Measuring political polarization: Twitter shows the two sides of Venezuela

A. J. Morales, J. Borondo, J. C. Losada, and R. M. Benito

Articles you may be interested in

Resource Letter MPCVW-1: Modeling Political Conflict, Violence, and Wars: A Survey
Am. J. Phys. **81**, 805 (2013); 10.1119/1.4820892

Temporal dynamics and impact of event interactions in cyber-social populations
Chaos **23**, 013131 (2013); 10.1063/1.4793540

Characterizing and modeling an electoral campaign in the context of Twitter: 2011 Spanish Presidential election as a case study
Chaos **22**, 023138 (2012); 10.1063/1.4729139

Plutonium, Politics and Policy: A New "Atoms for Peace"
AIP Conf. Proc. **673**, 35 (2003); 10.1063/1.1594542

Memoirs: A Twentieth-Century Journey in Science and Politics
Am. J. Phys. **70**, 462 (2002); 10.1119/1.1456079

Measuring political polarization: Twitter shows the two sides of Venezuela

A. J. Morales, J. Borondo, J. C. Losada, and R. M. Benito
Grupo de Sistemas Complejos, Universidad Politécnica de Madrid, ETSI Agrónomos, 28040 Madrid, Spain

We say that a population is perfectly polarized when divided in two groups of the same size and opposite opinions. In this paper, we propose a methodology to study and measure the emergence of polarization from social interactions. We begin by proposing a model to estimate opinions in which a minority of influential individuals propagate their opinions through a social network. The result of the model is an opinion probability density function. Next, we propose an index to quantify the extent to which the resulting distribution is polarized. Finally, we apply the proposed methodology to a Twitter conversation about the late Venezuelan president, Hugo Chávez, finding a good agreement between our results and offline data. Hence, we show that our methodology can detect different degrees of polarization, depending on the structure of the network. © 2015 AIP Publishing LLC. [http://dx.doi.org/10.1063/1.4913758]

Political polarization is a social phenomenon that has several consequences in peoples' lives and whose nature is not completely understood. Here, we develop a methodology to infer political polarization from Twitter conversations and propose a general index to measure polarization that can be applied to several cases. We apply our methods to a Venezuelan case study and show the relationship of the polarized users and the offline world. We demonstrate that the polarization in social media emerges from the influence of opinion leaders that dominate interactions.

I. INTRODUCTION

From a sociological point of view, polarization is a social phenomenon that appears when individuals align their beliefs in extreme and conflicting positions, with few individuals holding neutral or moderate opinions.[1,2] Thus, as a process, it is the increase of such divergence over time when people evaluate issues of diverse nature,[3–5] like politics or religion. In words of John Turner: 'Like polarized molecules, group members become even more aligned in the direction they were already tending.'[6]

In this paper, we propose a methodology to study the emergence of political polarization and quantify its effects. To this end, we introduce a model to estimate opinions, and a polarization index that quantifies to which extent the resulting distribution of opinions is polarized. We say that a population is perfectly polarized when divided in two groups of the same size and with opposite opinions. Hence, our measure of polarization is inspired by the electric dipole moment—a measure of the charge system's overall polarity. For two opposed point charges, the electric dipole moment increases with the distance between the charges. Analogously, the polarization of two equally populated groups depends on how distant their views are.

As Downs argued in 1957,[7] political discussion among individuals minimizes the cost of becoming politically informed. In other words, sensible individuals tend to rely in the opinions of experts instead of analyzing information by their own. In fact, several observational studies support this theory and suggest that the expertise distribution within a social network affects the political communication patterns.[8] Hence, by controlling the opinion of a minority of influential individuals and mapping the communication fluxes among the population, we can estimate their distribution of opinions. To this end, we propose a model based on DeGroot model.[9] The original model proposed by DeGroot describes how a group of individuals might reach a shared opinion, by iteratively updating their opinion as the average of their current opinion with the opinions of their neighbors. Such global coordination, without centralized control, can also be efficiently achieved when individuals adopt the majority state of their neighbors, even in the presence of noise or complex topologies.[10] Recently, the DeGroot model has been used to study the conditions under which consensus is achieved.[11–13] However, as consensus is rarely reached in real world,[14,15] variants of this model can held to a diversity of opinions.[16–19]

In contrast to opinion generation models, such as the voter model,[20–22] we do not aim to study the evolution of opinions, but to infer a distribution of opinions formed on a social network from which to measure polarization. In our model, a minority of influential individuals propagate their opinions through a directed network influencing the remaining individuals. Thus, each individual iteratively updates her opinion according to her incoming neighbors-those influencing her. Hence, by taking advantage of complex network analysis,[23] we are able to estimate the opinion of the whole majority that *a priori* was unknown. The behavior of the influential minority is similar to zealots in the voter model,[24,25] but their impact in the model's dynamics is different. In our model, zealots, rather than preventing consensus, allow us to infer the opinions of all the nodes in the network. Contrary to the voter model where opinions are binary (0 or 1), the opinions in our model represent a continuous distribution. In absence of polarization, the expected resulting distribution of opinions would be a narrow distribution centered at a

25, 033114-1

neutral opinion. However, as polarization emerges, the resulting distribution shifts to a bimodal distribution with two peaks emerging around the two dominant and confronted opinions.[26]

How can political polarization be detected and therefore be fixed? Nowadays, digital traces of human collective behavior[28] represent an opportunity to detect and measure in real time different phenomena, such as polarization. In fact political segregation has already been observed on political blogs[29] or Twitter.[30–32] Recent research has shown that the most prominent and politically active users mainly interact with their own partisans,[29–31] leaving little space for real debate and cross ideological interactions. However, segregation does not necessarily imply polarization, as two separated groups of people that share the same opinion cannot be considered as polarized. Hence, in order for a population to be polarized, the opinions of the two groups should also be conflicting or opposed.[33] In the latter part of this paper, we show how to apply our methodology to online data gathered from Twitter in order to estimate individuals opinions and to measure the emergent political polarization. Twitter provides an interesting context in which to study polarization as it represents a wide variety of different types of communications, going from personal to those coming from traditional mass media. In this platform, a minority of *elite* users concentrates much of the collective attention, but still a big fraction of the content they produce reaches the mass through intermediaries or 'opinion leaders.'[34] In other words, the 'two-step-flow' of communication is still valid on Twitter.[35]

We begin this paper by proposing a model to estimate opinions in which a minority of influential individuals propagates their opinion through a social network influencing the opinions of the remaining individuals. Thus, the result of the model is a probability density function $p(X)$, that determines the fraction of individuals holding an opinion X. Next, we introduce the polarization index to measure the political polarization from the resulting opinion distribution. To illustrate the power of the methodology, we apply it to a Twitter conversation regarding the death announcement of the Venezuelan President (Hugo Chávez). Finally, we contrast the results with offline data.

II. ESTIMATING OPINIONS

We present a model to estimate the opinions of individuals who interact on a social network, in order to obtain their opinions distribution. In it, we distinguish two types of individuals, *elite* and *listeners*. The first ones have a fixed opinion and act like seeds of influence, while the opinion of the second ones depends on their social interactions. The model is fully specified by the following assumptions:

(1) **Initial Conditions:** The world is abstracted by a directed network, G, in which each individual is represented by a node and links account for influence rather than friendship or other kind of relationship. We define two different subset of nodes, S accounting for *elite*; and L, accounting for *listeners*. Additionally, we endow each *elite* with a parameter, X_s, that determines her opinion value and that will remain constant for the duration of

the model. X_s lies in the range, $-1 \leq X_s \leq 1$, where 1 and -1 represent the two extreme and confronted poles. Finally, we set an initially neutral opinion, $X_l(0) = 0$ to all *listeners*.

(2) **Opinion Generation:** At each iteration, *elite* nodes, S, propagate their opinions through the established network, G, influencing *listeners*, L. Hence, each listener iteratively updates her opinion value as the mean opinion value of her incoming neighbors. Thus, the opinion at time step, t, of a given listener, i, is given by the following expression:

$$X_i(t) = \frac{\sum_j A_{ij} X_j(t-1)}{k_i^{in}}, \tag{1}$$

where A_{ij} represents the elements of the network adjacency matrix, which is 1 if and only if there is a link from j to i, and k_i^{in} corresponds to her indegree. The process is repeated until all nodes converge to their respective X_i value, lying in the range $-1 \leq X_i \leq 1$. Thus, the results of the model are given in a density distribution of nodes' opinion values $p(X)$. Note that the opinions of individuals do not depend on their opinion in the previous step. This is because we are estimating their opinion that *a priori* was unknown, rather than studying the evolution of opinions.

The dynamics of the model is illustrated in Fig. 1, where we present a schema of the influence spreading process. Panel (a) visualizes the instantiation of the model where each *elite* node has been colored according to her opinion (red, $X_s = -1$; and blue, $X_s = +1$). Panels (b)–(e) show the dynamics of the influence process from the initialization (b) to the final converged state (e). Panels (f) and (g) visualize two empirical networks corresponding to a non polarized (f) and a polarized (g) case.

III. INTRODUCING A NEW MEASURE OF POLARIZATION IN OPINION DISTRIBUTIONS: THE POLARIZATION INDEX

We say that a population is perfectly polarized when divided in two groups of the same size and with opposite opinions. Hence, we propose a measure of polarization that quantifies both effects for the resulting X distribution obtained from our model. This definition is inspired by the electric dipole moment—a measure of the charge system's overall polarity. In the simplest case of two point charges of opposite signs ($-q$ and $+q$), the electric dipole moment is proportional to the distance among the charges. This is analogous to a simple scenario consisting of two persons with different ideologies; thus, the polarization depends on how conflicting their points of view are (i.e., the distance among the two ideologies).

We begin by calculating the population associated with each opinion (positive and negative). To this end, we define A^- as the relative population of the negative opinions ($X < 0$). By the same token, we define A^+ as the relative population of the positive opinions ($X > 0$). Hence, both variables can be expressed as
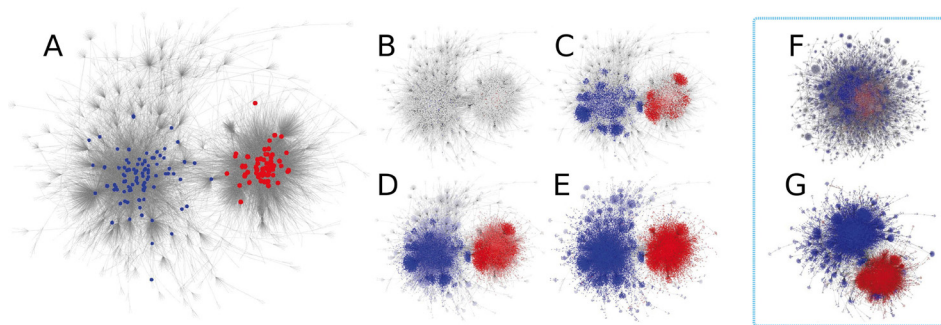
FIG. 1. Schema of the influence spreading process in the opinion estimation model. (a) displays the seed nodes in the network, colored according to their respective ideology. (b) displays the network at $t=0$, before seeds start to propagate their influence. (c) shows the state of the network at $t=1$. (d) shows the state of the network at $t=n/2$. (e) Displays the final state of the network at $t=n$. (f) and (g) Visualizations of two examples of the result of the opinion formation model to the Venezuelan dataset for non polarized (f) and polarized (g) days.

$$A^- = \int_{-1}^{0} p(X)dX = P(X < 0), \tag{2}$$

$$A^+ = \int_{0}^{1} p(X)dX = P(X > 0). \tag{3}$$

So we can express the normalized difference in population sizes, ΔA, as

$$\Delta A = |A^+ - A^-| = |P(X > 0) - P(X < 0)|. \tag{4}$$

Next, we quantify the distance between the positive and negative opinions. In other words, we measure how differing the opinions of the two sides are. To this end, we determine the gravity center of the positive and negative opinions that can be written as:

$$gc^- = \frac{\int_{-1}^{0} p(X)XdX}{\int_{-1}^{0} p(X)dX}, \tag{5}$$

$$gc^+ = \frac{\int_{0}^{1} p(X)XdX}{\int_{0}^{1} p(X)dX}, \tag{6}$$

and define the pole distance, d, as the normalized distance between the two gravity centers. Hence, it can be expressed as

$$d = \frac{|gc^+ - gc^-|}{|X_{max} - X_{min}|} = \frac{|gc^+ - gc^-|}{2}. \tag{7}$$

This formula gives $d=0$ when there is no separation between the gravity centers, i.e., there are no longer two differentiated groups and everyone shares a similar opinion; and $d=1$ when the two opinions are extreme and perfectly opposed.

Finally, we can use Eqs. (4) and (7) to write down a general formula to measure polarization as a function of the difference in size between both populations ΔA and the poles distance d. Thus, we define the *polarization index*, μ, as

$$\mu = (1 - \Delta A)d. \tag{8}$$

This formula gives $\mu = 1$ when the distribution is perfectly polarized. In this case, the opinion distribution function is two Dirac delta centered at -1 and $+1$, respectively. Conversely, $\mu = 0$ means that the opinions are not polarized at all, and the resulting distribution of opinions would either take the form of a single Dirac delta centered at a neutral opinion, or be entirely centered in one of the poles, implying that the population (A) of the other pole would be reduced to zero and $\Delta A = 1$. Notice that for non-uniform distributions centered in a neutral opinion, $|\mu| \ll 1$, but still presents a minimum polarization due to a small separation between gravity centers that depends on the standard deviation σ. In the case of a Gaussian distribution centered at zero, $\mu = \sigma\sqrt{2/\pi}$.

In between, polarization can lie within the range, $0 < \mu < 1$, for three reasons: (i) The population sizes associated to each opinion are equal, but the pole distance d is lower than 1. (ii) Despite d being equal to 1, the population sizes associated to each opinion are different, and therefore, there is a majority sharing a similar opinion. (iii) A combination of i and ii. Fig. 2(a) illustrates the basic concepts of the proposed index of polarization, as it visualizes the area associated to each opinion, their corresponding gravity centers, and the pole distance for a standard case of a perfect bimodal distribution. In panel (b) of this figure, we have visualized non polarized distributions ($\mu = 0$ and $|\mu| \ll 1$), a perfectly polarized one ($\mu = 1$), and a case in between.

IV. TWITTER DATA: THE VENEZUELAN CASE

In this section, we apply our model and polarization index to Twitter data regarding the late Venezuelan President Hugo Chávez. We downloaded over $16\,383\,490$ messages written by $3\,173\,090$ users from 02/04/2013 to 05/04/2013. This period covers one month preceding his death, the announcement of the death, and the schedule for new elections. We use retweets as a proxy for influence,[36–42] and build a weighted and directed network accounting for the adoption of ideas among Twitter users for each day. Whenever a user i retweets a message originally posted by user j, we assume that i is being influenced by j's ideas. Hence, a new directed link ($j \rightarrow i$) is created. We constructed an individual retweet network for each day of the observation period, which is a total of 56 networks. More details
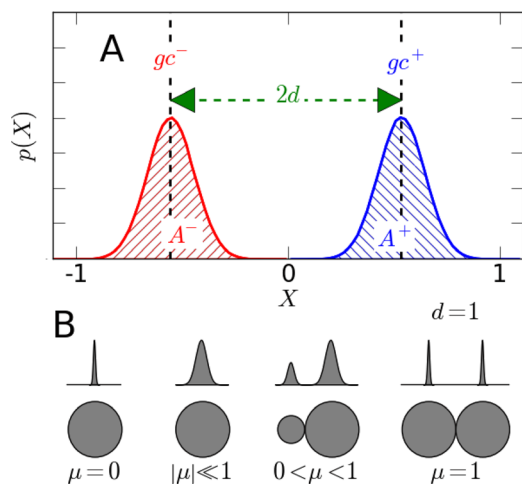
FIG. 2. Schema explaining polarization and the proposed index μ. (a) Density distribution of opinions. gc stands for the gravity center of each pole, A stands for the area associated to each ideology, and d stands for the pole distance. (b) Visualization of the polarization index, μ, given in Eq. (8), for four situations.
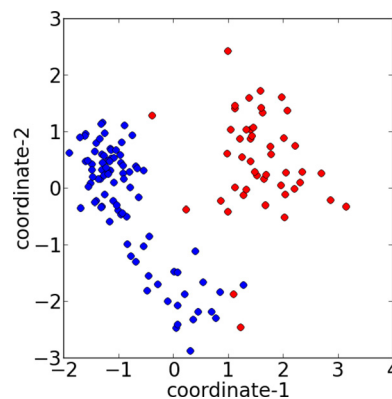

FIG. 3. Projection in a two-dimensional space of the distribution of *elite* users according to the similarity of their content. Dots represent users and colors indicate the community they belong to in the *elite* network: red for the officialism and blue for the opposition. The distance between users is inversely proportional to the similarity of their content.

about the dataset and the retweet networks can be found in Appendixes A and B, respectively.

In order to apply the model to these daily networks, we begin by defining a set of *elite* users. We denote as *elite* those users who gained a noticeable amount of retweets and actively participated in the conversation along the observation period. The distribution of users according to the total amount of retweets obtained (S_{out}) and participation rate (ρ) is shown in Fig. 9 of Appendix B. In this case, we considered a very small set (0.02%) of influential users who participated most of the observation period ($\rho > 89\%$) and obtained a very high number of retransmissions ($S_{out} > 1000$).

The *elite* users mainly correspond to politicians, journalists, and mass media accounts, whose political position and editorial tendency are publicly known and who belong to both sides of the Venezuelan political spectrum. In order to assign them an ideology value, X_s, we first studied their network of interactions. In the *elite* network, nodes represent the *elite* users, and links are created and accumulated whenever an *elite* user i retweets an *elite* user j. This network is polarized in a well defined two-community structure, with modularity $Q = 0.38$. In each community, users share political ideology and hardly interact with users from the other pole. In fact, the assortative mixing[43] by political ideology is very high ($r = 0.88$).

In order to further understand the *elite* polarization, we analyzed the content of their messages. For this purpose, we abstracted each *elite* user as a high-dimensional vector, where each element represents the number of times that the user posted each of the 500 mostly used words from all the *elite*'s messages. Then, we reduced the high-dimensional space into a two-dimensional one, by applying a multi-dimensional scaling algorithm.[44] In this algorithm, users are mapped into a new space by preserving the distance between them in the original one. This means that the distance between users is inversely proportional to the similarity of their posted contents. In Fig. 3, we present the projection of the users in the new two-dimensional space. Dots represent

users, and colors are assigned according to the community they belong to in the *elite* network. It can be noticed that these users are not homogeneously distributed in the new space. Instead, they are separated from each other in agreement with our previous classification. This means that the use of language is polarized among the *elite* users.

After identifying the *elite* users, we assigned them ideology values of $X_s = -1$ to the officialism side and of $X_s = 1$ to the opposition. The remaining users (99.98%) were assigned the role of *listeners* and $X_l = 0$. After running the model, we obtained an ideology probability density function $p(X)$ for each day. The resulting $p(X)$ for each network is presented in Fig. 4. The label indicates the day of observation, D representing the day of the death. The color indicates the network size in terms of the number of participants. As can be seen, the days with largest participation (purple and blue) correspond to the most important announcements: the presidents death (day D) and call for election (day $D + 6$). Next, we calculated the polarization index (μ), pole distance (d), and populations sizes for the resulting distributions of each day and plotted the results in Fig. 5.

We identify day D as a turning point which ended up polarizing even more the conversation. During the days preceding the announcement (from $D - 29$ to $D - 1$), X presents a bimodal distribution in which the officialism population (negative side of the X distribution) is considerably smaller than the opposition (positive side of the X distribution). This means that during this period, the conversation was still polarized, but practically monopolized by the opposition. Hence, despite the fact that the pole distance reached values over 0.9, the polarization index just averaged under 0.4. Then, a shift in the conversation emergent patterns took place on the day of the President's death announcement (day D). During this day, X lost its bimodal distribution, and the resulting $p(X)$ was centered around neutral values, minimizing the pole distance. All these meaning that the conversation was not so polarized and that the network does not have a two-island structure anymore. Therefore, the polarization index decreased, $\mu \approx 0.25$. This behavior is due to the bursty growth of the conversation at day D (see Fig. 7 in
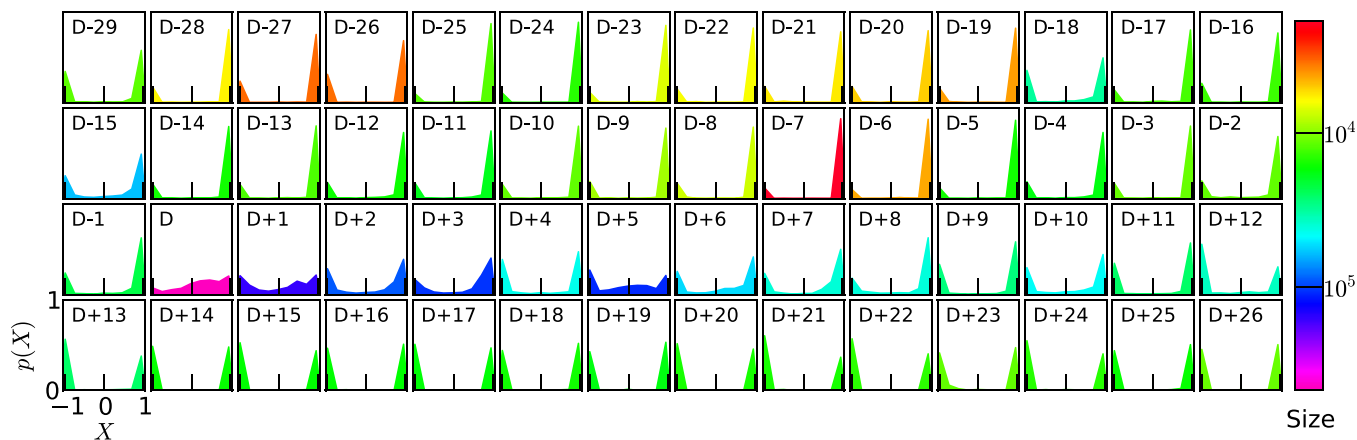
FIG. 4. Time evolution of ideological value (X_i) probability density functions ($p(X)$) for the Venezuelan conversation. Labels indicate the day of observation, D standing for the day of the Presidents death. Colors indicate the number of participants.

Appendix B). As a consequence, the previously segregated modules combined into a single-island structure, many times larger than the usual network size. Besides, a large amount of users from all around the globe joined to the conversation, making the topic international, rather than local from Venezuela. In fact, during this day, the percentage of users tweeting from Venezuela ($\approx 20\%$) was very low in comparison to the rest of the days (average around $> 80\%$). Hence, our set of Venezuelan *elite* was not capable of polarizing this majority of worldwide users. However, from thereon, the conversation recovered its bimodal distribution of opinions. Moreover, the polarization reached its maximum from day $D + 12$ (marked with the dashed line) onwards, day that the officialism new leader entered the conversation. From this day onwards, X presents a bimodal distribution, where the populations of both sides are similar. Therefore, the polarization index averaged values around 0.9.

V. TWITTER SHOWS THE TWO SIDES OF VENEZUELA

Next, we evaluate our model and the validity of Twitter data by comparing the geographic distribution of the polarized users with offline data regarding the Venezuelan socioeconomic and political landscape. More specifically, we analyze the geographical density of geolocated tweets in Caracas, the capital city of Venezuela, taking the results obtained from the most polarized days in Sec. IV as a proxy of their ideology. For this purpose, we have built the density functions that a tweet associated with the officialism or the opposition had been posted by a geolocated user at a given position (longitude and latitude). We considered a grid of 100 cells between longitudes $[-67.12°, -66.71°]$ and latitudes $[10.31°, 10.57°]$ and counted the number of tweets in each cell, identified with each ideology. Then, we normalized both counts by their respective total number of tweets. The resulting functions are two surfaces on top of the map, which we show in Fig. 6 as contour plots (red for the officialism and blue for the opposition) that indicate lines of equal
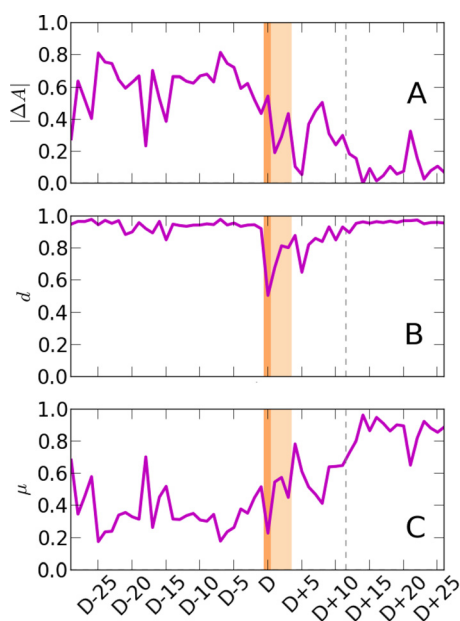

FIG. 5. Time evolution of the polarization index μ (c), and the variables associated with it: difference in population sizes (a) and pole distance d (b) for the Venezuelan conversation.
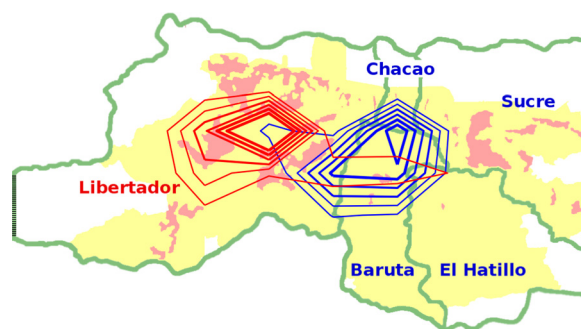

FIG. 6. Geographical polarization in the city of Caracas. Contour lines represent the density functions of the probability that a tweet associated with the officialism (red) or the opposition (blue) had been posted by a geolocated user at a given position (latitude and longitude). These contours have been superimposed to the map of Caracas, Venezuela. From inside out, contours indicate the following values: [0.175, 0.15, 0.0125 0.10, 0.075, 0.05]. The green lines border the five municipalities composing the city. Labels indicate the name of the municipality, and the color indicates the ruling party according to the 2013 Venezuelan local elections (red for the officialism party and blue for the opposition parties). White represents unpopulated areas, yellow urbanized areas, and pink the poorer neighborhoods.
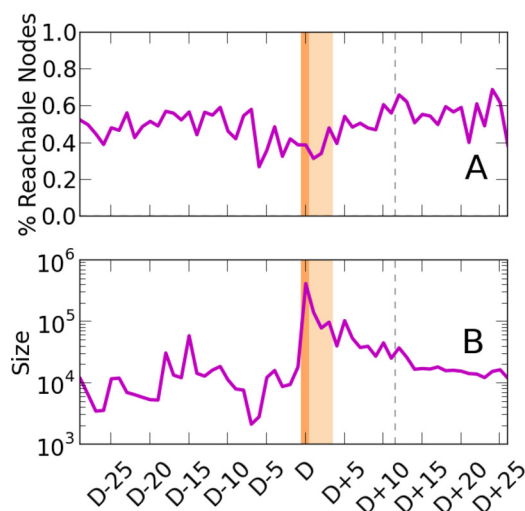
FIG. 7. Time evolution of the relative number of reachable nodes in comparison to the GC (a) and size of the reachable nodes' networks (b).

value in the 2-D probability density function. These contour lines are superimposed on a map of the municipalities composing the city of Caracas. There are five of them, bordered in green. The labels correspond to the municipality name, and the color indicates the ruling party-like the officialism in Libertador and the opposition in Chacao, Sucre, Baruta, and El Hatillo. Additionally, urbanized areas are colored in yellow and poorer regions (slums) in pink. Notice that the West region is characterized for having lower income and governed by the officialism, while the East part is wealthier and governed by the opposition.

It can be noticed that the regions where each pole concentrates most of their tweets are well separated from each other, showing that the city presents a clear geographical polarization. In fact, there is a good correspondence between the results of our model and offline evidence, such as electoral results or socioeconomic factors. Those municipalities governed by the opposition contain the highest concentration of users identified with this pole, and the same effect occurs for the officialism side of the political spectrum. We also have to remark that the areas with higher concentration of users aligned with the officialism correspond to the parts of the city with the largest concentration of poorer neighborhoods (pink areas). Conversely, the opposition users concentrate in urban developed regions. All these suggest that the basis of the Venezuelan popular polarization resides in socioeconomic factors and that the political conflict in Venezuela presents a strong territorial facet.

VI. CONCLUSIONS

Modern democracies have to represent the conflicts existing in our society, while at the same time maintain the social stability.[27] However, as polarization emerges, the few most powerful parties tend to capitalize the whole of the public attention and support, silencing the moderate opinions and under representing minorities. Consequently, todays' society is concerned about polarization, as a politically polarized society implies several risks. These risks include the appearance of radicalism or civil wars. In fact, one of the

actual challenges and a cutting edge topic is how to detect the emergence of political polarization and how to fix it.

We state that the possibility to gather user generated data from social media platforms,[28] together with network science,[46] represents an opportunity to detect political polarization. In this work, we have proposed a methodology to study and measure the emergence of polarization from social interactions. We have used it, to analyze the political polarization in one of the most polarized countries: Venezuela.[47,48] We have done this, by applying our methods to a Twitter conversation about the late Venezuelan president Hugo Chávez. We have shown that our methodology is able to detect different degrees of polarization in the conversation, depending on the participants' behavior, given by the structure of the network. Finally, we have contrasted our results against offline data, such as municipality governments or socioeconomic factors, finding a good correlation between the online and offline polarization. Hence, we conclude that online data seem to be a good proxy to detect politically polarized societies, as the online polarization that we found is a reflection of the Venezuelan political, territorial, and social polarization.

Another relevant question is: Can social media platforms help reduce political polarization as more voices could be heard? Although we do not answer this question, our results show that a minority of *elite* users were able to influence the whole online social network, resulting in a highly politically polarized conversation. However, these Venezuelan local influential accounts were not capable of polarizing the network when the conversation stopped being local of Venezuela and turned to be international. This opens two questions that can be studied from a social media analysis perspective: (i) How does online political polarization change at different scales-like city, country, continent or whole world? (ii) How could we target interventions in control strategies on social media that might be implemented to reduce polarization?

ACKNOWLEDGMENTS

APPENDIX A: DATASETS

In this work, we analyze messages from the online social network Twitter. We downloaded data from a temporal index of tweets managed by the Search API v1,[49] whose limitations are specified as the result of queries complexity and frequency, instead of fixing a percentage of the main stream. We queried for messages mentioning the name of the late Venezuelan President Hugo Chávez, during the events that surrounded his disease and death in 2013. We considered a two month period from February 4th, 2013 (29 days before the death announcement) to April 4th, 2013 (26 days after the death announcement). In summary, we downloaded 16 383 490 messages posted by 3 173 090 users from more than 159 countries (according to the 0.4% of geographically located messages). Our analyses are based on those messages

that represent retweets (49% of the downloaded content) and more specifically those that constitute the larger components of the communication networks, which were posted by 57% of original set of users.

The Venezuelan Internet penetration represents about 40% of the population, where most of users belong to middle and middle-low class.[50] Online social networks are very popular in this country. Around 33% of Venezuelans use Facebook[50] and almost 10% use Twitter.[52] In fact, Venezuela ranks thirteenth out of all countries in number of Twitter users.[52] Moreover, Venezuela has the highest proportion of mobile Internet in Latin America at over 30% of total connections, due to the popular use of social media from mobile phones.[51]

The political usage of Twitter in Venezuela is of great importance and has played a fundamental role in the recent Venezuelan history.[53,54] The late President Hugo Chávez was considered to be the second most influential world leader on Twitter,[55] preceded only by the US President Barack Obama. The collective who opposes the late President also finds on social media a channel to freely speak to their supporters and protest against the Government.[48]

APPENDIX B: NETWORKS

We have built one retweet network for each day of the observation period (56 networks). A retweet network emerges from user-to-user interactions during the message retransmission process provided by Twitter. Nodes represent users, and links are created between users i and j, when i forwards the content previously posted by j. Edges are weighted in proportion to the frequency that i retweeted j's messages, and directed in the sense of the flow of information from the message source j to the retweeter i.

A single network contains several retransmission cascades, seeded and propagated by the conversation participants. When these cascades are aggregated, several disconnected network components emerge. Among these components, there is a single one called Giant Component (GC) whose size is in the same order of the whole network. As part of the GC, there is a set of nodes that are reachable from the set of influential *elite* that represent about 50% of the GC's size (Fig. 7(a)). For most of days, the amount of reachable nodes fluctuated around 10 000 users and explosively grew to almost 500 000 users during day D (Fig. 7(b)). This behavior is typical of breaking news and critical events,[45,56] with a bursty increase during the main occurrence and a slow decay that may last for several days.

The retweet networks characterize the way that the collective attention is organized during an event on Twitter. The out strength (s_{out}) indicates the amount of retweets gained by a participant, while the in strength (s_{in}) indicates the number of retweets made by the participant. In Fig. 8, we have superimposed the out strength (top) and in strength (bottom) complementary cumulative density functions (CCDF) for each of the constructed networks, in log-log (left) and linear-log (right) scales. In both cases, the distributions display heterogeneous behavior, being the out strength distributions broader than the in strength distributions. In order to compare, whether
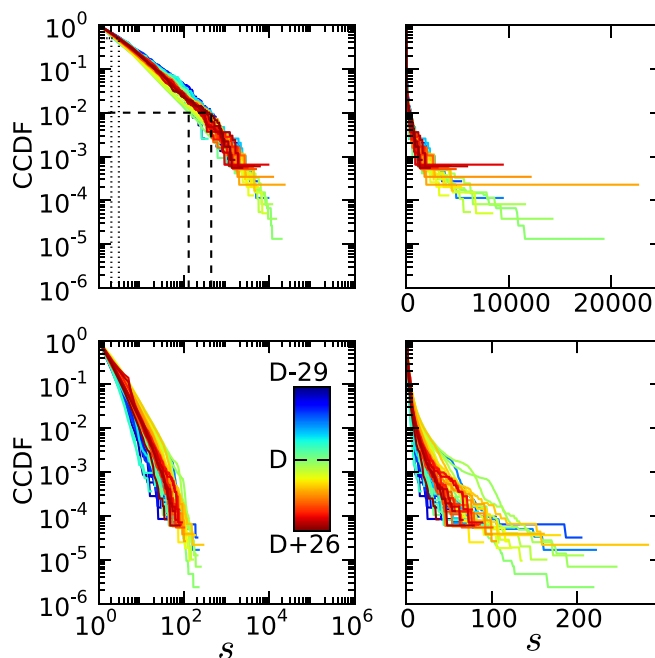

FIG. 8. Complementary cumulative density function (CCDF) of the retweet networks out strength s_{out} (top) and in strength s_{in} (bottom), from the Twitter conversation about the Venezuelan President Hugo Chávez, in log-log (left) and linear-log (right) scale. The colors indicate the corresponding day of the observation period. The dotted line indicates the range of s_{out} for 50% of the population, while the dashed lines indicate the range of s_{out} for 1% of the population.

these distributions behave like an exponential rather than a power law, we calculated the likelihood ratio statistical test.[57,58] We found that the probability of these distributions to follow an exponential curve, instead of a power law, has a p-value < 0.01 for more than 98% of the outgoing distributions and 75% of the incoming distributions, where over 87% of the distributions have a p-value < 0.05.

From a dynamical point of view, the power law distributions imply a preferential attachment mechanism,[46] where the chances of being retweeted increases with the number of retweets previously gained. These dynamics result in heterogeneous distributions where the great majority of users receive a very small amount of the collective attention, while some scarce users receive a disproportionally larger amount of it. For example, at all days, 50% of the population gained between 2 or 3 retweet at most (dotted lines in the top left panel of Fig. 8), while the 1% of most retweeted participants gained from 130 to 430 retweets as minimum (dashed lines in the top left panel of Fig. 8).

To further understand the relationship between the individual activity and the attention received, we will aggregate the observation period by characterizing the individuals according to their rate of participation and total amount of retweets gained. The participation rate is defined as

$$\rho = \rho_i/T, \tag{B1}$$

where ρ_i is the number of days that the user i actively participated in the retweet process and T is the total length of the observation period. The total number of retweets gained by user is measured as
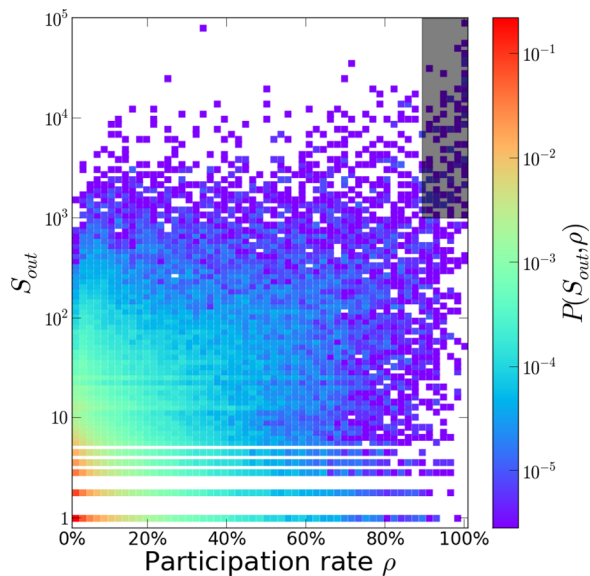
FIG. 9. Joint probability density function of the accumulated out strength (S_{out}) and the participation rate (ρ), from the Twitter conversation about the Venezuelan President Hugo Chávez. Colors correspond to the density of users. The grey reactangle marked at the top right corner indicates the *elite* users defined in Sec. IV.

$$S_{out} = \sum_{t=0}^{T} s_{out}(t), \tag{B2}$$

where $s_{out}(t)$ is the out strength of the node i at day t. If the user did not actively participate at day t, then $s_{out}(t) = 0$.

The joint probability density function of the accumulated out strength S_{out} and the participation rate ρ, $P(S_{out}, \rho)$, is shown in Fig. 9. This distribution indicates the total amount of attention received by users according to their participation rate. It can be noticed that the largest density of users (red and orange dots in Fig. 9) participated less than 20% ($\rho < 0.2$) of the days and present a small out strength value ($S_{out} < 10$), which means that most of them received a little amount of the collective attention. However, there is a very small set of users at the upper right corner in Fig. 9, who participated almost every day and present an extremely high S_{out}. This minority of highly influential users captured most of the collective attention throughout the observation period, and defines the *elite* users considered in Sec. IV.

[1] D. J. Isenberg, "Group polarization: A critical review and meta-analysis," J. Person. Soc. Psychol. **50**(6), 1141 (1986).

[2] C. R. Sunstein, "The law of group polarization," J. Polit. Phil. **10**(2), 175–195 (2002).

[3] D. Baldassarri and P. Bearman, "Dynamics of political polarization," Am. Sociol. Rev. **72**(5), 784–811 (2007).

[4] D. Baldassarri and A. Gelman, "Partisans without constraint: Political polarization and trends in American public opinion," Am. J. Sociol. **114**(2), 408–446 (2008).

[5] D. T. L. P. Dandekara and A. Goelb, "Biased assimilation, homophily, and the dynamics of polarization," Proc. Natl. Acad. Sci. USA **110**(15), 5791–5796 (2013).

[6] J. C. Turner, *Rediscovering the Social Group: A Self-categorization Theory* (B. Blackwell, 1987).

[7] A. Downs, "An economic theory of political action in a democracy," J. Polit. Econ. **65**(2), 135–150 (1957).

[8] R. Huckfeldt, "The social communication of political expertise," Am. J. Polit. Sci. **45**(2), 425–438 (2001).

[9] M. H. DeGroot, "Reaching a consensus," J. Am. Stat. Assoc. **69**(345), 118–121 (1974).

[10] A. A. Moreira, A. Mathur, D. Diermeier, and L. A. N. Amaral, "Efficient system-wide coordination in noisy environments," Proc. Natl. Acad. Sci. U.S.A **101**, 12085–12090 (2004).

[11] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," Dyn. Games Appl. **1**(1), 3–49 (2011).

[12] B. Golub and M. O. Jackson, "Naive learning in social networks and the wisdom of crowds," Am. Econ. J.: Microecon. **2**(1), 112–149 (2010).

[13] M. O. Jackson, *Social and Economic Networks* (Princeton University Press, 2010).

[14] D. Krackhardt, "A plunge into networks," Science **326**, 47–48 (2009).

[15] I. J. Benczik, S. Z. Benczik, B. Schmittmann, and R. K. P. Zia, "Lack of consensus in social systems," Europhys. Lett. **82**, 48006 (2008).

[16] D. Bindel, J. Kleinberg, and S. Oren, "How bad is forming your own opinion?," *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium* (IEEE, 2011), pp. 57–66.

[17] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar, "Opinion fluctuations and disagreement in social networks," Math. Oper. Res. **38**(1), 1–27 (2013).

[18] U. Krause, "A discrete nonlinear and non-autonomous model of consensus formation," Commun. Diff. Eqns **2000**, 227–236 (2000).

[19] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," J. Math. Sociol. **15**(3–4), 193–206 (1990).

[20] P. Lifford and A. Sudbury, "A model for spatial conflict," Biometrika **60**(3), 581–588 (1973).

[21] R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," Ann. Probab. **3**, 643–663 (1975).

[22] K. Suchecki, V. M. Eguíluz, and M. San Miguel, "Voter model dynamics in complex networks: Role of dimensionality, disorder, and degree distribution," Phys. Rev. E **72**(3), 036132 (2005).

[23] M. E. J. Newman, "The structure and function of complex networks," SIAM Rev. **45**(2), 167–256 (2003).

[24] M. Mobilia, A. Petersen, and S. Redner, "On the role of zealotry in the voter model," J. Stat. Mech. **8**, 08029 (2007).

[25] M. Mobilia, "Does a single zealot affect an infinite group of voters?," Phys. Rev. Lette. **91**(2), 028701 (2003).

[26] A. K. Dixit and J. W. Weibull, "Political polarization," Proc. Natl. Acad. Sci. U.S.A **104**, 7351–7356 (2007).

[27] L. J. Diamond, "Three paradoxes of democracy," J. Democracy **1**(3), 48–60 (1990).

[28] D. Lazer *et al.*, "Social science: Computational social science," Science **323**, 721–723 (2009).

[29] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 U.S. election: Divided they blog," Proceedings of LinkKDD, 2005.

[30] M. D. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, and F. Menczer, "Political Polarization on Twitter," *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*, pp. 89–96 (2011).

[31] J. Borondo, A. J. Morales, J. C. Losada, and R. M. Benito, "Characterizing and modeling an electoral campaign in the context of Twitter: 2011 spanish presidential election as a case study," Chaos **22**(2), 023138 (2012).

[32] J. Borondo, A. J. Morales, R. M. Benito, and J. C. Losada, "Multiple leaders on a multilayer social media,," Chaos, Solitons & Fractals **72**, 90–98 (2015).

[33] P. H. C. Guerra, W. Meira, Jr., C. Cardie, and R. Kleinberg, "A measure of polarization on social media networks based on community boundaries," Seventh International AAAI Conference on Weblogs and Social Media, 2013.

[34] R. S. Burt, "The social capital of opinion leaders," Ann. Am. Acad. Polit. Soc. Sci. **566**(1), 37–54 (1999).

[35] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts, "Who says what to whom on Twitter," Proceedings of the 20th international conference on World wide web, pp. 705–714, ACM, 2011.

[36] D. Boyd, S. Golder, and G. Lotan, "Tweet, tweet, retweet: Conversational aspects of retweeting on twitter," in *2010 43rd Hawaii International Conference on System Sciences (HICSS)*, Honolulu, HI, 5–8 Jan. 2010 (IEEE, 2010), pp. 1–10.

[37] S. Dann, "Twitter content classification," First Monday **15**(12), 2745 (2010).

[38] J. Kim and J. Yoo, "Role of sentiment in message propagation: Reply vs. retweet behavior in political communication," *2012 International Conference on Social Informatics*, pp. 131–136 (2012).

[39] D. Stewart, "When Retweets attack: are Twitter users liable for republishing the defamatory tweets of others?," J. Mass Commun. Quarter. **90**(2), 233–247 (2013).

[40] D. van Liere, "How far does a tweet travel?: Information brokers in the twitterverse," In Proceedings of the International Workshop on Modeling Social Media (MSM'10), ACM, New York, NY, USA, 2010.

[41] C.-T. Li, S.-D. Lin, and M.-K. Shan, "Exploiting endorsement information and social influence for item recommendation," In Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval (SIGIR'11), ACM, New York, NY, USA (ACM, 2011), pp. 1131–1132.

[42] Z. Liu and I. Weber, "Predicting ideological friends and foes in Twitter conflicts," In Proceedings of the companion publication of the 23rd international conference on World wide web companion (WWW Companion'14, 2014), pp. 575–576.

[43] M. E. Newman, "Mixing patterns in networks," Phys. Rev. E **67**(2), 026126 (2003).

[44] I. Borg and P. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, 2nd ed. (Springer-Verlag, New York, 2005), pp. 207–212.

[45] J. P. Bagrow, D. Wang, and A.-L. Barabási, PLOS ONE **6**, e17680 (2011).

[46] A. L. Barabási and J. Frangos, *Linked: The New Science Of Networks* (Basic Books, 2002).

[47] S. Ellner, *Venezuelan Politics in the Chávez Era: Class, Polarization and Conflict* (Lynne Rienner Publishers, 2002).

[48] A. Morales, J. Losada, and R. Benito, "Users structure and behavior on an online social network during a political protest," Physica A **391**(21), 5244–5253 (2012).

[49] See https://dev.twitter.com/docs/using-search for more information about the Twitter Search Application Programming Interface (API).

[50] Tendencias Digitales, "La penetración de Internet en Venezuela," 2012. http://tendenciasdigitales.com/1433/la-penetracion-de-internet-en-venezuela-alcanza-40-de-la-poblacion/.

[51] GSMA Intelligence, "Mobile Economy Latin America 2013," http://www.gsmamobileeconomylatinamerica.com/ENG_LatAmME_v10_WEB_FINAL.pdf.

[52] Semiocast, "Twitter reaches half a billion accounts more than 140 millions in the U.S.," 2012, see http://semiocast.com/en/publications/2012_07_30_Twitter_reaches_half_a_billion_accounts_140m_in_the_US.

[53] E. Minaya and K. Vyas, "When Chávez tweets, Venezuelans listen," The Wall Street Journal, April 25, 2012.

[54] F. T. Juan Nigel, "Facebook gives a platform to the challenger of Chávez," Foreign Policy, July 26, 2012.

[55] D. P. Council, "World leader rankings on twitter," Research Note of digitaldaya.com (2012), see http://www.digitaldaya.com/admin/modulos/galeria/pdfs/69/156_biqz7730.pdf.

[56] J. Yang and J. Leskovec, "Patterns of temporal variation in online media," in *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM'11 (New York, NY, USA)* (ACM, 2011), pp. 177–186.

[57] A. Clauset, C. R. Shalizi, and M. E. J. Newman, "Power-law distributions in empirical data," SIAM Review **51**, 661–703 (2009).

[58] J. Alstott, E. Bullmore, and D. Plenz, "Powerlaw: A python package for analysis of heavy-tailed distributions," PLoS ONE **9**(1), e85777 (2014).