



An empirical investigation of network polarization

Ruben Interian*, Celso C. Ribeiro

Institute of Computing, Universidade Federal Fluminense, Niterói, RJ 24210-346, Brazil



ARTICLE INFO

MSC:
05C78
62F03
91F99

Keywords:
Polarization
Homophily
Network
Graph
p-Value

ABSTRACT

This paper proposes and explores a new quantitative characterization of the polarization phenomenon in networks. New tools for evaluating the polarization of a network are presented. We first characterize the homophily of each node individually. We depart from the definition of a new measure of the homophily of the nodes of a network and we consider the homophily distribution over the nodes as a primary indicator of the strength of polarization. Next, to address the polarization of the network as a whole, a probabilistic approach is developed. The approach is based on the straightforward computation of empirical cumulative distribution functions of sampled data from the network. These empirical distributions provide a more insightful understanding of the status of the network. They may be used not only to compare the polarization of groups of nodes or entire networks, but also to estimate the impacts of external interventions in terms of node polarization. The usefulness of the approach is illustrated on several case studies associated with real-life data sets from different sources.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

Polarization is a widely known phenomenon that has been discussed by politicians, media, and researchers in recent years [1,2]. This subject has also attracted the attention of thinkers throughout history.

Since the 19th century, Mill, an important philosopher and political theorist, claimed that dialogue across lines of political difference is a key prerequisite for sustaining a democratic citizenry [3]. Arendt also asseverated that debate is irreplaceable for forming enlightened opinions that reach beyond the limits of one's own subjectivity to incorporate the standpoints of others [4]. More recently, several world leaders have often expressed concern about polarization problems caused by social media [5,6]. From sociologists to economists, many are interested in studying the behavior and interactions in social networks that rule the opinion formation process.

According to the Oxford Dictionaries, polarization is the *division into sharply contrasting groups or sets of opinions or beliefs* [7].

It is not a widely accepted fact that social media increase polarization. Although there are many arguments supporting this thesis [8,9], there are also opposing views [10]. However, it is known that social networks and mass media, like newspapers and blogs, are the place where this phenomenon manifests itself in a more strong way. Even if social and mass media do not contribute to increase polarization in modern societies, it is important to identify the mechanisms by which polarization arises, as well as the characteristics and the peculiarities of polarized networks.

* Corresponding author.

E-mail addresses: rinterian@ic.uff.br, rubenus@yandex.ru (R. Interian), celso@ic.uff.br (C.C. Ribeiro).




Fig. 1. Network of political blogs during the 2004 U.S. presidential election. Democratic and republican blogs are represented by blue and red circles, respectively. It is clear that this network is strongly polarized, since there are a huge number of edges connecting democratic blogs among themselves, a huge number of edges connecting republican blogs as well, and relatively few edges connecting a democratic blog and a republican blog. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Polarization is closely related to homophily (from Ancient Greek: *homo* = “self” and *philia* = “love”, love to oneself), the tendency of individuals to associate with others that are similar to themselves. To avoid misunderstandings, we regard polarization as an extreme expression of homophily.

Homophily, also called assortativity [11,12], has been widely studied by researchers. Several measures and models exist to characterize homophily. However, most of the work in the literature make use of a single numerical measure to define the homophily of some network in a given moment of time. An assortativity coefficient is proposed in [12] and applied to several networks, showing that homophily is a universal phenomenon. The probability of creating a new link between two individuals as a function of their similarity is studied in [13]. The inbreeding homophily measure that reflects the amount of bias towards same-type relationships is mentioned in [14]. The Pearson correlation coefficient is used in [11] as a measure of the preference of high-degree nodes to attach to other high-degree nodes.

The characterization and the properties of polarized networks are very helpful to provide a better understanding of the behavior of individuals and societies. In addition, they can also help policymakers to define policies based on these characteristics. Although quantitative methods are often difficult to use in real-life situations because they tackle models that are abstractions of concrete cases, they are very useful in moderation systems used to detect suspicious events or users, as well as for economists and marketing experts when dealing with social and complex networks [15].

In this work, we develop tools that can be used for evaluating the polarization of a network in a more deep way than that offered by single valued measures. We first characterize the homophily of each node individually and, then, the polarization of the network as a whole. Section 2 introduces some real-life test instances that will be used as case studies. The homophily of each node is defined in Section 3 and we consider the distribution of the homophily values over the nodes of the network as a primary indicator of the strength of polarization, which is analyzed and illustrated with some case studies. The probability of a node to be influenced by homophily is derived in Section 4 as an improved measure of network polarization. This new measure is used to assess the statistical relevance of the homophily value. Section 5 develops a probabilistic approach to compare the polarization of groups of nodes or entire networks. The approach is based on the straightforward computation of empirical cumulative distribution functions of sampled data from the network, which provide a more insightful understanding of the status of the network. The usefulness of the approach is illustrated on several case studies associated with real-life data sets. Concluding remarks are drawn in the last section.

2. Case studies

We consider as case studies a number of test data sets selected from very different sources. We assume that some of them may be polarized to some extent.

- *Books* – A network of books about U.S. politics sold by Amazon.com [16]. Edges between books represent frequent co-purchasing of those books by the same buyers. Most of the books are classified as conservative or liberal, and a small number of them as neutral.
- *Blogs* – A network of political blogs that emerged during the 2004 U.S. presidential election [17]. Blogs are divided into two groups: republican and democratic. Fig. 1 represents democratic and republican blogs by blue and red circles, respectively [18].
- *Trade* – A world trade network, built using information about bilateral trade data between countries. The data compiled by Aller et al. [19] was obtained from the United Nations COMTRADE database. We consider that two countries are connected by an edge if the amount of trade between them is at least 5% of the total traded by any two countries. In other words, we discard the edges that do not represent a significant trade for either country. We analyzed two groups




Fig. 2. Range of variation of the homophily value. The extreme case $h_i(v) = 0$ corresponds to the absence of polarization (heterophily). As the homophily value increases it reaches the other extreme case $h_i(v) = 1$, which corresponds to extreme polarization.

of countries: those that formerly belonged to Eastern bloc (full COMECON members in 1990) and those who were part of Western bloc (NATO members in 1990, mainly the same countries that participated in the Marshall Plan). Most of the available data refers to 2010 and we analyze whether there are still strong ties between countries in both groups, about 20 years after the end of the Cold War.

- *Game of Thrones* – A network of interactions between characters of the well known novel “A Song of Ice and Fire”, sometimes called “Game of Thrones” (GoT), by the name of the first book and the television series. A more detailed description of this instance can be found in [20]. Characters are classified according to the house (clan) to which they belong or of which they are vassals. There is also a classification of the characters according to their gender (to see the role of gender in network organization, see [21]).

All data for these four case studies are available in [22].

3. Node homophily

Let $G = (V, A)$ be a directed graph, where $V = \{v_1, v_2, \dots, v_n\}$ is the set of nodes and $A = \{a_1, a_2, \dots, a_m\} \subseteq V \times V$ is the set of arcs. Moreover, let $\mathcal{A} = \{A_1, A_2, \dots, A_q\}$ be a set of node groups defined on V , i.e., each $A_i \subseteq V$ for any $i = 1, \dots, q$. Each node $v \in V$ has an associated state $S(v) = \{i = 1, \dots, q : v \in A_i\}$, that reflects the set of groups to which node v belongs.

The groups may be interpreted as communities (as in social networks) or classes defined by some criterion. In case the groups define a partition of the node set V , i.e., $A_1 \cup A_2 \cup \dots \cup A_q = V$ and $A_i \cap A_j = \emptyset$ for any $i, j = 1, \dots, q : i \neq j$, each node belongs to one single group and its state is also referred to as its type.

For any node $v \in V$, we denote by $\mathcal{N}(v)$ the set of successors of v in G . We consider that any successor $u \in \mathcal{N}(v)$ wins influence on v (and, therefore, v is influenced by u), in the sense that v receives information from u .

The cardinality of $\mathcal{N}(v)$ is the out-degree $d(v)$ of node v . For any group $A_i \in S(v)$, the successors of v that belong to A_i form the set $\mathcal{N}_i(v) = \mathcal{N}(v) \cap A_i \subset \mathcal{N}(v)$. The cardinality of $\mathcal{N}_i(v)$ is the i -degree $d_i(v)$ of node v . The homophily of any node $v \in V$, with $d(v) > 0$, with respect to any group A_i , for $i = 1, \dots, q$, is defined by

$$h_i(v) \equiv \frac{d_i(v)}{d(v)}. \tag{1}$$

In other words, the homophily of a node $v \in V$ with respect to a group A_i , for $i = 1, \dots, q$, is the ratio between the number of successors of v that belong to the same group A_i as v and the number of successors of v . Of course, this definition only makes sense if $d(v) > 0$. The value of the homophily is a real number in the $[0,1]$ interval, where 0 suggests heterophily (preference for the opposite), while 1 indicates extreme homophily. If a node v belongs to only one group, i.e., $|S(v)| = 1$ and $S(v) = \{i\}$ for some $i = 1, \dots, q$, then the index i can be omitted and we denote $h(v) = h_i(v)$.

The homophily measure $h_i(v)$ is similar to the homophily index H_i defined by Currarian et al. [14], which denotes the average i -degree of the nodes in A_i , divided by the average degree of all nodes in this same group. Following our notation,

$$H_i = \frac{\sum_{j=1}^{|A_i|} d_i(v_j)}{\sum_{j=1}^{|A_i|} d(v_j)}. \tag{2}$$

The newly proposed node homophily measure $h_i(v)$ is more useful to fully describe the polarization of a network than the single value H_i . Therefore, we will use the $h_i(v)$ values for a node-centered homophily analysis of the network.

Let $w_i = |A_i|/|V|$ be the fraction of the nodes of the graph G that belong to group A_i , for $i = 1, \dots, q$. If $h_i(v)$ is close to w_i for a node $v \in V$, then $d_i(v) \approx w_i d(v)$, i.e., $d_i(v)$ is proportional to the number of nodes in group A_i . This means that if we randomly choose a node u from among the successors of v , then the probability that u belongs to A_i is approximately w_i , which is the same probability that a randomly selected node of the network belongs to A_i . In this case, we say that node v is *balanced*.

In case $d_i(v)$ is much greater than $w_i d(v)$, then the homophily of v is much larger than w_i and there is polarization of v . Finally, if $d_i(v)$ is significantly smaller than $w_i d(v)$, then we are in the presence of heterophily. Fig. 2 illustrates these concepts and the range of variation of the homophily value.

This analysis can be made for each node in the network. Therefore, the distribution of the homophily over the nodes of the network (or over some subset of them, such as one of the groups A_1, A_2, \dots, A_p) can be used to characterize its polarization.

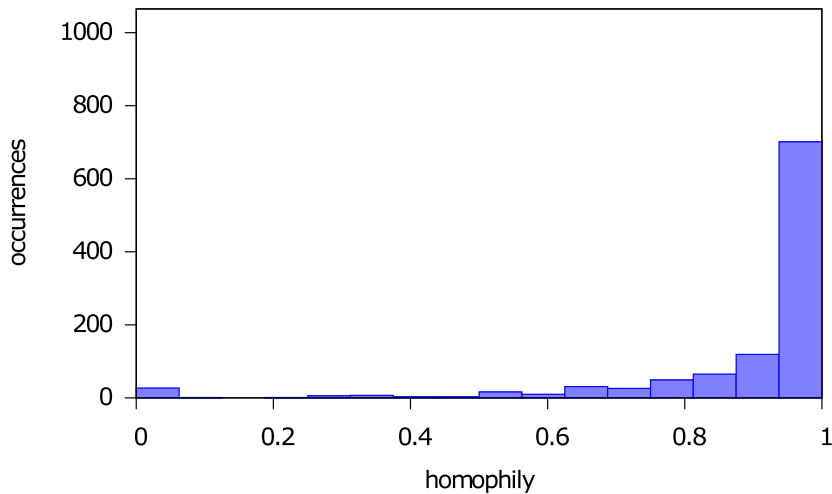


Fig. 3. Histogram representing the distribution of homophily values over the nodes of the political blogs network. The vertical axis indicates the number of occurrences of the homophily values in each interval of the horizontal axis.




Fig. 4. Histogram that represents the distribution of homophily values over the male character nodes of the Game of Thrones network. The vertical axis indicates the number of occurrences of the homophily values in each interval of the horizontal axis.

The distribution of the homophily values over the nodes of the network can be represented by a histogram, with the frequency of occurrences in each interval displayed in the vertical axis. Fig. 3 shows the case for the political blogs data set [17]. In this network, both groups (republican-affiliated and democrat-affiliated) form a partition of the node set and are of similar size ($w_{republican} = 0.491$ and $w_{democrat} = 0.509$). Homophily values in the histogram are mostly clustered near the value 1, far from 0.5. Therefore, the neighborhoods of a significant number of nodes of this network are mainly formed by same-type nodes, indicating polarization.

Next, as a second example, we consider the case of homophily by gender in the Game of Thrones network. There are two groups in this network, formed by male and female characters. We analyze the homophily values exclusively for the male characters. We are interested in knowing if male characters are more prone to interact with other male characters or not. Fig. 4 illustrates the occurrences of the homophily values in this network. If the network is not polarized, then the homophily values should average approximately $w_{males} = 0.757$, corresponding to the proportion of male characters in the node set V . Although we observe some clusterization of the histogram around this value, there is also a concentration of nodes with a high homophily value leaning towards 1, in the right extremity of the histogram. We abstain for now from assuming any further hypothesis in this case, and we return to this example in the next section.




Fig. 5. Example of a graph with five black and five grey nodes for which high values of the homophily measure for some specific nodes do not necessarily indicate polarization. For example, the fact that $h(v_3) = 1$ does not indicate “polarization” of node v_3 .

4. Polarization test

We note that although the homophily values are relevant, they may somehow be incomplete. Let us consider the graph in Fig. 5, with five black and five grey nodes. Each edge represents two arcs, one in each direction. The number of arcs between same-type nodes is equal to the number of arcs between different-type nodes. Therefore, this graph is not polarized. However, many nodes of the graph have different numbers of black and grey neighbors. In this example, both nodes v_1 and v_2 have two same-type neighbors and one neighbor of the other type. Therefore, $h(v_1) = h(v_2) = \frac{2}{3}$. Since v_3 has only one neighbor, $h(v_3) = 1$. The homophily value can be high for some nodes simply by chance, either because of their small degrees, or because it is impossible to divide an odd number by two, often lacking a greater meaning.

Therefore, “high” homophily values for some specific nodes do not necessarily indicate polarization.

In order to avoid the misuse of the raw homophily values we propose to use the binomial test to assess the statistical significance of $d_i(v)$ and $d(v)$ values as indicators of the polarization of node v .

Two definitions are relevant to establish whether a network is polarized or not. We assume that the networks considered in this work are such that, whenever a new arc (v, u) originating from node v is added to the graph, there is a probability $p_i \in [0, 1]$ that arc (v, u) is formed by same-type neighbors, where p_i depends only on the group A_i , $i = 1, \dots, q$, to which node v belongs to. More formally:

Definition 4.1 (Linking probability). Let $G = (V, A)$ be a directed graph with a set $\mathcal{A} = \{A_1, A_2, \dots, A_q\}$ of node groups defined on V , with $A_i \cap A_j = \emptyset$ for any $i, j = 1, \dots, q : i \neq j$. We say that $p_i \in [0, 1]$ is the linking probability of group A_i if, for each $v \in A_i$ and for any arc (v, u) that is added to the graph, the probability that node u also belongs to A_i is equal to p_i .

We observe that the linking probability p_i is the same for all nodes in group A_i . The concept of linking probabilities can be illustrated and better understood with some examples. Consider e.g. the case of the political blogs in the internet, associated with some ideological or political groups that establish new hyperlinks reflecting the information sources that exert influence on them. If a blog is biased to some extent, it is more likely (i.e., it happens with higher probability) to establish links with other same-type blogs. If we consider the trade network, we should expect that countries often establish commercial ties with other countries that are ideologically more similar to themselves. Therefore, we made the assumption that, in each case, there is a linking probability of establishing a same-type link from any node of a given group, and this probability is somehow related to the strength of the bias of this group.

Definition 4.2 (Balanced graph). A directed graph $G = (V, A)$ with a set $\mathcal{A} = \{A_1, \dots, A_q\}$ of node groups defined on V , with $A_i \cap A_j = \emptyset$ for any $i, j = 1, \dots, q : i \neq j$, is said to be balanced if, for each A_i , the linking probability p_i is equal to w_i , where $w_i = |A_i|/|V|$.

Now, let $G = (V, A)$ be a directed graph with a set $\mathcal{A} = \{A_1, A_2, \dots, A_q\}$ of node groups defined on V , with $A_i \cap A_j = \emptyset$ for any $i, j = 1, \dots, q : i \neq j$. Moreover, let $v \in V$ be a node of type $i \in \{1, \dots, q\}$ with $n = d(v)$ successors and p_i be the linking probability of group A_i , $i = 1, \dots, q$. Then, the number $d_i(v)$ of same-type successors of v follows a binomial distribution with parameters n and p_i :

$$P(d_i(v) = k) = \binom{n}{k} \cdot p_i^k (1 - p_i)^{n-k}.$$

We observe that if $G = (V, A)$ is a balanced graph, then arcs are formed indistinctly: in other words, the probability that a successor of node $v \in A_i$ be a same-type node is $w_i = |A_i|/|V|$. Therefore, the null hypothesis for the binomial test is

$$H_0 : p_i = w_i.$$

The alternative hypothesis is

$$H_1 : p_i > w_i$$

and, therefore, G is not a balanced graph. Consequently, node v is polarized to some extent.




Fig. 6. Histogram that represents the distribution of p -values over the nodes of the political blogs network. The vertical axis indicates the number of occurrences of the p -values in each interval of the horizontal axis.

We do not consider the case $p_i < w_i$, since in typical networks it is not expected to find a significantly large number of arcs with different-type extremities. Even if for some node we find a disproportionately large number of such inter-group arcs, and a too small number of intra-group arcs, the goal of the p -value measure is to evaluate the strength of the polarization of that node. Therefore, we consider that the possible presence of heterophily in some nodes is irrelevant to our analysis. Consequently, the one-tailed binomial test applies in this case.

If we choose some level of significance, such as 5%, we can decide if the null hypothesis is rejected by calculating the probability $p\text{-value}(v)$ of observing a number of same-type successors that is greater than or equal to $d_i(v)$. If $p\text{-value}(v)$ is less than or equal to the significance level, then the null hypothesis should be rejected. If $p\text{-value}(v)$ is greater than the significance level, then the null hypothesis cannot be rejected.

In the example of Fig. 5, the probability of having two or more same-type neighbors among three nodes is 0.5. Therefore, we cannot reject the null hypothesis at the 5% significance level. There is not enough information to conclude that the bias we observe is due to polarization and not to chance. However, if instead there was a node with 30 successors, with 20 of them having its same type, then the p -value would be approximately 0.04937 and we could reject the null hypothesis at the 5% level.

Therefore, we can see that two nodes with the same homophily measure $2/3$ computed by Eq. (1) can have very different probabilities of being truly influenced by homophily. The homophily measure $h(\cdot)$ shows that some bias is present, but the p -value is a stronger measure that indicates how significant the observed bias is.

This is partially explained by the fact that the measure $h(v)$ summarizes two informative variables in a single one, normalizing the number of same-type successors by dividing it by the total number of successors, thus losing usable information.

4.1. Distribution of p -values

Given the null hypothesis, it is well known [23] that the p -values are uniformly distributed in the interval $[0,1]$. Since the number of nodes in real-life networks is finite and some node degrees are small, the p -value distribution may be not perfectly uniform. Given their interpretation, the occurrence of small p -values indicate that the network is unbalanced. Therefore, a significant concentration of p -values towards zero means that the graph is polarized.

Fig. 6 shows the distribution of p -values over the nodes of the political blogs data set [17]. The values are mostly clustered close to zero, indicating polarization.

In contrast, Fig. 7 illustrates the distribution of p -values for the male characters of the Game of Thrones network, which is much more uniform than the precedent. The complete absence of small p -values (lower than any of the commonly used significance levels, such as 5% or even 10%) is striking, meaning that there is no statistically relevant homophily values that justify the rejection of the null hypothesis, therefore indicating the absence of polarization.

4.2. Evaluating the homophily of a network

The distribution of p -values summarizes the information about polarization of some group of nodes or network. In addition to the p -value histograms, we can present this information by choosing some values for the significance level α and counting the number of nodes (or the fraction of nodes) whose p -values are smaller than each value of α . A large fraction

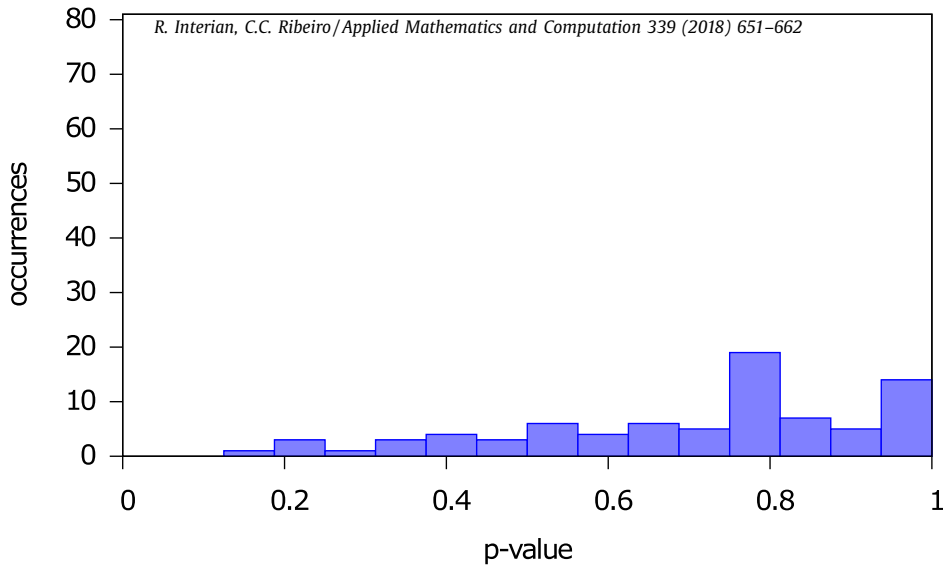


Fig. 7. Histogram that represents the distribution of p -values over the male character nodes of the Game of Thrones network. The vertical axis indicates the number of occurrences of the p -values in each interval of the horizontal axis.

Table 1

Fraction (in percent) of the nodes in each group whose p -value is below the significance level α , corresponding to the more polarized nodes.

Instance name	Group	Significance level, α			
		0.01	0.02	0.05	0.10
Political books	Conservative	40.8	49.0	73.5	73.5
	Liberal	53.5	69.8	81.4	83.7
	Neutral (*)	0.0	0.0	15.4	23.1
	All groups	41.0	51.4	69.5	71.4
Political blogs	Republican	53.8	57.4	62.3	68.6
	Democratic	44.2	48.5	53.2	58.5
	All groups	49.2	53.1	57.9	63.8
World trade	Eastern bloc	60.0	68.0	76.0	80.0
	Western bloc	68.8	75.0	81.3	87.5
Game of Thrones (houses)	Stark-Arryn	25.0	58.3	62.5	79.2
	Lannister	41.7	50.0	54.2	66.7
Game of Thrones (gender)	Male (*)	0.0	0.0	0.0	0.0
	Female (*)	0.0	0.0	0.0	7.7
	All groups (*)	0.0	0.0	0.0	1.9

of nodes with p -values smaller than the significance level α corresponds to a large number of polarized nodes. The higher the former is, the more polarized the network is.

Table 1 shows the fraction of the nodes of the network that satisfy this condition for the most common significance level values 1%, 2%, 5%, and 10%. If the groups cover the entire set of nodes, then every node in V has an associated p -value. In this case, we also show the ‘all groups’ row, that contains the fraction p -values that are smaller than the significance level α in the whole network.

We observe that, in most cases in this table, a substantial part of the network is polarized. The only exceptions are the groups in the instance corresponding to gender classification in the Game of Thrones network and the small neutral group in the political books network (both marked with an ‘*’). This means, for the case of the Game of Thrones network, that there is no gender assortativity over the characters, and this can be interpreted as the absence of any sexist pose of the author. Similarly, for the case of the political books network, buyers of neutral books often buy books of different tendencies besides the neutral ones.

5. Comparing polarized groups and networks

We have shown that the p -values are a good indicator of how likely it is to observe polarization and that smaller p -values point to more polarization. In this section, we use the tools proposed in this work to compare the strength of polarization of groups and networks. In order to compare two groups of nodes in terms of their polarization, we make use of the empirical distribution function [24] of the p -values over the nodes.

Table 2

Probability that the p -value of some group X be smaller than the p -value of another group Y .

Instance name	X	Y	$P(p\text{-value}(X) < p\text{-value}(Y))$
Political books	Liberal	Conservative	0.575
	Conservative	Liberal	0.425
Political blogs	Republican	Democratic	0.567
	Democratic	Republican	0.432
World trade	Eastern bloc	Western bloc	0.435
	Western bloc	Eastern bloc	0.565
Game of Thrones (houses)	Stark-Arryn	Lannister	0.536
	Lannister	Stark-Arryn	0.453
Game of Thrones (gender)	Male	Female	0.449
	Female	Male	0.538

The empirical distribution function (CDF) of a sample X_1, \dots, X_n of a real-valued random variable X is defined as

$$F_n(x) = \frac{1}{n} \cdot \sum_{i=1}^{i=n} 1 \cdot \{X_i \leq x\}, \quad \forall x \geq 0. \quad (3)$$

The empirical distribution function $F_n(x)$ of the random sample of size n is an estimator of the unknown distribution function $F_X(x)$ of the random variable X . This estimator has good statistical properties: it is unbiased and consistent, among other properties [24].

In the following, we explore this concept for comparing the polarization of groups of nodes in the same network, as well as the polarization of two different networks.

5.1. Statistical comparison of p -values

For each pair of groups X and Y of the networks considered in our case studies, Table 2 displays the probability $P(p\text{-value}(X) < p\text{-value}(Y))$ that $p\text{-value}(x)$ of a randomly selected node $x \in X$ be smaller than $p\text{-value}(y)$ of a randomly selected node $y \in Y$.

To compute $P(p\text{-value}(X) < p\text{-value}(Y))$, we simply divide the number of pairs of nodes (x, y) : $x \in X, y \in Y$ such that $p\text{-value}(x) < p\text{-value}(y)$ by the total number $|X| \cdot |Y|$ of pairs (x, y) : $x \in X, y \in Y$. Considering e.g. the case of the political books, let X be the set formed by 43 liberal books and Y be the set of 49 conservative books, amounting to $43 \times 49 = 2107$ (x, y) pairs. Since there are 1212 (x, y) pairs with $p\text{-value}(x) < p\text{-value}(y)$, then $P(p\text{-value}(X) < p\text{-value}(Y)) = \frac{1212}{2107} = 0.575$. Probabilities in the same instance do not necessarily add up to one, since there are pairs (x, y) : $x \in X, y \in Y$ of nodes with $p\text{-value}(x) = p\text{-value}(y)$.

For any of the instances in Table 2, we observe that $P(p\text{-value}(X) < p\text{-value}(Y)) \neq P(p\text{-value}(Y) < p\text{-value}(X))$. In principle, we could take the fact that $P(p\text{-value}(X) < p\text{-value}(Y)) > P(p\text{-value}(Y) < p\text{-value}(X))$ as an indication that group X is more polarized than group Y . However, we do not know a priori whether this lack of symmetry is sufficient or conclusive to compare the strength of polarization of the two groups X and Y . This issue will be addressed in the next section.

5.2. Comparing the polarization of two groups

In this section, we present and discuss the empirical distribution functions of the p -values associated to the groups in which some of our data sets are clustered.

Fig. 8 displays the p -value empirical CDFs (3) for the nodes of the groups representing conservative and liberal books. We observe that the conservative books p -value CDF first-order stochastically dominates [25] the liberal books p -value CDF, i.e., conservative books are less likely to be polarized. However, both groups are extremely polarized, since more than 70% of the p -values are below the 0.1 threshold. In fact, the CDF at 0.1 is approximately 0.735 for the conservative books and approximately 0.837 for the liberal books.

Similarly, Fig. 9 shows the p -value empirical CDFs for the nodes of the groups representing democratic and republican political blogs. In this case, the democratic blogs appear to be somewhat less polarized than the republican blogs although, once again, both are strongly polarized.

Fig. 10 compares the p -value CDFs of the two largest clans in the Game of Thrones network. There is no first-order stochastic dominance between them. However, most of the Lannister p -value CDF is below the Stark-Arryn p -value CDF, indicating more polarization on the Stark-Arryn clan, as already indicated by Table 2.

Finally, Fig. 11 presents the p -value CDFs for the gender groups in the Game of Thrones network. Although the values $P(p\text{-value}(X) < p\text{-value}(Y))$ and $P(p\text{-value}(Y) < p\text{-value}(X))$ in Table 2 are different for this instance, we do not observe first-order stochastic dominance. Instead, we observe that the plots are much closer to uniform distributions (corresponding

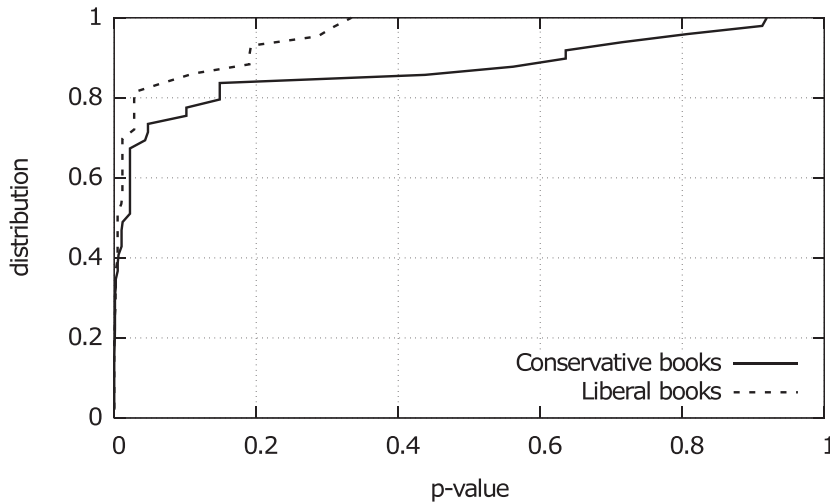


Fig. 8. Polarization in the network of conservative and liberal books: comparison of the p -value cumulative distribution functions. Both groups are extremely polarized, although the conservative books group less likely to be polarized.




Fig. 9. Polarization in the network of democratic and republican blogs: comparison of the p -value cumulative distribution functions. Both groups are extremely polarized, although the democratic blogs appear to be somewhat less polarized than the republican blogs.

to straight lines through the origin) and the p -values are greater than those in the previous figures, corresponding to the expected absence of polarization.

5.3. World trade case study

It is well established that there was very little trade between countries in the West and East countries before the nineties. We can use the tools proposed in this work to analyze how things have changed ever since.

We summarize in tabular form some relevant p -value information for the World trade network. We use this particular representation to explicitly represent p -values in order to show the five most open and the five least independent economies in 2010, with regard to these groups. Table 3 presents the results.

The results indicate that, in 2010, Russia and Ukraine remained commercially intertwined with other former Soviet Union states. Cuba and Vietnam, which were on the periphery of the Eastern bloc, no longer have any relevant commercial links with other former COMECON states. In the Western bloc, United States and Canada have more diversified economies, in the sense that they do not depend on other countries from the same group.

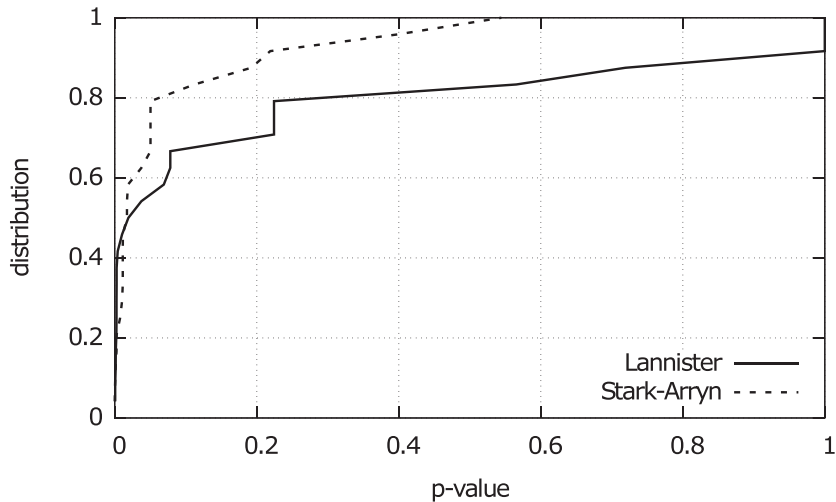


Fig. 10. Polarization in the Lannister and Stark-Arryn clans in the Game of Thrones network: comparison of the p -value cumulative distribution functions. Most of the Lannister p -value CDF is below the Stark-Arryn p -value CDF, indicating more polarization on the Stark-Arryn clan.




Fig. 11. Polarization by gender in the Game of Thrones network: comparison of the p -value cumulative distribution functions. There is no first-order stochastic dominance between the two plots, that are much closer to uniform distributions, corresponding to the expected absence of polarization.

5.4. Comparing the polarization of two networks

In the same way as we compared the polarization of two groups of nodes in Section 5.2, in some situations it is possible to compare the polarization of two different networks. In this case, it is required that the groups define a partition of the node set corresponding to each network: every node should belong to exactly one group and the p -value CDFs can be computed for each entire network.

The usefulness of this type of comparison can be illustrated with practical and realistic applications in two scenarios. First, consider a social network that is subject to some intervention (such as a marketing campaign) or, alternatively, severely modified by some external circumstances. If complete information on the network before and after the intervention (or external circumstance) is available, the p -value CDF provide a tool to compare the polarization of the networks representing the two states of the same physical entity, before and after the intervention (or external circumstance) and to evaluate its impacts. A second interesting scenario of application involves the comparison of the polarization of two different networks that represent collective behavior of the same group of people, such as two different social networks in the same country at the same time.

To illustrate this application, once again we consider the political books data set. Its node set V is partitioned into three groups: $A_{conservative}$, $A_{liberal}$, and $A_{neutral}$. Their average different-type out-degrees are, respectively, $\bar{d}_{conservative} = 0.94$,

Table 3
Openness of economies in terms of their p -values. The table presents the countries with the five largest and the five lowest p -values for each bloc.

Bloc	Feature	Country	p -value
East	Smaller p -values	Russia	3.681×10^{-11}
		Ukraine	1.459×10^{-6}
		Lithuania	1.106×10^{-5}
		Moldova	6.073×10^{-5}
		Latvia	1.084×10^{-4}
	Larger p -values	Turkmenistan	0.118
		Azerbaijan	0.155
		Mongolia	0.194
		Cuba	1.000
		Vietnam	1.000
West	Smaller p -values	Iceland	5.126×10^{-7}
		Norway	5.483×10^{-6}
		Germany	4.514×10^{-5}
		Luxembourg	8.619×10^{-5}
		Denmark	2.265×10^{-4}
	Larger p -values	Greece	0.017
		Turkey	0.022
		Italy	0.059
		Canada	0.491
		USA	0.563

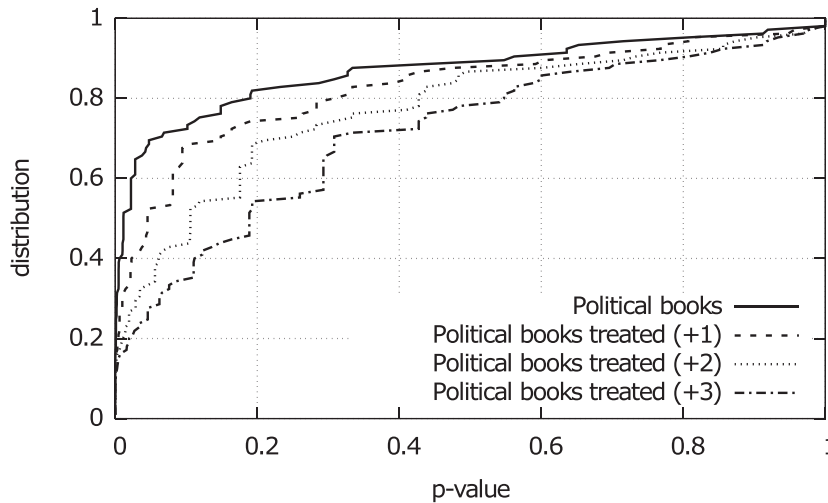


Fig. 12. Comparison of the p -value cumulative distribution functions of the original and treated political books instances with three levels of intervention. The strength of polarization decreases as the intervention level increases.

$\bar{d}_{liberal} = 0.84$, and $\bar{d}_{neutral} = 4.5$, with the overall average of different-type out-degrees being 1.30 and the average node degree being 8.40. This instance was previously shown to be strongly polarized in Section 5.2.

We assume now that this network is subject to some intervention and additional arcs connecting different-type nodes are inserted. Fig. 12 displays the p -value CDFs of the results observed by the simulation of three levels of intervention: with a low-level intervention, one new arc connecting each node to a randomly selected different-type node is added to the network. In the case of medium- or high-level interventions, the network is treated by the addition of two or three new arcs in each case, respectively.

We observe that as the level of the intervention increases and more arcs connecting nodes of different groups are added, the p -values CDF quickly becomes more uniform and the network becomes balanced, eliminating polarization. Considering e.g. the medium level intervention, the overall average of different-type out-degrees raises to 3.30 and the average node degree to 10.40. Since the fraction $2/10.40 \approx 19.2\%$ of new arcs in the resulting network is relatively low, we observe that a medium-level intervention might be sufficient to eliminate polarization. These results show that the p -value CDF is a powerful tool not only to evaluate the polarization of actual networks, but also the impact of interventions.

6. Discussion and concluding remarks

In this study, we proposed and explored a new quantitative characterization of the polarization phenomenon in networks. First, we defined the homophily of the node of a network and we analyzed the distribution of the homophily in a directed graph as a primary indicator of polarization.

Next, we used a probabilistic approach to define a new and improved polarization measure, which is based on the calculation, for each node, of the probability (p -value) of observing a number of same-type successors that is greater than or equal to the actual number of same-type successors observed for this node. We used the distribution of the obtained p -values for evaluating the strength of polarization of the network. The empirical distribution function of the p -values can be used to compare, in a more informative way, the polarization of different groups of nodes and even of entire networks. Several real-life networks from different sources have been used as case studies to illustrate the usefulness of the proposed tools.

The approach proposed in this work is generic and may be applied to a variety of real networks and situations. In particular, the p -value distributions can be used to estimate the impacts of interventions in the polarization of the nodes of a network. Also, the tool developed in this work could be applied into the identification of influential spreaders in networks [15] and for the promotion of collective cooperation within a regular cooperation network [26].

We consider that these tools are relevant in a world characterized by extreme political and ideological polarization.

Competing interests

The authors declare no competing interests.

Acknowledgments

Work of Ruben Interian was sponsored by a CAPES doctorate scholarship. Work of Celso C. Ribeiro was partially supported by CNPq research grants 303958/2015-4 and 425778/2016-9 and by FAPERJ research grant E-26/202.854/2017. We are also grateful for the constructive remarks of one anonymous referee.

References

- [1] New York Times, How we Became Bitter Political Enemies, 2017, <https://www.nytimes.com/2017/06/15/upshot/how-we-became-bitter-political-enemies.html>, last access on 2017-11-13.
- [2] The Economist, Not Persuaded – Political Polarization, 2017, <https://www.economist.com/blogs/democracyinamerica/2015/05/political-polarisation>, last access on 2017-11-13.
- [3] J.S. Mill, *On Liberty*, J. W. Parker and Son, London, 1859.
- [4] H. Arendt, *Between Past and Future*, Viking Press, New York, 1968.
- [5] B. News, Angela Merkel: The Big Internet Platforms Have Become an Eye of a Needle, 2017, <http://www.bbc.com/news/technology-37798762>, last access on 2017-09-26.
- [6] T.N. Yorker, Obama Reckons with a Trump Presidency, 2016, <http://www.newyorker.com/magazine/2016/11/28/obama-reckons-with-a-trump-presidency>, last access on 2017-06-01.
- [7] Oxford Dictionaries, Definition of Polarization, 2017, <https://en.oxforddictionaries.com/definition/polarization>, last access on 2017-09-06.
- [8] V. Almeida, D. Doneda, Mecanismos Invisíveis de Polarização Política, 2016, <http://www.valor.com.br/opiniaio/4768549/mecanismos-invisiveis-de-polarizacao-politica>, last access on 2017-11-13.
- [9] S. Hong, S.H. Kim, Political polarization on twitter: implications for the use of social media in digital governments, *Gov. Inf. Q.* 33 (4) (2016) 777–782.
- [10] P. Barberá, How social media reduces mass political polarization. Evidence from Germany, Spain, and the U.S., in: *Proceedings of the 2015 American Political Science Association Annual Meeting*, San Francisco, 2015.
- [11] M.E.J. Newman, Assortative mixing in networks, *Phys. Rev. Lett.* 89 (20) (2002) 208701.
- [12] M.E.J. Newman, Mixing patterns in networks, *Phys. Rev. E* 67 (2) (2003) 026126.
- [13] G. Kossinets, D.J. Watts, Empirical analysis of an evolving social network, *Science* 311 (2006) 88–90.
- [14] S. Currarini, M.O. Jackson, P. Pin, An economic model of friendship: homophily, minorities, and segregation, *Econometrica* 77 (2009) 1003–1045.
- [15] C. Li, L. Wang, S. Sun, C. Xia, Identification of influential spreaders based on classified neighbors in real-world complex networks, *Appl. Math. Comput.* 320 (2018) 512–523.
- [16] M.E.J. Newman, *Network Data – Books About US Politics*, 2017, <http://www-personal.umich.edu/~mejn/netdata/>, last access on 2017-12-11.
- [17] L.A. Adamic, N. Glance, The political blogosphere and the 2004 U.S. election: divided they blog, in: *Proceedings of the Third International Workshop on Link Discovery*, ACM, New York, 2005, pp. 36–43.
- [18] C. Jones, Visualizing Polarization in Political Blogs, 2018, <http://allthingsgraphed.com/2014/10/09/visualizing-political-polarization/>, last access on 2018-01-13.
- [19] C. Aller, L. Ductor, M.J. Herreras, The world trade network and the environment, *Energy Econ.* 52 (2015) 55–68.
- [20] A. Beveridge, J. Shan, Network of thrones, *Math Horiz. Mag.* 23 (4) (2016) 18–22.
- [21] I. Psylla, P. Sapiezynski, E. Mones, S. Lehmann, The role of gender in social network organization, *PLOS One* 12 (12) (2017) 1–21.
- [22] R. Interian, C.C. Ribeiro, Data for Case Studies in ‘An Empirical Investigation of Network Polarization’, 2017, <https://data.mendeley.com/datasets/xxwtd92pb/1.1017632/xxwtd92pb.1>
- [23] C. Blocker, J. Conway, L. Demortier, J. Heinrich, T. Junk, L. Lyons, G. Punzig, Simple facts about p -values, Technical Report, Laboratory of Experimental High Energy Physics, Rockefeller University, New York, 2006.
- [24] A.W. van der Vaart, *Asymptotic Statistics*, first, Cambridge University Press, 1998.
- [25] J. Hadar, W.R. Russell, Rules for ordering uncertain prospects, *Am. Econ. Rev.* 59 (1969) 25–34.
- [26] C. Xia, S. Ding, C. Wang, J. Wang, Z. Chen, Risk analysis and enhancement of cooperation yielded by the individual reputation in the spatial public goods game, *IEEE Syst. J.* 11 (2017) 1516–1525.