Social Networks 29 (2007) 173–191

SOCIAL NETWORKS

www.elsevier.com/locate/socnet

An introduction to exponential random graph (*) models for social networks

Garry Robins, Pip Pattison, Yuval Kalish, Dean Lusher

D , B , . 3010, A

A

This article provides an introductory summary to the formulation and application of exponential random graph models for social networks. The possible ties among nodes of a network are regarded as random variables, and assumptions about dependencies among these random tie variables determine the general form of the exponential random graph model for the network. Examples of different dependence assumptions and their associated models are given, including Bernoulli, dyad-independent and Markov random graph models. The incorporation of actor attributes in social selection models is also reviewed. Newer, more complex dependence assumptions are briefly outlined. Estimation procedures are discussed, including new methods for Monte Carlo maximum likelihood estimation. We foreshadow the discussion taken up in other papers in this special edition: that the homogeneous Markov random graph models of Frank and Strauss [Frank, O., Strauss, D., 1986. Markov graphs. Journal of the American Statistical Association 81, 832–842] are not appropriate for many observed networks, whereas the new model specifications of Snijders et al. [Snijders, T.A.B., Pattison, P., Robins, G.L., Handock, M. New specifications for exponential random graph models. Sociological Methodology, in press] offer substantial improvement.

K : Exponential random graph models; Statistical models for social networks; * models

In recent years, there has been growing interest in exponential random graph models for social networks, commonly called the * class of models (Frank and Strauss, 1986; Pattison and Wasserman, 1999; Robins et al., 1999; Wasserman and Pattison, 1996). These probability models for networks on a given set of actors allow generalization beyond the restrictive dyadic independence assumption of the earlier 1 model class (Holland and Leinhardt, 1981). Accordingly, they permit models to be built from a more realistic construal of the structural foundations of social behavior. The usefulness of these models as vehicles for examining multilevel and multitheoretical hypotheses has been emphasized (e.g., Contractor et al., 2006).

Corresponding author.
E : garrylr@unimelb.edu.au (G. Robins).

0378-8733/$ – see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.socnet.2006.08.002

There have been a number of major theoretical and technical developments since Anderson et al. (1999) presented their well-known primer on * models. We summarize these advances in this paper. In particular, we consider it important to ground these models conceptually in their derivation from dependence assumptions, as the underlying basis of a model is then made explicit and more readily linked with hypotheses about (unobserved) social processes underlying network formation. It is through such an approach that new models can be developed in a principled way, including models that incorporate actor attributes. Recent developments in model specification and estimation need to be noted, as do new technical steps regarding setting structures and partial dependence assumptions that not only expand the class of models but have important conceptual implications. In particular, we now have a much better understanding of the properties of Markov random graphs, and promising new specifications have been proposed to overcome some of their deficiencies.

This article describes the models and summarizes current methodological developments with an extended conceptual exposition. (More technical recent summaries are given by Wasserman and Robins, 2005; Robins and Pattison, 2005; Snijders et al., in press.) We begin by briefly describing the rationale for analyzing social networks with statistical models (Section 1). We then provide an overview of the underlying logic of exponential random graph models and outline our general framework for model construction (Section 2). In Section 3, we discuss the important concept of a at the heart of the modeling approach. In Section 4, we present a range of different dependence assumptions and models. For model estimation (Section 5), we briefly summarize the pseudo-likelihood estimation (PLE) approach, and review recent developments in Monte Carlo Markov chain maximum likelihood estimation techniques. In Section 6, we present a short example of fitting a model to network data. In conclusion, we note the importance of the new model specifications that are the focus of attention in other papers in this special edition.

1. W ?

There are many well-known techniques that measure properties of a network, of the nodes, or of subsets of nodes (e.g., density, centrality and cohesive subsets). These techniques serve valuable purposes in describing and understanding network features that might bear on particular research questions. Why, then, might we want to go beyond these techniques and search for a well-fitting of an observed social network, and in particular a model? Reasons

(2) Statistical models also allow inferences about whether certain network substructures – often represented in the model by one or a small number of parameters – are more commonly observed in the network than might be expected by chance. We can then develop hypotheses about the social processes that might produce these structural properties.

whether the observed network shows a strong tendency for reciprocity, over and above the chance appearance of a number of reciprocated ties if relationships occurred completely at random. In other words, do actors in the observed network tend to reciprocate relationship choices? Here the structural characteristic (reciprocated ties) is the outcome of a social process (individuals choosing to reciprocate the choices of others). Thus, as a simple example, we might posit a stochastic network model with two parameters, one that reflects the propensity for ties to occur at random and one that reflects an additional propensity for reciprocation to occur.

In general, the structural characteristics in question help to shape the form of the model. An assumption of a reciprocity process leads us to propose a model in which an index of the level

instance possible homophily effects in the classroom. Notice that each of these processes can be represented as a small-scale graph configuration: for instance, a reciprocated tie, or a tie between two girls.

2.1.3. 3: “ “ “

It can be proven that well-specified dependence assumptions(assumptreplie)]T7ip(na p2 TD[(he)-(p2 TD[

where (i) the summation is over all configurations A; (ii) η_A is the parameter corresponding to the configuration A (and is non-zero only if all pairs of variables in A are assumed to be); [2] (iii) $_A( ) =$ $_A$ is the corresponding to configuration A; $_A( ) = 1$ if the configuration is observed in the network , and is 0 otherwise; [3] (iv) κ is a normalizing quantity which ensures that (1) is a proper probability distribution. [4]

All exponential random graph models are of the form of Eq. (1) which describes a general probability distribution of graphs on nodes. The probability of observing any particular graph in this distribution is given by the equation, and this probability is dependent both on the statistics $_A( )$ in the network on the various non-zero parameters η_A for all configurations A in the model. Configurations might include reciprocated ties, transitive triads and so on, so the model enables us to examine a variety of possible structural regularities.

So why are dependence assumptions important here? Dependence assumptions have the consequence of picking out different types of configurations as relevant to the model. Note from point (ii) above, parameters are zero whenever variables in a configuration are conditionally independent of each other. In other words, the only configurations that are relevant to the model are those in which all possible ties in the configuration are mutually contingent on each other. [5]

It is worth noting that if a set of possible edges represents a configuration in the model, then (1) implies that any subset of possible edges is also a configuration. Thus, single edges are always configurations, as demonstrated in Section 4.

So the dependence assumption is crucial in constraining which configurations are possible in the model. We will discuss particular examples in Section 4. A configuration A refers to a subset of tie variables, and corresponds to a small network substructure. For instance, if for a directed network we apply a dyadic dependence assumption (see Section 4) it will follow that reciprocity parameters will be in the model. In this case, one configuration in the model is the set of variables $\{ _{12}, _{21}\}$, another is $\{ _{13}, _{31}\}$, and so on, with every dyad providing its own configuration. Obviously for any of these configurations, if both of the ties are present in the observed graph, we see a reciprocated tie, so the configuration represents a type of network substructure that may be observed in the graph . We can think of this configuration diagrammatically as that substructure, i.e. a reciprocated tie.

But of course there is no guarantee that all possible edges in a given configuration will be present in a realized graph , so we will observe some of these possible substructures but not others. Some ties will be reciprocated, some will not. Configurations represent possibilities. The graph statistic,

the network. If we could observe the evolution of the network, and if the network started with few reciprocated ties, we might expect to see more reciprocated ties emerge over time. In thinking this way, though, we need to bear in mind that as a particular tie emerges through an imagined process of generation, its presence may affect other potential neighboring ties. So there is an implicitly dynamic and self-organizing quality227.5(qubo)-202.7(bo)-20h92.6(bo)-2hypbobo

4.2. D :

A somewhat more complicated (but not usually very realistic) assumption for directed networks is that dyads, rather than edges, are independent of one another. With this dependence assumption we have two types of configurations in the model, single edges and reciprocated edges. With homogeneity imposed, the model then becomes:

$$\Pr(\quad = \quad) = \frac{1}{\kappa} \exp\left(\theta \quad + \rho \quad\right) = \frac{1}{\kappa} \exp(\theta \quad (\quad) + \rho \quad (\quad)) \qquad (3)$$

where ( ) is the number of ties in and ( ) = , is the number of mutual ties in . A slightly more complex homogeneity assumption results in the 1 model of Holland and Leinhardt (1981).

Related but more complex and realistic models include the 2 model (Lazega and van Duijn, 1997; Van Duijn et al., 2004) which assumes dyadic independence but conditional on node-level attribute effects. The 2 model is appropriate when structure is expected to arise from attributes. It is an extension of the 1 model with sender and receiver effects treated as random effects and with actor and dyadic effects included. The more complex assumptions underpinning this model make it more realistic for actual network data, especially when attribute effects are expected to be strong. It differs from usual exponential random graph models in the incorporation of random effects.Of course, in the case of non-directed networks, Bernoulli and dyad dependence models are identical: for non-directed networks, the reciprocity parameter ρ in Eq. (3) is irrelevant and the model reduces to that of Eq. (2).

.

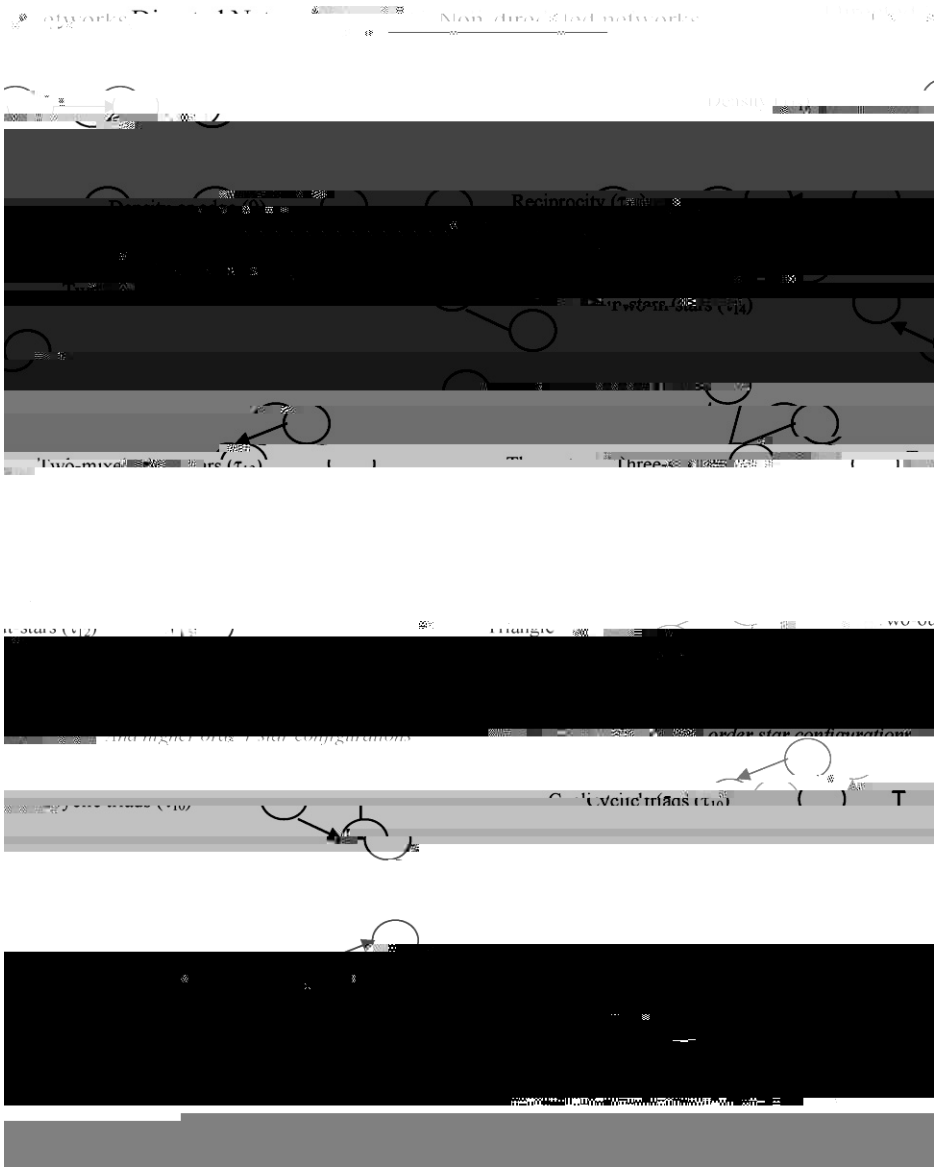

Fig. 1. Configurations and parameters for Markov random graph models.

(τ_{13}) relates to two-paths, and the two-in-star parameter (τ_{14}) relates to popularity. Note the important transitivity and cyclic configurations (τ_9 and τ_{10}). The inclusion of these parameters is a strength of these models because there is a paucity of network models that incorporate these effects (Newman, 2003

include a non-zero parameter for at least the three-star effect in models for many social networks (Robins et al., 2004, 2005). An alternative approach (see below) includes all higher-order star parameters but imposes constraints on the relationships between higher-order star parameters and lower-order ones.

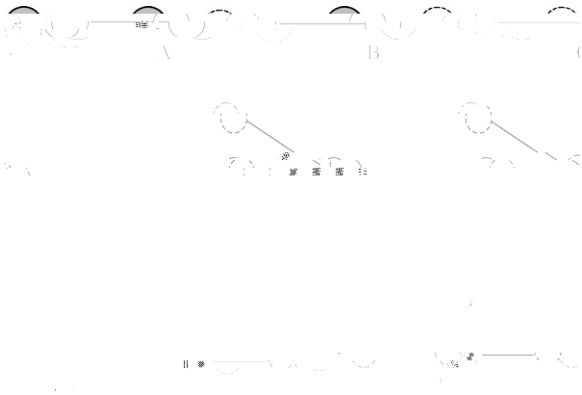

needs to be collected for a full understanding of a social network. For further elaborations, see also Schweinberger and Snijders (2003).

A second direction presented by Pattison and Robins (2002) was to propose non-Markov dependencies among ties that did not share an actor but might be interdependent through third party links. For instance, may be conditionally dependent on for four distinct actors if there is an observed tie between either or and either or . These - can be developed through what Pattison and Robins (2002) described as . These models also permit the introduction of triangles involving attribute effects.

4.6. fi

There is mounting evidence that homogeneous Markov random graph models are not good models for many observed social networks (see Section 5.2 below), so these models are not always useful in practical terms. Based on realization-dependence structures, Snijders et al. (in press) developed new specifications for exponential random graph models that include new higher order terms. These models introduce constraints on -star parameters, as well as new higher-order -triangle configurations which allow for the measurement of highly clustered regions of the network where two individuals may be connected to a large number of others (a -triangle). For these models, many higher order star and triangle effects are included (rather than set to zero) but they are constrained in the form of a weighted sum with alternating signs. The motivation behind these innovations, and the success of these new model specifications, are discussed in other papers in this special edition.

5. E

Anderson et al. (1999) in their * primer used pseudo-likelihood estimation introduced by Strauss and Ikeda (1990) in order to estimate the parameters of Markov models. We now know that, depending on the data, there may be serious problems with pseudo-likelihood estimates for these models. But for Markov random graph models, standard maximum likelihood estimation is not tractable for any but very small networks, because of the difficulties in calculating the normalizing constant in Eq. (1). What this means is that standard statistical techniques cannot be applied to these models. These problems have been overcome in recent times by the development of new Monte Carlo maximum likelihood techniques. We begin by making some rather brief comments about pseudo-likelihood and then introduce the new estimation approaches.

5.1. :

detail):

$$\log \frac{\Pr(\quad = 1| \quad)}{\Pr(\quad = 0| \quad)} = \sum_{A(\quad)} \eta_A \quad _A(\quad) \tag{5}$$

where (1) the sum is over all configurations A that contain ; (2) η_A is the parameter corresponding to configuration A; (3) $_A(\quad)$ is the ; the change in the value of the network statistic $_A(\quad)$ when changes from 1 to 0; (4) is all the observations of ties in except the observation .

The calculation of the change statistic has been discussed extensively by a number of authors (Anderson et al., 1999; Pattison and Robins, 2002; Wasserman and Pattison, 1996; Wasserman and Robins, 2005), so we do not go into it further here. With the change statistics calculated, to produce the pseudo-likelihood estimates, each possible tie becomes a case in a standard logistic regression procedure, with predicted from the set of change statistics (Anderson et al., 1999).

This procedure looks like a logistic regression – or indeed, a loglinear model – but . Logistic regression assumes independent observations, an assumption we explicitly do not make with Markov and higher order models. So the parameter estimates may be biased; and the standard errors are approximate at best, and may be too small. One should not rely on the Wald statistic as a means to decide whether a parameter is significant or not. As well, one cannot assume that the pseudo-likelihood deviance is asymptotically distributed as Chi-squared (which would be the case in normal logistic regression). When the dependence among observations is not so strong, it is generally the case that PL estimates will be more accurate. Pseudo-likelihood estimation has been used to date as a pragmatic convenience (given that alternatives have not hitherto been readily available) and the method does not have a principled basis. Whenever possible, the preferred option is to use Monte Carlo estimation procedures.

5.2. C ‘ ‘ ‘ ‘ (C C E)

Important recent developments in Monte Carlo estimation techniques for exponential random graph models have been presented and reviewed by a number of authors (see Snijders, 2002; Handcock et al., 2006; Snijders et al., in press; Wasserman and Robins, 2005), and are further discussed in other articles in this special edition, so we include only a brief summary here.

To begin, we note that simulation of these models can be implemented in a relatively straightforward way. Without going into details, simulation of the graph distribution for a given set of parameter values can be achieved through a number of algorithms (e.g., algorithms well-known in statistics more generally, such as the Metropolis algorithm). Simulation is at the heart of Monte Carlo maximum likelihood estimation. Procedures for simulating exponential random graph distributions have been described by Strauss (1986), Snijders (2002) and Robins et al. (2005).

Although there are variations between different Monte Carlo estimation techniques (Snijders, 2002; Hunter and Handcock, 2006), they are based on the same central approach: simulation of a distribution of random graphs from a starting set of parameter values, and subsequent refinement of the parameter values by comparing the distribution of graphs against the observed graph, with this process repeated until the parameter estimates stabilize. Recent software that implements Monte Carlo maximum likelihood estimation for exponential random graph models is reviewed in other papers in this special edition.

Both estimation and simulation studies have raised issues of model specification for Markov random graphs. Handcock (2003) defined as occurring when a model implied that only a few graphs had other than very low probability (often these were the full graph or the empty graph). If a model implies only these rather uninteresting outcomes, it will not be useful for modeling real networks. Simulation studies suggest that Markov graph models that contain at least non-zero three-star parameters tend to exhibit less near degeneracy than those with two-stars as the highest order non-zero star parameter (Robins et al., 2005). But the inclusion of three-star parameters often is not sufficient to remove near degeneracy behavior in Markov graph models, particularly when attempting to find models that reproduce the high levels of transitivity often observed in human social structures (there is an extended discussion in Snijders et al.,

Table 1
Parameter estimates for Markov graph model: Florentine families business network (maximum likelihood estimates with standard errors in brackets)

Parameter	Configuration	Estimate (standard error)
θ	⌒ ⌒	−4.27 (1.13)
σ_2		

Contractor, N., Wasserman, S., Faust, K., 2006. Testing multi-theoretical multilevel hypotheses about organizational networks: an analytic framework and empirical example. Academy of Management Journal 31, 681–703.

Erdös, P., Renyi, A., 1959. On random graphs. I. Publicationes Mathematicae (Debrecen) 6, 290–297.

Feld, S., 1981. The focused organization of social ties. American Journal of Sociology 86, 1015–1035.

Frank, O., Nowicki, K., 1993. Exploratory statistical analysis of networks. Annals of Discrete Mathematics 55, 349–366.

Frank, O., Strauss, D., 1986. Markov graphs. Journal of the American Statistical Association 81, 832–842.

Handcock, M.S., 2003. Statistical models for social networks: degeneracy and inference. In: Breiger, R., Carley, K., Pattison, P. (Eds.), Dynamic Social Network Modeling and Analysis. National Academies Press, Washington, DC, pp. 229–240.

Wasserman, S., Pattison, P.E., 1996. Logit models and logistic regressions for social networks. I. An introduction to Markov graphs and *. Psychometrika 61, 401–425.

Wasserman, S., Robins, G.L., 2005. An introduction to random graphs, dependence graphs, and *. In: Carrington, P., Scott, J., Wasserman, S. (Eds.), Models and Methods in Social Network Analysis. Cambridge University Press, New York, pp. 148–161.