

Received XXXX

(www.interscience.wiley.com) DOI: 10.1002/sim.0000

Bayesian exponential random graph modelling of interhospital patient referral networks

Alberto Caimo^{a*}, Francesca Pallotti^b and Alessandro Lomi^c

Using original data that we have collected on referral relations between 110 hospitals serving a large regional community, we show how recently derived Bayesian exponential random graph models may be adopted to illuminate core empirical issues in research on relational coordination among health care organisations. We show how a rigorous Bayesian computation approach supports a fully probabilistic analytical framework that alleviates well-known problems in the estimation of model parameters of exponential random graph models. We also show how the main structural features of interhospital patient referral networks that prior studies have described, can be reproduced with accuracy by specifying the system of local dependencies that produce – but at the same time are induced by – decentralised collaborative arrangements between hospitals.

Copyright © 2010 John Wiley & Sons, Ltd.

Keywords: Bayesian inference; exponential random graph models; interorganisational networks; interhospital patient referral networks; Monte Carlo methods; statistical models for social networks.

1. Introduction

Recent advances in inferential methods for observations characterised by complex dependencies are making network data increasingly amenable to statistical modelling and analysis [1, 2]. Exponential random graph models (ERGMs) are perhaps the best illustration of this trend [3]. ERGMs emerged as one of the main families of models capable of capturing the complex dependence structure that is typical of network data. According to Snijders et al. [4], for example, ERGMs are (2006: 99): “*The most promising class of statistical models for expressing structural properties of social networks observed at one moment in time.*”

In recent years, ERGMs have found wide application in empirical research in the field of health and medical care [5, 6]. Recent examples include the study of referral networks among cancer services organisations [7], genetic variation in human social networks [8], the effects of collaboration among institutions on the promotion of physical activity [9], the impact of physicians’ collaboration networks on hospitalisation cost and hospital readmission rate [10], and the relation between competition for patients and relational coordination among hospitals [11]. While addressing very different empirical settings, and interested in very different problems, what these studies have in common is a clear methodological commitment to *modelling network mechanisms directly*, rather than just attempting to “control for” unspecified dependence among the observations.

Exponential random graph (occasionally known as p -star, or p^*) models are not a recent discovery. Building on earlier work on the p_1 model by Holland and Leinhardt in 1981 [12], the major breakthrough came with the work of Frank and Strauss Markov random graph models [13]. A decade later, Wasserman and Pattison [14] introduced the p -star model as an extension of Markov random graph models (see also [15] and [16]). Another decade was necessary for Snijders et al. [4] to introduce new specifications of ERGMs that alleviated some known problems with Markov random graphs (see also

^aSchool of Mathematical Sciences, Dublin Institute of Technology, Ireland

^bCentre for Business Network Analysis, University of Greenwich, United Kingdom

^cSocial Network Analysis Research (SoNAR) Center, University of Lugano, Switzerland

* Correspondence to: Dublin Institute of Technology, Kevin Street, Dublin 2, Ireland. E-mail: alberto.caimo@dit.ie

[17]). These new specifications introduced higher order effects for closure-based and degree-based mechanisms leading to a significant improvement in model applicability and performance [18, 19, 20].

This innovation paved the way to the flexible general specifications that are now becoming increasingly common in the analysis of social networks [21].

Despite remarkable methodological developments and significant empirical success, the diffusion of ERGMs as a general analytical framework for network data is limited by ERGMs inherent computational complexity which makes model parameters difficult to estimate and interpret. In general, recent studies tend to agree that the intractability and degeneracy are the main limitations of ERGM as general models for network data [22, 23]. The first problem, intractability, derives from the impossibility of computing the normalising constant in non-trivial cases. Degeneracy refers to the existence of probability models that, even under the maximum likelihood estimator, place most of their mass on a small number of extreme graphs such as, for example, empty or complete graphs.

As part of a more general diffusion of Bayesian approaches in the social sciences [24, 25, 26], major advances in Bayesian computational methods allow to alleviate these problems. The breakthrough computational technology that has facilitated the application of Bayesian models is Markov Chain Monte Carlo (MCMC) simulation [27]. Recent theoretical developments in MCMC simulation algorithms facilitates computation of intractable likelihood models [28, 29] and show considerable promise of solving the main computational problems that still limit the diffusion of ERGMs [30, 31, 32, 33, 34]. More specifically, the methodology developed by Caimo and Friel [32] makes use of the approximate exchange algorithm to estimate the double-intractable posteriors. This approach has been shown to perform well and to have better convergence properties when compared with classical methods.

Against the backdrop of this general introduction, the main objective of this paper is to introduce BERGMs, discuss their properties and the new analytical possibilities they offer, and confront issues of model specification and parameter estimation in the context of a realistically complex empirical setting. In particular, we show how information about prior expectations about parameters of ERGM models can be specified, estimated and interpreted – and how overall models can be evaluated. Hence, the primary contribution of the paper is to the literature on statistical models for network data.

The empirical part of the study is based on data that we have collected on patient referral relations among 110 hospitals serving a large regional community with more than 5,000,000 potential patients. The data refer to the year 2007 and are analysed for the first time in this paper. Interhospital referral relations have been frequently framed as a particularly revealing example of how collaboration among health care organisations may affect the quality of care that patients receive [35, 36]. The network structure generated by dyadic interhospital relations has properties that affect the effectiveness and efficiency of the delivery of health care services [35, 37]. For this reason, it is surprising that only a handful of recent studies are available that have recognised the role of local relational mechanisms in explaining global structural features of these networks [11]. With the exceptions mentioned at the outset, available studies of interhospital networks are almost exclusively descriptive [38], or concerned with the transfer of critically ill patients where decentralised dyadic interorganisational relations are less important because of the overwhelming presence of formalised patient-transfer routines [39]. Studies based on ERGMs are beginning to appear that have looked at patient-sharing relations among physicians [40, 41], but we are not aware of studies of interhospital patient referral networks that have exploited the empirical advantage offered by Bayesian approaches. Hence, the secondary contribution of the paper is to the health care literature on interhospital collaboration with special reference to interhospital patient referral relations.

2. Exponential random graph models (ERGMs)

2.1. Definition and notation

The main research objective in social network analysis is to examine the structure of relationships among a set of actors represented as network nodes. Example of actors include individuals, organisations or even countries [42]. Relational structures can be described by network graphs. Formally, a network graph consists of a set of n nodes and a set of m edges which define a relationship between pairs of nodes called dyads. The connectivity pattern of a graph can be described by an $n \times n$ adjacency matrix y encoding the presence or absence of an edge between node i and j : $y_{ij} = 1$ if the dyad (i, j) is connected, $y_{ij} = 0$ otherwise. Covariates consisting of nodal attributes x_i and dyadic attributes w_{ij} between nodes can be used to model the effects of nodal and dyadic data on the formation of edges between nodes.

Most standard statistical modelling approaches, such as logistic regression, assume independence between the observations – in this case represented by dyadic values. However, for social networks this assumption is typically untenable. Ignoring dependence among edge variables has major consequences for the inference that we can draw from network data. For example, edges between nodes tend to “self-organise” to produce triadic structures. ERGMs have been specifically designed to capture these complex relational tendencies, such as transitivity expressed by triad closure, that are not modelled by most probability models. ERGMs can represent directly specific forms of dependence between

edges. Conventional regression-like models recognise the presence of forms of dependence in network data, but treat them only as sources of violation of assumptions of independence that need to be controlled for, rather than directly specified. ERGMs may be specified to address specific research questions and to obtain a good fit of the model to the data. Specifically, ERGMs assume that the relational structure in an observed network – encoded by its adjacency matrix y – can be explained by the relative prevalence of a set of overlapping sub-graph configurations $s(\cdot)$ also called *network statistics* [43] such as the number of edges, mutual edges, and transitive triadic closures. Each i -th network statistic $s_i(\cdot)$ has an associated unknown parameter θ_i . A positive value for a certain θ_i indicates that the edges involved in the formation of the corresponding network statistic $s_i(\cdot)$ are more likely to be observed relative to edges that are not involved in the formation of that network statistic, and vice versa.

The likelihood of an ERGM belongs to the exponential family of distributions and represents the probability distribution of a network graph y given a vector of parameters θ and, if available, some covariate information (x, w) :

$$f(y|\theta) = \frac{f(y; \theta)}{z(\theta)} = \frac{\exp\{\theta^T s(y; x, w)\}}{z(\theta)} \quad (1)$$

This equation states that the probability of observing a given network graph y is equal to the exponent of the observed graph statistics $s(y; x, w)$ which is a function of the network data y and covariate information (x, w) multiplied by parameter vector θ divided by a normalising constant term $z(\theta)$. The latter is calculated over the sum of all possible graphs on n nodes and it is therefore extremely difficult to evaluate for all but trivially small graphs since this sum involves $2^{\binom{n}{2}}$ terms (for undirected graphs). The intractable normalising constant makes inference difficult and standard statistical inferential tools cannot be used in this context.

2.2. Model specification components

The distinctive feature of network data is the presence of complex local dependencies among dyadic variables. In fact, a social network could be seen as a set of hypotheses about dependencies among edge variables. How could such dependencies be conceptualised? One possibility is to think in terms of a hierarchy of dependence relations generated and sustained by underlying network mechanisms. ERGMs can describe the structure of social networks by accommodating a hierarchy of network statistics reflecting dependence assumptions at different local levels such as dyadic effects involving 2-node subgraphs configurations (e.g., the number of edges), triadic effects involving 3-node subgraph configurations (e.g., the number of 2-stars), and extra-triadic effects involving subgraph configurations of more than 3 nodes (e.g., in- or out-coming edge distribution).

Social networks characterised by multiple dependencies are both empirically common, as well as theoretically relevant for understanding the mechanisms underlying observed network data. The relevance of these dependencies can be captured by analysing the effect of any network statistic by estimating its corresponding parameter. As argued in the previous section the specification of an ERGM can potentially accommodate any kind of network statistic which may be either based exclusively on connectivity informations (i.e., edge configurations) or it can be based on the interaction between network and nodal or dyadic information.

According to O'Malley [5] (2012: 547): “An attractive feature of ERGMs is their flexibility in allowing a wide range of hypotheses and sociological constructs to be tested through the inclusion of the appropriate network statistics.” In the Appendix, we define and discuss some of the main network statistics used in empirical research [4, 43, 44, 45, 46] – and included in the model that we specify in the empirical part of the paper. We describe classes of endogenous and covariate-based network statistics capturing increasing levels of dependence: dyadic, triadic, extra-triadic.

2.3. Parameter interpretation

The parameter estimates associated with the network effects expressed by the network statistics provide insights about the contribution of each network statistic to edge formation. ERGMs allow to establish a relationship between a binary outcome variable (presence/absence of a link between nodes) and a group of predictor variables (network statistics). It models the logit-transformed probability as a linear relationship with the predictor variables.

A negative (positive) parameter estimate means that the probability of observing a network with a higher value of the corresponding statistic is lower (higher) than the probability of observing the baseline network i.e., a random network generated by 0 valued parameters so that the probability of observing an edge between two nodes is 0.5:

$$\Pr(y_{ij} = 1 | \theta = 0) = \frac{\exp\{0\}}{1 + \exp\{0\}} = 0.5.$$

In other words, the value of an ERGM parameter represents the conditional log-odds of an edge between two nodes i and j , so that we have that the probability of observing a link between i and j given a negative value for the parameter θ will be

less than 0.5: $\Pr(y_{ij} = 1 | \theta < 0) < 0.5$ and, vice versa, the probability of observing a link between i and j given a positive value for the parameter θ will be greater than 0.5: $\Pr(y_{ij} = 1 | \theta > 0) > 0.5$.

3. A Bayesian approach to ERGMs

3.1. The promises of Bayesian approaches

The growing interest in Bayesian techniques for the analysis of social networks can be attributed to the development of efficient computational tools and the availability of fast user-friendly software. Bayesian analysis is a promising approach to social network analysis because it yields a rich picture of the uncertain quantities which is essential when dealing with complex social network models and heterogenous relational data. Using a Bayesian framework leads directly to the inclusion of prior information about the network effects into the modelling framework, and provides immediate access to the uncertainties by evaluating the posterior distribution of the parameters associated with the network effects.

In fact, the ease and flexibility with which prior information can be incorporated represent a major advantage of the Bayesian approach. The Bayesian approach permits the researcher to use both sample (data) and prior (expert-judgement) information in a consistent manner. For example, information from previous studies (e.g., estimates of the ERGM parameters) can easily be incorporated through an informative prior distribution. This can be done by simply placing prior restrictions on the possible values of the unknowns or by assigning a prior probability distribution based on data or summary statistics from previous studies.

In addition to this, one of the primary advantages of the Bayesian approach is the use of posterior probability. For example, based on fitting of a ERGM using an MCMC algorithm, one can obtain estimates of the posterior distributions of not only the ERGM parameters θ but also any function of the ERGM parameters (e.g., the odds ratio $\exp\{\theta\}$). Posterior means and 95% credible intervals can be used to summarise these posteriors. One can also estimate the posterior probability that a ERGM parameter is positive (or negative) or equivalently that the odds ratio is greater (or less) than 1. The posterior probability of an odds ratio less than 1 can be used in place of the p -value. This posterior probability is more intuitive than the p -value, which is the chance of observing a value as extreme as the observed value given repeated sampling under the null hypothesis.

The posterior distribution of the parameters θ given the observed network data y is defined as follows:

$$p(\theta|y) = \frac{f(y|\theta) p(\theta)}{p(y)} \quad (2)$$

This distribution cannot be evaluated directly as both the likelihood $f(y|\theta)$ defined in Equation 1 and the model evidence $p(y)$ are generally not available. For this reason the posterior defined in Equation 2 is often termed doubly-intractable distribution [47].

3.2. Prior specification

One of the most important differences between Bayesian and frequentist approaches is that the former combines available information about the unknown quantities in the prior distribution. Bayesian inference is based on the combination of prior and current information – included in the likelihood – which is reflected in the posterior distribution. In the absence of strong a priori information, prior specification should be done at the appropriate scale of interest. This often requires incorporation of weak prior information based on a vague probability distribution.

In the ERGM context, we suggest the normal distribution $\theta \sim N_d(\mu_{\text{prior}}, \Sigma_{\text{prior}})$ as a suitable prior model for the model parameters of interests, where the dimension d corresponds to the number of parameters, μ is mean vector and Σ_{prior} is a $d \times d$ covariance matrix.

3.3. Parameter estimation

As discussed earlier, the posterior distribution $p(\theta|y)$ is doubly intractable as both the likelihood normalising constant $z(\theta)$ defined in Equation (1) and model evidence $p(y)$ defined in Equation 2 cannot be directly evaluated or handled by exact sampling or standard inferential techniques such as MCMC procedures [47]. Recently, important progress has been made in Bayesian inference for doubly-intractable distributions. Approximate inference methods for doubly-intractable ERGM posterior distributions are based on the use of simulation from the likelihood defined in Equation 1 [32, 48, 49, 50]. The methodology behind these computational methods is based on the exchange algorithm proposed by [47] which can be described as follows:

- 1 - Gibbs update of (θ', y')

- i. Draw $\theta' \sim h(\cdot|\theta)$
- ii. Draw $y' \sim f(\cdot|\theta')$

2 - Accept the move from θ to θ' with probability:

$$\alpha = \min \left(1, \frac{f(y; \theta') p(\theta') h(\theta|\theta')}{f(y; \theta) p(\theta) h(\theta'|\theta)} \times \frac{f(y'; \theta)}{f(y'; \theta')} \times \frac{z(\theta) z(\theta')}{z(\theta) z(\theta')} \right)$$

where h is a proposal distribution, the distributions $f(y; \theta)$, $f(y; \theta')$, $f(y'; \theta)$, $f(y'; \theta')$ are un-normalised likelihoods of the ERGM and y' are auxiliary network data drawn from $f(\cdot|\theta')$. It is easy to see that the acceptance ratio α can be calculated since the ratio involving the normalising constants cancels out.

3.4. Computational aspects

Since the ERGM likelihood is not available, most commonly used methods for estimating ERGM parameters are simulation-based. MCMC algorithms offer the possibility of performing network simulation which is a pragmatic alternative to exact/perfect sampling [51] which, in the ERGM context, is feasible in a few cases only [52].

In practice, the simulation of network graphs from ERGMs can be generally provided by MCMC algorithms. One of the convenient ways to generate random graphs is by the Metropolis-Hastings algorithm applied to an initial adjacency matrix y_0 whose elements dyads are stochastically updated so that at each iteration t , the procedure implies that y_{t-1} and y_t differs in only one dyad. This mechanism cycles through the whole matrix so as to produce a distribution y_T which tends asymptotically to the desired random graph distribution. Recent improvements consisting in a new MLE approximation methods for moving toward a maximum likelihood estimator from an arbitrary starting parameter value have been proposed by [53].

The classical techniques may be organised into two broad categories: stochastic approximation [54] and Monte Carlo maximum likelihood estimation (MCMLE) [55, 17]. Techniques in either category may be affected by a number of estimation issues. For example, the convergence of these algorithms is often hindered by a poor choice of the starting parameter value leading to slow convergence, or failure of the algorithm to converge altogether.

The Bayesian inference based on the approximate exchange algorithm proposed by Caimo and Friel [32] has proven to be particularly suitable in the ERGM context. As described in the previous section, the algorithm consists in simulating auxiliary network data y' via MCMC and use them as an approximate draw from the ERGM likelihood (step 1-ii). However, unlike classical procedures, the approximate exchange algorithm is not sensitive to the choice of the starting parameter value, even if this value lies in the degenerate parameter region [32]. For this reason, in this paper we adopt the maximum pseudo likelihood estimate (MPLE) of the parameters (which is very fast to compute) as starting value of the main MCMC chain of the approximate exchange algorithm. This strategy turns out to be particularly useful in practice in order to reduce the initial burn-in iterations. Further advances have been recently proposed to reduce the computational effort of the algorithms for intractable likelihoods and improve their convergence properties through variance reduction techniques [50, 56] (see also [57] for a recent and comprehensive review of the computational methods for social network models).

In this paper, the computational procedures was performed by the **bergm** package for **R** [58] version 3.0.1. This package uses the **ergm** package [45, 59] as its backend for simulation and model specification.

For the analysis of the interhospital patient referral network, we used the parallel ADS procedure described in [32]. We set the number of chains to fairly large number. For each chain, we set the number of burn-in iterations equal to 100 and the number of iterations after the burn-in to 3, 000. The number of iterations used to simulate a network y' at each iteration is set to 10, 000.

3.5. Model assessment

Assessing the plausibility of an estimated model is of fundamental importance in any statistical modelling framework. The main issue concerning model assessment of social network models is the lack of standard large sample asymptotic theory. In 2006, Hunter et al. [19] proposed systematic simulation-based goodness of fit (GOF) diagnostics for ERGMs, comparing several high-level statistics of observed networks with those of corresponding networks simulated from the estimated network. The basic idea is that a fitted ERGM should reproduce structural properties similar to the observed one. The choice of the set of network statistics for constructing these GOF procedures depends on both empirical and theoretical questions.

In the Bayesian framework, in order to evaluate the model's goodness of fit in terms of posterior predictive assessment [60], the observed network is compared to a set of networks simulated from the estimated posterior distribution of the parameters of the model [32]. This comparison is carried out in terms of general network statistics distributions [19] in order to check how well the estimated posterior parameter distribution is able to reproduce networks resembling the

general structural features of the observed network. GOF distributions commonly used in statistical analysis to describe directed network structures include:

- *in-degree distribution* – the distribution of incoming edges for each node of a directed network (see $ID_d(y)$ defined in the Appendix);
- *out-degree distribution* – the distribution of outgoing edges for each node of a directed network (see $OD_d(y)$ defined in the Appendix);
- *minimum geodesic distance distribution* – the distribution of the length of shortest path distance between two nodes;
- *edgewise shared partner distribution* – the distribution of the number of unordered pairs of connected nodes having exactly k common neighbours (see $EP_d(y)$ defined in the Appendix);

Model assessment involves comparing the distributions of these various GOF diagnostic statistics $g(y)$ observed in the network data y and the set of S GOF statistics distributions $g(y_1), \dots, g(y_S)$ calculated on a set of networks y_1, \dots, y_S simulated from the model based on a sample $\theta^{(1)}, \dots, \theta^{(S)}$ drawn from the estimated posterior density.

4. Application

4.1. Research design and setting

The opportunity to illustrate the empirical value of BERGMs is provided by data that we have collected on a community of hospital organisations providing health care coverage to Lazio – one of the largest Italian regions with a population of approximately 5,700,000 inhabitants. Health care in Italy is organised on a regional basis according to federal principles. Thus, the regional health system in Lazio is part of the Italian National Health Service (INHS), a publicly funded health care system providing universal coverage free of charge at the point of service. The health system in Lazio is partitioned into twelve Local Health Units (LHUs). LHUs are vertically integrated public organisations funded by the region through a capitated budget and responsible for the provision of a wide range of services in geographical areas with target populations of approximately 500,000. LHUs represent the reference markets from which hospitals derive basic input resources – namely, patients and budgetary funds – and to which hospitals sell their services. The majority of hospitals in Italy are publicly owned, but a significant number of investor-owned, and not-for-profit hospitals also receive public funding through a contracting system with the INHS. We concentrate on patient referral relations established among 110 hospitals in Lazio in 2007. Patient referral is one of the most important forms of interhospital collaboration that has been only recently investigated in the health care literature on interorganisational network [36, 61, 62, 35, 11]. A patient referral typically occurs when (in)patients discharged from one (sender) hospital are admitted to another (receiver) hospital. Building on our fieldwork and extant research, we treat the presence of patient referral relations as the observable counterpart of the latent propensity of hospitals to collaborate via the creation of network edges.

Using publicly available data on transferred patients during the year 2007, we constructed a patient mobility adjacency matrix (v) of size 110×110 . The matrix contains in each row/column the hospital sending/receiving patients, and in the intersection cells (v_{ij}) the number of patients transferred from the row to the column hospital. The overall patient flow between hospitals is 16,557. The volume of transferred patients within dyads ranges from 0 to 360 patients, with an overall average of 1.37 (standard deviation = 9.46). On average, hospitals transferred patients out to 20 other hospitals. The matrix of patient referral relations is asymmetric, since for any hospital in the sample the number of patients sent typically differs from the number of patients received. Because we are interested in the existence of network edges, rather than their intensity, we dichotomised the matrix by using the overall mean (non-zero) number of transferred patients as cut-off value.

We model the probability of observing network edges between two hospitals as a function of a number of organisational, institutional and physical factors that may affect the propensity of organisations to collaborate. As we explain below, the modelling framework that we adopt allows us to also take important local dependencies into account and specify how they might affect endogenously the presence of network edges. Therefore this framework provides more reliable evidence on key questions of collaboration between organisations.

4.2. Data

In our empirical model specification, we control for the effects of three broad categories of exogenous factors that we think may provide meaningful explanation of the tendency of hospitals to collaborate by establishing network edges. The first category involves organisation-specific covariates. Among these covariates, the number of beds in hospitals (`n.beds`) is included to control for the effect of differences in size on the tendency of hospitals to exchange patients. We also include occupancy rate (`occ.rate`) to account for the potentially higher propensity of hospitals close to full capacity

to refer patients to partner hospitals. We include the average length of stay (`alos`) as a general measure of operational performance. We include this variable to control for the possible tendency of hospitals in our sample to move patients from low to high performance hospitals. We include the *case mix index* (`case.mix`) to capture differences in the complexity of cases treated by hospitals. Finally, we control for differences (or complementarities) in the typology of services offered to patients by reconstructing a 2-mode matrix of hospitals by the clinical specialties they contain. We then computed the *Jaccard distance* (`jacc.dist`) coefficients between hospitals in the space spanned by all the clinical specialties. The Jaccard distance is a measure of difference in terms of clinical competencies and health care services that hospitals provide. We know from prior research [63] that collaborative relations are more likely between similar organisations, that is, between organisations with overlapping competences. We formulate a prior consistent with this result and expect the corresponding parameter to be negative.

The second category involves institutional factors that may influence patterns of collaboration between hospitals. The variable *organisational form* (`org.form`) captures the institutional diversity of hospitals in the community and reflects the official classification adopted by national health authorities. Organisational form is a categorical variable ranging from 1 to 6 (1 = LHU Hospital; 2 = Hospital trust; 3 = University hospital; 4 = Hospital of the National Institute for Scientific Research; 5 = Classified hospital; 6 = Private accredited hospital). The boundaries of these categories reflect fundamental differences in institutional constraints operating on hospitals, as well as broad differences in forms of ownership and governance. The variable *LHU membership* (`LHU`) reflects another powerful categorical distinction between hospitals. LHU membership is a categorical variable, uniquely assigning each hospital to its reference administrative unit. Because of administrative, spatial and financial reasons membership in the same LHU greatly facilitates exchange relationships between hospitals and may be considered also as a general proxy for other forms of collaborations that we have not observed in our sample.

The third category involves physical factors such as *geographical distance* (`geo.dist`) that we include to control for the tendencies of hospitals closer to each other to exchange patients.

Table 1 reports the basic descriptive statistics of the network structure. Table 2 summarises essential information on the variables included in our empirical model specifications. Some variables are defined at the organisational level (i.e., monadic variables) while others refer to hospital dyads (i.e., dyadic variables).

Table 1. Descriptive statistics of the network.

Density	0.18
Mutuality	0.26
Average in-/out-degree	19.94
Average path length	1.95
Custering coefficient	0.53

Table 2. Descriptive statistics of the continuous covariates observed on the nodes of the network.

Covariate	Type	Controls for	Operationalisation	Mean (Sd)	Range (min–max)
<code>n.beds</code>	Monadic	Size of hospital	Total number of staffed beds in each hospital	197.6 (240.3)	15 – 1572
<code>occ.rate</code>	Monadic	Available capacity	Average percentage of hospital beds occupied	0.63 (0.23)	0.03 – 1
<code>alos</code>	Monadic	Operational performance	Average length of patients' stay after admission	7.06 (2.35)	2.20 – 17.60
<code>case.mix</code>	Monadic	Hospital complexity	Average Diagnosis Related Groups weights	0.97 (0.15)	0.20 – 1.60
<code>geo.dist</code>	Dyadic	Geographical distance	Distance (in km) between each pair of hospitals	49.9 (40.30)	0 – 222.60
<code>jacc.dist</code>	Dyadic	Difference in internal structure of the hospitals	Jaccard Index	0.13 (0.18)	0 – 0.79

4.3. Results

In the analysis we focus on the comparison between the dyad-independent model (which we use as a null model) and the partial conditional dependence model that incorporates fundamental network mechanisms [64, 4]. We are interested in understanding how the model reproduces known institutional features of the organisational community under examination and in comparing the ability of the different models to capture network structure. Finally, we want the analysis to illustrate specific features of the Bayesian approach to ERGMs that we believe add considerable value to the empirical analysis.

In our analysis, we incorporate prior information about the low density; high level of reciprocity; low level of covariate effect based on the geographical distance between hospitals; low level of covariate effect based on the Jaccard distance; high tendency of multiple triangles to close simultaneously. This prior information derives from empirical research that has revealed that interorganisational networks are characterised by low density, and by significant levels of reciprocity [65] and clustering [66]. Prior empirical research has also shown that collaborative relations are more likely to be formed between physically proximate organisations [67] and between organisations that are similar in terms of competences [63].

We specify a vague prior information for all the other network effects included in our models. To do that, we set all the parameters to be following a normal distribution with mean 0 and standard deviation equal to 3 except for the following parameters, where we set: $\theta_1 \sim N(-3, 1)$ assuming negative density (`edges`) effect; $\theta_2 \sim N(1, 1)$ assuming positive reciprocation (`mutual`) effect; $\theta_{11} \sim N(-1, 1)$ assuming negative `geo.dist-edgecov` effect; $\theta_{12} \sim N(-1, 1)$ assuming negative `jacc.dist-edgecov` effect; $\theta_{18} \sim N(-1, 1)$ assuming negative `gwnsp` effect; $\theta_{19} \sim N(1, 1)$ assuming positive `gwesp` effect. The covariance matrix is defined as $\Sigma_{\text{prior}} = I_d \sigma_{\text{prior}}$ where I_d is the $d \times d$ identity matrix and σ_{prior} a vector containing the diagonal entries defined above. Table 3 reports the estimated posterior means of a dyad-independent model containing only `edges` (to account for average density) and `mutual` (to capture the tendency toward reciprocation in interhospital patient referrals). We consider this specification of Model 1 as a “null model” because it involves only the simplest possible dependence assumptions restricting dependence to be only within individual dyads. The covariates are included to control for the preferential tendency of hospitals members in the same administrative unit (`LHU-homophily`), and hospitals with similar institutional identities (`org.form-homophily`) to select each other as partners. The former effect is significantly positive while the latter effect is not significant. The null model also controls for four hospital-specific covariates that may affect the propensity of hospitals to establish patient transfer relations: hospital size (as measured by number of staffed beds: `n.beds`), available capacity (as measured by the occupancy rate: `occ.rate`), average operational efficiency (as measured by the average length of stay `also`), and hospital complexity (`case.mix`).

With the exception of sharing the same institutional form (`org.form-homophily`) – which does not affect the probability of observing patient transfer relations between two hospitals – the other covariates have effects that are significant in the expected direction (i.e., the credible interval does not contain zero). Larger hospitals tend to rely on a larger number of network partners (`n.beds-sender`), and so do hospitals operating closer to their full capacity (`occ.rate`), and hospitals dealing with patients with a higher average length of stay (`alos`). The negative sign of the parameter associated with the case mix index suggests that hospitals capable of providing higher levels of care (because their complexity is high) tend to retain patients (because the associated sender effect is negative). The dyadic covariates geographical distance (`geo.dist`) and complementarity (`jacc.dist`) complete the model specification. The estimated parameter associated with the former covariate (`geo.dist`) is negative to indicate the tendency of hospitals to prefer less distant exchange partners to reduce the risks and costs inherent in patient transfer. The estimated effect of the latter (`jacc.dist`) is also negative indicating that joint problem solving is facilitated by similarity in knowledge bases – a result that resonates clearly with absorptive capacity arguments [68]. This result seems to indicate that similarity, rather than complementarity, supports collaboration between hospitals. We note that this effect is consistent with [41]. However the study of [41] is conducted at the individual rather than the organisational level.

We now focus on the estimates of the “full model” (Model 2) which incorporates the endogenous network effects of interest. The estimates are reported in Table 3 and four posterior density parameter plots associated with `edges`, `geo.dist`, `2-in-star`, and `GWESP` effects and corresponding credible intervals are displayed in Figure 1. The negative `mixed-2-star` parameter indicates a negative correlation between in-degrees and out-degrees, suggesting at least some degree of inter-hospital division of labor: it is unlikely to observe hospitals that simultaneously send and receive patients. In this model both `2-in-stars`, `2-out-stars`, geometrically weighted in-degree (`GWID`) and out-degree (`GWOD`) effects are included. `GWID` and `GWOD` give us the possibility of modelling the overall shape of in-degree and out-degree distributions whereas `2-in-stars` and `2-out-stars` can control for the presence of local hubs (upper tail of the distributions) and therefore giving us an indication of the heterogeneity of the two distributions. The positive `2-in-star` and `GWID` effects reveals a significant overall tendency toward the centralisation of the in-degrees: a limited number of hospitals seem to attract a large number of patient transfer relations. The opposite is the case for the `GWOD`: the estimate reveals a tendency away from centralisation in out-degrees balanced by the positive `2-in-star` parameter revealing the presence of out-degree heterogeneity. Interpreted together, receiving activities tend to be centralised–or

concentrated in a limited number of hospitals. This result reveals the presence of a limited number of “hubs.” On the contrary, there is no tendency toward centralisation in the out-degrees (GWOD effect) suggesting a diffuse activity of sending patients. Together these two effects seem to suggest the presence of preferential attachment – whereby popular hospitals (hospitals receiving patients from many other hospitals) tend to become ever more popular. The positive sign associated with GWESP, together with the negative sign of GWNSP suggest a tendency of hospitals in our sample to group together in closed transitive structures. In our empirical context, clustering implies a tendency of interhospital partnership to be organised in closed structures of collaborating hospitals. Organisations with multiple partners in common are more likely to be directed connected. As consequence for our sample “partners of partners are partners.” The analysis is clearly able to identify basic principles – or mechanisms of “organisational bonding” [70] underlying the observed inter-organisational community structure [71].

Noteworthy in Model 2 (Table 3) accounting for endogenous network structures is the general decrease of the effects corresponding to hospital-specific attributes included in Model 1. For example, the effect of hospital size (`n.beds`) decreases by more than three times from Model 1 to Model 2. We interpret this results as evidence that the effect of nodal attributes on network edges operates, at least in part, through network effects. The effect of geographical distance in Model 2 (`geo.dist`) is halved (in absolute value). We take this as evidence that the negative effects of physical distance on patient transfer relations is moderated by the presence of collaborative relations between partner hospitals. An exception to the general tendency of hospital-specific attributes to become weaker once endogenous effects of network structure are specified, is represented by membership in the same administrative unit (`LHU-homophily`) which more than doubles in Model 2. We take this as evidence that network-based processes operate more strongly within the boundaries of administrative units.

Table 3. Estimated posterior means, medians and 95% credible intervals for Model 1 and Model 2.

Parameter (statistic)	Model 1				Model 2			
	Mean	0.025	Median	0.975	Mean	0.025	Median	0.975
θ_1 (edges)	-4.47	-4.89	-4.45	-4.11	-5.24	-5.28	-5.24	-5.20
θ_2 (mutual)	0.87	0.71	0.88	1.02	1.04	0.99	1.04	1.08
θ_3 (LHU-homophily)	0.96	0.78	0.96	1.12	1.45	1.41	1.45	1.49
θ_4 (org.form-homophily)	0.43	0.27	0.43	0.59	0.41	0.36	0.41	0.45
θ_5 (n.beds-sender)	0.09	0.08	0.09	0.10	0.03	0.02	0.03	0.05
θ_6 (occ.rate-sender)	2.14	1.83	2.16	2.36	0.72	0.68	0.72	0.77
θ_7 (n.beds-receiver)	0.14	0.13	0.14	0.15	0.03	0.01	0.03	0.05
θ_8 (occ.rate-receiver)	1.00	0.69	1.00	1.28	0.24	0.20	0.24	0.29
θ_9 (alos-sender)	0.03	-0.01	0.03	0.07	0.01	-0.02	0.01	0.03
θ_{10} (case.mix-sender)	-0.95	-1.24	-0.95	-0.70	-0.49	-0.53	-0.49	-0.45
θ_{11} (geo.dist-edgecov)	-0.21	-0.24	-0.21	-0.18	-0.10	-0.12	-0.10	-0.07
θ_{12} (jacc.dist-edgecov)	-1.13	-1.37	-1.14	-0.86	-0.56	-0.60	-0.56	-0.52
θ_{13} (2-in-star)					0.06	0.06	0.06	0.07
θ_{14} (2-out-star)					0.06	0.06	0.06	0.07
θ_{15} (mixed-2-star)					-0.01	-0.02	-0.01	-0.01
θ_{16} (GWID, $\alpha = 0.25$)					0.61	0.57	0.60	0.65
θ_{17} (GWOD, $\alpha = 0.25$)					-0.46	-0.50	-0.46	-0.41
θ_{18} (GWNSP, $\alpha = 0.25$)					-0.06	-0.07	-0.06	-0.05
θ_{19} (GWESP, $\alpha = 0.25$)					0.44	0.40	0.44	0.49

From Table 3, we can notice that the posterior means tend to be close to the posterior medians for most of the parameters meaning that the posterior distribution is close to be symmetric.

The ERGM parameters represent conditional log-odds of observing an edge between two nodes. We can notice that the probability of observing an edge y_{ij} based on the density parameter θ_1 , conditional to all the other effects, is $\Pr(y_{ij} = 1 | \theta_1 = -4.47) = \frac{\exp\{-4.47\}}{1 + \exp\{-4.47\}} \approx 0.011$ meaning that the conditional probability of observing an edge (not involved in the creation of other network statistics included in the model) is very low as the network is sparse. Analogous calculations may be made, for example, for the homophily effect between hospitals belonging to the same LHU (LHU-homophily): $\Pr(y_{ij} = 1 | \theta_3 = 0.96) \approx 0.723$ meaning that, within LHUs, the probability of an edge between hospitals is high; and the sender effect based on the occupancy rate covariate of the hospitals (`occ.rate-sender`): $\Pr(y_{ij} = 1 | \theta_6 = 2.14) \approx 0.895$ meaning that there is high probability of observing an out-edge sent by hospitals with high occupancy rate.

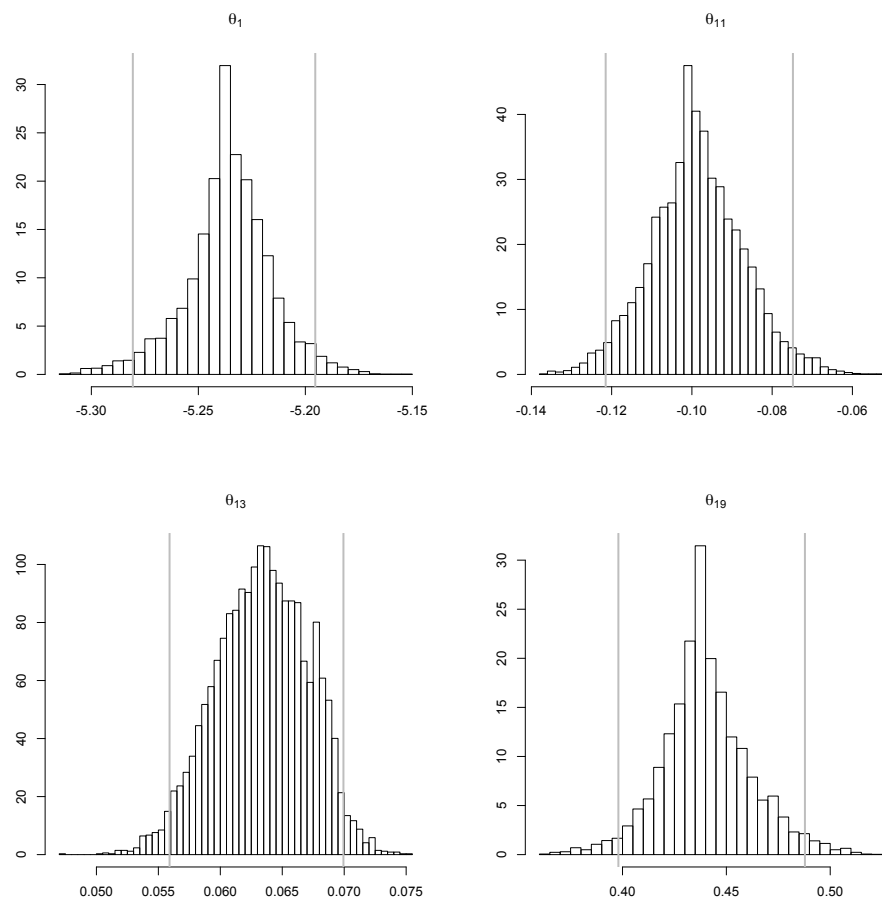


Figure 1. Posterior density plots and 95% credible intervals (vertical grey lines) for the parameters associated with edges, geo.dist, 2-in-star, and GWESP.

If we consider the edge covariate effects `geo.dist-edgecov` and `jacc.dist-edgecov`, we can calculate the conditional odds of observing a link between two nodes as a function of the *covariate value range*: $R(x) = x_{max} - x_{min}$ where x_{max} and x_{min} are respectively the maximum and minimum value of the covariate. For example, if we consider the edge covariate effect based on the geographical distance between hospitals (`geo.dist-edgecov`) we can calculate the following conditional odds: $\exp(\theta_{11} \times R(\text{geo.dist} - \text{edgecov})) = \exp(-0.21 \times 222.6) \approx 4.994 \times 10^{-21}$. This means that the probability of observing an edge connecting two hospitals that have maximum geographical distance (i.e., 222.6 km, see Table 2), conditional to all the other effects, is 4.994×10^{-21} times lower than the probability of observing an edge between two hospitals that are less than 1 km apart.

One distinctive feature of the Bayesian approach is that it allows computation of correlation between model parameters. Table 4 reports the posterior correlation matrix for the parameters in the full model.

The strong negative correlation (-0.5) between the parameter associated with the sender effect of size (θ_5) and the parameter associated with the 2-out-stars effect (θ_{14}) may be interpreted as evidence supporting the claim that these two effects are competing in explaining the out-degree structure of the network. In fact, the (dyadic) sender effect produced by big hospitals (`n.beds-sender`) tend to be “embedded” within the (triadic) out-stars effect so that the overall out-degree structure can actually be described by either effects.

A sensitivity analysis has been carried out to understand the impact of the informative prior specification. We have estimated the parameter posterior distribution by setting all the priors to be vague normal distributions (i.e., normal distributions centred at 0 with variance equal to 100). As expected, the use of informative prior distribution on some of the parameters improved the convergence of the MCMC algorithm by reducing the autocorrelation and the effective sample size. In terms of estimates we did not observe significant changes in the posterior estimates.

Table 4. Posterior correlation matrix for Model 2.

	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9	θ_{10}	θ_{11}	θ_{12}	θ_{13}	θ_{14}	θ_{15}	θ_{16}	θ_{17}	θ_{18}	θ_{19}
θ_1 (edges)	0.1	-0.1	-0.1	-0.0	0.1	-0.0	0.0	0.0	0.0	-0.1	0.1	-0.0	-0.1	-0.0	-0.1	-0.0	-0.0	-0.1
θ_2 (mutual)	1.0	-0.0	0.0	0.0	0.0	0.1	0.0	0.0	-0.1	-0.1	-0.1	-0.0	0.0	-0.1	0.0	-0.1	0.0	0.0
θ_3 (LHU-h)		1.0	0.1	0.0	0.0	-0.0	0.0	-0.1	-0.0	-0.0	0.1	0.0	0.0	-0.0	0.1	-0.0	-0.0	0.0
θ_4 (org. form-h)			1.0	-0.0	0.0	0.0	0.1	0.0	-0.1	-0.0	-0.1	-0.0	-0.0	0.0	0.0	0.0	-0.0	-0.0
θ_5 (n.beds-s)				1.0	0.0	0.1	0.0	-0.2	-0.0	-0.2	-0.1	-0.1	-0.5	-0.4	0.0	0.0	-0.1	0.1
θ_6 (occ.rate-s)					1.0	0.0	0.0	-0.1	-0.0	-0.1	-0.0	-0.0	-0.0	-0.0	-0.1	0.0	-0.0	0.0
θ_7 (n.beds-r)						1.0	-0.0	-0.2	-0.0	-0.3	-0.0	-0.7	-0.1	-0.2	-0.1	0.0	0.0	-0.0
θ_8 (occ.rate-r)							1.0	-0.1	-0.1	-0.0	-0.0	-0.0	-0.0	-0.0	0.0	0.0	-0.0	0.0
θ_9 (alos-s)								1.0	0.0	-0.1	-0.0	-0.1	-0.2	-0.1	-0.0	-0.0	-0.1	-0.1
θ_{10} (case.mix-s)									1.0	-0.0	-0.1	-0.0	-0.0	-0.0	0.1	0.1	0.1	-0.1
θ_{11} (geo.dist-ec)										1.0	-0.1	-0.1	-0.2	0.2	0.0	0.0	-0.1	-0.1
θ_{12} (jacc.dist-ec)											1.0	-0.0	0.0	0.0	0.0	-0.0	0.0	0.0
θ_{13} (2-in-star)												1.0	0.4	-0.2	0.0	-0.1	0.1	-0.0
θ_{14} (2-out-star)													1.0	-0.3	-0.0	-0.1	0.1	-0.1
θ_{15} (m-2-star)														1.0	0.0	0.0	-0.1	-0.1
θ_{16} (GWID)															1.0	-0.0	-0.0	-0.0
θ_{17} (GWOD)																1.0	-0.1	0.0
θ_{18} (GWNSP)																	1.0	-0.0
θ_{19} (GWESP)																		1.0

4.4. Goodness of fit diagnostics

Bayesian goodness-of-fit diagnostics plots are displayed in Figures 2 and 3. The red lines represent the distributions of the observed data. The boxplots represent the goodness-of-fit distributions calculated on 100 network graphs simulated from ERGMs based on values sampled from the estimated posterior distribution. The solid light grey lines mark the 95% intervals. An estimated ERGM is fitting perfectly a certain observed network if the red line falls inside this interval – a result that is very difficult to obtain in practice.

Model 2 does a much better job than Model 1 at fitting the observed network. This is particularly clear when we compare the fit of the two models in terms of the minimum geodesic distance distribution and edgewise shared partners distribution: The red line in Figure 3 (Model 2) is almost always inside the 95% interval, whereas in Figure 2 (Model 1) the red line falls outside the 95% interval for many network statistics. The number of GOF statistics of the observed network falling outside the 95% interval is considerably lower for Model 2. The better fit of Model 2 to high-level transitivity configurations (described by the edgewise shared partner distribution) is due to the presence of the extra-triadic endogenous network statistics (GWID, GWOD, GWNSP, GWESP) that are included in the model.

In Table 5 we compare the first three moments (mean, variance and skewness) of each observed GOF distributions (represented by the red lines in Figures 2 and 3) and the corresponding mean simulated GOF distributions. As we can easily see the estimates associated with Model 2 are closer to all the moments of the observed GOF distributions. In more general terms, the results of the goodness of fit diagnostic analysis that we have reported demonstrate that the estimates of Model 2 support our prediction and reproduce with accuracy important structural features of the observed network.

Table 5. Summary statistics calculated on the observed and simulated GOF distributions for Model 1 and Model 2.

	In-/Out-degree Distribution			Minimum Geodesic Distance			Edge-wise Shared Partners		
	Observed	Model 1	Model 2	Observed	Model 1	Model 2	Observed	Model 1	Model 2
Mean	20	20	20	11	2	12	14	11	13
Variance	2202	2199	2206	3022	3827	2890	2576	3824	2633
Skewness	130	129	130	117	117	117	131	147	135

Bayesian goodness-of-fit diagnostics

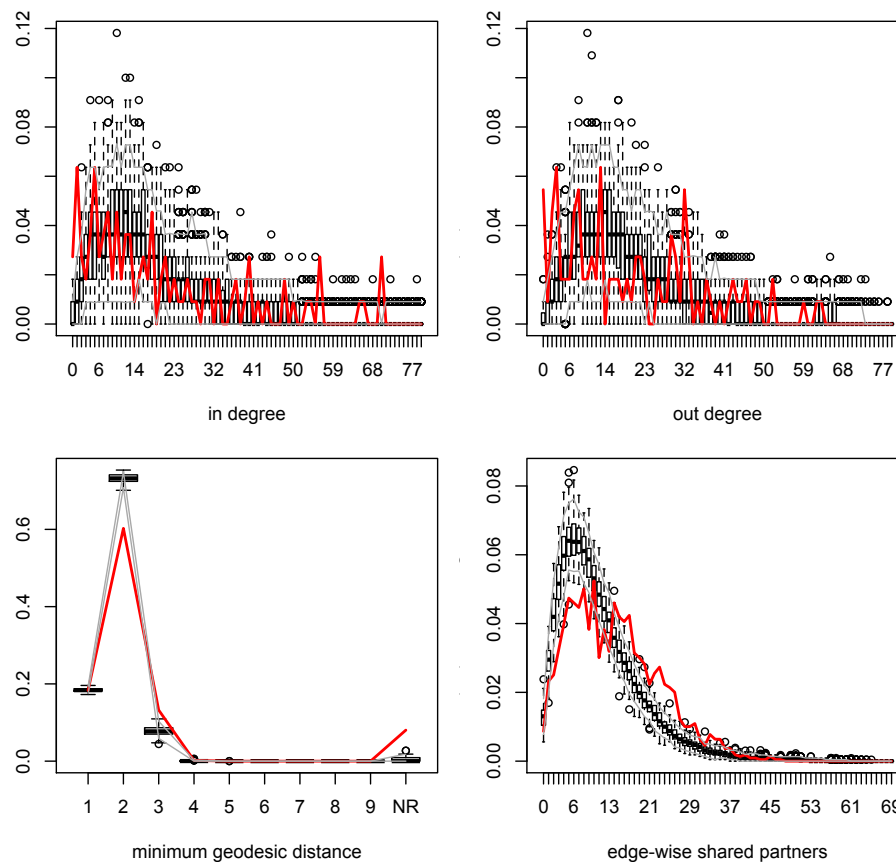


Figure 2. Goodness-of-fit diagnostics for Model 1.

5. Discussion and conclusions

A comprehensive and widely cited review published in 2003 lamented the absence of statistical models capable of accounting for some of the most common and meaningful features of networks – social or otherwise [70]. Examples of such features include: high levels of centralisation (the tendency of a limited number of nodes to generate or attract a disproportionate number of relations); high levels of clustering (the tendency of nodes sharing common contacts to be directly connected); homophily (the tendency of connected nodes to be similar along relevant dimensions), and “small-worldliness” (the tendency of networks to have low density, high clustering, and short average path-length). Until relatively recent times, available statistical models were unable to reproduce these recurrent features of networks and to identify the specific mechanisms underlying these aggregate regularities.

During the last decade, exponential random graph models (ERGMs) have been progressively developed and refined to address these issues. Today ERGMs represent perhaps the most promising models for network observed at one moment in time [8]. Yet, despite their considerable promises, adoption of ERGMs has been slowed down by issues of computational tractability, degeneracy and difficulties in parameter estimation and interpretation. Recent work on Bayesian exponential random graph models (BERGMs) promises to alleviate these problems [32, 33].

In this paper we build on this work to illustrate how BERGMs may be specified and estimated to reveal specific aspects of interhospital patient referral networks – interorganisational networks that have attracted considerable attention in studies of organisational sociology [11], public health [61], health care management [72], critical care [35] and emergency medicine [37]. We are not aware of any study of interhospital patient referral that has implemented the analytical approach we propose and illustrate in this paper. This is probably due, at least in part, to the novelty of statistical models for social networks based on BERGMs. In fact, the software implementing the estimation procedures used in this paper has only recently been made publicly available online.

We have shown that models incorporating general dependence assumptions between edge variables are more informative and more consistent with the data than models that assume strictly local (dyadic) forms of dependence

Bayesian goodness-of-fit diagnostics

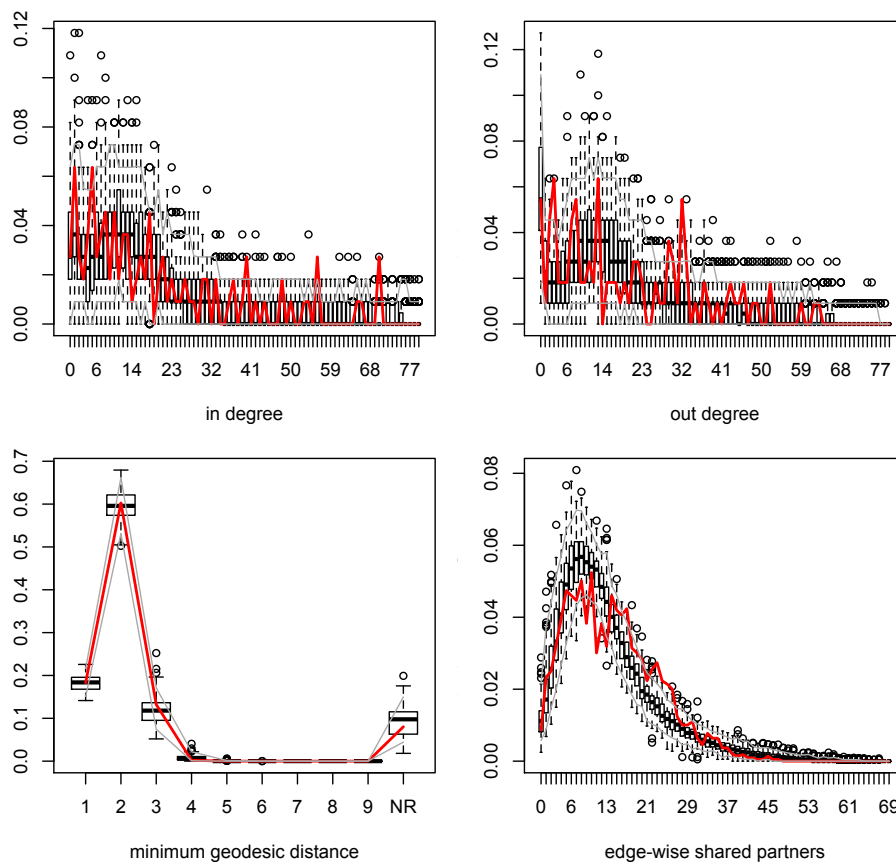


Figure 3. Goodness-of-fit diagnostics for Model 2.

between network edges. We know of no other analytical framework that could help specifying the specific forms that these network dependencies may take. We have shown that the model reproduces basic intuitions. For example, patient referral relations between distant hospitals are unlikely to be observed: as the distance between partner hospitals increases from the minimum to the maximum, the probability of interhospital relations drops asymptotically to zero. We have also shown that network edges between hospitals in the sample are produced by closure of multiple triangles simultaneously – or by multiple transitive closure. This tendency towards network closure is consistent with the idea that patient transfer relations require a considerable level of trust and social control between partner hospitals. Finally, we have shown that it is possible to formulate priors and compute posteriors for a selected number of network statistics. While the results of the analysis are somewhat specific to our sample, the method that we have illustrated is general and its application is not restricted to specific empirical settings. The joint advances in parallel computing and snowball sampling are making the characterisation of very large network datasets amenable to analysis through ERGMs [73]. In terms of future work many issues remain unsettled despite the recent important computational advances in social network analysis. These include the lack of standard criteria for assessing the goodness of fit of social network models, the lack of inferential approaches for imputing missing data (despite the recent advances made in [74] and [31]), and the lack of computational algorithms for improving the efficiency of network sampling procedures to alleviate the computational burden in carrying out parameter estimation and model selection procedures.

An alternative approach to ERGMs concerns the family of latent variable random graph models [75]. These models accommodate explanatory variables in addition to the latent distance, letting the probability of a tie between two nodes depend on the distance and on explanatory variables through a logistic (or other) link function.

Another important aspect concerns the longitudinal nature of network relations and the challenge of dealing with complex network dependencies and dynamics. While ERGMs for longitudinal data are becoming available (see, for example, [42] and [76]), there are many exciting open challenges regarding the possibility of improving the modelling tools to account for time-dependence in the ERGM parameters. In the current state of the art, model selection methods for the estimation of Bayes factor or model evidence [33, 49] struggle with high-dimensional models. Computational

methods for intractable likelihood problems based on Monte Carlo algorithms and approximations as well as likelihood-free approaches have become increasingly efficient but they are still unable to work for social networks on more than a few thousands of nodes. Recently proposed ERGM extensions for valued networks [77] still require considerable computational effort and resources to be estimated. Any development along these directions will also have a major impact on the implementation of powerful and less demanding software for end-users and practitioners.

Acknowledgement

This research was supported by the Swiss National Science Foundation (SNSF), under grant: CR1211-156229.

References

1. Snijders TA. Statistical models for social networks. *Annual Review of Sociology* 2011; **37**:131–153.
2. Salter-Townshend M, White A, Gollini I, Murphy TB. Review of statistical network analysis: Models, algorithms, and software. *Statistical Analysis and Data Mining* 2012; **5**(4):243–264.
3. O'Malley AJ, Marsden PV. The analysis of social networks. *Health Services and Outcomes Research Methodology* 2008; **8**(4):222–269.
4. Snijders TAB, Pattison PE, Robins GL, Handcock MS. New specifications for exponential random graph models. *Sociological Methodology* 2006; **36**:99–153.
5. O'Malley AJ. The analysis of social network data: An exciting frontier for statisticians. *Statistics in Medicine* 2013; **32**(4):539–555.
6. O'Malley AJ, Marsden PV. The analysis of social networks. *Health Services and Outcomes Research Methodology* 2008; **8**(4):222–269.
7. Harris JK, Cyr J, Carothers BJ, Mueller NB, Anwuri VV, James AI. Referrals among cancer services organizations serving underserved cancer patients in an urban area. *American Journal of Public Health* 2011; **101**(7):1248.
8. Fowler JH, Dawes CT, Christakis NA. Model of genetic variation in human social networks. *Proceedings of the National Academy of Sciences* 2009; **106**(6):1720–1724.
9. Parra DC, Dauti M, Harris JK, Reyes L, Malta DC, Brownson RC, Quintero MA, Pratt M. How does network structure affect partnerships for promoting physical activity? Evidence from Brazil and Colombia. *Social Science & Medicine* 2011; **73**(9):1365–1370.
10. Uddin S, Hamra J, Hossain L. Mapping and modeling of physician collaboration network. *Statistics in Medicine* 2013; **32**(20):3539–3551.
11. Lomi A, Pallotti F. Relational collaboration among spatial multipoint competitors. *Social Networks* 2012; **34**(1):101–111.
12. Holland PW, Leinhardt S. An exponential family of probability distributions for directed graphs (with discussion). *Journal of the American Statistical Association* 1981; **76**:33–65.
13. Frank O, Strauss D. Markov graphs. *Journal of the American Statistical Association* 1986; **81**:832–842.
14. Wasserman S, Pattison P. Logit models and logistic regression for social networks: I. An introduction to Markov graphs and p^* . *Psychometrika* 1996; **61**:401–425.
15. Pattison P, Wasserman S. Logit models and logistic regressions for social networks: II. Multivariate relations. *British Journal of Mathematical and Statistical Psychology* 1999; **52**(2):169–194.
16. Robins G, Pattison P, Wasserman S. Logit models and logistic regressions for social networks: III. Valued relations. *Psychometrika* 1999; **64**(3):371–394.
17. Hunter DR, Handcock MS. Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics* 2006; **15**:565–583.
18. Goodreau SM. Advances in exponential random graph (p^*) models applied to a large social network. *Social Networks* 2007; **29**(2):231–248.
19. Hunter DR, Goodreau SM, Handcock MS. Goodness of fit of social network models. *Journal of the American Statistical Association* 2008; **103**(481):248–258.
20. Robins G, Pattison P, Wang P. Closure, connectivity and degree distributions: Exponential random graph (p^*) models for directed social networks. *Social Networks* 2009; **31**(2):105–117.
21. Lusher D, Koskinen J, Robins G. *Exponential random graph models for social networks: Theory, methods, and applications*. Cambridge University Press, 2013.
22. Chatterjee S, Diaconis P. Estimating and understanding exponential random graph models. *The Annals of Statistics* 2013; **41**(5):2428–2461.
23. Schweinberger M. Instability, sensitivity, and degeneracy of discrete exponential families. *Journal of the American Statistical Association* 2011; **106**(496):1361–1370.
24. Gill J. *Bayesian methods: A social and behavioral sciences approach*, vol. 20. CRC press, 2014.
25. Caimo A, Lomi A. Knowledge sharing in organizations: A Bayesian analysis of the role of reciprocity and formal structure. *Journal of Management* 2014; **41**:665–691.
26. Zyphur MJ, Oswald FL, Rupp DE. Rendezvous overdue: Bayes analysis meets organizational research. *Journal of Management* 2015; **41**(2):387–389.
27. Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian data analysis*, vol. 2. Taylor & Francis, 2014.
28. Andrieu C, Roberts GO. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* 2009; **37**:697–725.
29. Møller J, Pettitt AN, Reeves R, Berthelsen KK. An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika* 2006; **93**(2):451–458.
30. Koskinen JH. Bayesian analysis of exponential random graphs – Estimation of parameters and model selection. *Technical Report*, Department of Statistics, Stockholm University 2004.
31. Koskinen JH, Robins GL, Pattison PE. Analysing exponential random graph (p^*) models with missing data using Bayesian data augmentation. *Statistical Methodology* 2010; **7**(3):366–384.
32. Caimo A, Friel N. Bayesian inference for exponential random graph models. *Social Networks* 2011; **33**(1):41–55.
33. Caimo A, Friel N. Bayesian model selection for exponential random graph models. *Social Networks* 2013; **35**(1):11–24.

34. Thiemichen S, Friel N, Caimo A, Kauermann G. Bayesian exponential random graph models with nodal random effects. *Social Networks* 2016; **46**:11–28.
35. Iwashyna TJ, Christie JD, Moody J, Kahn JM, Asch DA. The structure of critical care transfer networks. *Medical Care* 2009; **47**(7):787.
36. Lomi A, Mascia D, Vu DQ, Pallotti F, Conaldi G, Iwashyna TJ. Quality of care and interhospital collaboration: A study of patient transfers in Italy. *Medical Care* 2014; **52**(5):407.
37. Iwashyna TJ. The incomplete infrastructure for interhospital patient transfer. *Critical Care Medicine* 2012; **40**(8):2470–2478.
38. Lee BY, Song Y, Bartsch SM, Kim DS, Singh A, Avery TR, Brown ST, Yilmaz SL, Wong KF, Potter MA, *et al.* Long-term care facilities: Important participants of the acute care facility social network? *PLoS one* 2011; **6**(12):e29342.
39. Veinot TC, Bosk EA, Unnikrishnan K, Iwashyna TJ. Revenue, relationships and routines: The social organization of acute myocardial infarction patient transfers in the united states. *Social science & medicine* 2012; **75**(10):1800–1810.
40. Moen EL, Austin AM, Bynum JP, Skinner JS, O'Malley AJ. An analysis of patient-sharing physician networks and implantable cardioverter defibrillator therapy. *Health Services and Outcomes Research Methodology* 2016; **16**(3):132–153.
41. Landon BE, Keating NL, Barnett ML, Onnela JP, Paul S, O'Malley AJ, Keegan T, Christakis NA. Variation in patient-sharing networks of physicians across the United States. *Jama* 2012; **308**(3):265–273.
42. Koskinen J, Caimo A, Lomi A. Simultaneous modeling of initial conditions and time heterogeneity in dynamic networks: An application to foreign direct investments. *Network Science* 2015; **3**(01):58–77.
43. Robins G, Snijders T, Wang P, Handcock M, Pattison P. Recent developments in exponential random graph (p^*) models for social networks. *Social Networks* 2007; **29**(2):192–215.
44. Handcock MS, Hunter DR, Butts CT, Goodreau SM, Morris M. statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of Statistical Software* 2007; **24**(1):1–11. URL <http://www.jstatsoft.org/v24/i01>.
45. Hunter DR, Handcock MS, Butts CT, Goodreau SM, Morris M. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software* 2008; **24**(3):1–29. URL <http://www.jstatsoft.org/v24/i03>.
46. Ripley RM, Snijders TA, Preciado P. *Manual for RSIENA* 2011.
47. Murray I, Ghahramani Z, MacKay D. MCMC for doubly-intractable distributions. *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, AUAI Press: Arlington, Virginia, 2006.
48. Everitt RG. Bayesian parameter estimation for latent Markov random fields and social networks. *Journal of Computational and Graphical Statistics* 2012; **21**(4):940–960.
49. Alquier P, Friel N, Everitt R, Boland A. Noisy monte carlo: Convergence of markov chains with approximate transition kernels. *Statistics and Computing* 2016; **26**(1-2):29–47.
50. Caimo A, Mira A. Efficient computational strategies for doubly intractable problems with applications to Bayesian social networks. *Statistics and Computing* 2015; **25**:113–125.
51. Propp JG, Wilson DB. Exactly sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms* 1996; **9**:223–252.
52. Butts CT. *A Perfect Sampling Method for Exponential Random Graph Models* 2012.
53. Hummel RM, Hunter DR, Handcock MS. Improving simulation-based algorithms for fitting ergms. *Journal of Computational and Graphical Statistics* 2012; **21**(4):920–939.
54. Snijders TAB. Markov chain Monte Carlo estimation of exponential random graph models. *Journal of Social Structure* 2002; **3**.
55. Geyer CJ, Thompson EA. Constrained Monte Carlo maximum likelihood for dependent data (with discussion). *Journal of the Royal Statistical Society, Series B* 1992; **54**:657–699.
56. Friel N, Mira A, Oates CJ. Exploiting multi-core architectures for reduced-variance estimation with intractable likelihoods. *Bayesian Analysis* 2016; **11**(1):215–245.
57. Hunter DR, Krivitsky PN, Schweinberger M. Computational statistical methods for social network models. *Journal of Computational and Graphical Statistics* 2012; **21**(4):856–882.
58. Caimo A, Friel N. Bergm: Bayesian exponential random graphs in R. *Journal of Statistical Software* 2014; **61**(2):1–25.
59. Morris M, Handcock MS, Hunter DR. Specification of exponential-family random graph models: terms and computational aspects. *Journal of statistical software* 2008; **24**(4):15–48.
60. Gelman A, Meng XL, Stern H. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica* 1996; **6**(4):733–760.
61. Lee BY, McGlone SM, Song Y, Avery TR, Eubank S, Chang CC, Bailey RR, Wagener DK, Burke DS, Platt R, *et al.* Social network analysis of patient sharing among hospitals in Orange County, California. *American Journal of Public Health* 2011; **101**(4):707–713.
62. Iwashyna TJ, Courey AJ. Guided transfer of critically ill patients: Where patients are transferred can be an informed choice. *Current Opinion in Critical Care* 2011; **17**(6):641–647.
63. Powell WW, White DR, Koput KW, Owen-Smith J. Network dynamics and field evolution: The growth of interorganizational collaboration in the life sciences I. *American journal of sociology* 2005; **110**(4):1132–1205.
64. Pattison P, Robins GL. Neighbourhood-based models for social networks. *Sociological Methodology* 2002; **32**:301–337.
65. Larson A. Network dyads in entrepreneurial settings: A study of the governance of exchange relationships. *Administrative science quarterly* 1992; **37**:76–104.
66. Lomi A, Pallotti F. How to close a hole: Exploring alternative closure mechanisms in interorganizational networks 2013.
67. Whittington KB, Owen-Smith J, Powell WW. Networks, propinquity, and innovation in knowledge-intensive industries. *Administrative science quarterly* 2009; **54**(1):90–122.
68. Cohen WM, Levinthal DA. Absorptive capacity: A new perspective on learning and innovation. *Administrative science quarterly* 1990; **35**:128–152.
69. Coleman JS. Social capital in the creation of human capital. *American Journal of Sociology* 1988; **94**:S95–S120.
70. Laumann EO, Marsden PV. Microstructural analysis in interorganizational systems. *Social Networks* 1982; **4**(4):329–348.
71. Laumann EO, Galaskiewicz J, Marsden PV. Community structure as interorganizational linkages. *Annual Review of Sociology* 1978; **4**:455–484.
72. Veinot TC, Bosk EA, Unnikrishnan K, Iwashyna TJ. Revenue, relationships and routines: The social organization of acute myocardial infarction patient transfers in the United States. *Social Science & Medicine* 2012; **75**(10):1800–1810.
73. Byshkin M, Stivala A, Mira A, Krause R, Robins G, Lomi A. Auxiliary parameter MCMC for exponential random graph models. *Journal of Statistical Physics* 2016; **165**(4):740–754.
74. Handcock MS, Gile KJ. Modeling social networks from sampled data. *The Annals of Applied Statistics* 2010; **4**(1):5.

75. Hoff PD, Raftery AE, Handcock MS. Latent space approaches to social network analysis. *Journal of the American Statistical Association* 2002; **97**:1090–1098.
76. Krivitsky PN, Handcock MS. A separable model for dynamic networks. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 2014; **76**(1):29–46.
77. Krivitsky PN. Exponential-family random graph models for valued networks. *Electronic Journal of Statistics* 2012; **6**:1100–1128.

Appendix

Dyadic endogenous network statistics. Dyadic endogenous network statistics are statistics that are based exclusively on network information between pairs of nodes i and j i.e., dyads. This family of statistics represent the lower level of interactions that can be described in a network.

- *Edge statistic* (edges) is the number of edges in the network:

$$E(y) = \sum_{i \neq j} y_{ij}.$$

This statistic provides a measure of the density of the network, i.e., the proportion of potential edges between nodes in the network that are actual edges.

- *Mutual edge statistic* (mutual) is the number of mutual edges in the network:

$$M(y) = \sum_{i \neq j} y_{ij}y_{ji}.$$

This statistic reflects tendencies toward reciprocity, i.e., the portion of potential reciprocated edges between nodes in the network that are actual reciprocated edges.

Triadic endogenous network statistics. Triadic endogenous network statistics are statistics that are based exclusively on network information $s(y)$ involving three nodes i.e., triads. In this section we consider only 2-star based configurations i.e., subgraphs concerning directed edges of two nodes j and k with a third node i .

- *2-in-star statistic* (2-in-star) is the number of 2-in-stars in the network:

$$2IS(y) = \sum_{i \neq j \neq k} y_{ji}y_{ki}.$$

This statistic provides a measure of the tendency of nodes to be selected by more than one exchange partner in the network.

- *2-out-star statistic* (2-out-star) is the number of 2-out-stars in the network:

$$2OS(y) = \sum_{i \neq j \neq k} y_{ij}y_{ik}.$$

This statistic provides a measure of the tendency of partner selection toward non-exclusivity.

- *Mixed-2-star statistic* (mixed-2-star) is the number of mixed-2-stars in the network:

$$M2S(y) = \sum_{i \neq j \neq k} y_{ji}y_{ik}.$$

This statistic provides a measure of the tendency of the number of incoming edges into a node (*in-degree*) and outgoing edges from a node (*out-degree*) to be correlated.

Extra-triadic endogenous network statistics. Simulation-based inferential procedures have brought to light model degeneracy problems related to ERGMs defined by traditional specification [54, 23]. In order to overcome these issues, a new specification of network statistics based on *geometrically weighted* functions of extra-triadic network statistics distributions such as degree distributions have been proposed [4]. Below we describe some of the geometrically weighted network statistics that are common in empirical research and that we incorporate in the models discussed in the empirical part of the paper.

- *Geometrically weighted degree statistic* (GWD) is a function of the degree counts $D_d(y)$ defined as the number of nodes in y with degree d :

$$GWD(y) = \sum_d g_d(\alpha) D_d(y),$$

where $g_d(y)$ is an exponential weight function defined as:

$$g_d(\alpha) = e^\alpha \left\{ 1 - (1 - e^{-\alpha})^d \right\}.$$

In directed networks we consider the *geometrically weighted in-degree statistic* (GWID) which is a function of the in-degree counts $ID_d(y)$ defined as the number of nodes in y with in-degree d ; and the *geometrically weighted out-degree statistic* (GWOD) which is a function of the out-degree counts $OD_d(y)$ defined as the number of nodes in y with out-degree d .

- *Geometrically weighted edgewise shared partner statistics* (GWESP) is a function of the edgewise shared partner statistics $EP_d(y)$ defined as the number of unordered connected pairs (i, j) (partners) that are both connected to exactly d other nodes:

$$GWESP(y) = \sum_d g_d(\alpha) EP_d(y) = \sum_d g_d(\alpha) \sum_{i < j} y_{ij} \mathbb{1}_{\{\sum_k y_{ik} y_{jk} = d\}},$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function. For directed networks the geometric weighting is over homogeneous shared partners only (i.e., only partners on a directed two-path connecting the nodes in the edge and in the same direction).

- *Geometrically weighted non-edgewise shared partner statistics* (GWNSP) is a function of the non-edgewise shared partner statistics $NP_d(y)$ defined as the number of unordered unconnected pairs (i, j) to exactly d other nodes [45]:

$$GWNSP(y) = \sum_d g_d(\alpha) DP_d(y) = \sum_d g_d(\alpha) \sum_{i < j} \mathbb{1}_{\{\sum_k y_{ik} y_{jk} = d\}}.$$

For directed networks the count is over homogeneous shared partners only (i.e., only partners on a directed two-path connecting the nodes in the dyad). The edgewise and non-edgewise shared partner statistics reflect tendencies toward transitivity.

Dyadic covariate-based network statistics. Dyadic covariate-based network statistics are statistics equivalent to the ones defined in Section 5 which are capable of describing the network information between dyads as a function of exogenous covariate information (x, w) .

- The homophily effect (`homophily`) based on nodal covariate information (categorical) is a function of the edge statistic and it is defined as the number of connected pairs (i, j) having the same nodal information:

$$H(y; x) = \sum_{i \neq j} y_{ij} \mathbb{1}_{\{x_i = x_j\}}.$$

This statistic provides a measure of the density within groups of nodes.

- The sender effect (`sender`) based on nodal information is a function of the out-degree counts and it is defined as the sum of the number of out-edges from each node i to nodes multiplied by the sender's covariate value:

$$S(y; x) = \sum_{i < j} y_{ij} x_i.$$

This statistic may be interpreted as a measure of the interaction between nodal out-degree and covariate value.

- The covariate effect (`edgecov`) based on edge information is a function of the edge statistic and it is defined as the sum of the covariate values for each edge appearing in the network graph:

$$EC(y; w) = \sum_{i \neq j} y_{ij} w_{ij}.$$

This statistic may be interpreted as a measure of the interaction between density and dyadic covariate information.