

Identifying Similar Opinions in News Comments using a Community Detection Algorithm

Jonathan Scott¹, David Millard¹ and Pauline Leonard²

¹ Web and Internet Science, School of Electronics and Computer Science, University of Southampton, UK

² Sociology, School of Social Sciences, University of Southampton, UK

Abstract. Despite playing many important roles in society, the news media have been frequently criticised for failing to represent a wide range of viewpoints. Online news systems have the potential to allow readers to add additional information and perspectives. However, due to the simplicity of the filtering mechanisms typically employed, these systems can themselves be prone to over-promoting popular viewpoints at the expense of others. Previous research has attempted to diversify news comments through the use of content similarity, sentiment analysis, named entity recognition, and other factors. In this paper we propose the use of a commonly used community detection algorithm on a network of voting data to identify sentiment groups in news discussion threads, with the eventual goal that these groups may be used to present diverse content. In a controlled experiment with 154 participants, we verify that the Louvain Community Detection algorithm is able to group users with accuracy comparable to an average human. This produces groups containing users who share similar sentiment on a given topic. This is an important step towards ensuring that each group is represented, as by using this method future news systems can ensure that more diverse views are represented in open comment threads.

Keywords: community detection, news, comments, discussion, sentiment, viewpoint

1 Introduction

In 2008, Stromback claimed that the media have become “the most important source of information for most people in advanced democracies around the world” [29]. They fill many roles in modern democratic society: As an agenda-setter the media influences the focus of public opinion [30, 18], as the “fourth estate” they are expected to hold the powerful to account [27], and as an information provider they are tasked with ensuring the population are informed about the processes and decisions which concern their lives [2, 16].

Criticism of the news media has been frequent however, with many complaints of an over-reliance on elite sources (e.g. [8, 13, 4]). They have also been shown to be failing in their role as information provider. In 2013, an Ipsos MORI

survey [14] found that, among other misperceptions, British people estimated the amount of benefit fraud as 34 times higher than official estimates, that 24% of the population were Muslim (compared to the official figure of 5%) and 31% of the population were immigrants (compared to the official figure of 13%). A similar report in 2014 found very similar figures [24]. With welfare and immigration regularly ranking high on election priorities [31] this shows that the media have not adequately informed the public of the information they need to properly participate in democracy.

Part of this issue may be explained by the lack of perspectives represented in the media. In 1972, McCombs found that amongst local daily newspapers, national newspapers, and national news broadcasts, there existed a high degree of similarity of news agenda [19]. More recently, Nick Davies found that 60% of the stories in four chosen “quality newspapers”³ comprised wholly or mainly of material from newswires and public relations groups, and only 12% could be confidently attributed to a named reporter [8]. This reveals a situation where much of the information the public receive is provided by very few sources.

The rise of online news has allowed for new methods of citizen involvement in the news process. Once restricted to passive consumption of journalist-produced content, readers now regularly contribute to news, through online commenting and discussion systems, submitting media directly to journalists, and on some systems through the creation of their own stories. In 2014 it was found that half of social network users shared news on their social network accounts, 46% of users discussed news on these sites, and roughly 10% of users had published news videos they made themselves [22].

However, with this much increased public contribution comes the difficult problem of sorting and filtering this content into a form that can be easily consumed, a job performed in traditional media by a news editor. One common solution is to allow other readers to vote for high quality content and then show the highest rated prominently. This is the solution employed by a number of social media platforms, as well as the websites of many newspapers [28].

One issue with this method of filtering is that it can encourage groupthink [21]. Jokingly referred to as the “hive mind” within Reddit communities, this refers to the common phenomenon whereby members of a group tend to reach a consensus and avoid dissenting opinions. Muchnik et al. performed an experiment on an un-named social news aggregation website and found that both prior positive votes and prior negative votes increased the number of future positive votes [23], and Mills found that on Reddit, minority opinions are “slightly marginalised but not excluded” [21].

There is evidence that being responded to is a key determinant in the decision to continue contributing to a community [3, 6]. Interfaces which make less highly rated comments less visible can lead to discouraging future contributions from those who do not agree with the majority viewpoints. Over time this will lead to less diversity in news discussions, undermining the potential of online news to solve the issues identified by critics of traditional journalism. To avoid this, it is

³ The Times, The Guardian, The Independent, and The Daily Telegraph

important to ensure that news discussions present diverse comments and fairly represent the viewpoints of the different groups involved.

This paper provides an overview of existing attempts at diversifying comments online, proposes a method of diversification using voting data and community detection algorithms, and presents an experiment to test the ability of community detection algorithms to automatically detect groups of commenters with similar opinions and views. If comment systems could automatically group users in this way then it would be possible to design new commentary systems that move beyond simple chronology or popularity in order to highlight a more diverse set of comments and opinions to users.

2 Background

Previous work by Giannopoulos, et al. attempted to diversify comments by building on textual diversity algorithms used in other domains, and added additional criteria specific to comments. They produced a system for selecting diverse comments using measures of content similarity, sentiment analysis, named entity recognition, and comment quality [11]. We will instead approach the problem by concentrating on groups of users rather than individual comments, relying on homophily and an assumption that people will tend to maintain a consistent viewpoint within the timeframe of a single discussion thread.

Previous research has attempted to identify communities on social networks and in online discussions without the focus on diversification. This has primarily used explicit friend/follower relationships or text-mining. Jaffali, et al. look at community identification using tweets in the context of friend recommender systems. They propose an algorithm which performs text-mining and sentiment analysis, producing communities of users which share sentiment towards a given entity [15]. Abu-Jbara, et al. utilise opinion mining and sentiment analysis on online Arabic discussions to identify groups who share an opinion by their use of subjective language [1]. Parau, et al. investigate the use of multidimensional data in identifying sentiment communities [26]. Each of these approaches has merit, but they do not leverage the very common interaction of users voting for content they agree with.

Voting is one of the more common interactions in online communities, including Reddit, Facebook [12], and Twitter [20], and we propose the use of this data to identify users who share sentiment towards a topic. It should be noted that by “sentiment” we mean “an attitude toward something; opinion”⁴. For this voting data to be useful for our purposes, we require voting be performed for a particular reason: to indicate agreement or disagreement with a piece of content. Kriplean, et al. view the “like” buttons on social websites as overloading two functionalities: providing a way to recognise and appreciate a speaker, but also including an implicit agreement with the content [17]. Using the buttons in

⁴ “sentiment” definition from dictionary.com. Accessed 14/10/15

this way is contrary to the rules of many popular social news systems⁵ but as mentioned, Reddit users regularly discuss the “hive mind” and it has been shown that there is at least some marginalisation of minority views in social news [21]. This indicates that to some extent, users are voting for content with which they agree rather than only for content which they think is of high quality.

In this paper we will use this voting data to create a network of users with edge weights representing the number of times they’ve mutually “liked” a piece of content. Once the network is created, it can be partitioned using one of many existing community detection algorithms.

2.1 Community Detection Algorithms

A community is “a subgraph of a network whose nodes are more tightly connected with each other than with nodes outside the subgraph” [10]. Community detection techniques are commonly used for analysing networks such as community organisations and scientific collaborations [9]. They have also been used for tag disambiguation, user profiling, and event detection [25] and to improve friend recommendation systems and collaborative filtering techniques.

There are some specific requirements of a community detection algorithm to be used with online discussions. First, it is required that the algorithm be able to run in near-real-time on average news discussions. This is so that the results can be immediately presented to users to maximize the impact. Second, due to the nature of online discussions, there is no way of predicting in advance how many communities will exist or what the size of those communities will be, techniques that require this information to be provided will not be suitable for the task.

Community detection algorithms can be broadly grouped into three categories: divisive algorithms, which detect and remove inter-community links; agglomerative algorithms, which recursively merge communities; and optimization algorithms, which attempt to maximize some function [5]. They are typically judged on the “modularity” of the communities detected, a value between -1 and 1 that represents the density of links within communities compared to the density of links between communities.

There are a number of commonly used algorithms which optimize for modularity, though Fortunato and Barthelemy identified a problem with these techniques, showing that modularity “contains an intrinsic scale that depends on the total number of links in the network”, and that communities which are smaller than this scale may not be detected at all [10]. This is particularly a problem in very large networks where smaller communities will not be detected.

The method used in this paper is the Louvain Community Detection algorithm, as proposed in 2008 by Blondel, et al. [5]. It is a heuristics-based method which optimises for modularity. It is simple to implement, performs well (achieving high modularity in low computing time), and runs well even on large net-

⁵ e.g. Reddit “Moderate based on quality, not opinion” from <http://www.reddit.com/wiki/reddiquette> or Slashdot “simply disagreeing with a comment is not a valid reason to mark it down” from <http://slashdot.org/faq>

works. It also partially deals with the resolution limit problem identified by Fortunato and Barthelemy.

2.2 Louvain Community Detection

The Lovain Community Detection algorithm is very fast, with linear complexity on typical data [5]. The algorithm begins with each node assigned to its own community, and repeatedly executes two phases. The first phase is described by Blondel, et al. as:

For each node i we consider the neighbours j of i and we evaluate the gain of modularity that would take place by removing i from its community and by placing it in the community of j . The node i is then placed in the community for which this gain is maximum (in case of a tie we use a breaking rule), but only if this gain is positive. If no positive gain is possible, i stays in its original community. This process is applied repeatedly and sequentially for all nodes until no further improvement can be achieved and the first phase is then complete.

The second phase involves creating a new network whose nodes are the communities created during the first phase. The weights of edges between the new nodes are the sum of the edge weights between the two communities. These two phases are repeated until there are no more changes.

Due to the recursive nature of the algorithm, it provides intermediate stages which allow for different levels of granularity in the communities. This feature helps to partially avoid the resolution limit identified in by Fortunato and Barthelemy as selecting different levels will result in communities of different sizes appearing. These intermediate stages provide a hierarchy of communities which may be useful when identifying groups of users with similar sentiment, though in this study we will focus only on the level with the highest modularity.

In their paper, Blondel et al. test the algorithm’s performance by running it against a small social network, a network of scientific papers and their citations, a sub-network of the internet, a network of webpages, and other datasets [5]. It performs well in all cases (having a high level of modularity with a small computation time). To verify that the Lovain Community Detection algorithm will be able to group together news discussion participants who share similar sentiment, we created a network from existing news discussions and employed a web-based study to explore the accuracy of the detected communities.

3 Methodology

We looked to the discussion systems of the websites analysed in our previous work [28] for systems which 1) have enough activity to generate a large amount of data and 2) have APIs which allow access to votes. Four of the systems provided the data needed to create a network for use with the Louvain algorithm: Facebook, Twitter, CNN, and The Telegraph.

Of these, The Telegraph and CNN use a common discussion system (Disqus), and therefore have a common API. They also, as outright news sites, have a concept of a news “story” which the other two systems do not. This makes it possible to investigate grouping users based on sentiment without needing to first separate them based on topic. For these reasons, this experiment uses data from The Telegraph and CNN websites.

3.1 Data Collection

We collected every news story which, between 22/05/2014 and 27/05/2014, featured in the “Most Popular” section of CNN or the “Most Viewed” section of The Telegraph. This resulted in 44 stories: 12 from CNN and 32 from The Telegraph. The difference can be attributed to the fact that The Telegraph’s “Most Viewed” section shows 10 items at a time, whereas CNN’s “Most Popular” section shows only 6. We then gathered all comments on each of these stories, including which users had voted for which comments. This resulted in 23,655 comments and 49,486 votes.

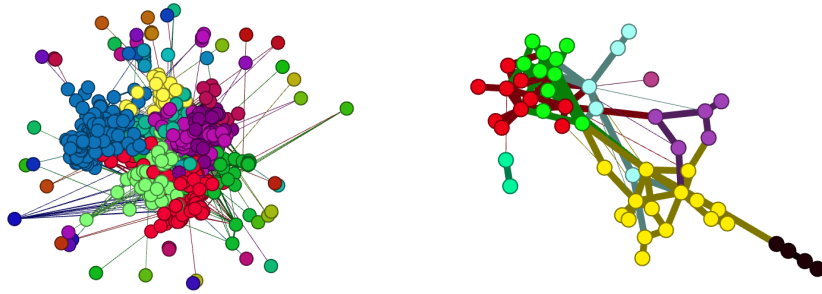


Fig. 1: Detected communities for two randomly selected stories: “Bulls take down bullfighters but still lose in Spain”, and “How ITV missed the best moment of the FA Cup final”

From this data, for each story, a weighted graph was produced of relationships between the people who interacted with that story. This gave us an average of 207 nodes ($s = 262$) and 841 edges ($s = 1621$) per graph.

In each graph, nodes represent users who have interacted with the story, and edges represent the level of agreement between users (the number of times they voted the same way minus the number of times they voted differently). As the Louvain algorithm does not support negative edge weights, values less than 0 are discarded.

When used on these networks, the Louvain algorithm generates a total of 877 groups (an average of 19 groups per story, $s = 24$) with an average of 10 members per group ($s = 29$), and an average modularity of 0.516. We have calculated these

numbers for each category of news on each of the news systems, which can be seen in Table 1.

Figure 1 shows the communities generated for two randomly selected stories. On the left is the CNN story “Bulls take down bullfighters but still lose in Spain” about a bullfighter being injured. The generated communities have a modularity of 0.63. There are 56 communities generated in total, though the majority of these are single-user communities which are not connected to any other nodes and are not shown in the figure. The largest community (with 23% of users) is shown in blue and consists of users supportive of bull fighting, and the second largest (19% of users) shown in red with comments typically supportive of the bulls. However the third largest community (12.4% of users), shown in light green is also supportive of the bulls and a difference in content between the two groups is not clear.

On the right is the Telegraph story “How ITV missed the best moment of the FA Cup final”, about a football match between Hull City and Arsenal. This network has 9 communities and a modularity of 0.49. The largest community (27.3% of users), is shown in yellow and consists primarily of Hull supporters. The second largest (18.2% of users) is shown in green and consists primarily of supporters of Arsenal. The third largest in this case (18.2%) is shown in red and does not seem to specify a team preference, instead consisting of users criticising the article and author.

Table 1: Number of stories; average number of nodes, edges, and communities; average community size; and average modularity for each category of story, separated by news website

Category	Stories	Nodes	Edges	Communities	Community Size	Modularity
The Telegraph						
Sport	15	49.27 (s=32.15)	83.33 (s=65.59)	9.07 (s=4.46)	4.8 (s=2.14)	0.5 (s=0.09)
World	6	300.83 (s=210.76)	1562.33 (s=1451.01)	15 (s=3.74)	17.67 (s=10.27)	0.42 (s=0.06)
Domestic	5	618.4 (s=372.59)	3726.6 (s=2850.92)	23.4 (s=12.44)	25.6 (s=9.18)	0.46 (s=0.04)
Technology	2	184 (s=109)	353 (s=222)	26.5 (s=16.5)	6.5 (s=0.5)	0.59 (s=0.11)
Obituaries	1	32	54	8	4	0.46
Finance	1	367	1371	12	30	0.47
Business	1	123	193	17	7	0.61
Culture	1	118	291	16	7	0.55
CNN						
Showbusiness	4	119.5 (s=82.89)	212 (s=157.56)	22.5 (s=11.43)	4.5 (s=0.87)	0.54 (s=0.09)
Travel	3	72.67 (s=4.03)	111.33 (s=21.68)	20.67 (s=4.5)	3 (s=0.82)	0.61 (s=0.04)
World	3	527 (s=290.25)	1195 (s=863.81)	80.33 (s=55.76)	7.33 (s=4.78)	0.65 (s=0.13)
Business	1	68	110	13	5	0.47
US	1	127	257	22	5	0.67
Overall	44	207.18 (s=262.95)	841.05 (s=1621.38)	19.93 (s=23.8)	9.68 (s=9.52)	0.52 (s=0.1)

3.2 Comparison with Manual Classification

Having gathered this dataset of comments and groups, we performed an experiment to compare the results of the Louvain algorithm with human classification. This was undertaken in the form of a web-based assessment, which showed users example comments and asked them to make a judgement as to whether they expressed similar sentiment. Our hypothesis is that pairs from users in the same group will be viewed as “similar” more often than pairs from users in different groups.

ERGO/FPSE/10814

3 answered of 20

The 'desert woman' who crossed Australia's 'dead heart'

OriginalCommenter

Women cannot 'do it just like men'. Certainly, women and men are equally intelligent, but men in general are undeniably physically stronger, and therefore would have an advantage over women in a scenario such as this.

SubCommenter1

Men have greater upper-body strength. Women have greater endurance. Plus, being smaller, they'd need less food and water. In a scenario like this, I'd give the nod to the ladies.

SubCommenter2

But sometimes perseverance (or plain ole madness :P dunno what it is) outwits physical strength .

Instructions

On the left is a comment by **disqus_HY8K0...** which was posted in response to an article titled **"The 'desert woman' who crossed Australia's 'dead heart'"**

The **green** comment (by **SubCommenter1**) and the **blue** comment (by **SubCommenter2**) are responses to that comment.

Do you believe that the author of the **green** comment and the author of the **blue** comment share a similar sentiment regarding the original comment?

Dissimilar

Can't tell

Similar

Fig. 2: The sentiment experiment interface

To generate the questions for this experiment, we identified every pair of comments posted in reply to a common parent (this was to make sure that the comments were broadly on the same topic, which makes manual classification easier for the participants, and reduces the number of times they skip the question - thus maximising the number of results received for the number of questions asked). These pairs were filtered to ensure that there was only a single question for each combination of groups and parent comment. That is, for comment i , and groups j and k there is only a single pair of comments with parent i , and one author from each of j and k . We also ensured that there were no pairs with both comments from the same author. After this filtering, 3554 pairs remained, 1658 with both authors from the same group, and 1896 with authors from different groups. For the purposes of this study, 50 pairs with same group authors and 50

pairs with different group authors were randomly selected. This gave 56 pairs from The Telegraph and 44 pairs from CNN.

Participants were then invited to view these comment pairs through an online interface. The study was shared on social media and participants were invited to share the study with their own contacts. This resulted in 154 respondents, though 20 did not complete the experiment and so the ratings they contributed were removed.

The participants were presented with 20 pairs of comments and instructed to decide for each pair of comments if their sentiments towards the parent comment were similar, dissimilar, or if they could not tell. The interface for deciding the similarity of sentiment can be seen in Figure 2.

4 Results

From the 134 participants who completed the experiment, we received 2376 ratings, 1098 votes for “similar”, 838 votes for “dissimilar”, and 404 votes for “can’t tell”. The “can’t tell” votes will be treated as an inability to answer, which leaves 1936 usable votes. The distribution of the ratings can be seen in Figure 3, and an example question can be seen in Table 2.

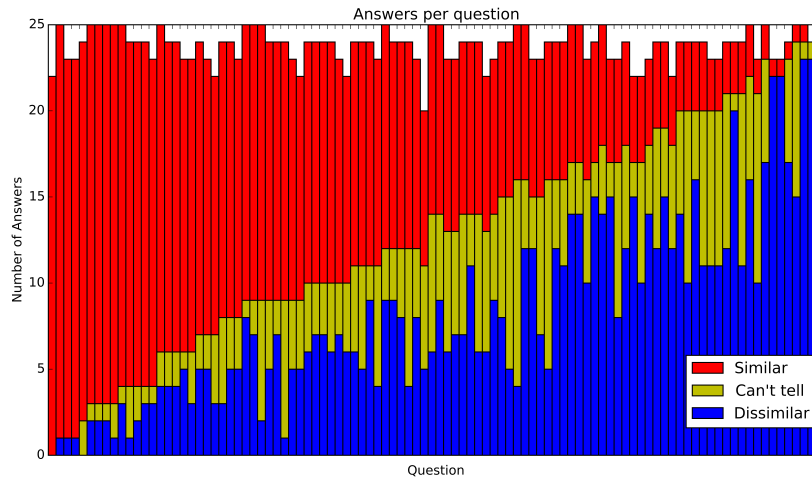


Fig. 3: The distribution of sentiment ratings

For each question, we calculated the percentage of respondents that said the sentiments expressed were “similar” (see Figure 4). As the data are not normally distributed, we used a Mann-Whitney test to compare the percentage between questions with comments from authors in the same group and questions with comments from authors in different groups. On average, comments by authors

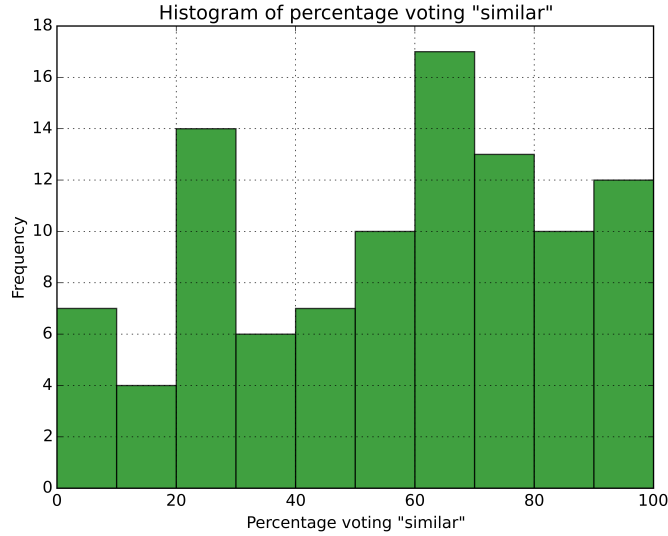


Fig. 4: The distribution of questions by percentage voting “similar”

who were allocated to the same group had a higher percentage of “similar” answers ($M = 66.91$, $SE = 3.11$) than comments by authors allocated to different groups ($M = 43.6$, $SE = 3.91$). This difference is significant ($U = 653$, $z = -4.116$, $p < .001$, $r = -.4116$).

To evaluate the algorithm’s success when compared to a human, we then used the grouping status (same group or different groups) as a vote by the system (for similar and dissimilar respectively), and calculated how often the system’s vote agreed with the votes of participants (62.09%). We also calculated this for each participant ($M = 66.37\%$, $s = 8.73$), and generated 10,000 randomly allocated groupings to calculate how often these random groupings agreed with the votes of participants ($M = 44.61\%$, $s = 1.82$). This can be seen in Figure 5.

5 Discussion

There is a significant difference between the percentage of respondents voting “similar” on questions where the authors were allocated to the same group, when compared to questions where the authors were allocated to different groups. This indicates that the Louvain Community Detection algorithm, when applied to a network produced using voting data, is able to produce groups of users who share sentiment on a given topic, in that the groups produced are more likely than random groupings to group together users judged to have similar sentiment. This is evidence that supports our original assumptions (that people will maintain consistent viewpoints within a single discussion, and that they will vote for content that they agree with).

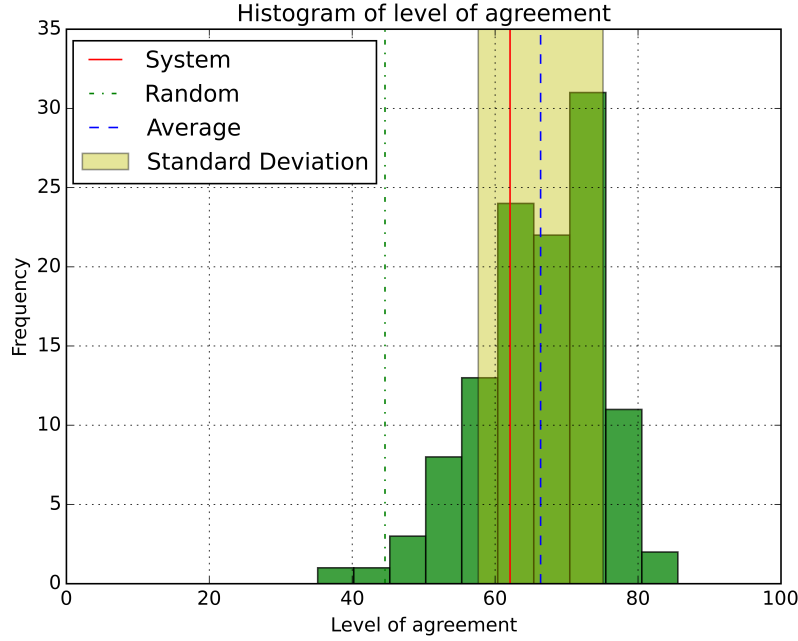


Fig. 5: The level of agreement between participants

To compare the system to the potential ability of a human to categorise the commenters by sentiment, we calculated how often the system agreed with each participant in their assessment of the similarity of comments, and also how often participants agree with each other. The average agreement of the system (62.09%) is within a single standard deviation (8.73) of the participants' average level of agreement (66.37%). This indicates that the system has comparable success to an average human when deciding if the sentiment of two comments is similar. These agreement levels are much higher than that which would be expected given random groupings (44.61%).

To further explore the quality of the sentiment groupings would require access to existing news editors who would be willing to submit their own votes on the comment pairs used in this study. This would allow a comparison between the use of community detection algorithms and best practice for other types of content (such as letters and emails to the editor), where manual filtering is used to ensure content diversity.

The relatively low levels of average agreement between participants may indicate that there is disagreement over what level of difference should be considered as “dissimilar sentiment”. This should be investigated in future work and could lead to the development of systems which use the views of the user to present a broad range of viewpoints, rather than presenting a single set of content for

Table 2: Example question from sentiment experiment

Question ID	36682
Groups	Different
Parent Comment	<p>OriginalCommenter: If the King did not approve the coup that kick Thaksin out, but let democracy does its job back then, Thailand would not be divided as much as today.</p> <p>where is the King now? It's time for him to clean up the mess he and his people created. "I believe Thai both side still listen to him". it is not too late for him to make it right.</p> <p>He need to say and do something now. the longer he wait the worse it will get for his country and his people. And he can only blame himself for the problem he took part in creating it.</p>
Comment 1	SubCommenter1: The King and Queen are extremely old, impaired and live in seclusion in a hospital. He only makes the occasional public appearance on his birthday etc and no longer has the capacity to influence events.Their son and heir apparent is in the military so...
Comment 2	SubCommenter2: the scenario is simple: The king or his son ordered the general to move... they are going to twist the arm of the acting Pm to resign and appoint an interim party.. Then the army will descend on the Red and kill 100 or 200 on the excuse is that they have started something. some crazy excuse to satisfy the US three years later, they will have a controlled election.. and that is the end of that...Thailand has no democracy and never will SO the part I do not understand is that why the Japanese are still doing business there or the Germans, or the Koreans.. it is a corrupted infested state.
Similar	5
Dissimilar	15
Can't Tell	4
Percentage Similar	25%

all users, or, as in the case of many “personalised” news systems, primarily presenting viewpoints which the user identifies with.

6 Conclusions

This paper proposes using community detection algorithms to identify comments that express similar views or opinions, the long term goal being to use these groupings to reduce issues of group-think in open news systems, by presenting more diverse opinions. The success of this technique relies on the assumption that news discussion participants would maintain a consistent viewpoint within the timeframe of a single discussion, and would not strictly obey the instructions to “vote based on quality”, instead allowing their biases to show through their voting behaviour.

After identifying a suitable algorithm (Louvain Community Detection), an experiment was conducted which gathered 1936 votes on the similarity of pairs of comments, and it was found that there was a significant relationship between

group status (if the algorithm placed the authors in the same group or not) and the percentage of respondents who believe the pair shared similar sentiment. The system was then compared to each participant, and was found to have a level of agreement comparable to the average level of agreement between participants.

Future work may investigate the possibility of adding vote-based community detection to the promising work using explicit friend/follower relationships and textual analysis, or apply different community detection algorithms to voting networks. Using these techniques may allow for identification of more fine-grained sentiment groups.

In our own future work, we will focus on formatting and presenting full news discussions using the algorithm explored here, with a view to discovering if presenting comments in the context of like-minded groups causes people to interpret and interact with the content in a different way. The hope is that this work will lead to online news discussions which represent a wider range of viewpoints, are more welcoming of dissenting opinion, and more informative for readers.

References

1. Abu-Jbara, A., King, B., Diab, M. T., and Radev, D. R. (2013). Identifying opinion subgroups in arabic online discussions. In *ACL (2)*, pages 829–835.
2. Anderson, P. J., Williams, M., and Ogola, G. (2013). *The Future of Quality News Journalism: A Cross-continental Analysis*, volume 7. Routledge.
3. Arguello, J., Butler, B. S., Joyce, E., Kraut, R., Ling, K. S., Rosé, C., and Wang, X. (2006). Talk to me: foundations for successful individual-group interactions in online communities. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 959–968. ACM.
4. Bennett, W. L. (1988). *News: The Politics of Illusion*. Longman, New York.
5. Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10).
6. Cheng, J., Danescu-Niculescu-Mizil, C., and Leskovec, J. (2014). How community feedback shapes user behavior. *ICWSM 2014*.
7. Clauset, A., Newman, M. E., and Moore, C. (2004). Finding community structure in very large networks. *Physical review E*, 70(6):066111.
8. Davies, N. (2011). *Flat Earth news: an award-winning reporter exposes falsehood, distortion and propaganda in the global media*. Random House, London.
9. Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, 486(3):75–174.
10. Fortunato, S. and Barthelemy, M. (2007). Resolution limit in community detection. *Proceedings of the National Academy of Sciences*, 104(1):36–41.
11. Giannopoulos, G., Koniaris, M., Weber, I., Jaimes, A., and Sellis, T. (2014). Algorithms and criteria for diversification of news article comments. *Journal of Intelligent Information Systems*, 44(1):1–47.
12. Hampton, K., Goulet, L. S., Rainie, L., and Purcell, K. (2011). Social networking sites and our lives. Technical report, Pew Internet and American Life Project.
13. Herman, E. S. and Chomsky, N. (1988). *Manufacturing consent: The political economy of the mass media*. Random House, London.

14. Ipsos MORI (2013). Perceptions are not reality: The top 10 we get wrong. Accessed at <https://www.ipsos-mori.com/researchpublications/researcharchive/3188/Perceptions-are-not-reality-the-top-10-we-get-wrong.aspx>.
15. Jaffali, S., Jamoussi, S., and Hamadou, A. B. (2014). Grouping like-minded users based on text and sentiment analysis. In *Computational Collective Intelligence. Technologies and Applications*, pages 83–93. Springer.
16. Kovach, B. and Rosenstiel, T. (2007). *The elements of journalism: What newspeople should know and the public should expect*. Random House LLC.
17. Kriplean, T., Toomim, M., Morgan, J., Borning, A., and Ko, A. (2012). Is this what you meant?: promoting listening on the web with reflect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1559–1568. ACM.
18. McCombs, M. (2005). A look at agenda-setting: Past, present and future. *Journalism studies*, 6(4):543–557.
19. McCombs, M. E. and Shaw, D. L. (1972). The agenda-setting function of mass media. *Public opinion quarterly*, 36(2):176–187.
20. Meier, F., Elweiler, D., and Wilson, M. L. (2014). More than liking and bookmarking? towards understanding twitter favouriting behaviour. *ICWSM 2014*.
21. Mills, R. (2011). Researching social news – is reddit.com a mouthpiece for the ‘hive mind’, or a collective intelligence approach to information overload? In *ETHICOMP 2011 Proceedings*. Sheffield Hallam University.
22. A. Mitchell. State of the news media 2014: Overview. *Pew Research Journalism Project*, 2014.
23. Muchnik, L., Aral, S., and Taylor, S. J. (2013). Social influence bias: A randomized experiment. *Science*, 341(6146):647–651.
24. Nardelli, A. and Arnett, G. (2014). You are probably wrong about almost everything. *The Guardian*, 29 October 2014. Accessed at <http://www.theguardian.com/news/datablog/2014/oct/29/todays-key-fact-you-are-probably-wrong-about-almost-everything>.
25. Papadopoulos, S., Kompatsiaris, Y., Vakali, A., and Spyridonos, P. (2012). Community detection in social media. *Data Mining and Knowledge Discovery*, 24(3):515–554.
26. Parau, P., Stef, A., Lemnaru, C., Dinsoreanu, M., and Potolea, R. (2013). Using community detection for sentiment analysis. In *2013 IEEE International Conference on Intelligent Computer Communication and Processing*, pages 51–54. IEEE.
27. Schultz, J. (1998). *Reviving the fourth estate: Democracy, accountability and the media*. Cambridge University Press.
28. Scott, J., Millard, D., and Leonard, P. (2014). Citizen participation in news: An analysis of the landscape of online journalism. *Digital Journalism*, (ahead-of-print):1–22.
29. Strömbäck, J. (2008). Four phases of mediatization: An analysis of the mediatization of politics. *The International Journal of Press/Politics*, 13(3):228–246.
30. Wanta, W., Golan, G., and Lee, C. (2004). Agenda setting and international news: Media influence on public perceptions of foreign nations. *Journalism & Mass Communication Quarterly*, 81(2):364–377.
31. Whiteley, P., Clarke, H. D., Sanders, D., and Stewart, M. (2013). The economic and electoral consequences of austerity policies in britain. In *EPSA 2013 Annual General Conference Paper*, volume 138.