



Title	A multilayer exponential random graph modelling approach for weighted networks
Authors(s)	Caimo, Alberto, Gollini, Isabella
Publication date	2020-02
Publication information	Caimo, Alberto, and Isabella Gollini. "A Multilayer Exponential Random Graph Modelling Approach for Weighted Networks." Elsevier, February 2020. https://doi.org/10.1016/j.csda.2019.106825 .
Publisher	Elsevier
Item record/more information	http://hdl.handle.net/10197/10992
Publisher's statement	This is the author's version of a work that was accepted for publication in Computational Statistics & Data Analysis. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in Computational Statistics & Data Analysis (142 (2020)) https://doi.org/10.1016/j.csda.2019.106825
Publisher's version (DOI)	10.1016/j.csda.2019.106825

Downloaded 2026-05-01 23:37:26

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

A multilayer exponential random graph modelling approach for weighted networks

Alberto Caimo¹ and Isabella Gollini²

¹Technological University Dublin, Ireland; alberto.caimo@dit.ie

²University College Dublin, Ireland; isabella.gollini@ucd.ie

August 2, 2019

Abstract

A new modelling approach for the analysis of weighted networks with ordinal/polytomous dyadic values is introduced. Specifically, it is proposed to model the weighted network connectivity structure using a hierarchical multilayer exponential random graph model (ERGM) generative process where each network layer represents a different ordinal dyadic category. The network layers are assumed to be generated by an ERGM process conditional on their closest lower network layers. A crucial advantage of the proposed method is the possibility of adopting the binary network statistics specification to describe both the between-layer and across-layer network processes and thus facilitating the interpretation of the parameter estimates associated to the network effects included in the model. The Bayesian approach provides a natural way to quantify the uncertainty associated to the model parameters. From a computational point of view, an extension of the approximate exchange algorithm is proposed to sample from the doubly-intractable parameter posterior distribution. A simulation study is carried out on artificial data and applications of the methodology are illustrated on well-known datasets. Finally, a goodness-of-fit diagnostic procedure for model assessment is proposed.

Keywords— Statistical network models; Bayesian analysis; Weighted networks; Intractable models.

1 Introduction

Statistical network analysis concerns modelling relationships defined by edges between nodes [Salter-Townshend et al., 2012]. In many empirical contexts these relationships have a strength

associated with their edges [Barrat et al., 2004]. The nature of the variation in the strength of an edge between two nodes may be determined by a variety of aspects depending on the application context; for example, the amount of traffic flowing along connections in transportation networks [Opsahl et al., 2008], the functional connectivity levels in brain networks [Bullmore and Sporns, 2009], and interactions in cellular and genetic networks [Horvath, 2011].

In the context of social network analysis, one of the most important families of models used to describe the connectivity network structure is represented by exponential random graph models (ERGMs) [Holland and Leinhardt, 1981, Strauss and Ikeda, 1990]. These are generative models for networks postulating an exponential family over the sample space of networks on the fixed set of nodes and are specified by a set of sufficient network statistics [Besag, 1974] representing sub-graph configurations that are believed to be important to the generative process of the observed network structure. ERGMs are very flexible models as they can potentially incorporate any type of network statistic and were originally defined for networks with binary edges encoding the presence or absence of an edge between two nodes. Commonly used network statistics include summaries of density (e.g., number of edges), homophily (e.g., number of edges among nodes with the same nodal attribute), degree-based statistics (e.g., number of k-stars), and triad-based statistics (e.g., number of triangles) [Snijders et al., 2006].

Several modelling extensions based on ERGMs have been proposed. These include generalisations of binary ERGMs to polytomous forms of weighted edges [Robins et al., 1999, Pattison and Wasserman, 1999]; the multi-valued curved ERGMs [Wyatt et al., 2010]; ERGMs for inference on networks with continuous edge values [Desmarais and Cranmer, 2012, Wilson et al., 2017]; the Geometric/Poisson reference ERGMs for ordinal/count networks [Krivitsky, 2012]; and the rank-order-edge ERGMs [Krivitsky and Butts, 2017].

In this paper, we introduce a Bayesian hierarchical multilayer ERGM approach for ordinal/polytomous network data in order to simplify model specification and provide a substantial improvement in terms of interpretability by making use of binary ERGMs to capture the dependence structure between ordinal categories (network layers) of the weighted network structure. The paper is organised as follows. In Section 2, we review the main features of ERGMs. In Section 3 we show how multilayer graphs can be used to represent weighted network structures. In Section 4 we introduce the multilayer ERGM approach. In Section 4.3 we focus on the interpretation of the multilayer ERGM framework as a dissolution process which is able to capture the between-layer generative process between the network layers. In Section 5 we generalise the modelling framework using a Bayesian hierarchical modelling approach and we propose to extend the approximate exchange algorithm [Caimo and Friel, 2011] to sample from the doubly-intractable parameter posterior distribution (see Park and Haran [2018] for a recent review). In Section 6, we test our methodology on simulated data and, in Section 7 we demonstrate the usefulness of the multilayer ERGM approach by analysing two very well-known network datasets and assessing model fit by comparing goodness of fit diagnostics based on the weighted degree distribution. Final remarks are provided in the Conclusions in Section 8.

2 Exponential random graph models (ERGMs)

Typically networks consist of a set of actors and relationships between pairs of them, for example, social interactions between individuals. The relational structure of a network graph is described by a random adjacency matrix \mathbf{Y} of a graph on N nodes (actors) and a set of edges (relationships) $\{Y_{i,j} : i = 1, \dots, N; j = 1, \dots, N\}$ where $Y_{i,j} = 1$ if nodes i and j are connected and $Y_{i,j} = 0$ otherwise. The network graph \mathbf{Y} may be directed or undirected ($Y_{i,j} = Y_{j,i}$) depending on the nature of the relationships between the actors and \mathbf{y} a realisation of \mathbf{Y} . Self-loops are generally not allowed: $Y_{i,i} = 0$.

Exponential random graph models (ERGMs) are a particular class of discrete linear exponential families which represents the probability distribution of \mathbf{Y} as:

$$p(\mathbf{y}|\boldsymbol{\theta}) = \frac{\exp\{\boldsymbol{\theta}^\top s(\mathbf{y})\}}{c(\boldsymbol{\theta})}, \quad (1)$$

where $s(\mathbf{y}) \in \mathbb{R}^r$ is a known vector of r sufficient statistics, $\boldsymbol{\theta} \in \mathbb{R}^r$ is the associated parameter vector, and $c(\boldsymbol{\theta})$ a normalising constant which is difficult to evaluate for all but trivially small graphs. The dependence hypothesis at the basis of these models is that edges self organise into small structures called configurations. There is a wide range of possible network configurations [Robins et al., 2007] which gives flexibility to adapt to different contexts. A positive parameter value for θ_i results in a tendency for the certain configuration corresponding to $s_i(\mathbf{y})$ to be observed in the data more frequently than would otherwise be expected by chance.

2.1 Curved ERGMs

The ERGM likelihood defined in Equation 1 can be generalised by assuming that the parameter $\boldsymbol{\beta}(\boldsymbol{\theta})$ is nonlinear in the exponential family of distribution

$$p(\mathbf{y}|\boldsymbol{\theta}) = \exp\left\{\boldsymbol{\beta}(\boldsymbol{\theta})^\top s(\mathbf{y}) - c(\boldsymbol{\beta}(\boldsymbol{\theta}))\right\},$$

which is therefore a curved exponential family [Hunter, 2007]. Curved ERGMs are implemented by specifying geometrically weighted network statistics [Snijders et al., 2006] and are commonly used to alleviate ERGM degeneracy issues [Handcock, 2003]. However estimating the decay parameters associated to these network effects is a challenging issue and therefore these parameters are generally fixed rather than estimated. Recent contributions by Bomiriyi et al. [2016] and Stewart et al. [2019] have shown how curved ERGMs can be estimated in the context of Bayesian inference of bipartite and classical inference for multilevel network data, respectively.

3 Weighted networks as multilayer graphs

A weighted network can be described by an $N \times N$ weighted adjacency matrix \mathbf{Y} where $Y_{i,j} \neq 0$ if nodes i and j are connected and $Y_{i,j} = 0$ otherwise. The value of $Y_{i,j}$ can be positive or negative, ordinal, count, bounded continuous and unbounded continuous. We denote with D and E the number of dyads and (weighted) edges of the network respectively.

In this paper, we focus on ordinal and polytomous networks and we describe their connectivity structure through three-dimensional adjacency arrays [Robins et al., 1999] corresponding to what we define as a *multilayer network structure*. Let \mathbf{Y} be a weighted network on N nodes with edges that are ordinal random variables taking value $w = 1, 2, \dots, W$. We consider a three-dimensional array $N \times N \times W$ that we denote with $\mathbf{y}_{\{W\}}$, consisting of a set of *network layers* $\{\mathbf{Y}_w, w = 1, \dots, W\}$ each of which represents a binary adjacency matrix with dyads defined as follows:

$$Y_{i,j,w} = \begin{cases} 1, & y_{i,j} \geq w; \\ 0, & y_{i,j} < w. \end{cases}$$

The set of edges observed in the w -th layer, \mathcal{E}_w , is a subset of the edges observed in the lower network layers, i.e.: $\mathcal{D}_1 \supset \mathcal{E}_1 \supset \mathcal{E}_2 \supset \dots \supset \mathcal{E}_W$, where \mathcal{D}_1 is the set of dyads of the weighted network (Figure 1).




Figure 1: An example of multilayer structure of a weighted network. The network $\mathbf{y}_{\{3\}}$ consists of edges that can take three ordinal values and its connectivity structure can be described by three overlapping binary network layers.

4 Multilayer ERGMs

The basic assumption of the multilayer network model is that the probability of observing the overall multilayer relational structure $\mathbf{y}_{\{W\}}$ corresponds to the product of conditional binary network models describing the transition processes between consecutive network layers \mathbf{y}_w and \mathbf{y}_{w+1} :

$$p(\mathbf{y}_{\{W\}}|\boldsymbol{\theta}) = p(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_W|\boldsymbol{\theta}) = p(\mathbf{y}_1|\boldsymbol{\theta}) \prod_{w=1}^{W-1} p(\mathbf{y}_{w+1}|\mathbf{y}_w, \boldsymbol{\theta}), \quad (2)$$

where $\boldsymbol{\theta}$ is a vector of parameters, $p(\mathbf{y}_1|\boldsymbol{\theta})$ is the unconditional network model for the first layer \mathbf{y}_1 , $p(\mathbf{y}_{w+1}|\mathbf{y}_w, \boldsymbol{\theta})$ is the transition probability from layer \mathbf{y}_w to layer \mathbf{y}_{w+1} by assuming

first order Markov dependence between layers (this assumption may be generalised to higher order Markov dependence), and $p(\mathbf{y}_{W+1}|\mathbf{y}_W, \boldsymbol{\theta})$ is the transition probability from the last layer \mathbf{y}_W to the completely empty network graph \mathbf{y}_{W+1} (i.e., $y_{i,j,W+1} = 0, \forall i, j \in \mathcal{D}_1$).

4.1 The multilayer random graph model

The simplest multilayer network model is the multilayer random graph model which is based on the assumption of dyadic independence. Let p denote the probability that two nodes i and j are connected in a network layer w and $q = 1 - p$. The multilayer random graph model can be defined as:

$$\begin{cases} p(\mathbf{y}_1|p) = p^{E_1} q^{D_1 - E_1}; \\ p(\mathbf{y}_{w+1}|\mathbf{y}_w, p) = p^{E_{w+1}} q^{E_w - E_{w+1}} = p^{E_{w+1}} q^{D_{w+1} - E_{w+1}}, \quad w = 1, \dots, W; \end{cases} \quad (3)$$

where D_1 is the overall number of dyads in the weighted network and E_w is the number of edges in the network layer \mathbf{y}_w which corresponds to the number of random dyads D_{w+1} allowed to take value 1 in the network layer \mathbf{y}_{w+1} .

The multilayer random graph model defined in Equation 3 corresponds to the weighted random graph model defined by Garlaschelli [2009] which is equivalent to the geometric-reference random graph model defined by Krivitsky [2012]. In fact, merging Equations 2 and 3, we have that:

$$p(\mathbf{y}_{\{W\}}|p) = \prod_{w=1}^W p^{E_w} q^{D_w - E_w} q^{E_W} = p^{\sum_{w=1}^W E_w} q^{\sum_{w=1}^W D_w - \sum_{w=1}^{W-1} E_w}.$$

Since:

$$\sum_{w=1}^W D_w = D_1 + \sum_{w=2}^W D_w = D_1 + \sum_{w=1}^{W-1} E_w,$$

we have that:

$$p(\mathbf{y}_{\{W\}}|p) = p^{\sum_{w=1}^W E_w} q^{D_1} = p^{\sum_{w=1}^W \sum_{i,j \in \mathcal{D}_1} y_{i,j,w}} q^{D_1} = \prod_{(i,j) \in \mathcal{D}_1} p^{\sum_{w=1}^W y_{i,j,w}} q,$$

which corresponds to the product of the probabilities that any two nodes (i, j) are connected by an edge of weight $\sum_{w=1}^W y_{i,j,w} = y_{i,j}$ and therefore each weighted dyad $Y_{i,j} \stackrel{i.i.d.}{\sim} \text{Geometric}(q)$.

The exponential form of the geometric-reference random graph model defined in Equation 3 is:

$$p(\mathbf{y}_{\{W\}}|\theta) = \prod_{(i,j) \in \mathcal{D}_1} \frac{\exp\{\theta y_{i,j}\}}{1 - \exp(\theta)} = \frac{\exp\left\{\theta \sum_{(i,j) \in \mathcal{D}_1} y_{i,j}\right\}}{c(\theta)},$$

where $c(\theta)$ is a normalising constant, $1 - \exp(\theta) = q$ and $\exp(\theta) = p$ and therefore $\theta = \ln(p) \leq 0$.

4.2 Multilayer ERGMs as dissolution processes

The multilayer random graph model defined in Equation 3 can be easily extended by adopting a general multilayer ERGM generative process incorporating extra-dyadic network effects (such as degree-based statistics or transitive configurations) and thus relaxing the simplistic dyadic dependence assumption of the multilayer random graph model. In Table 1, we can observe that edges between two nodes in layer \mathbf{y}_w may or may not ‘survive’ to the next upper layer \mathbf{y}_{w+1} whereas empty edges in layer \mathbf{y}_w remain empty edges in all the upper layers.

Table 1: Possible transitions of a single dyadic variable between two consecutive network layers.

$Y_{i,j,w}$	\rightarrow	$Y_{i,j,w+1}$
0	\rightarrow	0
1	\rightarrow	0
1	\rightarrow	1

We assume that the transition process between consecutive network layers is a conditional ERGM model

$$p(\mathbf{y}_{w+1} | \mathbf{y}_w, \phi) = \frac{\exp\{\phi^\top s(\mathbf{y}_{w+1}; \mathbf{y}_w)\}}{c(\phi, \mathbf{y}_w)},$$

where ϕ is a parameter vector associated to $s(\mathbf{y}_{w+1}; \mathbf{y}_w)$ that is the vector of network statistics that do not dissolve in the transition process from \mathbf{y}_w to \mathbf{y}_{w+1} . This vector consists of standard binary network statistics such as the number of edges, stars, triangles, etc. that are present in both network layers \mathbf{y}_w and \mathbf{y}_{w+1} .

The transition process from the last network layer \mathbf{y}_W to the empty network graph \mathbf{y}_{W+1} is deterministic as $Y_{i,j,W+1} = 0, \forall i, j \in \mathcal{D}_1$ regardless the value of $Y_{i,j,W}$ so it is not included in the model.

The conditional log-odds of an edge between nodes i and j in layer \mathbf{y}_{w+1} , given the presence of an edge between i and j in layer \mathbf{y}_w while keeping all the rest of the network layer \mathbf{y}_{w+1} fixed is:

$$\ln \left(\frac{\Pr(Y_{i,j,w+1} = 1 | Y_{i,j,w} = 1, \mathbf{Y}_{-(i,j,w+1)})}{\Pr(Y_{i,j,w+1} = 0 | Y_{i,j,w} = 1, \mathbf{Y}_{-(i,j,w+1)})} \right) = \phi^\top \Delta(\mathbf{y})_{i,j,w+1}, \quad (4)$$

where ϕ is a vector of parameters and $\Delta(\mathbf{y})_{i,j,w+1}$ is the vector of change network statistics from \mathbf{y}_w to \mathbf{y}_{w+1} , i.e., the change in the value of the network statistic $s(\cdot)$ that would occur if $y_{i,j,w+1}$ were changed from 0 to 1 while leaving all of the rest of \mathbf{y}_{w+1} fixed.

From Table 1, we can notice that the formation process is not allowed when moving from layer \mathbf{y}_w to layer \mathbf{y}_{w+1} , so the conditional log-odds of an edge between nodes i and j in layer \mathbf{y}_{w+1} , given the absence of an edge between nodes i and j in layer \mathbf{y}_w is:

$$\ln \left(\frac{\Pr(Y_{i,j,w+1} = 1 | Y_{i,j,w} = 0, \mathbf{Y}_{-(i,j,w+1)})}{\Pr(Y_{i,j,w+1} = 0 | Y_{i,j,w} = 0, \mathbf{Y}_{-(i,j,w+1)})} \right) = -\infty, \quad (5)$$

as $\Pr(Y_{i,j,w+1} = 1 | Y_{i,j,w} = 0, \mathbf{Y}_{-(i,j,w+1)}) = 0$ and $\Pr(Y_{i,j,w+1} = 0 | Y_{i,j,w} = 0, \mathbf{Y}_{-(i,j,w+1)}) = 1$. Therefore the parameter ϕ associated with the network effects expressed by the network

statistics $s(\cdot)$ provides insights about the contribution of each network statistic to edge dissolution between consecutive network layers.

If we consider the multilayer random graph model defined in Section 4.1, where only the number of edges is included in the model, we obtain $\Delta(\mathbf{y})_{i,j,w+1} = 1$ and therefore the relationship between ϕ and θ is: $\phi = \ln(p) - \ln(q) = \theta - \ln(1 - \exp(\theta))$.

The likelihood of multilayer ERGMs defined by Equations 4 and 5 can be written as:

$$p(\mathbf{y}_{\{W\}}|\phi) = \frac{\exp\{\phi^\top s(\mathbf{y}_{\{W\}})\}}{c(\phi, \mathbf{y}_{\{W\}})} = \frac{\exp\{\phi^\top s(\mathbf{y}_1)\}}{c(\phi, \mathbf{y}_1)} \prod_{w=1}^{W-1} \frac{\exp\{\phi^\top s(\mathbf{y}_{w+1}; \mathbf{y}_w)\}}{c(\phi, \mathbf{y}_w)}.$$

The interpretation of the multilayer ERGM as a dissolution process implies that many network statistics developed for ERGMs can be readily used within this modelling framework, retaining much of their interpretation. A positive value for the parameter ϕ_i corresponding to a particular network statistic $s_i(\cdot)$ increases the probability of observing that network statistic in the next upper layer or vice versa. This means that the dissolution process between two consecutive layers can be interpreted exactly like a standard binary ERGM process defined in Section 2. An analogous representation of multilayer ERGMs consists in considering a formation process from layer W to 1 by just reversing the direction of the dependence structure between layers.

The multilayer ERGMs is similar to a static version of discrete separable temporal ERGMs [Krivitsky and Handcock, 2014] which is including initial conditions (i.e., the baseline network connectivity structure represented by the first network layer \mathbf{y}_1 , see also the longitudinal ERGM approach developed by Koskinen et al. [2015]) but without formation process (which is not needed in this weighted network context).

4.3 Relaxing the homogeneity assumption across layers

We now generalise the multilayer model introduced above by relaxing the parameter homogeneity assumption across network layers by considering layer-specific ERGM processes:

$$p(\mathbf{y}_{\{W\}}|\phi_1, \dots, \phi_W) = p(\mathbf{y}_1|\phi_1) \prod_{w=1}^{W-1} p(\mathbf{y}_{w+1}|\mathbf{y}_w, \phi_{w+1}).$$

where ϕ_1, \dots, ϕ_W are the between-layer parameters capturing specific network effects that might characterise the transition processes between consecutive network layers. In fact, the behaviour of some network effects might vary depending on the layer values: we can for example imagine that for lower layers some network effects are positive and for higher layers the same effects are negative, or vice versa.

It is important to notice that the first ERGM transition $p(\mathbf{y}_1|\phi_1)$ is a standard binary ERGM that is equivalent to an ERGM dissolution process defined in Section 4.2 conditional on the fully connected binary network graph that we denote by \mathbf{y}_0 defined on the same set of nodes, i.e., $y_{i,j,0} = 1, \forall i, j \in \mathcal{D}_1$. Consequently, we have that $p(\mathbf{y}_1|\phi_1) = p(\mathbf{y}_1|\mathbf{y}_0, \phi_1)$, as:

$$\ln \left(\frac{\Pr(Y_{i,j} \geq 1 \mid \mathbf{Y}_{-(i,j)})}{\Pr(Y_{i,j} = 0 \mid \mathbf{Y}_{-(i,j)})} \right) = \ln \left(\frac{\Pr(Y_{i,j,1} = 1 \mid y_{i,j,0} = 1, \mathbf{Y}_{-(i,j,1)})}{\Pr(Y_{i,j,1} = 0 \mid y_{i,j,0} = 1, \mathbf{Y}_{-(i,j,1)})} \right).$$

This modelling approach yields a flexible way to specify between-layer transition processes. In fact, it is possible to specify different network statistics for modelling different network transitions.

The conditional log-odds of an edge with weight w^* between nodes i and j is:

$$\ln \left(\frac{\Pr(Y_{i,j} = w^* \mid \mathbf{Y}_{-(i,j)})}{\Pr(Y_{i,j} = 0 \mid \mathbf{Y}_{-(i,j)})} \right) = \sum_{w=1}^{w^*} \phi_w^\top \Delta(\mathbf{y})_{i,j,w} .$$

5 Bayesian inference

Bayesian analysis is a promising approach to social network analysis because it yields a rich fully-probabilistic evaluation of uncertainty which is essential when dealing with complex and heterogenous relational data.

The growing popularity of Bayesian techniques for ERGMs can be attributed to the development of efficient computational methods [Caimo and Friel, 2011, Alquier et al., 2016, Caimo and Mira, 2015, Bouranis et al., 2017, 2018] and the availability of user-friendly software tools such as the `Bergm` [Caimo and Friel, 2014, Caimo et al., 2019] and `hergm` [Schweinberger and Luna, 2018, Schweinberger et al., 2019] packages for R [R Core Team, 2019].

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 social network analysis the Bayesian approach leads to the possibility of specifying informative parameter prior distribution consistent with some a priori expectation, for example, in terms of low density and high transitivity [Caimo et al., 2017]. In fact, parameter prior distribution can be concentrated on negative values for the density parameter and positive values for transitivity parameters and/or positive correlation between density and transitivity parameters.

In the following sections we will be extending the modelling framework introduced in Section 4.3 and describing a Bayesian parameter estimation procedure based on the approximate exchange algorithm [Murray and Ghahramani, 2004, Caimo and Friel, 2011] which can sample from the ERGM posterior distribution which is doubly-intractable as both ERGM likelihood and marginal likelihood are not available.

5.1 A hierarchical framework

The Geometric-reference ERGM approach [Garlaschelli, 2009, Krivitsky, 2012] is not capable of modelling every between-layer dependence process in the network and, on the other hand, the multilayer modelling approach defined in Section 4.3 is not capable of capturing overall trends of the network effects across the entire multilayer structure. For this reason, in order to model both the relational processes simultaneously and improve the goodness of fit of our model, we specify a hierarchical multilayer ERGM where layer specific parameters ϕ_1, \dots, ϕ_W are coupled through a random variable $\boldsymbol{\eta}$ representing the overall across-layer trend of the r network effects of interest with prior distribution $p(\boldsymbol{\eta} \mid \boldsymbol{\gamma}_0)$ defined by the hyper-parameters $\boldsymbol{\gamma}_0$.

The benefit of using Bayesian hierarchical approaches for ERGMs has been shown in several contexts. In particular, [Slaughter and Koehly \[2016\]](#) demonstrated how these methods can be useful to describe the systematic patterns within groups and how these structural patterns differ across groups in multilevel networks.

To develop a hierarchical multilayer approach, we define the following model:

$$p(\phi_1, \dots, \phi_W, \boldsymbol{\eta} | \mathbf{y}_{\{W\}}) \propto p(\mathbf{y}_{\{W\}} | \phi_1, \dots, \phi_W) p(\phi_1, \dots, \phi_W | \boldsymbol{\eta}) p(\boldsymbol{\eta} | \gamma_0) \quad (6)$$

and its structure is displayed in Figure 2.




Figure 2: Graphical structure of the hierarchical multilayer ERGM defined in Equation 6.

We define a r -dimensional model assuming that the between layer parameters are independent realisations from a Normal distribution, i.e.:

$$p(\phi_1, \dots, \phi_W | \boldsymbol{\eta}) = \prod_{w=1}^W p(\phi_w | \boldsymbol{\mu}, \boldsymbol{\Sigma}),$$

where: $\phi_w | \boldsymbol{\mu}, \boldsymbol{\Sigma} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

We assume a Normal-Inverse-Wishart prior setting where:

$$\boldsymbol{\Sigma} \sim \mathcal{W}^{-1}(\boldsymbol{\Lambda}_0, \nu_0); \quad \boldsymbol{\mu} | \boldsymbol{\Sigma} \sim \mathcal{N}\left(\boldsymbol{\mu}_0, \frac{\boldsymbol{\Sigma}}{\kappa_0}\right),$$

so that the joint prior distribution for $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ is:

$$\boldsymbol{\mu}, \boldsymbol{\Sigma} \sim \mathcal{NW}^{-1}(\boldsymbol{\mu}_0, \kappa_0, \boldsymbol{\Lambda}_0, \nu_0).$$

The full conditional distribution of $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ is:

$$\boldsymbol{\mu}, \boldsymbol{\Sigma} | \phi_1, \dots, \phi_W \sim \mathcal{NW}^{-1}(\boldsymbol{\mu}_1, \kappa_1, \boldsymbol{\Lambda}_1, \nu_1);$$

with parameters:

$$\boldsymbol{\mu}_1 = \frac{\kappa_0}{\kappa_0 + W} \boldsymbol{\mu}_0 + \frac{W}{\kappa_0 + W} \bar{\boldsymbol{\phi}}; \quad \kappa_1 = \kappa_0 + W;$$

$$\boldsymbol{\Lambda}_1 = \boldsymbol{\Lambda}_0 + \boldsymbol{S} + \frac{\kappa_0}{\kappa_0 + W} W (\bar{\boldsymbol{\phi}} - \boldsymbol{\mu}_0)(\bar{\boldsymbol{\phi}} - \boldsymbol{\mu}_0)^\top; \quad \nu_1 = \nu_0 + W;$$

where:

$$\bar{\boldsymbol{\phi}} = \frac{1}{W} \sum_{w=1}^W \phi_w; \quad \boldsymbol{S} = \sum_{w=1}^W (\phi_w - \bar{\boldsymbol{\phi}})(\phi_w - \bar{\boldsymbol{\phi}})^\top.$$

5.2 Posterior estimation

To estimate the posterior distribution $p(\phi_1, \dots, \phi_W, \boldsymbol{\mu}, \boldsymbol{\Sigma} | \mathbf{y}_{\{W\}})$ defined in Equation 6, we extend the approximate exchange algorithm of [Caimo and Friel \[2011\]](#) to sample from the doubly-intractable distribution $p(\phi_1, \dots, \phi_W | \mathbf{y}_{\{W\}}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ and then we sample $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ from the full conditional distribution via Gibbs sampling.

The approximate exchange algorithm, described in Algorithm 1, is an asymptotically exact MCMC algorithm as it guarantees asymptotically exact recovery of the posterior distribution as the number of posterior samples increases [[Everitt, 2012](#)].

Algorithm 1 Approximate exchange algorithm for $p(\phi_1, \dots, \phi_W | \mathbf{y}_{\{W\}}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$

Initialise $\phi_1^{(1)}, \dots, \phi_W^{(1)}$

for $i = 1, \dots, I$ **do**

for $w = 1, \dots, W$ **do**

 1) $\phi'_w \sim h(\cdot | \phi_w^{(i)})$

 2) $\mathbf{y}'_w \sim p(\cdot | \phi'_w, \mathbf{y}_{w-1})$ via MCMC (see Algorithm 2)

 3) Set $\phi_w^{(i+1)} = \phi'_w$ with probability:

$$\min \left(1, \frac{q_{\phi_w^{(i)}}(\mathbf{y}'_w; \mathbf{y}_{w-1}) p(\phi_1, \dots, \phi'_w, \dots, \phi_W | \boldsymbol{\mu}, \boldsymbol{\Sigma}) q_{\phi'_w}(\mathbf{y}_w; \mathbf{y}_{w-1})}{q_{\phi_w^{(i)}}(\mathbf{y}_w; \mathbf{y}_{w-1}) p(\phi_1, \dots, \phi_w^{(i)}, \dots, \phi_W | \boldsymbol{\mu}, \boldsymbol{\Sigma}) q_{\phi_w^{(i)}}(\mathbf{y}'_w; \mathbf{y}_{w-1})} \right)$$

end for

end for

The unnormalised between-layer ERGM likelihood is defined as:

$$q_{\phi_w}(\mathbf{y}_w; \mathbf{y}_{w-1}) = \exp\{\phi_w^\top s(\mathbf{y}_{w+1}; \mathbf{y}_w)\},$$

and $h(\cdot)$ is a proposal distribution for updating the model parameters. Adaptive strategies [[Caimo and Friel, 2011](#), [Caimo and Mira, 2015](#)] and approximate transition kernels approaches [[Alquier et al., 2016](#)] have been successfully implemented in this context.

From a computational viewpoint, the algorithm becomes increasingly expensive as the network size and the number of layers increase. In fact, the number of MCMC iterations needed to simulate the auxiliary network layers at Step 2 of Algorithm 1 is proportional to the number of dyads in the network [[Everitt, 2012](#)] and the MCMC sampling is repeated for generating the binary network layers of the observed adjacency array as many times as the value of maximum weight (W) that an edge can take.

However, it is important to notice that the number of iterations needed for simulating the network layer \mathbf{y}'_{w+1} should be just proportional to the number of edges (E_w) in the lower network layer \mathbf{y}_w which is a subset of the overall number of dyads (\mathcal{D}_1) in the network. In fact, according to Equation 4, only the subset of dyads belonging to \mathcal{D}_{w+1} and corresponding to the non-zero dyads of the network layer \mathbf{y}_w can take value $w + 1$ and they are the only dyads involved in the network simulation of layer $\mathbf{y}'_{w+1} \sim p(\cdot | \phi'_{w+1}, \mathbf{y}_w)$.

The constrained ERGM simulation for the network layer \mathbf{y}_{w+1} corresponding to Step 2 of Algorithm 1 is described in detail in Algorithm 2. The number of MCMC steps I_{w+1} required for simulating a network layer \mathbf{y}_{w+1} can be smaller than the number of steps required for simulating any lower layer and consequently the computational cost of Step 2 of Algorithm 1 decreases for increasing values of w .

Algorithm 2 Constrained network layer simulation for \mathbf{y}_{w+1}

Initialise $\mathbf{y}_{w+1}^{(1)} = \mathbf{y}_w$

for $i = 1, \dots, I_{w+1}$ **do**

1) $\mathbf{y}'_{w+1} \sim h(\cdot | \mathbf{y}_{w+1}^{(i)}, \mathbf{y}_w)$

2) Set $\mathbf{y}_{w+1}^{(i+1)} = \mathbf{y}'_{w+1}$ with probability:

$$\min \left\{ 1, \frac{p(\mathbf{y}'_{w+1} | \phi_{w+1}, \mathbf{y}_w) h(\mathbf{y}'_{w+1} | \mathbf{y}_{w+1}^{(i)}, \mathbf{y}_w)}{p(\mathbf{y}_{w+1}^{(i)} | \phi_{w+1}, \mathbf{y}_w) h(\mathbf{y}_{w+1}^{(i)} | \mathbf{y}'_{w+1}, \mathbf{y}_w)} \right\}$$

end for

In this paper, we adopt the ‘tie-no-tie’ sampler for the proposal distribution $h(\cdot)$ which is the default procedure in the `ergm` package [Hunter et al., 2008b, Handcock et al., 2019] for R at Step 1 of Algorithm 2.

6 Simulation study

We propose two simulation studies to assess the performance of the proposed modelling framework and the benefits of the hierarchical Bayesian approach. We consider two 50-node weighted networks with 3 network layers generated by two distinct plausible network processes. We specify the same set of ERGM network statistics for both datasets:

s_1 : Edge statistic (`edges`) is the number of edges in the network and captures the network density effect.

s_2 : Geometrically weighted edgewise shared partner statistic (`gwesp`) is a function of the edgewise shared partner statistics $EP_d(\mathbf{y})$ defined as the number of unordered connected pairs (i, j) (partners) that are both connected to exactly d other nodes: $\sum_d g_d(\alpha) EP_d(\mathbf{y})$, where $g_d(\alpha)$ is an exponential weight function with decay parameter α defined as:

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

This statistic captures the tendency towards transitivity, i.e., the tendency of edges to be connected through multiple triadic relations simultaneously.

As mentioned in Section 2.1 estimating the decay parameters α of the geometrically weighted network statistics from a single network is in general challenging. For this reason, we fix $\alpha = \log(2)$ for all the geometrically weighted network statistics used in this paper.

The between layer parameters ϕ of the two examples are respectively generated from the following prior distributions:

- (a) $\phi|\mu, \Sigma \sim \mathcal{N}\left(\mu = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}\right)$, meaning that both density (μ_1) and transitivity (μ_2) overall trends are null.
- (b) $\phi|\mu, \Sigma \sim \mathcal{N}\left(\mu = \begin{pmatrix} -2 \\ 0.5 \end{pmatrix}, \Sigma = \begin{pmatrix} 2 & 0 \\ 0 & 0.5 \end{pmatrix}\right)$, corresponding to negative density and positive transitivity overall trends.




Figure 3: Estimated posterior distribution for μ for each simulated network.

For both the simulated datasets we used the adaptive direction sampling strategy based on 4 MCMC chains consisting of 10,000 iterations each. We tuned the algorithm parameters in order to get about 20% acceptance rate. In terms of hyper-parameters for the hierarchical model, we set $\kappa_0 = 1$, $\mathbf{\Lambda}_0 = I_r$, $\nu_0 = r + 2$, where I_r is the $r \times r$ identity matrix ($r = 2$ in this case). An analogous hyper-prior specification will be used for the applications in Section 7.

As shown in Figure 3, the estimated posterior modes correspond to the true values of μ for both examples. The estimated predictive posterior densities (displayed in Figure 4) are in agreement with the simulated parameters ϕ for both the experiments. Moreover, it is possible to observe how the hierarchical multilayer model allows for the accounting of the structural between-layer variation and dependencies. In particular, in simulation study (a) even though the overall trends are null ($\mu_1 = \mu_2 = 0$) we can see that between-layer processes are heterogeneous and in the conditional process between network layer 2 and 3 both density and transitivity effects are negative in contrast to what happens in the lower between-layer processes where the transitivity effects are positive and the density effects are null.

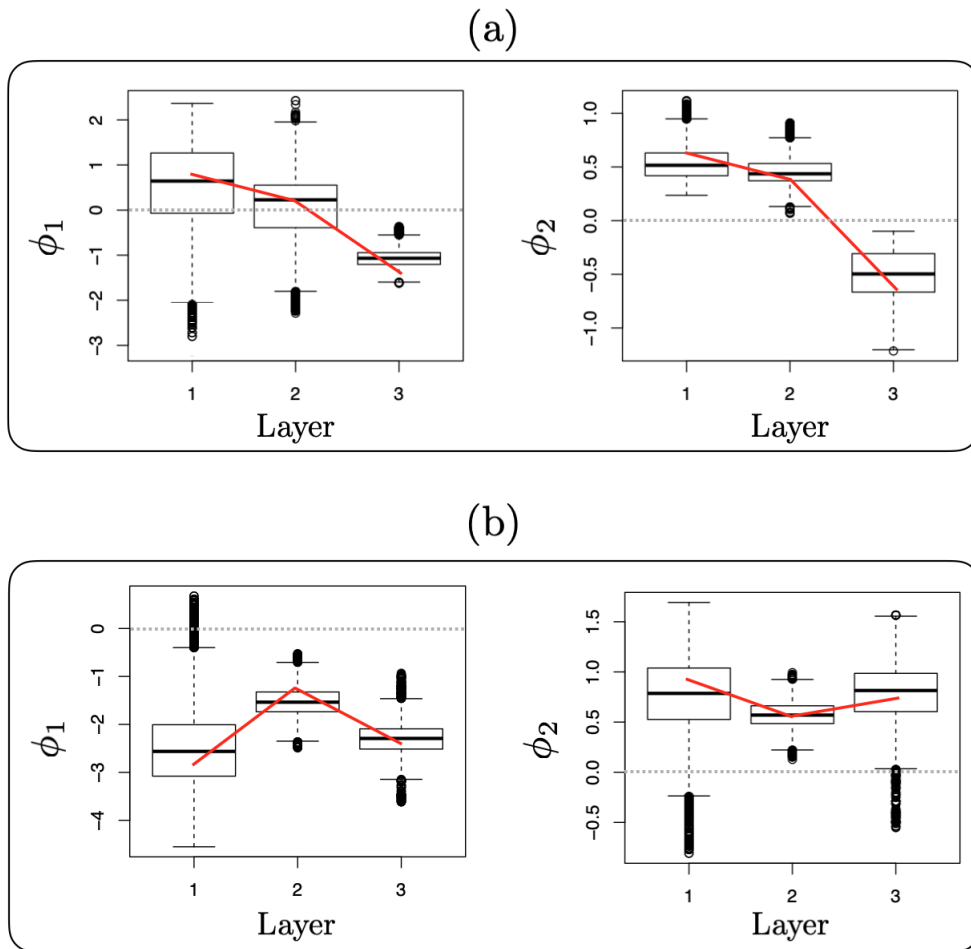


Figure 4: The solid red line represents the simulated values of ϕ_1 (edges) and ϕ_2 (gwesp) for simulation study (a) and (b). The boxplots represent the estimated predictive posterior for ϕ for every network layer.

7 Applications

In this section we present two examples of application of the hierarchical multilayer modelling framework defined in Section 5.1. We will be considering two well-known datasets: the Bernard and Killworth office network [Killworth and Bernard, 1977] and the Zachary karate club network [Zachary, 1977]. We will focus on a polytomous transformation of weighted networks by thresholding the weighted edges at 3 different values so as to obtain 3-layer network structures where the strength of the dyadic relations can be interpreted as either low, medium, or high. This procedure has been carried out by arbitrarily defining threshold values, based on the quantiles of the edge weight distribution.

7.1 Bernard and Killworth office network

The Bernard and Killworth office network dataset concerns the observed frequency of interactions between $N = 40$ individuals in a small business office.




Figure 5: Weighted structure of the Bernard and Killworth office network derived by thresholding the original weights at 2, 4, and 8. The width of the edge lines is proportional to the edge weight.

The network graph and the multilayer visualisation of the Bernard and Killworth office dataset obtained by selecting 3 layers of the weighted graph are shown in Figure 5 and Figure 6 respectively.

7.1.1 Model specification

As mentioned in Section 4.2, an important aspect of the multilayer modelling approach is that it can accommodate binary network statistics [Snijders et al., 2006] to describe the weighted network topology. For the Bernard and Killworth office dataset we propose the following model specification:




Figure 6: The multilayer structure of the Bernard and Killworth office network. The number of edges in each network layer is: $E_1 = 123, E_2 = 44, E_3 = 15$.

s_1 : Edge statistic (**edges**).

s_2 : Geometrically weighted degree statistic (**gwdegree**): $\sum_d g_d(\alpha) D_d(\mathbf{y})$, where $g_d(\alpha)$ is the exponential weight function with decay parameter α . This statistic captures the tendency towards centralisation in the degree distribution $D_d(\mathbf{y})$ of the network.

s_3 : Geometrically weighted edgewise shared partner statistic (**gwesp**).

7.1.2 Posterior analysis

We used the adaptive direction sampling strategy based on 6 MCMC chains consisting of 10,000 iterations each for improving the mixing of the approximate exchange. Traceplots for the model parameters can be found in the Appendix.

The posterior distribution for $\boldsymbol{\mu}$, displayed in Figure 7 and summarised in Table 2, shows the general tendencies of the three local effects across the network layers. A large part of the probability density of μ_1 (corresponding to the **edges** statistic) is concentrated on negative values. This is compensated by a large part of the probability density of μ_3 (corresponding to the **gwesp** statistic) mostly concentrated on positive values. These two tendencies explain the overall increasing sparsity and transitivity of the multilayer process across the network layers.

Table 2: Posterior estimates for $\boldsymbol{\mu}$ for the Bernard and Killworth office network.

Parameter (Effect)	Mean	SD
μ_1 (edges)	-1.10	1.00
μ_2 (gwdegree)	0.14	0.92
μ_3 (gwesp)	0.38	0.78

The predictive posterior for ϕ_1, ϕ_2, ϕ_3 , displayed in Figure 8 and summarised in Table 3, allows us to understand the between-layer ERGM processes. In particular we can notice that most of the density dissolution process described above is concerning the generation of the first network layer as the 99% credible interval of ϕ_1 (corresponding to the **edges** statistic) falls completely on negative values. On the other hand the increase of transitivity captured




Figure 7: Posterior distribution of μ_1 (edges), μ_2 (gwdegree), and μ_3 (gwesp) for the Bernard and Killworth office network.

by ϕ_3 (corresponding to the gwesp statistic) is observed between each layer meaning that the dissolution ERGM process across layers affects edges that are not embedded into transitive structures. The tendency towards centralisation in the degree distribution represented by μ_2 and ϕ_2 (corresponding to the gwdegree statistic) does not seem to be important in explaining both the overall and the between-layer weighted structure of the network meaning that the strong edges are not necessarily centralised or dispersed in the degree distribution.




Figure 8: Predictive posterior of ϕ_1 (edges), ϕ_2 (gwdegree), ϕ_3 (gwesp) for the Bernard and Killworth office network.

7.2 Zachary karate club network

The Zachary karate club network concerns social relations in a university karate club involving 34 individuals. The network graph in Figure 9 shows the relative strength of the associations, i.e., the number of situations in and outside the club in which interactions occurred between individuals.

A multilayer visualisation of the karate club network obtained by selecting 3 layers of the weighted graph is shown in Figure 10.

Table 3: Predictive posterior estimates of ϕ_1, ϕ_2, ϕ_3 for the Bernard and Killworth office network.

Parameter (Effect)	Layer 1		Layer 2		Layer 3	
	Mean	SD	Mean	SD	Mean	SD
ϕ_1 (edges)	-2.94	0.34	-0.53	0.47	-0.94	0.85
ϕ_2 (gwdegree)	1.17	0.96	-0.56	0.48	-0.08	0.62
ϕ_3 (gwesp)	0.56	0.13	0.31	0.18	0.70	0.41




Figure 9: Weighted graph structure of the Zachary karate club network derived by thresholding the original weights at 1, 3, and 4. The width of the edge lines is proportional to the edge weight.




Figure 10: The multilayer structure of the Zachary karate club weighted network. The number of edges in each layer is: $E_1 = 78, E_2 = 48, E_3 = 21$.

7.2.1 Model specification

We included in the model the `edges` (s_1) and `gwesp` (s_2) statistics as in the previous example plus the following statistics:

s_3 : Geometrically weighted non-edgewise shared partner statistic (`gwensp`) is a function of the non-edgewise shared partner statistics $NP_d(\mathbf{y})$ defined as the number of unordered unconnected pairs (i, j) to exactly d other nodes: $\sum_d g_d(\alpha) NP_d(\mathbf{y})$. This statistic captures the tendency of non-directly-connected nodes to be connected through multiple others.

s_4 : Homophily statistic (`nodematch`) is the number of edges between actors having the same nodal attribute X . For this example, X is the faction alignment of the club members so that this statistic captures the density of edges between nodes within the same faction.

7.2.2 Posterior analysis

We used the adaptive direction sampling strategy based on 8 MCMC chains consisting of 10,000 iterations each. In terms of hyper-parameters for the hierarchical model, we used the same set-up of the previous application.

The posterior distribution for $\boldsymbol{\mu}$, displayed in Figure 11 and summarised in Table 4, shows the general tendencies of the 4 local effects across the 3 network layers. A large part of the probability density of μ_1 (corresponding to the `edges` statistic) is concentrated on negative values. This is compensated by a large part of the probability density of μ_2 (corresponding to the `gwesp` statistic) and μ_4 (corresponding to the `nodematch` statistic) mostly concentrated on positive values. The probability density of μ_3 (corresponding to the `gwensp` statistic) is concentrated on values around 0. These tendencies suggest that the overall dissolution process mainly concerns edges that do not connect nodes within factions and/or are not embedded in transitive triads.

Table 4: Posterior estimates of $\boldsymbol{\mu}$ for the Zachary karate club network.

Parameter (Effect)	Mean	SD
μ_1 (<code>edges</code>)	-1.20	1.16
μ_2 (<code>gwesp</code>)	0.40	0.77
μ_3 (<code>gwensp</code>)	0.02	0.76
μ_4 (<code>nodematch</code>)	0.51	0.82

More specifically the estimated predictive posterior for $\phi_1, \phi_2, \phi_3, \phi_4$, displayed in Figure 12 and summarised in Table 5, indicates that most of the density dissolution process described above and represented by the parameter ϕ_1 (corresponding to the `edges` statistic) is stronger in the first and between the second and the third network layer. We can also observe the importance of the transitivity effect captured by ϕ_2 (corresponding to the `gwesp` statistic) between all the layers meaning that the dissolution ERGM process across layers affects mainly edges that are not embedded into triadic transitive structures. We can therefore realise that the




Figure 11: Posterior distribution of μ_1 (edges), μ_2 (gwesp), μ_3 (gwmsp), and μ_4 (nodematch) for the Zachary karate club network.

strongest edges, i.e., the ones surviving the multilayer ERGM dissolution process, are mainly the ones involved in transitive triads. It is interesting to notice how the parameter ϕ_3 (corresponding to the gwmsp statistic) contributes positively to the generation of the first layer and negatively to the generation of the second layer conditional on the first layer. This can be explained by the fact that connectivity structure of layer 1 is very much influenced by the presence of open transitive triads involving the two competing leaders of the karate club factions (John and Mr. Hi) who are not directly connected but share many partners among the members of the club. However when we consider the transition from layer 1 to layer 2 many connections involved in these open triads dissolve meaning that they are mostly weak edges. The faction density representing homophily between members of the same faction, captured by ϕ_4 (corresponding to the nodematch statistic), is important in explaining the generation of the first layer but not for explaining the dissolution process between the layers.

7.3 Model assessment

A way to examine the fit of the data to the estimated posterior distribution of the parameters is to implement a graphical Bayesian goodness-of-fit procedure [Hunter et al., 2008a]. In the Bayesian context, network data are simulated from a sample of parameter values drawn from the estimated posterior distribution and compared to the observed data in terms of high-level network characteristics that are not explicitly included as sufficient statistics in the model. Since




Figure 12: Predictive posterior of ϕ_1 (edges), ϕ_2 (gwesp), ϕ_3 (gwnsp), ϕ_4 (nodematch) for the Zachary karate club network.

Table 5: Predictive posterior estimates of $\phi_1, \phi_2, \phi_3, \phi_4$ for the Zachary karate club network.

Parameter (Effect)	Layer 1		Layer 2		Layer 3	
	Mean	SD	Mean	SD	Mean	SD
ϕ_1 (edges)	-3.96	0.27	0.27	0.68	-1.09	0.63
ϕ_2 (gwesp)	0.53	0.14	0.53	0.24	0.51	0.25
ϕ_3 (gwnsp)	0.17	0.02	-0.18	0.09	0.06	0.17
ϕ_4 (nodematch)	1.29	0.27	0.56	0.43	0.21	0.52

we are dealing with weighted networks, we focus on the weighted degree distribution. The black solid lines represent the distribution of the weighted degrees in the observed data, the grey lines represent the distribution of the weighted degrees calculated on network graphs simulated from the estimated posterior density.

The plots in Figure 13 suggest that both models are a reasonable fit to their respective datasets as the black lines representing the observed weighted degree distributions are lying on the high posterior predictive distribution despite the absence of a degree-based statistic in the Zachary karate club model (this absence is only partially compensated by the presence of the open triadic effect corresponding to the `gwnsp` statistic.)

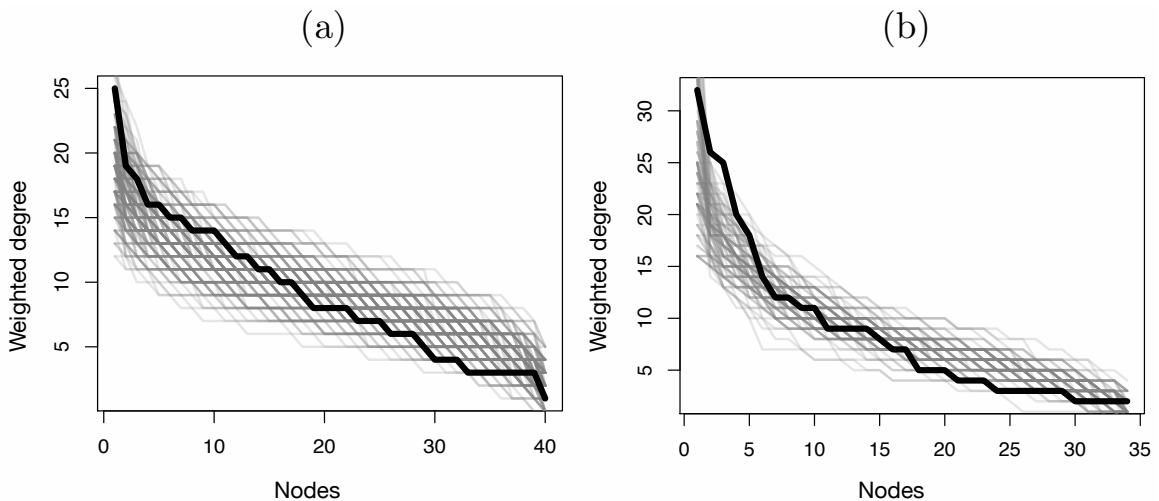


Figure 13: Weighted degrees for each node in (a) Bernard and Killworth office network and (b) Zachary karate club network. The black line and the gray lines represent the weighted degree values for the observed network and simulated networks, respectively.

8 Conclusions

This paper has introduced a new Bayesian hierarchical ERGM framework for the analysis of weighted networks with ordinal/polytomous edges which complements the recent advances proposed by Wyatt et al. [2010], Krivitsky [2012], Desmarais and Cranmer [2012], Krivitsky and Butts [2017], Wilson et al. [2017]. The modelling approach is based on a flexible multilayer ERGM process able to describe the ERGM dissolution process which leads to the generation of the strength of the network weighted edges. The multilayer ERGM process is parametrised using binary network statistics and it is therefore providing a natural interpretation of the network effects that are assumed to be at the basis of the generative process.

A fully-probabilistic Bayesian approach has been adopted to provide the possibility of specifying prior information of the network effects and analysing their posterior distribution given the observed data. An extension of the approximate exchange algorithm [Caimo and Friel, 2011] has been used in order to carry out inference on the doubly-intractable posterior distribution.

Probabilistic goodness-of-fit diagnostics based on the weighted degree distribution have been proposed to assess the models used in two applications on well-known datasets.

The modelling framework proposed in this paper can be extended to deal with weighted signed networks by, for example, making a distinction between a network process generating the baseline interaction layer encoding the presence or absence of any signed relation between nodes and a joint conditional multilayer ERGM process describing the formation of positive and negative weighted relations between the interacting nodes.

Moreover, the affinity between separable temporal ERGMs and multilayer ERGMs can be exploited to propose a joint modelling approach for temporal weighted networks.

References

- P. Alquier, N. Friel, R. Everitt, and A. Boland. Noisy Monte Carlo: Convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, 26(1-2):29–47, 2016.
- A. Barrat, M. Barthélemy, R. Pastor-Satorras, and A. Vespignani. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences*, 101(11):3747–3752, 2004.
- J. E. Besag. Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society, Series B*, 36:192–236, 1974.
- R. P. Bomiriyá, S. Bansal, and D. R. Hunter. Modeling homophily in ERGMs for bipartite networks. In *International Conference on Robust Statistics*, 2016.
- L. Bouranis, N. Friel, and F. Maire. Efficient Bayesian inference for exponential random graph models by correcting the pseudo-posterior distribution. *Social Networks*, 50:98–108, 2017.
- L. Bouranis, N. Friel, and F. Maire. Bayesian model selection for exponential random graph models via adjusted pseudolikelihoods. *Journal of Computational and Graphical Statistics*, pages 1–13, 2018.
- E. Bullmore and O. Sporns. Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186, 2009.
- A. Caimo and N. Friel. Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41 – 55, 2011.
- A. Caimo and N. Friel. Bergm: Bayesian exponential random graphs in R. *Journal of Statistical Software*, 61(2):1–25, 2014. URL <http://www.jstatsoft.org/v61/i02/>.
- A. Caimo and A. Mira. Efficient computational strategies for doubly intractable problems with applications to Bayesian social networks. *Statistics and Computing*, 25:113–125, 2015.

- A. Caimo, F. Pallotti, and A. Lomi. Bayesian exponential random graph modelling of interhospital patient referral networks. *Statistics in Medicine*, 36(18):2902–2920, 2017.
- A. Caimo, L. Bouranis, R. Krause, and N. Friel. *Bergm: Bayesian Exponential Random Graph Models.*, 2019. URL <https://CRAN.R-project.org/package=Bergm>. R package version 5.0.0.
- B. A. Desmarais and S. J. Cranmer. Statistical inference for valued-edge networks: The generalized exponential random graph model. *PLOS ONE*, 7(1):e30136, 2012.
- R. G. Everitt. Bayesian parameter estimation for latent Markov random fields and social networks. *Journal of Computational and Graphical Statistics*, 21(4):940–960, 2012.
- D. Garlaschelli. The weighted random graph model. *New Journal of Physics*, 11(7):073005, 2009.
- M. S. Handcock. Assessing degeneracy in statistical models of social networks. Technical report, Working Paper 39, Center for Statistics and the Social Sciences, University of Washington, 2003.
- M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, P. N. Krivitsky, and M. Morris. *ergm: Fit, Simulate and Diagnose Exponential-Family Models for Networks.* The Statnet Project (<https://statnet.org>), 2019. URL <https://CRAN.R-project.org/package=ergm>. R package version 3.10.4.
- P. W. Holland and S. Leinhardt. An exponential family of probability distributions for directed graphs (with discussion). *Journal of the American Statistical Association*, 76:33–65, 1981.
- S. Horvath. *Weighted network analysis: Applications in genomics and systems biology.* Springer Science & Business Media, 2011.
- D. R. Hunter. Curved exponential family models for social networks. *Social Networks*, 29(2):216–230, 2007.
- D. R. Hunter, S. M. Goodreau, and M. S. Handcock. Goodness of fit of social network models. *Journal of the American Statistical Association*, 103(481):248–258, 2008a.
- D. R. Hunter, M. S. Handcock, C. T. Butts, S. M. Goodreau, and M. Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3):1–29, 2008b. URL <http://www.jstatsoft.org/v24/i03>.
- P. D. Killworth and H. Bernard. Informant accuracy in social network data II. *Human Communication Research*, 4(1):3–18, 1977.
- J. Koskinen, A. Caimo, and A. Lomi. Simultaneous modeling of initial conditions and time heterogeneity in dynamic networks: An application to foreign direct investments. *Network Science*, 3(1):58–77, 2015.

- P. N. Krivitsky. Exponential-family random graph models for valued networks. *Electronic Journal of Statistics*, 6:1100, 2012.
- P. N. Krivitsky and C. T. Butts. Exponential-family random graph models for rank-order relational data. *Sociological Methodology*, 47(1):68–112, 2017.
- P. N. Krivitsky and M. S. Handcock. A separable model for dynamic networks. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):29–46, 2014.
- I. Murray and Z. Ghahramani. Bayesian learning in undirected graphical models: Approximate MCMC algorithms. In *Uncertainty in Artificial Intelligence (UIA-2004)*, pages 392–399, 2004.
- T. Opsahl, V. Colizza, P. Panzarasa, and J. J. Ramasco. Prominence and control: The weighted rich-club effect. *Physical Review Letters*, 101(16):168702, 2008.
- J. Park and M. Haran. Bayesian inference in the presence of intractable normalizing functions. *Journal of the American Statistical Association*, 113(523):1372–1390, 2018.
- P. Pattison and S. Wasserman. Logit models and logistic regressions for social networks: II. Multivariate relations. *British Journal of Mathematical and Statistical Psychology*, 52(2):169–193, 1999.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019. URL <https://www.R-project.org/>.
- G. Robins, P. Pattison, and S. Wasserman. Logit models and logistic regressions for social networks: III. Valued relations. *Psychometrika*, 64(3):371–394, 1999.
- G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph models for social networks. *Social Networks*, 29(2):169–348, 2007.
- M. Salter-Townshend, A. White, I. Gollini, and T. B. Murphy. Review of statistical network analysis: Models, algorithms, and software. *Statistical Analysis and Data Mining*, 5(4):243–264, 2012.
- M. Schweinberger and P. Luna. hergm: Hierarchical exponential-family random graph models. *Journal of Statistical Software*, 85(1), 2018.
- M. Schweinberger, M. S. Handcock, S. Babkin, J. Stewart, D. Vu, and P. Luna. *hergm: Hierarchical Exponential-Family Random Graph Models.*, 2019. URL <https://CRAN.R-project.org/package=hergm>. R package version 4.1-3.
- A. J. Slaughter and L. M. Koehly. Multilevel models for social networks: Hierarchical bayesian approaches to exponential random graph modeling. *Social Networks*, 44:334–345, 2016.

- T. A. B. Snijders, P. E. Pattison, G. L. Robins, and H. M. S. New specifications for exponential random graph models. *Sociological Methodology*, 36:99–153, 2006.
- J. Stewart, M. Schweinberger, M. Bojanowski, and M. Morris. Multilevel network data facilitate statistical inference for curved ERGMs with geometrically weighted terms. *Social Networks*, 59:98 – 119, 2019. ISSN 0378-8733.
- D. Strauss and M. Ikeda. Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association*, 5:204–212, 1990.
- J. D. Wilson, M. J. Denny, S. Bhamidi, S. J. Cranmer, and B. A. Desmarais. Stochastic weighted graphs: Flexible model specification and simulation. *Social Networks*, 49:37–47, 2017.
- D. Wyatt, T. Choudhury, and J. A. Bilmes. Discovering long range properties of social networks with multi-valued time-inhomogeneous models. In *AAAI*, 2010.
- W. Zachary. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33:452–473, 1977.

9 Appendix




Figure 14: Bernard and Killworth office network; MCMC traces for μ (thinning factor = 100).




Figure 15: Bernard and Killworth office network; MCMC traces for ϕ (thinning factor = 100).

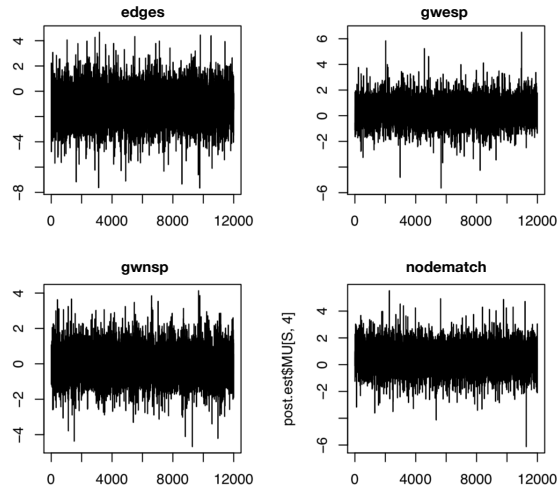


Figure 16: Zachary karate club network; MCMC traces for μ (thinning factor = 100).

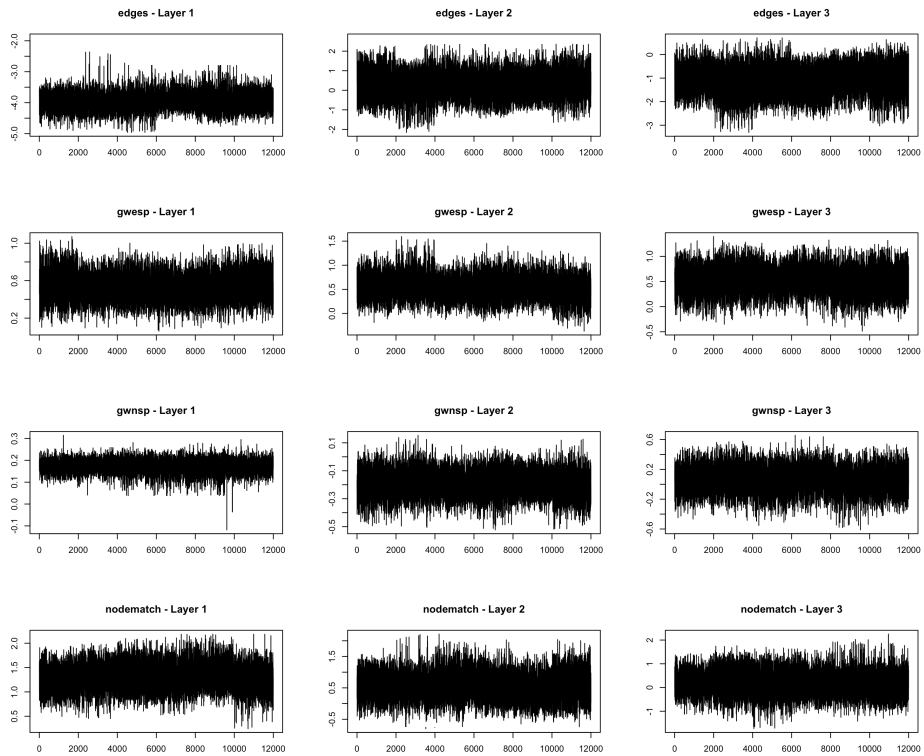


Figure 17: Zachary karate club network; MCMC traces for ϕ (thinning factor = 100).