Social Networks 44 (2016) 334–345

Contents lists available at ScienceDirect

Social Networks

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

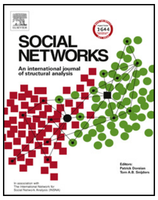

Multilevel models for social networks: Hierarchical Bayesian approaches to exponential random graph modeling

Andrew J. Slaughter[a,*], Laura M. Koehly[b,1]

[a] US Army Research Institute for the Behavioral and Social Sciences, USA
[b] National Human Genome Research Institute, USA

ARTICLE INFO

Article history:
Available online 7 December 2015

Keywords:
Multilevel networks
Hierarchical models
Exponential random graph models
Autoregressive models

ABSTRACT

In many applications, researchers may be interested in studying patterns of dyadic relationships that involve multiple groups, with a focus on modeling the systematic patterns within groups and how these structural patterns differ across groups. A number of different models – many of them potentially quite powerful – have been developed that allow for researchers to study these differences. However, as with any set of models, these are limited in ways that constrain the types of questions researchers may ask, such as those involving the variance in group-wise structural features. In this paper, we demonstrate some of the ways in which multilevel models based on a hierarchical Bayesian approach might be used to further develop and extend existing exponential random graph models to address such constraints. These include random coefficient extensions to the standard ERGM for sets of multiple unconnected or connected networks and examples of multilevel models that allow for the estimation of structural entrainment among connected groups. We demonstrate the application of these models to real-world and simulated data sets.

Published by Elsevier B.V.

1. Introduction

One important goal of network analysis is the ability to statistically model ties in social or other networks. Typically, this is done by making inferences about potential types of dependencies, characterized by particular structural features or patterns within the network. These patterns can be used to help understand the various processes – such as social or cognitive processes – that networks exhibit in particular contexts. Although researchers are typically interested in understanding general processes across networks and contexts, network research has often focused on the analysis of individual networks.

There are a number of possible reasons for this tendency. One is the relative difficulty in collecting data from multiple networks – especially networks large enough to display more than the very simplest types of structure. Another is the fact that much of network analysis has historical roots in areas like sociology and anthropology – disciplines with prominent histories of research involving in-depth qualitative and quantitative studies of specific groups and social contexts. Thus, these studies, while aimed at understanding general social phenomena, tend to focus on individual network systems, rather than multiple networks. The third – and perhaps primary reason – is the (relative) lack of statistical tools to appropriately handle data from multiple networks of varying sizes.

The focus on analyzing individual networks is not at all unreasonable. However, it is becoming progressively easier to collect network data from large numbers of groups. Similarly, network methods have continued to proliferate in disciplines where scientific methods are not grounded in making general inferences based on the individual analysis of a small number of ad-hoc networks. Rather, such disciplines aim to understand systematic structural patterns observed across many groups, and differences observed between groups. There have been a number of advances in methods for the analysis of collections of networks (Lubbers and Snijders, 2007; Wang et al., 2013; Zijlstra et al., 2006); in their current form, these methods are useful in a number of ways, and even quite powerful, but have limitations in how they can be used to model variation in relational tendencies within and across groups.

The purpose of the current paper is to present some new options for researchers interested in studying networks in a multilevel context, based on extensions to the usual exponential random graph

---

[*] Corresponding author at: 6000 6th St, Ft. Belvoir, VA 22060, USA. Tel.: +1 703 545 2353.
E-mail addresses: andrew.j.slaughter.civ@mail.mil (A.J. Slaughter), koehlyl@mail.nih.gov (L.M. Koehly).
[1] Address: Building 31, Room B1B54, 31 Center Dr, MSC 2073, Bethesda, MD 20892, USA.

http://dx.doi.org/10.1016/j.socnet.2015.11.002
0378-8733/Published by Elsevier B.V.

model (ERGM) using a hierarchical Bayesian approach. We first describe models for the special case in which multiple independent sets of networks have been observed. Then, options that account for increased complexity illustrated by true "multilevel" networks (e.g., networks between nodes at different levels of analysis) are presented. Several empirical examples are then presented, and potential applications and future extensions are discussed.

2. Extending the ERGM

There are a wide variety of methods and metrics available to researchers who are interested in exploring the structural properties of networks. However, for those interested in making inferences about such properties, stochastic approaches, which provide a basis for statistical tests, are required. Perhaps the most widely used approach is the ERGM.

The ERGM is a conceptually powerful modeling framework for the analysis of network data. Its development, rationale, and use has been covered extensively elsewhere (for example, see Lusher et al., 2013 for a general overview), but briefly, ERGMs allow researchers to specify sufficient statistics in the form of counts of so-called "graph motifs." These motifs are typically defined based on the researchers' beliefs about the potential nature and sources of dependencies between and across particular sets of actors and their dyadic relations.

The basic ERGM may be written in terms of the following equation:

$$P(G = g|\theta) = \frac{e^{\theta z(G)}}{k(\theta)} \tag{1}$$

This equation defines the probability of observing any particular graph realization as a function of a vector of sufficient statistics defined on the observed graph, $z(g)$, a vector of ERGM parameters, θ, and an intractable normalizing constant, $k(\theta)$. For notational convenience, we define the distribution of a set of network ties given by the above ERGM using the following:

$$g \sim F(\theta, z) \tag{2}$$

ERGMs are tremendously flexible models, which are easily extended to much more complicated structures than simple graphs. Extensions include models for networks with multiple types of ties, various types of hypergraphs (which include two-mode networks and cognitive social structures as special cases), networks with temporal dependencies, and multilevel "networks of networks" (cf., Desmarais and Cranmer, 2012; Koehly and Pattison, 2005; Wang et al., 2009, 2013).

Such models are extremely powerful, but with a handful of notable exceptions, they have generally been oriented towards the study of single networks. They have not been designed to model the variability in structure across collections of (possibly independent) networks. However, for researchers in many fields, the ability to model between-group variability may be of particular interest. For example, in the area of team and small-group research, researchers may be especially interested in understanding how group-level contextual factors such as training and leadership relate to variability in relational tendencies at the group level. In the area of public health, where intervention efforts might be focused on activating interpersonal relationships within a group, approaches that can compare intervention and control groups with respect to interpersonal interactions, or network structure, are necessary.

Traditionally, extensions to ERGMs have been difficult to develop and apply, due in large part to the presence of an intractable normalizing constant in the likelihood function. However, recent work in Markov Chain Monte Carlo (MCMC) estimation now offers a variety of tools to deal with the issue. These include algorithms for approximate Bayesian computation (ABC; Marin et al., 2012) and MCMC approaches based on auxiliary variables (e.g., Caimo and Friel, 2011; Jin et al., 2013; Liang, 2010).

These Bayesian techniques provide an elegant way of handling estimation for statistical models like ERGMs. Perhaps more importantly, these algorithms fit naturally into the hierarchical Bayesian paradigm and applications to multilevel models. In so doing, these algorithms afford an opportunity to more easily develop new types of multilevel models that take advantage of the utility of the ERGM framework, while offering additional flexibility for exploring variability in structure across networks and levels. Several simple examples are described in the next sections.

2.1. *Modeling independent samples of networks*

When faced with the analysis of multiple independent networks, two primary approaches have been used. One approach is the random effects p2 model (Zijlstra et al., 2006). Although this model allows for the estimation of a single comprehensive "random effects" model for networks via IGLS (iterative generalized least squares) or MCMC, the underlying p2 model is limited in the type of structures and dependencies that can be modeled, allowing for modeling the effects of density, reciprocity, as well as actor-specific effects on sending and receiving ties, and covariates which can predict the actor-specific effects.

The second, and more widely used approach, is the meta-analytic approach (see Lubbers and Snijders, 2007, for example). The meta-analytic approach is essentially a two-stage method. That is, given some set of n individual networks, the analyst first estimates n different ERGMs using standard statistical techniques. For each network i, this results in a vector of ERGM parameter estimates θ_i and an estimate of the uncertainty for those estimates $\Sigma_{se(i)}$.

Different variations on the meta-analytic approach are possible, including both fixed-effects and random-effects approaches. In the random-effects meta-analytic approach, we assume that there are some unobserved "true" random effects μ_i given by the following set of distributions:

$$\begin{aligned} &\theta_i \sim MVN(\mu_i, \Sigma_{se(i)}) \\ &\mu_i \sim MVN(\theta_f, \Sigma_r) \end{aligned} \tag{3}$$

In this model, θ_f represents the vector of fixed effects and Σ_r is the variance–covariance matrix associated with the random effects. In a Bayesian context, we place appropriate priors on θ_f and Σ_r to complete the model specification. The posterior distribution of the parameters is then given by

$$P(\mu_i, \theta_f, \Sigma_r|\theta_i, \Sigma_{se(i)}) \propto \prod_{i=1}^{n} \{P(\theta_i|\mu_i, \Sigma_{se(i)})P(\mu_i|\theta_f, \Sigma_r)\}P(\theta_f)P(\Sigma_r) \tag{4}$$

Despite the extreme simplicity of the two-step approach, there is much to recommend its use. For applied researchers, it has the benefit of being conceptually straightforward, and can be easily implemented using current software packages such as statnet (Handcock et al., 2003), WinBUGS (Lunn et al., 2000) or Stan (Stan Development Team, 2013). Moreover, estimates from such an approach are not expected to be unreasonable: the use of random effects approaches to the analysis of aggregate data have been extensively developed and widely used in many disciplines. That being said, two-step approaches do have potential drawbacks.

First, it is not immediately clear that the two-stage method based on random effects meta-analysis will, in practice, always produce identical results to simultaneous estimation across the full data set. This question (in a non-network context) has received some attention in the medical literature, where several researchers