Detecting Communities through Network Data

**Jeroen Bruggeman,[a] V. A. Traag,[b]
and Justus Uitermark[c]**

Abstract

Social life coalesces into communities through cooperation and conflict. As a case in point, Shwed and Bearman (2010) studied consensus and contention in scientific communities. They used a sophisticated modularity method to detect communities on the basis of scientific citations, which they then interpreted as directed positive network ties. They assumed that a lack of citations implies disagreement. Some scientific citations, however, are contentious and should therefore be represented by negative ties, like conflicting relations in general. After expanding the modularity method to incorporate negative ties, we show that a small proportion of negative ties, commonly present in science, is sufficient to significantly alter the community structure. In addition, our research suggests that without distinguishing negative ties, scientific communities actually represent specialized subfields, not contentious groups. Finally, we cast doubt on the assumption that lack of cites would signal disagreement. To show the general importance of discerning negative ties for understanding conflict and its impact on communities, we also analyze a public debate.

Social life clusters into many kinds of groups, a prime topic in the social sciences. Groups often coalesce around common activities, or more generally around social foci (Feld 1981; Kossinets and Watts 2009). Exchanges of information and resources are more frequent within than between groups, which tend to be connected by relatively weaker ties (Granovetter 1973). Some groups may have conflicting relations between them. Conflicts can also exist within groups, but if conflict escalates, groups typically split into opposing factions. In a given population, when neither groups nor group memberships are known beforehand, both can be inferred from social network data. Today, to detect groups or communities in networks, researchers typically use modularity optimization (Fortunato 2010; Reichardt

2009), a method that builds on block modeling (White, Boorman, and Breiger 1976).

Network analysts typically assume ties are positive, even though they know not all social relations are positive. Science, for instance, is characterized by cooperation and benign disagreement, but also by epistemic rivalry. In democratic politics, disagreement with

[a]Universiteit van Amsterdam
[b]Université Catholique de Louvain
[c]Erasmus University Rotterdam

Corresponding Author:
Jeroen Bruggeman, Department of Sociology and Anthropology, Universiteit van Amsterdam Oudezijds Achterburgwal 185, 1012 DK Amsterdam, Netherlands
E-mail: j.p.bruggeman@uva.nl

opponents is endemic as it is vital to political identity and to attract voters. Other cases in point include military alliances and conflicts, and economic collaboration and competition. In social fields in general, actors are embedded in a variety of cooperative and conflicting relationships (Elias 1978; Simmel 1986) that originate from or lead to (possibly overlapping) groups. Incorporating contention is thus a major challenge for the detection of groups through network data.

In the December 2010 *American Sociological Review*, Shwed and Bearman used the modularity approach to study consensus formation in scientific communities. They claimed their approach enables them to distinguish consensus and benign criticism from epistemic rivalry (Shwed and Bearman 2010:818). Their data were scientific journal citations, which they interpreted for their modularity analysis as positive ties. On the basis of these citation data, they determined scientific communities and their salience, that is, the extent to which communities stood out from a random network. However, once we take into account not only the topology of ties but also an important aspect of their content—agreement versus contention—community detection yields different outcomes. This is important because mechanisms of network formation and social outcomes in *signed* networks (with both positive and negative ties) are different as well.

We discuss three important assumptions that Shwed and Bearman made to get their results. Their first assumption is that in normal science (Kuhn 1970), most citations signal agreement (Shwed and Bearman 2010:820). We agree with this assumption (see also Hanney et al. 2005) and found additional support for it in the literature (Case and Higgins 2000; White 2004). Their second assumption is that the comparatively few citations that represent disagreement have no ramifications for the communities detected (Shwed and Bearman 2010:820). We demonstrate, in contrast, that a small proportion of negative citations can substantially alter the results. Their third assumption is that epistemic rivalries between

communities are marked by a lack of cross-community citations (Shwed and Bearman 2010:821). In other words, contending factions largely ignore each other. They infer from this assumption that if the salience of communities diminishes, consensus increases (Shwed and Bearman 2010:818, 821, 822). However, a lack of citations between groups does not necessarily imply opposing views. We argue there are strong reasons to believe that groups detected in scientific citation networks are *thematic* communities, that is, groups of scholars specializing in and writing about the same themes or topics. These groups are less likely to be *positional* communities of scholars who agree with community members and disagree with other communities' views. This is a consequence of scholars mostly citing articles they consider relevant, regardless of their agreement or disagreement with those articles.

In the following, we first provide a general discussion on modularity, which is a quality measure used to compare different partitions of a network into groups. Along with pointing out a serious technical flaw made by Shwed and Bearman, our main claim is that it is not possible to analyze contention and consensus within or between communities when treating all ties as positive (Shwed and Bearman 2010:817). We offer an alternative approach that takes both positive and negative ties into account. Our approach is very general and can be applied to all networks, social and others, and is not specialized for scientific networks. Second, we analyze patterns of scientific citations on smoking is cancerous and solar radiation is cancerous, the latter is the same data used by Shwed and Bearman. As a baseline, we treat all ties as positive and test whether salient community differences arise when a small, randomly chosen portion of ties is coded as negative. Third, we study the evolution of the smoking is cancerous field over time by combining a new longitudinal approach to community detection with automated content analysis of abstracts of pertinent articles. Fourth, to illustrate that our approach is general, we analyze a dataset on a political public

debate wherein we distinguished positive and negative ties during data coding. For these data we show that the community structure obtained when—falsely—assuming all ties are positive is radically different from the community structure obtained when we distinguish between positive and negative ties. We conclude by arguing that, in general, researchers have to explicate negative relations (e.g., criticism, repel, competition, or violence) when analyzing fields where conflict is a mechanism of group formation. Only then can communities be validly detected.

MODULARITY AND COMMUNITY DETECTION

The basic idea of modularity is that one compares an actual network to a network with the same distribution of ties over the nodes (i.e., degree distribution) but that is otherwise random (Newman and Girvan 2004). For positive ties, groups (or communities) should have higher densities than would be expected on the basis of random chance, and they should be separated by cleavages that are sparser than in the randomized network. When tie strength is taken into account, strongly connected nodes are more likely to be in the same group than are weakly connected nodes, in line with Granovetter's (1973) well-known argument. Out of multiple possible assignments of nodes to groups that fulfill these requirements, the one with the maximal difference between actual and expected densities—maximal modularity—is arguably the best. In larger networks where the number of groups is not known at the start (such as the networks we study) there are in fact so many different assignments that comparing all of them is not feasible. Numerous algorithms exist that search heuristically for a good solution in ways that are efficient and effective (for an overview, see Fortunato 2010). Because most of these algorithms randomly reassign nodes to other groups in their search for modularity improvement, researchers might find a slightly different assignment after each run of community detection. Salient communities are robust, but the less salient communities are, and the

more the network resembles a random network, the less stable the assignment will be. This is not simply a limitation of the software: for non-salient communities with overlaps, different assignments can have (almost) the same maximal modularity value. To detect group overlaps and smaller subgroups initially overlooked, a parameter (of the *spinglass* model) can be fine-tuned to overcome some of the resolution limit inherent to this method, although several challenges remain (Fortunato 2010; Kumpula et al. 2007; Traag, Van Dooren, and Nesterov 2011). This is modularity optimization in a nutshell.

To assess Shwed and Bearman's application of modularity, it is important to stress at the outset that modularity scores are by definition normalized and maximally equal one. This is done by dividing the summed differences, of actual minus expected intra-community ties, by the total number of ties in a network. Shwed and Bearman normalized their modularity scores once more, however, by dividing the—already normalized—modularity of a given network by the logarithm of the network size (see their online supplement [http://asr.sagepub.com/content/ 75/6.toc]). They did not infer renormalization within the theoretical framework on which modularity is built, but inferred it on the basis of an arbitrary comparison with a simulated network. Note that none of the 457 expert papers reviewed by Fortunato (2010) renormalized modularity. Before Shwed and Bearman did so, most of their observation period saw modularity increases when the size of the corpus of scientific papers increased (left-hand side of Figure S2 in their online supplement). This finding could have been sociologically meaningful in itself, namely as a sign of increasing specialization. By their renormalization, however, their finding was turned into its opposite, namely stable or slightly decreasing modularity, which they interpreted as a sign of increasing consensus.

Furthermore, random graphs can sometimes exhibit surprisingly high modularity scores due to random fluctuations of ties (Reichardt and Bornholdt 2007), contrary to Shwed and Bearman's (2010:823) claim that

modularity would equal zero in a random network. This makes it difficult to compare modularity scores across different networks. To establish significance, one would have to generate large numbers of simulated random networks with the same degree distribution, and then compare them with a large number of community detection runs on the empirical network. This does not imply that the modularity approach is flawed, and in fact it has been validated in many test cases (Fortunato 2010), but when interpreting modularity results, one should take these considerations into account.

Although Shwed and Bearman's study is longitudinal, the authors did not consider how groups changed over time, only the overall modularity score. Recently, Mucha and colleagues (2010) crafted an effective method, termed multi-slice modularity, where each slice represents a network for a single year. An advantage over earlier approaches is that its longitudinal modularity depends not only on individual slices, like standard modularity, but also on communities' temporal stability. By using multi-slice modularity we muster additional analytic power for our approach.

To detect communities in signed graphs, negative ties should lie mostly between communities rather than within them, opposite to positive ties. Although cooperation, such as exchange, binds together people who may then be assigned to the same community if they are connected by more (or stronger) positive ties than would be expected by random chance, conflict pushes them apart. Consequently, if two or more people have more (or stronger) conflicting relationships with each other than would be expected by random chance, they should not be assigned to the same community. Combining these two arguments in modularity maximization, densities of within-group positive ties should be maximized, just as above, and negative within-group ties should be minimized (Traag and Bruggeman 2009). This conception generalizes social balance theory (Harary 1953; Wasserman and Faust 1994) by loosening its rigid community assignment, in which negative ties were not allowed to exist within communities.

The algorithm we use to detect groups with positive and negative ties is based on the Louvain method (Blondel et al. 2008), which performs well in test settings and is available online.[1] An earlier application of this algorithm to scientific citation data appears in Wallas, Gingras, and Duhon (2009).

SCIENTIFIC COMMUNITIES

Let us now scrutinize Shwed and Bearman's assumption that the comparatively few citations that represent disagreement have no substantial impact on the communities detected (p. 820). Although it is certainly true that some scientific citations are critical, perhaps the proportion of negative references is so low that it is safe to assume the comparatively few citations representing disagreement have no impact on the communities found?

To see if this is the case we examine two cases: solar radiation is cancerous and smoking is cancerous datasets. We received solar radiation is cancerous data from Shwed, so these citations are exactly the same as those used in their article. We collected the smoking is cancerous dataset from the ISI Web of Science using the same procedure Shwed and Bearman followed. For the latter data we also have most of the articles' abstracts, allowing us to analyze, to some extent, the scientific content of the communities, which we cannot do for Shwed and Bearman's data.

To distinguish negative from positive references, we would have to acquaint ourselves with the vernacular of cancer researchers and read thousands of articles, which is beyond our capability here. To test the impact of negative ties, we therefore set up the following procedure. We took a random sample of ties in the corpus, turned them into negative ties, and performed community detection on that network. We repeated this procedure a hundred times for each year, and each time we measured the difference between the negatively spiced assignment of nodes into communities and the original assignment; we did this for 5 and 10 percent of negative ties, respectively. Such small percentages of negative ties are

normally present in science (Case and Higgins 2000; Hanney et al. 2005; White 2004). To quantify the similarity of the assignments, we used the measure of Normalized Mutual Information (NMI) (Lancichinetti and Fortunato 2009). The NMI measure indicates, given one assignment, how much can be inferred about the other assignment. If one assignment completely predicts the other, the NMI score equals one; if nothing can be inferred, it equals zero. To make sure observed differences do not arise because of the algorithm's heuristic nature, which may lead to somewhat different outcomes in subsequent runs, we performed the same comparison but without changing any of the ties to negative.

Figure 1 displays the results; vertical bars indicate 95 percent confidence intervals and mean over 100 runs of the NMI score, and the comparison treatment with only positive ties is called independent. The figure shows that even a low proportion of negative ties can cause assignments to differ more strongly than when all ties are positive. This is the case for smoking is cancerous data (Figure 1a) and solar radiation is cancerous data (Figure 1b). Of course, these differences become larger and more salient when the percentage of negative ties increases. Our findings suggest Shwed and Bearman's assumption that the comparatively few citations representing disagreement have no impact on the communities detected (p. 820) is incorrect. Figure 1 shows that negative ties do have an impact and cannot be ignored if one wants to study contention.

One might expect that, in actuality, negative ties would lead to even more salient differences. The reason is that by sampling a certain percentage of ties randomly, we ignored any pattern in the negative ties, although we know from social balance theory (Wasserman and Faust 1994) and empirical studies (Szell, Lambiotte, and Thurner 2010) that negative ties tend to be present in-between specific communities and not randomly throughout a network. Networks with a small percentage of negative ties are thus likely to have a more salient community

structure than we find here. During periods of epistemic rivalry, when the percentage of negative ties is higher, the difference will usually be larger. As mentioned earlier, the actual pattern of negative ties is unknown to us and remains an empirical question. Nonetheless, our analysis shows that researchers are likely to detect a different community structure when negative ties are explicated.

DISSENSUS OR SPECIALIZATION?

We now focus on Shwed and Bearman's assumption that epistemic rivalry, that is, a lack of consensus, is characterized by a lack of cross-community citations (p. 818). Does this imply that if one finds a lack of cross-community citations, one could infer contention? Shwed and Bearman suggest this is so, saying that "when different communities are salient to the global structure, the field is contentious" (p. 822). But this inference is logically unsound; a lack of cross-community citations could also mean scholars within communities specialize in their proper subfields rather than they disagree with scholars in other subfields. Interestingly, Shwed and Bearman also speak about *changes* of modularity and consensus. Accordingly, an alternative reading of their assumption seems to be that over time, "consensus formation exhibits a decline in community salience" (p. 822). Again, it does not follow that if community salience declines, consensus increases, as Shwed and Bearman say when they discuss their findings: "We view such a significant decline [of modularity] over several years as consensus formation" (p. 830). Although neither decreasing nor increasing modularity in itself indicates consensus formation, looking at it over time is a potentially useful idea.

To examine what modularity over time could mean with respect to a scientific field, we used the smoking is cancerous dataset that contains abstracts and citations. We first extracted all words used in all abstracts of the corpus. We assumed that a group of articles that uses a shared vocabulary distinct from
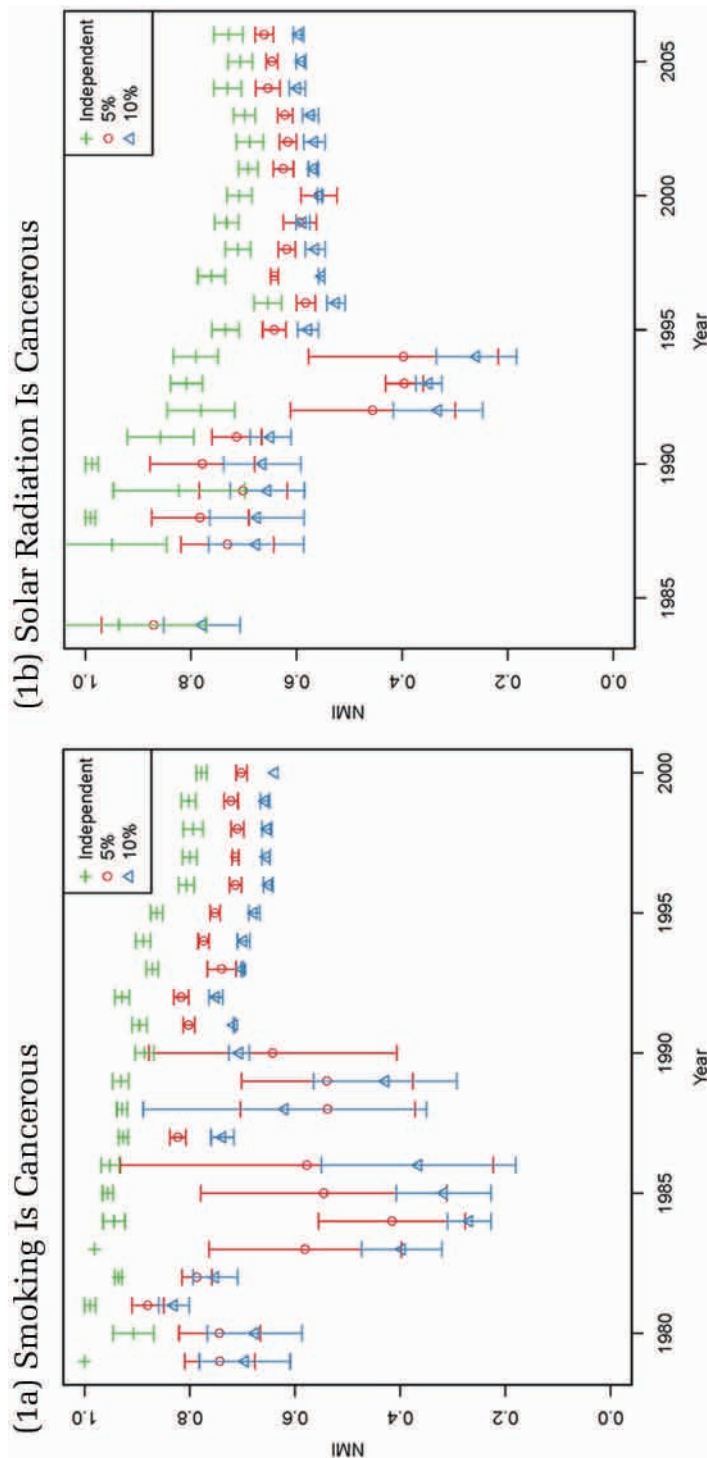
Figure 1. Different Proportions of Randomly Chosen Negative Ties in the (1a) "Smoking Is Cancerous" Corpus and (1b) "Solar Radiation Is Cancerous" Corpus
Note: The figure is best seen in color (see the article online [doi: 10.1177/0003122412463574]).

other groups discusses similar topics or methods. The common technique for extracting terms specific to a set of documents is the *term frequency-inverse document frequency* (tf-idf). The underlying principle is that for a certain term to be of specific interest or salience in a document, it should be frequently mentioned in that specific document and not mentioned much elsewhere (Salton and McGill 1983). For the groups here (that we first detected through modularity of citation patterns), if terms are common in a specific group and rare elsewhere, this indicates that articles in that group concern similar themes or topics. At the group level, we focus on the five most salient terms according to this tf-idf measure. Moreover, by using the multi-slice modularity method (Mucha et al. 2010), we obtained a dynamic view of the evolution of the groups, displayed in Figure 2, together with the five most salient terms for each group. This graphical representation of group dynamics as an alluvial diagram was invented by Rosvall and Bergstrom (2010). To avoid a cluttered image, we show only the 12 largest groups, which over the period of observation had at least 1,000 articles. Within a group, it is possible that scholars criticize each other, but we cannot detect contention because we do not know which citations are negative. We can, however, analyze how community structure changes over time with respect to common themes or topics.

Our approach makes it possible to provide a more substantive description of the evolution of the community structure. Figure 2 provides evidence that the field self-organizes into thematic groups in a process of ongoing scientific specialization, net of possible disagreements within these groups. This is a consequence of scientists citing articles that they consider relevant, regardless of possible disagreements. It seems less likely that the groups detected are positional, that is, consisting of scholars who mutually agree while disagreeing with other groups' views. During most of the observation period, modularity stayed more or less the same (see Figure 2).

Shwed and Bearman renormalized their modularity scores and found stable or decreasing modularity, which they took for increasing consensus. As we point out, without arbitrarily renormalizing modularity, it first increases but then remains relatively stable. Without information about negative ties, not much can be said with any certainty about consensus or dissensus. It is possible that within thematic communities there is disagreement such that, once negative ties are explicated, they turn out to be further partitioned into positional communities. In that case, single colored bands in Figure 2 would split into multiple colored bands (see the article online for figures in color).

A PUBLIC DEBATE

We now look at data that assess how large the difference between community assignments can be when ignoring the distinction between positive and negative ties. Because our method is very general and applicable to any network, we use data for which we know which ties are negative and are not scientific citations. Our dataset contains references between opinion makers in the debate over minority integration in the Netherlands. We focus on longer articles published in two broadsheet newspapers (*NRC Handelsblad* and *De Volkskrant*) between the assassination of the populist politician Pim Fortuyn (May 6, 2002) and the assassination of filmmaker Theo van Gogh (November 2, 2004). We selected articles from the Lexis-Nexis database through the key word integration in conjunction with foreigners, Muslims, or minorities. During this turbulent period in Dutch political history (Uitermark 2012), the newspapers ran 149 long articles (more than 1,000 words) on integration.

We manually coded references to individuals (both Dutch and foreign, dead and alive), institutions (e.g., political parties), and think tanks. Manual coding is laborious and requires some knowledge of the debate under study but it is currently the only way to properly distinguish the positive, neutral, or negative
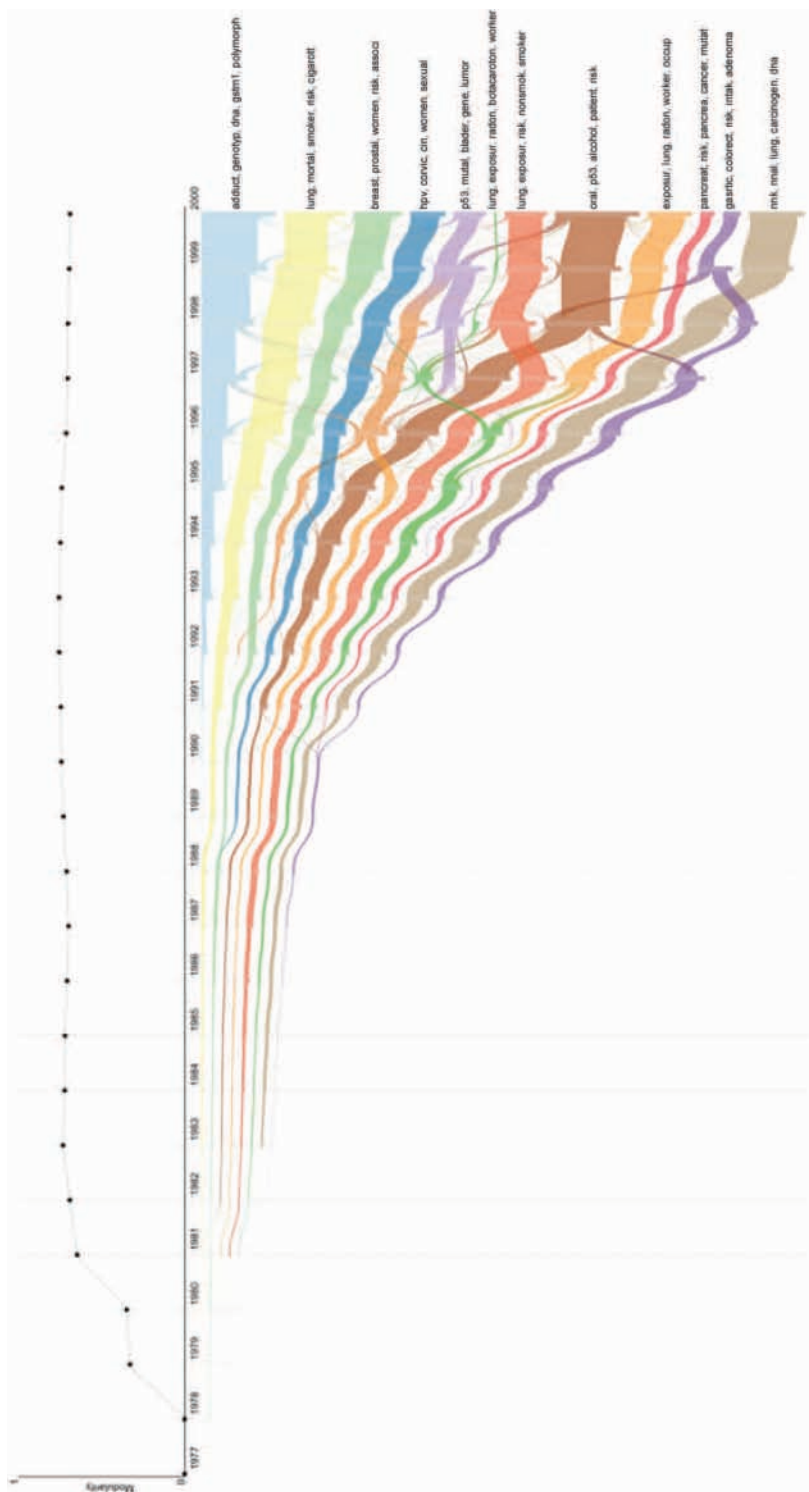
Figure 2. Alluvial Diagram of Scientific Specialization in Thematic Groups
Note: The figure is best seen in color (see the article online).

1057

content of references. Progress in automated coding is slower than expected a couple decades ago, and high validity coding is currently limited to word frequencies (e.g., Michel et al. 2011). In our 149 articles, we distinguished references according to their tone: positive, neutral, or negative. As a rule, we assigned positive and negative codes only if references were unambiguous. Because we coded references at the level of paragraphs, it was possible for one article to contain several references to the same actor, with each paragraph coded by the evaluation implicit or explicit in the references.

In total, we coded 1,779 references by authors commenting on others; among these directed ties, 318 were positive, 930 were neutral, and 531 were negative. Here we include only positive and negative ties and we consider only the largest component of the network, which has 323 actors. We defined tie weights between two actors by subtracting the number of negative references from the number of positive references. In Figures 3, 4, and 5, thicker lines denote stronger (either positive or negative) references, and gray shadings denote communities of referring actors.

First, we identify communities while assuming that all references are positive. As a result, the network in Figure 3 has a number of relatively dense groups of actors referring to each other. When we distinguish positive and negative ties in Figure 4, consistent with the data, we find two large communities, each with quite different membership from any of the communities in Figure 3. These two large communities have many references between them, but they are mostly negative. The two communities clearly disagree, and community membership now corresponds to a large extent to ideological identification. The communities are positional rather than thematic, and contention is a key mechanism of group formation in this field. The large community on the left consists mostly of actors who argue against stigmatization of Islam and other minorities, and the large community on the right contains a majority of actors who argue that mass migration and (radical) Islam

present a threat to Western civilization and to the Netherlands in particular.

Figure 5 shows the actual disagreement within the communities from Figure 3; here, we map the positional communities from Figure 4 onto the result from Figure 3. Gray shadings of the nodes show that disagreement dissects the communities initially found. The NMI score for Figures 3 and 4 is .34, which is relatively low given that many positively connected actors who were together in Figure 3 stay together in Figure 4. Our key point here is that if one assumes all ties to be positive, one obtains a very different result than if negative ties are explicated. Once both positive and negative ties are taken into account, it then becomes possible to analyze contention.

CONCLUSIONS

In recent decades, network analysts have used various approaches and measures to identify groups through network data. It has been difficult, however, to incorporate negative ties into analyses. Consequently, conflictual aspects of social life often escape from view. Shwed and Bearman proposed one way of using network data to study consensus and dissensus in scientific debates. Their main claim was that they can detect contention and consensus through scientific citation data by using modularity optimization. Their empirical finding was that consensus increases, indicated by decreasing renormalized modularity.

We broadly encourage studying consensus and dissensus through network data by using modern community detection methods. Nevertheless, we believe that Shwed and Bearman's claims and findings do not stand up to scrutiny. Without distinguishing negative ties explicitly, little can be said about contention within the field under study. Moreover, without renormalization, modularity does not decrease but stays more or less the same. If anything, the network data combined with word frequency analysis suggest that a key mechanism of group formation is specialization into subfields.
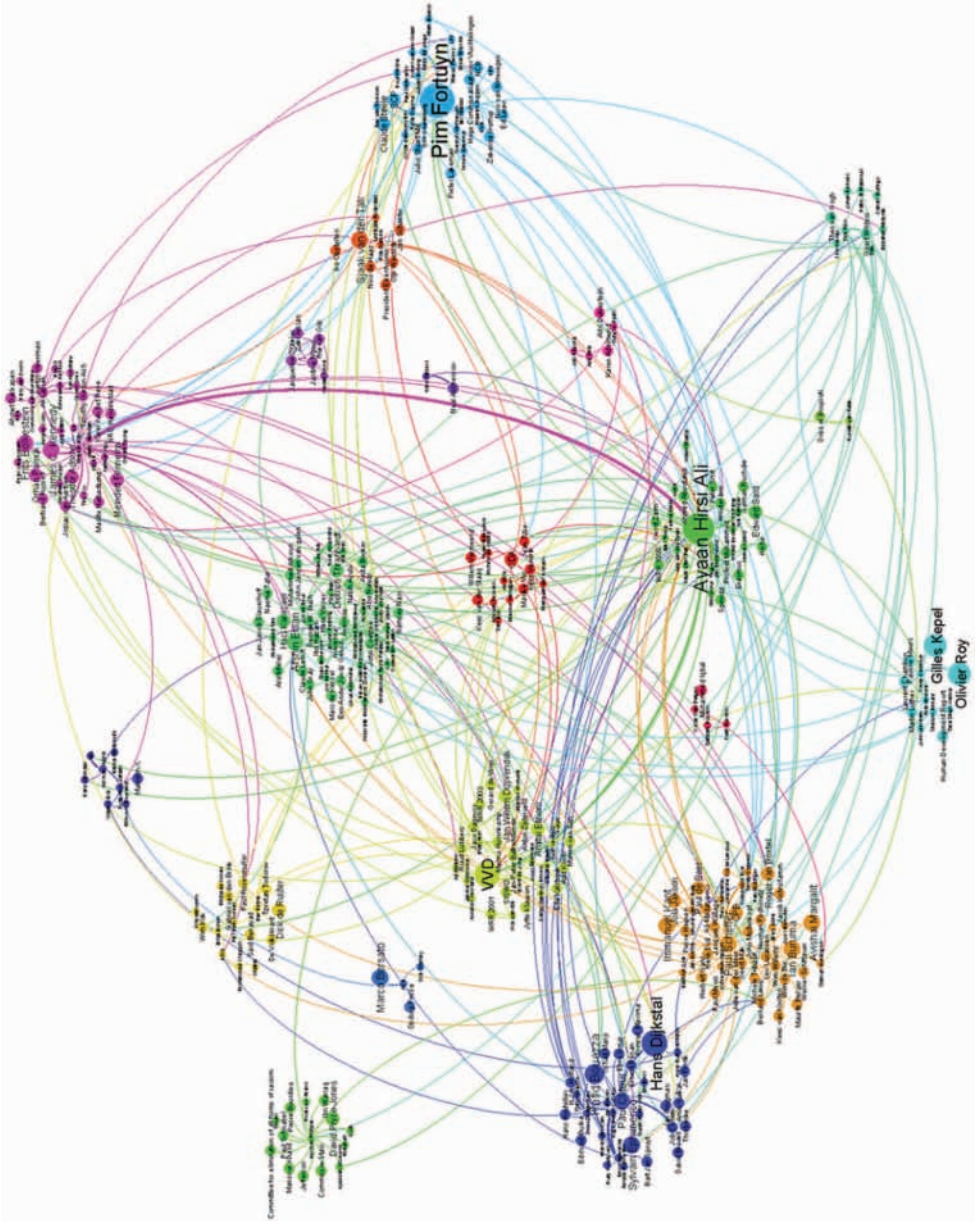
Figure 3. Public Debate Assuming All Ties to Be Positive
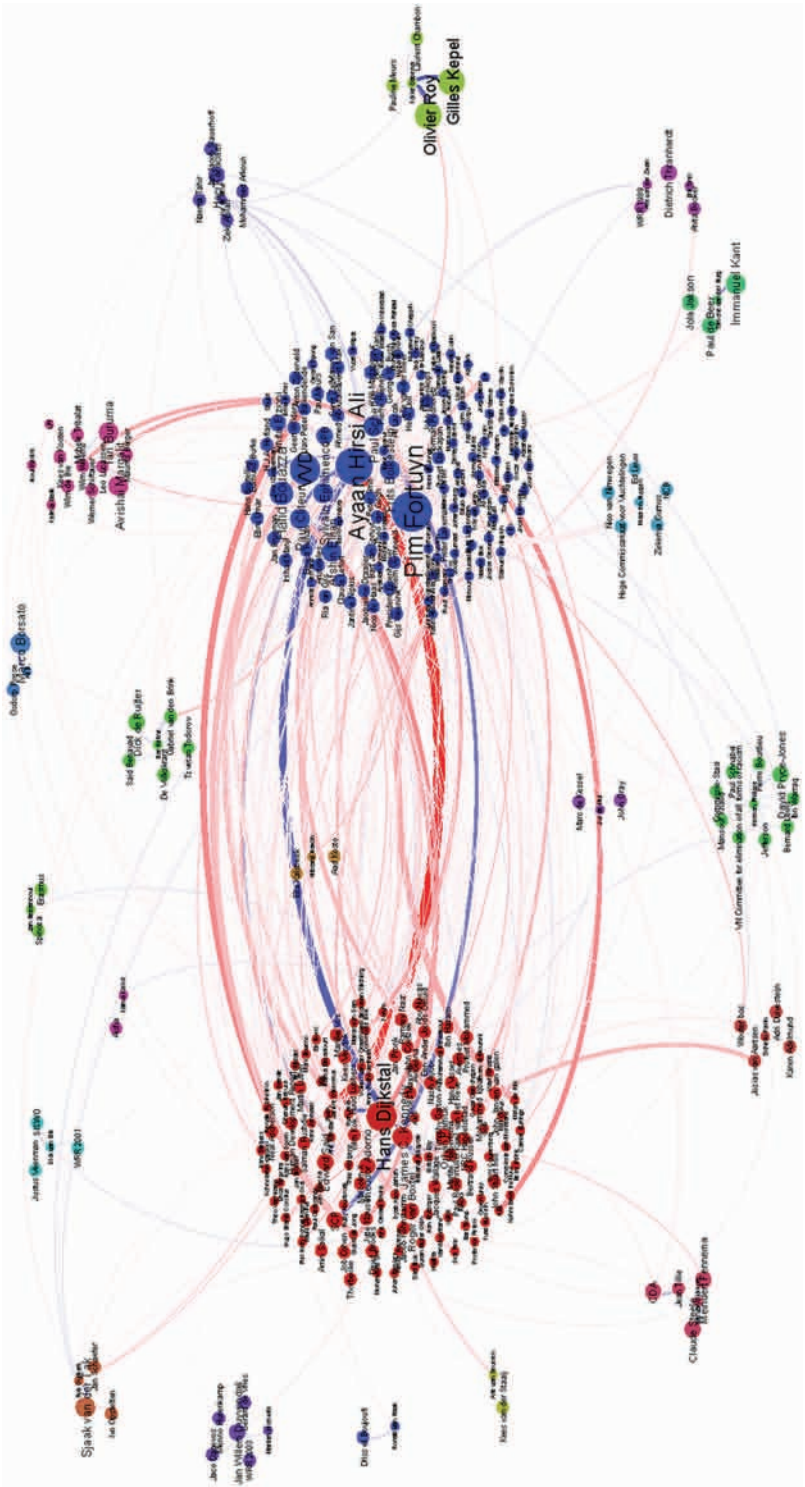*Note:* The figure is best seen in color (see the article online).

1059

Figure 4. Public Debate Distinguishing Positive and Negative Ties
Note: The figure is best seen in color (see the article online).

1060

Figure 5. Dissensus within Communities from Figure 3
Note: The figure is best seen in color (see the article online).

1061

Our several examples showed that distinguishing positive and negative ties is essential when studying conflict in networks. Incorporating the signs of ties has a substantial impact on the communities detected, even in fields where, as is the case for science, interaction is highly civilized and the proportion of negative references is low. We also showed that a community structure can be better understood by (1) using multi-slice modularity to map changes in community structure over time and (2) combining analysis of network topology with network content. For the latter, one can count word frequencies, along with incorporating tie signs. We hope this comment serves as a guideline for researchers interested in using community detection, especially in fields where negative ties are present. We believe this approach opens up exciting new possibilities to analyze group formation and change in scientific, political, and many other fields.

Acknowledgments

Funding

Note

1. The Louvain algorithm is available at http://launch pad.net/Louvain/.

References

Blondel, Vincent D., Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. "Fast Unfolding of Communities in Large Networks." *Journal of Statistical Mechanics: Theory and Experiment* 10:P10008.

Case, Donald O. and Georgeann M. Higgins. 2000. "How Can We Investigate Citation Behavior? A Study of Reasons for Citing Literature in Communication." *Journal of the American Society for Information Science* 51:635–45.

Elias, Norbert. 1978. *What Is Sociology?* New York: Columbia University Press.

Feld, Scott L. 1981. "The Focused Organization of Social Ties." *American Journal of Sociology* 86:1015–1035.

Fortunato, Santo. 2010. "Community Detection in Graphs." *Physics Reports* 486:75–174.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360–80.

Hanney, Steve, Lain Frame, Jonathan Grant, Martin Buxton, Tracy Young, and Grant Lewison. 2005. "Using Categorizations of Citations when Assessing the Outcomes from Health Research." *Scientometrics* 65:357–79.

Harary, Frank. 1953. "On the Notion of Balance of a Signed Graph." *Michigan Mathematical Journal* 2:143–46.

Kossinets, Gueorgi and Duncan J. Watts. 2009. "Origins of Homophily in an Evolving Social Network." *American Journal of Sociology* 115:405–450.

Kuhn, Thomas S. 1970. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Kumpula, J. M., J. Saramäki, K. Kaski, and J. Kertész. 2007. "Limited Resolution in Complex Network Community Detection with Potts Model Approach." *European Physical Journal B* 56:41–45.

Lancichinetti, Andrea and Santo Fortunato. 2009. "Community Detection Algorithms: A Comparative Analysis." *Physical Review E* 80:056117.

Michel, Jean-Baptiste, Yuan K. Shen, Aviva P. Aiden, Adrian Veres, Matthew K. Gray, Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez L. Aiden. 2011. "Quantitative Analysis of Culture Using Millions of Digitized Books." *Science* 331:176–82.

Mucha, Peter J., Thomas Richardson, Kevin Macon, Mason A. Porter, and Jukka-Pekka Onnela. 2010. "Community Structure in Time-Dependent, Multiscale, and Multiplex Networks." *Science* 328:876–78.

Newman, Mark E. J. and Michelle Girvan. 2004. "Finding and Evaluating Community Structure in Networks." *Physical Review E* 69:026113.

Reichardt, Jörg. 2009. *Structure in Complex Networks* (Lecture Notes in Physics). New York: Springer.

Reichardt, Jörg and Stefan Bornholdt. 2007. "Partitioning and Modularity of Graphs with Arbitrary Degree Distribution." *Physical Review E* 76:015102.

Rosvall, Martin and Carl T. Bergstrom. 2010. "Mapping Change in Large Networks." *PLoS one* 5:e8694.

Salton, Gerard and Michael J. McGill.1983. *Introduction to Modern Information Retrieval*. New York: McGraw-Hill.

Shwed, Uri and Peter S. Bearman. 2010. "The Temporal Structure of Scientific Consensus Formation." *American Sociological Review* 75:817–40.

Simmel, Georg. 1986. *Conflict & the Web of Group-Affiliations*. New York: Free Press.

Szell, Michael, Renaud Lambiotte, and Stefan Thurner. 2010. "Multirelational Organization of Large-Scale Social Networks in an Online World." *Proceedings of the National Academy of Sciences USA* 107:13636–641.

Traag, V. A. and Jeroen Bruggeman. 2009. "Community Detection in Networks with Positive and Negative Links." *Physical Review E* 80:036115.

Traag, V. A., P. Van Dooren, and Y. Nesterov. 2011. "Narrow Scope for Resolution-Limit-Free Community Detection." *Physical Review E* 84:016114.

Uitermark, Justus. 2012. *Dynamics of Power in Dutch Integration Politics: From Accommodation to Confrontation*. Amsterdam: University of Amsterdam Press.

Wallas, Matthew L., Yves Gingras, and Russell Duhon. 2009. "A New Approach for Detecting Scientific Specialties from Raw Cocitation Networks." *Journal of the American Society of Information Science and Technology* 60:240–46.

Wasserman, Stanley and Katherine Faust. 1994. *Social Network Analysis*. Cambridge, UK: Cambridge University Press.

White, Howard D. 2004. "Citation Analysis and Discourse Analysis Revisited." *Applied Linguistics* 25:89–116.

White, Harrison, Scott A. Boorman, and Ronald L. Breiger. 1976. "Social Structure from Multiple Networks: I. Blockmodels of Roles and Positions." *American Journal of Sociology* 81:730–80.

Jeroen Bruggeman works at the University of Amsterdam in the fields of organizational sociology and social networks. He wrote the textbook *Social Networks: An Introduction* (2008). He connects these and other fields in a broader evolutionary perspective.

V. A. Traag is a PhD candidate in applied mathematics at the Université Catholique de Louvain. He obtained his master's in sociology at the University of Amsterdam (2008). He currently works on community detection and social balance dynamics.

Justus Uitermark is Assistant Professor of Sociology at the Erasmus University Rotterdam. He is an urban and political sociologist and has published on social movements and urban governance. His recent book, *Dynamics of Power in Dutch Integration Politics: From Accommodation to Confrontation*, examines the drastic transformations in how the Dutch deal with ethnic and religious differences.