (2014). Exponential Random Graph Models (ERGMs; e.g., Frank and Strauss 1986; Wasserman and Pattison 1996; Pattison and Wasserman 1999; Snijders 2002; Lusher et al. 2012; Koskinen and Melas 2009; Robins 2011; Chatterjee and Diaconis 2013) use graph statistics, such as the number of closed triangles, as sufficient statistics in an exponential family. Schweinberger and Handcock (2014) defined local dependence on a graph as a decoupling of the dependence structure such that dependence exists only among ties within the same community and among ties between the same two communities. Schweinberger and Handcock (2014) drew parallels to local (dependence) neighborhoods in spatial statistics and M-dependence in time series.

A key feature of locally dependent ERGMs is that they are composed of ERGMs. Shalizi and Rinaldo (2013) showed that ERGMs are not projective when the sufficient statistics involve more than two nodes. Lacking projectivity implies that the value of the model parameters changes meaning depending on the sample size. As a result, when partial communities are observed (i.e., some of the actors in various communities are not included in the sample), the parameter estimates from locally dependent ERGMs are difficult to interpret. Further, since the ties within each community are modeled using an ERGM, it is not possible to compare parameters across communities within the same graph unless the communities happen to be the same size. Locally dependent ERGMs are in fact projective if the sampling units are taken to be *communities* rather than actors. Schweinberger and Handcock (2014) called this limited form of projectivity *domain consistency*. While our proposed model uses a similar decomposition across subgraphs as Schweinberger and Handcock (2014), we model within- and between-community structure using latent variable mixture models and show the model class we define is projective when indexed by actors, broadening the notion of local dependence and alleviating the challenges with interpretation and comparison.

2. Multiresolution Network Model

In this section, we propose a general modeling framework that reflects the global sparsity and local density, “chain of islands” (Cross et al. 2001), structure observed in many large networks.

Consider a hypothetical infinite population of actors and communities, where each actor is a member of a single community. Define $\gamma: \mathbb{N} \rightarrow \mathbb{N}$ to be the community membership map, which partitions the actors into disjoint communities. That is, $\gamma(i) = \gamma_i$ is the community of actor i . Define $K_N = |\gamma(\{1, \dots, N\})|$ to be the number of unique communities among actors $\{1, \dots, N\}$. The community map is only meaningful up to relabelings of the communities. Without loss of generality, we require that $\gamma_i = K_{i-1} + 1$ (defining $K_0 = 0$) so that actors $1, \dots, N$ span communities $1, \dots, K_N$. In practice the community memberships are typically estimated from the data, though in certain circumstances it can be defined a priori from known structural breaks in the network (see Sweet, Thomas, and Junker 2013, for an example).

Let $S_k = \{i: \gamma_i = k\}$ be the collection of actors in community k in the population. We assume that the number of actors in each community, for example, $|S_k|$, is bounded, implying that $K_N = O(N)$ as $N \rightarrow \infty$. Our assumption of bounded communities is supported by empirical evidence that suggests the “best” communities contain small sets of actors, which are almost disconnected from the rest of the network (Leskovec et al. 2009) and by psychologists and primatologists who have proposed a limit on the size of human social networks (e.g., Dunbar 1998). This fact allows us to strategically allocate modeling effort within a large graph to be concentrated on a relatively small portion of dyads. Since we postulate that the structure of within-community relations will typically be most complex and interesting, we desire a model flexible enough to differentially devote modeling effort to those relations.

Restricting community sizes to be bounded is also consistent with Schweinberger and Handcock (2014), though their motivation is quite different. Following similar justification as in the time series and spatial contexts, Schweinberger and Handcock (2014) defined a decomposition of the graph such that the propensity to form ties between any set of nodes depends only on a finite number of other nodes. Along with global sparsity, the finite communities assumption in Schweinberger and Handcock (2014) facilitates their asymptotic normality results for graph statistics.

The network ties among a sample of N actors can be represented as an $N \times N$ symmetric matrix $\mathbb{Y}_N$, with (i, j) entry $y_{ij} \in \{0, 1\}$ denoting the absence or presence of a tie between actors i and j . We focus on undirected relations, restricting $y_{ij} = y_{ji}$ and assume the relation between an actor and itself, y_{ii} , is undefined. The partition γ of the actors then induces a partition of the network $\mathbb{Y}_N$ into *blocks* $\mathbb{Y}_{N,kl} = \{y_{ij}: 1 \leq i < j \leq N, i \in S_k, j \in S_l\}$. We call $\mathbb{Y}_{N,kk}$ a *within-community* block and $\mathbb{Y}_{N,kl}$, where $k \neq l$, a *between-community* block. We define *multiresolution network models* as the class of distributions over $\mathbb{Y}_N$ such that, for a specific vector of γ ’s, each distribution in the class can be expressed as

$$P_{\gamma, \alpha, \omega, N}(\mathbb{Y}_N) = \prod_{k=1}^{K_N} W_{\alpha}(\mathbb{Y}_{N,kk}) \prod_{k=1}^{K_N-1} \prod_{l=k+1}^{K_N} B_{\omega}(\mathbb{Y}_{N,kl}), \quad (1)$$

where W is the probability distribution depending on α associated with the within-community model, and B is the probability distribution depending on ω associated with the between-community model. The population parameters α and ω characterize the distribution of within-block and between-block structure, respectively. Furthermore, we express the within-community and between-community probability distributions as mixture distributions

$$W_{\alpha}(\mathbb{Y}_{N,kk}) = \int W(\mathbb{Y}_{N,kk} | \eta_k) dR_{\alpha}(\eta_k) \quad (2)$$

$$B_{\omega}(\mathbb{Y}_{N,kl}) = \int B(\mathbb{Y}_{N,kl} | \tau_{kl}) dS_{\omega}(\tau_{kl}) \quad (3)$$

and require the functional form of $B_{\omega}(\cdot)$ and $W_{\alpha}(\cdot)$ do not depend on the size of the network, N . Thus η_k is the within-community random effect for community k with random effect distribution R_{α} and τ_{kl} is the between-community random effect for the pair of communities k, l with distribution S_{ω} . The dimension of both η_k and τ_{kl} do not depend on the sizes of the respective blocks as they come from common distributions R_{α} and S_{ω} . We assume that both W and B are projective, which is in contrast to the Schweinberger and Handcock (2014) strategy of assuming ERGMs as the between and within block distributions, which are generally not projective. Projectivity of W and B is essential for coherent inference based on the model; the importance of this property is detailed in Section 4.

In our probabilistic framework, we assume communities are exchangeable. That is, we assume that the community labels can be arbitrarily permuted and the probability of the network remains unchanged, or equivalently that there is no information about the social structure of the network contained in the specific values of the community labels. We also assume that the node labels within each community are exchangeable. A familiar case where exchangeability does not hold is network data collected via snowball sampling, where nodes are progressively sampled by following ties in the network and nodes close together in the sampling order are likely to be connected.

We call the models in (1) *multiresolution* because the model parameters and random effects correspond to parameters at three resolutions. At the coarsest (global) level, γ defines the distribution of community sizes. Related literature on community detection defines groups based on subgraph densities, such that actors have a higher propensity to interact within the group than between groups (Newman 2006), or based on *stochastic equivalence*, where groups include actors that display similar interaction patterns to the rest of the network (Lorrain and White 1971). The most well-known models for community identification is the stochastic block model (SBM) and its variants (Nowicki and Snijders 2001; Airoldi et al. 2008; Rohe et al. 2011; Choi, Wolfe, and Airoldi 2012; Amini et al. 2013). While these methods distinguish between clusters of actors and their aggregate structure at a macro-level, they lack the ability to encode low-level structure such as transitivity, which manifests as triangles in the network. Recent extensions of SBMs include “multiscale” versions (e.g.,

Peixoto 2014; Lyzinski et al. 2015), which repeatedly subdivide clusters. These models also fall within our general class.

At the local level, α represents a (multivariate) population parameter determining the distribution of within-community structure across the population, where we might expect the richest structure. Recall the discussion in Section 1.1 about the LPCM and the inherent constraints on the model tie propensities due to the latent space embedding. When multiple dense pockets exist in a large, overall sparse network, the LPCM can sacrifice accuracy in characterizing within cluster structure to distinguish between the clusters. Vivar and Banks (2012) documented a case involving baboon interactions where the separation of the dense communities in a troop dominates the latent space positioning, crowding out within group structure. Our model resolves this issue by disentangling the modeling of within and between group relations. In particular, ω characterizes the distribution of between-community relational structure across the population, and hence controls the overall sparsity and small-world properties of the network.

While γ , α , and ω summarize global structure, at a finer resolution, η_k and τ_{kl} represent the unique structure present within and between specific communities. Finally, at the finest level of resolution, any actor-specific latent variables apart of each within-community distribution W or each between-community relation distribution B provide local representations of any involutive structure. In the next section, we provide more concrete examples of these parameters in the context of popular network models.

The general class of multiresolution models defined in (1) contains a diverse set of possible model specifications. The stochastic blockmodel (Holland, Laskey, and Leinhardt 1983), for example, is a special case. The stochastic blockmodel decomposes a network into communities and models the probability of a tie between any two actors as solely a function of their community memberships. In our formulation, η_k would denote the tie probability between two individuals within community k and τ_{kl} would denote the tie probability between an actor in community k and an actor in community l . Viewing these block-level probabilities as random effects, R_α and S_ω represent the mixing distribution governing the distributions of these effects, and W and B represent products of independent and identically distributed Bernoulli distributions. We could also construct a model that nests stochastic blockmodels within one-another (Peixoto 2014). With this approach, we would represent both within- and between-community structure as a stochastic blockmodel. Greater nesting depths might be specified for the within-community stochastic blockmodels to capture more complex patterns within communities. Furthermore, separate random effects for sender and receiver effects could be added to each block, as in the social relations model (Kenny and La Voie 1984). In the following section, we explore another example of our model class, which we call the Latent Space Stochastic Blockmodel.

3. Latent Space Stochastic Blockmodel

In this section, we introduce a particular multiresolution network model, called the Latent Space Stochastic Blockmodel (LS-SBM). In the LS-SBM, the propensity for within-block ties is modeled with a latent space model and the between-block ties are modeled as in a

stochastic blockmodel. We denote the probability an actor belongs to block k as π_k , and let $\boldsymbol{\pi} = (\pi_1, \dots, \pi_{K_N})$ denote the vector of membership probabilities, where $\sum_{i=1}^{K_N} \pi_i = 1$.

The LS-SBM uses a latent Euclidean distance model (Hoff, Raftery, and Handcock 2002) for the within-community distribution W . In this model, the edges in $\mathbb{Y}_{N,kk}$ are conditionally independent given the latent positions of the actors in S_k . Specifically, given $\mathbf{Z}_i$ and $\mathbf{Z}_j$ where $i, j \in S_k$, Y_{ij} is Bernoulli with probability $\text{logit}^{-1}(\beta_k - \|\mathbf{Z}_i - \mathbf{Z}_j\|)$, where $\|\cdot\|$ denotes the ℓ_2 -norm, that is, Euclidean distance. The latent positions in group k are themselves independent and identically distributed (IID) as spherically normal with mean zero and variance $\sigma_k^2: N_D(\mathbf{Z}_i; 0, \sigma_k^2 I_D)$. Thus,

$$W(\mathbb{Y}_{N,kk} | \eta_k) = \int \left(\prod_{i,j} G(Y_{ij}; \beta_k, \mathbf{Z}_i, \mathbf{Z}_j) \right) \times \left(\prod_i dN_D(\mathbf{Z}_i; 0, \sigma_k^2 I_D) \right), \quad (4)$$

where G is the Bernoulli distribution stated above and the products are taken with respect to all nodes in block k .

In the terminology of multiresolution models, $\eta_k \equiv (\beta_k, \log \sigma_k)$ is the within-community random effect governing the network structure within community k . β_k can be interpreted as the maximum logit-probability of a relation in block k : two nodes i and j are stochastically equivalent in block k if and only if $\mathbf{Z}_i = \mathbf{Z}_j$ in which case $P(Y_{ij} | \mathbf{Z}_i = \mathbf{Z}_j) = \text{logit}^{-1}(\beta_k)$. σ_k is a measure of heterogeneity in block k , as $\sigma_k = 0$ is equivalent to an Erdős-Rényi model with tie probability $\text{logit}^{-1}(\beta_k)$. If $\sigma_k = 0$ for all blocks, then the LS-SBM model reduces to the stochastic blockmodel. We model η_k for all $k = 1, \dots, K_N$ as samples from a bivariate normal with parameters $\boldsymbol{\alpha} = \{\boldsymbol{\mu}, \boldsymbol{\Sigma}\}$. Thus, $R_{\boldsymbol{\alpha}}(\eta_k) = N_2((\beta_k, \log \sigma_k); \boldsymbol{\mu}, \boldsymbol{\Sigma})$.

For the between-community distribution B , we use an Erdős-Rényi model. That is, all edges between communities k and l are IID with probability τ_{kl} . Thus,

$$B(\mathbb{Y}_{N,kl} | \tau_{kl}) = \prod_{i \in S_k} \prod_{j \in S_l} \tau_{kl}^{Y_{ij}} (1 - \tau_{kl})^{1 - Y_{ij}}. \quad (5)$$

We model τ_{kl} as Beta distributed with parameters $\boldsymbol{\omega} = (a_0, b_0)$: $S_{\boldsymbol{\omega}}(\tau_{kl}) = \text{Beta}(\tau_{kl}; a_0, b_0)$. The model maintains a parsimonious structure in modeling relationships between blocks, requiring only a single parameter, but is flexible in modeling ties within each block, allowing tie prevalence to depend on the distance between actors in the unobserved social space.

There are two key differences between the proposed LS-SBM and the latent position cluster model (LPCM) introduced in Handcock, Raftery, and Tantrum (2007). The first key distinction is that the probability of a tie between actors belonging to different communities in the LS-SBM is a function of only their community memberships, whereas in the LPCM it

is a function of the distance between the actor positions in the latent space. This means that in the LPCM, all of an actor's ties and nonties are used in determining the latent positions. In contrast, in the LS-SBM, the latent space only affects within-community connections. As a result, the structure of between-community connections are not constrained by the dimension and geometry of the latent space in the LS-SBM like they are in the LPCM. The second distinction between the models is that the LPCM contains a single intercept parameter and the LS-SBM contains block-specific intercepts, β_k . Each intercept β_k can be interpreted as the maximum logit-probability of a tie in community k . In practice, we find there is often large heterogeneity in this maximum probability across communities, suggesting having different intercepts is a critical piece of model flexibility in the LS-SBM.

3.1. Prior Specification

Here, we discuss the prior distributions for α and ω . Our intended application of the LS-SBM is to networks where the within-community ties are denser than the between-community ties. We use the prior on α to reflect this knowledge. Specifically, we set the prior on μ and Σ to be a conjugate Normal-Inverse-Wishart distribution, with parameters $\{\mathbf{m}_0, s_0, \Psi_0, \nu_0\}$, subject to an additional *assortativity restriction*. Given a_0 and b_0 , the prior can be expressed as

$$P(\alpha | a_0, b_0, \mathbf{m}_0, s_0, \Psi_0, \nu_0) \propto N_2(\mu; \mathbf{m}_0, \Sigma_0/s_0) \times \text{Inv. Wish}(\Sigma; \Psi_0, \nu_0) \mathbf{1}(a_0, b_0, \mu), \quad (6)$$

where $\mathbf{1}(a_0, b_0, \mu)$ is the indicator function enforcing the assortativity condition. We fix a_0 and b_0 based on the observed density of the graph. The assortativity condition we require is that the (logit) marginal probability of a within-community tie for the *average block*, induced by μ , be larger than the (logit) average between-block probability of a tie, $a_0/(a_0 + b_0)$:

$$E[\text{logit}(Pr(Y_{ij} = 1)) | \gamma_i = \gamma_j] \geq E[\text{logit}(Pr(Y_{ij} = 1)) | \gamma_i \neq \gamma_j]. \quad (7)$$

Calculating these expectations (see the web-based supplementary materials for details), the restriction on the population parameter space we wish to enforce is

$$\mu_1 - 2e^{\mu_2 \frac{\Gamma(\frac{D+1}{2})}{\Gamma(\frac{D}{2})}} \geq \psi(a_0) - \psi(b_0),$$

where $\psi(x)$ is the digamma function defined as $\psi(x) = \frac{d \log(\Gamma(x))}{dx}$. The digamma function is not available in closed form but is easily approximated with most standard statistical software packages. We proceed with estimation using this global assortativity restriction.

3.2. Estimation and Block Number Selection

Here, we provide a brief summary of our estimation procedure for the LS-SBM. A full description of the model specification and algorithm are provided in the web-based supplementary materials. The posterior for our Bayesian model is not available in closed form, so instead we approximate the posterior using draws obtained via Markov chain Monte Carlo (MCMC). The MCMC algorithm performs the estimation with the number of blocks K fixed. Thus, we first describe a procedure for choosing K and then outline the MCMC procedure given the number of blocks.

We suggest comparing different K using a series of 10-fold cross-validation procedures. For each repetition of the procedure, randomly partition the unordered node pairs into 10 folds. For each fold, use assortative spectral clustering (Saade, Krzakala, and Zdeborová 2014) on the adjacency matrix, excluding that fold, to partition the nodes into numbers of blocks K from 2 to $\lfloor N/4 \rfloor$. To adapt the spectral clustering algorithm to deal with the held out, missing at random, edges, we propose using an iterative EM-like scheme where first the missing values are imputed using observed degrees, clustering is performed on the imputed data, and the missing values from hold-out are reimputed using the predicted probabilities. This should be repeated, until convergence. This procedure does not require computing the full posterior and can be done in parallel for each value of K and each validation fold.

For each repetition and K value, we propose calculating three metrics of predictive performance: area under the ROC curve (AUC), mean squared error (MSE), and mean predictive information (MPI). Calculate the mean value and 95% CIs for the mean of these criteria over the repetitions for each K . Then, for each criteria, we suggest finding the *smallest* K such that the mean value of the criteria for K falls in the 95% CI of the mean value of the K with the best mean value (either maximal or minimal depending on the criterion).

Once we have selected a value of K , we use a Metropolis-within-Gibbs algorithm to approximate the joint posterior distribution. Our use of conjugate priors allows Gibbs updates of $\boldsymbol{\pi}$, $\boldsymbol{\tau}$, $\boldsymbol{\mu}$, and $\boldsymbol{\Sigma}$. We update each η_k with a Metropolis step, using a bivariate normal proposal distribution.

Each node is assigned a single block membership at every iteration of the chain. This block membership is jointly updated with the node's latent position using a Metropolis step. We also take additional (unsaved) Metropolis steps for the latent positions to allow nodes which have switched blocks to find higher likelihood points in the latent space.

The likelihood is invariant to permutations of the block memberships and to rotations and reflections of the latent spaces. We address these nonidentifiabilities by post-processing the posterior samples using equivalence classes defined over the parameter spaces. See the supplementary materials for additional details.

In our experiments, computation was feasible for networks with 300 nodes in under 2 hr using a personal computer with a 2GHz processor and 8GB of RAM. The most computationally expensive piece is the Metropolis step that jointly updates block

membership and latent positions and then subsequently takes additional draws from latent spaces. In our experiments, however, using a joint update substantially improved mixing. Since we do not use MCMC to compute K , this is not a limiting step computationally. To scale our method beyond what is possible with MCMC, or for cases where the full posterior is not of interest, we provide a two-stage fitting procedure in the supplementary material. This procedure uses an assortative graph clustering algorithm to quickly estimate block membership, and variational inference to estimate parameters within each block. We have used this two-stage procedure to estimate our model in a sparse network with 13,000 nodes, which took about 3 min on a standard personal computer. The results of this analysis are provided in Section 3.3. Additional details about the two-stage procedure are provided in the supplementary material.

3.3. Karnataka Villages

We estimate the proposed LS-SBM on data from a social network study consisting of households in villages in Karnataka, India to illustrate the utility of the model. These data were collected as part of an experiment to evaluate a micro-finance program performed by Banerjee et al. (2013). Data consist of multiple undirected relationships between individuals and households in 75 villages. Relationship types include social and familial interactions (e.g., being related or attending temple together) and views related to economic activity (e.g., lending money or borrowing rice/kerosene).

We used the household-level “visit” relation from village 59, which has $N = 293$ households with nonzero degree. We estimate the LS-SBM on the data with $K = 6$. Details of the cross-validation selection procedure for K and LS-SBM estimation on this data are provided in the supplementary materials. Codes to replicate the results we present here are available at https://github.com/tedwestling/multiresolution_networks.git. Data are available at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=hdl:1902.1/21538>.

The results of the estimation algorithm are displayed in Figures 1–3, and in Table 1. Figure 1 shows the observed adjacency matrix organized by marginal posterior mode block membership. The estimated block memberships result in an assortative network structure: there are more ties between households in the same community than between households in different communities. Table 1 shows the posterior mean estimates of the between block connectivity parameters, again illustrating the assortative patterns. The within-community ties seen in Figure 1 are fairly clearly nonuniform, further justifying our departure from an SBM for the within-community model.

Figure 2 shows the estimated block-level parameters $\eta_k = (\beta_k, \log \sigma_k)$ and the global mean $\mu = (\mu_1, \mu_2)$. Recall that β_k is the intercept parameter for each block and $\log \sigma_k$ describes the variation in the latent space. The μ_1 and μ_2 terms describe the mean of the distribution of β_k and $\log \sigma_k$, respectively. Since the multiresolution framework is projective, we can compare parameters between block-level parameters. Blocks two, five, and six, denoted as B2, B5, and B6, respectively, display strong a posteriori differences from the mean block. In B2, the posterior distribution is shifted toward a larger intercept parameter, with posterior probability approximately 97% that $\beta_2 > \mu_1$. The larger intercept indicates that the overall propensity for ties is larger in B2. In addition, the posterior probability that $\log \sigma_2$ is greater

than the overall mean μ_2 is greater than 99%, suggesting there is greater heterogeneity in tie probabilities in block two compared to the global mean. In B6, the β_6 term is shifted substantially lower than the overall mean μ_1 , indicating a smaller maximum tie probability. The posterior probability that $\beta_6 < \mu_1$ is greater than 99%.

We now further explore the structure implied within each block by examining the within-block latent positions. Figure 3 shows the latent positions (after accounting for nonidentifiability issues) obtained from multidimensional scaling on the posterior mean distance matrix, along with the observed edges within each block. The posterior places a distribution over block memberships for each node, however in Figure 3 we show nodes only in the block for which they have the largest posterior probability of inclusion. Individuals that are unlikely to belong to any specific block (with inclusion probabilities less than 30%) are omitted. Shading represents the concentration of the posterior over block memberships, with darker colors indicating higher assignment probabilities.

Moving to the structure within the blocks, we investigate the relation between the household memberships, positions, and caste which is a formalized social class system in India. Castes are denoted in Figure 3 using different colors. We see strong sorting by caste, with almost all of the members of schedule castes and schedule tribes (the two lowest castes) being grouped into B6. Members of the slightly higher class OBC (Other Backwards Caste) are the most common in the network and are spread throughout the remaining blocks. An extensive literature in economics (e.g., Townsend 1994; Munshi and Rosenzweig 2009; Mazzocco 2012; Ambrus, Chandrasekhar, and Elliott 2014) explores on the role of the caste system in individuals' financial decisions. In particular, this literature focuses on informal credit markets. That is, the social structures that provide financial support in times of need without a formal, corporate credit structure. Recent work by Ambrus, Chandrasekhar, and Elliott (2014) present a theoretical argument for the importance of ties that bridge otherwise disconnected groups. In our results, these individuals would be individuals whose block assignment based on their social interactions does not match that of others in their caste. For example, this group of bridging individuals would include members of schedule tribes or castes that are in blocks other than B6.

We also used our two-stage procedure (described in detail in the supplementary material) to estimate our model for all 75 village networks combined. We formed an undirected network of $N = 13,009$ nodes by combining all 75 household-level “visit” relation networks from the Karnataka village data. We estimated the block structure using label propagation (Raghavan, Albert, and Kumara 2007), which returned 534 blocks. Every block contained only households from a single village—that is, there were no blocks containing households from multiple villages. This was expected since by design there are no between-village edges. There were a median of six blocks per village, with as few as one block per village (i.e., the entire village constitutes a single block) and as many as 20 blocks per village. The number of nodes per block varied considerably, with a median of 14, mean of 24.4, and maximum of 233. The density of edges within a block and nodes per block were well-described by a linear function on the log-log scale with intercept 0.34 and slope -0.82 , as shown in the left panel of Figure 4. The estimated within-block latent space parameters $\log \sigma$ and β are shown

in the right panel of Figure 4. The larger blocks tend to be sparser and more heterogeneous, while the smaller blocks are more homogenous.

3.4. Simulation Study

In this section, we detail a simulation study illustrating the advantages of using a multiresolution model like the LS-SBM over existing models such as the latent space model, LPCM and SBM.

Binary, undirected network data were generated for 300 nodes from the LS-SBM model with five equally sized blocks such that the between-block tie probabilities were either 0.2 or 0.02. Within-block tie probabilities stemmed from a heterogeneous set of two-dimensional block-specific latent spaces. Further details about the simulation parameters are provided in the supplementary materials. One thousand simulations were performed where 10% of the undirected dyads in the network were held out in each simulation and the models were fit to the remaining 90% of the data. Predictions were then made for the held out portion of the network and the accuracy of these predictions quantified by computing the area under the precision-recall curves. The results are shown in Figure 5.

From the leftmost panel of Figure 5, it is evident that the LS-SBM outperforms all three existing models as the relative AUPRCs are greater than one for all simulations and all network models. Separate precision-recall curves were constructed for the held out portions of the network corresponding to relationships between nodes within the same block (rightmost panel of Figure 5) and those portions between nodes that reside in different blocks (middle panel of Figure 5). These illustrate that while the LS-SBM appears to predict edges between nodes in different blocks as well as existing models, there are notable improvements in predictions for ties between nodes within the same block.

4. Projectivity of Multiresolution Network Models

Focusing on inference, we seek to understand which features of the hypothetical infinite population we can reasonably expect to learn from a sample of N nodes. Projectivity is essential for inference as it facilitates comparison of model parameters across networks of different sizes. This notion of “different sizes” naturally arises in multiple network samples, which are almost certainly never be of the same size (because of the complexities of sampling networks and prevalence of missing data). In the case of the multiresolution framework, these sizes may also refer to the inferred block sizes.

Shalizi and Rinaldo (2013) investigated the projectivity of families of statistical network models, where a model family $\{P_{\theta, N} : N \in \mathbb{N}, \theta \in \Theta\}$ is deemed projective if distribution $P_{\theta, N}$ for a sample of N actors can be recovered by marginalizing the distribution $P_{\theta, M}$ for $N < M$, over actors $\{N+1, \dots, M\}$. Stated more formally, a family of network models is projective if $P_{\theta, N} = P_{\theta, M} \circ \pi_{M \rightarrow N}^{-1}$ for all $N < M < \infty$, where $\pi_{M \rightarrow N}$ is the natural projection map that selects the subgraph on the first N nodes from the full graph on M nodes and $\circ$ denotes function composition. Letting $\mathbb{Y}_{M \setminus N}$ be $\mathbb{Y}_M$ after removing the $\mathbb{Y}_N$ sub-graph and $\mathcal{Y}_{M \setminus N}$ be its sample space, we can write

$$\begin{aligned}
P_{\theta, N}(\mathbb{Y}_N) &= P_{\theta, M}(\pi_{M \mapsto N}^{-1}(\mathbb{Y}_N)) = \sum_{\mathbb{Y}_M \in \pi_{M \mapsto N}^{-1}(\mathbb{Y}_N)} P_{\theta, M}(\mathbb{Y}_M) \\
&= P_{\theta, M}(\mathbb{Y}_N, \mathbb{Y}_{M \setminus N} \in \mathcal{Y}_{M \setminus N}),
\end{aligned}$$

where $\pi_{M \mapsto N}^{-1}(\mathbb{Y}_N)$ is the set of graphs on $\{1, \dots, M\}$ that have $\mathbb{Y}_N$ as the subgraph on the first N actors.

To see why projectivity is crucial for comparisons across networks of different sizes, or equivalently blocks, recall the Karnataka dataset. Suppose we have a model family and consider two village networks: village A network containing 100 households and village B network containing 100,000 households. Upon observing these networks, a researcher wishes to formally compare them by fitting the statistical model to each one and comparing the parameter estimates. In order for this comparison to be meaningful, the statistical model *must* be projective. Suppose the parameter associated with the generation of network A is θ_A , the parameter generating network B is θ_B , and $\theta_B = \theta_A$. The statistical model is projective if, when 90,900 households are marginalized over in the network model fit to network B, the resulting probability model on the remaining 100 households is equal to the model on network A. (Note in this discussion, we assume the probability model is row and column exchangeable, that is, node exchangeable as in von Plato (1991).)

Our multiresolution framework proposes projective models, W and B , for capturing within- and between-community relations, and combines these to form a model for the entire network. We demonstrate below that a model class defined using combinations of projective distributions forms a projective family of models. This permits researchers to make coherent comparisons across communities within the same network, even if communities are of different sizes.

We start by proving that mixtures of projective models are projective. In the definition of multiresolution network models in (1)–(3), we assume that W and B are projective. Further, since W_α and B_ω are not indexed by N , R_α , and S_ω must be the same regardless of the number of nodes in the graph. Below we show that this implies that W_α and B_ω are projective by showing that, in general, a mixture of projective models is also projective.

Theorem 1. Suppose $\{\tilde{P}_{\theta, M} : M \in \mathbb{N}, \theta \in \Theta\}$ is a projective collection of statistical models over networks and latent variables, such that $\tilde{P}_{\theta, M}$ is a distribution on $(\mathbb{Y}_{M, \eta})$ supported over $\mathcal{Y}_M \times \mathcal{N}$ where the dimension of $\mathcal{N}$ does not depend on M . Let $P_{\theta, M} = \tilde{P}_{\theta, M} \circ \tau_M^{-1}$ for $\tau_M : \mathcal{Y}_M \times \mathcal{N} \rightarrow \mathcal{Y}_M$ the projection map. Then $\{P_{\theta, M} : M \in \mathbb{N}, \theta \in \Theta\}$ is a projective family as well.

Proof. Let $N < M$. Further, let $\tilde{\pi}_{M \mapsto N}$ be the projection map from $\mathcal{Y}_M \times \mathcal{N} \rightarrow \mathcal{Y}_N \times \mathcal{N}$ and $\pi_{M \mapsto N}$ be the projection map from $\mathcal{Y}_M$ to $\mathcal{Y}_N$. Let us first suppose that

$$\tilde{\pi}_{M \mapsto N}^{-1} \circ \tau_N^{-1} = \tau_M^{-1} \circ \pi_{M \mapsto N}^{-1}. \quad (8)$$

Then, $P_{\theta, N} = \tilde{P}_{\theta, N} \circ \tau_N^{-1} = \tilde{P}_{\theta, M} \circ \tilde{\pi}_{M \mapsto N}^{-1} \circ \tau_N^{-1} = \tilde{P}_{\theta, M} \circ \tau_M^{-1} \circ \pi_{M \mapsto N}^{-1} = P_{\theta, M} \circ \pi_{M \mapsto N}^{-1}$, where the second equality follows from the projectivity of $\{\tilde{P}_{\theta, M} : M \in \mathbb{N}, \theta \in \Theta\}$ and third equality is a consequence of (8). Thus, if (8) holds, $\{P_{\theta, M} : M \in \mathbb{N}, \theta \in \Theta\}$ is projective by definition.

Verifying (8) is straightforward. Let $\mathbb{Y}_N \in \mathcal{Y}_N$. Then

$$\begin{aligned} (\tilde{\pi}_{M \mapsto N}^{-1} \circ \tau_N^{-1})(\mathbb{Y}_N) &= \tilde{\pi}_{M \mapsto N}^{-1}(\{(\mathbb{Y}_N, \eta) : \eta \in \mathcal{N}\}) \\ &= \{(\mathbb{Y}_N, \mathbb{Y}_{M \setminus N}, \eta) : \mathbb{Y}_{M \setminus N} \in \mathcal{Y}_{M \setminus N}, \eta \in \mathcal{N}\}. \end{aligned}$$

Similarly,

$$\begin{aligned} (\tau_M^{-1} \circ \pi_{M \mapsto N}^{-1})(\mathbb{Y}_N) &= \tau_M^{-1}(\{(\mathbb{Y}_N, \mathbb{Y}_{M \setminus N}) : \mathbb{Y}_{M \setminus N} \in \mathcal{Y}_{M \setminus N}\}) \\ &= \{(\mathbb{Y}_N, \mathbb{Y}_{M \setminus N}, \eta) : \mathbb{Y}_{M \setminus N} \in \mathcal{Y}_{M \setminus N}, \eta \in \mathcal{N}\}. \end{aligned}$$

□

By a similar argument, we can show that node-level latent variable models, such as the latent space network model specified for the within-block ties in the LS-SBM, are also projective. Using Theorem 1, we now show that the class of multiresolution models is projective.

Theorem 2. Multiresolution network models are projective.

Proof. Since W and B are projective models, the models $W_\alpha(\mathbb{Y}_{N, kk})$ and $B_\omega(\mathbb{Y}_{N, kl})$ are then projective by Theorem 1 because their distributions do not depend on N and they are mixtures over projective models. Consider

$$\begin{aligned} P_{\gamma, \alpha, \omega, M}(\mathbb{Y}_N, \mathbb{Y}_{M \setminus N} \in \mathcal{Y}_{M \setminus N}) &= \prod_{k=1}^{K_M} W_\alpha(\mathbb{Y}_{N, kk}, \mathbb{Y}_{M \setminus N, kk} \in \mathcal{Y}_{M \setminus N, kk}) \quad (9) \\ &\times \prod_{k=1}^{K_M-1} \prod_{l=k+1}^{K_M} B_\omega(\mathbb{Y}_{N, kl}, \mathbb{Y}_{M \setminus N, kl} \in \mathcal{Y}_{M \setminus N, kl}). \end{aligned}$$

For any k, l such that none of the nodes in blocks k, l are in $\mathbb{Y}_N$, we have

$W_\alpha(\mathbb{Y}_{N, kk}, \mathbb{Y}_{M \setminus N, kk} \in \mathcal{Y}_{M \setminus N, kk}) = W_\alpha(\mathbb{Y}_{M \setminus N, kk} \in \mathcal{Y}_{M \setminus N, kk}) = 1$ and similarly for B_ω . Hence, we only need consider blocks with at least one node from $\mathbb{Y}_N$, and the right-hand side of (9) is equal to

$$\prod_{k=1}^{K_N} W_{\alpha}(\mathbb{Y}_{N,kk}, \mathbb{Y}_{M \setminus N,kk}) \prod_{k=1}^{K_N-1} \prod_{l=k+1}^{K_N} B_{\omega}(\mathbb{Y}_{N,kl}, \mathbb{Y}_{M \setminus N,kl} \in \mathcal{Y}_{M \setminus N,kl}).$$

Since W_{α} and B_{ω} are projective, $W_{\alpha}(\mathbb{Y}_{N,kk}, \mathbb{Y}_{M \setminus N,kk} \in \mathcal{Y}_{M \setminus N,kk}) = W_{\alpha}(\mathbb{Y}_{N,kk})$, and similarly for B_{ω} . Thus, the right-hand side of (9) is equal to

$$\prod_{k=1}^{K_N} W_{\alpha}(\mathbb{Y}_{N,kk}) \prod_{k=1}^{K_N-1} \prod_{l=k+1}^{K_N} B_{\omega}(\mathbb{Y}_{N,kl}),$$

which equals $P_{\gamma, \alpha, \omega, N}(\mathbb{Y}_N)$. $\square$

We emphasize that projectivity of the multiresolution model only occurs when B_{ω} and W_{α} are both projective, and hence these distributions do not depend on N . If the between-block tie probabilities depend on the size of the observed graph, then the model is not projective overall but is projective within each block. By “projective within each block” we mean that the within-block parameters are still comparable across blocks and networks of different sizes, although it is meaningless to, in general, compare the between-block parameters. An advantage of the between-block parameters depending on N is that the model can be sparse in the limit. That is, as the number of nodes grows from N to infinity, the expected density (i.e., proportion of edges present in the graph) goes to zero. Equivalently, a graph is asymptotically sparse if the average expected degree grows sub-linearly with N . The Aldous-Hoover Theorem implies that infinitely exchangeable sequences of nodes correspond to dense graphs (Aldous 1981; Orbanz and Roy 2015). In our case, however, for a fixed N and γ , the multiresolution model is only exchangeable modulo γ , meaning that nodes within the same block are exchangeable (similar to that in regression; see McCullagh 2005). If we assume each node has an equal probability of belonging to any block (e.g., placing a uniform distribution on each γ_i), the model is finitely exchangeable (von Plato 1991). However, we are unaware of a prior distribution on γ that gives both finite exchangeability and projectivity, but does not assume knowledge about a bound on the size of blocks. In the web-based supplementary material we show that families of projective models which are finitely exchangeable are asymptotically dense. Nevertheless, for practical purposes, our model can be arbitrarily sparse for any network with a finite number of nodes and, since it is projective, permit comparison between networks of different sizes.

Our results on exchangeability and projectivity relate to recent work on nonparametric generative models for networks (e.g., Caron and Fox 2014; Veitch and Roy 2015; Crane and Dempsey 2015, 2016; or Broderick and Cai 2016). The objectives in our framework are subtly but critically different, however. In recent developments in nonparametrics, the objective is to understand the structure of the graph that is implied by a probability model as the network grows from the observed size N to its limit. These recent works define new notions of exchangeability that imply critically different graph properties than node exchangeability discussed here. Our work, in contrast, focuses on the inverse inference

paradigm: given a sample from an infinite population, our goal is to understand which properties of the population could be feasibly estimated using the multiresolution model. This perspective leads to a focus on projectivity.

5. Discussion

In this article, we present a multiresolution model for social network data. Our model is well-suited for graphs that are overall very sparse, but contain pockets of local density. Our model uses mixtures of projective models to separately characterize tie structure within and between dense pockets in the graph. Our proposed framework is substantially more flexible than existing latent variable approaches (such as the LPCM) and supports meaningful comparisons of parameter estimates across communities and networks of varying sizes.

We introduced the LS-SBM as one example of a model within the multiresolution class. However, alternative multiresolution models could be defined by replacing the latent space model representing within-community relations with LPCMs or LS-SBMs. This would add complexity to the model, allowing for sub-community structure within the global-community structure. A key distinction between the general multiresolution framework propose here and previous approaches of, for example, Peixoto and Lyzinski et al. is that our approach allows and suggests using different network models to capture structure at different levels. This allows models to be constructed that leverage the advantages (e.g., parsimony, detailed structure) of multiple models simultaneously.

Our work also contributes to an active discussion in the statistics literature emphasizing the importance of understanding the relationship between sampling and modeling social networks. Crane and Dempsey (2015), for example, presented a general framework for sampling and inference for network models. Our work emphasizes the importance of “consistency under sampling” for comparison across networks and across communities within the same graph.

Our projectivity result for multiresolution models means that our framework can be used to compare across communities within a graph, even if the communities are different sizes. Schweinberger and Handcock (2014) defined the concept of domain consistency, which can be thought of as projectivity over communities, and proposed a class of models that satisfy this property. If, for example, a member of a community is missing, then the interpretation of the parameters describing behavior in that community will be fundamentally different than if the member were present. Our model also has this property, but is more general. In particular, the notion of projectivity over communities requires that data be collected by using cluster sampling over communities. Here, we strengthened and generalized this framework by proving that our model is projective at the actor level, that is, even if complete communities are not observed. Using our framework, it is possible to sample at the actor rather than the community level while estimating parameters that are comparable across communities.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Funding

This work was partially supported by the National Science Foundation under grant number SES-1461495 to Fosdick and grant number SES-1559778 to McCormick. McCormick is also supported by grant number K01 HD078452 from the National Institute of Child Health and Human Development (NICHD). This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF-12-1-0379. Murphy and Ng are supported by the Science Foundation Ireland funded Insight Research Centre (SFI/12/RC/2289). The authors would also like to thank the Isaac Newton Institute Program on Theoretical Foundations for Statistical Network Analysis workshop on Bayesian Models for Networks, supported by EPSRC grant number EP/K032208/1.

References

- Airoldi EM, Blei DM, Fienberg SE, and Xing EP (2008), “Mixed Membership Stochastic Blockmodels,” *Journal of Machine Learning Research*, 9, 1981–2014. [PubMed: 21701698]
- Aldous DJ (1981), “Representations for Partially Exchangeable Arrays of Random Variables,” *Journal of Multivariate Analysis*, 11, 581–598.
- Ambrus A, Chandrasekhar AG, and Elliott M (2014), “Social Investments, Informal Risk Sharing, and Inequality,” Technical Report, National Bureau of Economic Research, Working Paper No. 20669.
- Amini AA, Chen A, Bickel PJ, and Levina E (2013), “Pseudo-Likelihood Methods for Community Detection in Large Sparse Networks,” *The Annals of Statistics*, 41, 2097–2122.
- Banerjee A, Chandrasekhar AG, Duflo E, and Jackson MO (2013), “The Diffusion of Microfinance,” *Science*, 341, , doi: 10.1126/science.1236498 [PubMed: 1236498]
- Borgatti SP, Mehra A, Brass DJ, and Labianca G (2009), “Network Analysis in the Social Sciences,” *Science*, 323, 892–895. [PubMed: 19213908]
- Broderick T, and Cai D (2016), “Edge-Exchangeable Graphs and Sparsity,” arXiv preprint arXiv:1603.06898.
- Caron F, and Fox EB (2014), “Sparse Graphs using Exchangeable Random Measures,” arXiv preprint arXiv:1401.1137.
- Chatterjee S, and Diaconis P (2013), “Estimating and Understanding Exponential Random Graph Models,” *The Annals of Statistics*, 41, 2428–2461.
- Choi DS, Wolfe PJ, and Airoldi EM (2012), “Stochastic Blockmodels with a Growing Number of Classes,” *Biometrika*, 99, 273–284. [PubMed: 23843660]
- Crane H, and Dempsey W (2015), “A Framework for Statistical Network Modeling,” arXiv preprint arXiv:1509.08185.
- Crane H, and Dempsey W (2016), “Edge Exchangeable Models for Network Data,” arXivpreprint arXiv:1603.04571.
- Cross R, Parker A, Prusak L, and Borgatti SP (2001), “Knowing What We Know,” *Organizational Dynamics*, 30, 100–120.
- Dunbar R (1998), “The Social Brain Hypothesis,” *Brain*, 9, 178–190.
- Fraley C, and Raftery AE (2002), “Model-Based Clustering, Discriminant Analysis, and Density Estimation,” *Journal of the American Statistical Association*, 97, 611–631.
- Frank O, and Strauss D (1986), “Markov Graphs,” *Journal of the American Statistical Association*, 81, 832–842.
- Handcock MS, Raftery AE, and Tantrum JM (2007), “Model-Based Clustering for Social Networks,” *Journal of the Royal Statistical Society, Series A*, 170, 301–354.
- Hoff PD, Raftery AE, and Handcock MS (2002), “Latent Space Approaches to Social Network Analysis,” *Journal of the American Statistical Association*, 97, 1090–1098.
- Holland PW, Laskey KB, and Leinhardt S (1983), “Stochastic Block-models: First Steps,” *Social Networks*, 5, 109–137.
- Kenny DA, and La Voie L (1984), “The Social Relations Model,” *Advances in Experimental Social Psychology*, 18, 141–182.

- Koskinen J, and Melas V (2009), “Using Latent Variables to Account for Heterogeneity in Exponential Family Random Graph Models,” in Proceedings of the 6th St. Petersburg Workshop on Simulation, St Petersburg State University, pp. 845–849.
- Leskovec J, Lang KJ, Dasgupta A, and Mahoney MW (2009), “Community Structure in Large Networks: Natural Cluster Sizes and the Absence of Large Well-Defined Clusters,” *Internet Mathematics*, 6, 29–123.
- Lorrain F, and White HC (1971), “Structural Equivalence of Individuals in Social Networks,” *The Journal of Mathematical Sociology*, 1, 49–80.
- Lusher D, Koskinen J, and Robins G (Eds.). (2012), *Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications (Structural Analysis in the Social Sciences)*, Cambridge: Cambridge University Press.
- Lyzinski V, Tang M, Athreya A, Park Y, and Priebe CE (2015), “Community Detection and Classification in Hierarchical Stochastic Block-models,” arXiv preprint arXiv:1503.02115.
- Mazzocco M (2012), “Testing Efficient Risk Sharing with Heterogeneous Risk Preferences,” *The American Economic Review*, 102, 428–468.
- McCullagh P (2005), “Exchangeability and Regression Models,” *Oxford Statistical Science Series*, 33, 89.
- Munshi K, and Rosenzweig M (2009), “Why is Mobility in India So Low? Social Insurance, Inequality, and Growth,” Technical Report, National Bureau of Economic Research, Working Paper No. 14850.
- Newman ME (2006), “Modularity and Community Structure in Networks,” *Proceedings of the National Academy of Sciences*, 103, 8577–8582.
- Nowicki K, and Snijders TAB (2001), “Estimation and Prediction for Stochastic Blockstructures,” *Journal of the American Statistical Association*, 96, 1077–1087.
- Orbanz P, and Roy DM (2015), “Bayesian Models of Graphs, Arrays and Other Exchangeable Random Structures,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37, 437–461. [PubMed: 26353253]
- Page GL, and Quintana FA (2016), “Spatial Product Partition Models,” *Bayesian Analysis*, 11, 265–298.
- Pattison P, and Wasserman S (1999), “Logit Models and Logistic Regressions for Social Networks: II. Multivariate Relations,” *British Journal of Mathematical and Statistical Psychology*, 52, 169–193. [PubMed: 10613111]
- Peixoto TP (2014), “Hierarchical Block Structures and High-Resolution Model Selection in large Networks,” *Physical Review X*, 4, 011047.
- Raghavan UN, Albert R, and Kumara S (2007), “Near Linear Time Algorithm to Detect Community Structures in Large-Scale Networks,” *Physical Review E*, 76, 036106.
- Robins G (2011), “Exponential Random Graph Models for Social Networks,” in *Encyclopaedia of Complexity and System Science*, New York: Springer.
- Rohe K, Chatterjee S, and Yu B (2011), “Spectral Clustering and the High-Dimensional Stochastic Blockmodel,” *The Annals of Statistics*, 39, 1878–1915.
- Saade A, Krzakala F, and Zdeborová L (2014), “Spectral Clustering of Graphs with the Bethe Hessian,” *Advances in Neural Information Processing Systems*, 406–414.
- Salter-Townshend M, White A, Gollini I, and Murphy TB (2012), “Review of Statistical Network Analysis: Models, Algorithms, and Software,” *Statistical Analysis and Data Mining*, 5, 243–264.
- Schweinberger M, and Handcock MS (2014), “Local Dependence in Random Graph Models: Characterization, Properties and Statistical Inference,” *Journal of the Royal Statistical Society, Series B*, 77, 647–676.
- Shalizi CR, and Rinaldo A (2013), “Consistency Under Sampling of Exponential Random Graph Models,” *The Annals of Statistics*, 41, 508–535. [PubMed: 26166910]
- Snijders TA (2002), “Markov Chain Monte Carlo Estimation of Exponential Random Graph Models,” *Journal of Social Structure*, 3, 1–40.

- Sweet TM, Thomas AC, and Junker BW (2013), “Hierarchical Network Models for Education Research Hierarchical Latent Space Models,” *Journal of Educational and Behavioral Statistics*, 38, 295–318.
- Townsend RM (1994), “Risk and Insurance in Village India,” *Econometrica: Journal of the Econometric Society*, 62, 539–591.
- Ugander J, Karrer B, Backstrom L, and Marlow C (2011), “The Anatomy of the Facebook Social Graph,” *arXiv preprint arXiv:1111.4503*.
- Veitch V, and Roy DM (2015), “The Class of Random Graphs Arising from Exchangeable Random Measures,” *arXiv preprint arXiv:1512.03099*.
- Vivar JC, and Banks D (2012), “Models for Networks: A Cross-Disciplinary Science,” *Wiley Interdisciplinary Reviews: Computational Statistics*, 4, 13–27.
- von Plato J (1991), “Finite Partial Exchangeability,” *Statistics & Probability Letters*, 11, 99–102.
- Wasserman S, and Pattison P (1996), “Logit Models and Logistic Regressions for Social Networks: I. An Introduction to Markov Graphs and p^* ,” *Psychometrika*, 61, 401–425.

Figure 1.
Household-level “visit” relation adjacency matrix of village number 59 from the Karnataka village dataset. Nodes are grouped by marginal posterior mode block membership.

Figure 2.

Block-level and global parameter estimates and 95% highest posterior density (HPD) regions. Each panel contains the posterior mean and 95% HPD region for a single block parameter (β , $\log \sigma$) (solid point and solid line) and the posterior mean and 95% HPD of the global parameter (μ_1 , μ_2) (triangle and dashed line).

Figure 3.

Latent positions within each block. Shading represents posterior probability of inclusion in a given block, with the lightest shading representing a posterior estimate of around 30% and darkest colors representing values near unity. Colors and plot symbols differentiate the five caste categories.

Figure 4.

The left panel shows the $\log_{10}$ block density as a function of $\log_{10}$ block size. The blue line is the OLS linear regression fit. The right panel shows the block-level latent space parameters β_k and $\log(\sigma_k)$, where point size corresponds to block size and point color corresponds to block density. In both plots, each point is an estimated block from the model fit to all 75 Karnataka villages using the two-stage procedure.

Figure 5.

Relative area under the precision-recall curve (AUPRC) based on out-of-sample predictions for the LS-SBM and three existing models: latent space (LS) model, latent position cluster model (LPCM), and stochastic blockmodel (SBM). The leftmost panel shows the relative AUPRC for all held out edges, the middle panel shows the results for edges that are between nodes that are in different blocks, and the right panel shows that for edges between nodes within the same block.

Table 1.

Between block probability matrix. The off-diagonal elements are the posterior mean probabilities of a tie between individuals in different blocks, $\tau_{\gamma_i \gamma_j}$. The diagonal elements represent the maximum probability of a tie within each block based on the block-level parameter β_k posterior mean: $e^{\hat{\beta}_k} / (1 + e^{\hat{\beta}_k})$. Values less than 0.01 are grayed out.

	1	2	3	4	5	6
1	0.719	0.006	0.004	0.044	0.010	0.003
2	0.006	0.951	0.018	0.016	0.006	0.003
3	0.004	0.018	0.802	0.011	0.003	0.002
4	0.044	0.016	0.011	0.727	0.079	0.021
5	0.010	0.006	0.003	0.079	0.838	0.002
6	0.003	0.003	0.002	0.021	0.002	0.356