



GREThA

Groupe de Recherche en
Économie Théorique et Appliquée

**The modelling of networks using
Exponential Random Graph Models: an introduction**

Johannes VAN DER POL

GREThA, CNRS, UMR 5113, Université de Bordeaux

johannes.van-der-pol@u-bordeaux.fr

Cahiers du GREThA

n° 2016-22

Octobre

GREThA UMR CNRS 5113

Université de Bordeaux

Avenue Léon Duguit - 33608 PESSAC - FRANCE

Tel : +33 (0)5.56.84.25.75 - Fax : +33 (0)5.56.84.86.47 - www.gretha.fr

La modélisation des réseaux par les Modèles à Graph Aléatoire Exponentiel: une introduction

Résumé

Un réseau est une représentation de données relationnelles. Que ce soient des données sur des relations sociales, des coopérations voire des dépendances interbancaires, la structure d'un réseau est une représentation des décisions stratégiques prises par les agents. Ces décisions sont basées sur une variété de facteurs (relations sociales communes, proximité technologique, proximité géographique). Un des objectifs de l'analyse des réseaux est d'identifier quels sont ces facteurs qui poussent un agent à interagir avec un autre agent en particulier. L'analyse de la structure avec un ERGM permet d'identifier ces facteurs. En d'autres termes, avec un ERGM on est capable d'effectuer une analyse économétrique sur des données réseaux. Dans un réseau, la présence d'un lien est souvent dépendant de la structure existante (i se connecte avec j car ils ont un collaborateur commun). Les méthodes économétriques sont alors inefficaces car les données réseau violent l'hypothèse d'indépendance des observations. Les ERGM sont capables de gérer ce genre de dépendance.

L'objectif de ce papier est de présenter la théorie statistique derrière les ERGM et de proposer une application avec R-project.

Mots-clés : Analyse réseau ; ERGM ; Structure de réseau ; R

The modelling of networks using Exponential Random Graph Models: an introduction

Abstract

Networks are representations of relational data. Whether the data used represents social interactions, cooperations or inter-bank dependencies, the structure of the network reflect a decision-making process based on many factors (common friends, technological proximity, geographical proximity). One of the objectives of network analysis is to identify these factors. The analysis of the structure using Exponential Random Graph Models (ERGM) can help in the identification of these factors by answering why two particular agents interact with one another, or why a specific agent has a particular position inside a network. In other words ERGMs allow us to perform an econometric analysis on network data. In the case of networks it is possible that a link depends upon the structure of the network. Usual econometric methods cannot be used because the dependence violates the hypothesis of independence of observations. ERGMs take into account this particularity of network data.

The aim of this paper is to present the statistical theory behind ERGM models and present an application using R-Project.

Keywords: Network analysis ; ERGM ; Network structure ; R

JEL: C59;C13;D85

Reference to this paper: VAN DER POL Johannes (2016) The modelling of networks using Exponential Random Graph Models: an introduction, <i>Cahiers du GREThA</i> , n°2016-22.

http://ideas.repec.org/p/grt/wpegrt/2016-22.html

Introduction

Networks are representations of relational data. Nodes represent entities while the links connecting them represent any form of interaction or connection between the entities. A large diversity in the types of networks exists ranging from networks of social contacts between individuals to inventor networks, collaboration networks, financial networks and so on. The trouble with networks, especially when represented by a graph, is that it looks like a large heap of lines, it resembles pure chaos. However, interactions between individuals, firms or banks rarely appear at random. The motivations for link creation cannot be observed directly from a graph, nor are they clear from glimpsing at a database containing relational data. In order to identify the motivations for entities to create links and identify the global network structure, a more in-depth analysis is required. Econometric analysis could shed more light on the motivations behind an observed link through logistic regressions. The probability of a link could be explained by a number of variables. There is one important limitation to this method. Due to the hypothesis of independence of the observations the probability of a link between two nodes can never be explained by the presence of another link inside the network. It is feasible that a link between two nodes exists only because of the presence of other links in the network. Take for instance the idea that John and Mary are connected solely because they have a common contact: Paul. ERGM models are modified logistic regressions that allow for the probability of a link to depend upon the presence of other links inside the network (amongst other variables of course). An ERGM identifies the probability distribution of a network so that it can generate large samples of networks. The samples are then used to infer on the odds of a particular link inside a network. Applications for this method are numerous in many fields of research as shown by the increasing trend in the number of publications using ERGM models (see figure 1). In




Figure 1: Evolution of the number of publications involving ERG models for all disciplines (statistics included) (source: Scopus)

Support from IdEx Bordeaux is gratefully acknowledged.

economics the number of published papers appears to be relatively low when compared to the other social sciences. Only 19 published papers could be found in the Scopus database (and even less in the web of science database). The topics are however quite diverse: Link between money and inflation (Tang et al., 2015; Özdemir, Saygili, 2009; ?; Belke, Polleit, 2006; Price, Nasim, 1998), knowledge sharing in organizations (Caimo, Lomi, 2015), GDP targeting (Belongia, Ireland, 2014), alliance networks (Cranmer et al., 2012; Lomi, Pallotti, 2012; Lomi, Fonti, 2012; Broekel, Hartog, 2013) and geographic proximity (Ter Wal, 2013).

The growing interest, and development of a theory of economic networks, provides a fertile ground for the use of ERGM models from the geography of innovation to venture capital investments. The aim of this paper is to provide an overview of the basic statistical theory behind ERGM models, which will be dealt with in the first section. Section 2 discusses the concept of dependence and the explanatory variables that can be included in the models. Section 3 discusses estimation methods while section 4 provides the R scripts and the interpretation of an example using data for the French aerospace sector alliance network.

1 Theory

1.1 The canonical form of ERGM models




Figure 2: Network G

The aim of an ERGM is to identify the processes that influence link creation. The researcher includes variables in the model that are hypothesized to explain the observed network, the ERGM will provide information relative to the statistical significance of the included variable much like a standard linear regression.

It is useful at this point to explain that sub-structures of a network can (and are predominantly) used as explanatory variables. Substructures are small graphs contained inside the network. Examples can be found in figure 28. The presence of some of these structures reflects certain link creation phenomena. A random network, i.e a network in which links are created at random, show a low number of triangles. A triangle is an interconnection of three nodes, the smallest possible complete subgraph. The presence of triangles in an empiric network bares witness that there is a process that generates triangles that is not the result of random link creation e.g, a tendency to create a link between common friends. A network with a small number of large stars and a large number of small stars can be the results of having a small number of very popular nodes. This is found in citation networks as well as lexicographical networks. Including sub-structures allows the modeling of certain processes as would any other variable. In an ERGM we can find two types of explanatory variables: structural and node or

edge-level variables. The latter come from other data sources and can be for example age, size of a firm, proximity, gender and so forth. The structural variables can contain indicators such as triadic closure, degree distribution and subgraphs.

1.2 The odds of a link

With this in mind the probability that one would observe a link between any two nodes i and j in a given network is proportional to a given set of explanatory variables:

$$p(G_{ij} = 1) = \theta_1 \cdot X_1 + \theta_2 \cdot X_2 + \dots + \theta_n \cdot X_n \quad (1)$$

We note G a graph and ij the focal nodes. $G_{ij} = 1$ means that a link exists between nodes i and j in graph G , $G_{ij} = 0$ implies the absence of a link between nodes i and j ¹. θ is a vector of parameters and X a vector containing the variables.

Equation 1 gives us the probability of a *single link* in graph G . Since nothing guarantees that the probability stays within $[0, 1]$ the equation needs to be somewhat adjusted. We start by transforming the probability into an odds ratio:

$$odds(G_{ij} = 1) = \frac{p(G_{ij} = 1)}{1 - p(G_{ij} = 1)} = \frac{p(G_{ij} = 1)}{p(G_{ij} = 0)} \quad (2)$$

Now, in equation 2 we notice that the odds of a link tend towards zero when the probability of a link tends towards one, while tending towards $-\infty$ when the probability tends towards zero. A final modification is required to ensure that the probability remains between the bounds of 0 and 1, this is accomplished using the natural logarithm:

$$\log(odds(G_{ij} = 1)) = \text{logit}(p(G_{ij} = 1)) = \log\left(\frac{p(G_{ij} = 1)}{p(G_{ij} = 0)}\right) \quad (3)$$

The probability is now bounded by 0 and 1. If we suppose that the probability of a link is explained by a vector of n variables accompanied by their respective parameters $(\theta_1 \dots \theta_n)$ then we can write:

$$\text{logit}(p(G_{ij} = 1)) = \theta \cdot X \quad (4)$$

So in our example we have:

$$\text{logit}(p(G_{ij} = 1)) = \theta_1 \cdot X_1 + \dots \theta_n \cdot X_n \quad (5)$$

$$p(G_{ij} = 1) = \exp\{\theta_1 \cdot X_1 + \dots \theta_n \cdot X_n\} \quad (6)$$

This gives us the logit which supplies the marginal probability of a tie between nodes i and j . Using these equations the probability of a tie between i and j would

¹the values 0 and 1 refer to values found in an adjacency matrix, 1 indicating the presence of a link, 0 the absence

be independent from the probability between i and k . Indeed, a logistic regression works under the hypothesis of independence of observations. In the case of networks observations are not independent. For instance, common friends tend to connect more in social networks, common collaborators have a higher tendency towards collaboration. A model that aims at explaining a network structure should be able to include these tie formation processes.

We hence modify the initial equations to include the network structure as observed before the link. This modification is introduced by Strauss, Ikeda (1990). We note G_{ij}^c the network without link ij :

$$\text{odds}(G_{ij} = 1) = \frac{p(G_{ij} = 1|G_{ij}^c)}{1 - p(G_{ij} = 1|G_{ij}^c)} = \frac{p(G_{ij} = 1|G_{ij}^c)}{p(G_{ij} = 0|G_{ij}^c)} \quad (7)$$

In equation 7 the odds of a link between nodes i and j now depends on the structure of the network before a link between i and j is created (noted by $|G_{ij}^c$). The probabilities are now conditional.

We discussed previously that some of the variables in the model can be subgraphs. The manner in which these are included in the model is simply by the count of these sub-structures. In other words, the value of the variable triangles is the number of triangles in the network. The same is true for stars, circuits and shared partners. This has as a consequence that the counts of these variables are not the same when a link between two nodes is present or absent. For instance the number of edges changes by one. This means that we need to differentiate between the value of the variables when a link is present and when it is absent. We hence note the vector of variables $v(G_{ij}^+)$ when a link between i and j is added (hence the " + ") and $v(G_{ij}^-)$ when the link is absent. By including this differentiation we can rewrite equation 7 using the result in equation 6:

$$\text{odds}(G_{ij} = 1) = \frac{p(G_{ij} = 1|G_{ij}^c)}{p(G_{ij} = 0|G_{ij}^c)} = \frac{\exp\{\theta' \cdot v(G_{ij}^+)\}}{\exp\{\theta' \cdot v(G_{ij}^-)\}} \quad (8)$$

Where $v(G_{ij}^+)$ represents the vector of variables in the network with the link between i and j present and $v(G_{ij}^-)$ the vector of variables with no link between i and j . With some basic algebra we can develop the previous equation a bit further:

$$\frac{\exp\{\theta' \cdot v(G_{ij}^+)\}}{\exp\{\theta' \cdot v(G_{ij}^-)\}} = \exp\{\theta' \cdot v(G_{ij}^+)\} \cdot \exp\{-\theta' \cdot v(G_{ij}^-)\} \quad (9)$$

$$= \exp\{\theta' (v(G_{ij}^+) - v(G_{ij}^-))\} \quad (10)$$

When developing the vector of variables we have:

$$= \exp\{\theta'_1 \cdot (v_1(G_{ij}^+) - v_1(G_{ij}^-)) + \dots + \theta'_n \cdot (v_n(G_{ij}^+) - v_n(G_{ij}^-))\} \quad (11)$$

Equation 11 shows that each parameter of the model is associated not with the counts of sub-structure but with the difference in counts. The difference from having an extra link, and the absence of said link. In essence $(v_1(G_{ij}^+) - v_1(G_{ij}^-))$ represents the variation in the number of counts of network statistic 1 that result from the additional link. The variables are hence referred to as "change statistics". In order to remove the exponential from the right hand side of the equation we apply the logarithm:

$$\log\left(\frac{\exp\{\theta \cdot v(G_{ij}^+)\}}{\exp\{\theta \cdot v(G_{ij}^-)\}}\right) = \theta'_1 \cdot (v_1(G_{ij}^+) - v_1(G_{ij}^-)) + \dots + \theta'_n \cdot (v_n(G_{ij}^+) - v_n(G_{ij}^-)) \quad (12)$$

So we can rewrite equation 12 noting $v_1(\Delta_1 G_{ij})$ the change statistic for a link between i and j for variable 1 as follows:

$$\log\left(\frac{\exp\{\theta \cdot v(G_{ij}^+)\}}{\exp\{\theta \cdot v(G_{ij}^-)\}}\right) = \theta'_1 \cdot v_1(\Delta_1 G_{ij}) + \dots + \theta'_n \cdot v_n(\Delta_n G_{ij}) \quad (13)$$

Each variable now accounts for the change in counts of network statistics. It is important to remind us that equation 13 accounts for the odds of one edge in the network while we are interested in the probability for the whole network. Following (Besag, 1972) we can invoke here the Hammersley-Clifford theorem. Since this theorem is based on concepts out of the reach and the purpose of this document we will not detail the theorem. For a detailed explanation please refer to (Hammersley, Clifford, 1971). The theorem states that the probability of a network can be defined solely by the counts of subgraphs. This is important because it tells us that all we have to do is identify the correct subgraphs to ensure that a model of the network structure can be found. The more accurate the subgraphs to more reliable the inference of additional covariates.

1.3 The probability distribution of a network

The Hammersley-Clifford theorem states that the probability of a graph can be identified solely by counts of subgraphs. As such, we know that the probability is proportional to these variables. Since we have an observed network that we wish to replicate we look for the probability that the network generated by the model (X) is identical to the observed network (x). The logarithm is applied to bound the probability:

$$\log(p(X = x)) \propto \theta \cdot v(G) \quad (14)$$

$$p(X = x) \propto \exp\{\theta \cdot v(G)\} \quad (15)$$

The right-hand side of the equation now needs to be normalized in order to obtain a proper probability distribution. The normalization is not straightforward, indeed, in order to normalize the probability of a network one needs to normalize by all possible networks with the same number of nodes:

$$p(X = x) = \frac{\exp\{\theta \cdot v(G)\}}{\sum_{y \in Y} \exp\{\theta \cdot v(G)\}} \quad (16)$$

With y a possible network structure in the set of all possible networks Y . The numerator is normalized by the sum of the parameters over all possible network structures. Note that this number is large. For a network with n nodes the number of possible graphs is $2^{\frac{n(n-1)}{2}}$. So even for a graph with 10 nodes there are 35184372088832 possible graphs. The major problem to overcome with ERGMs is exactly this normalizing constant.

With some simple algebra we find a general form for this model. Using θ as a vector of parameters and $v(G)$ a vector of variables for network G :

$$p(X = x) = \frac{\exp\{\theta \cdot v(G)\}}{\exp\{\log(\sum_{y \in Y} \exp\{\theta \cdot v(G)\})\}} \quad (17)$$

$$p(X = x) = \frac{1}{\psi(\theta)} \cdot \exp\{\theta \cdot v(G)\} = \exp\{\theta \cdot v(G) - \psi(\theta)\} \quad (18)$$

Equation 18 is the most general and commonly used form of the model (Lusher et al., 2012). Equation 18 also gives the canonical form of an ERGM model. Since the density of the random variable (the network structure) has the particular form in equation 18 it is referred to as an exponential family. In addition, since the structures are represented by a random variable, they are random graphs.

Putting both elements together and this results in an Exponential Family Random Graph. Since ERGM is easier to pronounce than EFRGM, the models are referred to as ERGM².

The canonical form gives the equation we wish to estimate. However, before we tackle the question of estimation we need to explore in more detail the variables that we would want to include in the model. We have stated previously that counts of subgraphs can be used as explanatory variables. The following section will explain which particular subgraphs are to be included in a model.

2 The dependence assumption

The previous section has shown that ERGMs are capable of providing conditional probabilities for links. This dependence assumption is important because it allows the researcher to study different phenomena that rule the formation of networks. This section will show how the hypothesis of dependence of links is connected to the choice of subgraphs that may be included in an analysis.

2.1 The Markovian hypothesis

Links between nodes rarely appear at random, agents have specific reasons to connect with one agent rather than another. The motivations behind interactions are numerous

²This type of model is also referred to as the P^* family of models (Anderson et al., 1999; Lusher et al., 2012).

and complex and have been subject to scrutiny from researchers in different strands of the social sciences. This research shows that people or firms with common acquaintances or collaborators will have a higher tendency to cooperate for example. This makes sense for two reasons, first, having common partners means common working practices which have a positive impact on collaboration. Second, when searching for collaborations firms tend to rely on referrals. A collaboration already in place allows firms to observe in detail the efficiency of other firms, referrals that result from cooperation should hence be trustworthy. In addition cooperators of a firms have a higher probability to be in contact with each other since they are more likely to meet during social or professional events.

The odds of a link depend on the neighborhood of the node and not on the entire rest of the graph. In more formal terms: two potential links X_{ij} and X_{kl} are in the same neighborhood if $\{i, j\} \cap \{k, l\} \neq \emptyset$. For instance, the dotted lines in Figure 3 represent potential links. The odds of this link will depend upon the nodes f and d have in common, in this case node a . This shows that not all subgraphs are compatible with analyzing this phenomenon. Any substructure that cannot account for a common node should not be included.




Figure 3: Node level dependence illustration: the Markovian neighborhood

Since it is hypothesized that only neighboring nodes impact the odds of a link, we seek node-level dependence. This level of dependence is also referred to as a nearest-neighbors level of dependence³ or dyadic dependence (Harris, 2013).

For the purpose of the identification of relevant structures we need to find all structures that are subject to a Markov (or nearest-neighbor) level of dependence.

Suppose we have the social network depicted in figure 4. The graph shows social interactions between four agents, a, b, c and d . Markovian dependence suggests that a link between a and b depends on the connections between common nodes. The nodes a and b have in common are d and c . A link between a and b hence depends upon connections between $a - c, a - d, b - c$ and $b - d$.

When one identifies all the possible dependencies one can generate a dependence graph. The dependence graph for the complete Markov graph in our example can be found in figure 6. In red we find the dependence links for link $a - b$, it show the links on which $a - b$ depends.

³Nearest-neighbor systems have been studied by (Besag, 1972).




Figure 4: Markov graph




Figure 5: Dependence graph




Figure 6: Dependence graph and configuration identification




Figure 7: 2-star identification in the dependence graph

From this graph one can identify the substructures that comply with Markovian dependence. All subgraphs in the dependence graph can be included in a model to add Markovian dependence. For example, the dotted line between ab and ad represents a 2-star centered on agent a (figure 7).

Using the same method as for the 2-star one can also identify a triangle between the three agents on the links bd , bc and dc . The Markov model hence includes three configurations: edges, 2-stars and triangles. With the inclusion of these configurations the Markovian model takes the form:

In more complex graphs one could also identify 3-stars, 4-stars etc.

It should be obvious here that the number of distinct 2-stars is large and it is near impossible to add a parameter for each distinct 2-star in the dependence graph. To reduce the number of variables a hypothesis is made that each type of configuration has the same probability of appearance, this allows for the inclusion of one parameter per substructure. In the case of Markov dependence the ERGM model would have the following form:

$$p(x = X|\theta) = \frac{1}{\psi(\theta)} \exp\{\theta_E \cdot v_E(x) + \theta_{S_2} \cdot v_{S_2} + \dots + \theta_{S_{n-1}} \cdot v_{S_{n-1}} + \theta_{\Delta} \cdot v_{\Delta}\} \quad (19)$$

Where θ_E is the parameter for the number of edges, θ_{S_2} the parameter for the number of 2-stars and θ_{Δ} the parameter for the number of triangles. Note here that the model does not include simultaneously a 1-star and an edge parameter since they would be the same variable. With this model one is able to study if common nodes have a positive impact on the odds of link creation.

In addition the combination of the 2-star parameter and the triangle parameter account for triadic closure effects. In other words, are triangles created because 3 nodes are connected at the same time, or are triangles formed by the closing of a 2-star.

Of course, Markovian dependence is only one of the possible levels of dependence. One can imagine higher levels of dependence, or even any level of dependence to be empirically relevant. One could suggest that firms evolving on the periphery of a network to have a higher probability to connect with firms in the center of the network than between them. The previous model was hence extended by Wasserman, Pattison (1996) to allow for a general level of conditional dependence giving the researcher a total liberty in the theories to test. Whatever the level of dependence chosen, the dependence graph gives the substructure that may be included (Frank, Strauss, 1986).

2.1.1 Higher levels of dependence

It is possible to assume that the Markovian level of dependence is not adequate or does not capture the full complexity of mechanisms of link creation. Links can be dependent without there being a common node involved. For instance, consider a case in which people work on the same floor in a company. The probability of a social link does not depend upon a common node but simply on the fact that they are geographically close, belong to the same community or have common cultural aspects (White, 1992). In order to be able to model more complex aspects of social interactions and indeed even strategic interactions, one needs to be able to account for more structural aspects than stars and triangles (however potent in explanatory power these might be). The latter implies that the links are only dependent on each-other if nodes are part of a same neighborhood (neighborhood takes a broad definition here, it can be social, geographical or cultural). Due to the inclusion of general dependence the model is transformed to take the form:

$$p(x = X) = \frac{1}{\psi(\theta)} \exp\left\{ \sum_{A \in M} \lambda_A \cdot z_A(x) \right\} \quad (20)$$

Where A indicates the neighborhood as part of the ensemble M of all possible neighborhoods. The parameter λ_A will take a positive value when the probability of observing

network x is increased. With the broad definition of "neighborhood" this model is able is almost limitless. The latter results in a problem, the model is too general.

We have seen in the previous subsection that the structures that can included in the ERGM model are defined by the dependency graph. In the case of a generalization of the dependency assumption, i.e all ties may be dependent upon all other ties, the dependency graph is a complete graph and all possible subgraphs can be considered as variables. This leaves a tremendous amount of parameters to be estimated.

Pattison, Robins (2002) and Carrington et al. (2005) offer two solutions to this problem. Their aim is to find a way to reduce the number of subgraphs to be included in the model. The only way to achieve this is to reduce the level of dependency from general to a more restricted level. A first step is to simply fix a level of dependency which will automatically switch all other parameters to 0. This means that once one defines a condition under which links are dependent upon each other a *setting* is defined. Defining a level of dependency can be simply supporting the hypothesis that links between firms depend upon a common region, or sector, size or any other group. s is a setting, s being a subset of the set of nodes M : $s \in M$. The restriction gives a new dependence graph which will contain a restricted number of subgraphs to include. All parameters for substructures that are not part of the dependency graph are equal to 0. Obviously, defining the settings oneself required extensive knowledge about the network at hand. The inclusion of these settings results in what Pattison, Robins (2002) refer to as *partial conditional dependence*.

Of course, one can also include other types of variables to a model, such as age of the firm, geographic location, amount of public funds received etc. These variables are referred to as node variates or node attributes. The addition of these attributes is introduced by Robins et al. (2001). The idea here is that links depend upon the attributes of the nodes they are connecting. In other words the probability that two nodes are connected depends upon the attribute of the node. These models are also called *social selection models* and take the following form:

$$p(Y = y | X = x) = \frac{1}{\psi(\theta)} \cdot \exp\left\{\sum_i \theta z(x) + \theta^a \cdot z^a(x, y)\right\} \quad (21)$$

Where the exponent a indicates parameters and structures relative to the inclusion of dyadic dependence for the purpose of social selection.

In the same paper Robins et al. also described a *social influence model* in which the attributes of the nodes are a result of the structure of the network (nodes take a particular attribute according to the nodes in the neighborhood for example). In other words, the probability that a node variables takes a particular value depends upon the structure of the network and the values of this (or indeed any other) node-level variable.

We hence need to add a node variables to the model. Suppose we note Y_i the value of a node variable for node i . This variable can be anything from a geographic region to the amount of public investment received by firms to the number of publications or patents. When this variable is included in the ERGM the model is written:

$$p(Y = y | X = x) = \frac{1}{\psi(\theta_i)} \cdot \exp\left\{\sum_i \theta z(y, x)\right\} \quad (22)$$

Just as it is possible to put values on nodes it is also possible to values on dyads. The question then is to know if the value of the dyad increases the probability of nodes being connected. Think of situations where we would like to know if the amount of investment between firms is related to cooperation or if technological proximity between firms induces collaboration. Note here that the difference between dyadic covariates and actor attributes resides in the value on the link between two nodes. In the case of proximity is refers to the proximity of both firms, it is hence not a firm-level variable. The value only makes sense when we consider firms two-by-two.

All the extensions made to the ERGM framework allow researchers to answer a large variety of questions about social and economic processes. Many other extensions which are beyond the scope of this document, but worth noting, are multivariate relations in networks (Pattison, Wasserman, 1999) and dynamic networks in which new ties depend upon the structure of the network at time $t - 1$. It is also possible to model multivariate networks using ERGM. The idea is then that each link can exist in different matrixes, each corresponding to a different type of link (social, work, geography etc.). This extension allows researchers to study interplay between different networks and how each network is affected by the other networks.

Before looking at estimation methods for ERGM models one problem needs to be addressed: the degeneracy problem.

2.2 Solving degeneracy: Curved ERGMs

ERGM models are prone to degeneracy issues. When estimating the model the change statistics can behave in such a way that the large majority of the probability distribution is placed on either an empty or a full graph. As we will discuss in more detail a bit later, a simulation is performed to identify the probability distribution of a graph. This is done on a step by step basis, starting with an empty network and adding links one by one until a maximum likelihood is achieved. This probability distribution is a function of the change statistics and is thus impacted by the change in log-odds for an additional tie (for a given effect). In other words if an additional edge would create two new 2-stars, then the log-odds of observing that tie would increase by two multiplied by the parameter of the 2-star variable. A random graph model is considered stable if small changes in the parameter values result in small changes in the probabilistic structure of the model (Handcock et al., 2003). When a new edge is added to the graph this not only increases the number of edges but might also increase the number of other sub-configurations that might be included in the model. The parameters of the model control the number of sub-graphs of a particular type that are expected in the graph. For instance a 2-star might be transformed into a triangle by the addition of an edge which also adds two 2-stars. A 2-star can become at 3-star and so on. This cascading effect will result in the simulation procedure jumping over the MLE and converge to a full graph. Lusher et al.

(2012) [chapter 6] show that, in the case of the Markov (or triad model), the number of expected triangles increases as the parameter for triangles increases. They highlight a phase transition for certain values of the parameter where the number of expected triangles increases more than exponentially. This transition explain that the probability density distribution has high weights either on a (close to) empty graph or on a complete graph⁴. This problem is increasingly present as the number of nodes increases. The larger the network the higher the number of substructures one can have.

In order to avoid the model to put too much weight on the full graphs, Snijders et al. (2006) propose to add several variables based on their concept of *partial conditional dependence*. The idea is to include a weighted degree distribution to the model, giving a high weight to low density while decreasing the weights as the degree increases. This reduces the impact of the high density variables responsible for the degeneracy of the initial model. Mathematically we can then write (using the notations of the initial paper):

$$u_{\alpha}^{(d)} = \sum_{k=0}^{n-1} e^{-\alpha k} d_k(y) \quad (23)$$

Where $d_k(y)$ is the number of nodes with degree k and α the parameter of the weights. This is referred to as the *geometrically weighted degree distribution*. The degree distribution can also be written as a function of the stars in the network. After all, a degree distribution is nothing more than a distribution of stars. Nodes with a degree of five are 5-stars, degree two are 2-stars and so forth. We can hence formulate the distribution as follows:

$$u_{\lambda}^s = S_2 - \frac{S_3}{\lambda} + \frac{S_4}{\lambda^2} - \dots + (-1)^{n-2} \cdot \frac{S_{n-1}}{\lambda^{n-3}} = \sum_{k=2}^{n-1} (-1)^k \cdot \frac{S_k}{\lambda^{k-2}} \quad (24)$$

Where S_k is a the number of stars of degree k and λ the parameter. This method is referred to as *Alternating k-stars*. The difference between the geometrically weighted degree distribution and the K-stars is resides in the alternating signs. A large value of 3-stars is counterbalanced by a negative value for 4-stars due to the inverse sign of the parameter. The addition of the weights ensures that the change in change statistics stays small. Indeed Snijders et al. (2006) show that the change statistics can be written:

$$z_{ij} = -(1 - e^{-\alpha})(e^{-\alpha y_{i+}} + e^{-\alpha y_{j+}}) \quad (25)$$

Where y_{i+} is the density of firm i when the link between i and j is added. Equation 25 shows that the value of the change statistic is reduced by the factor $-(1 - e^{-\alpha})$. This factor hence ensures that the change statistics do not take too high values and result in a nested probability distribution. The inclusion of either the alternating k-stars or the geometrically weighted degrees transform the ERGM model into a Curved Exponential

⁴In addition, the degeneracy of the model can result in problems with the estimation procedures.

Random Graph Model (Efron, 1975).

All that is left to do now is estimate the model.

3 Estimation

Estimation allows for the identification of the parameters that maximize the likelihood of a graph. Since we only have one observation (the observed graph) a set of graphs from the same distribution is required. The set of graphs that may be generated by this procedure should have the observed graph as a central element to ensure a valid sample.

3.1 Markov Chain Monte Carlo

In the first section we identified the general form of an ERGM model (see equation 18). The odds of a graph were normalized by the sum of the parameters of all possible graphs. This leaves us with a constant to estimate which is near impossible. A workaround has to be found for ERGMs to be useful. A first development by Besag (1975); Strauss, Ikeda (1990) was to estimate the model using pseudo-likelihood estimation. The properties of this method are however not clear (Snijders et al., 2006; Robins et al., 2007) we shall hence focus here on more recent methods that are better understood.

A method for estimating the parameters of ERGMs using a sample is developed by Anderson et al. (1999); Snijders (2002); Geyer, Thompson (1992). They estimate the model by Markov Chain Monte Carlo (MCMC) to find the maximum likelihood estimates (MLE). The idea is to extract a sample from a distribution that follows equation 18 asymptotically, not requiring the direct computation of the normalizing constant. Their paper points out that almost any maximum likelihood can be accomplished by a MCMC.

A Markov chain is a sequence of random variables such that the value taken by the random variable only depends upon the value taken by the previous variable. We can hence consider a network in the form of an adjacency matrix in which each entry is a random variable. By switching the values of these variables to 0 or to 1 (adding or removing a link from the network) one can generate a sequence of graphs such that each graph only depends upon the previous graph. This would be a Markov chain. The hypothesis is then that if the value at step t is drawn from the correct distribution than so will the value at step $t + 1$. Unlike regular Monte-Carlo methods, the observations that are sampled are close to each-other since they vary by a single link. However, one would need a method for selecting which variable should change state in order to get closer to the MLE, this is done using the Metropolis-Hastings algorithm or the Gibbs sampler.

3.2 Metropolis-Hastings algorithm and the Gibbs sampler

The Metropolis-Hastings algorithm picks a variable at random and changes it's state. This results in either a new edge in the network or in the disappearance of an edge. The

probability of the graph is then computed and only if the probability of the altered graph is higher than the previous one is the new graph retained for the next step. In other words the new graph is retained as long as the likelihood is increased:

$$\min\left\{1, \frac{p_{\theta}(x^*)}{p_{\theta}(x^{m-1})}\right\} \quad (26)$$

This decision rule is called the *Hastings ratio*. The advantage of this ratio is that it does not include the normalizing constant $\psi(\theta)$.

Since the Markov chain starts at 0, a burn-in is needed to remove part of the chain to identify if the chain has converged or not (the burn-in can be parameterized in most software).

The steps taken by the Metropolis-Hastings algorithm are quite small. These small steps are implemented in order to avoid overstepping the global optimum which can easily happen in the case of larger parameter spaces. Other methods allow for bigger steps and as such converge faster and need a lower burn-in. The risk of larger steps is however overstepping the global optimum and convergence towards other local optima. The Metropolis-Hastings algorithm may be slower than others but is more precise in its estimation.

Some programs use the Gibbs sampler, which is a special case of the Metropolis-Hastings algorithm (Hunter, Handcock, 2006). The difference between Gibbs and Metropolis-Hastings resides in the chosen step. In the case of the Gibbs sampler, the state of each element in the vector of parameters is chosen and updated conditionally on the state of the other parameters. This means that if this decision rule was implemented in the Metropolis-Hastings algorithm the probability that the change is retained is equal to one. This makes the Gibbs sampler a relatively fast method. This sampling method is used by the different algorithms that are used to estimate ERGM models.

Both methods allow for the generation of a sample of graphs that can be used for inference. The sample of graphs is obtained by varying not the parameters but the variables of the model until it is centered around the observed graph. Now that a sample of graphs has been found we need to estimate the parameters of the model. Two of the most widely used estimation algorithms, the "Stepping" and "Robbins-Monro" algorithm will now be reviewed.

3.3 The "Stepping algorithm"

This method introduced by Hummel et al. (2012). It has the advantage of approaching the MLE directly while "Robbins-Monro" does not. ERGM models are indeed estimated using the maximum likelihood method. Starting from the canonical ERGM form we define the log likelihood function as:

$$L(\theta) = \theta \cdot v(G) - \log(\psi(\theta)) \quad (27)$$

The problem here is the presence of the normalizing constant which cannot be computed. The improvements of this method over the previously one resides in the use

of a log-normal approximation. The algorithm proposed here will converge towards the log-likelihood using a step-by-step method. The sampler used in with this estimation procedure is the metropolis-hastings sampler discussed previously. Once a sample of graphs has been identified the estimation algorithm is launched. Since the normalizing constant in equation 27 cannot be compute a workaround has to be found. The idea is to give starting parameters (θ_0).The log-likelihood ratio can then be written (Hummel et al., 2012):

$$L(\theta) - L(\theta_0) = (\theta - \theta_0)^T V(G) - \log E_{\theta_0}[\exp(\theta - \theta_0)^T V(g)] \quad (28)$$

Geyer, Thompson (1992) point out that maximizing this ration by the means of a sample distribution of graphs generated with θ_0 only behaves well when θ is close to θ_0 . In other words one has to choose the correct starting point for the algorithm to find the MLE. The MLE solves the equation:

$$E_{\hat{\theta}} v(G) = v(G^{obs}) \quad (29)$$

The idea is to suppose that the MLE is not the observed value of the parameters but some point between the mean value parameterization and the observed value. A parameter γ defines the steps taken:

$$\hat{\omega}_t = \gamma_t \cdot v(G) + (1 - \gamma_t)\bar{\omega} \quad (30)$$

Where ω_t represents the estimate in the mean parameter space⁵. Ideally then, we would want $\gamma = 1$ so that the expected value of the parameters is the observed value. If this is the case the algorithm is considered to converge, this is shown in figure ?? which is the output of the R code. Once convergence is detected a large sample based on the parameters is computed and the MLE are estimated and gives as the final values.

- Step 1 : Set the iteration number equal to 0 and choose initial parameters for vector η_0 .
- Step 2 : Use MCMC to simulate graphs from the probability function for parameter vector η_0 .
- Step 3 : Compute the mean of the sample.
- Step 4 : Define a pseudo observation that is a convex combination of the mean of the sample and the observed value.
- Step 5 : Replace the observed value by the pseudo observation.

⁵In the algorithm the initial values are chosen to be the MPLE

Robbins-Monro The Robbins-Monro algorithm is a stochastic approximation method introduced by Robbins, Monro (1951) which is used by Snijders (2001) and Snijders (2002) to estimate ERGM models. Typically the method estimates:

$$E\{Z_\theta\} = 0 \quad (31)$$

Where Z_θ is a vector of parameters equal to $u(Y) - u_0$ where u_0 is the observed value of the statistics. This allows us to rewrite the equation as a moment equation. The algorithm gives starting parameter values equal to the average of the parameter values. The initial parameters that will launch phase two are defined by (Lusher et al., 2012):

$$\theta^{t+1} = \theta^t - a_t \cdot D^{-1}(z(x^m) - z(x_{obs})) \quad (32)$$

Where D is the co-variance matrix, the diagonal of this matrix will be used as the scaling matrix. a defines the convergence, it is set to $a_t = \frac{a_{t-1}}{2}$. The idea is that each step brings the values closer to the MLE. Hence large steps might result in the exceeding the MLE and divergence. The fact that the a_t reduces in value with each step allows a smooth path to the MLE. As we move closer to the observed values of the statistics $z(x^m) - z(x_{obs})$ tends towards zero. The R output in figure ?? shows how the steps (a) start at a value of 0.1 and tend towards 0 with each iteration of the second phase of the algorithm. At the start of each step the starting parameters are considered to be the average values of the previous step. The number of iterations varies from model to model. The iterations stop once the trajectories of the generated statistics cross those of the observed ones (Lusher et al., 2012).

The burn-in represents the number of simulations that are removed from the MCMC in order to make the chain "forget" the starting point. In other words it is to make sure the starting values do not impact the final result.

Finally the algorithm checks for convergence using a convergence statistic. Just as in the case of the stepping algorithm one supposes that the MLE is reached when the distance between the observed values and the average of the simulated ones is close to 0. If there is no convergence than one can relaunch the estimation with as starting parameters the results of the previous simulation (Lusher et al., 2012).

The largest difference between this method and the stepping method resides in two factors. First this method approaches an estimate of the MLE and does not evaluate the MLE function directly. Second, the steps are of a higher magnitude and can exceed the MLE if the starting values are close to the MLE. The use of either of the algorithms purely depends upon the model to be estimated. One algorithm might have better convergence in one case while the opposite can be true in another case. Note however that both use the Metropolis-Hastings method for the simulation of the MC.

Now that we have discussed which variables can be included and how to estimate the parameters we will turn to an application using R.

4 Code R and example

We use here different R packages (Hunter et al., 2008a) , (Hunter et al., 2008b; Handcock et al., 2008; Butts, 2008).

The data (and hence the results of the estimations) are from the French aerospace collaboration network. Using patent data collaborations were identified which resulted in a network. The aim of the study is to understand if technological proximity played a significant role in the structuring of the collaboration network. We hence used a dyadic covariate called "proximity". The network contains 176 firms.

```
1 #Import data.
2 Network<-read.table("ADJ_MATRIX.csv", sep=";")
```

The data used here was already in the form of an adjacency matrix and hence could be used directly. It is however also possible to use edgelist. Since the data needs to be transformed into a network object the network package will be needed. The latter is able to transform edgelist into adjacency matrices.

```
1 #Import the dyadic covariates
2 proximity<-read.table("Proximity_matrix.csv", sep=";")
3 proximity.e<-read.table("Proximity_matrix_exp.csv", sep=";", dec=",")
4 citation<-read.table("Citation_matrix.csv", sep=";")
```

Since I'm using two measures of proximity I have two matrices and a matrix that includes the number of citations between firms. These need to be imported in the same manner as the network itself.

```
1 #We now transform the imported data into network objects with the
  package 'network'
2 Network<-as.matrix(Network)
3 proximity<-as.matrix(proximity)
4 proximity.e<-as.matrix(proximity.e)
5
6 #Transform to network format
7 Network<-as.network(Network, directed = FALSE)
8 proximity<-as.network(proximity, directed=FALSE)
9 exp_proximity<-as.network(exp_proximity, directed=FALSE)
10 citation<-as.network(citation, directed=TRUE)
```

We now have different objects to work with: the network and dyadic covariates in the form of networks. Note that the model constructed here is for an undirected network as shown by the option *directed = FALSE* in the *as.network* function.

We now have to decide which variables to include in the model. Let's start with a simple model containing only edges. We invoke here the *ergm()* function from the *ergm* package:

```
1 model<-ergm(Network ~ edges)
```

This gives us the most basic form of an ERGM model, the estimation method defaults to Monte Carlo MLE (Geyer, Thompson, 1992).

```
1 =====
2 Summary of model fit
```

```

3 =====
4
5 Formula:   Network ~ edges
6
7 Iterations: 7 out of 20
8
9 Monte Carlo MLE Results:
10      Estimate   Std. Error   MCMC   \%p-value
11 edges -3.85280   0.05649     0      <1e-04 ***
12 -----
13 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
14
15      Null Deviance: 21349   on 15400 degrees of freedom
16 Residual Deviance: 3113   on 15399 degrees of freedom
17
18 AIC: 3115   BIC: 3122   (Smaller is better.)

```

The parameter for the variable *edges* has an estimated value of -3.8528. This means that the probability of two ties connecting is:

$$p(i \rightarrow j) = \frac{\exp(-3.85)}{1 - \exp(-3.85)} = 0.02174241 \quad (33)$$

Recall equation 11, this equation stated that the variables were change statistics. The parameter should hence be multiplied by the change in the number of subgraphs. In other words, if an additional edge creates three triangles the parameter should be multiplied by three. Since an additional edge only creates one new edge we do not multiply. Let's try the same but with only triangles as explanatory variable.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula:   Network ~ triangles
6
7 Iterations: NA
8
9 Stepping MLE Results:
10      Estimate   Std. Error   MCMC   \% p-value
11 triangle -1.91496   0.01712     0      <1e-04 ***
12 -----
13 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
14
15      Null Deviance: 21349   on 15400 degrees of freedom
16 Residual Deviance: 8935   on 15399 degrees of freedom
17
18 AIC: 8937   BIC: 8945   (Smaller is better.)

```

The parameter for the variable *triangles* has an estimated value of -1.9146. This means that the log-odds of two nodes connecting is:

$$-1.9146 * \delta \text{ triangles} \quad (34)$$

Where $\delta \text{ triangles}$ gives the change in the number of triangles. Hence the log-odds depend upon the number of triangles that will be created by an additional tie:

- > If the link creates 1 triangle, the log-odds are $1 \cdot -1.9146$. The probability is then 0.1284
- > If the link creates 2 triangles, the log-odds are $2 \cdot -1.9146$. The probability is then 0.0213

Let's see how we interpret estimates when we have more than one variable: the Markov model.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula:   Network ~ edges + kstar(2) + triangles
6
7 Iterations: NA
8
9 Monte Carlo MLE Results:
10      Estimate Std. Error MCMC % p-value
11 edges      -28.30222    0.22246      0 <1e-04 ***
12 kstar2       7.68159    0.07356      0 <1e-04 ***
13 triangle    57.09972    0.79544      0 <1e-04 ***
14
15 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
16
17 Null Deviance:    21349 on 15400 degrees of freedom
18 Residual Deviance: 35811780 on 15397 degrees of freedom
19
20 AIC: 35811786    BIC: 35811809    (Smaller is better.)

```

In this case an additional edge, if it creates x 2-stars and y triangles, has log-odds:

$$-28.30 + x \cdot 7.68 + y \cdot 57.099$$

Since the model includes lower and higher order subgraphs (2-stars are a substructure of triangles) we can conclude here that triadic closure is significant in the network. In other words, 2-stars are closed to form triangles.

Note however the values of the information criteria AIC and BIC, stating these values are high is an understatement. The model must suffer from some kind of problem. Subsection 4.3 will show how to deal with degeneracy and other problems.

Now we want to see how proximity (technological in the case of this study) influences the structuring of the network. Let's start with a simple model using the edge covariate "proximity". The comment for adding an edge covariate is simple *edgescov*. Similarly node covariates (age, gender, type of firm, country...) are added using the *nodecov()* or *nodefactor* command. The node factor command is particularly useful since it allows to compare log-odds to a reference point. For example, one could categorize the R&D expenses of firms into *low*, *average* and *high*. The ERGM would then give the log-odd of 2 categories as compared to a third. In other words is a link more likely for firm with average expenses than for low? This feature seems not to be available for edge covariates however.

The following model only contains edges and the edge covariate. A first point we can notice is that the AIC and BIC criteria are lower with the addition of the edg covariate. The model is hence improved with the addition of the proximity parameter. Firms with proximity are 2.13 times more likely to connect in this network (*ceteris paribus*)⁶. The probability of an additional edge is then positively impacted by the technological proximity. More specifically the average degree of the network is positively impacted by technological proximity.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula: Network ~ edges + edg cov (proximity)
6
7 Iterations: 7 out of 20
8
9 Monte Carlo MLE Results :
10 Estimate Std. Error MCMC % p-value
11 edges -4.5222 0.1972 0 < 1e-04 ***
12 edg cov . proximity 0.7575 0.2058 0 0.000234 ***
13
14 Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
15
16 Null Deviance: 21349 on 15400 degrees of freedom
17 Residual Deviance: 3096 on 15398 degrees of freedom
18
19 AIC: 3100 BIC: 3115 (Smaller is better.)

```

A network analysis performed on this network showed that the network has a scale-free structure. This information can be helpful in the modeling of the ERGM as we have information on the distribution of the degrees. The same is valid for any other information about the network, small-world properties, level of clustering or centrality distribution. The information provided by SNA allows a first understanding of the structural properties of the network that will allow for a more robust model once edge and nodal covariates are added.

So if the network has a scale-free structure the structure should be explained by the degree distribution. To check this we can add different degrees to the model. This can be done by using the command *degree()*. One can simply add one statistic for one degree, i.e. *degree(3)*, for the impact of nodes with degree 3, or add multiple degrees as was done in the following model. Note that the addition of multiple degrees is achieved by writing *degree(a : b)* to add all the degrees between *a* and *b*.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula: Network ~ edges + degree(2:6) + edg cov (proximity)
6
7 Iterations: 7 out of 20
8
9 Monte Carlo MLE Results :
10 Estimate Std. Error MCMC % p-value

```

⁶ $e^{0.7575}$ since this is the odds and not the log-odds

```

11 edges          -4.2417    0.1606    0 < 1e-04 ***
12 degree2       -0.8945    0.2233    0 < 1e-04 ***
13 degree3       -1.4608    0.2381    0 < 1e-04 ***
14 degree4       -1.8040    0.2720    0 < 1e-04 ***
15 degree5       -1.6668    0.2699    0 < 1e-04 ***
16 degree6       -2.0959    0.4109    0 < 1e-04 ***
17 edgcov.proximity 0.5451    0.1726    0 0.00159 **
18 -----
19 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
20
21 Null Deviance: 21349 on 15400 degrees of freedom
22 Residual Deviance: 2989 on 15393 degrees of freedom
23
24 AIC: 3003    BIC: 3056    (Smaller is better.)

```

The addition of these variables to the model once again decreases the AIC and the BIC, the model is hence enhanced. The structure of the network can be explained by a degree distribution.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula:    Network ~ edges + degree(2:6) + edgcov(proximity2)
6
7 Iterations: 7 out of 20
8
9 Monte Carlo MLE Results:
10
11           Estimate Std. Error MCMC % p-value
12 edges          -4.2382    0.1815    0 < 1e-04 ***
13 degree2       -0.8764    0.2053    0 < 1e-04 ***
14 degree3       -1.4438    0.2229    0 < 1e-04 ***
15 degree4       -1.7876    0.2728    0 < 1e-04 ***
16 degree5       -1.6763    0.2874    0 < 1e-04 ***
17 degree6       -2.1068    0.4595    0 < 1e-04 ***
18 edgcov.proximity2 0.5446    0.1941    0 0.00502 **
19 -----
20 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
21
22 Null Deviance: 21349 on 15400 degrees of freedom
23 Residual Deviance: 2989 on 15393 degrees of freedom
24
25 AIC: 3003    BIC: 3057    (Smaller is better.)

```

These models seem to work quite nicely. We discussed in the previous sections that models were prone to degeneracy and the solutions to this problem. Let's have a look at how we can model ERGMs with alternating k-stars and a geometrically weighted degree distribution.

4.1 Curved Exponential Random Graph Models

Studies show that the addition of weights on the degree distribution helps to avoid bi-modal distributions in the parameter space, i.e. avoids the generated networks from being either full or close to empty. Different forms can be added to the R code. Since we have here an undirected network we can use either the alternating k-stars `altkstar()`

or the geometrically weighted degree distribution *gwdegree*. For directed graph there are additional commands which work in similar manner as what we show here. We include here a statistic for the gwdegree with the option *fixed = TRUE*. The latter means that we do not make an estimation of the scaling parameter, we want it to be equal to 1. The resulting model is hence not a curved ERGM.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula:   Network ~ edges + triangles + edg cov(proximity2) + gwdegree
6           (1,
7           fixed = TRUE)
8
9 Iterations: NA
10
11 Stepping MLE Results:
12           Estimate Std. Error MCMC % p-value
13 edges      -5.717e+00  1.829e-01    0 < 1e-04 ***
14 triangle   1.802e+00  3.026e-05    0 < 1e-04 ***
15 edg cov. proximity2  6.811e-01  2.159e-01    0 0.001607 **
16 gwdegree   2.917e-01  8.380e-02    0 0.000502 ***
17
18 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
19
20 Null Deviance: 21349 on 15400 degrees of freedom
21 Residual Deviance: 4170 on 15396 degrees of freedom
22
23 AIC: 4178    BIC: 4208    (Smaller is better.)

```

In order to have a curved exponential random graph model, the parameter that defines that we fix in the previous code has to be estimated as well. In the following code we estimated the model with an edgewise shared partners variable (see figure 24). This variable is used to check for transitivity. By switching the option *fixed = TRUE* to *fixed = FALSE* the model becomes curved. The results now include an estimate for the parameter alpha of the model. Note here that the parameter can only be interpreted if the *gwesp* statistic is significant.

```

1 =====
2 Summary of model fit
3 =====
4
5 Formula:   Network ~ edges + edg cov(proximity) + gwesp(alpha = 1,
6           fixed = FALSE)
7
8 Iterations: NA
9
10 Stepping MLE Results:
11           Estimate Std. Error MCMC % p-value
12 edges      -5.39343   0.45601    0 <1e-04 ***
13 edg cov. proximity  0.48255   0.51641    0 0.05 ***
14 gwesp       1.19503   0.08333    0 <1e-04 ***
15 gwesp.alpha  0.88784   0.09792    0 <1e-04 ***
16
17 Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```



```

18 Null Deviance: 21349 on 15400 degrees of freedom
19 Residual Deviance: 2847 on 15396 degrees of freedom
20
21 AIC: 2855 BIC: 2885 (Smaller is better.)

```

The interpretation of these estimates are much more complex than previously. The parameters need to be exponentiated to find λ that we saw in the equations (Hunter, 2007). We hence find $e^{0.88784} = 2.42$. Since the parameter is positive we can conclude that transitivity is present.

4.2 Goodness of fit diagnostics

In order to check if a model is a good fit we use the `mcmc.diagnostics` command. This gives us a number of outputs, notable the matrix of correlations and p-values for both the individual parameters and the model as a whole.

```

1 Sample statistics cross-correlations:
2           kstar2  edg cov . proximity
3 kstar2           1.0000000      0.5494967
4 edg cov . proximity 0.5494967      1.0000000
5
6 Individual P-values (lower = worse):
7           kstar2  edg cov . proximity
8           0.3126642      0.9294963
9 Joint P-value (lower = worse): 0.6994233 .

```

The p-values are high for the parameters and the model, we can hence conclude that the model is globally significant. In order to go into a bit more detail when it comes to the estimates, the commands also provides us with plots, see figures 8 and 9. We stated that an ERGM should fit the observed network perfectly, on average. This means that from the simulated networks we expect the average values to be those of the observed network. If this is not the case then the sample the model is based on does not come from the same distribution as the observed network.

Figure 8 shows an example of what we want to observe. in the first graphs the values oscillate around the mean which is what we want. The graphs on the right hand show a centered distribution of the values, we hence conclude that the model is a good fit. A bad example can be found in figure 9. The graphs show that the distribution is not at all centered, and there is no oscillation around the mean. This model is hence a bad fit.

One can also study the goodness of fit using the `gof()` command. Using a plot command this provides a box plot (see figure10).

The `gof` command can receive different parameters, one can chose to plot any number of variables and decide to increase or decrease the number of simulations to refine the results. The following code provides the GOF of the whole model ($GOF = Model$) using 20 simulations ($nsim = 10$).

```

1 gof_model<-gof(model, GOF=~Model, nsim=20)

```

The results give us a box plot per variable and a black line representing the observations on the empirical network. Since we want the mean of the simulations to be equal to the observed network, the dark line should coincide with the center of the boxplots (the vertical line in the boxplot representing the median of the distribution). This is the case here, the model is hence a good fit.

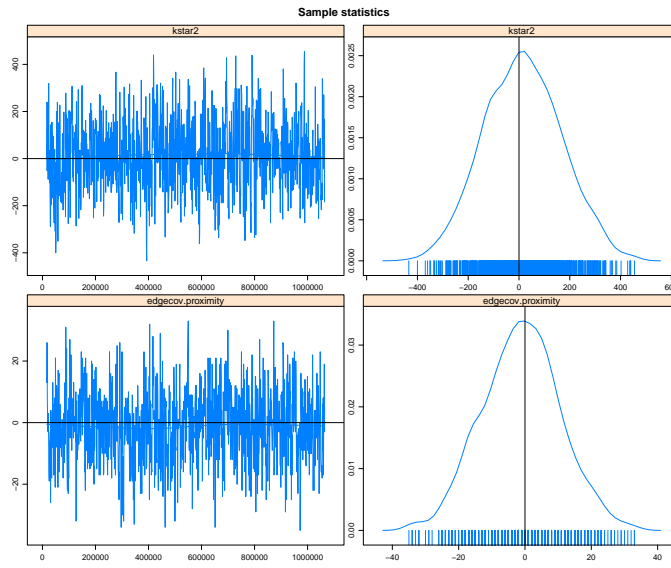


Figure 8: Goodness of Fit diagnostics

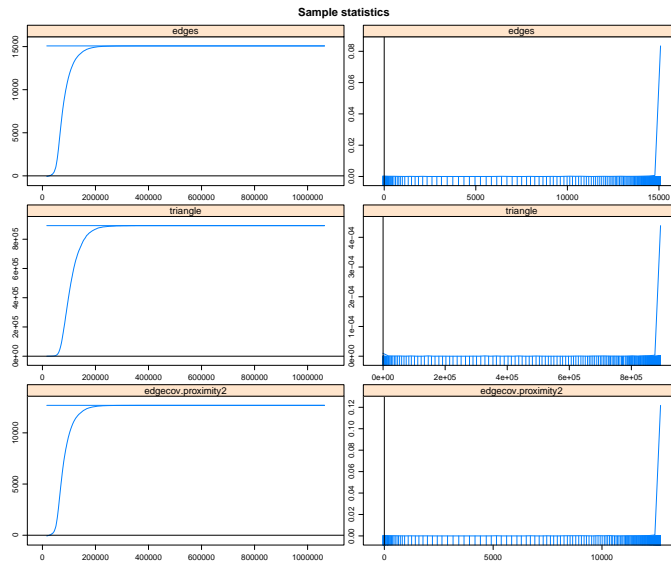


Figure 9: Goodness of Fit diagnostics, bad example

4.3 Improve bad models

The fitting of an ERGM is a trial and error procedure. If a model behaves badly there are a couple of parameters to change in order to improve results. Of course one should

only start these procedures once the variables chosen are stabilized. Starting with a null model containing only edges and adding on to this model while comparing the AIC and BIC values to find the variables of importance.

Once this is done and degeneracy is still observed one can start by switching estimation methods. One method might work better in one case than the other.

The estimation algorithm can be chosen in the control arguments of the `ergm()` command.

```
1 model<-ergm(Network~ edges , control=control.ergm(main.method = "  
  Stepping"))  
2 model<-ergm(Network~ edges , control=control.ergm(main.method = "  
  Stepping" ,MCMC.samplesize=70000, MCMC.interval=5000))
```

The burn-in can also solve problems, the burnin represents the number of iterations that are removed from the simulation. In other words, the higher the burnin the more the procedure forgets about its initial parameters. Increasing this value hence allows for keeping only the latest values which might represent the real values better. This can be achieved by adding an option to the `ergm`.

A second method of improving estimation would be to increase the sample size by changing the parameter `MCMC.samplesize`. This increase will result in having more precise estimates by an increase in the number of statistics drawn from the sample. This, of course, increases computation time. The `ergm` package includes multicore features that can help reduce computation time drastically. All this requires is the addition of some parameters to the `control.ergm` argument. Adding `parallel = 4` notifies the package that the computer has 4 cores, `parallel.type` sets the type of multicore. For a regular computer this should be fixed to `"PSOCK"`. This will distribute the computations over the 4 cores of the computer and hence increase speed.

```
1 a<-ergm(Network~edgecov(proximity2)+triangles+altkstar(1.812, fixed=  
  FALSE) , control=control.ergm(main.method=c("Stepping") , parallel=4,  
  parallel.type="PSOCK" , MCMC.samplesize=20000))
```

5 Conclusion

ERGM models are able to analyze the structure of network to an extent that other methods of network analysis are unable to reach. A basic analysis of the network structure can come in handy when defining the ERGM model and can be used as a verification procedure. Whether they are used to analyse social networks, collaboration networks, trade networks or financial networks, ERGMs can provide vital insights into the understanding of network dynamics. Even though the models are powerful, the tools used for their analysis still need to be improved. The tweaking required in the estimation procedures should be reduced for the models to be used by a larger audience.




Figure 10: GOF: boxplot analysis




Figure 11: Empty




Figure 12: Dyad




Figure 13: 2-star




Figure 14: Triad

Figure 15: Triads




Figure 16: 1-star




Figure 17: 2-star




Figure 18: 3-star




Figure 19: 4-star

Figure 20: k-stars




Figure 21: 1 partner




Figure 22: 2 partners




Figure 23: 3 partners

Figure 24: Shared partners (edgewise)




Figure 25: 1 partner




Figure 26: 2 partners




Figure 27: 3 partners

Figure 28: Shared partners (dyadic)

References

- Anderson Carolyn J, Wasserman Stanley, Crouch Bradley.* A p* primer: Logit models for social networks // *Social Networks*. 1999. 21, 1. 37–66.
- Belke A, Polleit T.* Money and Swedish inflation // *Journal of Policy Modeling*. 2006. 28, 8. 931–942.
- Belongia M.T, Ireland P.N.* A working solution to the question of nominal GDP targeting // *Macroeconomic Dynamics*. 2014. 31. cited By 0.
- Besag Julian E.* Nearest-neighbour systems and the auto-logistic model for binary data // *Journal of the Royal Statistical Society. Series B (Methodological)*. 1972. 75–83.
- Besag Julian E.* Statistical analysis of non-lattice data // *The statistician*. 1975. 179–195.
- Broekel Tom, Hartog Matt e.* Explaining the structure of inter-organizational networks using exponential random graph models // *Industry and Innovation*. 2013. 20, 3. 277–295.
- Butts Carter T.* network: a package for Managing Relational Data in R. // *Journal of Statistical Software*. 2008. 24, 2.
- Caimo A., Lomi A.* Knowledge Sharing in Organizations: A Bayesian Analysis of the Role of Reciprocity and Formal Structure // *Journal of Management*. 2015. 41, 2. 665–691. cited By 2.
- Carrington Peter J, Scott John, Wasserman Stanley.* Models and methods in social network analysis. 28. 2005.
- Cranmer S.J., Desmarais B.A., Menninga E.J.* Complex Dependencies in the Alliance Network // *Conflict Management and Peace Science*. 2012. 29, 3. 279–313. cited By 13.
- Czudaj R.* P-star in times of crisis - Forecasting inflation for the Euro area // *Economic Systems*. 2011. 35, 3. 390–407.
- Efron Bradley.* Defining the curvature of a statistical problem (with applications to second order efficiency) // *The Annals of Statistics*. 1975. 1189–1242.
- Frank Ove, Strauss David.* Markov graphs // *Journal of the American Statistical Association*. 1986. 81, 395. 832–842.
- Geyer Charles J, Thompson Elizabeth A.* Constrained Monte Carlo maximum likelihood for dependent data // *Journal of the Royal Statistical Society. Series B (Methodological)*. 1992. 657–699.
- Markov fields on finite graphs and lattices. // . 1971.
- Handcock Mark S., Hunter David R., Butts Carter T., Goodreau Steven M., Krivitsky Pavel N., Bender-deMoll Skye, Morris Martina.* statnet: Software Tools for Statistical Analysis of Network Data // *Journal of Statistical Software*. 2008. 24, 1. 1 – 11.

- Handcock Mark S, Robins Garry, Snijders Tom AB, Moody Jim, Besag Julian.* Assessing degeneracy in statistical models of social networks. 2003.
- Harris Jenine K.* An introduction to exponential random graph modeling. 173. 2013.
- Hummel Ruth M, Hunter David R, Handcock Mark S.* Improving simulation-based algorithms for fitting ERGMs // *Journal of Computational and Graphical Statistics.* 2012. 21, 4. 920–939.
- Hunter David R.* Curved exponential family models for social networks // *Social networks.* 2007. 29, 2. 216–230.
- Hunter David R, Handcock Mark S.* Inference in curved exponential family models for networks // *Journal of Computational and Graphical Statistics.* 2006. 15, 3.
- Hunter David R., Handcock Mark S., Butts Carter T., Goodreau Steven M., Morris , Martina .* ergm: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks // *Journal of Statistical Software.* 2008a. 24, 3. 1 – 29.
- Hunter David R, Handcock Mark S, Butts Carter T, Goodreau Steven M, Morris Martina.* ergm: A package to fit, simulate and diagnose exponential-family models for networks // *Journal of statistical software.* 2008b. 24, 3. nihpa54860.
- Lomi A, Fonti F.* Networks in markets and the propensity of companies to collaborate: An empirical test of three mechanisms // *Economics Letters.* 2012. 114, 2. 216–220.
- Lomi A., Pallotti F.* Relational collaboration among spatial multipoint competitors // *Social Networks.* 2012. 34, 1. 101–111. cited By 14.
- Lusher Dean, Koskinen Johan, Robins Garry.* Exponential random graph models for social networks: Theory, methods, and applications. 2012.
- Özdemir K.A, Saygili M.* Monetary pressures and inflation dynamics in Turkey: Evidence from P-star model // *Emerging Markets Finance and Trade.* 2009. 45, 6. 69–86. cited By 5.
- Pattison Philippa, Robins Garry.* Neighborhood-based models for social networks // *Sociological Methodology.* 2002. 32, 1. 301–337.
- Pattison Philippa, Wasserman Stanley.* Logit models and logistic regressions for social networks: II. Multivariate relations // *British Journal of Mathematical and Statistical Psychology.* 1999. 52, 2. 169–194.
- Price S, Nasim A.* Modelling inflation and the demand for money in Pakistan; Cointegration and the causal structure // *Economic Modelling.* 1998. 16, 1. 87–103. cited By 4.
- Robbins Herbert, Monro Sutton.* A stochastic approximation method // *The annals of mathematical statistics.* 1951. 400–407.

- Robins Garry, Pattison Philippa, Elliott Peter.* Network models for social influence processes // *Psychometrika*. 2001. 66, 2. 161–189.
- Robins Garry, Pattison Pip, Kalish Yuval, Lusher Dean.* An introduction to exponential random graph (p*) models for social networks // *Social networks*. 2007. 29, 2. 173–191.
- Snijders Tom AB.* The statistical evaluation of social network dynamics // *Sociological methodology*. 2001. 31, 1. 361–395.
- Snijders Tom AB.* Markov chain Monte Carlo estimation of exponential random graph models // *Journal of Social Structure*. 2002. 3, 2. 1–40.
- Snijders Tom AB, Pattison Philippa E, Robins Garry L, Handcock Mark S.* New specifications for exponential random graph models // *Sociological methodology*. 2006. 36, 1. 99–153.
- Strauss David, Ikeda Michael.* Pseudolikelihood estimation for social networks // *Journal of the American Statistical Association*. 1990. 85, 409. 204–212.
- Tang M.M.-J., Puah C.-H., Affendy Arip M., Dayang-Affizzah A.M.* Forecasting performance of the P-star model: The case of Indonesia // *Journal of International Business and Economics*. 2015. 15, 2. 7–12.
- Ter Wal Anne LJ.* The dynamics of the inventor network in German biotechnology: geographic proximity versus triadic closure // *Journal of Economic Geography*. 2013. lbs063.
- Wasserman Stanley, Pattison Philippa.* Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p // *Psychometrika*. 1996. 61, 3. 401–425.
- White Harrison C.* Identity and control: A structural theory of social action. 1992.

Cahiers du GREThA
Working papers of GREThA

GREThA UMR CNRS 5113

Université de Bordeaux

Avenue Léon Duguit
33608 PESSAC - FRANCE
Tel : +33 (0)5.56.84.25.75
Fax : +33 (0)5.56.84.86.47

<http://gretha.u-bordeaux.fr/>

Cahiers du GREThA (derniers numéros – last issues)

- 2016-08 : LORENZON Emmanuel, *Collusion with a Greedy Center in Position Auctions*
- 2016-09 : DELFAUD Pierre, *Retour sur une expérience unique d'interventionnisme en économie de marché: La Politique agricole commune européenne (1955-2015)*
- 2016-10 : MUTASCU Mihai, PEREAU Jean-Christophe, URSU Eugen : *A Wavelet Analysis of the Environmental Kuznets Curve in France*
- 2016-11 : LISSONI Francesco : *Migration and Innovation Diffusion: An Eclectic Survey*
- 2016-12 : BONIN Hubert : *Les paradoxes de l'économie maritime française au début du XXI^e siècle*
- 2016-13 : CASTIGLIONESI Fabio, NAVARRO Noémi : *(In)Efficient Interbank Networks*
- 2016-14 : PEREAU Jean-Christophe, MOUYSSSET Lauriane, DOYEN Luc : *Groundwater management in food security context*
- 2016-15 : FAURE Yves-André: *Institutions locales et résistances au test du VIH/sida. Quelques leçons d'une enquête dans la ville de Fortaleza, Brésil*
- 2016-16 : FAUVET Pierre : *Autorisation de mise sur le marché, défaillance de la responsabilité, et groupes de pression*
- 2016-17 : BECUWE Stéphane, BLANCHETON Bertrand : *French Textile Specialisation in Long Run Perspective (1836-1938) : Trade Policy as Industrial Policy*
- 2016-18 : LEVASSEUR Pierre: *The effects of bodyweight on wages in urban Mexico*
- 2016-19 : LUNG Yannick, JULLIEN Bernard, MIDLER Christophe : *De la Logan à la Kwid. Ambidextrie, innovation inversée et fractale, design-to-cost : les recettes de la stratégie Entry de Renault*
- 2016-20 : TROMEUR Eric, DOYEN Luc : *Optimal biodiversity erosion in multispecies fisheries*
- 2016-21 : FAURE Yves-A : *Institutions locales et résistances au test du VIH/sida. Quelques leçons d'une enquête dans la ville de Fortaleza, Brésil*

La coordination scientifique des Cahiers du GREThA est assurée par Emmanuel PETIT.

La mise en page est assurée par Julie VISSAGUET