G Model
SON-1040; No. of Pages 12
ARTICLE IN PRESS

Social Networks xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Social Networks

journal homepage: www.elsevier.com/locate/socnet

Change we can believe in: Comparing longitudinal network models on consistency, interpretability and predictive power

Per Block[a,*], Johan Koskinen[b], James Hollway[c], Christian Steglich[d], Christoph Stadtfeld[a]

[a] Chair of Social Networks, ETH Zürich, Switzerland
[b] The Mitchell Centre for SNA, and Social Statistics Discipline Area, University of Manchester, United Kingdom
[c] Department of International Relations/Political Science, Graduate Institute Geneva, Switzerland
[d] The Institute for Analytical Sociology, Linköping University, Sweden AND Department of Sociology/ICS, University of Groningen, The Netherlands

ARTICLE INFO

Article history:
Received 3 January 2017
Received in revised form 26 July 2017
Accepted 4 August 2017
Available online xxx

Keywords:
SAOM
ERGM
TERGM
Longitudinal network models
Network evolution
Out-of-sample testing
Parameter interpretation
Auto-regressive network models
Process-based network models

ABSTRACT

While several models for analysing longitudinal network data have been proposed, their main differences, especially regarding the treatment of time, have not been discussed extensively in the literature. However, differences in treatment of time strongly impact the conclusions that can be drawn from data. In this article we compare auto-regressive network models using the example of TERGMs – a temporal extensions of ERGMs – and process-based models using SAOMs as an example. We conclude that the TERGM has, in contrast to the ERGM, no consistent interpretation on tie-level probabilities, as well as no consistent interpretation on processes of network change. Further, parameters in the TERGM are strongly dependent on the interval length between two time-points. Neither limitation is true for process-based network models such as the SAOM. Finally, both compared models perform poorly in out-of-sample prediction compared to trivial predictive models.

1. Introduction

The study of social networks is increasingly concerned with modelling network change over time, as longitudinal analysis is usually better equipped for finding explanations and testing theories about the evolution of networks as well as the impact its structure has on constituent nodes (e.g. Steglich et al., 2010). Network analysis over time commonly uses network panel data: a network structure among the same set of nodes that is observed at two or more time points. By now (this is written in early 2017), several statistical approaches are available to analyse such data sets. The most widely used are the stochastic actor-oriented model (SAOM; Snijders et al., 2010b) and several extensions to the exponential random graph model (ERGM; Lusher et al., 2013). These models and variations may appear almost indistinguishable to scientists interested in applying inferential methods to network panel data. However, they rest on quite different statistical assumptions that strongly affect the kind of inference one can draw from the estimated model parameters and, thus, the kind of questions that can be answered with each method.

While statistical models can be compared on many dimensions, we mainly focus in this article on differences in how they treat time. In particular, we discuss the difference between discrete-time, auto-regressive models and continuous-time, process-based models. Due to its increased use (e.g. in McFarland et al., 2014) and recent claims about its advantage relative to other models for network panel data (Desmarais and Cranmer 2012; Leifeld and Cranmer 2016), we choose the TERGM (or temporal ERGM) for this comparison case to represent auto-regressive models.[1] The continuous-time model we discuss for comparison is the SAOM. Note that for ERGMs both continuous-time and auto-regressive extensions have been proposed – we focus on the latter group. The purpose is to compare the principles of auto-regressive and continuous-time network models and not the relative merits of either particular model – the two cases can be seen as represen-

* Corresponding author.
E-mail address: per.block@gess.ethz.ch (P. Block).

[1] We point out commonalities and differences between the TERGM (as defined in Desmarais and Cranmer 2012; Leifeld and Cranmer 2016) and other longitudinal variants of the ERGM where appropriate.

http://dx.doi.org/10.1016/j.socnet.2017.08.001

tations of their respective model classes. This article highlights, by way of illustration, the most important differences in assumptions and their interpretive implications between these approaches and thus facilitates the applied researcher's decision which to use in their own research.

1.1. Dimensions of comparison

When comparing statistical models it is tempting to ask which model is "better". However, "better" implicates at least two quite different dimensions: explanation and prediction. On the one hand, it has been argued that accurate prediction is a chief criterion of a "good" model (Friedman 1953; Jasso 1988). Intuitively a "good" model should be able to extrapolate accurately into the future, which can be tested for a single dataset by simple out-of-sample prediction. At the same time, the criterion for what should be predicted correctly in a model with dependent data (such as networks) is not trivial, as a network is more than just an series of independent tie observations but also the structures that these ties form (see discussion in Section 5).

On the other hand, it has been argued that the endeavour of social science is not to predict, but to explain and understand the world (Hedström 2005; Elster 2007). Models with absurd assumptions or intractable algorithms can generate fairly accurate predictions, but teach us little about the world. Social mechanisms, by contrast, can help us explain the social world and inform our understanding of our own and others' behaviour, but their concatenation in complex ways means that only in the simplest of systems can we expect this to result in accurate prediction at a micro-level. Indeed, even models with poor predictive power can generate valuable insights (see also Epstein 2008). In this line of reasoning, a good model is characterised by reasonable assumptions, as well as by clear interpretability of parameters in light of social mechanisms.

In this paper, we do not necessarily advocate for one or the other position, but investigate how different model assumptions make them applicable to different questions and thus to different empirical problems. As such, we elaborate what conclusions can be drawn from estimated parameters using the SAOM or the TERGM.

The remainder of the article is organised as follows. We first introduce the two different longitudinal/temporal network models (Section 2), and highlight their main features from a statistical point of view. The first main distinguishing feature of the model that is discussed concerns whether it is actor-oriented or tie-oriented (Section 3). Subsequently, the treatment of time is examined. Focus is on the interpretation of parameters and model consistency with regards to the differences between auto-regressive compared to process-based modelling (Section 4). The different treatment of time and how that influences parameters is shown in an empirical example. Finally, we demonstrate that both models perform poorly in out-of-sample prediction (Section 5) across two datasets, suggesting that we need to be careful as to the purposes of longitudinal network research.

2. The models

A social network needs to be understood as a system of interdependent units. Whether one is interested in the details of network dependencies or just needs to control for them, research on networks requires statistical tools that can adequately deal with this challenge. The model families that most explicitly deal with dependencies for such inferential-statistical analysis of social network data are *exponential random graph models* (ERGMs; Frank and Strauss, 1986; Pattison and Wasserman, 1999; Snijders et al., 2006; Lusher et al., 2013) and *stochastic actor-oriented models* for network evolution (SAOMs; Snijders 2001, 2005; Snijders et al., 2010b).[2]

2.1. Exponential random graph models

ERGMs were originally formulated for cross-sectional data, i.e., a single observation of a network. The guiding idea behind the model family is to express the probability of observing a given network as a function of subgraphs in this network (called statistics denoted $z(x)$ on the network x), e.g. the reciprocated dyad, or the transitive triplet. These subgraphs express local dependencies between tie variables (reciprocity and transitive clustering, respectively). At the heart of the ERGM lies a linear predictor that weighs the prevalence of statistics in the network by the parameter vector θ:

$$\sum_k \theta_k z_k(x)$$

What is considered local differs between model specifications, with the general rule that specifications including more complex subgraphs instantiate more dependency (Pattison and Snijders, 2013). Model parameters θ_k can be interpreted as expressing, on the tie-level, the probability of observing a specific tie, given the rest of the graph, or on the network-level, as indicating tendencies of a graph to exhibit certain sub-structures relative to what would be expected from a model not containing this parameter (this is discussed further in Section 4).

Longitudinal variants of the ERGM come in two forms, the continuous-time and the discrete-time version. The first, called longitudinal exponential random graph models (LERGMs; Snijders and Koskinen 2013; Koskinen et al., 2015), is a longitudinal, continuous-time form of the ERGM, in the sense that changes to the network are modelled using the conditional probabilities of the ERGM and the process has the cross-sectional ERGM as its limiting distribution. In its treatment of time, the LERGM is identical to the SAOM, thus, we do not focus on the LERGM in this article – the interested reader can generalise from our discussion.

The most prominent discrete-time variant of the ERGM is the temporal exponential random graph model (TERGM; Robins and Pattison 2001, Hanneke, Fu and Xing 2010; Desmarais and Cranmer 2012).[3] The model is based on the idea of panel regression. In a sequence of observations, lagged earlier observations or derived information thereof can be used as predictors for later observations. In other words, some of the statistics $z(x)$ are direct functions of an earlier realisation of the network. In its most basic form, the TERGM is a conditional ERGM with an earlier observation of the network occurring among the predictors. It is this basic TERGM (as presented in Desmarais and Cranmer 2012; Leifeld and Cranmer 2016) that we focus on in this article. While other statistics of a previous network realisation (e.g. past two paths) can be included in the model as predictors (e.g. to model transitivity over time), this does not change the fundamental challenges of parameter interpretability or time dependence of parameters modelling dependence as discussed in Section 4; consequently, we only deal with these extended specifications, when necessary, in footnotes. The interested reader can generalise.[4]

[2] There is a host of models that allow for dependent network ties (such as the p2 model, van Duijn et al., 2004; and an ever-expanding class of latent variable models, see for example the review by Salter-Townshend et al., 2012) that we do not discuss here.

[3] Given the limited space in one article, we do not discuss other discrete-time models, such as the StERGM (Krivitsky and Handcock 2014), even though they deserve a similar comparison elsewhere that might give different results.

[4] It should be noted that the TERGM might only include transformations of an earlier network as predictors of the network, as presented in Hanneke et al. (2010). In this case, all dependence between ties is assumed to be captured by the previous