Political Polarization on Twitter

M. D. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, F. Menczer
Center for Complex Networks and Systems Research
School of Informatics and Computing
Indiana University, Bloomington, IN, USA

Abstract

In this study we investigate how social media shape the networked public sphere and facilitate communication between communities with different political orientations. We examine two networks of political communication on Twitter, comprised of more than 250,000 tweets from the six weeks leading up to the 2010 U.S. congressional midterm elections. Using a combination of network clustering algorithms and manually-annotated data we demonstrate that the network of political retweets exhibits a highly segregated partisan structure, with extremely limited connectivity between left- and right-leaning users. Surprisingly this is not the case for the user-to-user mention network, which is dominated by a single politically heterogeneous cluster of users in which ideologically-opposed individuals interact at a much higher rate compared to the network of retweets. To explain the distinct topologies of the retweet and mention networks we conjecture that politically motivated individuals provoke interaction by injecting partisan content into information streams whose primary audience consists of ideologically-opposed users. We conclude with statistical evidence in support of this hypothesis.

1 Introduction

Social media play an important role in shaping political discourse in the U.S. and around the world (Bennett 2003; Benkler 2006; Sunstein 2007; Farrell and Drezner 2008; Aday et al. 2010; Tumasjan et al. 2010; O'Connor et al. 2010). According to the Pew Internet and American Life Project, six in ten U.S. internet users, nearly 44% of American adults, went online to get news or information about politics in 2008. Additionally, Americans are taking an active role in online political discourse, with 20% of internet users contributing comments or questions about the political process to social networking sites, blogs or other online forums (Pew Internet and American Life Project 2008).

Despite this, some empirical evidence suggests that politically active web users tend to organize into insular, homogenous communities segregated along partisan lines. Adamic and Glance (2005) famously demonstrated that political blogs preferentially link to other blogs of the same political ideology, a finding supported by the work of Hargittai,

Gallo, and Kane (2007). Consumers of online political information tend to behave similarly, choosing to read blogs that share their political beliefs, with 26% more users doing so in 2008 than 2004 (Pew Internet and American Life Project 2008).

In its own right, the formation of online communities is not necessarily a serious problem. The concern is that when politically active individuals can avoid people and information they would not have chosen in advance, their opinions are likely to become increasingly extreme as a result of being exposed to more homogeneous viewpoints and fewer credible opposing opinions. The implications for the political process in this case are clear. A deliberative democracy relies on a broadly informed public and a healthy ecosystem of competing ideas. If individuals are exposed exclusively to people or facts that reinforce their pre-existing beliefs, democracy suffers (Sunstein 2002; 2007).

In this study we examine networks of political communication on the Twitter microblogging service during the six weeks prior to the 2010 U.S. midterm elections. Sampling data from the Twitter 'gardenhose' API, we identified 250,000 politically relevant messages (*tweets*) produced by more than 45,000 users. From these tweets we isolated two networks of political communication — the *retweet* network, in which users are connected if one has rebroadcast content produced by another, and the *mention* network, where users are connected if one has mentioned another in a post, including the case of tweet replies.

We demonstrate that the retweet network exhibits a highly modular structure, segregating users into two homogenous communities corresponding to the political left and right. In contrast, we find that the mention network does not exhibit this kind of political segregation, resulting in users being exposed to individuals and information they would not have been likely to choose in advance.

Finally, we provide evidence that these network structures result in part from politically motivated individuals annotating tweets with multiple *hashtags* whose primary audiences consist of ideologically-opposed users, a behavior also documented in the work of Yardi and boyd (2010). We argue that this process results in users being exposed to content they are not likely to rebroadcast, but to which they may respond using mentions, and provide statistical evidence in support of this hypothesis.

The major contributions of this work are:

- Creation and release of a network and text dataset derived from more than 250,000 politically-related Twitter posts authored in the weeks preceeding the 2010 U.S. midterm elections (§ 2).

- Cluster analysis of networks derived from this corpus showing that the network of retweet exhibits clear segregation, while the mention network is dominated by a single large community (§ 3.1).

- Manual classification of Twitter users by political alignment, demonstrating that the retweet network clusters correspond to the political left and right. These data also show the mention network to be politically heterogeneous, with users of opposing political views interacting at a much higher rate than in the retweet network (§ 3.3).

- An interpretation of the observed community structures based on injection of partisan content into ideologically opposed hashtag information streams (§ 4).

2 Data and Methods

2.1 The Twitter Platform

Twitter is a popular social networking and microblogging site where users can post 140-character messages, or *tweets*. Apart from broadcasting tweets to an audience of *followers*, Twitter users can interact with one another in two primary public ways: retweets and mentions. Retweets act as a form of endorsement, allowing individuals to rebroadcast content generated by other users, thereby raising the content's visibility (boyd, Golder, and Lotan 2008). Mentions function differently, allowing someone to address a specific user directly through the public feed, or, to a lesser extent, refer to an individual in the third person (Honeycutt and Herring 2008). These two means of communication —retweets and mentions— serve distinct and complementary purposes, together acting as the primary mechanisms for explicit, public user-user interaction on Twitter.

Hashtags are another important feature of the Twitter platform. They allow users to annotate tweets with metadata specifying the topic or intended audience of a communication. For example, #dadt stands for "Don't Ask Don't Tell" and #jlot for "Jewish Libertarians on Twitter." Each hashtag identifies a stream of content, with users' tag choices denoting participation in different information channels.

The present analysis leverages data collected from the Twitter 'gardenhose' API (dev.twitter.com/pages/streaming_api) between September 14th and November 1st, 2010 — the run-up to the November 4th U.S. congressional midterm elections. During the six weeks of data collection we observed approximately 355 million tweets. Our analysis utilizes an infrastructure and website (truthy.indiana.edu) designed to analyze the spread of information on Twitter, with special focus on political content (Ratkiewicz et al. 2011).

2.2 Identifying Political Content

Let us define a political communication as any tweet containing at least one politically relevant hashtag. To identify

Table 1: Hashtags related to #p2, #tcot, or both. Tweets containing any of these were included in our sample.

Just #p2	#casen #dadt #dc10210 #democrats #du1 #fem2 #gotv #kysen #lgf #ofa #onenation #p2b #pledge #rebelleft #truthout #vote #vote2010 #whyimvotingdemocrat #youcut
Both	#cspj #dem #dems #desen #gop #hcr #nvsen #obama #ocra #p2 #p21 #phnm #politics #sgp #tcot #teaparty #tlot #topprog #tpp #twisters #votedem
Just #tcot	#912 #ampat #ftrs #glennbeck #hhrs #iamthemob #ma04 #mapoli #palin #palin12 #spwbt #tsot #tweetcongress #ucot #wethepeople

Table 2: Hashtags excluded from the analysis due to ambiguous or overly broad meaning.

Excl. from #p2	#economy #gay #glbt #us #wc #lgbt
Excl. from both	#israel #rs
Excl. from #tcot	#news #qsn #politicalhumor

an appropriate set of political hashtags and to avoid introducing bias into the sample, we performed a simple tag co-occurrence discovery procedure. We began by seeding our sample with the two most popular political hashtags, #p2 ("Progressives 2.0") and #tcot ("Top Conservatives on Twitter"). For each seed we identified the set of hashtags with which it co-occurred in at least one tweet, and ranked the results using the Jaccard coefficient. For a set of tweets S containing a seed hashtag, and a set of tweets T containing another hashtag, the Jaccard coefficient between S and T is

$$\sigma(S,T) = \frac{|S \cap T|}{|S \cup T|}. \qquad (1)$$

Thus, when the tweets in which both seed and hashtag occur make up a large portion of the tweets in which either occurs, the two are deemed to be related. Using a similarity threshold of 0.005 we identified 66 unique hashtags (Table 1), eleven of which we excluded due to overly-broad or ambiguous meaning (Table 2). This process resulted in a corpus of 252,300 politically relevant tweets. There is substantial overlap between streams associated with different political hashtags because many tweets contain multiple hashtags. As a result, lowering the similarity threshold leads to only modest increases in the number of political tweets in our sample — which do not substantially affect the results of our analysis — while introducing unrelated hashtags.

2.3 Political Communication Networks

From the tweets containing any of the politically relevant hashtags we constructed networks representing political communication among Twitter users. Focusing on the two primary modes of public user-user interaction, mentions and retweets, we define communication links in the following ways. In the retweet network an edge runs from a node representing user A to a node representing user B if B retweets

content originally broadcast by A, indicating that information has propagated from A to B. In the mention network an edge runs from A to B if A mentions B in a tweet, indicating that information may have propagated from A to B (a tweet mentioning B is visible in B's *timeline*). Both networks therefore represent potential pathways for information to flow between users.

The retweet network consists of 23,766 non-isolated nodes among a total of 45,365. The largest connected component accounts for 18,470 nodes, with 102 nodes in the next-largest component. The mention network is smaller, consisting of 10,142 non-isolated nodes out of 17,752 total. It has 7,175 nodes in its largest connected component, and 119 in the next-largest. Because of their dominance we focus on the largest connected components for the rest of our analysis. We observe that the retweet and mention networks exhibit very similar scale-free topology (power-law degree distribution not shown), with a number of users receiving or spreading a huge amount of information.

3 Cluster Analysis

Initial inspection of the retweet suggested that users preferentially retweet other users with whom they agree politically, while the mention network appeared to form a bridge between users of different ideologies. We explore this hypothesis in several stages. In § 3.1 we use network clustering algorithms to demonstrate that the retweet network exhibits two highly segregated communities of users, while the mention network does not. In § 3.2 we describe a statistical analysis of political tweet content, showing that messages produced by members of the same community are more similar to each other than messages produced by users in different communities. Finally, in § 3.3, by manually annotating users, we show that the retweet network is polarized on a partisan basis, while the mention network is much more politically heterogeneous.

3.1 Community Structure

To establish the large-scale political structure of the retweet and mention networks we performed community detection using a label propagation method for two communities.[1] Label propagation (Raghavan, Albert, and Kumara 2007) works by assigning an initial arbitrary cluster membership to each node and then iteratively updating each node's label according to the label that is shared by most of its neighbors. Ties are broken randomly when they occur. Label propagation is a greedy hill-climbing algorithm. As such it is extremely efficient, but can easily converge to different suboptimal clusters dependent on initial label assignments and random tie breaking. To improve its effectiveness and stability, we seeded the algorithm with initial node labels determined by the leading-eigenvector modularity maximization method for two clusters (Newman 2006).

To confirm that we can produce consistent clusters across different runs we executed the algorithm one hundred times

Table 3: Minimum, maximum, and average ARI similarities between 4,950 pairs of cluster assignments computed by label propagation on the mention and retweet networks.

Network	Min	Max	Mean
Mention	0.80	1.0	0.89
Retweet	0.94	0.98	0.96

for each network and compared the label assignments produced by every run. Table 3 reports the high average agreement between the resulting cluster assignments for each graph, as computed by the Adjusted Rand Index (Hubert and Arabie 1985). Such a high agreement suggests that the clusters are consistent, and therefore we avoid resorting to consensus clustering for simplicity.

Figure 1 shows the retweet and mention networks, laid out using a force-directed layout algorithm (Fruchterman and Reingold 1991), with node colors determined by the assigned communities. The retweet network exhibits two distinct communities of users, while the mention network is dominated by a single massive cluster of interconnected users. *Modularity* (Newman and Girvan 2004) resulting from the cluster assignments offers a first measure of segregation, and reinforces the qualitative finding above. The modularity induced by the communities in the retweet and mention networks have values of 0.48 and 0.17, respectively.

A direct comparison of the modularity values is however problematic because of the different size and overall connectivity of the two networks. We need a way to compare the 'goodness' of cluster assignments across different graphs. To this end we generate, for both retweet and mention graphs, $N = 1000$ *shuffled* versions of the graph that preserve the original *degree sequence*.

Each randomized network is clustered with the method described above for the original graphs and associated with the resulting modularity value. We use the distribution of these values as a baseline against which to compare the quality of the clusters in the original graph. The intuition behind this approach is that the degree to which the actual graphs are more modular than the shuffled graphs tells us how amenable each is to being split into two clusters — a measure of segregation. The modularities of the shuffled graphs can be viewed as observed values of a random variable. We can use these values to compute z-scores for the modularities of the original networks; they are $z_r = 11.02$ and $z_m = 2.06$ for the retweet and mention networks, respectively. We conclude that the community structure found in the retweet network is significantly more segregated than that found in the mention network.[2]

In summary, the retweet network contains two clusters of users who preferentially propagate content within their own communities. However, we do not find such a structure in the network of mentions and replies among politically ac-

---

[1]While the partisan nature of U.S. political discourse makes two a natural number of clusters, in § 3.3 we describe the effect on our analysis of increasing the target number of communities.

[2]This discussion assumes that the modularities of the shuffled graph cluster assignments are distributed normally, which is not true in general. See Appendix A for an argument that does not need this assumption, and reaches the same conclusion.

Figure 1: The political retweet (left) and mention (right) networks, laid out using a force-directed algorithm. Node colors reflect cluster assignments (see § 3.1). Community structure is evident in the retweet network, but less so in the mention network. We show in § 3.3 that in the retweet network, the red cluster A is made of 93% right-leaning users, while the blue cluster B is made of 80% left-leaning users.

tive Twitter users. This structural difference is of particular importance with respect to political communication, as we now have statistical evidence to suggest that mentions and replies may serve as a conduit through which users are exposed to information and opinions they might not choose in advance. Despite this promising finding, the work of Yardi and boyd (2010) suggests that cross-ideological interactions may reinforce pre-existing in-group/out-group identities, exacerbating the problem of political polarization.

3.2 Content Homogeneity

The clustering described above was based only on the network properties of the retweet and mention graphs. An interesting question, therefore, is whether it has any significance in terms of the actual *content* of the discussions involved. To address this issue we associate each user with a profile vector containing all the hashtags in her tweets, weighted by their frequencies. We can then compute the cosine similarities between each pair of user profiles, separately for users in the same cluster and users in different clusters. Figure 2 shows that in the mention network, users placed in the same cluster are not likely to be much more similar to each other than users in different clusters. On the other hand, in the retweet network, users in cluster A are more likely to have very similar profiles than users in cluster B, and users in different clusters are the least similar to each other. As a result the average similarity within retweet clusters is higher than across clusters. Further, we note that in both mention and retweet networks, one of the clusters is more cohesive than the other — meaning the tag usage within one community is more homogeneous.

|	Retweet	Mention
A$\leftrightarrow$A	0.31	0.31
B$\leftrightarrow$B	0.20	0.22
A$\leftrightarrow$B	0.13	0.26

Figure 2: Cosine similarities among user profiles. The table on the left shows the average similarities in the retweet and mention networks for pairs of users both in cluster A, both in cluster B, and for users in different clusters. All differences are significant at the 95% confidence level. The plot on the right displays the actual distributions of cosine similarities for the retweet network.

3.3 Political Polarization

Given the communities of the retweet network identified in § 3.1, their content homogeneity uncovered in § 3.2, and the findings of previous studies, it is natural to investigate whether the clusters in the retweet network correspond to groups of users of similar political alignment.

To accomplish this in a systematic, reproducible way we used a set of techniques from the social sciences known as *qualitative content analysis* (Krippendorff 2004; Kolbe 1991). Similar to assigning class labels to training data in supervised machine learning, content analysis defines a set of practices that enable social scientists to define reproducible categories for qualitative features of text. Next we outline our annotation categories, and then explain the procedures

used to establish the rigor of these category definitions.

Our coding goals were simple: for a given user we wanted to identify whether his tweets express a 'left' or 'right' political identity, or if his identity is 'undecidable.' The groups primarily associated with a 'left' political identity are democrats and progressives; those primarily associated with a 'right' political identity are republicans, conservatives, libertarians and the Tea Party. A user coded as 'undecidable' may be taking part in a political dialogue, but from the content of her tweets it is difficult to make a clear determination about political alignment. Irrelevant non-English and spam accounts constitute less than 3% of the total corpus and were excluded from this analysis. We experimented with more detailed categorization rubrics but the simple definitions described above yielded the highest inter-annotator agreement in early trials of the coding process.

Using this coding scheme one author first annotated 1,000 random users who appeared in both the retweet and mention networks. Annotations were determined solely on the basis of the tweets present in the six week sample. In line with the standards of the field, we had a non-author judge with a broad knowledge of politics annotate 200 random users from the set of 1,000 to establish the reproducibility of this annotation scheme. The judge was provided a brief overview of the study and introduced to the coding guidelines described above, but did not have any other interaction with the authors during the coding process.

The statistic typically used in the social sciences to measure the extent to which a coders' annotations agree with an objective judge is Cohen's Kappa, defined as

$$\kappa = \frac{P(\alpha) - P(\epsilon)}{1 - P(\epsilon)} \qquad (2)$$

where $P(\alpha)$ is the observed rate of agreement between annotators, and $P(\epsilon)$ is the expected rate of random agreement given the relative frequency of each class label (Krippendorff 2004; Kolbe 1991). For agreement between the 'left' and 'right' categories we report $\kappa = 0.80$ and $\kappa = 0.82$ respectively, both of which fall in the "nearly perfect agreement" range (Landis and Koch 1977). For the undecidable category we found "fair to moderate" agreement ($\kappa = 0.42$), indicating that there are users for whom a political identity might be discernible in the context of specific domain knowledge. To address this issue of context-sensitive ambiguity we had a second author also annotate the entire set of 1,000 users. This allowed us to assign a label to a user when either author was able to determine a political alignment, resolving ambiguity in 15.4% of users.

For completeness we also report binomial p-values for observed agreement, treating annotation pairs as observations from a series of Bernoulli trials. Similar to the Kappa statistic results, inter-annotator agreement for the 'left' and 'right' categories is very high ($p < 10^{-12}$). Agreement on the 'undecidable' category is again lower ($p = 0.18$).

Based on this analysis it is clear that a majority of politically active users on Twitter express a political identity in their tweets. Both annotators were unable to determine a political identity in only 8% of users. A more conservative approach to label assignment does not change this story much;

Table 4: Partisan composition and size of network clusters as determined by manual inspection of 1,000 random user profiles.

Network	Clust.	Left	Right	Undec.	Nodes
Retweet	A	1.19%	93.4%	5.36%	7,115
	B	80.1%	8.71%	11.1%	11,355
Mention	A	39.5%	52.2%	8.18%	7,021
	B	9.52%	85.7%	4.76%	154

if we assign a political identity only to users for whom both annotators agree, we report unambiguous political valences for more than 75% of users.

Using these annotations we can infer the expected political makeup of the network communities identified in § 3.1. As shown in Table 4, the network of political retweets exhibits a highly partisan community structure with two homogenous clusters of users who tend to share the same political identity. Surprisingly, the mention network does not exhibit a clear partisan community structure. Instead we find that it is dominated by a politically heterogeneous cluster accounting for more than 97% of the users, suggests that politically active Twitter users may be exposed to views with which they do not agree in the form of cross-ideological mentions.

Increasing the number of target communities in the mention network does not reveal a more fine-grained ideological structure, but instead results in smaller yet politically heterogeneous clusters. Similarly, the retweet network communities are maximally homogenous in the case of two clusters.

4 Interaction Analysis

The strong segregation evident in the retweet network and the fact that the two clusters correspond to political ideologies suggest that, when engaging in political discourse, users often retweet just other users with whom they agree politically. The dominance of the mention network by a single heterogeneous cluster of users, however, suggests that individuals of different political alignments may interact with one another much more frequently using mentions. Let us test these conjectures, and propose an explanation based on selective hashtag use by politically motivated individuals.

4.1 Cross-Ideological Interactions

To investigate cross-ideological mentions, we compare the observed number of links between *manually-annotated* users with the value we would expect in a graph where users connect to one another without any knowledge of political alignment. The intuition for the expected number of links is as follows: for a set of users with k directed edges among them, we preserve the source of each edge and assign the target vertex to a random user in the graph, simulating a scenario in which users connected irrespective of political ideology. For example, if there are a total of k_R links originating from right-leaning users, and the numbers of left-leaning and right-leaning users are U_L and U_R respectively, then the

Table 5: Ratios between observed and expected number of links between users of different political alignments in the mention and retweet networks.

	Mention		Retweet	
	→ Left	→ Right	→ Left	→ Right
Left	1.23	0.68	1.70	0.05
Right	0.77	1.31	0.03	2.32

expected number of edges going from right-leaning to left-leaning users is given by:

$$E[R \to L] = k_R \cdot \frac{U_L}{U_L + U_R}. \qquad (3)$$

We compute the other expected numbers of edges ($R \to R$, $L \to R$, $L \to L$) in the same way.

In Table 5 we report the ratio between the observed and expected numbers of links between users of each political alignment. We see that for both means of communication, users are more likely to engage people with whom they agree. This effect, however, is far less pronounced in the mention network, where we observe significant amounts of cross-ideological interaction.

4.2 Content Injection

Any Twitter user can select arbitrary hashtags to annotate his or her tweets. We observe that users frequently produce tweets containing hashtags that target multiple politically opposed audiences, and we propose that this phenomenon may be responsible in part for the network structures described in this study.

As a thought experiment, consider an individual who prefers to read tweets produced by users from the political left. This user would frequently see the popular hashtag #p2 ("Progressives 2.0") in the body of tweets produced by other left-leaning users, as shown in Table 6. However, if this user clicked on the #p2 hashtag hyperlink in one of these tweets, or searched for it explicitly, she would be exposed to content from users on both sides of the political spectrum. In fact, because of the disproportionate number of tweets produced by left- and right-leaning users, nearly 30% of the tweets in the #p2 search feed would originate from right-leaning users.

A natural question is why a user would annotate tweets with hashtags strongly associated with ideologically opposed users. One explanation might be that he seeks to expose those users to information that reinforces his political views. Consider the following tweets:

```
User A: Please follow @Username for
an outstanding progressive voice! #p2
#dems #prog #democrats #tcot

User B: Couple Aborts Twin Boys For
Being Wrong Gender..http://bit.ly/xyz
#tcot #hhrs #christian #tlot #teaparty
#sgp #p2 #prolife
```

Table 6: The ten most popular hashtags produced by left- and right-leaning users in the manually annotated set of users, including frequency of use in the two retweet communities and ideological valence.

Rank	Hashtag	Left	Right	Valence
1	#tcot	2,949	13,574	0.384
2	#p2	6,269	3,153	-0.605
3	#teaparty	1,261	5,368	0.350
4	#tlot	725	2,156	0.184
5	#gop	736	1,951	0.128
6	#sgp	226	2,563	0.694
7	#ocra	434	1,649	0.323
8	#dems	953	194	-0.818
9	#twisters	41	990	0.843
10	#palin	200	838	0.343
	Total	26,341	53,880	

These tweets were selected from the first page of the realtime search results for the #tcot ("Top Conservatives on Twitter") and #p2 hashtags, respectively, and messages in this style make up a substantial portion of the results.

This behavior does not go unnoticed by users, as underscored by the emergence of the left-leaning hashtag #p21. According to a crowdsourced hashtag definition site (www.tagdef.com), #p21 is a hashtag for "Progressives sans RWNJs" and "Political progressives w/o all the RWNJ spam that #p2 has," where RWNJ is an acronym for "Right Wing NutJob." This tag appears to have emerged in response to the efforts by right-leaning users to inject messages into the high-profile #p2 content stream, and ostensibly serves as a place where progressives can once again be exposed only to content aligned with their views.

We propose that when a user is exposed to ideologically opposed content in this way, she will be unlikely to rebroadcast it, but may choose to respond directly to the originator in the form of a mention. Consequently, the network of retweets would exhibit ideologically segregated community structure, while the network of mentions would not.

4.3 Political Valence

To explore the content injection phenomenon in more detail let us introduce the notion of *political valence*, a measure that encodes the relative prominence of a tag among left- and right-leaning users. Let $N(t, L)$ and $N(t, R)$ be the numbers of occurrences of tag t in tweets produced by left- and right-leaning users, respectively. Then define the valence of t as

$$V(t) = 2 \frac{N(t, R)/N(R)}{[N(t, L)/N(L)] + [N(t, R)/N(R)]} - 1 \qquad (4)$$

where $N(R) = \sum_t N(t, R)$ is the total number of occurrences of all tags in tweets by right-leaning users and $N(L)$ is defined analogously for left-leaning users. The translation and scaling constants serve to bound the measure between -1 for a tag only used by the left, and $+1$ for a tag only used by the right. Table 7 illustrates the usefulness of this measure by listing hashtags sampled from valence quintiles ranging

Table 7: Hashtags in tweets by users across the political spectrum, grouped by valence quintiles.

| Far Left	Moderate Left	Center	Moderate Right	Far Right
#healthcare	#aarp #women	#democrats #social	#rangel #waste	#912project #twisters
#judaism #hollywood	#citizensunited	#seniors #dnc	#saveamerica	#gop2112 #israel
#2010elections	#democratic	#budget #political	#american #gold	#foxnews #mediabias
#capitalism #recession	#banksters #energy	#goproud #christian	#repeal #mexico	#constitution
#security #dreamact	#sarahpalin	#media #nobel	#terrorism #gopleader	#patriots #rednov
#publicoption	#progressives		#palin12	#abortion
#topprogs	#stopbeck #iraq			
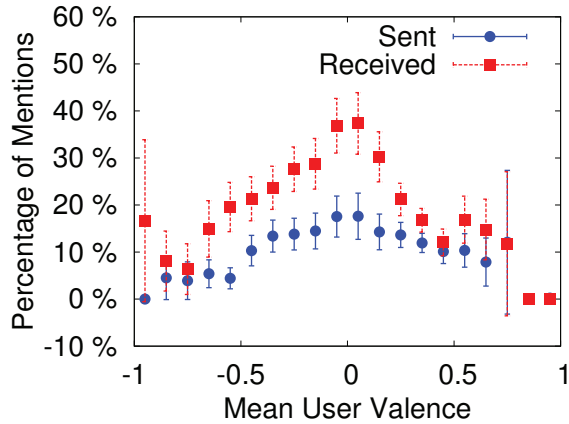
Figure 3: Proportion of mentions a user sends and receives to and from ideologically-opposed users relative to her valence. Points represent binned averages. Error bars denote 95% confidence intervals.

from the far left to the far right, where valence is computed only for hashtags produced by manually-annotated users.

If hashtag-based content injection is related to the comparatively high levels of cross-ideological communication observed in the mention network, we expect users who use hashtags in this way to receive proportionally more mentions from users with opposing political views. Using community identities in the retweet network as a proxy for political alignment, we plot in Figure 3 the average proportions of mentions users receive from and direct toward members of the other community versus the mean valence of all tags produced by those users. A key finding of this study, these results indicate that users contributing to a politically balanced combination of content streams on average receive and produce more inter-ideological communication than those who use mostly partisan hashtags. Moreover, Table 6 shows that the most popular hashtags do not have neutral valence, ruling out that neutral-valence users are simply using the most popular hashtags.

5 Conclusions

In this study we have demonstrated that the two major mechanisms for public political interaction on Twitter — mentions and retweets — induce distinct network topologies. The retweet network is highly polarized, while the mention network is not. To explain these observations we highlight the role of hashtags in exposing users to content they would not likely choose in advance. Specifically, users who apply hashtags with neutral or mixed valence are more likely to engage in communication with opposing communities.

Although our findings could be interpreted as encouraging evidence of cross-ideological political discourse, we emphasize that these interactions are almost certainly not a panacea for the problem of political polarization. While we know for certain that ideologically-opposed users interact with one another, either through mentions or content injection, they very rarely share information from across the divide with other members of their community. It is possible that these users are unswayed by opposing arguments and facts, or that the social pressures that lead to group polarization are too strong for most users to overcome (Sunstein 2002). Whatever the case, political segregation, as manifested in the topology of the retweet network, persists in spite of substantial cross-ideological interaction.

Qualitatively speaking, our experience with this body of data suggests that the content of political discourse on Twitter remains highly partisan. Many messages contain sentiments more extreme than you would expect to encounter in face-to-face interactions, and the content is frequently disparaging of the identities and views associated with users across the partisan divide. If Yardi and boyd (2010) are correct, and our experience suggests this may be the case, these interactions might actually serve to exacerbate the problem of polarization by reinforcing pre-existing political biases. Further study of the content of inter-ideological communication, including sentiment analysis, as well as studies of network topology that include the follower network, could help to illuminate this issue.

The fractured nature of political discourse seems to be worsening, and understanding the social and technological dynamics underlying this trend will be essential to attenuating its effect on the public sphere. We have released a public dataset based on the information accumulated during the course of this study, in hopes that it will help others explore the role of technologically-mediated political interaction in deliberative democracy. The dataset is available at cnets.indiana.edu/groups/nan/truthy.

A Modularity Distributions

In § 3.1 we computed the z-score of our modularity values with respect to the distribution of modularity values that arise from clustering degree-preserving shuffles of the original graph. We used these z-scores to argue that the retweet network is significantly more segregated than the mention network. However, this argument relied on the assumption that the modularities are normally distributed. By testing the sampled modularities with the omnibus K^2 statistic (D'Agostino 1971), we find that this assumption is not true: we have $p < 10^{-10}$ that variations between the observed data and the best-fit normal distribution are due to random chance. This is the case for the modularities sampled based on both the retweet and mention networks.

Fortunately we can reach the same statistical conclusion without relying on the assumption of normality. Let us use Chebyshev's inequality:

$$P(|X - \mu| \geq k\sigma) = P(|z| \geq k) \leq \frac{1}{k^2} , \qquad (5)$$

which gives us a very conservative bound on the probability that the random variable X, being the modularity of a sampled graph, will take on the value of the original graph's modularity. We therefore compute z_m and z_r for the modularities of clusters in the 1,000 shuffled versions of each of the the mention and retweet networks, respectively. Using Equation 5 and the modularities of the original graph clusters, we find:

$$P(z \geq z_m) \leq \frac{1}{z_m^2} = 0.24 \qquad (6)$$

$$P(z \geq z_r) \leq \frac{1}{z_r^2} = 0.008 \qquad (7)$$

for $z_m = 2.06$ and $z_r = 11.02$ (§ 3.1). Thus, since a network that can be clustered as well as the retweet network is much less likely to arise randomly (relative to the mention network), we confirm that the retweet network is much more segregated.

References

Adamic, L., and Glance, N. 2005. The political blogosphere and the 2004 U.S. election: Divided they blog. In *Proc. 3rd Intl. Workshop on Link Discovery (LinkKDD)*, 36–43.

Aday, S.; Farrel, H.; Lynch, M.; Sides, J.; Kelly, J.; and Zuckerman, E. 2010. Blogs and bullets: New media in contentious politics. Technical report, U.S. Institute of Peace.

Benkler, Y. 2006. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press.

Bennett, L. 2003. New media power: The Internet and global activism. In Couldry, N., and Curran, J., eds., *Contesting media power: Alternative media in a networked world*. Rowman and Littlefield. 17–37.

boyd, d.; Golder, S.; and Lotan, G. 2008. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. *Proc. Hawaii Intl. Conf. on Systems Sciences* 1–10.

D'Agostino, R. 1971. An omnibus test of normality for moderate and large size samples. *Biometrika* 58(2):341–348.

Farrell, H., and Drezner, D. 2008. The power and politics of blogs. *Public Choice* 134(1):15–30.

Fruchterman, T. M. J., and Reingold, E. M. 1991. Graph drawing by force-directed placement. *Software Practice and Experience* 21(11):1129–1164.

Hargittai, E.; Gallo, J.; and Kane, M. 2007. Cross-ideological discussions among conservative and liberal bloggers. *Public Choice* 134(1):67–86.

Honeycutt, C., and Herring, S. C. 2008. Beyond microblogging: Conversation and collaboration via Twitter. In *Proc. 42nd Hawaii Intl Conf. on System Sciences*.

Hubert, L., and Arabie, P. 1985. Comparing partitions. *Journal of Classification* 2(1):193–218.

Kolbe, R. H. 1991. Content analysis research: An examination of applications with directives for improving research reliability and objectivity. *The Journal of Consumer Research* 18(2):243–250.

Krippendorff, K., ed. 2004. *Content Analysis: An Introduction to Its Methodology*. Sage Publications.

Landis, J., and Koch, G. 1977. The measurement of observer agreement for categorical data. *Biometrics* 33(1):154–165.

Newman, M. E. J., and Girvan, M. 2004. Finding and evaluating community structure in networks. *Physical review E* 69(2):26113.

Newman, M. E. J. 2006. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E* 74(3):036104.

O'Connor, B.; Balasubramanyan, R.; Routledge, B. R.; and Smith, N. A. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Proc. 4th Intl. AAAI Conf. on Weblogs and Social Media (ICWSM)*.

Pew Internet and American Life Project. 2008. Social networking and online videos take off: Internet's broader role in campaign 2008. Technical report, Pew Research Center.

Raghavan, U. N.; Albert, R.; and Kumara, S. 2007. Near linear time algorithm to detect community structures in large-scale networks. *Physical Review E* 76(3):036106.

Ratkiewicz, J.; Conover, M.; Meiss, M.; Gonçalves, B.; Patil, S.; Flammini, A.; and Menczer, F. 2011. Truthy: Mapping the spread of astroturf in microblog streams. In *Proc. 20th Intl. World Wide Web Conf. (WWW)*.

Sunstein, C. 2002. The law of group polarization. *Journal of Political Philosophy* 10(2):175–195.

Sunstein, C. R. 2007. *Republic.com 2.0*. Princeton University Press.

Tumasjan, A.; Sprenger, T. O.; Sandner, P. G.; and Welpe, I. M. 2010. Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. In *Proc. 4th Intl. AAAI Conf. on Weblogs and Social Media (ICWSM)*.

Yardi, S., and boyd, d. 2010. Dynamic debates: An analysis of group polarization over time on twitter. *Bulletin of Science, Technology and Society* 20:S1–S8.