This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Garimella, Kiran; Gionis, Aristides; De Francisci Morales, Gianmarco; Mathioudakis, Michael
The effect of collective attention on controversial debates on social media

Published in:
WebSci 2017 - Proceedings of the 2017 ACM Web Science Conference

DOI:
[10.1145/3091478.3091486](https://doi.org/10.1145/3091478.3091486)

Published: 25/06/2017

Document Version
Peer reviewed version

Please cite the original version:
Garimella, K., Gionis, A., De Francisci Morales, G., & Mathioudakis, M. (2017). The effect of collective attention on controversial debates on social media. In *WebSci 2017 - Proceedings of the 2017 ACM Web Science Conference* (pp. 43-52). ACM. <https://doi.org/10.1145/3091478.3091486>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

The Effect of Collective Attention on Controversial Debates on Social Media

Kiran Garimella
Aalto University
Helsinki, Finland
kiran.garimella@aalto.fi

Aristides Gionis
Aalto University
Helsinki, Finland
aristides.gionis@aalto.fi

Gianmarco De Francisci Morales
Qatar Computing Research Institute
Doha, Qatar
gdfm@acm.org

Michael Mathioudakis
Aalto University
Helsinki, Finland
michael.mathioudakis@aalto.fi

ABSTRACT

We study the evolution of long-lived controversial debates as manifested on Twitter from 2011 to 2016. Specifically, we explore how the *structure of interactions* and *content of discussion* varies with the level of collective attention, as evidenced by the number of users discussing a topic. Spikes in the volume of users typically correspond to external events that increase the public attention on the topic, for instance, discussions about ‘gun control’ often erupt after a mass shooting.

This work is the first to study the dynamic evolution of polarized online debates at such scale. By employing a wide array of network and content analysis measures, we find consistent evidence that increased collective attention is associated with increased network polarization and network concentration within each side of the debate; and overall more uniform lexicon usage across all users.

ACM Reference format:

Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. The Effect of Collective Attention on Controversial Debates on Social Media. In *Proceedings of WebSci '17, Troy, NY, USA, June 25-28, 2017*, 10 pages.

DOI: <http://dx.doi.org/10.1145/3091478.3091486>

1 INTRODUCTION

Social media are a major venue of public discourse today, hosting the opinions of hundreds of millions of individuals. Due to their prevalence they have become an invaluable instrument in the study of social phenomena and a fundamental subject of *computational social science*. In this work, we study discussions around issues that are deemed important at a societal level — and in particular, ones that are controversial. This work is a step towards understanding how the discussion about controversial topics on social media evolves, and more broadly how these topics shape the discussion at a societal and political level [1, 15, 22].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '17, Troy, NY, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM.
978-1-4503-4896-6/17/06...\$15.00

DOI: <http://dx.doi.org/10.1145/3091478.3091486>

We study how online discussions around controversial topics change as interest in them increases and decreases. We are motivated by the observation that interest in enduring controversial issues is re-kindled by external events, e.g., when a major related story is reported. One typical example is the gun control debate in U.S., which is revived whenever a mass shooting occurs.¹ The occurrence of such an event commonly causes an increase in collective attention, e.g., in volume of related activity in social media.

Given a controversial topic, our focus is to analyze the interactions among users involved in the discussion, and quantify how certain structural properties of the interaction network vary with the change in volume of activity. Our main finding is that the polarization reflected in the network structure of online interactions is correlated with the increase in the popularity of a topic.

Differently from previous studies, we study the *dynamic* aspects of controversial topics on social media. While the evolution of networks and polarization on social media have been studied in the past [7, 20], they have not been studied in conjunction before. In addition, we seek to understand the *response* of social media to stimuli that cause increased interest in the topics, an issue that only very recently has seen some attention [28].

We take a longitudinal approach and collect data from Twitter that covers approximately five years. This dataset gives us a very fine-grained view of the activity on social media, including the structure of the interactions among users, and the content they produced during this period. We track four topics of discussion that are controversial in the U.S., recurring, and have seen considerable attention during the 2016 U.S. elections.

Our methodology relies on recent advances in quantifying controversy on social media [11]. We build two types of networks: an *endorsement* network from the retweet information on Twitter, and a *communication* network from the replies. We aggregate the data at a daily level, thus giving rise to a time series of interaction graphs. Then, we identify the sides of a controversy via graph clustering, and find the *core* of the network, i.e., the users who are consistently participating to the online discussion about the topic. Finally, we employ a wide array of measures that characterize the discussion about a topic on social media, both from the point of view of the network structure and of the actual content of the posts.

Apart from our main result — an increase in polarization linked to increased interest — we also report on several other findings.

¹See, e.g., <http://slate.me/1NswLLD>.

We find that most of the interactions during events of interest happen within the different controversy sides, and replies do not cross sides very often, in line with previous observations [29]. In addition, increased interest does not alter the fundamental structure of the endorsement network, which is hierarchical, with a disproportionately large fraction of edges linking the periphery to the core. This finding suggests that most casual users, who seldom participate in the discussion, endorse opinions from the core of the side they belong to. When looking at the content of the posts on the two sides of a controversy, we find a consistent trend of *convergence*, as the lexicons become both more uniform and more similar to each other. This finding indicates that while the discussion is still controversial, both sides of the debate focus over the same fundamental issues brought under the spotlight by the event at hand. Conversely, we do not find a consistent long-term trend in the polarization of discussions, which contradicts the common narrative that our society is becoming more divided over time. Finally, we perform similar measurements for a set of topics that are non-political and non-controversial, and highlight differences with the results for controversial discussions.²

2 RELATED WORK

A few studies exist on the topic of controversy in online news and social media. In one of the first papers, Adamic and Glance [2] study linking patterns and topic coverage of political bloggers, focusing on blog posts on the U.S. presidential election of 2004. They measure the degree of interaction between liberal and conservative blogs, and provide evidence that conservative blogs are linking to each other more frequently and in a denser pattern. These findings are confirmed by a more recent study of Conover et al. [7], who focus on political communication regarding congressional midterm elections. Using data from Twitter, they identify a highly segregated partisan structure (present in the retweet graph, but not in the mention graph), with limited connectivity between left- and right-leaning users. In another recent work, Mejova et al. [25] consider discussions of controversial and non-controversial news over a span of 7 months. They find a significant correlation between controversial issues and the use of negative affect and biased language. More recently, Garimella et al. [11] show that controversial discussions on social media have a well-defined structure, when looking at the *endorsement* network. They propose a measure based on random walks (RWC), which is able to identify controversial topics, and *quantify* the level of controversy of a given discussion via its network structure alone.

The aforementioned studies focus on static networks, which are a snapshot of the underlying dynamic networks. Instead, we are interested in network dynamics and, specifically, in how it responds to increased collective attention in the controversial topic.

Several studies have looked at how networks evolve, and proposed models of network formation [19, 20]. Densification over time is a pattern often observed [20], i.e., social networks gain more edges as the number of nodes grows. A change in the scaling behavior of the degree distribution has also been observed [3]. Newman et al. [27] offer a comprehensive review. Most of these studies focus on social networks, and in particular, on the friendship relationship.

In our work, we are interested in studying an *interaction* network, which has markedly different characteristics.

There is a large amount of literature devoted to studying the evolution of networks. For an overview, see the book by Dorogovtsev and Mendes [10]. However, none of these previous studies has devoted much attention to the evolution of interaction networks for controversial topics, especially when tracking topics for a long period of time.

DiFonzo et al. [9] report on a user study that shows how the network structure affects the formation of stereotypes when discussing controversial topics. They find that segregation and clustering lead to a stronger “echo chamber” effect, with higher polarization of opinions. Our study examines a similar correlation between polarization and network structure, although in a much wider context, and focusing on the influence of external events.

Garimella and Weber [14] study polarization on Twitter over a long period of time, using content and network related measures for polarization and find that over the past decade, polarization on Twitter has increased considerably.

Perhaps the closest work to this paper is the work by Smith et al. [29], who study the role of social media in the discussion of controversial topics. They try to understand how positions on controversial issues are communicated via social media, mostly by looking at user level features such as retweet and reply rates, url sharing behavior, etc. They find that users spread information quicker if it agrees with their position and that Twitter debates may not play a big role in deciding the outcome of a controversial issue (like a vote).

However, there are a few differences with our work: (i) they study one local topic (California ballot), over a small period of time, while we study a wide range of popular topics, spanning multiple years; and (ii) their analysis is mostly user centric, where as we take a global viewpoint, constructing and analyzing networks of user interaction.

The effect of external events on social networks. A few studies have examined the effects of events on social networks. Romero et al. [28] study the behavior of a hedge-fund company via the communication network of their instant messaging systems. They find that in response to external shocks, i.e., when stock prices change significantly, the network “turtles up,” strong ties become more important, and the clustering coefficient increases. In our case, we examine both a communication network and an endorsement network, and we focus on controversial issues. Given the different setting, many of our findings are quite different.

Other works, such as the ones by Lehmann et al. [18] and Wu and Huberman [30], examine how collective attention focuses on individual topics or items and evolves over time. Lehmann et al. [18] examine spikes in the frequency of hashtags and whether most frequency volume appears before or after the spike. They find that the observed patterns point to a classification of hashtags, that agrees with whether the hashtags correspond to topics that are endogenously or exogenously driven. Wu and Huberman [30], on the other hand, examine items posted on digg.com and how their popularity decreases over time.

Morales et al. [26] study polarization over time for a single event, the death of Hugo Chavez. Our analysis has a more broad spectrum,

²A limited subset of our results appeared in a poster at ICWSM 2017 [12].

Table 1: Keywords for the controversial topics.

Topic	Keywords	#Tweets	#Users
Obamacare	obamacare, #aca	866 484	148 571
Abortion	abortion, prolife, prochoice, anti-abortion, pro-abortion, planned parenthood	1 571 363	327 702
Gun Control	gun control, gun right, pro gun, anti gun, gun free, gun law, gun safety, gun violence	824 364	224 270
Fracking	fracking, #frack, hydraulic fracturing, shale, horizontal drilling	2 117 945	170 835

as we establish common trends across several topics, and find strong signals linking the volume of interest to the degree of polarization in the discussion.

Andris et al. [4] study the partisanship of the U.S. congress over a long period of time. They find that partisanship (or non-cooperation) in the U.S. congress has been increasing exponentially for over 60 years. Our study suggests that increased controversy is linked to an increase in attention on a topic, while we do not see a global trend over time.

3 DATASET

Our study uses data collected from Twitter. Using the repositories of the Internet Archive,³ we collect a 1% sample of tweets from September 2011 to August 2016,⁴ for four topics of discussion, related to ‘Obamacare’, ‘Abortion’, ‘Gun Control’, and ‘Fracking’⁵. These topics constitute long-standing controversial issues in the U.S.⁶ and have been used in previous work [24]. For each topic, we use a keyword list as proposed by Lu et al. [24] (shown in Table 1), and extract a base set of tweets containing at least one topic-related keyword. To enrich this original dataset, we use the Twitter REST API to obtain all tweets of users who have participated in the discussion at least once.⁷ Note that since we are only looking for peaks in volumes of interest for a topic, not being able to obtain all the tweets of a user does not introduce any biases in our results. Table 1 shows the final statistics for the dataset.

We infer two types of interaction network from the dataset: (i) a retweet network — a directed endorsement network of users, where there is an edge between two users ($u \rightarrow v$) if u retweets v , and (ii) a reply network — a directed communication network of users, where an edge ($u \rightarrow v$) indicates that user u has replied to a tweet by user v . Note that replies are characterized by a tweet starting with ‘@username’ and do not include mentions and retweets.⁸

³<https://archive.org/details/twitterstream>

⁴To be precise, we have data for 57 months from that period, as our data source has no data available for three months.

⁵The term ‘fracking’ is also sometimes used as an expletive, which might explain the effects we measure are not as pronounced for this topic as the other ones.

⁶According to <http://2016election.procon.org>.

⁷Up to 3200 due to limits imposed by the Twitter API.

⁸See also <https://support.twitter.com/articles/14023> for terminology related to different types of Twitter messages.

Polarized networks, especially the ones considered here, can be broadly characterized by two opposing *sides*, which express different opinions on the topic at hand. It is commonly understood that retweets indicate endorsement, and endorsement networks for controversial topics have been shown to have a bi-clustered structure [7, 11], i.e., they consist of two well-separated clusters that correspond to the opposing points of view on the topic. Conversely, replies can indicate discussion, and several studies have reported that users tend to use replies to talk across the sides of a controversy [6, 23]. These two types of network capture different dynamics of activity, and allow us to tease apart the processes that generate these interaction networks.

In this paper, we build upon the observation that the clustering structure of retweet networks reveals the opposing sides of a topic. In particular, following an approach from previous work [11], we collapse all retweets contained in the dataset of each topic into a single large static retweet network. Then, we use the METIS clustering algorithm [16] to identify two clusters that correspond to the two opposing sides. This process allows us to identify more consistent sides for the topic. We evaluate the sides by manual inspection of the top retweeted users, URLs, and hashtags. The results are consistent and accurate, and can be inspected online.⁹

Let us now consider the temporal dynamics of these interaction networks. Given the traditional daily news reporting cycle, we build the time series of networks with the same daily granularity. This high resolution allows us to easily discern the level of interest in the topic, and possibly identify spikes of interest linked to real world external events, as shown in Figure 1. These spikes usually correspond to external newsworthy events, as shown by the annotations. These results support the observation that Twitter is used as an *agorá* to discuss the daily matters of public interest [8].

As shown in Figure 1, the size of the active network for each day varies significantly. There is, however, a *hard core* set of active users who are involved in the discussion of these controversial topics most of the time. Therefore, to understand the role of these more engaged users, we define the ‘core network’ as the one induced by users who are active for more than $\frac{3}{4}$ of the observation time. Specifically, to build a *core* set of users, we first identify two subsets — one consisting of those users who generated or received a retweet at least once per month for 45 months; and another one defined similarly for replies. We define the core set of users as the union of the aforementioned two sets. Nodes of any network that do not belong to the core are said to belong to the *periphery* of that network. The size of the core ranges from around 600 to 2800 nodes for the four topics. For any given day, the core accounts for at most around 10% of the active users.

3.1 Notation

The set of retweets that occur within a single day d gives rise to one retweet network N_d^{rt} . Each user associated with a retweet is represented with one node in the network. There is a directed edge from user u to user v only when user u has retweeted at least one tweet authored by user v . Correspondingly, the set of replies that occur within a single day give rise to a reply network N_d^{re} . In addition, each node u in the network is associated with a binary

⁹<https://mmathioudakis.github.io/polarization/>

Figure 1: Daily trends for number of active users for the four controversial topics under study. Clear spikes occur at several points in the timeline. Manually chosen labels describing related events reported in the news on the same day are shown in blue for some of the spikes.

attribute $c(u) \in \{\text{true}, \text{false}\}$ that indicates whether the node is part of the core, and an attribute $s(u) \in \{1, 2\}$ that represents the side the node belongs to. In some cases, we consider undirected versions of the networks defined above. In such cases, we write G_d^{rt}, G_d^{re} to denote the undirected graphs corresponding to N_d^{rt}, N_d^{re} , respectively.

Besides these two types of network, for each day we consider the set of tweets that were generated on that day. Every tweet m is associated with an attribute $s(m) \in \{1, 2\}$ that indicates the side its author belongs to. Moreover, every tweet m is associated with the list of words $w(m)$ that occur in its text. This information gives rise to two unigram distributions W_d^1 and W_d^2 , one for each side. Each distribution expresses the number of times each word appears in the tweets of nodes from each side.

4 MEASURES

For each day d , we employ a set of measures on the associated networks N_d^{rt}, N_d^{re} , and unigram distributions W_d^1 and W_d^2 . We describe them below.

Polarization. We quantify the polarization of a network N_d by using the random-walk controversy (RWC) score introduced in previous work [11]. Intuitively, the score captures whether the network consists of two well-separated clusters.

Clustering coefficient. In an undirected graph, the clustering coefficient $cc(u)$ of a node u is defined as the fraction of closed triangles in its immediate neighborhood. Specifically, let d be the degree of node u , and T be the number of closed triangles involving u and two of its neighbors, then

$$cc(u) = \frac{2T}{d(d-1)}.$$

In our case, we consider the undirected graph G_d and compute the average clustering coefficient of all nodes that belong to each *side* – then take the mean of the two averages as the clustering coefficient of the network.

In order to control for scale effects, i.e., correlation between the size of the network (as determined by the volume of users active on day d) and the clustering coefficient, we employ a normalizer for the score. More in detail, we use an Erdős-Rényi graph as null model (with edges drawn at random among pair of nodes), and normalize the score by the expected value for a null-model graph of the same size. Unless otherwise specified, we apply the same type of normalization for all the methods defined below.

Tie strength. For each node u in a graph G_d , we consider all nodes v it is connected to across all days, and order them decreasingly by the number of occurrences $|\{d : (u, v) \in G_d\}|$. That is, the node v at the top of the list for u is the node to which v connects the most consistently throughout the time span of the dataset. Then,

we define the *strong ties* of a node u as the top 10% of the nodes ordered as described above. For a given day d , we define the tie strength of a node as the number of strong ties it is connected to in the corresponding graph G_d . The tie-strength measure for the day is defined as the average tie strength for all nodes on either side. As described for the previous measure above, we normalize the reported measure by the expected value for a random graph with the same number of nodes and edges.

Cross-side openness. This measure reports the number of edges that connect nodes from opposing sides, and captures the inter-side interaction happening in the network on a given day. Formally, it is defined as

$$CSO = |\{(u, v) \in G_d : s(u) \neq s(v)\}|.$$

We apply the same normalization as described above.

Sides edge composition. For a given network, we distinguish two types of edges: *within-sides*, where both adjacent nodes belong to the same side, and *across-sides*, where the adjacent nodes belong to different sides. For each day and network, we track the fraction of the two types of edges.

Core-periphery openness. This measure is defined as the number of edges that connect a node from the core to the periphery. It captures the amount of interaction between the hard core users and the casual ones. Formally,

$$CPO = |\{(u, v) \in G_d : c(u) \wedge \neg c(v)\}|.$$

Bimotif. For a network N_d , we define the bimotif measure as the number of directed edges $(u, v) \in N_d$ for which the opposite edge (v, u) also appears in the network

$$Bimotif = |\{(u, v) \in N_d : (v, u) \in N_d\}|.$$

This measure captures the mutual interactions happening within the network.¹⁰

Core Density. This measure captures the number of edges that connect exclusively members of the core

$$CoreDens = |\{(u, v) \in N_d : c(u) \wedge c(v)\}|.$$

Core-periphery edge composition. For a given network, we distinguish three types of edges: *core-core*, where both adjacent nodes belong to the core we have identified, *core-periphery*, where one node belongs to the core and one to the periphery, and *periphery-periphery*, where both nodes belong to the periphery. For each day and network, we track the fraction of each type of edges.

Cross-side content divergence. This measure captures the difference between the word distributions W_d^1 and W_d^2 , and is based on the Jensen-Shannon divergence [21]. The Jensen-Shannon divergence is undefined when one of the two distributions is zero at a point where the other is not. Thus, we smooth the distributions by adding Laplace counts $\beta = 10^{-5}$ to avoid zero entries in either distribution.

The traffic volume on a given day can increase the vocabulary size, and thus induce an unwanted bias in the measure. In order to counter this bias, we employ a sampling procedure similar to bootstrapping from the two distributions. For each smoothed distribution W_d^1 and W_d^2 , we sample with replacement $k = 10\,000$ words

at random, and compute the Jensen-Shannon divergence of these equal-sized samples. We repeat the process 100 times and report the average sample Jensen-Shannon divergence as the ‘cross-side content divergence’ for day d . Intuitively, the higher its value, the more different the word distributions across the two sides.

Within-side entropy. This measure captures how ‘concentrated’ each of the two distributions W_d^1 and W_d^2 is. For each side, we compute the entropy for each distribution. The higher its value, the more widely spread is the corresponding distribution. We use the same bootstrap sampling method described above to avoid bias due to activity volume.

Topic variance. This measure captures, to some extent, *what* is being talked on the two sides of the discussion. We extract a large number of topics by using Latent Dirichlet Allocation ($k=100$) on the complete tweet corpus. We then compute the distribution of topics in each bucket. This distribution gives an estimate on which of the 100 topics are being talked about in the bucket. We report the variance of this distribution. If the distribution is focused on a small number of topics, the variance is high. Conversely, a low variance indicates a uniform distribution of topics.

Sentiment variance. This measure captures the variance of sentiment valence (positive versus negative) in all the tweets of one day d [11].

Psychometric analysis. To understand if there are behavioral changes in terms of content generated and shared by users with increasing activity, we use the Linguistic Inquiry and Word Count (LIWC) dictionary,¹¹ which identifies emotions in words [17]. We measure the fraction of tweets containing the LIWC categories: anger, sadness, posemo, negemo, and anxiety.

4.1 Analysis

We explore how the aforementioned measures vary with the number of active users in the networks, which is a proxy for the amount of collective attention the topic attracts. We sort the time series of networks by volume of active users, and partition it into ten quantiles (each having an equal number of days), so that days of bucket i are associated with smaller volume than those of bucket j , for $j > i$. For each bucket, we report the mean and standard deviation of the values for each measure, and observe the trend from lower to higher volume.

Note that the measures presented in this section are carefully defined so that their expected value does not depend on the volume of underlying activity (i.e., number of network nodes and edges or vocabulary size).

5 FINDINGS

In what follows, we report our findings on the measures defined in Section 4 — starting from the ones related to the retweet and reply networks (Section 5.1), then proceeding to the ones related to content (Section 5.2) and network cores (Section 5.3). We provide additional analysis for the periods around the spikes in interest (Section 5.4), as well as for the evolution of measures over time (Section 5.5).

¹⁰The measure is also referred to as ‘reciprocity’ in the literature.

¹¹<http://liwc.net>

Figure 2: RWC score as a function of the activity in the retweet network. An increase in interest in the controversial topic corresponds to an increase in the controversy score of the retweet network.

5.1 Network

We observe a significant correlation between RWC score and interest in the topic. Figure 2 shows the RWC score as a function of the quantiles of the network by retweet volume (as explained in the previous section). There is a clear increasing trend, which is consistent across topics. This trend suggests that increased interest in the topic is correlated with an increase in controversy of the debate, and increased polarization of the retweet networks for the two sides. Conversely, reply networks are sparser and more disconnected, thus, the RWC score is not meaningful in this case (not shown due to space constraints). This difference is expected, and was already observed in the work that introduced RWC [11].

A similar result can be observed for the clustering coefficient, as shown in Figure 3. As the interest in the topic increases, the two sides tend to *turtle up*, and form a more close-knit retweet network. This result suggests that the *echo chamber* phenomenon gets stronger when the discussion sparks. Our finding is also consistent with results by Romero et al. [28]. As for the previous measure, the clustering coefficient does not show a significant pattern for the reply networks. Replies are often linked to dyadic interactions, while the clustering coefficient measures triadic ones, so we expect such a difference between the two types of network.

In line with the above results, tie strength is correlated with retweet volume, as indicated by Figure 4. When the discussion intensifies, users tend to endorse the opinions of their closest friends, or their trusted sources of information. Again, this observation indicates a closing up of both sides when the debate gets heated. Interestingly, a similar trend is present for the reply network, as shown in Figure 5. Differently from previous work, we find an increase of communication of users with their strong ties, rather than with weak ties or users of the opposing side. We also observe an increase in back-and-forth communication, indicating a dialogue between users of the same side. Figure 6 shows an increase in bimotifs in the reply network when the discussion intensifies. This measure is inconclusive for the retweet network, for the reasons mentioned above.

Finally, when calculating the fractions of *within-side edges* and *across-side edges* for *across sides edge composition*, we find that reply networks typically contain higher proportions of across-side activity compared to retweet networks, consistently with earlier work. In fact, for retweet networks, almost all edges are classified as *within-side edges*. Interestingly, we also find that these proportions do not change significantly as the volume increases. The same is true for the *cross-side openness* measure (not shown).

Figure 3: Average clustering coefficient as a function of the activity in the retweet network. Spikes in interest correspond to an increase in the clustering coefficient on both sides of the discussion, which indicates the retweet networks tend to close up.

Figure 4: Tie strength as a function of the activity in the retweet network. Spikes in activity correspond to more interaction with stronger ties, which indicates a closing up of the retweet network.

Figure 5: Tie strength as a function of the activity in the reply network. Users tend to communicate proportionally more with closer ties when interest spikes, which reveals a further closing up of the network.

Figure 6: Bimotifs as a function of the activity in the reply network. Users tend to reciprocate the communication more as the discussion intensifies.

5.2 Content

Let us now switch our attention to the content measures. Recall that for these measures we do not distinguish between retweet and reply networks, but only between the two sides of the discussion. The main observation is that the Jensen-Shannon divergence between the two sides decreases, as shown by Figure 7. This decrease indicates that the lexicon of the two sides tend to converge. The cause of this phenomenon might be the participation of casual users to the discussions, who contribute a more general lexicon to the discussion. Alternatively, the cause might be in the event that sparks the discussion, which brings the whole network to adopt similar lexicon to speak about it, i.e., there is an event-based convergence.

To further examine the cause of the convergence of lexicon, we report the entropy of the unigram distribution. Figure 8 shows that the entropy for one of the sides increases as interest increases

Figure 7: Jensen-Shannon divergence of the lexicon between the two sides as a function of network activity. As the interest in the topic rises, the lexicon used by the two sides tends to converge.

Figure 8: Entropy of the distribution over the lexicon for one side of the discussion as a function of the activity in the network (the other side shows similar patterns). As the interest increases, the entropy increases, thus indicating the use of a wider lexicon.

Figure 9: Variance of the topic distribution. As the interest increases, variance decreases, indicating that a wider range of topics are being discussed.

(results for the other side show similar trends). Thus, we find that the lexicon is more uniform and less skewed, which supports the hypothesis that a larger group of users brings a more general lexicon to the discussion, rather than the alternative hypothesis of event-based convergence.

To investigate *what* causes the lexicon to be generalized, we compute the variance of the topic distribution for each bucket. As we see from Figure 9, the variance decreases with increased activity, meaning that the topic distribution becomes more uniform. This results provides evidence that users do indeed discuss a wider range of topics when there is a spike in activity.

Finally, we also examine how the sentiment and other linguistic cues change with interest. We measure the variance in sentiment, fraction of tweets containing various LIWC categories, such as anger, sadness, positive and negative emotion, and anxiety. Previous work shows that sentiment variance is a measure able to separate controversial from non-controversial topics [11] and linguistic patterns of communication change during shocks [28]. However, we do not see any consistent trend. We hypothesise that this might be due to the noise in language (slang, sarcasm, short text, etc) on social media.

Figure 10: Edge composition as a function of network activity in the retweet network. As the interest increases, there are no major changes in the fractions of core-core (blue), core-periphery (green), and periphery-periphery (red) edges.

Figure 11: Edge composition as a function of network activity in the reply network. As the interest increases, there are no major changes in the fractions of core-core (blue), core-periphery (green), and periphery-periphery (red) edges.

5.3 Core

Looking at the fractions of the different types of edges (core–core, core–periphery, and periphery–periphery) across the volume buckets in Figures 10 and 11, we see that the composition of edges does not change significantly with increase collective attention. This result suggests that the discussion grows in a self-similar way.

A disproportionately large fraction of edges link the periphery to the core, when taking into account the core size, as seen in Figure 10. This finding suggests that, during a spike in interest, most casual users, who seldom participate in the discussion, endorse opinions from the core of the side they belong to. For replies, we see a similar trend with respect to activity volume in Figure 11. In general, the core is less prevalent in the discussion, as shown by the lower fraction of core-periphery edges (green bars).

However, when looking at the *core–periphery openness* (Figures 12 and 13), we see that the *normalized* number of edges between core and periphery increases, i.e., the number of edges between core and periphery increases compared to the expected number based on a random-graph null model. To interpret this result, note that when the network grows, given that the periphery is much larger than the core, most edges for the null model are among periphery nodes. Therefore, the interaction networks show a clear hierarchical structure when growing.

5.4 Local analysis

So far, we have analyzed global trends across the time series. We now focus on local trends, to drill down on what happens around the spikes, and look at local variations of the measures just before and after the spike. We mark a day in the time series as a spike if the volume of active users is at least two standard deviations above the mean. Table 2 shows the Pearson correlation between various measures and network activity, one week before and after the spike. The trends observed globally still hold. There is a positive correlation of RWC with activity, which adds more evidence to our

Figure 12: Core-periphery openness as a function of activity in the retweet network. As the interest increases, the number of core-periphery edges, normalized by the expected number of edges in a random network, increases. This suggests a propensity of periphery nodes to connect with the core nodes when interest increases.

Figure 13: Core-periphery openness as a function of activity in the reply network. As the interest increases, the number of core-periphery edges, normalized by the expected number of edges in a random network, increases for most topics. This suggests a propensity of periphery nodes to connect with the core nodes when interest increases.

Table 2: Pearson correlation of various measures with volume one week before, during and after a spike in interest. All values except those marked with an asterisk (*) are significant at $p < 0.05$.

Measure	Obamacare	Gun Control	Abortion	Fracking
RWC	0.20	0.21	0.19	0.23
Openness	-0.09*	0.81	0.23	0.08
Bimotif	0.27	0.36	0.33	0.23
Tie Strength	0.96	0.98	0.95	0.86
JSD	-0.66	-0.86	-0.63	-0.46
Entropy	0.42	0.46	0.67	0.26
Frac. RT	0.15*	0.6	0.59	0.56
Frac. Men.	0.20	0.71	0.54	0.51
Frac. URL	0.32	0.36	0.39	0.40

finding that polarization increases during spikes. The trends for bimotif, tie strength, and content divergence also persist, and are much stronger locally.

In addition to the previous measures, we also analyze other content features, such as the fraction of retweets, replies, mentions, and URLs around the spike. Interestingly, we find strong positive correlation of retweets, mentions, and URLs with volume, which indicates that discussion and endorsement increase during a spike. This finding is consistent with the ones by Smith et al. [29], who find that users tend to add URLs to their tweets when discussing controversial topics. Note that these additional content measures are only indicative for the local analysis, and do not produce consistent results at the global level.

Figure 14: Long-term trends of RWC (controversy) score in our dataset. No consistent trend can be observed, which contradicts the narrative that social media is making our society more divided.

5.5 Evolution over time

Let us now focus on how the measures change throughout time. The longitudinal span of the dataset of five years allows us to track the long-term evolution of discussion on controversial topics. A common point of view holds that social media is aggravating the polarization of society and exacerbating the divisions in it [5]. At the same time, the political debate (in U.S.) itself has become more polarized in recent years [4]. However, we do not find conclusive evidence for this argument with our analysis on this dataset.

Figure 14 shows the long-term trends of the RWC measure for the four topics. The trend is downwards for ‘abortion’ and ‘fracking’, while it is upwards for ‘obamacare’ and ‘gun control’. One could argue that the latter topics are more politically linked to the current administration in U.S., and for this reason have received increasing attention with the elections approaching. However, the only safe conclusion that can be drawn from this dataset is that there is no clear signal. The figure suggests that social media, and in particular Twitter, are better suited at capturing the ‘twitch’ response of the public to events and news. In addition, while our dataset spans a quite long time span for typical social media studies, it is still much shorter than other ones used typically in social science (coming from, e.g., census, polls, congress votes). This limit is intrinsic of the tool, given that social media have risen in popularity only relatively recently (e.g., Twitter is 10 years old).

5.6 Non-controversial topics

For comparison, we perform measurements over a set of non-controversial topics, defined by the hashtags *#ff*, standing for ‘Follow Friday’, used every Friday by users to recommend interesting accounts to follow; *#nba* and *#nfl*, used to discuss sports games; *#sxsw*, used to comment on the *South-by-South-West* conference; *#tbt*, standing for ‘Throwback Thursday’, used every Thursday by users to share memories (news, pictures, stories) from the past.

Figure 15: Non-controversial topics: RWC score as a function of the activity in the retweet network.

Figure 16: Non-controversial topics: Jensen-Shannon divergence of the lexicon between the two sides as a function of network activity. As the interest in the topic rises, the lexicon used by the two sides tends to converge.

We find that several structural measures, namely *clustering coefficient*, *tie strength*, and *bimotif*, behave similarly to the controversial topics, in that they obtain increased values for increased volume of activity. This result is in accordance with the ones by Romero et al. [28]. Conversely, the values of the RWC measure typically remain in ranges that indicate low presence of controversy, even as the volume of activity spikes (Figure 15). Additionally, with the definition of ‘core’ introduced above, we could only identify a negligibly small core for these topics (i.e., found very few users who were consistently active on these topics).

Finally, in terms of content measures we find that, as for the controversial topics, the entropy of the lexicon increases with volume (Figure 17). Topic variance also decreases with volume in most cases, meaning that a wider range of topics are discussed (Figure 18). On the contrary, the Jensen-Shannon divergence stays at relatively constant values across volume levels (Figure 16). It thus behaves differently compared to controversial topics (Figure 7). This result is to be expected, as the two ‘sides’ identified by METIS on the networks of non-controversial topics are not as well defined as they are in the case of controversial topics.

6 CONCLUSION

The evolution of networks is a well-studied phenomenon in social sciences, physics, and computer science. However, the evolution

Figure 17: Non-controversial topics: Entropy of the distribution over the lexicon for one side of the discussion as a function of the activity in the network (the other side shows similar patterns).

Figure 18: Non-controversial topics: Variance of the topic distribution. As the interest increases, variance decreases, indicating that a wider range of topics are being discussed.

of *interaction* networks has received substantially less attention so far. In particular, interaction networks related to discussions of controversial topics, which are important from a sociological point of view, have not been analyzed before. This study is a first step towards understanding this important social phenomenon.

We analyzed four highly controversial topics of discussion on Twitter for a period of five years. By examining the endorsement and communication networks of users involved in these discussions, we found that spikes in interest correspond to an increase in the controversy of the discussion. This result is supported by a wide array of network analysis measures, and is consistent across topics. We also found that interest spikes correspond to a convergence of the lexicon used by the opposite sides of a controversy, and a more uniform lexicon overall. Upon publication, the code and datasets used in the paper will be made available on the project website.⁹

Implications of this work relate to the understanding of how our society evolves via continuous debates, and how culture wars develop [1, 15, 22]. It is often argued that technology, and social media in particular, is having a negative impact on our ability to relate to the unfamiliar [5], due to the “echo chamber” and “filter bubble” effects. However, while we found instantaneous temporary increase in controversy in relation to external events, our study did

not find evidence of long term increase in polarization of the discussions, neither after these events nor as a general longitudinal trend. At the same time, investigating how to reduce the polarization of these discussions on controversial topics is a research-worthy problem [13], and taking into account the dynamics of the process is a promising direction to explore.

Our observations pave the way to the development of models of evolution for controversial interaction networks, similarly to how studies about measuring the Web and social media were the stepping stone to developing models for them. The logical next step for this line of work is to investigate how to use early signals from social media network structure and content to predict the impact of an event. Equally of interest is whether the observations made in this study translate to other social media beside Twitter, for instance, Facebook or Reddit. Finally, while we did not find any consistent long-term trend in the polarization of the discussions, it is worth continuing this line of investigation, as the effects of increased polarization might not be easily discoverable from social-media analysis alone.

7 ACKNOWLEDGEMENTS

This work has been supported by the Academy of Finland project “Nestor” (286211) and the EC H2020 RIA project “SoBigData” (654024).

8 REFERENCES

- [1] Alan Abramowitz and Kyle Saunders. 2005. Why can’t we all just get along? The reality of a polarized America. In *The Forum*, Vol. 3. bepress, 1–22.
- [2] Lada A Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *LinkKDD*. 36–43.
- [3] Yong-Yeol Ahn, Seungyeop Han, Haewoon Kwak, Sue Moon, and Hawoong Jeong. 2007. Analysis of topological characteristics of huge online social networking services. In *WWW*. ACM, 835–844.
- [4] Clio Andris, David Lee, Marcus J Hamilton, Mauro Martino, Christian E Gunning, and John Armistead Selden. 2015. The rise of partisanship and super-cooperators in the US House of Representatives. *PLoS one* 10, 4 (2015), e0123507.
- [5] Yochai Benkler. 2006. *The wealth of networks: How social production transforms markets and freedom*. Yale University Press.
- [6] Alessandro Bessi, Guido Caldarelli, Michela Del Vicario, Antonio Scala, and Walter Quattrociocchi. 2014. Social determinants of content selection in the age of (mis) information. In *SocInfo*. Springer, 259–268.
- [7] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political Polarization on Twitter. In *ICWSM*.
- [8] Gianmarco De Francisci Morales, Aristides Gionis, and Claudio Lucchese. 2012. From Chatter to Headlines: Harnessing the Real-Time Web for Personalized News Recommendation. In *WSDM*. 153–162.
- [9] Nicholas DiFonzo, Jerry Suls, Jason W Beckstead, Martin J Bourgeois, Christopher M Homan, Samuel Brougher, Andrew J Younge, and Nicholas Terpstra-Schwab. 2014. Network structure moderates intergroup differentiation of stereotyped rumors. *Social Cognition* 32, 5 (2014), 409.
- [10] Sergei N Dorogovtsev and José FF Mendes. 2013. *Evolution of networks: From biological nets to the Internet and WWW*. OUP Oxford.
- [11] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2016. Quantifying Controversy in Social Media. In *WSDM*. ACM, 33–42.
- [12] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. The Ebb and Flow of Controversial Debates on Social Media. In *ICWSM*.
- [13] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing Controversy by Connecting Opposing Views. In *WSDM*. ACM, 81–90.
- [14] Kiran Garimella and Ingmar Weber. 2017. A Long-Term Analysis of Polarization on Twitter. In *ICWSM*.
- [15] Benjamin Highton and Cindy D Kam. 2011. The long-term dynamics of partisanship and issue orientations. *The Journal of Politics* 73, 01 (2011), 202–215.
- [16] George Karypis and Vipin Kumar. 1995. METIS - Unstructured Graph Partitioning and Sparse Matrix Ordering System. (1995).
- [17] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111, 24 (2014), 8788–8790.
- [18] Janette Lehmann, Bruno Gonçalves, José J. Ramasco, and Ciro Cattuto. 2012. Dynamical Classes of Collective Attention in Twitter. In *WWW*. ACM, 251–260.
- [19] Jure Leskovec, Lars Backstrom, Ravi Kumar, and Andrew Tomkins. 2008. Microscopic evolution of social networks. In *KDD*. ACM, 462–470.
- [20] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2005. Graphs over time: densification laws, shrinking diameters and possible explanations. In *KDD*. ACM, 177–187.
- [21] Jianhua Lin. 1991. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory* 37, 1 (1991), 145–151.
- [22] Kara Lindaman and Donald P Haider-Markel. 2002. Issue evolution, political parties, and the culture wars. *Political Research Quarterly* 55, 1 (2002), 91–110.
- [23] Zhe Liu and Ingmar Weber. 2014. Is Twitter a public sphere for online conflicts? A cross-ideological and cross-hierarchical look. In *SocInfo*. Springer, 336–347.
- [24] Haokai Lu, James Caverlee, and Wei Niu. 2015. BiasWatch: A Lightweight System for Discovering and Tracking Topic-Sensitive Opinion Bias in Social Media. In *CIKM*. ACM, 213–222.
- [25] Yelena Mejova, Amy X Zhang, Nicholas Diakopoulos, and Carlos Castillo. 2014. Controversy and Sentiment in Online News. *Symposium on Computation + Journalism* (2014).
- [26] AJ Morales, J Borondo, JC Losada, and RM Benito. 2015. Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos* 25, 3 (2015).
- [27] Mark Newman, Albert-Laszlo Barabasi, and Duncan J Watts. 2011. *The structure and dynamics of networks*. Princeton University Press.
- [28] Daniel M Romero, Brian Uzzi, and Jon Kleinberg. 2016. Social Networks Under Stress. In *WWW*. 9–20.
- [29] Laura M Smith, Linhong Zhu, Kristina Lerman, and Zornitsa Kozareva. 2013. The role of social media in the discussion of controversial topics. In *SocialCom*. IEEE, 236–243.
- [30] Fang Wu and Bernardo A. Huberman. 2007. Novelty and collective attention. *Proceedings of the National Academy of Sciences* 104, 45 (2007), 17599–17601.