

A Measure of Polarization on Social Media Networks Based on Community Boundaries

Pedro H. Calais Guerra, Wagner Meira Jr.

Dept. of Computer Science
Universidade Federal de Minas Gerais (UFMG)
Belo Horizonte, MG, Brazil

Claire Cardie, Robert Kleinberg

Dept. of Computer Science
Cornell University
Ithaca, NY 14853 USA

Abstract

Polarization in social media networks is a fact in several scenarios such as political debates and other contexts such as same-sex marriage, abortion and gun control. Understanding and quantifying polarization is a long-term challenge to researchers from several areas, also being a key information for tasks such as opinion analysis. In this paper, we perform a systematic comparison between social networks that arise from both polarized and non-polarized contexts. This comparison shows that the traditional polarization metric – modularity – is not a direct measure of antagonism between groups, since non-polarized networks may be also divided into fairly modular communities. To bridge this conceptual gap, we propose a novel polarization metric based on the analysis of the boundary of a pair of (potentially polarized) communities, which better captures the notions of antagonism and polarization. We then characterize polarized and non-polarized social networks according to the concentration of high-degree nodes in the boundary of communities, and found that polarized networks tend to exhibit low concentration of popular nodes along the boundary. To demonstrate the usefulness of our polarization measures, we analyze opinions expressed on Twitter on the gun control issue in the United States, and conclude that our novel metrics help making sense of opinions expressed on online media.

1 Introduction

In Social Sciences, polarization is the social process whereby a social or political group is divided into two opposing sub-groups having conflicting and contrasting positions, goals and viewpoints, with few individuals remaining neutral or holding an intermediate position (Sunstein 2002; Isenberg 1986). A typical domain where polarization is witnessed is Politics (Waugh et al. 2009; Dixit and Weibull 2007), although a range of other issues are known to induce in the society a divisive debate that often makes a fraction of people to have very extreme opinions, such as global warming (McCright and Dunlap 2011), gun control, same-sex marriage and abortion (Mouw and Sobel 2001; Hunter 1992).

As expected, the rise of social media systems quickly turned the web into a platform of lively discussions and debates (Weinberger 2011); now online battles are fought over polarizing, polemic issues, specially when new evidences that support one side of the discussion arise, such as the intensity increase associated with the gun control debate after the shootings in Newtown, Connecticut. Social and computer scientists are paying increasing attention to such discussions, seeking for patterns that unveil the dynamics of online debate and the bursts of opinionated content generated in reaction to real-life events. Thus, identifying whether a topic induces polarization on individuals is important for at least three reasons:

1. It is a relevant question from the sociological point of view, since polarization causes segregation and political conflict in the society, as a consequence of the increase of extreme opinions over time and the high degree of bias of polarized opinions (Paul DiMaggio 1996; Mouw and Sobel 2001).
2. Polarization may be a key information for tasks such as opinion analysis. A biased opinion holder is likely to keep the same, extreme opinion over time, and the knowledge of the side in a discussion an opinion holder is (in favor or against an issue) can help predict the polarity of his/her opinions (Calais et al. 2011; Tan et al. 2011).
3. In polarized debate, the strong bias of opinions suggests that they should not be taken into consideration without considering *who* is issuing the opinion (Walton 1991). In other words, two equivalent opinions may have different interpretations and impact if issued by people from opposite sides, and new opinion mining tasks may apply, such as monitoring people who changed their previous viewpoint over a topic.

Given the relevance of the contexts in which polarization is witnessed, many works from CS (and more specifically, Social Network Analysis) have investigated online social networks induced by polarized debate, specially in the political domain (Adamic and Glance 2005; Livne et al. 2011). In general, the Computer Science literature assumes (either implicitly or explicitly) that a social network is *polarized* if nodes can be partitioned into two highly cohesive sub-groups, reflecting, possibly, two contrasting viewpoints. In particular, the well-known community quality metric known

as *modularity* (Newman 2006) is commonly used to measure the level of segregation of two particular groups; such as democrats vs republicans, people in favor or against abortion etc. A network segmentation with high modularity indicates that the social graph may be divided into clusters having many internal connections among nodes and few connections to the other group, what is widely accepted as an indication of polarization (Conover et al. 2011).

Our main claim in this paper is that, although we acknowledge that modularity is correlated to the social phenomenon of polarization, and highly modular networks are certainly linked with an increased likelihood of polarization of positions expressed by users who are part of the network, modularity is not a *direct* measure of polarization. We draw this observation from the fact that it is not clear how much modularity is “enough” to state that a social network is polarized. For instance, people may be divided into those that like basketball and those that like football, even though there is no notion of opposition among the two groups – they are just two different preferences which are not mutually exclusive, since some individuals can be practitioners of both sports and thus belong to both communities. Although the existence of two segregated social groups is certainly a *necessary* condition for polarization, the modularity measured for any network divided into two cohesive communities will be a value different from zero, even if no polarization *at all* is present among nodes.

Our goals in this paper are twofold. First, we perform a systematic comparison of social media networks emerging from both polarized and non-polarized contexts, by collecting a diversity of social networks from social media systems such as Twitter, Facebook and blogs. Our goal is to avoid the bias of current works, which focus on networks from domains that are previously known to be polarized, specially Politics. We then identify communities in these networks and verify that their modularity is not sufficiently clear to state that polarization is an ongoing phenomenon or not; although polarized social networks tend to be more modular than non-polarized networks, the determination of a threshold of polarization is a challenging task that depends on factors such as social media platform and nature of interactions.

Motivated by this observation, we focus on the following question: since polarization is a strong, remarkable sociological phenomenon, are there structural patterns which better capture the differences between polarized and non-polarized networks, rather than the level of modularity between communities? We propose an analysis of the boundary between the two potentially polarizing communities – the portion of the social graph comprising nodes from one community which link to one or more nodes of the other community. Our hypothesis is that, in such community boundaries, one group unveils what they “think” about the other group, and thus it is the place where we seek for evidences of antagonism. Our metric considers a null model of polarization that assumes that, on a non-polarized network, cross-group interactions should be at least as frequent as interactions with internal nodes on the community. The model considers nodes’ likelihood into connecting to users which belong to the other (potentially opposing) group, in comparison to the likelihood of

connecting to members from its own group.

We also empirically demonstrate that polarized and non-polarized social networks tend to differ according to another structural property: the concentration of popular (high-degree) nodes outside community boundaries. On non-polarized contexts, we observed a concentration of popular nodes along the boundary, since the sharing of similarities between members of the boundary increase the popularity of such nodes (e.g., users that like both football *and* basketball). On the other hand, we found that polarized networks tend to have a lower concentration of popular nodes in the boundary, since the antagonism between both sides decrease the likelihood of existence of nodes that are popular in both groups.

To show the applicability of our findings on the interpretation of opinions expressed on social media, we employ our metrics to perform an analysis of opinions expressed on Twitter on the gun control issue in the United States. We demonstrate that our metrics based on community boundaries are a useful complement to the traditional modularity measure in helping to understand how the structure of a social network links with the viewpoints and opinions expressed in online social environments.

Our work is organized as follows. In §2 we discuss related work. In §3 we evaluate the modularity of a range of polarized and non-polarized networks. We then propose a new metric to measure polarization based on community boundaries, in §4. In §5, we employ our metric to understand opinions expressed on Twitter on the gun control issue on America. Next, we compare polarized and non-polarized networks in terms of another structural property – concentration of popular nodes in the boundary – in §6. Finally, we present the conclusions in §7.

2 Related Work

From the sociological perspective, polarization can be formally understood as a state that “refers to the extent to which opinions on an issue are opposed in relation to some theoretical maximum”, and, as a process, it is the increase in such opposition over time (Mouw and Sobel 2001; Paul DiMaggio 1996). A typical sign that polarization is playing a role in a society with regard to an issue is when opinions become more extreme over time even after opposing sides examine the *same* evidence (Yardi and Body 2010), as demonstrated on a classical experiment in the 70’s, where people against and in favor of death penalty have become more convinced of their conflicting positions after reading the *same* essay on the topic (Lord, Ross, and Lepper 1979). In that direction, a common approach to reason about such setting is to model the polarization phenomenon with bayesian probabilities (Dixit and Weibull 2007; Bengio et al. 2009): the previous belief each group or individual has on a topic is the *prior*, and, if the updated beliefs of the opposing groups become more divergent after both examine the same evidence, then it is likely that polarization is happening.

Sociologists usually resort to polls and elections data to assess the presence of extreme opinions on the public opinion (Abramowitz and Saunders 2005). When information

on the relationships among people is available, polarization is commonly accepted as an ongoing phenomenon if people can be divided into highly cohesive communities; each community represents a distinct position or preference: liberal versus conservative parties, pro-gun and anti-gun voices, for instance. The segregation of people into groups seems to be a remarkable characteristic of social networks induced by polarized debate as an immediate consequence of the homophily principle, which states that people with similar beliefs and opinions tend to establish social ties with higher probability (McPherson, Smith-Lovin, and Cook 2001; Yardi and Body 2010). As a general concept, a community is cohesive if the internal connectivity of the group is high and if the connectivity of members of the group with members from outside the community is low. Group cohesion is usually measured with community quality metrics such as *conductance* (Kannan, Vempala, and Vetta 2004) and *modularity* (Newman 2006). For instance, (Zhang et al. 2008) and (Waugh et al. 2009) argue that modularity may be used to study partisan polarization in U.S. Congress. On the online world, modularity has been used as evidence of segregation between political groups in a diversity of online media such as blogs (Adamic and Glance 2005) and Twitter (Conover et al. 2011; Livne et al. 2011).

However, these works analyze contexts and domains which are previously known and expected to induce polarization – in particular, Politics. As a consequence, they do not examine networks from non-polarized domains, and it remains an open question which are the necessary and sufficient conditions of polarization between groups of individuals, in terms of the structure of the induced social network. In order to precisely understand how polarization affects the social network structure, we need to inspect *both* polarized and non-polarized domains, thus avoiding the “sampling bias” of examining only highly-polarized networks. Previous research provides strong evidence that the existence of highly modular groups is a *necessary* condition in order to observe polarization; our work contributes to the better understanding of how polarization affects the structure of social networks by performing a systematic comparison between both polarized and non-polarized networks. Instead of using modularity, we propose two new measures of polarization based on the analysis of the structure of the boundary of two potentially polarizing groups, which we will detail and evaluate on the next sections.

3 Comparing Modularity of Polarized and Non-Polarized Social Networks

We consider existing and publicly available social networks and additional networks we collected from Facebook and Twitter in order to compare the structural characteristic of social networks with varying degrees of polarization, including total absence of polarization. We work with the following social networks:

1. **University Friendships’ Network:** This social network comprises the social relationships established on Facebook by professors, undergraduate and graduate students of a large department at a Brazilian University.
2. **Brazilian Soccer Supporters:** We collected mentions, on Twitter, to two of the most popular soccer teams in Brazil – Cruzeiro and Atletico Mineiro, known by being the fiercest rivals in the country. Nodes are Twitter users, and a direct edge connects users involved in any *retweet*. A *retweet* usually means an endorsement (Calais et al. 2011), and thus it is a good evidence of sharing of similar viewpoints between two individuals.
3. **New York City Sports Teams:** We collected mentions, on Twitter, to two sports teams hosted in New York City: New York Giants (football) and New York Knicks (basketball). The network is induced by *retweets*; we restrict the network to nodes that mentioned both teams at least once, to guarantee that we are taking into account only users which are interested in both teams. Note that, differently from the previous network, we do not expect polarization here, since the two potential communities represent supporters of teams from different sports.
4. **Karate’s Club:** This is a social network of friendship ties established between 34 members of a karate club at a U.S. university in the 1970s, and the emergence of two communities was a result of a disagreement developed between the administrator of the club and the club’s instructor, which ultimately resulted in the instructor’s leaving and starting a new club, taking about a half of the original club’s members with him (Zachary 1977; Easley and Kleinberg 2010).
5. **2004 U.S. Political Blogosphere:** This dataset was among the first that showed that political blogs on the U.S. are divided into two dense communities – representing liberals and conservatives (Adamic and Glance 2005). Directed edges are links between two blogs.
6. **Gun Control:** We collected tweets mentioning gun control issues since the shootings on Sandy Hook Elementary School in Newtown, Connecticut, on December 14, 2012. We considered the following keywords to collect data: gun control, guns, mass shootings and NRA¹. As in other networks obtained from Twitter, users are linked through *retweets*.

Note that all the aforementioned networks have a semantic unicity, in the sense that users interacting and expressing opinions are restricted to a single domain or topic. In Table 1 we provide a summary of the main characteristics of these networks, including number of nodes and edges. For each network, we split nodes into communities, in order to assess the structural patterns that arise from the segmentation of the graph into groups. In the case of the networks University, Brazilian-Soccer, Political-Blogs and Gun-Control, we have run the community detection algorithm from (Blondel et al. 2008), a simple modularity maximization approach provided by the *Gephi*² software package. In the case of the network NYC-Teams, we separated users into the community of NY Giants or NY Knicks according to the number of hash-tags each user posted referring to each team. For network

¹NRA is the National Rifle Association.

²*Gephi* is available at <http://www.gephi.org>.

Karate-Club, we employed the ground-truth separation provided by (Easley and Kleinberg 2010). The gun control debate graph was divided into three large communities, and we leave its analysis to §5, after we introduce our novel polarization metric.


For each pair of communities, we calculate modularity Q . The modularity of a network quantifies the extent, relative to a random network, to which vertices cluster into community groups, and the higher its value, more modular the network is (Newman 2006). Modularity is traditionally formulated as Equation 1; m is the number of edges, A is the adjacency matrix, k_i and k_j are node degrees and $s_i s_j = 1$ if nodes i and j belong to the same community and -1 otherwise. Values of Q obtained for the datasets we consider in this paper are shown in Table 1.

$$Q = \frac{1}{2m} \sum_{ij} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \frac{s_i s_j + 1}{2} \quad (1)$$

We first observe that networks induced by domains for which we expect polarization (networks 3, 4 and 5) exhibit a high measure of modularity when compared to networks 1 and 2. This observation is in accordance with previous works that associate high modularity to polarization (Conover et al. 2011). However, we point out three drawbacks on mapping modularity to the sociological behavior of polarization:

1. On communities that arise from contexts where we do not expect polarization, the modularity value is still a positive, moderate value, as in the case of University and NYC-Teams networks. Modularity for the University network is 0.24 (shown in Figure 1(a)), what suggests a network less polarized than Political-Blogs, which exhibits a modularity value of 0.42 (Figure 1(b)). However, from the sociological standpoint, we do not expect to observe *any* antagonism at all between undergraduate and graduate students.
2. The direct mapping of modularity values into degrees of polarization shows some inconsistencies when we compare modularity measures obtained by independent researches working with different data. (Zhang et al. 2008), for instance, have found modularity values not higher than 0.18 from the examination of networks induced by voting agreement on the U.S. Congress. Although the authors' goal is to evaluate the increase of modularity over time to conclude that polarization was rising among politicians over the decades, the maximum modularity measure they found is just 0.01 higher than the value that (Conover et al. 2011) found to conclude that $Q = 0.17$ is not associated with an evident community structure on a communication network in Twitter. In previous researches, modularity is used more to *confirm* an early suspicion of polarization, rather than find whether polarization exists or not in an unknown domain.
3. Modularity has a known resolution limit problem caused by the fact that its null model assumes that each node may connect to any other node of the network, what is not realistic for large graphs (Good, de Montjoye, and Clauset 2010). Therefore, comparing the modularity value across

different networks is not a good practice if the graphs' size are very different (Fortunato and Barthélemy 2007), which is the case of the graphs compared in Table 1.



(a) Network of Facebook Friends (b) 2004 U.S. Political Blogosphere from an University Department

Figure 1: Two social graphs showing a non-polarized network (Facebook Friends) and a polarized network (Politics). Although the political network is more modular, it is not clear what is the minimum level of modularity associated with the sociological phenomenon of polarization.

The conceptual gap between the modularity measure and the sociological behavior of polarization, evidenced on these extreme cases, limits the understanding of networks and contexts where it is less clear whether polarization is taking place. In the next section, we will provide details about a novel structural pattern we propose, in order to better capture the presence and absence of polarization in communities formed around a given domain or topic of discussion.

4 Measuring Polarization on Community Boundaries

It is known that a significant portion of the structure of a social network is affected by the context and the behavior of the nodes (Easley and Kleinberg 2010). Behavioral patterns such as homophily (McPherson, Smith-Lovin, and Cook 2001), social influence (Friedkin 1998) and social balance (Heider 1958) directly affect the likelihood that specific pairs of users will establish a tie in a social environment. Since polarization is a strong, remarkable sociological phenomenon, we expect that a social network embedded in such a context of opposing and conflicting relationships will induce structural patterns which are not observed on general, non-polarized networks.

The link between high modularity and polarization carries the implicit assumption that the absence of positive interactions between nodes (e.g., *retweets* and friendship ties) is a sign of *antagonism*, i.e., a segmentation of social groups due to opposition and clash of viewpoints. Modularity compares the internal and external connectivity of two groups G_i and G_j ; it quantifies both homophily (nodes from a community establishing ties due similarity) and antagonism (nodes avoiding establishing ties with the alternate community) through the same equation, and limits the understanding of

Table 1: General Description of Social Networks and derived communities.

network	media	# nodes	edge type	# edges	communities	modularity Q
1 - NYC Teams	Twitter	19,585	directed	201,691	NY Giants fans and NY Knicks fans	0.15
2 - University	Facebook	133	undirected	2,241	graduate and undergraduate students	0.24
3 - Karate's Club	friendships	34	undirected	78	followers of instructors 1 and 2	0.35
4 - Brazilian Soccer Teams	Twitter	27,415	directed	156,489	Cruzeiro and Atletico fans	0.39
5 - US Political Blogs	blogs	1,224	directed	16,715	liberals and conservatives	0.42
6 - Gun Control	Twitter	61,740	directed	342,449	analyzed in §5	–

antagonism in isolation. To better understand polarization, we propose to seek for social structures that highlight the presence (or absence) of antagonism, since homophily is a pattern present both on non-polarized and polarized networks, but antagonism is expected only on the latter.

With this idea in mind, we focus our analysis on nodes that effectively interact with the (potentially) opposing group. These nodes are part of a *community boundary*, which we define, for a group/community G_i , as the subset of nodes $B_{i,j}$ that satisfies two conditions:

1. A node $v \in G_i$ has at least one edge connecting to community G_j ;
2. A node $v \in G_i$ has at least one edge connecting to a member of G_i which is not connected to G_j .

Equation 2 formally defines boundary $B_{i,j}$. In Figure 2 we show a toy example of a network divided into communities G_1 (dark) and G_2 (white). According to our definition, $B_{1,2} = \{b, d\}$ and $B_{2,1} = \{1, 2\}$. Note that node e does not belong to $B_{1,2}$ because it does not meet condition 2.

$$B_{i,j} = \{v_i : v_i \in G_i, \exists e_{ik} | v_k \in G_j, \exists e_{ik} | (v_k \in G_i, \nexists e_{kl} | v_l \in G_j), i \neq j\} \quad (2)$$

Nodes from G_i which do not belong to $B_{i,j}$ are named *internal nodes* and are grouped in the set I_i , defined by Equation 3. In Figure 2, $I_1 = \{a, c\}$ and $I_2 = \{3, 4\}$.

$$I_i = G_i - B_{i,j} \quad (3)$$

We perform our analysis of polarization by analyzing the connectivity between I_i , $B_{j,i}$, I_j and $B_{i,j}$. These four sets allow us to compare nodes' choices in connecting to nodes from a very different nature. Due to condition 1, we can assess the connections of $B_{i,j}$ with $B_{j,i}$, i.e., with nodes that belong to a potentially opposing group. Due to condition 2, nodes from $B_{i,j}$ also establish contact with a set of nodes which do not connect to any member of the potentially opposing group. Nodes from I_i avoid any connection to the alternate group and restrict their connections to nodes from their own community, representing individuals that, theoretically, are very different from nodes from the other group, in the case of a polarizing domain.

We focus on I_i , I_j , $B_{i,j}$ and $B_{j,i}$ as groups that better represent the (potential) distinct nature of (potentially) polarized individuals, in comparison to the division between G_i and G_j that is analyzed by modularity. Our proposal is to compare the degree of preference of each node in $B_{1,2}$




Figure 2: Toy example of a graph divided into two communities G_1 and G_2 . Sets I_1 , $B_{1,2}$, $B_{2,1}$ and I_2 are defined according to Equations 2 and 3.

to connect to members from I_1 or $B_{2,1}$, and of each node in $B_{2,1}$ to connect to members from I_2 or $B_{1,2}$. To perform such comparison, we define two sets of edges. The first set is E_B , which is the set of edges that connect members from G_i to members from G_j :

$$E_B = \{e_{mn} : v_m \in B_{i,j} \wedge v_n \in B_{j,i}\} \quad (4)$$

In Figure 2, $E_B = \{(b, 1), (d, 2)\}$. These edges are evidence of interaction between the two distinct groups. To contrast with these interactions, we also define E_{int} as the set of edges that connect boundary nodes to internal nodes:

$$E_{int} = \{e_{mn} : v_m \in (B_{1,2} \cup B_{2,1}) \wedge v_n \in (I_1 \cup I_2)\} \quad (5)$$

In the example, $E_{int} = \{(a, b), (c, d), (1, 3), (2, 4)\}$. The modularity for this community configuration is $Q = 0.30$, what indicates a reasonable level of segregation among the two communities. However, let us examine the decisions taken by each node at the boundary in establishing their connections. Consider node b , which has a node degree $d(b) = 3$:

1. $(b, 1)$ is a cross-group edge and belongs to E_B ;
2. (b, a) is an internal edge and belongs to E_{int} ;
3. (b, d) is neither an internal edge, nor a cross-group edge.

We consider that b did not exhibit any type of antagonism to members of the other group; since it established the

same number of connections to $B_{2,1}$ and I_1 . Note that the same reasoning is applicable to the remaining members of the boundary, d , 1 and 2. The network from Figure 2, according to our principle, does not exhibit polarization. Note that edges (b, d) , (a, c) , $(1, 2)$, $(3, 4)$ and $(e, 2)$ are intentionally not included in this evaluation, since they capture more homophily between nodes than antagonism between groups.

Equation 6 generalizes the comparison among the connectivity choices that nodes in $B_{i,j}$ make while connecting to members from I_i or $B_{j,i}$. For each node v belonging to the boundary B , we compute the ratio between the number of edges it has in E_{int} (which we call $d_i(v)$) and the total number of edges in E_{int} and edges in E_B (which we call $d_b(v)$). We compare such ratio with the following null hypothesis: each node spreads its edges equally between internal nodes and nodes from the other community. P lies in the range $(-1/2, +1/2)$; a P value below 0 indicates not only lack of polarization, but also that nodes in the boundary are more likely to connect to the other side. Conversely, a P value greater than zero indicates that, on average, nodes on the boundary tend to connect to internal nodes rather than to nodes from the other group, indicating that antagonism is likely to be present. In the case of the communities shown on Figure 2, $P = 0$, since all boundary nodes established the same number of connections to internal nodes and to nodes from the alternate community.

$$P = \frac{1}{|B|} \sum_{v \in B} \left[\frac{d_i(v)}{d_b(v) + d_i(v)} - 0.5 \right] \quad (6)$$

Absence of Boundary. While traditional community quality measures such as modularity are relatively high for a network comprised of two isolated communities, our polarization metric cannot be computed when $B = \emptyset$. While this case can be interpreted as a network of very high polarization, we consider that it is more reasonable to state that it is not possible to assess polarization between two isolated communities, since it can be the case that each group does not know each other at all. The intuition here is that the hypothesis is not verifiable, since the groups do not have any interaction and we cannot guarantee that there is any polarization. It corresponds to asking if there is polarization between human beings and extraterrestrials.

In Table 2 we compare values of modularity Q and polarization P for the set of datasets we consider in this work; networks are sorted according to their modularity values. For the network comprising supporters of New York City football and basketball teams (NY Giants and NY Knicks), our metric P detects absence of polarization ($P = -0.002$), suggesting that although fans are divided into two groups, they do not oppose each other. This is different from network 4, which comprises fans of two rival soccer teams from Brazil; in this case our metric indicates that there is, indeed, polarization among such fans ($P = 0.20$). The University network exhibits a negative value $P = -0.24$. This result is consistent with recent work that examine the overlap between communities in social networks and concluded that the overlap tend to be denser, in terms of number of edges, than the group themselves (Yang and

Leskovec 2012). The boundary connects users that share common interests and background, such as supporting both NY Knicks and NY Giants or having attended high school and college together. In the case of polarized communities, such pluralistic homophily is not present.

In order to highlight the differences in the structure of large polarized and non-polarized online social networks, we compare in Figure 3 the node-specific values of $\frac{d_i(v)}{d_b(v) + d_i(v)} - 0.5$, which we call P_v , for each node v on the boundary of each network. The number of nodes with $P_v < 0.5$ is very limited on the polarized network of Brazilian soccer rivals, indicating their likelihood to connect to internal nodes rather than endorsing (retweeting) adversaries. Note, also, that the slope of the curve formed by points with $P_v < 0$ on the polarized network is more inclined, reflecting that nodes face resistance to connect to the boundary. We interpret such difference w.r.t. slope as a genuine manifestation of antagonism. In the curve of the non-polarized network, however, the slope before and after $P = 0$ is roughly the same, indicating that nodes presents the same likelihood to establish connections, what we interpret as a sign of absence of polarization.

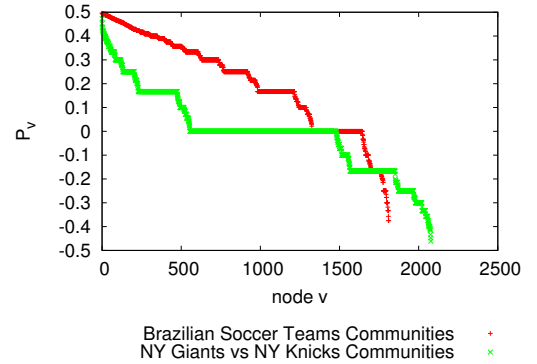


Figure 3: Distribution of P_v for Twitter communities debating Sports. A polarized social network is characterized by a small number of nodes preferring cross-boundary connections ($P_v < 0$).

5 Opinion Analysis on the Gun Control Debate

In this section we use the polarization metric P we introduced in the last section to analyze opinions expressed on the gun control issue in Twitter. The debate around gun control laws has long history in the United States and is often present in political debates (Blendon, Young, and Hemenway 1996). Events related to the issue, such as the shootings in the Sandy Hook Elementary School in Connecticut, on December 14, 2012, unleash bursts of strong opinions on the topic. From that date until February 10, 2013 we collected 3,816,137 tweets mentioning gun control-related keywords. Since gun control is a typically polarizing topic, we attempt to use the network structure to interpret, predict and analyze opinions expressed regarding the issue.

Table 2: Modularity Q and Polarization P for networks described in Table 1.

network	media	modularity Q	polarization P
1 - NYC Teams	Twitter	0.15	-0.002
2 - University	Facebook	0.24	-0.24
3 - Karate's Club	friendships	0.35	0.17
4 - Brazilian Soccer Teams	Twitter	0.39	0.20
5 - US Political Blogs	blogs	0.42	0.18



Figure 4: Communities obtained from gun control debate on Twitter. Nodes are users and edges represent *retweets*. From the left to the right, we refer to them as communities 1 (green), 2 (blue) and 3 (yellow).

When plotting the social network induced by *retweets* on *Gephi* and executing the modularity maximization algorithm from (Blondel et al. 2008), we got the three communities shown in Figure 4. We start by computing modularity Q for each of the three pairs of communities: the modularity between the leftmost group (colored in green) and the rightmost group (in yellow) is $Q = 0.47$; while modularity for communities 1 and 2 is $Q = 0.31$. Finally, the modularity between groups 2 and 3 is $Q = 0.26$. Although we expect the most distant groups to have conflicting opinions and viewpoints, the lack of a more precise measurement of how polarization limits the understanding of the opinion sharing patterns among nodes. Does group 2 has a different, third opinion in comparison to group 3, or do they share a common viewpoint, and the division into two communities is caused by other factors? This answer is not provided by the analysis of the modularity metric by itself, because we do not know in advance whether $Q = 0.26$ is high enough to state that there is antagonism between community members, or if such threshold exists and is dependent of the social media platform or the nature of the interactions.

To gain insights on the relationships between the groups, we calculate the metric P we proposed for each pair of communities. Results are shown in Table 3. By analyzing the Q values, it is not immediately obvious what are the sharing and conflicting opinions between groups. However, our polarization metric P provides better clues on the opinion sharing patterns. Community 1 is predicted to be polarized with communities 2 and 3, with $P = +0.23$ and $P = +0.32$, re-

spectively. On the other hand, our metric predicts that communities 2 and 3 have no polarization at all ($P = -0.14$). On the contrary, a P value significantly below zero means that nodes in the boundary tend to establish more cross-group connections than expected. By manual verification of the profiles of users belonging to each group, we concluded that group 1 is dominated by conservative voices, while liberals are concentrated on group 3. Group 2 is dominated by independent opinion holders.

Table 3: Modularity Q and Polarization P for Gun Control debate.

communities	modularity Q	polarization P
GC-1 and GC-2	0.31	+0.23
GC-1 and GC-3	0.47	+0.32
GC-2 and GC-3	0.26	-0.14

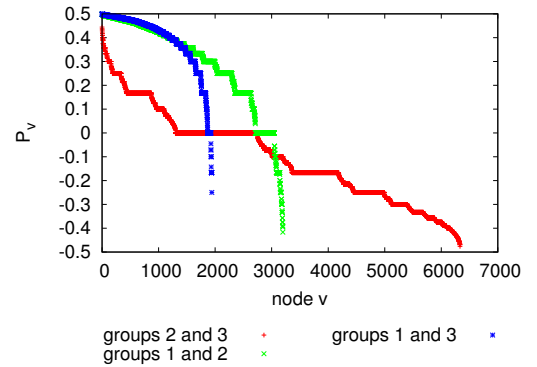


Figure 5: Distribution of P_v for communities debating Gun Control. On the pairs of polarized communities (1–2 and 1–3), few nodes establish more connections with the alternate group than with internal nodes ($P_v < 0$).

In Figure 5, we plot the distribution of P_v for the boundary nodes for each pair of communities. We can note a clear difference in the shape of the curve corresponding to the pair of communities 2–3 in comparison to 1–2 and 1–3: in addition to a significant number of nodes with $P_v < 0$ on the non-polarized network, the smooth transition from nodes that are more likely to connect to internal nodes from nodes that are more likely to connect to boundary nodes contrasts with the quickly decrease in the polarized curves, indicating that the boundary reduces the likelihood of connections, acting as a barrier. Moreover, the difference in modularity Q for pairs

Table 4: Popular Tweets on Gun Control debate since Newtown Shootings, on December 14, 2012, for each community shown in Figure 4. Confirmation bias on polarized debate makes people focus on facts that confirm their previous opinion.

group	Twitter user	tweet	#RTs
1	@tillerylakelady	2 those of you whining about #gun control-a madman used a KNIFE to stab 20+ kids in China today. Its not about guns,its about mental health.	71
1	@JohnGaltTx	Since the Australia Gun Ban, the following increased: Armed Robberies +69%, Assaults with Guns +28%, Gun Murders +19%, Home Invasions +21%	12
1	@Gere341	TAKE NOTE LIBS- Canada is a gun controlled country. Yet there was a deadly Mall Shooting last June Someone wants 2 get access2guns, theyWill	16
1	@RightCentrist	466 violent crimes per 100K ppl in the US, 2034 violent crimes per 100K ppl in UK - Statistic says it all - gun control fails.	14
2	@alexblagg	22 children in China attacked with a knife today, no deaths. Senseless violence can't be prevented. Gun violence can.	182
2	@jasonwstein	18 school shootings in Australia before 1981. They banned semi-automatic weapons. No big school shootings since. via @cnbc	19
2	@igorvolsky	9,000 people killed with guns last year. In similar countries like Germany, 170. Canada , 150. There is a reason for that.	12
2	@Tinkerbell_	51 people were killed by guns in the UK in 2011. 51. In the ENTIRE year. USA 8,583. Now say gun control does not work.	10
3	@EstherKramer1	22 hurt by knife attack in China vs. 20 kids dead via gun in USA. Crazy people kill but guns help the m out a lot. #GunControlNow	14
3	@rationalists	1996: Gunman kills 35 in Port Arthur, Australia . 1997: Australia bans guns. 2012: Massacres since 1997: NONE!	57
3	@sean_dixon	In Canada we watch the same films as in the US. We play the same games, have the same mental health issues. So whats the difference? #NRA	11
3	@Good.Bear	5 times as many murders per head of population in USA as in UK . Do you think thats because Americans are 5 times as evil or have more guns?	53

2–3 and 1–2 is just 0.05, however their structure is fundamentally different, as Figure 5 shows. We believe that the different distributions we found may support the building of graph generation models that better represent polarized and non-polarized social networks.

In Table 4, we present some of the most popular tweets posted by members from each community comparing statistics, facts and gun regulations from three other countries – China, Australia, and Canada, in addition to the United Kingdom. Tweets from group 1 show a clear pro-gun rationale, and an anti-gun position for groups 2 and 3 is also evident. The content posted by users, therefore, is in accordance to the measurement of our polarization metric. More interestingly, each group attempted to use statistics from each country in their favor; the same country is used as a case that favors gun rights (group 1) and as a case that favors gun control (groups 2 and 3). The focus on evidences that reinforce previous opinions is a cognitive bias known as confirmation bias (Nickerson 1998). In the case of China, Twitter users comment on the *same* fact – the attack of children with a knife-armed man – and yet they use the fact to reinforce contrasting opinions. Such phenomenon, known in the social psychology literature as *belief polarization* (Lord, Ross, and Lepper 1979), is one of the strongest evidences that a group of individuals is divided into polarized groups. Note that our understanding of the relationship between the three groups provided by our polarization metric P allowed us to quickly find such contradicting opinions, and our methodology may support sociological studies on polarization of opinions based on social media data.

6 Concentration of Popular Nodes along the Boundary

In this section we investigate another structural characteristic that may help on the identification of polarization – the concentration of popular (high-degree) nodes in the boundary. Since polarization is associated with antagonism, we expect popular nodes to be present far from the boundary, as strong representatives of their group viewpoints that do not find endorsement from the opposing side. On the other hand, we expect non-polarized communities to promote the existence of high-degree nodes in the boundaries, since such nodes are more prone to enjoy popularity from both sides.

To measure the concentration of high-degree nodes in the boundary, we build, for each social network, two ranks r

and r_b . r is a rank of all nodes in the graph sorted by degree, in descending order of popularity. r_b ranks the same nodes, but according to d_b , i.e., the number of cross-boundary connections. We then use Spearman’s rank correlation coefficient (Spearman 1987) ρ to capture the statistical dependence between r and r_b . Spearman’s correlation captures how well the relationship between two variables can be described by a monotonic function and its value ranges from -1 to $+1$. $\rho(X, Y) = 1$ means that variable Y is a perfectly monotonic function of X . In our context, a high ρ means that high-ranked nodes in the graph tend to be also high-ranked in the boundary, indicating a concentration of high-degree nodes along the boundary. A low ρ indicates that many high-degree nodes in the graph are low-ranked in r_b , what indicates that there is a significant number of popular nodes which do not belong to the boundary.

Figure 6 compares overall and boundary ranking positions for nodes in the University social network. Note that high-ranked nodes in r tend to also be high-ranked in r_b , and $\rho = 0.84$ indicates that the network promotes a convergence of popular nodes to the boundary. We interpret this result as a strong indication of absence of polarization.

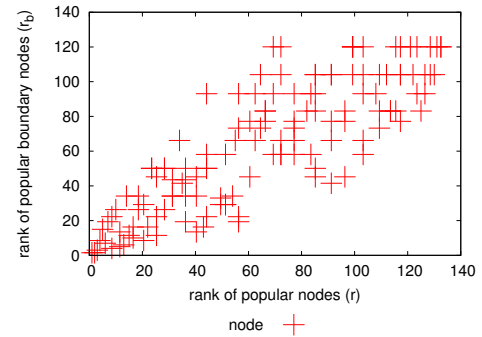


Figure 6: Concentration of popular nodes in the boundary for University network – $\rho = 0.84$. The high value of ρ indicates absence of polarization.

In Figure 7 we show r and r_b for the nodes belonging to non-polarized communities 2–3 in the gun control network. This graph is better interpreted when compared to Figures 8 and 9, which exhibit the corresponding results for polarized communities 1–2 and 1–3, respectively. Since nodes that exhibit the same degree are tied in the rankings, we added to

each rank position a random value between 0 and 5% of its absolute value to allow a better visualization of point density. Note that, in Figure 7, a large number of high-ranked nodes in r are also high-ranked boundary nodes in r_b . A large concentration of nodes is observed in the range 1–5000 of r and r_b in this pair of communities, in comparison to Figures 8 and 9. The ρ value is also significantly higher in the case of Figure 7 ($\rho = 0.70$), supporting our intuition w.r.t. the relationship between the concentration of high-degree nodes along the boundary and the existence of polarization.

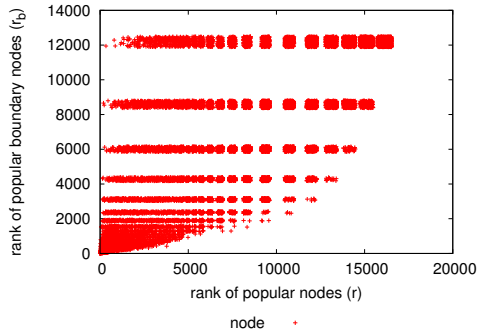


Figure 7: Concentration of popular nodes on the boundary – Gun Control communities 2–3 – $\rho = 0.70$. The high value of ρ indicates absence of polarization.

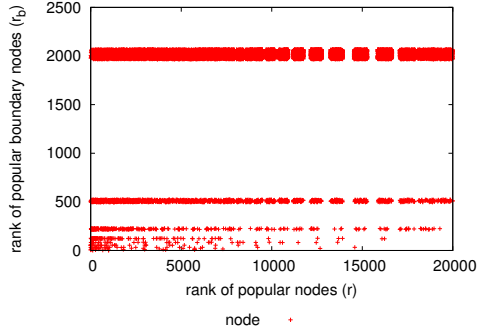


Figure 8: Concentration of popular nodes on the boundary – Gun Control communities 1–2 – $\rho = 0.21$. The low value of ρ indicates presence of polarization.

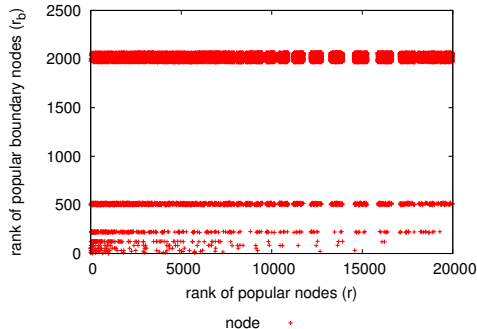


Figure 9: Concentration of popular nodes on the boundary – Gun Control communities 1–3 – $\rho = 0.23$. The low value of ρ indicates presence of polarization.

Table 5 shows ρ measurements for the other social networks we consider in this work. We note that, although polarized networks tend to exhibit low values of ρ , this is not always true. The U.S. Political blogs has a concentration of popular nodes in the boundary which is equivalent to the NYC-Teams network, despite of the differences in both Q and P . A possible explanation for such differences is that the political domains count with many media outlets that connect to both sides and thus gain popularity in the boundary, despite the polarized context.

7 Conclusions and Future Work

Literature of polarization of opinions in social networks has focused attention on domains previously known to induce polarization; as a consequence, the necessary and sufficient structural characteristics of polarized social networks are unclear. In this paper we perform a comparison between polarized and non-polarized networks and propose a new metric designed to measure the degree of polarization between two communities. Unlike modularity, which simultaneously measures homophily and antagonism between groups, our metric focus on the existence (or absence) of antagonism between the groups. We consider nodes’ decisions towards connecting to users who belong to the other (potentially opposing) group, in comparison to connect to members of its own group. Furthermore, we have shown that polarized networks tend to exhibit a low concentration of high-degree nodes in the boundary between two communities.

In practical applications, we believe that modularity and our metrics P and ρ can be used together and complementarily. As future work, we plan to build models based on the observation that the lack of positive interaction between nodes can also indicate *indifference* among them, rather than antagonism. Accounting for these effects is crucial in scenarios of multipolarization, which include multi-partisan political systems and sports competitions (for instance, FIFA Soccer World Club has 32 competing nations). In scenarios of multipolarization, we begin to observe more complex relationships among sides, rather than the duality support/antagonism – such as indifference between specific groups. Moreover, we will also observe social networks under a temporal perspective and look for “tipping points” of polarization, and work on graph generation models built upon the distributions we found in this work.

Acknowledgments

This work was supported in part by InWeb - National Institute of Science and Technology for the Web, CNPq, Fapemig, and by UOL (www.uol.com.br) through its Research Scholarship Program, Proc. Number 20120130151500.

References

- Abramowitz, A., and Saunders, K. 2005. Why can’t we all just get along? the reality of a polarized america. *The Forum: A Journal of Applied Research in Contemporary Politics* 2.
- Adamic, L. A., and Glance, N. 2005. The political blogosphere and the 2004 u.s. election: divided they blog. In *Pro-*

Table 5: Modularity Q , Polarization P and Spearman's Correlation ρ for networks described in Table 1.

network	media	modularity Q	polarization P	ρ
1 - NYC Teams	Twitter	0.15	-0.002	0.65
2 - University	Facebook	0.24	-0.24	0.84
3 - Karate's Club	friendships	0.35	0.17	0.62
4 - Brazilian Soccer Teams	Twitter	0.39	0.20	0.39
5 - US Political Blogs	blogs	0.42	0.18	0.65

ceedings of the 3rd international workshop on Link discovery, LinkKDD '05, 36–43. New York, NY, USA: ACM.

Bengio, Y.; Schuurmans, D.; Lafferty, J. D.; Williams, C. K. I.; and Culotta, A., eds. 2009. *Bayesian Belief Polarization*. Curran Associates, Inc.

Blendon; Young; and Hemenway. 1996. The american public and the gun control debate. *JAMA* 275(22):1719–1722.

Blondel, V. D.; Guillaume, J.-L.; Lambiotte, R.; and Lefebvre, E. 2008. Fast unfolding of community hierarchies in large networks. *CoRR* abs/0803.0476.

Calais, P. H.; Veloso, A.; Meira, Jr, W.; and Almeida, V. 2011. From bias to opinion: A transfer-learning approach to real-time sentiment analysis. In *Proc. of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.

Conover, M.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2011. Political polarization on twitter. In *Proc. 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

Dixit, A. K., and Weibull, J. W. 2007. Political polarization. *Proceedings of the National Academy of Sciences* 104(18):7351–7356.

Easley, D., and Kleinberg, J. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.

Fortunato, S., and Barthélemy, M. 2007. Resolution limit in community detection. *Proceedings of the National Academy of Sciences* 104(1):36–41.

Friedkin, N. E. 1998. *A Structural Theory of Social Influence (Structural Analysis in the Social Sciences)*. Cambridge University Press.

Good, B. H.; de Montjoye, Y. A.; and Clauset, A. 2010. Performance of modularity maximization in practical contexts. *Physical Review E* 81(4):046106+.

Heider, F. 1958. *The Psychology of Interpersonal Relations*. New York: John Wiley & Sons.

Hunter, J. 1992. *Culture Wars: The Struggle to Define America*. Sociology. Religion. BasicBooks.

Isenberg, D. J. 1986. Group polarization: A critical review and meta-analysis. *Journal of Personality and Social Psychology* 50(6):1141–1151.

Kannan, R.; Vempala, S.; and Vetta, A. 2004. On clusterings: Good, bad and spectral. *J. ACM* 51(3):497–515.

Livne, A.; Simmons, M. P.; Adar, E.; and Adamic, L. A. 2011. The party is over here: Structure and content in the 2010 election. In *Proceedings of the 5th Int'l AAAI Conference on Weblogs and Social Media (ICWSM)*.

Lord, C.; Ross, L.; and Lepper, M. 1979. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37(11):2098–2109.

McCright, A. M., and Dunlap, R. E. 2011. The politicization of climate change and polarization in the american pub-

lic's views of global warming, 2001–2010. *The Sociological Quarterly*.

McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27(1):415–444.

Mouw, T., and Sobel, M. 2001. Culture wars and opinion polarization: The case of abortion. *American Journal of Sociology* 106(4):913–943.

Newman, M. E. J. 2006. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* 103(23):8577–8582.

Nickerson, R. S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2:175–220.

Paul DiMaggio, John Evans, B. B. 1996. Have American's Social Attitudes Become More Polarized? *American Journal of Sociology* 102(3):690–755.

Spearman, C. 1987. The proof and measurement of association between two things. By C. Spearman, 1904. *The American journal of psychology* 100(3-4):441–471.

Sunstein, C. R. 2002. The Law of Group Polarization. *Journal of Political Philosophy* 10(2):175–195.

Tan, C.; Lee, L.; Tang, J.; Jiang, L.; Zhou, M.; and Li, P. 2011. User-level sentiment analysis incorporating social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '11*, 1397–1405. New York, NY, USA: ACM.

Walton, D. 1991. Bias, critical doubt, and fallacies. Number 28, 1–22. *Argumentation and Advocacy*.

Waugh, A. S.; Pei, L.; Fowler, J. H.; Mucha, P. J.; and Porter, M. A. 2009. Party polarization in congress: A social networks approach.

Weinberger, D. 2011. *Too Big to Know: Rethinking Knowledge Now That the Facts Aren't the Facts, Experts Are Everywhere, and the Smartest Person in the Room Is the Room*. Basic Books.

Yang, J., and Leskovec, J. 2012. Structure and overlaps of communities in networks. *CoRR* abs/1205.6228.

Yardi, S., and Body, D. 2010. Dynamic debates: An analysis of group polarization over time on twitter. *Bulletin of Science, Technology and Society* 5(30).

Zachary, W. 1977. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* 33:452–473.

Zhang, Y.; Friend, A. J.; Traud, A. L.; Porter, M. A.; Fowler, J. H.; and Mucha, P. J. 2008. Community structure in congressional cosponsorship networks. *Physica A: Statistical Mechanics and its Applications* 387(7):1705–1712.