
A simple model for complex networks with arbitrary degree distribution and clustering

keywords: Random graph models; Markov Chain Monte Carlo; Statistical Exponential Families

Mark S. Handcock

HANDCOCK@STAT.WASHINGTON.EDU

Department of Statistics, C14-B Padelford Hall, Seattle, WA 98195-4322 USA

Martina Morris

MORRISM@U.WASHINGTON.EDU

Center for Studies in Demography and Ecology, 206A Raitt Hall, Seattle, WA 98195-3412 USA

Abstract

We present a stochastic model for networks with arbitrary degree distributions and average clustering coefficient. Many descriptions of networks are based solely on their computed degree distribution and clustering coefficient. We propose a statistical model based on these characterizations. This model generalizes models based solely on the degree distribution. We present alternative parameterizations of the model. Each parameterization of the model is interpretable and tunable. We present a simple Markov Chain Monte Carlo (MCMC) algorithm to generate networks with the specified characteristics. We provide a algorithm based on MCMC to infer the network properties from network data and develop statistical inference for the model. The model is generalizable to include mixing based on attributes and other complex social structure.

1. Introduction

In this paper we consider models where individuals are represented as nodes in a network and edges represent some form of social contact. We assume that the network is a realization of a stochastic process characterized by random mixing between individuals conditional on the individual activity levels (i.e., the nodal degrees) and clustering (Newman, 2002b; Dezső & Barabási, 2002). One popular class are those that exhibit power-law behavior, often loosely referred to as “scale-free” distributions. We also consider models for the network degree distributions in which the

variance can greatly exceed the mean.

In Section 2 we develop the general form of the model and models for the degree distribution. In Section 3 we give an simple algorithm for the generation of random networks from the model. In Section 4 we provide an algorithm for approximating the likelihood function for the model as a basis for inference. In Section 5 we apply the model to a protein-protein interaction network. Finally, in Section 6, we discuss generalizations of the model for more complex structures.

2. Models for Social Networks

2.1. Exponential Family Models

Let the random matrix X represent the adjacency matrix of an unvalued network on n individuals. We assume that the diagonal elements of X are 0 – that self-partnerships are disallowed. Suppose that $\mathbb{X}$ denotes the set of all possible networks on the given n individuals. The multivariate distribution of X can be parameterized in the form:

$$P_{\eta, \mathbb{X}}(X = x) = \frac{\exp[\eta \cdot t(x)]}{c(\eta, \mathbb{X})} \quad x \in \mathbb{X} \quad (1)$$

where $\eta \in \Upsilon \subseteq \mathbb{R}^q$ is the model parameter and $t: \mathbb{X} \rightarrow \mathbb{R}^q$ are statistics based on the adjacency matrix (Frank & Strauss, 1986; Handcock, 2002). There is an extensive literature on descriptive statistics for networks (Wasserman & Faust, 1994; Borgatti et al., 1999). This statistics are often crafted to capture features of the network (e.g., centrality, mutuality and betweenness) of primary substantive interest to the researcher. Often a researcher has have specific set of statistics in mind. The above model then has the property of maximizing the entropy within the family of all distributions with given expectation of $t(X)$ (Barndorff-Nielsen, 1978). Paired with the flexibility of the choice of t this property does provide some justification for the model (1)

that will vary from application to application.

The denominator $c(\eta, \mathbb{X})$ is the normalizing function that ensures the distribution sums to one: $c(\eta, \mathbb{X}) = \sum_{y \in \mathbb{X}} \exp[\eta \cdot t(y)]$. This factor varies with both η and the support $\mathbb{X}$ and is the primary barrier to simulation and inference under this modeling scheme.

The most commonly used class of random network models exhibit Markov dependence in the sense of Frank and Strauss (1986). For these models, dyads that do not share an individual are conditionally independent; this is an idea analogous to the nearest neighbor concept in spatial statistics. Typically a homogeneity condition is also added: all isomorphic networks have the same probability under the model. Frank and Strauss (1986) show that homogeneous Markov networks are exactly those having the *degree parameterization*:

$$d_k(x) = \frac{\text{the proportion of nodes with degree exactly } k}{k = 0, \dots, n-1}$$

$$N_\Delta(x) = \frac{1}{6} \sum_{i,j,k} x_{ij} x_{jk} x_{kl},$$

where $d_k(x)$ counts the proportion of individuals with degree k and $N_\Delta(x)$ is a count of the complete triads. This model can be reexpressed in the notation of model (1) by setting $t_k(x) = d_k(x)$, $k = 1, \dots, n-1$, $t_n = N_\Delta(x)$, $q = n$, $\eta \in \Upsilon = \mathbb{R}^n$. This parameterization has the advantage that it is directly interpretable in terms of concurrency of partnerships (i.e. $d_m(x)$ for $m > 0$ is the proportion of individuals with exactly m concurrent partners).

A popular variant of the statistic $N_\Delta(x)$ is the clustering coefficient defined as

$$C(x) = \frac{3N_\Delta(x)}{N_3(x)}$$

where $N_3(x)$ is the number of connected triples of nodes (i.e., 2-stars Frank & Strauss, 1986). This describes the proportion of complete triads in the networks out of the total number of possible triads.

In the remainder of this paper we focus on the following novel model

$$\log[P_\theta(X = x)] = \eta(\phi) \cdot d(x) + \nu C(x) - \log c(\phi, \nu, \mathbb{X}), \quad (2)$$

where $x \in \mathbb{X}$, $\theta = (\phi, \nu)$, $\Theta \subset \mathbb{R}^n$, $d(x) = \{d_1(x), \dots, d_{n-1}(x)\}$. The parameters ϕ and ν represent the network degree distribution and clustering, respectively. Specifically, the ratio of the probability of a given network to a network with the same degree distribution and correlation coefficient 1% less is $0.01 \times \exp(\nu)$. Alternatively, consider the conditional probability of a tie existing given the rest of the network. If the formation of

the tie increases the correlation coefficient by $\alpha\%$ (relative to the same network without the tie) then the log-odds of the tie existing is $\alpha\nu\%$. The degree distribution parameters have similar interpretations: $\eta_k(\phi)$ is the ratio of the log-probability of a given network to a network with the same clustering coefficient and one less node of degree k and one more isolate. An important property of the model is the variational independence of the parameters (Barndorff-Nielsen, 1978).

This model is a curved exponential family if Θ is a smooth curve in $\Upsilon = \mathbb{R}^n$ (Hunter & Handcock, 2005; Handcock, 2003a). Any degree distribution can be specified by $n-1$ or less independent parameters. Typically the number of parameters is small. As we shall see, this is true for the models considered below.

If $\nu = 0$ the model corresponds to random networks with arbitrary degree distributions, as considered by many researchers (Newman et al., 2001). If $\eta_k(\phi) = \phi k$, $k = 1, \dots, n-1$ the value of ϕ is interpretable as the log-probability of a given network to a network with one less tie and the same clustering coefficient (Hunter & Handcock, 2005). If both $\nu = 0$ and $\eta_k(\phi) = \phi k$, $k = 1, \dots, n-1$ it is the classical random network model of Renyi and Erdos (Bollobas, 1985).

The model (1) has a generative interpretation, which we illustrate with model (2). Consider a dynamic process for the network $\{X(t): t \geq 0\}$ developing according to the local rules

$$\begin{aligned} & \logit[P(X_{ij}(t) = 1 | X_{ij}(t^-) = x_{ij})] \\ &= \eta(\phi) \cdot [d(x_{ij}^+) - d(x_{ij}^-)] + \nu [C(x_{ij}^+) - C(x_{ij}^-)] \end{aligned}$$

where x_{ij}^+ is the network with a tie between i and j and the rest of the network equal to x_{ij} . x_{ij}^- is similar with no tie between i and j . Based on the theory of continuous-time Markov Chains, the equilibrium distribution is model (2). Ties are formed (or broken) based on their propensity to change the network characteristics. This also provides another interpretation of the parameters ϕ and ν and their joint effects.

An alternative parameterization that is usually more interpretable is: (ϕ, ρ) where the mapping is:

$$\rho = \mathbb{E}_{\phi, \rho} [C(X)] = \sum_{y \in \mathbb{X}} C(y) \exp[\eta(\phi) \cdot d(y) + \nu C(y)] \geq 0 \quad (3)$$

Thus ρ is the mean clustering coefficient over networks in $\mathbb{X}$. Thus models with higher ρ have higher clustering coefficients on average. Note that models with $\rho = 0$ will not have any complete triads. The range of ρ is a subset of $[0, 1]$ and depends on the other parameters and $\mathbb{X}$.

The two parameterizations represent the same model class

(Handcock, 2003a). Translating between equivalent parameters is achieved using the MCMC algorithm given in Section III (Handcock, 2003a; Hunter & Handcock, 2005).

2.2. Models for Degree Distributions

Let $P_\theta(K = k)$ be the probability mass function of K , the number of ties that a randomly chosen node in the network has. Based on the model (2)

$$P_\theta(K = k) = \mathbb{E}_\theta [d_k(X)] \quad k = 0, \dots, n-1$$

Clearly for a given network of size n nodes, the distribution of K has finite range with upper bound $n-1$. In some cases this distribution is approximated by an idealized distribution with infinite range. Let K^* be the degree of a node in a (possibly hypothetical) infinite population of nodes. Then K can be thought of as the degree of the node restricted to nodes in the network. In cases where this conceptualization is used we will consider the case

$$P_\theta(K = k) = P(K^* = k | K^* < n) \quad k = 0, \dots, n-1,$$

While the model (2) has arbitrary degree distribution, of particular interest are the various “scale-free,” preferential attachment and power-law models popular in the physics literature (see, e.g., Newman, 2003). These models assume that all networks with the same degree distribution are equally likely. We say $P(K^* = k)$ has *power-law behavior* with scaling exponent $\phi > 1$ if there exist constants c_1, c_2 , and M such that $0 < c_1 \leq P(K^* = k)k^\phi \leq c_2 < \infty$ for $k > M$.

We focus on a stochastic mechanisms for the formation of the social networks that is a variation on a preferential attachment process, such as those advocated by several recent authors (Barabási & Albert, 1999; Pastor-Satorras & Vespignani, 2001). The limiting distributions of this mechanism can be characterized by long tails.

2.3. Simple Preferential Attachment Models

A mechanism that has been suggested for the formation of power-law social networks is preferential attachment (Albert & Barabási, 2000; Liljeros et al., 2001; Dezső & Barabási, 2002). This and related stochastic processes have a long history in applied statistics (Simon, 1955; Kendall, 1961; Irwin, 1963). Consider a population of r people in in which (1) there is a constant probability p that the $r + 1$ st partnership in the population will be initiated from a randomly chosen person to a previously sexually inactive person, and (2) otherwise the probability that the $r + 1$ st partnership will be to a person with exactly k partners is proportional to $k f(k|r)$, where $f(k|r)$ is the frequency of nodes with exactly k connections out of the r total links in the population. The limiting distribution of this process is known as the Waring distribution (Irwin, 1963).

The Yule distribution discussed by Simon (1955) and used by Jones and Handcock (2003) to model degree distributions is a special case of the Waring distribution with $p = (\phi_2 - 2)/(\phi_2 - 1)$.

The probability mass function (PMF) of the Waring distribution (Johnson et al., 1992) is:

$$P(K^* = k) = \frac{(\phi_2 - 1)\Gamma(\phi_2 + \phi_1)}{\Gamma(\phi_1 + 1)} \cdot \frac{\Gamma(k + \phi_1)}{\Gamma(k + \phi_1 + \phi_2)}, \quad (4)$$

$$\phi_1 > -1, \phi_2 > 2$$

where $\Gamma(\cdot)$ is the Gamma function and the mixing parameter ϕ_1 is related to p via:

$$p = \frac{\phi_2 - 2}{\phi_2 + \phi_1 - 1}. \quad (5)$$

The Waring distribution has power-law behavior with scaling exponent ϕ_2 . The mean and variance of the Waring distribution are:

$$\mathbb{E}(K^*) = \frac{1}{p}, \quad \mathbb{V}(K^*) = \frac{(1-p)(\phi_2 - 1)}{p^2(\phi_2 - 3)}, \quad \phi_2 > 3$$

Thus, the expected value of the Waring distribution is simply the inverse of the probability of forming a tie to an individual lacking existing ties. Both the Waring and the Yule distributions have been re-discovered, apparently without awareness of their historical antecedents, by Levene et al. (2002) and Dorogovtsev et al. (2000) respectively in the context of modeling growth of the Internet.

3. Generating Random Networks with Specified Structure

Markov Chain Monte Carlo (MCMC) algorithms for generating from the model (1) have a long history and been well studied (see Geyer & Thompson, 1992 for a review). The basic idea is to generate a Markov chain whose stationary distribution is given by equation (1). The simplest Markov chain proceeds by choosing (by some method, either stochastic or deterministic) a dyad (i, j) and then deciding whether to set $X_{ij} = 1$ or $X_{ij} = 0$ at the next step of the chain. One way to do this is using Gibbs sampling, whereby the new value of X_{ij} is sampled from the conditional distribution of X_{ij} conditional on the rest of the network. Denote “the rest of the network” by X_{ij}^c . Then $X_{ij}|X_{ij}^c = x_{ij}^c$ has a Bernoulli distribution, with odds given by

$$\frac{P(X_{ij} = 1 | X_{ij}^c = x_{ij}^c)}{P(X_{ij} = 0 | X_{ij}^c = x_{ij}^c)} = \exp\{\eta \cdot \Delta(t(x))_{ij}\},$$

where $\Delta(t(x))_{ij}$ denotes the difference between $t(x)$ when x_{ij} is set to 1 and $t(x)$ when x_{ij} is set to 0. A simple variant to the Gibbs sampler (which is an instance of a Metropolis-Hastings algorithm) is a pure Metropolis algorithm in which the proposal is always to change the value of x_{ij} . This proposal is accepted with probability $\min\{1, \pi\}$, where

$$\pi = \frac{P(X_{ij} = 1 - x_{ij} | X_{ij}^c = x_{ij}^c)}{P(X_{ij} = x_{ij} | X_{ij}^c = x_{ij}^c)} \quad (6)$$

$$= \begin{cases} \exp\{\eta \cdot \Delta(t(x))_{ij}\} & \text{if } x_{ij} = 0; \\ \exp\{-\eta \cdot \Delta(t(x))_{ij}\} & \text{if } x_{ij} = 1. \end{cases}$$

The vector $\Delta(t(x))_{ij}$ used by these MCMC schemes is often much easier to calculate directly than as the difference of two separate values of $t(x)$. For instance, if one of the components of the $t(x)$ vector is the total number of edges in the network, then the corresponding component of $\Delta(t(x))_{ij}$ is always equal to 1.

The Metropolis scheme is usually preferred over the Gibbs scheme because it results in a greater probability of changing the value of x_{ij} , a property thought to produce better-mixing chains. However, it is well known that these simple MCMC schemes often fail for various reasons to produce well-mixed chains (Snijders, 2002; Handcock, 2000; Snijders et al., 2005). More sophisticated MCMC schemes have been developed and are a topic of ongoing research (Hunter & Handcock, 2005).

A variant of this algorithm proceeds in two steps:

1. Generate $d_k \stackrel{\text{i.i.d.}}{\sim} P_\theta(K = k)$, $k = 0, 1, \dots, n - 1$.
2. Generate a random network conditional on this degree distribution:

$$P_v(X = x | d_k(X) = d_k) = \frac{\exp[vC(x)]}{c(v, d_k, \mathbb{X})} \quad x \in \mathbb{X}(d_k)$$

where $\mathbb{X}(d_k) = \{x \in \mathbb{X} : d_k(x) = d_k\}$.

The first generates individual degrees from an arbitrary distribution, and the second generates networks condition on those degrees. Note that the structure of the exponential family in (1) ensure that the samples are from the correct distribution (Barndorff-Nielsen, 1978). The first step can be simulated easily as we know $P_\theta(K = k)$. Note that not all degree sequences will be consistent with a network of size n . For example, sequences with an odd total number of ties are not realizable. However we can construct a compatible sequence $\{d_k\}_{k=0}^{n-1}$ via a simple rejection algorithm. The second step is also straightforward: we can conditionally simulate values using a MCMC holding the degree distribution fixed by using a Metropolis proposal consistent with

this restriction. It is convenient for this algorithm to have a starting network with the given degree distribution. This network is easy to construct by a finite algorithm (as it need not be a draw from a random distribution). An important property of this the second step is the independence of the distribution from ϕ . It is a simple parameter distribution depending only on v (Barndorff-Nielsen, 1978).

As an application of this algorithm, consider a network model for $n = 50$ nodes. We choose a degree distribution which is Yule with scaling exponent $\phi_2 = 3$. This corresponds to a “scale-free” degree model. If $v = 0$ the network is random with the given degree distribution. This corresponds to a mean clustering coefficient $\rho = 3\%$. A realization of this model is given in the left-hand panel of Fig. 1. The clustering coefficient for this network is 2%. The right-hand side of the panel is a realization from the model with mean clustering coefficient $\rho = 15\%$ (corresponding to a clustering parameter of $v = 0.46$.) The centralization of the clustering is apparent relative to the first network.

As an second application we generate a network model for $n = 1000$ nodes with the same degree distribution ($\phi_2 = 3$). A realization of this model is given in the left-hand panel of Fig. 2. The clustering coefficient for this network is 2%. The right-hand side of the panel is a realization from the model with mean clustering coefficient chosen to be $\rho = 15\%$ (corresponding to a clustering parameter of $v = 27$.) The elongated nature of the resulting network is apparent as is the centralization of the clustering.

4. Statistical Inference for Network Models

As we have specified the full joint distribution of the network through (1), we choose to conduct inference within the likelihood framework (Besag, 1975; Geyer & Thompson, 1992). For economy of notation, in this section, we use ϕ to represent either η in (1) or the curved exponential family form (ϕ, v) in (2). Differentiating the loglikelihood function:

$$\ell(\phi; x) \equiv \log[P_\eta(X = x)] = \eta(\phi) \cdot t(x) - \log[c(\phi, \mathbb{X})] \quad (7)$$

shows that the maximum likelihood estimate $\hat{\phi}$ satisfies

$$\nabla \ell(\hat{\phi}) = \nabla \eta(\hat{\phi}) \cdot [t(x_{\text{obs}}) - E_{\eta(\hat{\phi})} t(X)], \quad (8)$$

where $\nabla \eta(\phi)$ is the $p \times q$ matrix of partial derivatives of η with respect to ϕ and $t(x_{\text{obs}})$ is the observed network statistics. We may search for a solution to equation (8) using an iterative technique such as Newton-Raphson; however, the exponential family form of the model makes the Fisher information matrix

$$I(\phi) = \nabla \eta(\phi) \cdot [\text{Cov}_{\eta(\phi)} t(X)] \nabla \eta(\phi) \quad (9)$$




Figure 1. Two examples of networks generated from model (2) with $n = 50$. Both networks have degree distribution draw from the Yule distribution (equation 4) with scaling exponent $\phi_2 = 3$. Left: Random network from the model with mean clustering coefficient $\rho = 3\%$. The example network has clustering coefficient $C(x) = 2\%$. Right: Random network from the model with mean clustering coefficient $\rho = 15\%$. The example network has clustering coefficient $C(x) = 18\%$.




Figure 2. Random networks from the model with $n = 1000$ and the same degree distribution. The largest component is visualized. Left: Random network from the model with mean clustering coefficient $\rho = 3\%$. The example network has clustering coefficient $C(x) = 1\%$. Right: Random network from the model with mean clustering coefficient $\rho = 15\%$. The example network has clustering coefficient $C(x) = 14\%$.

easier to calculate than the Hessian matrix of second derivatives required for Newton-Raphson. For more about equations (8) and (9), see Hunter and Handcock (2006). The method of Fisher scoring is an iterative method analogous to Newton-Raphson except that the negative Fisher information is used in place of the Hessian matrix.

Direct calculation of the log-likelihood by enumerating $\mathbb{X}$ is infeasible for all but the smallest networks. As an alternative, we approximate the likelihood equations (8) by replacing the expectations by (weighted) averages over a sample of networks generated from a known distribution. This procedure is described in Geyer and Thompson (1992). To generate the sample we use the MCMC algorithm of Section 3.

5. Application to a Protein-Protein Interaction Network

As an application of these methods, we fit the model to a biological network of protein-protein interactions found in cells. By interact is meant that two amino acid chains were experimentally identified to bind to each other. The network is for *E. Coli* and is drawn from the “Database of Interacting Proteins (DIP)” (Xenarios et al., 2002). The DIP database lists protein pairs that are known to interact with each other. The dataset we use is `Ecoli20050403`. We have chosen *E. Coli* as it is well studied and this will minimize the number of false-negative interactions (that is, two proteins that interact but are not in the database). For simplicity we focus on proteins that interact with themselves and have at least one other interaction. We do not represent the self-interactions as part of the network. This results in a network in Figure 3 with 108 proteins and 94 interactions.

We consider the model (2) with a clustering coefficient term and the degree distribution model by a preferential attachment process (the Yule distribution with scaling exponent ϕ). We choose the Yule as it represents the simple version of preferential attachment that is common in the literature. The estimates are given in Table 1. They are derived using the algorithm in Section 4.

Parameter	est.	s.e.
Scaling decay rate (ϕ)	3.034	0.3108
Correlation Coefficient (ν)	1.176	0.1457

Table 1. MCMC maximum likelihood parameter estimates for the protein-protein interaction network.

The estimate of the preferential attachment scaling decay rate of about three suggests that the network is close to the so-called “scale-free” range (that is, $\phi \leq 3$). The standard error of the scaling rate of suggests some uncertainty.

However the parameter of the correlation coefficient is very positive. This indicates strong clustering (given the degree sequence) and hence so-called “small world” behavior in the network. Thus, this model provides a statistical valid means to test for small-world characteristics of a network using the statistics commonly used to characterize small-world networks.

Finally, we can test if the network is generated by this preferential attachment model. If preferential attachment among proteins generated this network then the parameter ν of the clustering coefficient will be zero. However we see that the estimate is positive. We can test this more rigorously by comparing the log-likelihood values for the maximum likelihood fit in Table 1 to the model where ν is constrained to be zero. The change in the log-likelihood is 52.3, so that the change in deviance is 104.6. This indicates that deviation from the preferential attachment model is statistically significant.

6. Discussion

We have presented a simple stochastic model for random networks that has arbitrary degree distribution and average clustering coefficient. The clustering component of the model is directly interpretable via the clustering coefficient of the realizations from the model. The model places positive probability over the set of possible networks. Conditional on the degree sequence, the clustering coefficient covers the full range of values possible. The distribution over this range is tuned as a monotone function of the clustering parameter.

The model form (1) is very general, and can incorporate general social structure (Frank & Strauss, 1986; Strauss & Ikeda, 1990; Handcock, 2003a; Hunter & Handcock, 2005). For example, in disease epidemiology, the two-sex random network epidemic model is a commonly used to represent the contact structure of pathogens transmitted by intimate contact (Newman, 2002a). This model is the model (2) with $\rho = 0$ and $\mathbb{X}$ is restricted to heterosexual networks. However, this model contains a major weakness which ultimately limits its utility. Specifically, it assumes random mixing conditional on degree. The model (2) is a simple extension of that allows tunable correlation coefficient. More generally, (1) can be used to include nodal attributes and other structural characteristics. Such models have proven to be valuable in epidemiology (Morris, 2003; Handcock, 2003b).

References

- Albert, R., & Barabási, A. L. (2000). Topology of evolving networks: Local events and universality. *Physical Review Letters*, 85, 5234–5237.




Figure 3. A protein - protein interaction network for *E. Coli*. The nodes represent proteins and the ties indicate that the two proteins are known to interact with each other.

- Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286, 509–512.
- Barndorff-Nielsen, O. E. (1978). *Information and exponential families in statistical theory*. New York: Wiley.
- Besag, J. (1975). Statistical analysis of non-lattice data. *The Statistician*, 24, 179–95.
- Bollobas, B. (1985). *Random graphs*. Academic Press, NY, USA.
- Borgatti, S. P., Everett, M. G., & Freeman, L. C. (1999). *Ucinet 5.0 for windows*. Natick: Analytic Technologies.
- Dezső, Z., & Barabási, A. L. (2002). Halting viruses in scale-free networks. *Physical Review E*, 65, art. no. 055103.
- Dorogovtsev, S. N., Mendes, J. F. F., & Samukhin, A. N. (2000). Structure of growing networks with preferential linking. *Physical Review Letters*, 85, 4633–4636.
- Frank, O., & Strauss, D. (1986). Markov graphs. *Journal of the American Statistical Association*, 81, 832–842.
- Geyer, C. J., & Thompson, E. A. (1992). Constrained monte carlo maximum likelihood calculations (with discussion). *Journal of the Royal Statistical Society, Series B*, 54, 657–699.
- Handcock, M. S. (2000). *Progress in statistical modeling of drug user and sexual networks* (Manuscript). Center for Statistics and the Social Sciences, University of Washington.
- Handcock, M. S. (2002). Degeneracy and inference for social network models. *Paper presented at the Sunbelt XXII International Social Network Conference in New Orleans, LA*.
- Handcock, M. S. (2003a). *Assessing degeneracy in statistical models of social networks* (Working Paper). Center for Statistics and the Social Sciences, University of Washington.
- Handcock, M. S. (2003b). Statistical models for social networks: Inference and degeneracy. In R. Breiger, K. Carley and P. Pattison (Eds.), *Dynamic social network modeling and analysis*, 229–240. Committee on Human Factors, Board on Behavioral, Cognitive, and Sensory Sciences. National Academy Press: Washington, DC.
- Hunter, D. R., & Handcock, M. S. (2005). Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, to appear.
- Irwin, J. (1963). The place of mathematics in medical and biological statistics. *Journal of the Royal Statistical Society Series A (General)*, 126, 1–45.
- Johnson, N., Kotz, S., & Kemp, A. (1992). *Univariate discrete distributions*. Wiley series in probability and mathematical statistics. New York: Wiley. 2nd edition.

- Jones, J., & Handcock, M. S. (2003). An assessment of preferential attachment as a mechanism for human sexual network formation. *Proceedings of the Royal Society of London, B*, 270, 1123–1128.
- Kendall, M. (1961). Natural law in the social sciences: Presidential address. *Journal of the Royal Statistical Society Series A-Statistics in Society*, 124, 1–16.
- Levene, M., Fenner, T., Loizou, G., & Wheeldon, R. (2002). A stochastic model for the evolution of the web. *Computer Networks*, 39, 277–287.
- Liljeros, F., Edling, C. R., Amaral, L. A. N., Stanley, H. E., & Åberg, Y. (2001). The web of human sexual contacts. *Nature*, 411, 907–908.
- Morris, M. (2003). Local rules and global properties: Modeling the emergence of network structure. In R. Breiger, K. Carley and P. Pattison (Eds.), *Dynamic social network modeling and analysis*, 174–186. Committee on Human Factors, Board on Behavioral, Cognitive, and Sensory Sciences. National Academy Press: Washington, DC.
- Newman, M. E. J. (2002a). Assortative mixing in networks. *Physical Review Letters*, 89, art no. – 208701.
- Newman, M. E. J. (2002b). Spread of epidemic disease on networks. *Physical Review E*, 66, art. no.–016128.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45, 167–256.
- Newman, M. E. J., Strogatz, S. H., & Watts, D. J. (2001). Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64, 026118.
- Pastor-Satorras, R., & Vespignani, A. (2001). Epidemic dynamics and endemic states in complex networks. *Physical Review E*, 63, art. no.–066117.
- Simon, H. (1955). On a class of skew distribution functions. *Biometrika*, 42, 435–440.
- Snijders, T. A. B. (2002). Markov chain monte carlo estimation of exponential random graph models. *Journal of Social Structure*, 3.
- Snijders, T. A. B., Pattison, P., Robins, G. L., & Handcock, M. S. (2005). New specifications for exponential random graph models. *Sociological Methodology*, 35, to appear.
- Strauss, D., & Ikeda, M. (1990). Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association*, 85, 204–212.
- Wasserman, S. S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge: Cambridge University Press.
- Xenarios, I., Salwinski, L., Duan, X. J., Higney P., Kim, S., & Eisenberg, D. (2002). Dip, the database of interacting proteins: a research tool for studying cellular networks of protein interactions. *Nucleic Acids Res*, 28, 289–291.