

Public discourse and news consumption on online social media: A quantitative, cross-platform analysis of the Italian Referendum

Michela Del Vicario
 IMT School for Advanced Studies
 Lucca, Italy
 michela.delvicario@imtlucca.it

Sabrina Gaito
 University of Milan
 Dept. of Computer Science
 Milan, Italy
 sabrina.gaito@unimi.it

Walter Quattrociocchi
 IMT School for Advanced Studies
 Lucca, Italy
 walter.quattrociocchi@imtlucca.it

Matteo Zignani
 University of Milan
 Dept. of Computer Science
 Milan, Italy
 matteo.zignani@unimi.it

Fabiana Zollo
 Ca' Foscari University of Venice
 DAIS
 Venice, Italy
 fabiana.zollo@unive.it

ABSTRACT

The rising attention to the spreading of fake news and unsubstantiated rumors on online social media and the pivotal role played by confirmation bias led researchers to investigate different aspects of the phenomenon. Experimental evidence showed that confirmatory information gets accepted even if containing deliberately false claims while dissenting information is mainly ignored or might even increase group polarization. It seems reasonable that, to address misinformation problem properly, we have to understand the main determinants behind content consumption and the emergence of narratives on online social media. In this paper we address such a challenge by focusing on the discussion around the Italian Constitutional Referendum by conducting a quantitative, cross-platform analysis on both Facebook public pages and Twitter accounts. We observe the spontaneous emergence of well-separated communities on both platforms. Such a segregation is completely spontaneous, since no categorization of contents was performed a priori. By exploring the dynamics behind the discussion, we find that users tend to restrict their attention to a specific set of Facebook pages/Twitter accounts. Finally, taking advantage of automatic topic extraction and sentiment analysis techniques, we are able to identify the most controversial topics inside and across both platforms. We measure the distance between how a certain topic is presented in the posts/tweets and the related emotional response of users. Our results provide interesting insights for the understanding of the evolution of the core narratives behind different echo chambers and for the early detection of massive viral phenomena around false claims.

CCS CONCEPTS

•Human-centered computing → *Empirical studies in HCI*; •Applied computing → *Sociology*;

KEYWORDS

information spreading, news consumption, cross-platform comparison, online social media

ACM Reference format:

Michela Del Vicario, Sabrina Gaito, Walter Quattrociocchi, Matteo Zignani, and Fabiana Zollo. 2017. Public discourse and news consumption on online social media: A quantitative, cross-platform analysis of the Italian Referendum. In *Proceedings of KDD, Halifax, Nova Scotia - Canada, August 2017 (KDD'17)*, 9 pages.
 DOI: 10.475/123_4

1 INTRODUCTION

Social media have radically changed the paradigm of news consumption and their impact on the information spreading and diffusion has been largely addressed [4, 6, 10, 19, 23]. In particular, several studies focused on the prediction of social dynamics [4, 6], with a special emphasis on information flows patterns [24] and on the emergence of specific community structures [18, 27]. Indeed, social media are always more involved in both the distribution and consumption of news [26]. According to a recent report [25], approximately 63% of users access news directly from social media, and such information undergoes the same popularity dynamics as other forms of online contents, such as selfies or kitty pictures. In such a disintermediated environment, where users actively participate to contents production and information is no longer mediated by journalists or experts, communication strategies have changed, both in the way in which messages are framed and are shared across the social networks. The general public drifts into dealing with a huge amount of information, but the quality may be poor. Social media do represent a great tool to inform, engage and mobilize people easily and rapidly. However, it also has the power to misinform, manipulate or control public opinion. Since 2013, the World Economic Forum (WEF) has been listing massive digital misinformation at the core of technological and geopolitical risks to our society, along with terrorism and cyber attacks [20].

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

KDD'17, Halifax, Nova Scotia - Canada

© 2016 ACM. 123-4567-24-567/08/06...\$15.00

DOI: 10.475/123_4

Indeed, online disintermediation elicits users' tendency to select information that adheres (and reinforces) their pre-existing beliefs, the so-called confirmation bias. In such a way, users tend to form groups of like-minded people where they polarize their opinion, i.e., echo chambers [8, 13, 30, 38]. Confirmation bias plays a pivotal role in informational cascades. Experimental evidence shows that confirmatory information gets accepted even if containing deliberately false claims [7, 8], while dissenting information is mainly ignored [38]. Moreover, debating with like-minded people has been shown to influence users' emotions negatively and may even increase group polarization [36, 39]. Therefore, current approaches such as debunking efforts or algorithmic-driven solutions based on the reputation of the source seem to be ineffective to contrast misinformation spreading [29].

The rising attention to the spreading of fake news and unsubstantiated rumors on online social media led researchers to investigate different aspects of the phenomenon, from the characterization of conversation threads [5], to the detection of bursty topics on microblogging platforms [15], to the mechanics of diffusion across different topics [33]. Misinformation spreading also motivated major corporations such as Google and Facebook to provide solutions to the problem [22]. Users' polarization is a dominating aspect of online discussions [2, 37]. Furthermore, it has been shown that on Facebook users tend to confine their attention on a limited set of pages, thus determining a sharp community structure among news outlets [35], and a similar pattern was also observed around the Brexit debate [14]. Therefore, it seems reasonable that, to address misinformation problem properly, we have to understand the main determinants behind content consumption and the emergence of narratives on online social media. In this paper we address such a challenge by focusing on the discussion around the Italian Constitutional Referendum by conducting a quantitative, cross-platform analysis on both Facebook public pages and Twitter accounts. First, we characterize the structural properties of the discussion by modeling users' interactions by means of a bipartite network where nodes are Facebook pages (respectively, Twitter accounts) and connections among pages (respectively, accounts) are the direct result of users' activity. By comparing the results of different community detection algorithms, we are able to determine the existence of well-separated communities on both platforms. Notice that such a segregation is completely spontaneous, since we did not perform any categorization of contents a priori. Then, we explore the dynamics behind the discussion by looking at the activity of the users within the most active communities, finding that users tend to restrict their attention to a specific set of pages/accounts. Finally, taking advantage of automatic topic extraction and sentiment analysis techniques, we are able to identify the most controversial topics inside and across both platforms. We measure the distance between how a certain topic is presented in the posts/tweets and the related emotional response of users. Summarizing, the novel contribution of this paper is double:

- (1) First, we provide a comparative, quantitative analysis of the way in which news and information get consumed on two very different and popular platforms, Facebook and Twitter;

- (2) Second, we provide automatic tools to detect the popularity of narratives online and their related polarization effects.

Such an approach may provide important insights for the understanding of the evolution of the core narratives behind different echo chambers and for the early detection of massive viral phenomena around false claims.

2 DATASET

Facebook and Twitter are two of the most popular online social media where people can access and consume news and information. However, their nature is different: Twitter is an information network, while Facebook is still a social network, despite its evolution into a "personalized portal to the online world" ¹. Such a difference highlights the importance of studying, analyzing, and comparing users' consumption patterns inside and between both platforms.

2.1 Facebook

Following the exhaustive list provided by *Accertamenti Diffusione Stampa* (ADS) [3], we identified a set of 57 Italian news sources and their respective Facebook pages. For each page, we downloaded all the posts from July 31st to December 12th, 2016, as well as all the related likes and comments. Then we filtered out posts about the Italian Constitutional Referendum (hold on December, 4th) by keeping those containing at least two words in the set *{Referendum, Riforma, Costituzionale}* in their textual content i.e., their description on Facebook or the URL content they pointed to. The exact breakdown of the dataset is provided in Table 1. Data collection was carried out using the Facebook Graph API [16], which is publicly available. For the analysis (according to the specification settings of the API) we only used publicly available data (thus users with privacy restrictions are not included in the dataset). The pages from which we downloaded data are public Facebook entities and can be accessed by anyone. Users' content contributing to such pages is also public unless users' privacy settings specify otherwise, and in that case it is not available to us.

Table 1: Breakdown of Facebook and Twitter datasets.

	Facebook	Twitter
Pages/Accounts	57	50
Posts/Tweets	6015	5483
Likes/Favorites	2034037	57567
Retweets	-	55699
Comments	719480	30079
Users	671875	29743
Likers/Favoriters	568987	16438
Retweeters	-	14189
Commenters	200459	8099

2.2 Twitter

Data collection on Twitter followed the same pipeline adopted for Facebook. First, we identified the official accounts of all the news

¹http://www.slate.com/articles/technology/technology/2016/04/facebook_isn_t_the_social_network_anymore_so_what_is_it.html

sources reported in [3]. Then, we gathered all the tweets posted by such accounts from July 31st to December 12th, 2016 through the Twitter Search API. Starting from that set, we selected only tweets pointing to news related to the Referendum debate. Specifically, we considered only the statutes² whose URL was present in the Facebook dataset. The resulting set of valid tweets consists of about 5,400 elements. To make the available information comparable to Facebook, from each tweet we collected information about the users who "favored" (left a favorite), retweeted or replied. Specifically, favorites and retweets express an interest in the content as Facebook likes do, while replies are similar to comments. To obtain such information we bypassed the Twitter API, since it does not return favorites or retweets. In particular, we scraped the tweet page and extracted the users who liked or retweeted the status, and the retweet/favorite counts. Although we are able to identify the users acting on a tweet, the information is partially incomplete since the tweet page shows at most 25 users who retweet or favor (retweeters or "favoriters"). However, such a restriction has a limited impact on the set of valid tweets. Indeed, our retweeters' and favoriters' sets capture the entire set of users acting on tweets for about 80% of statuses. As for the replies to a tweet, we were able to collect every user who commented on a tweet reporting a news about the Referendum and the reply tweet she wrote. The tweet page reports the entire discussion about the target tweet, including the replies to replies and so on. Here we limited our collection to the first level replies –i.e., direct replies to the news source tweet– since the target of the comment is identifiable, i.e. the news linked to the tweet³.

The breakdown of the Twitter dataset is reported in Table 1. The ratio between Facebook and Twitter volumes mirrors the current social media usage in the Italian online media scene. Indeed, news sources and newspapers exploit both media channels to spread their contents, as denoted by the similar number of posts. Nevertheless, activities on the posts are skewed towards Facebook, since in Italy its active users are 4/5-fold those of Twitter⁴. Our collection methodology and the high comparability between the two datasets provide a unique opportunity to investigate news consumption inside and between two different and important social media.

3 METHODS

Interactions between Facebook pages/Twitter accounts and users can be modeled by a bipartite network $G = (P, U, E)$, where P and U are disjoint sets of vertices representing news sources accounts and users, respectively. When a user $u \in U$ interacts with a post/tweet published by the news source $p \in P$, we draw a link $(u, p) \in E$. To each link we assign a weight w_p^u denoting how many times u has interacted with the posts published by the page p . Since we distinguish between comments/replies and likes/favorites-retweets, we have different types of weight corresponding to different kinds of interaction.

Starting from the bipartite network representation, we are then able to extract groups of news sources which are perceived as similar by social media users.

²Status and tweet are synonyms.

³For deeper levels the target could be the parent reply or the target tweet.

⁴<http://vincos.it/2016/04/01/social-media-in-italia-analisi-dei-flussi-di-utilizzo-del-2015/>

3.1 Bipartite Network Projections

When limiting to a specific topic –i.e., the constitutional Referendum– a single user interacting with different news sources can be interpreted as a signal of closeness or similarity between the pages, especially when interactions are likes, favorites or retweets. To this aim, we apply three different projections of the bipartite network of pages and users on the pages set, where in turn we assign a different weight to the inter-pages links, assessing diverse similarity scores:

- **Simple Edge Counting:** The weight of each link is given by the number of common users between two pages, in terms of either likes (Facebook) or favorites and retweets (Twitter);
- **Jaccard Similarity:** Given two pages a and b , and their respective sets of users A and B , the weight of their link is given by the Jaccard similarity coefficient of the two pages $J_{ab} = |A \cap B| / |A \cup B|$;
- **User Activity Jaccard Similarity :** We also consider a modification of the Jaccard Similarity by taking into account the weight of the page-user links. Let U_{ab} be the set of common users between pages a and b , whose cardinality is N_{ab} , U the set of all users, and P that of all pages, the page relative users weight w_{ab} is equal to:

$$w_{ab} = \frac{1}{N_{ab}} \sum_{u \in U_{ab}} \left(1 - \frac{|w_a^u - w_b^u|}{\max_{c \in P} w_c^u} \right) \frac{|w_a^u + w_b^u|}{2 \max_{v \in U, c \in P} w_c^v}.$$

The User Activity Jaccard Similarity coefficient is given by the product $w_{ab} J_{ab}$.

3.2 Community detection algorithms

Given the above projections, we are able to identify groups of newspapers which are perceived similar by the social media users. The identification of the groups turns into the well-known community detection problem. For this task, we choose algorithms representative of different approaches on community detection, ranging from optimizing the modularity function, to dynamical processes such as diffusion or random walks.

- **Fast Greedy [11]:** The algorithm is based on the maximization of the Newman and Girvan's modularity, a quality function which measures the goodness of a graph clustering w.r.t. its random counterpart with the same expected degree sequence. The algorithm starts from a set of isolated nodes and, at each step, continues to add links from the original graph to produce the largest modularity increase.
- **Louvain [9]:** The Louvain method was introduced by Blondel *et al.* and is one of the most popular greedy algorithms for *modularity* optimization. The method is very fast and can produce very high quality communities. The algorithm is iterative. Each iteration consists of two steps. In the first one, every node is initially set to a new community. Then, for every node i and its neighbors j , the algorithm calculates the gain in modularity moving i from its community to j 's community. The node i is then moved to the community with maximum gain. The second step is to group all the nodes in the same community into a macro node. A new graph, in which macro nodes are linked by

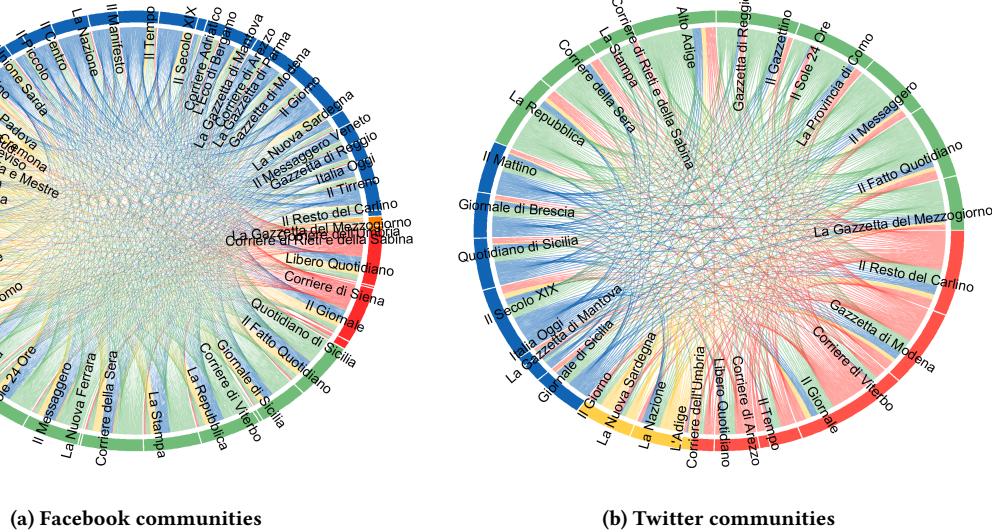


Figure 1: a) Community structure for the User Activity Jaccard Similarity projection of the pages-users graph on Facebook and b) on Twitter. Colors indicate the membership of Facebook pages or Twitter accounts in the different communities (green for C1, red for C2, blue for C3, yellow for C4, and orange for C5) detected by the Louvain algorithm.

an edge if there's an edge between two nodes belonging to the two macro nodes, is built and a new iteration starts.

- **Label propagation** [31]: The algorithm provides each node $v \in V$ to determine its community by choosing the most frequent label shared by its neighbors. Initially every node belongs to a different community. After some iterations, groups of nodes quickly reach a consensus on their label and they begin to contend those nodes that lay between groups.
- **Walk Trap** [28]: The algorithm is based on the tendency of random walks to be trapped into the densely connected regions of the graph, i.e. communities. Specifically, a random walk of length t induces a distance between two vertices depending on the probability the random walker has to go from a vertex to the other in t steps. The distance can be used in a hierarchical clustering algorithm to obtain a hierarchical community structure.
- **Infomap** [34]: Infomap turns the community detection problem into the problem of optimally compressing the information of a random walk taking place on the graph. The algorithm achieves the optimal compression by optimizing the minimum description length of the random walk.

To compare the partitions returned by the different algorithms, we use standard techniques that compute the similarity between different clustering methods by considering how nodes are assigned by each community detection algorithm [21, 32]. Specifically, we adopted the Rand Index [32], based on the relationships between pairs of vertices in both partitions, and the Normalized Mutual Information (NMI) [12], measuring the statistical dependence between the partitions. In both cases a value close to 1 indicates that the partitions are almost identical, whereas independent partitions lead the scores to 0.

4 RESULTS

4.1 Communities and Users Polarization

Online social media facilitate the aggregation of individuals in communities of interest, also known as echo chambers [30, 38], especially when considering users interaction around conflicting and contrasting narratives [13]. In this manuscript we do not perform any categorization of contents *a priori*. On the contrary, we only account for the connections created by users activities, and then observe the resulting, emerging communities.

We compare the results of the different community detection algorithms on all three bipartite projections, for the case of both Facebook and Twitter. Our analysis focuses on the the bipartite networks built around likes and favorites/retweets, since a) these activities represent more explicit positive feedbacks than comments or replies; b) the average Rand Index is greater for all the projections in both social media w.r.t. the bipartite networks built on comments; and c) in Twitter we lose fewer news sources accounts by applying the projections on the favorite bipartite network w.r.t. the comment one.

To compare each projection of the like/favorite bipartite networks we compute the average Rand Index and NMI among the partitions returned by the algorithms. For the case of both Facebook and Twitter the User Activity Jaccard Similarity obtains a better degree of concordance among the algorithms w.r.t. the simple edge counting projection. For instance, in Twitter the User Activity Jaccard Similarity gets 0.73 and 0.69 as average Rand index and NMI, while the simple edge counting gets 0.49 and 0.28, indicating a low concordance among the algorithms. Jaccard similarity performs even better (Rand = 0.8 and NMI=0.7), however we choose the User Activity Jaccard Similarity since it accounts for the users' activity frequency and how they allocate their interactions among

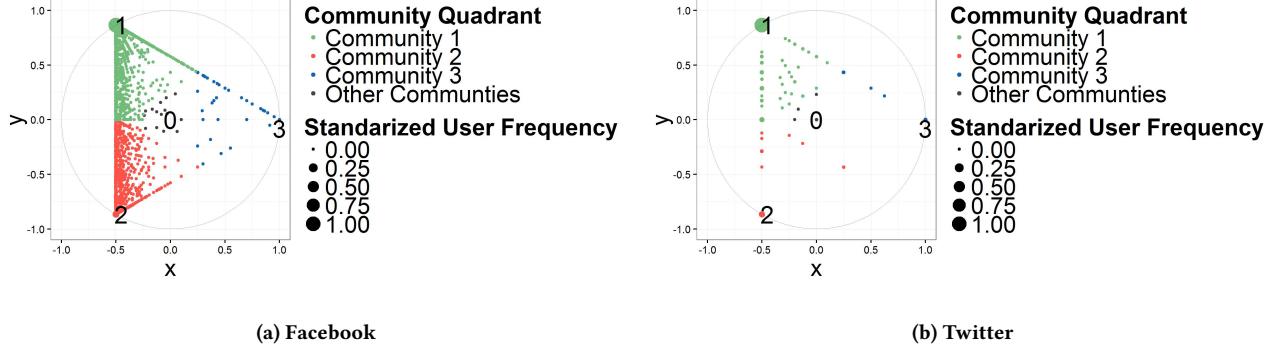


Figure 2: Users Polarization: Users activity across the three most active communities of Facebook a) and Twitter b). In the graph, the vertices of the triangle represent the three most active communities and the central point all the remaining ones. The position of each dot is determined by the number of communities the users interact with. The size indicates the number of users in that position. Hence, data denotes a very characteristic behavior: users are strongly polarized and tend to focus their attention on a single community of pages.

the pages they like. Similar results hold for the Facebook bipartite network, i.e., the User Activity Jaccard Similarity obtains 0.68 and 0.66 as Rand index and NMI, respectively. Finally, having fixed the bipartite network and its projection, we evaluate the concordance of each algorithm with the other ones. In Facebook the Louvain algorithm obtains the best average accordance (average Rand=0.76 and NMI=0.65), while in Twitter the algorithms based on dynamical processes get the best scores. The latter result is due to a high intra-approaches concordance and low inter-approaches scores, so that the biggest class drives the final score. To maintain a common setting between Facebook and Twitter, in the remaining of the paper we will consider and compare the communities extracted by the Louvain algorithm from the likes/favorites bipartite networks projected by the User Activity Jaccard Similarity .

In Facebook we identify five communities of news pages, while in Twitter the communities are four. In Figure 1 we show the structure of Facebook and Twitter networks where vertices are grouped according to their community membership, while a complete list of the pages/accounts and their membership is reported in Table 2. For the case of both Facebook and Twitter, pages are not equally distributed among the communities, but we observe a main community (C_1) including about 40% of vertices and two smaller communities (C_2 and C_3), each one formed by 20% of accounts. Activities on pages/accounts are even more skewed; the main community C_1 accounts for most of the activities (80%) in both Facebook and Twitter.

In order to characterize the relationship between the observed spontaneous emerging community structure of Facebook pages/Twitter accounts and users behavior, we quantify the fraction of the activity of any user in the three most active communities versus that in any other community. Figure 2 shows the activity of users across the three most active communities emerging from either Facebook pages (Fig. 2a) or Twitter accounts (Fig. 2b). In both cases, we find that users are strongly polarized and that their attention is confined to a single community of pages. Both dynamics are similar, however we observe a substantial differentiation in the

activity volume, with a smaller and even more polarized activity for the Twitter case.

4.2 News Presentation

We observed that users activity is mainly restricted to one echo chamber and that users show a limited interaction with other news sources. We are now interested in learning the process that drives such a pattern, which is observed in both Facebook and Twitter. We do that by measuring the distance between the sentiment with which the news is presented to the audience and the sentiment expressed by the users w.r.t. the same topic. To perform the analysis we make use of IBM Watson™ AlchemyLanguage service API [1], which allows to extract semantic meta-data from posts content. Such a procedure applies machine learning and natural language processing techniques aimed to analyze text by automatically extracting relevant entities, their semantic relationships, as well as the emotional sentiment they express [17]. In particular, we extract the sentiment and the main entities presented by each post of our datasets, whether it has a textual description or a link to an external document. Entities are represented by items such as persons, places, and organizations that are present in the input text. Entity extraction adds semantic knowledge to content to help understand the subject and context of the text that is being analyzed.

As a first step towards the identification and analysis of controversial topics, we look at how news are presented. Figure 3 shows the Probability Density Function of the posts sentiment score on the three most active communities for both Facebook (Figure 3a) and Twitter (Figure 3b). The sentiment score is defined in the range $[-1, 1]$, where -1 is negative, 0 is neutral, and 1 is positive. A neutral overall pattern is observed on both platforms meaning that topics are mainly presented in a neutral manner, although exhibiting a slightly higher probability of positive sentiment on the two largest communities (solid green and dashed red lines) with respect to the smaller one. Such a behavior is observed on both Facebook and Twitter. Notice that we consider how subjects are presented in a post; here we do not take into account the sentiment that the post

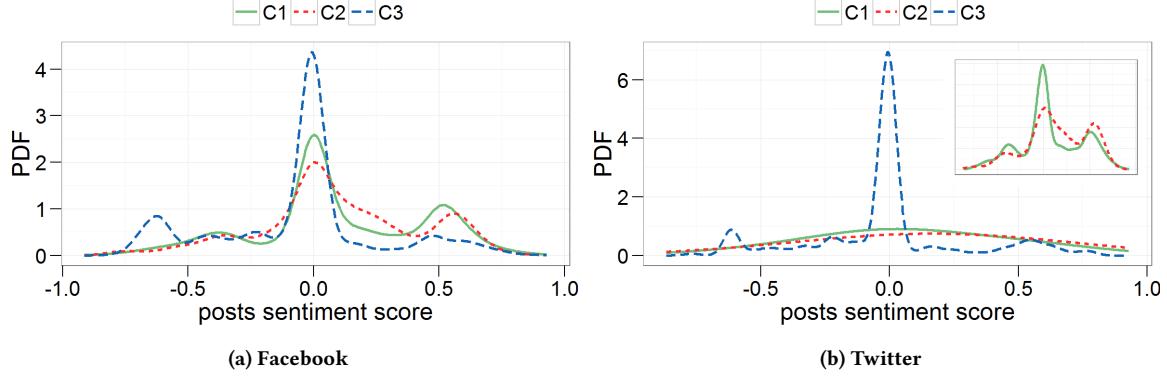


Figure 3: Probability Density Function (PDF) of posts sentiment score on C1 (solid green), C2 (dashed red), and C3 (dashed blue) of Facebook a) and on Twitter b). The sentiment score is defined in the range $[-1, 1]$, where -1 is negative, 0 is neutral, and 1 is positive.

may elicit in the reader, or the sentiment of users involved in the discussion.

We now focus on the entities shared across communities: on Facebook we identify 20 such entities appearing in posts from all the three most active communities, while on Twitter they are 21. For each entity we compute its average sentiment with respect to every community, i.e., we get three values representing the mean sentiment of the posts containing that entity in either C1, C2, or C2. The mean emotional distance of such entities is then the mean of the differences between the average sentiment pairwise computed. Figure 4a shows, for each entity, the average sentiment in each community (green dots for C1, red for C2, and blue for C3) and the mean emotional distance (yellow diamonds) among communities on Facebook, while Figure 4b reports the same information for Twitter. In both panels a) and b) entities are shown in descending order with respect to their mean emotional distance, with those on the left being the ones discussed with the greatest difference in sentiment, while those on the right the ones discussed in a much more similar way across the three echo chambers.

Figure 4c shows for each entity, the emotional distance between Facebook communities aggregated (green dots) and Twitter ones (red dots). Entities on the left are presented in a more positive way on Facebook, and vice versa. We consider only entities appearing on both platforms and at least in two communities for each platform.

4.3 Users Response to Controversial Topics

We have characterized how subjects are debated inside and across communities and we identified the emerging controversial topics. We are now interested in the attitude of users towards such topics. We then analyze the emotional response of users by looking at the sentiment they express when commenting. Specifically, we focus on comments left on posts containing entities shared by at least two posts and we then compute their sentiment score through AlchemyAPI. Thus, to each comment is associated a sentiment score in $[-1, 1]$, characterized as before, and for each post (respectively, user) we compute the average sentiment of its (respectively, her

comments i.e., the mean of the sentiment of all comments of the post (respectively, user). Then, for each entity, we consider the emotional distance between the average sentiment of the post and that of its users. Since we are interested in identifying the most controversial entities, we consider only those for which the emotional distance (in absolute value) between the post and user is greater than 0.2. Figure 5 shows the emotional response of users to posts debating one of the listed controversial topics, on both Facebook and Twitter. Figures 5a, 5c and 5e (respectively, 5b, 5d and 5f) refer to C1, C2, and C3 in Facebook (respectively, Twitter). In all panels a vertical dashed line denotes a change in users response: entities on the left are those for which users response is more negative than the sentiment expressed in the post, and vice versa for those on the right. We may notice that users tend to react to the content of the posts in C1 and C2 slightly more negatively on Facebook rather than Twitter, while the opposite is observed for C3.

5 CONCLUSIONS

Social media have radically changed the paradigm of news consumption and the rising attention to the spreading of fake news and unsubstantiated rumors on online social media led researchers to investigate different aspects of the phenomenon. In this paper we aim to understand the main determinants behind content consumption on social media by focusing on the online debate around the Italian Constitutional Referendum. Throughout a quantitative, cross-platform analysis on both Facebook public pages and Twitter accounts, we characterize the structural properties of the discussion, observing the emergence of well-separated communities on both platforms. Such a segregation is completely spontaneous; indeed, we find that users tend to restrict their attention to a specific set of pages/accounts. Finally, taking advantage of automatic topic extraction and sentiment analysis techniques, we are able to identify the most controversial topics inside and across both platforms, and to measure the distance between how a certain topic is presented in the posts/tweets and the related emotional response of

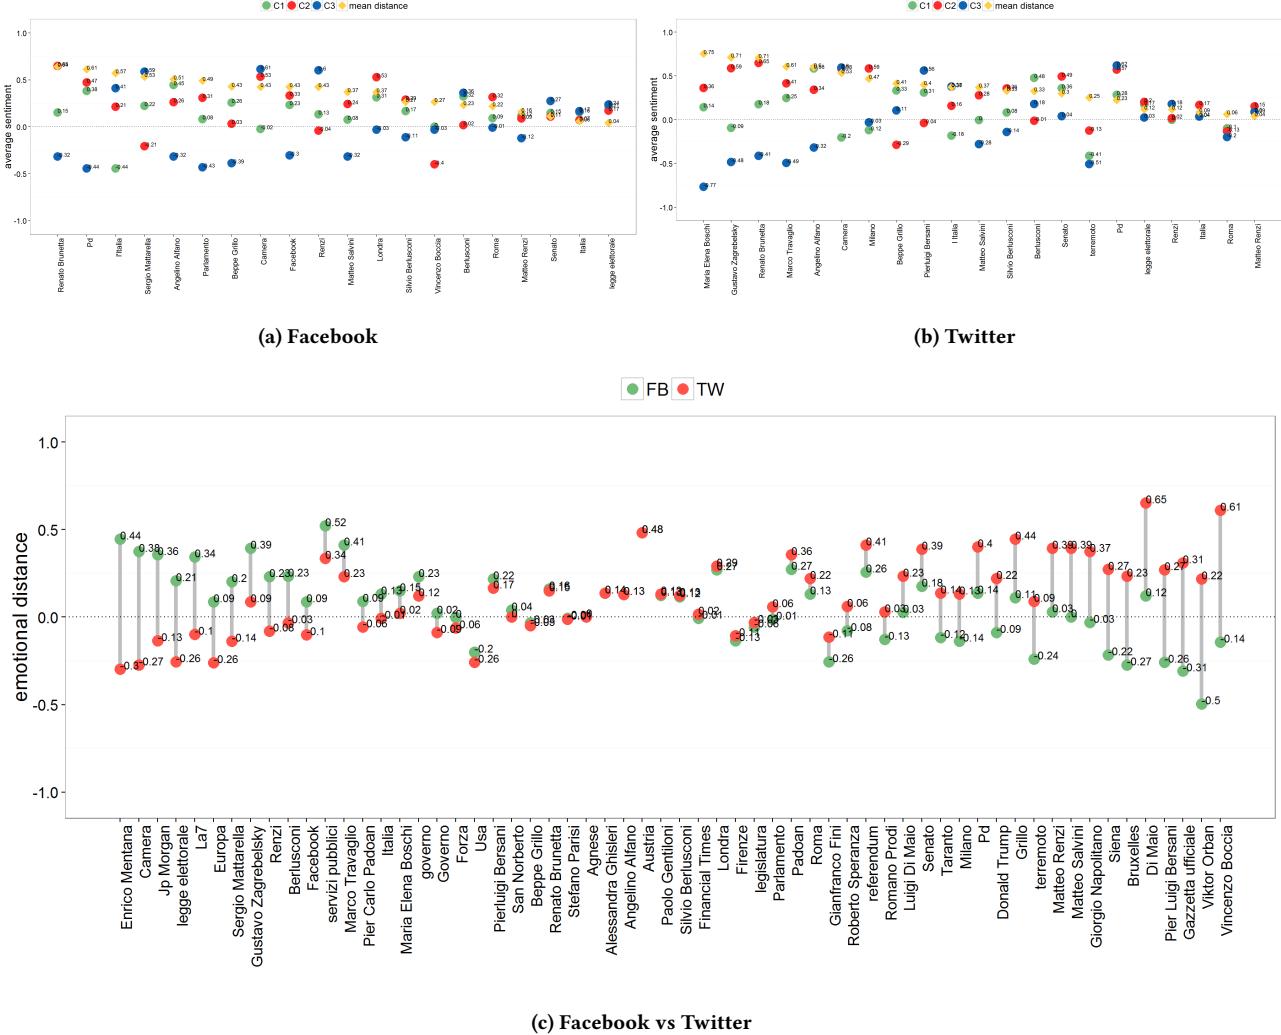


Figure 4: Emotional Distance Among Communities. Average sentiment of an entity on each community (green dots for C1, red for C2, and blue for C3) on Facebook a) and Twitter b). Yellow diamonds indicates the mean emotional distance between pairs of communities, i.e., the mean of the differences between the average sentiment of an entity on each pair of communities, for each entity debated in all three communities. Entities are shown in a descending order by the largest to the smallest mean emotional distance. c) Emotional distance between Facebook communities aggregated (green dots) and Twitter one (red dots).

users. Our novel approach may provide interesting insights for a) the understanding of the evolution of the core narratives behind different echo chambers and b) the early detection of massive viral phenomena around false claims on online social media.

REFERENCES

- [1] 2017. Alchemy API — Powering the New AI Economy. An IBM Company. (February 2017). <http://www.alchemyapi.com/>
- [2] Lada A Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*. ACM, 36–43.
- [3] ADS. 2016. Elenchi Testate. Website. (2016). http://www.adsnotizie.it/_testate.asp
- [4] Sitaram Asur and Bernardo A Huberman. 2010. Predicting the future with social media. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, Vol. 1. IEEE, 492–499.
- [5] Lars Backstrom, Jon Kleinberg, Lillian Lee, and Cristian Danescu-Niculescu-Mizil. 2013. Characterizing And Curating Conversation Threads: Expansion, Focus, Volume, Re-Entry. In *Proceedings of WSDM*. 13–22.
- [6] Hila Becker, Mor Naaman, and Luis Gravano. 2010. Learning similarity metrics for event identification in social media. In *Proceedings of the third ACM international conference on Web search and data mining*. ACM, 291–300.
- [7] Alessandro Bessi, Mauro Coletti, George Alexandru Davidescu, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2015. Science vs Conspiracy: collective narratives in the age of misinformation. *PloS one* 10, 2 (2015), 02.
- [8] Alessandro Bessi, Fabio Petroni, Michela Del Vicario, Fabiana Zollo, Aris Anagnostopoulos, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2015. Viral Misinformation: The Role of Homophily and Polarization. In *Proceedings of the 24th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee, 355–356.
- [9] Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008), P10008.

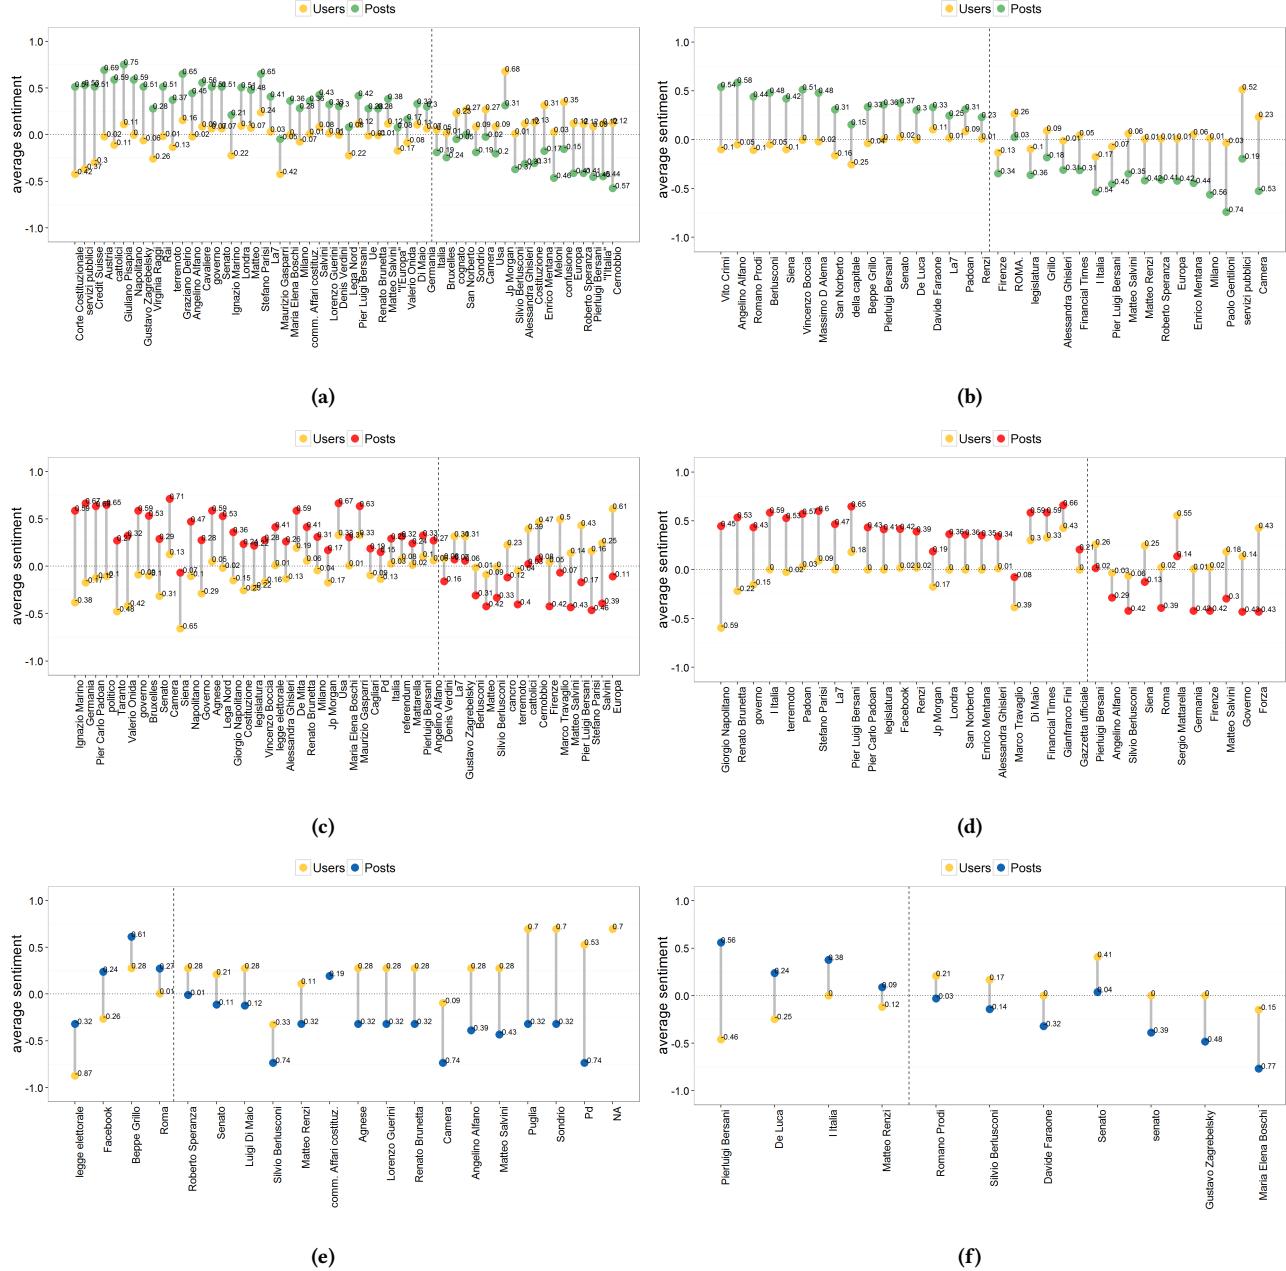


Figure 5: Emotional Response to Controversial Entities. Left panels refer to Facebook, right ones to Twitter. In the specific, panels (a-b) show the emotional response of users (yellow dots) to posts of C1 (green dots) debating one of the listed controversial entities, panels (c-d) show the emotional response of users (yellow dots) to posts of C2 (red dots), and panels (e-f) show the emotional response of users (yellow dots) to posts of C3 (blue dots). Only entities for which the emotional distance between the posts and users sentiment is greater than 0.2 are reported. The vertical dashed lines denote a change in users response.

- [10] Carlos Castillo, Mohammed El-Haddad, Jürgen Pfeffer, and Matt Stempeck. 2014. Characterizing the life cycle of online news stories using social media reactions. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 211–223.

[11] Aaron Clauset, M. E. J. Newman, and Cristopher Moore. 2004. Finding community structure in very large networks. *Phys. Rev. E* 70 (Dec 2004), 066111. Issue 6. DOI: <http://dx.doi.org/10.1103/PhysRevE.70.066111>

[12] Leon Danon, Albert Daz-Guilera, Jordi Duch, and Alex Arenas. 2005. Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment* 2005, 09 (2005), P09008. <http://stacks.iop.org/1742-5468/2005/i=09/a=P09008>

[13] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the National Academy of Sciences* 113, 20 (May 2016), 5649–5654. DOI: <https://doi.org/10.1073/pnas.1523036113>

Table 2: Facebook and Twitter News Sources with Community Membership. News sources that are not members of any community are denoted by N while those that are not available for Twitter are denoted by a dash.

ID	Page Name	Facebook Community	Twitter Community
1	Corriere della Sera	C1	C1
2	Il Fatto Quotidiano	C1	C1
3	Il Sole 24 Ore	C1	C1
4	La Repubblica	C1	C1
5	La Stampa	C1	C1
6	Corriere di Viterbo	C1	C2
7	La Nuova Ferrara	C1	-
8	La Provincia di Como	C2	C1
9	Il Giornale	C2	C2
10	Liber Quotidiano	C2	C2
11	Quotidiano di Sicilia	C2	C3
12	Corriere di Siena	C2	-
13	Corriere Adriatico	C3	N
14	Il Centro	C3	N
15	Il Piccolo	C3	N
16	Il Tirreno	C3	N
17	L'Eco di Bergamo	C3	N
18	La Gazzetta di Parma	C3	N
19	Il Messaggero	C3	C1
20	Corriere di Arezzo	C3	C2
21	Gazzetta di Modena	C3	C2
22	Il Resto del Carlino	C3	C2
23	Il Tempo	C3	C2
24	Giornale di Sicilia	C3	C3
25	Il Mattino	C3	C3
26	Il Secolo XIX	C3	C3
27	La Nazione	C3	C4
28	Il Manifesto	C3	-
29	La Sicilia	C3	N
30	Il Messaggero Veneto	C4	N
31	L'Unione Sarda	C4	N
32	La Provincia di Cremona	C4	N
33	Gazzetta di Reggio	C4	C1
34	La Gazzetta del Mezzogiorno	C4	C1
35	Giornale di Brescia	C4	C3
36	Italia Oggi	C4	C3
37	La Gazzetta di Mantova	C4	C3
38	Il Giorno	C4	C4
39	La Nuova Sardegna	C4	C4
40	Gazzetta del Sud	C4	-
41	Il Giornale di Vicenza	C5	N
42	Il Mattino di Padova	C5	N
43	La Nuova di Venezia e Mestre	C5	N
44	Trentino	C5	N
45	Alto Adige	C5	C1
46	Il Gazzettino	C5	C1
47	L'Adige	C5	C4
48	Corriere delle Alpi	C5	N
49	L'Arena	C5	-
50	La Provincia Pavese	C5	-
51	La Tribuna di Treviso	C5	N
52	Corriere di Rieti e della Sabina	C6	C1
53	Corriere dell'Umbria	C6	C2
54	La Provincia di Lecco	N	N
55	La Provincia di Sondrio	N	N
56	Quotidiano di Puglia	N	N
57	Avenire	N	-

- Sciences* 113, 3 (2016), 554–559.
- [14] Michela Del Vicario, Fabiana Zollo, Guido Caldarelli, Antonio Scala, and Walter Quattrociocchi. 2016. The Anatomy of Brexit Debate on Facebook. *arXiv preprint arXiv:1610.06809* (2016).
- [15] Qiming Diao, Jing Jiang, Feida Zhu, and Ee-Peng Lim. 2012. Finding bursty topics from microblogs. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. Association for Computational Linguistics, 536–544.
- [16] Facebook. 2013. Using the Graph API. Website. (8 2013). <https://developers.facebook.com/docs/graph-api/using-graph-api/> last checked: 19.01.2014.
- [17] Aldo Gangemi. 2013. A comparison of knowledge extraction tools for the semantic web. In *Extended Semantic Web Conference*. Springer, 351–366.
- [18] Przemyslaw A Grabowicz, José J Ramasco, Esteban Moro, Josep M Pujol, and Victor M Eguiluz. 2012. Social features of online networks: The strength of intermediary ties in online social media. *PLoS one* 7, 1 (2012), e29358.
- [19] Alfred Hermida. 2010. Twittering the news: The emergence of ambient journalism. *Journalism practice* 4, 3 (2010), 297–308.
- [20] W. Lee Howell. 2013. *Digital Wildfires in a Hyperconnected World*. Technical Report Global Risks 2013. World Economic Forum.
- [21] Lawrence Hubert and Phipps Arabie. 1985. Comparing partitions. *Journal of classification* 2, 1 (1985), 193–218.
- [22] Managing director: Jenni Sargent. 2017. First Draft Coalition. Website. (2017). <https://firstdraftnews.com>
- [23] Alice Ju, Sun Ho Jeong, and Hsiang Iris Chyi. 2014. Will social media save newspapers? Examining the effectiveness of Facebook and Twitter as news platforms. *Journalism Practice* 8, 1 (2014), 1–17.
- [24] Jure Leskovec. 2011. Social media analytics: tracking, modeling and predicting the flow of information through networks. In *Proceedings of the 20th international conference companion on World wide web*. ACM, 277–278.
- [25] Nic Newman, David AL Levy, and Rasmus Kleis Nielsen. 2015. Reuters Institute Digital News Report 2015. Available at SSRN 2619576 (2015).
- [26] Alexandra Olteanu, Carlos Castillo, Nicholas Diakopoulos, and Karl Aberer. 2015. Comparing Events Coverage in Online News and Social Media: The Case of Climate Change. *ICWSM* 15 (2015), 288–297.
- [27] Symeon Papadopoulos, Yiannis Kompatseris, Athena Vakali, and Ploutarchos Spyridonos. 2012. Community detection in social media. *Data Mining and Knowledge Discovery* 24, 3 (2012), 515–554.
- [28] Pascal Poma and Matthieu Latapy. 2006. Computing communities in large networks using random walks. *Journal of Graph Algorithms and Applications* 10, 2 (2006), 191–218.
- [29] Vahed Qazvinian, Emily Rosengren, Dragomir R Radev, and Qiaozhu Mei. 2011. Rumor has it: Identifying misinformation in microblogs. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1589–1599.
- [30] Walter Quattrociocchi, Antonio Scala, and Cass R Sunstein. 2016. Echo Chambers on Facebook. Available at SSRN (2016).
- [31] Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. 2007. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E* 76 (Sep 2007), 036106. Issue 3. DOI : <http://dx.doi.org/10.1103/PhysRevE.76.036106>
- [32] William M Rand. 1971. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association* 66, 336 (1971), 846–850.
- [33] Daniel M Romero, Brendan Meeder, and Jon Kleinberg. 2011. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th international conference on World wide web*. ACM, 695–704.
- [34] Martin Rosvall, Daniel Axelsson, and Carl T Bergstrom. 2009. The map equation. *The European Physical Journal Special Topics* 178, 1 (2009), 13–23.
- [35] Ana Lucia Schmidt, Fabiana Zollo, Michela Del Vicario, Alessandro Bessi, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. 2017. Anatomy of News Consumption on Facebook. (*to appear in*) *Proceedings of the National Academy of Sciences* (2017).
- [36] Cass R Sunstein. 2002. The law of group polarization. *Journal of political philosophy* 10, 2 (2002), 175–195.
- [37] Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. 2012. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences* 109, 16 (2012), 5962–5966.
- [38] Fabiana Zollo, Alessandro Bessi, Michela Del Vicario, Antonio Scala, Guido Caldarelli, Louis Shekhtman, Shlomo Havlin, and Walter Quattrociocchi. 2015. Debunking in a World of Tribes. *arXiv preprint arXiv:1510.04267* (2015).
- [39] Fabiana Zollo, Petra Kralj Novak, Michela Del Vicario, Alessandro Bessi, Igor Mozetič, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2015. Emotional dynamics in the age of misinformation. *PloS one* 10, 9 (2015), e0138740.