Social Networks 37 (2014) 42–55

Contents lists available at ScienceDirect

Social Networks

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

Disaster response on September 11, 2001 through the lens of statistical network analysis

Michael Schweinberger[a,*], Miruna Petrescu-Prahova[b], Duy Q. Vu[c]

[a] Department of Statistics, Rice University, Houston, TX, USA
[b] Department of Health Services, University of Washington, Seattle, WA, USA
[c] Department of Mathematics and Statistics, University of Melbourne, Melbourne, Australia

ARTICLE INFO

Keywords:
Discrete exponential families
Hierarchical models
Mixture models
Model-based clustering
Social networks
Stochastic block models

ABSTRACT

The rescue and relief operations triggered by the September 11, 2001 attacks on the World Trade Center in New York City demanded collaboration among hundreds of organisations. To shed light on the response to the September 11, 2001 attacks and help to plan and prepare the response to future disasters, we study the inter-organisational network that emerged in response to the attacks. Studying the inter-organisational network can help to shed light on (1) whether some organisations dominated the inter-organisational network and facilitated communication and coordination of the disaster response; (2) whether the dominating organisations were supposed to coordinate disaster response or emerged as coordinators in the wake of the disaster; and (3) the degree of network redundancy and sensitivity of the inter-organisational network to disturbances following the initial disaster. We introduce a Bayesian framework which can answer the substantive questions of interest while being as simple and parsimonious as possible. The framework allows organisations to have varying propensities to collaborate, while taking covariates into account, and allows to assess whether the inter-organisational network had network redundancy—in the form of transitivity—by using a test which may be regarded as a Bayesian score test. We discuss implications in terms of disaster management.

1. Introduction

Large-scale disasters can strike without warning, leaving death and destruction behind. Examples are natural disasters, such as earthquakes and tsunamis, and man-made disasters, such as industrial accidents and terrorist attacks (e.g., Topper and Carley, 1999; Butts et al., 2007; Petrescu-Prahova and Butts, 2008). When disasters strike, governmental as well as private organisations converge at the scene of the disaster to conduct rescue and relief operations (Auf der Heide, 1989). The scale of operations may overtax the resources of individual organisations and compel them to collaborate (Haas and Drabek, 1973). In addition, organisations may be forced to collaborate in order to coordinate operations. As a result, inter-organisational networks of collaboration tend to emerge in response to disasters. Studying the topological structure of inter-organisational networks is important with a view to understanding the weaknesses and strengths of the response to past disasters and planning and preparing the response to future disasters.

We consider here data on the massive rescue and relief operations triggered by the September 11, 2001 attacks on the World Trade Center in New York City (Tierney and Trainor, 2004; Bevc, 2010). The network of collaborations between the $n = 717$ organisations which responded to the attacks is interesting for at least three reasons. First, an important question is whether some organisations were more involved in the disaster response than others. The presence of such organisations is considered critical to the disaster response, because in the unstructured and changing environment in which rescue and relief operations are conducted such organisations facilitate communication and coordination of the disaster response (e.g., Auf der Heide, 1989). Second, if such organisations were present, then the question is whether those organisations were supposed to coordinate disaster response or emerged as coordinators in the wake of the disaster. The possible presence of emergent coordinators has implications in terms of resource allocation to responders (e.g., Petrescu-Prahova and Butts, 2008). Third, while the presence of established and emergent coordinators can improve the disaster response, it can make the disaster response more vulnerable to disturbances which follow the initial disaster and can disrupt the disaster response (e.g., Topper and Carley, 1999). Such disturbances are not uncommon, e.g., earthquakes may be followed by aftershocks, and terrorists may attack responders on purpose by first detonating explosive devices to attract by-standers and responders and then detonating more explosive devices to harm them (see, e.g., the Bali bombing in 2002 and the Madrid

* Corresponding author. Tel.: +1 713 348 2278; fax: +1 713 348 5476.
E-mail address: m.s@rice.edu (M. Schweinberger).

0378-8733/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.socnet.2013.12.001

bombing in 2004, Koschade, 2006). In particular, the impact of the second airplane into the South Tower of the World Trade Center on September 11, 2001 disrupted the response to the September 11, 2001 attacks. Therefore, it is important to assess how much network redundancy inter-organisational networks possess and how sensitive inter-organisational networks are to disturbances.

These questions can be framed as questions about the organisations and the topological structure of the network of organisations, first and foremost the propensities of organisations to collaborate as well as the redundancy of the network in the form of transitivity, while taking available covariates into account.

A simple approach to answering the first question—whether some organisations were more involved in the disaster response than others—is based on clustering organisations by using deterministic clustering methods in social network analysis (e.g., Borgatti et al., 2002), physics (e.g., Girvan and Newman, 2002), and computer science (e.g., Leskovec et al., 2008). However, deterministic clustering methods are not model-based and ignore the uncertainty about the clustering; neglect covariates and all other features of the network (e.g., transitivity); and lack desirable statistical properties: e.g., Bickel and Chen (2009) showed that the Girvan–Newman modularity (Girvan and Newman, 2002)—which has attracted much attention—is not consistent.

An alternative approach to answering the first question is based on modeling the sequence of degrees, corresponding to the numbers of collaborations of organisations (e.g., Erdös and Rényi, 1959; Barabàsi and Albert, 1999; Newman et al., 2001; Jones and Handcock, 2004). However, such one- and two-parameter models are not flexible models of the sequence of degrees and may not be able to account for the heterogeneity of the degrees. In addition, such models ignore covariates and all other features of networks (e.g., transitivity).

A second, alternative approach to answering the first question is based on mixture models and other latent variable models, such as stochastic block models (e.g., Nowicki and Snijders, 2001; Tallberg, 2005; Airoldi et al., 2008), random effect models (e.g., van Duijn et al., 2004; Hoff, 2005), and latent space models (e.g., Hoff et al., 2002; Schweinberger and Snijders, 2003; Handcock et al., 2007). However, some of those models make strong model assumptions about the data: e.g., latent space models assume that organisations are embedded in a latent, metric space, which is a strong model assumption and is not needed to answer the substantive questions of interest. In addition, the computing time of most of them scales with n^2. Indeed, most of the statistical software which implements these statistical models and methods is either not publicly available (e.g., Tallberg, 2005) or cannot handle networks with $n \gg 200$: e.g., neither the statistical software BLOCKS (Snijders and Nowicki, 2007) implementing the stochastic block models of Nowicki and Snijders (2001) nor the `R` package `lda` (Chang, 2011) implementing the extended stochastic block models of Airoldi et al. (2008) can handle networks with $n \gg 200$ organisations.

A third, alternative approach to answering the first question is based on exponential-family models with so-called *k*-stars as sufficient statistics (Frank and Strauss, 1986) and curved exponential-family models with so-called geometrically weighted degree terms (Hunter and Handcock, 2006; Snijders et al., 2006; Hunter, 2007). However, these exponential-family models are, at present, not well-understood and sometimes possess undesirable properties, e.g., model degeneracy, and in accordance have attracted much criticism (Strauss, 1986; Jonasson, 1999; Snijders, 2002; Handcock, 2003; Bhamidi et al., 2008; Rinaldo et al., 2009; Butts, 2011; Schweinberger, 2011; Chatterjee and Diaconis, 2013).

Last, an approach to answering the first and the third question—whether there is more network redundancy than expected—is based on clustering coefficients (e.g., Kolaczyk, 2009, pp. 96–97), though clustering coefficients neither allow to cluster organisations based on the propensity to collaborate nor allow to test whether there is more network redundancy than expected.

Table 1
Scale of operations and type of 717 organisations.

	Collective	Governmental	Non-profit	Profit	Total
Local	3	86	67	68	224
State	5	67	6	9	87
National	23	157	32	88	300
International	1	8	20	77	106
Total	32	318	125	242	717

To overcome the shortcomings of existing approaches, we introduce a Bayesian exponential-family framework which helps to answer the substantive questions of interest and is as simple and parsimonious as possible. The framework allows us to incorporate covariates as well as variation in the propensities of organisations to collaborate. In addition, the framework allows us to test whether there is network redundancy in the form of transitivity by using a test which can be considered to be a Bayesian analogue of the score test (Rao and Poti, 1946; Rao, 1948; Bera and Bilias, 2001a,b). The advantage of using a Bayesian score test is that models with transitivity—which are hard to estimate—need not be estimated, while the question of primary interest—whether there is an excess of transitivity—can be answered. An additional advantage of the framework is that the computing time scales with n rather than n^2 and thus the framework can be applied to the inter-organisational network with $n = 717$.

The paper is structured as follows. We describe the data and the substantive questions of interest in Section 2 and introduce Bayesian models in Section 3. We discuss Bayesian model estimation and model selection in Section 4. We shed light on the substantive questions of interest by using the Bayesian framework in Section 5 and discuss implications in terms of disaster management in Section 6.

2. Data

2.1. The inter-organisational network

On September 11, 2001, the United States experienced one of the deadliest and costliest events in its history: The attacks on the World Trade Center in New York City killed 2749 individuals and resulted in immeasurable psychological, social, and economic costs (Tierney and Trainor, 2004; Bevc, 2010).

We consider here data on the massive rescue and relief operations triggered by the September 11, 2001 attacks. A detailed description of the data can be found in Tierney and Trainor (2004) and Bevc (2010). The data were collected during the first 12 days following September 11, 2001 and represent the rescue and relief phase of operations as well as the transition into the less communication-intensive recovery phase of operations. The resulting data set includes 717 organisations and two organisational attributes: scale of operations, which includes the categories of local, state, national, and international; and type of organisation, which includes the categories of collective, governmental, non-profit, and profit. Table 1 reports the number of organisations with these two attributes. Data on the $\binom{717}{2} = 256,686$ (possible) collaborations among the 717 organisations were collected by content analysis of field documents and newspaper articles. The collaborations are binary—i.e., either absent or present—and undirected. Fig. 1 shows the collaborations among the 717 organisations.