

A Framework for Longitudinal Influence Measurement between Communication Content and Social Networks

Shenghui Wang and Paul Groth

VU University Amsterdam

De Boelelaan 1081a, Amsterdam, The Netherlands

{s.wang, p.t.groth}@vu.nl

Abstract

Artificial intelligence has a long history of learning from domain problems ranging from chess to jeopardy. In this work, we look at a problem stemming from social science, namely, how do social relationships influence communication content and vice versa. The tools used to study communication content (content analysis) have rarely been combined with those used to study social relationships (social network analysis). Furthermore, there is even less work addressing the longitudinal characteristics of such a combination. This paper presents a general framework for measuring the dynamic bi-directional influence between communication content and social networks. The framework leverages the idea that knowledge about both kinds of networks can be represented using the same knowledge representation. In particular, through the use of Semantic Web standards, the extraction of networks is made easier. The framework is applied to two use-cases: online forum discussions and conference publications. The results provide a new perspective over the dynamics involving both social networks and communication content.

1 Introduction

Does an informative post on a microblogging service lead to a user gaining followers? If a user is popular in a social network, will their new status updates be widely quoted? If a researcher identifies a new topic one year, does that result in the research having more coauthors the next? As an increasing amount of content is mediated through social networks, these types of questions are of great interest, in particular, to developers, social scientists, and businesses that aim to understand the link between content generation and social connection. A key aspect to answering these questions is to understand how the relationships between users influence the content of their communication and vice versa.

In this paper, we summarize our work [Wang and Groth, 2010], which extended our previous work [Oegema *et al.*, 2010] by proposing a general framework for measuring such influence over time. In our approach, we translate both user relationships and content into two corresponding networks:

a social network and a content network. The networks are then characterized using common network properties such as (in-/out-)degree and betweenness centrality. The influence is then measured using a set of multilevel time-series regression models producing what we term an influence network showing how these variables impact each other in time. Additionally, our Influence Framework can integrate other network properties tailored to a given problem domain. We show how the Influence Framework can be applied to networks obtained from a corpus of conference information and networks extracted from a Dutch political forum.

Importantly, the use of Semantic Web standards simplifies the extraction interrelated networks as these knowledge representation standards enable explicit connections between content and the users who exchange it. The ability to study the connection between people through their objects was posited as a key benefit to using the Semantic Web in conjunction with social networks [Bojars *et al.*, 2008]. This work is an example of where these benefits are coming to fruition.

We describe three core contributions of our work:

- A general framework for measuring the bi-directional influence between networks of people and the content associated with those people.
- A multilevel time-series regression model for measuring the longitudinal influences between the network properties of content and social networks.
- The generation of influence networks for both Dutch political forums and the World Wide Web conference series, which provide new material for social scientists to investigate these domains.

The rest of this paper is organized as follows. We begin by presenting the Influence Framework and its constituent parts. This is followed by a discussion of the application of the Framework to two use cases: one studying a conference series and the other studying data from a Dutch political forum. Related work is then discussed, which is followed by a conclusion.

2 Influence Framework

The Influence Framework is a three stage framework for measuring the influence between (and within) user relationships and the content they communicate. While such measures of

influence are clearly possible to perform on a case-by-case basis, a key realization in this work is that by representing content and user relationships as networks, standard network properties can provide a good initial insight into influence in different domains. We note that influence is a time-dependent notion and thus our framework requires time series data.

The three stages of the framework are:

1. Network Generation
2. Network Property Measurement
3. Time Series Analysis

We now discuss each of these stages.

2.1 Network Generation

The first stage of the framework is to generate a series of both content and social networks as well as bindings between those networks. The starting point is information about a set of actors who interact over time, *e.g.*, participants in online discussions, scientists who co-author, *etc.* From these data sets, a *series of social networks* representing the interaction of these actors over time can be produced. Then, a corpus of content related to each actor, over time, is needed *e.g.*, the textual content of online discussions a participant posted, the abstract a scientist wrote, *etc.* This content corpus should also have the property that pieces of content are stable across a group of actors. Based on some relations between content at each time step, a *series of content networks* can be generated. A key artifact for the framework is documentation of the relationship between actors and the content they produce at each time step. We term these relationships *bindings*.

The network generation stage is perhaps the most domain specific part of the framework as a decision must be made about which content and which sort of user relationship should be represented in the network. Furthermore, many domains have different data formats requiring specialized programs to generate the needed networks. This is where Social Semantic Web technologies are particularly important. By providing common query interfaces and data representations, the extraction of these networks is made easier.

2.2 Measuring Network Properties

Once the content networks and social networks have been produced, the properties of those networks that are of interest need to be defined (as variables) and then measured. Standard network properties (*e.g.*, various kinds of centrality in [Scott, 2000]) can be measured as initial choices. It is important to note that while these network properties can be measured for every graph, their underlying meaning with respect to the social reality needs to be defined on a per domain basis.

While these measures are a useful start, any network property that varies over time is allowable within the Influence Framework. Later in Section 3.3, we show how other domain specific network properties can be used to gain additional insight into the influence between content and social networks.

The output of this stage is a table mapping each actor to values for each property at each time step. This forms a *time series*, that is, data at successive time steps spaced at uniform time intervals. Next, we will introduce a *time series analysis*

to extract meaningful statistics of the data in order to better understand the underlying forces and structures that produced the observed data.

2.3 Multilevel time-series regression models

When modeling variations in the level of a process, one of the typical methods is to use the *autoregressive* (AR) models. Let \mathbf{X} be a time series: $\mathbf{X} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots\}$, where $\mathbf{x}^{(t)}$ is the data observation at time t . Here, $\mathbf{x}^{(t)}$ is a vector, *i.e.*, $\mathbf{x}^{(t)} = (x_1^{(t)}, x_2^{(t)}, \dots, x_m^{(t)})^T$, where m is the total number of variables we are modelling and each $x_i^{(t)}$, $i = 1, \dots, m$, is a variable we are interested in, such as the betweenness and degree centrality of a node in the social network or the centrality values of certain political or scientific topics. In this paper, we opt for a simple model for each variable x_i independently, which only includes the values from the last time-point as independent variables, *i.e.*, an AR(1)-process:

$$x_i^{(t)} = a_i + b_{1i} x_1^{(t-1)} + \dots + b_{mi} x_m^{(t-1)} + \varepsilon_i^{(t)}, \quad (1)$$

where a_i is a constant and $\varepsilon_i^{(t)}$ is Gaussian noise with zero mean and variance σ_ε^2 . The influence coefficients or *effects* $\{b_{1i}, \dots, b_{mi}\}$ quantify how variation in the predictor variable at time $t - 1$ is related to the variation of the predicted variable at time t , *i.e.*, the influence among different variables over time.

Generally, the above mentioned variables are referred to in statistics as *units of analysis*. In social reality, these variables are often from different levels, which are frequently hierarchically nested. For example, when studying research achievements, attributes of individual researchers, research groups, faculties and the universities as a whole can all be important units of measures. This stage applies the above introduced regressive model to study the influence between variables. The resulting coefficients are also called *fixed effects*. However, there exist variations among different actors, *i.e.*, *random effects* (actor-level errors). Therefore, such single-level statistical methods are no longer appropriate to study these so-called *complex data sets* [Snijders and Bosker, 1999]. We, thus, need to apply *multilevel analysis* to examine both fixed and random effects of variables measured at different levels [Hayes, 2006; Snijders and Bosker, 1999].

Formally, we define $\mathbf{x}_p^{(t)} = (x_{1,p}^{(t)}, \dots, x_{m,p}^{(t)})^T$, a vector containing the variables for actor p at time t . We can then rewrite (1) as

$$x_{i,p}^{(t)} = a_i + \mathbf{b}_i^T \mathbf{x}_p^{(t-1)} + \varepsilon_i^{(t)} + \mathbf{c}_{i,p}^T \mathbf{x}_p^{(t-1)} + \varepsilon_{i,p}^{(t)}, \quad (2)$$

where $\mathbf{b}_i = (b_{i1}, \dots, b_{im})^T$ and $\mathbf{c}_i = (c_{i1}, \dots, c_{im})^T$ are the fixed-effect coefficients and random-effects coefficients respectively.

The output of this stage is the set of statistics generated as a result of fitting the regression models as well as a diagram, called an *influence network*, that shows the statistically significant effects between variables.

3 Use Cases

3.1 Influence between co-authors of academic papers and the topics they address

Data Collection

We obtained a corpus of paper metadata about the World Wide Web conference from the Semantic Web Dogfood repository [Möller *et al.*, 2007]. The corpus spans four years of the conference from 2007 to 2010 using generally the same schema.

Generating social networks

We chose the co-author network as the social network of interest. For every year, we retrieved the co-author pairs for each article using a simple query database query. From these results, we built a weighted undirected graph for each year where nodes are authors, edges are shared authorship of an article and the weights on edges are the number of co-authorships between the two linked authors.

For each year, we measured the degree and betweenness centrality of each author. The degree centrality represents how active the author is in coauthoring with others, while the betweenness centrality indicates the author's role in connecting different authors. We also measure clustering coefficient which provides a measure of whether authors write with the same set of other authors.

Generating content networks

Here, we are interested in the topics under discussion at the conference each year. We use author assigned keywords as proxies for topics. Similar to the co-author network, we retrieved the keywords for each article in the conference again using a database query. To improve overlap between keywords assigned by different authors, keywords containing more than one word were split into separate words and then stemmed. Therefore, a weighted undirected graph is built for each year, where a node is a keyword and an edge is the co-occurrence between two keywords in the same articles. Edges are weighted by the number of co-occurrences.

We then compute several standard network properties. Again, the degree provides information about the popularity of a given topic. The betweenness centrality provides information about whether a keyword is a bridge between other keywords (*i.e.* topics).

Binding social content networks

We bind the two networks together via the papers within the conference. Thus, we know which author discusses a topic and what topics are associated with particular authors via their connection to papers.

Influence Network

For this use case, we study five network properties.

- Three social network properties: degree centrality, betweenness centrality, and clustering coefficient.
- Two content-wise properties: degree centrality, betweenness centrality.

The units of analysis are all year \times participant combinations. The multilevel time-series regression models are then




Figure 1: Influence network for WWW conference

constructed to study the influence network between topics of a conference and the co-authorship of papers. Figure 1 shows the resulting influence network. This network only shows effects which are statistically significant. Note, when reading such an influence network, the edges are directional in time. For example, in Figure 1, the edge between degree in the content network and clustering coefficient in the social network, should be read as the degree at some time t has large negative effect on the clustering coefficient in time $t + 1$.

The network suggests a number of avenues for investigation. First, there is strong negative effect between the degree centrality of a topic (*i.e.*, keyword) on itself, which suggests that a popular topic one year is likely to be less popular the next. Degree centrality of a topic also has strong negative effects on the degree centrality and clustering coefficient for an author. One interpretation of this result is that after a burst of collaboration on a hot topic, the topic becomes less exciting and the collaboration between authors around it dies down. There are strong positive effects of the betweenness centrality of a topic and the subsequent degree centrality and clustering coefficient of an author. A possible explanation for these effects is that if a topic bridges the gap between other topics in one conference year, it is likely to become the focus for new collaborations between authors concentrating on these normally separate topics. Such new collaborations would then come to the foreground in the next conference year.

3.2 Influence between social status of online forum participants and their political attention

Data collection

Our data is collected from one of the oldest Dutch forums, NL.politiek, which is entirely devoted to politics. Our dataset contains all the postings from October 2003 to December 2008, in total more than 1.1 million postings.

Generating social networks

All postings were divided into weekly subsets. In each subset, all postings were grouped by their threads. Within each thread, we record the replying action as the directed links between participants. Therefore, the social networks are formed as follows: each node represents a participant and there is a directed link between two participants if one replied to another in one thread, where the weights on the links indicate the number of such replying action within a thread. We




Figure 2: Standard influence network

then aggregated all the threads within one week into a bigger network, producing a series of 259 weekly social networks where 21,127 participants are involved.

For each week, we measured the in/outdegree and betweenness centrality of all participants. Here, the indegree centrality indicates the degree of *popularity* one participant has in the online community, while the outdegree centrality indicates how active he is. The betweenness centrality is an indicator of the *mediating/brokerage* role of a participant.

Generating content networks

In this use case, we are interested in the attention to the political parties that online participants have when they discuss in the forum. We thus extract the co-occurrence of parties as the content network. In the content network, the nodes are 19 Dutch parties and the edge indicates the two linked parties are mentioned in the same postings. The weight of the edge is calculated as the Jaccard similarity coefficient between two sets of postings which mentioned two parties individually. In this way, we also extracted 259 weekly content networks.

We then measured the betweenness and degree centrality of each party in each week. These centrality can tell us how one party's popularity and breakage role evolves over time. When a party has a higher degree centrality, then this party is more often mentioned while other parties are being discussed, *i.e.*, this party is more relevant or important. A party with a higher betweenness centrality is more often mentioned as a reference while more than two parties are mentioned.

Binding social and content networks

We bind two networks based on *who talked when, about what*. For each participant, we sum up the corresponding centrality values of the parties which he mentioned in his postings. Therefore, if one participant mentioned an important party often, then his degree centrality in terms of his discussion content is high.

Influence network

Similar to the conference case, we have five standard network variables to model:

- Three social network properties: in and out degree centrality and betweenness centrality
- Two content-wise properties: betweenness and degree centrality

As shown in Figure 2, the in/out degree and betweenness centrality have positive effects upon each other and to themselves. Looking at the effects between social network and the content network, the indegree centrality (*i.e.*, the popularity of a participant) has a positive effect on the degree centrality

of the content. This suggests that when a popular participant talks about certain parties, these parties are likely to become popular in the next week. When a participant becomes a broker, they tend to communicate with different opinion-holders, therefore they discuss more parties instead of only popular ones. This might be the reason for the negative effect from the social betweenness centrality to the content degree centrality. However, this needs to be further investigated.

3.3 Influence between user-defined content variables with social network properties

Content networks can be extracted in a manner that is more suitable to specific problems within a domain. Here we show how user-defined content variables can be studied in the same framework.

Extracting specific content variables

In this paper, we focus on two aspects related to the forum content. The first aspect is related to the agenda setting [Severin and Tankard, 2010]. We are interested in whether the social status of the participants is influenced by the extent to which they follow mass media. Therefore, we use a list of political issues and measure the weekly attention to these issues (the frequencies of occurrence of these issues) in forum postings and the newspaper articles from five biggest Dutch national newspaper, respectively. Then a correlation is calculated between these two lists of the attention, which gives the first content variable *NewspaperContagion*. A higher *NewspaperContagion* indicates that the participant more strongly follows the agenda of the newspapers.

Another interesting aspect is the emotion expressed in the forum discussions. We check whether the amount of emotion expressed in the online discussion influences the social status of the participants and his willingness to following the mass media. Similar to measuring the attention to political parties, the frequencies of occurrence of a list of emotional keywords were measured. We separated the emotion of disgust and hate as a separate variable as they are the major emotions the communication scientists are studying [Oegema *et al.*, 2008]. Therefore, we have two other content variables: *DisgustHate* and *OtherEmotions*.

Influence network

The five variables we investigate are

- Two network properties: *Popularity* (=indegree centrality), *Activity* (=outdegree centrality) and *Betweenness* (betweenness centrality)
- Three communication contents: *DisgustHate*, *OtherEmotions* and *NewspaperContagion*.

Figure 3 shows that both popularity and activity are involved in a positive feedback loop, in a spiral of centrality in the social network and the expression of disgust and hate. It suggests that central participants feel unhindered or even obliged to use rather crude words to maintain their position. We also find that the mass media agenda impacts the agendas of popular and active participants. Because their popularity is based on the mass media's agenda, that inspires them to continue to follow that agenda, which apparently helps them to become more popular and active, and so on.

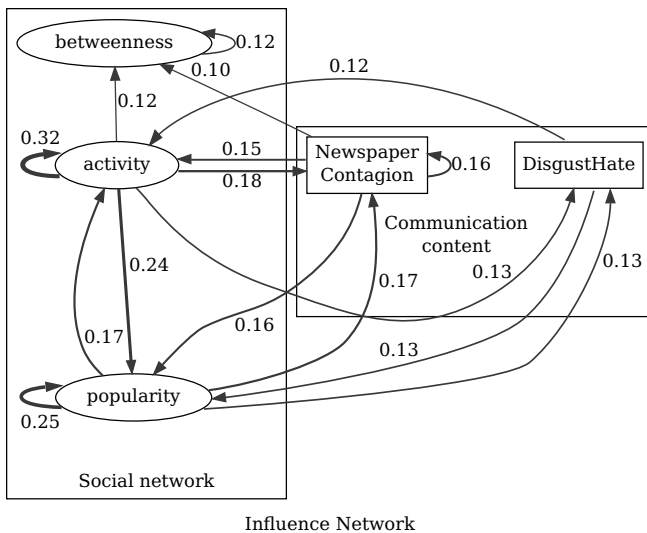


Figure 3: Regression model of user-defined variables

4 Related work

Social network analysis (SNA) has recently become a popular topic of study in organisation studies, communication studies, information science, *etc.* It views social relationships in terms of network theory consisting of nodes and ties. Using graph algorithms, SNA characterises the structure of social networks, strategic positions in these networks, specific sub-networks and decompositions of people and activities [Scott, 2000]. SNA has been applied not only to Web 2.0 platforms such as Facebook [Ackland, 2009] and wikis [Tomasev and Mladenec, 2009], but also directly to the whole Web, the blogosphere, ontologies and the Semantic Web [Jamali and Abolhassani, 2006; Kim *et al.*, 2005; Hoser *et al.*, 2006]. Recently, Semantic Web techniques have been adopted to facilitate standard SNA procedures [Mika, 2005; Martin and Gutierrez, 2009; Ereteo *et al.*, 2009].

On the other hand, content analysis is a research tool which has been used since the mid-1950's to determine the presence of certain words or concepts within texts [Krippendorff, 2003]. By quantifying and analysing the presence, meanings and relations of such words and concepts, social scientists can make inferences about the content of the texts. As it is applicable to any piece of writing or recorded communication, it has been widely used in many fields, such as media studies, literature, sociology and political science [Holsti, 1969; Budge *et al.*, 2001; Wimmer and Dominick, 2005]. Recently, many efforts have been focused on automated content analysis, such as [van Atteveldt *et al.*, 2008], which to a large degree improves the access to large corpora.

These two classes of analysis have been investigated and applied in a rather parallel style. Only until recently, social scientists started to combine social network analysis and content analysis, such as the *discourse network analysis* in [Leifeld and Haunss, 2010], and the work in [oliver and Montgomery, 2008].

Another focus of our paper is on the longitudinal analysis over content and social networks. Recognized as a *Holy Grail*

for network researchers, there has been a large degree of focus on the analysis of social networks over time [McCulloh and Carley, 2009]. However, there has not been much work with respect to the longitudinal analysis on the combination of social and content networks. The closest work is that of Gloor *et al.*, who use network analysis over social networks and corresponding content to identify trends, however, they concentrate on a time dependent betweenness measure and do not provide a general framework for a variety of network properties [Gloor *et al.*, 2009].

5 Conclusion

In this paper, we presented a general framework for analyzing the dynamic bi-directional influence between social relationships and the content produced with respect to those relationships. The Influence Framework leverages a key insight that by representing both social relationships and content as networks, common network properties can be used to bootstrap the analysis of influence. Based on these properties, the framework applies a time-series regression model to generate influence network diagrams representing the statistically significant effects of these properties. We applied our framework to two domains, dutch politics and a conference series, resulting in interesting conclusions about the influence of media on political forum participants and the impact of topics on academic collaboration. The data was acquired from both a web crawl and a Semantic Web source. We note that to acquire the data from the Semantic Web source required a simple database query whereas the Web Crawl required significant preprocessing.

By linking across both content and social networks, the Semantic Web is providing a new data source for understanding the relationship between users and the content that they produce [Bojrs *et al.*, 2008]. The framework described in this paper provides a new tool for analyzing these relationships from a longitudinal perspective.

References

- [Ackland, 2009] R. Ackland. Social network services as data sources and platforms for e-researching social networks. *Social Science Computer Review*, 27:481–492, 2009.
- [Bojars *et al.*, 2008] Uldis Bojars, John G. Breslin, Vassilios Peristeras, Giovanni Tummarello, and Stefan Decker. Interlinking the social web with semantics. *IEEE Intelligent Systems*, 23:29–40, 2008.
- [Bojrs *et al.*, 2008] U. Bojrs, J. G. Breslin, A. Finn, and S. Decker. Using the semantic web for linking and reusing data across web 2.0 communities. *Web Semant.*, 6(1):21–28, 2008.
- [Budge *et al.*, 2001] Ian Budge, Hans-Dieter Klingemann, Andrea Volkens, Judith Bara, and Eric Tanenbaum. *Mapping Policy Preferences. Estimates for Parties, Electors and Governments 1945-1998*. Oxford University Press, Oxford, 2001.
- [Ereteo *et al.*, 2009] Guillaume Ereteo, Michel Buffa, Fabien Gandon, and Olivier Corby. Analysis of a real online social network using semantic web frameworks. In

- 8th International Semantic Web Conference (ISWC2009), October 2009.
- [Gloor *et al.*, 2009] Peter A. Gloor, Jonas Krauss, Stefan Nann, Kai Fischbach, and Detlef Schoder. Web science 2.0: Identifying trends through semantic social network analysis. In *CSE (4)*, pages 215–222. IEEE Computer Society, 2009.
- [Hayes, 2006] Andrew F. Hayes. A Primer on Multilevel Modeling. *Human Communication Research*, 4:385–410, October 2006.
- [Holsti, 1969] Ole R. Holsti. *Content Analysis for the Social Sciences and Humanities*. Addison-Wesley, Reading, MA, 1969.
- [Hoser *et al.*, 2006] Bettina Hoser, Andreas Hotho, Robert Jäschke, Christoph Schmitz, and Gerd Stumme. Semantic network analysis of ontologies. In *Proceedings of 3rd European Semantic Web Conference*, pages 514–529, 2006.
- [Jamali and Abolhassani, 2006] M. Jamali and H. Abolhassani. Different aspects of social network analysis. In *IEEE/WIC/ACM International Conference on Web Intelligence*, pages 66 – 72, Hong Kong, 2006.
- [Kim *et al.*, 2005] Henry M. Kim, Markus Biehl, and John A. Buzacott. M-ci2: Modelling cyber interdependencies between critical infrastructures. In *Proceedings of 3rd IEEE International Conference on Industrial Informatics*, pages 644–648, 2005.
- [Krippendorff, 2003] Dr. Klaus H. Krippendorff. *Content Analysis: An Introduction to Its Methodology*. Sage Publications, Inc, 2003.
- [Leifeld and Haunss, 2010] Philip Leifeld and Sebastian Haunss. A comparison between political claims analysis and discourse network analysis: The case of software patents in the european union. In *MPI Collective Goods Preprint*, 2010/21, May 2010.
- [Martin and Gutierrez, 2009] M. S. Martin and C. Gutierrez. Representing, querying and transforming social networks with rdf/sparql. In L. Aroyo, P. Traverso, and F. Ciravegna, editors, *Semantic Web: Research and Applications*, pages 293–307, 2009.
- [McCulloh and Carley, 2009] Ian McCulloh and Kathleen Carley. Longitudinal dynamic network analysis, using the over time viewer feature in ora. Technical report, Institute for Software Research, School of Computer Science, Carnegie Mellon University, 2009.
- [Mika, 2005] Peter Mika. Flink: Semantic web technology for the extraction and analysis of social networks. *Journal of Web Semantics*, 3:211–223, 2005.
- [Möller *et al.*, 2007] Knud Möller, Tom Heath, Siegfried Handschuh, and John Domingue. Recipes for semantic web dog food: the eswc and iswc metadata projects. In *ISWC'07/ASWC'07: Proceedings of the 6th international The semantic web and 2nd Asian conference on Asian semantic web conference*, pages 802–815, Berlin, Heidelberg, 2007. Springer-Verlag.
- [Oegema *et al.*, 2008] D. Oegema, J. Kleinnijenhuis, K. Anderson, and A.M.J. Van Hoof. Flaming and blaming: The influence of mass media content on interactions in on-line discussions. In E.A. Konijn, M. Tanis, and S. Utz, editors, *Mediated Interpersonal Communication*. Mahwah: Erlbaum, 2008.
- [Oegema *et al.*, 2010] Dirk Oegema, Shenghui Wang, and Jan Kleinnijenhuis. Dynamics of online discussions about politics: a function of structural network properties, mass media attention or emotional utterances? In *Proceedings of the WebSci10: Extending the Frontiers of Society On-Line*, Raleigh, NC: US, April 2010.
- [oliver and Montgomery, 2008] Amalya L. oliver and Kathleen Montgomery. Using field-configuring events for sense-making: A cognitive network approach. *Journal of Management Studies*, 45:1147–1167, 2008.
- [Scott, 2000] John Scott. *Social Network Analysis: A Handbook. 2nd Ed.* Newberry Park, CA: Sage, 2000.
- [Severin and Tankard, 2010] W.J. Severin and J.W. Tankard. *Communication theories*. New York: Pearson, 2010.
- [Snijders and Bosker, 1999] T. A. B. Snijders and Roel J. Bosker. *Multilevel analysis: an introduction to basic and advanced multilevel modeling*. SAGE, 1999.
- [Tomasev and Mladenic, 2009] N. Tomasev and D. Mladenic. Semantic web wiki: Social network analysis of page editing. In V. LuzarStiffler, I. Jarec, and Z. Bekic, editors, *Proceedings of the Iti 2009 31st International Conference on Information Technology Interfaces*, pages 505–510, 2009.
- [van Atteveldt *et al.*, 2008] Wouter van Atteveldt, Jan Kleinnijenhuis, and Nel Ruigrok. Parsing, semantic networks, and political authority using syntactic analysis to extract semantic relations from dutch newspaper articles. *Political Analysis*, 16(4):428–446, 2008.
- [Wang and Groth, 2010] Shenghui Wang and Paul Groth. Measuring the dynamic bi-directional influence between content and social networks. In *9th International Semantic Web Conference (ISWC2010)*, November 2010.
- [Wimmer and Dominick, 2005] Roger D. Wimmer and Joseph R. Dominick. *Mass Media Research: An Introduction. 8th ed.* Belmont, CA: Wadsworth, 2005.