

OPEN

Topology comparison of Twitter diffusion networks effectively reveals misleading information

 Francesco Pierri , Carlo Piccardi  & Stefano Ceri *

In recent years, malicious information had an explosive growth in social media, with serious social and political backlashes. Recent important studies, featuring large-scale analyses, have produced deeper knowledge about this phenomenon, showing that misleading information spreads faster, deeper and more broadly than factual information on social media, where echo chambers, algorithmic and human biases play an important role in diffusion networks. Following these directions, we explore the possibility of classifying news articles circulating on social media based exclusively on a topological analysis of their diffusion networks. To this aim we collected a large dataset of diffusion networks on Twitter pertaining to news articles published on two distinct classes of sources, namely outlets that convey *mainstream*, reliable and objective information and those that fabricate and disseminate various kinds of *misleading* articles, including false news intended to harm, satire intended to make people laugh, click-bait news that may be entirely factual or rumors that are unproven. We carried out an extensive comparison of these networks using several alignment-free approaches including basic network properties, centrality measures distributions, and network distances. We accordingly evaluated to what extent these techniques allow to discriminate between the networks associated to the aforementioned news domains. Our results highlight that the communities of users spreading mainstream news, compared to those sharing misleading news, tend to shape diffusion networks with subtle yet systematic differences which might be effectively employed to identify misleading and harmful information.

In recent years social media have witnessed an explosive growth of malicious and deceptive information. The research community usually refers to it with a variety of terms, such as disinformation, misinformation and most often false (or "fake") news, hardly reaching agreement on a single definition^{1–5}. Several reasons explain the rise of such malicious phenomenon. First, barriers to enter the online media industry have dropped considerably and (dis)information websites are nowadays created faster than ever, generating revenues through advertisement without the need to adhere to traditional journalistic standards (as there is no third-party verification or editorial judgment for online news)⁶. Second, human factors such as confirmation biases⁷, algorithmic biases^{1,8} and naive realism⁹ have exacerbated the so-called *echo chamber* effect, i.e. the formation of homogeneous communities where people share and discuss about their opinions in a strongly polarized way, insulated from different and contrary perspectives^{6,10–13}. Third, direct intervention that could be put in place by platform government bodies for banning deceptive information is not encouraged, as it may raise ethical concerns about censorship^{1,5}.

The combat against online misinformation is challenged by: the massive rates at which malicious items are produced, and the impossibility to verify them all⁵; the adversarial setting in which they are created, as misinformation sources usually attempt to mimic traditional news outlets¹; the lack of gold-standard datasets and the limitations imposed by social media platforms on the collection of relevant data¹⁴.

Most methods for "fake news" detection are carried out by using features extracted from the news articles and their social context (notably textual features, users' profile, etc); existing techniques are built on this content-based evidence, using traditional machine learning or more elaborate deep neural networks¹⁴, but they are often applied to small, ad-hoc datasets which do not generalize to the real world¹⁴.

Recent important studies, featuring large-scale analyses, have produced deeper knowledge about the phenomenon, showing that: false news spread faster and more broadly than the truth on social media^{2,15}; social bots play an important role as "super-spreaders" in the core of diffusion networks⁵; echo chambers are primary drivers

Politecnico di Milano, Department of Electronics, Information and Bioengineering, 20133, Milano, Italy. *email: stefano.ceri@polimi.it

for the diffusion of true and false content¹⁰. In this work, we focus on analyzing the diffusion of misleading news along the direction pointed by these studies.

Leveraging the sole diffusion network allows to by-pass the intricate information related to individual news articles—such as content, style, editorship, audience, etc.—and to capture the overall diffusion properties for two distinct news domains: reputable outlets that produce *mainstream*, reliable and objective information, opposed to sources which notably fabricate and spread different kinds of *misleading* articles (as defined more precisely in the Materials and Methods section). We consider any article published on the former domain as a *proxy* for credible and factual information (although it might not be true in all cases) and all news published on the latter domain as proxies for misleading and/or inaccurate information (we do not investigate whether misleading news can be accurately distinguished also from factual but non-mainstream news originated from niche outlets).

The contribution of this work is manifold. We collected thousands of Twitter diffusion networks pertaining to the aforementioned news domains and we carried out an extensive network comparison using several alignment-free approaches. These include training a classifier on top of global network properties and centrality measures distributions, as well as computing network distances. Based on the results of the analysis, we show that it is possible to classify networks pertaining to the two different news domains with high levels of accuracy, using simple off-the-shelf machine learning classifiers. We furthermore provide an interpretation of classification results in terms of topological network properties, discussing why different features are the footprint of how the communities of users spread news from these two different news domains.

Materials and Methods

Mainstream versus misleading. As highlighted by recent research on the subject^{1–5}, there is not a general consensus on a definition for malicious and deceptive information, e.g. authors in¹⁶ define disinformation as information at the intersection between false information and information intended to harm whereas Wikipedia¹⁷ defines it as “false information spread deliberately to deceive”; consequently, to assess whether a news outlet is spreading unreliable or objective information is a controversial matter, subject to imprecision and individual judgment.

The consolidated strategy in the literature—which we follow in this work—consists of building a classification of websites, based on multiple sources (e.g. reputable third-party news and fact-checking organizations). Along this approach, we characterize a list of websites that notably produce *misleading* articles, i.e. low-credibility content, false and/or hyper-partisan news as well as hoaxes, conspiracy theories, click-bait and satire. We oppose to these malicious sources a set of traditional news outlets (defined as in⁴) which deliver *mainstream* reliable news, i.e. factual, objective and credible information. We are aware that this might not be always true as reported cases of misinformation on mainstream outlets are not rare¹, yet we adopt this approach as it is currently the best available proxy for a correct classification.

Data collection. We collected all tweets containing a Uniform Resource Locator (URL) pointing to websites (specified next) which belong either to a *misleading* or *mainstream* domain. Following the approach described in^{1,3–5,15,18} we assume that article labels are associated with the label of their source, i.e. all items published on a misleading (mainstream) website are misleading (mainstream) articles. We took into account censoring effects described in¹⁹, by retaining only diffusion cascades relative to articles that were published after the beginning of the collection process (*left censoring*), and observing each of them for at least one week (*right censoring*).

For what concerns misleading sources we referred to the curated list of 100+ news outlets provided by^{5,15,18}, which contains websites featured also in^{2–4}. Leveraging Hoaxy API, we obtained tweets pertaining to news items published in the period from Jan, 1st 2019 to March, 15th 2019, filtering articles with less than 50 associated tweets. The final collection comprises 5775 diffusion networks. In Fig. 1 we show the distribution of the number of networks per each source.

We replicated the collection procedure described in^{15,18} in order to gather reliable news articles by using the Twitter Streaming API. We referred to U.S. most trusted news sources described in²⁰ (the list is also available in the Supplementary Information); this includes websites described also in^{3,4}. We associated tweets to a given article after canonicalization of the attached URL(s), using tracking parameters as in^{15,18}, to handle duplicated hyperlinks. We collected the tweets during a window of three weeks, from February 25th 2019 to March 18th 2019; we restricted the period w.r.t the misleading collection in order to obtain a balanced dataset of the two news domains. At the end of the collection, we excluded articles for which the number of associated tweets was less than 50, obtaining 6978 diffusion networks; we show in Fig. 1 the distribution of the number of networks per each source. A different classification approach using several data sampling strategies on networks in the same 3-week period is available in the Supplementary Information. The number of misleading networks is only ~1200, resulting in an imbalanced dataset with misleading/mainstream proportion 1 to 5. Results are nonetheless in accordance with those provided in the main paper.

Furthermore, we assigned a *political bias* label to sources in both news domains, as to perform binary classification experiments considering separately *left*-biased and *right*-biased outlets. We derived labels following the procedure outlined in⁴. Overall, we obtained 4573 left-biased, 1079 centre leaning and 1292 right-biased mainstream diffusion networks; on the other side, we counted 1052 left-biased, 444 satire and 4194 right-biased misleading diffusion networks. Labels for each source in both news domains are shown in Fig. 1, and a more detailed description is provided in the Supplementary Information.

Eventually, mainstream news generated ~1.7 million tweets, corresponding to ~400 k independent cascades, ~680 k unique users and ~1.2 million edges; misleading news generated ~1.6 million tweets, ~210 k independent cascades, ~420 k unique users and ~1.4 million edges.

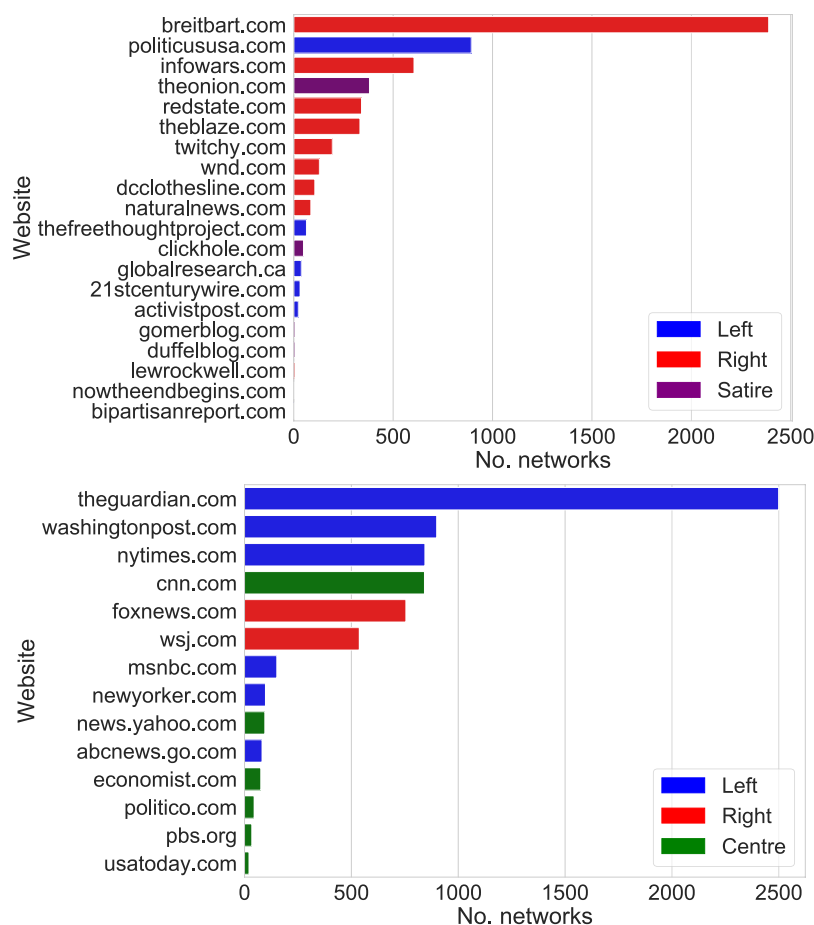


Figure 1. Distribution of the number of networks per each source for misleading (*top*) and mainstream (*bottom*) outlets; colors indicate different political bias labels as specified in the legend.

Twitter diffusion networks. We represent Twitter sharing diffusion networks as directed, unweighted graphs following^{5,15}: for each unique URL we process all tweets containing that hyperlink and build a graph where each node represents a unique user and a directed edge is built between two nodes whenever a user re-tweets/quotes, mentions or replies to another user. Edges between nodes are built only once and they all have weight equal to 1. Isolated nodes correspond to users who authored tweets which were never re-tweeted nor replied/quoted/mentioned.

An intrinsic yet unavoidable limitation in our methodology is that, as pointed out in^{2,15,19}, it is impossible on Twitter to retrieve *true* diffusion cascades because the re-tweeting functionality makes any re-tweet pointing to the original content, losing intermediate re-tweeting users. As such, the majority of Twitter cascades often ends up in *star* topologies. In contrast to², we consider as a single diffusion network the union of several cascades generated from different users which shared the same news article on the social network; thus such network is not necessarily a single connected component. Notice that our approach, although yielding a description of diffusion cascades which might be partial, is the only viable approach based on publicly available Twitter information.

Global network properties. We computed the following set of global network properties, allowing us to encode each network by a tuple of features: **(a)** the number of strongly connected components, **(b)** the size of the largest strongly connected component, **(c)** the number of weakly connected components, **(d)** the size and **(e)** the diameter of the largest weakly connected component, **(f)** the average clustering coefficient and **(g)** the main K-core number. We selected these basic features from the network science toolbox^{21,22}, because our goal is to show that even simple measures, manually selected, can be effectively used in the task of classifying diffusion networks, while an exhaustive search of global indicators is outside of our scope. More details on these features are available in the Supplementary Information.

We observed what follows: **(a)** is highly correlated with the size of the network (see Supplementary Information), as the diffusion flow of news mostly occurs in a broadcast manner, i.e. edges almost consist of re-tweets, and **(b)** allows to capture cases where the mono-directionality of the information diffusion is broken; **(c)** indicates approximately the number of distinct cascades, with exceptions corresponding to cases where two or more cascades are merged together via mentions/quotes/replies on Twitter; **(d)** and **(e)** represent respectively the size and the depth of the largest cascade of a given news article; **(f)** indicates the degree to which users in diffusion

No. nodes	Mainstream networks	Misleading networks	Label
all	6978	5575	D_{all}
[0, 100)	4177	2640	$D_{[0,100)}$
[100, 1000)	2605	2900	$D_{[100,1000)}$
[1000, $+\infty$)	196	235	$D_{[1000,+\infty)}$

Table 1. The number of analyzed diffusion networks, in total (first row) and after splitting based on size (second to last row). In the last column, the label used in the paper to denote the network subset.

networks tend to form local cliques whereas (**g**) is commonly employed in social networks to identify influential users and to describe the efficiency of information spreading¹⁵.

Network distances. In addition, we considered two alignment-free network distances that are commonly used in the literature to assess the topological similarity of networks, namely the Directed Graphlet Correlation Distance (DGCD) and the Portrait Divergence (PD).

The first distance²³ is based on directed graphlets²⁴. These are used to catch specific topological information and to build graph similarity measures; depending on the graphlets (and the orbits) considered, different DGCD can be obtained, e.g. DGCD-13 is the one that we employed in this work. Among all graphlet-based distances, which often yield a prohibitive computational cost to compute graphlets, DGCD has been demonstrated as the most effective at classifying networks from different domains.

The second distance²⁵, which was recently defined, is based on the network portrait²⁶, a graph invariant measure which yields the same value for all graph isomorphisms. This distance is purely topological, as it involves comparing, via Jensen-Shannon divergence, the distribution of all shortest path lengths of two graphs; moreover, it can handle disconnected networks and it is computationally efficient.

We also conducted experiments on several centrality measures distributions—such as total degree, eigenvector and betweenness centrality—and results are available in the Supplementary Information. They overall perform worse than the above methods, in accordance with current literature on network comparison techniques^{27,28}.

Dataset splitting. As we expect networks to exhibit different topological properties within different ranges of node sizes (see also Supplementary Information), prior to our analyses, we partitioned the original collection of networks into subsets of similar sizes. This simple heuristic criterion produced a splitting of the dataset into three subsets according to specific ranges of cardinalities (see Table 1); we also considered the entire original dataset for comparison. Splitting proved effective for improving the classification and also for highlighting interesting properties of diffusion networks.

Results

We employed a non-parametric statistical test, Kolmogorov-Smirnov (KS) test, to verify the null hypothesis (each individual feature has the same distribution in the two classes). Hypothesis is rejected ($\alpha = 0.05$) for all indicators in all data subsets, with a few exceptions on networks of larger size. More details are available in the Supplementary Information.

We then employed these features to train two traditional classifiers, namely Logistic Regression (LR) and K-Nearest Neighbors (K-NN) (with different choices of the number k of neighbors). Experiments on other state-of-the-art classifiers, which exhibit comparable results, are described in the Supplementary Information. Before training each model, we applied feature normalization, as commonly required in standard machine learning frameworks²⁹. Finally we evaluated performances of both classifiers using a 10-fold stratified-shuffle-split cross validation approach, with 90% of the samples as training set and 10% as test set in each fold. In Table 2 we show values for Precision, Recall and F1-score, and in Fig. 2 we show the resulting Receiver Operating Characteristic curve (ROC) for both classifiers with corresponding Area Under the Curve (AUC) values. The performance is in all cases much better than that of a random classifier.

Next, we considered two specific classifiers, Support Vector Machines (SVM) and K-NN, applied to the network similarity matrix computed considering network distances (DGCD and PD). In Fig. 3, we report the AUROC values for the K-NN classifier, trained on top of PD and DGCD similarity matrix; we excluded SVM as it was considerably outperformed (results are available in the Supplementary Information). DGCD was evaluated only on networks with less than 1000 nodes (which still account for over 95% of the data) as the computational cost for larger networks was prohibitive. We carried out the same cross validation procedure as previously described. Again, the performance of the classifiers is in all instances much better than the baseline random classifier value.

Finally, we carried out several tests to assess the robustness of our classification framework when taking into account the political bias of sources, by computing global network properties and evaluating the performances of several classifiers (including balanced versions of Random Forest and Adaboost classifiers using `imblearn` Python package³⁰). We first classified networks altogether excluding two specific sources of misleading articles, namely “breitbart.com” and “politicususa.com”, one at a time and both at the same time; we carried out these tests as these are very prolific sources (the former has by far the largest number of networks among right-biased sources, which is 4 times larger than “infowars.com”, the second uppermost right-biased source; similarly, the latter has 10 times the number of networks of “activistpost.com”, the second uppermost left-biased source). We then evaluated classifiers performance in two different scenarios, i.e. including in training data only mainstream

Classifier	Dataset	Recall	Precision	F1-Score
Logistic Regression	D_{all}	0.71 (sd 0.02)	0.74 (sd 0.02)	0.71 (sd 0.02)
	$D_{[0,100]}$	0.65 (sd 0.01)	0.70 (sd 0.01)	0.65 (sd 0.01)
	$D_{[100,1000]}$	0.75 (sd 0.02)	0.76 (sd 0.01)	0.74 (sd 0.02)
	$D_{[1000,+\infty]}$	0.85 (sd 0.06)	0.86 (sd 0.06)	0.85 (sd 0.06)
K-NN ($k = 10$)	D_{all}	0.70 (sd 0.01)	0.72 (sd 0.02)	0.70 (sd 0.01)
	$D_{[0,100]}$	0.67 (sd 0.02)	0.70 (sd 0.01)	0.67 (sd 0.02)
	$D_{[100,1000]}$	0.76 (sd 0.02)	0.76 (sd 0.02)	0.76 (sd 0.02)
	$D_{[1000,+\infty]}$	0.84 (sd 0.04)	0.84 (sd 0.04)	0.84 (sd 0.04)

Table 2. Evaluation metrics for Logistic Regression and Random Forest classifiers, built using global network properties. We report average values and standard deviations (in brackets) from 10-fold cross validation.

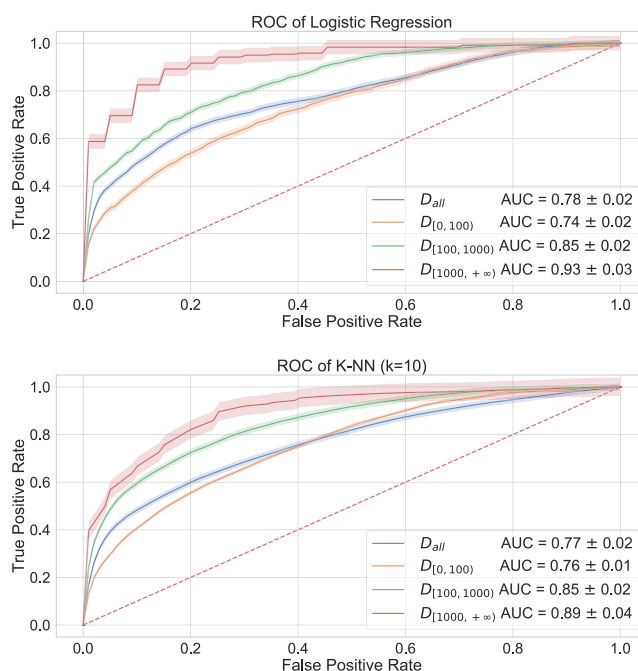


Figure 2. ROC curves for Logistic Regression and K-NN (with $k = 10$) classifiers evaluated using global network properties. The dashed line corresponds to the ROC of a random classifier baseline with $AUC = 0.5$.

and misleading networks with a specific bias (in turn left and right) and testing on the entire set of sources; in Fig. 4 we show the resulting ROC curves with corresponding AUC values. Results are in all cases better than those of a random classifier and in agreement with the result of the classifier developed without excluding any source from the training and test sets; a more detailed description of aforementioned classification results is available in the Supplementary Information.

Discussion

In a nutshell, we demonstrated that our choice of basic global network properties provides an accurate classification of news articles based solely on their Twitter diffusion networks—AUROC in the range 0.75–0.93 with basic K-NN and LR, and comparable or better performances with other state-of-the-art classifiers (see Supplementary Information); these results hold also when considering news based on the political bias of their sources. The use of more sophisticated network distances confirms the result, which is altogether in accordance with prior work on the detection of online political *astroturfing* campaigns³¹, and two more recent network-based contributions on *fake news* detection^{32,33}. However, we remark that, given the composition of our dataset, we did not test whether our methodology allows to accurately classify misleading information vs factual but non-mainstream news which is produced by niche outlets.

For what concerns global network properties, comparing networks with similar sizes turned out to be the right choice, yielding a general increase in all classification metrics (see Supplementary Information for more details). We experienced the worst performances when classifying networks with smaller sizes (with less than 100 nodes); we argue that small diffusion networks appear more similar and that differences across news domains emerge particularly when their size increases.

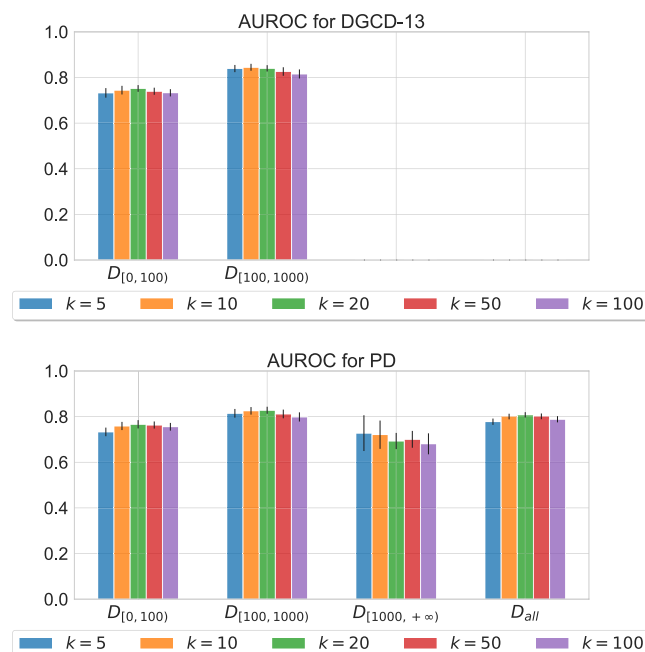


Figure 3. AUROC values for K-NN classifiers (with different choices of k) using PD and DGCD distances.

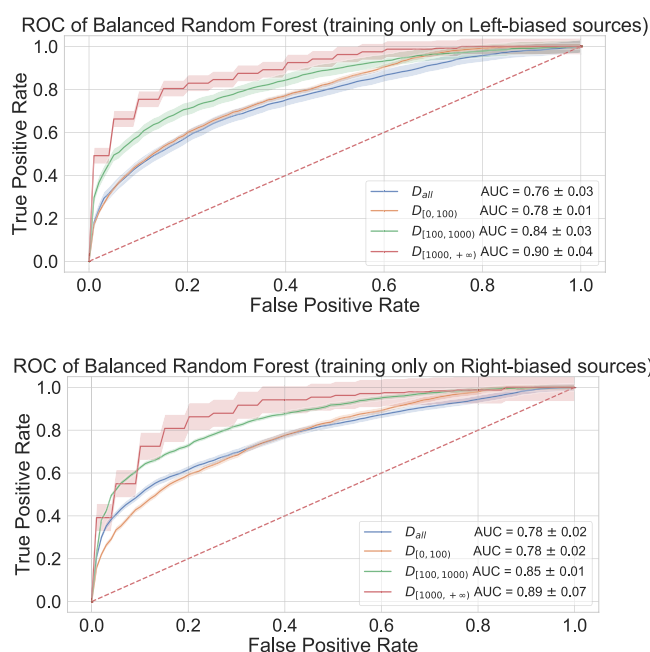


Figure 4. ROC curves for a balanced Random Forest classifier, evaluated using global network properties, training only on left-biased (*top*) or right-biased (*bottom*) sources and testing using all sources. The dashed line corresponds to the ROC of a random classifier baseline with AUC = 0.5.

For what concerns network distances, they overall exhibit a similar trend in classification performances, with worst results on networks with less than 100 nodes and a slight improvement when considering the entire dataset; accuracy in classifying networks with more than 1000 nodes is lower, perhaps due to data scarcity. DCGD and PD distances appear equivalent in our specific classification task; the former is generally used in biology to efficiently cluster together similar networks and identify associated biological functions²³. They reinforce the results of our more naive approach involving a manual selection of the input features.

Understanding classification results in terms of input features is notably a controversial problem in machine learning³⁴. In the following we give our own qualitative interpretation of the results in terms of global network properties.

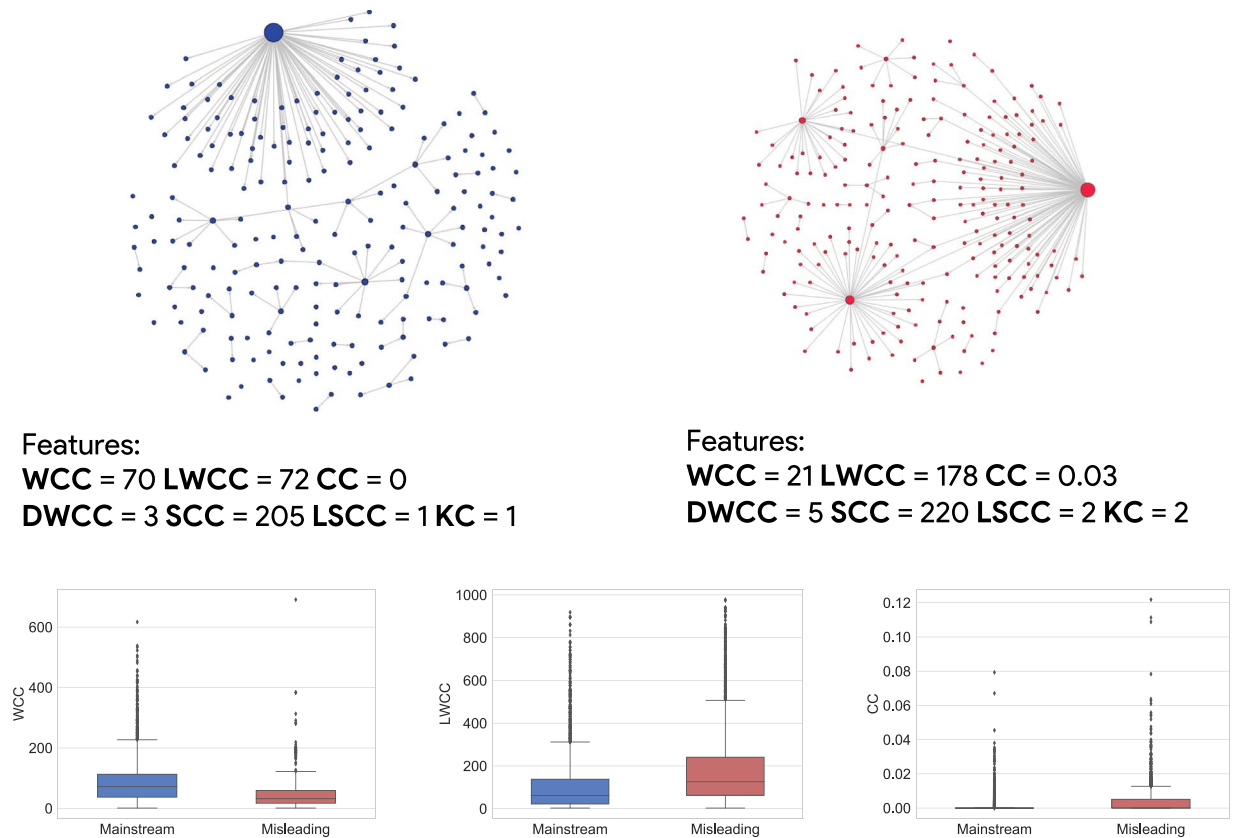


Figure 5. *Top.* Prototypical examples (the *nearest* individuals) of two diffusion networks in the subset $D_{[100,1000]}$ of mainstream (left) and misleading (right) networks. The size of nodes is adjusted according to their degree centrality, i.e. the higher the degree value the larger the node. *Middle.* Feature values corresponding to the two examples (**WCC** = Number of Weakly Connected Components; **LWCC** = Size of the Largest Weakly Connected Component; **CC** = Average Clustering Coefficient; **DWCC** = Diameter of the Largest Weakly Connected Components; **SCC** = Number of Strongly Connected Components; **LSCC** = Size of the Largest Strongly Connected Component; **KC** = Main K-Core Number). *Bottom.* Box-plots of values of the three most significant features—WCC, LWCC, CC—highlighting different distributions in the $D_{[100,1000]}$ subset of the two news domains.

For networks with less than 1000 nodes, we observed that misleading networks exhibit higher values of both size and diameter of the largest weakly connected components; recalling that the largest weakly connected component corresponds to the largest cascade, this result is in accordance with² where it is shown that false rumor cascades spread deeper and broader than true ones.

For networks with more than 100 nodes, we noticed higher values of both size of the largest strongly connected component and clustering coefficient in misleading networks compared to mainstream ones. This denotes that communities of users sharing misleading news tend to be more connected and clustered, with stronger interaction between users, whereas mainstream articles are shared in a more broadcast manner with less discussion between users. A similar result was reported in³⁵ where a sample of most shared news was inspected in the context of 2016 US presidential elections. Conversely, mainstream networks manifest a much larger number of weakly connected components (or cascades). This is not surprising since traditional outlets have a larger audience than websites sharing misleading news^{3,4}.

Finally, we observed that the main K-core number takes higher values in misleading networks rather than in mainstream ones. This result confirms considerations from¹⁵ where authors perform a K-core decomposition of a massive diffusion network produced on Twitter in the period of 2016 US presidential elections; they show that low-credibility content proliferates in the core of the network. More details on differences between news domains, according to the size of diffusion networks, are available in the Supplementary Information.

A pictorial representation of these properties is provided in Fig. 5, where we display two networks, with comparable size, which represent the *nearest* individuals pertaining to both news domains in the $D_{[100,1000]}$ subset, i.e. the network with the smallest Euclidean distance—in the feature space of global network properties—from all other individuals in the same domain. Although they may appear similar at first sight, they actually exhibit different global properties. In particular we observe that the misleading network has a non-zero clustering coefficient, and higher value of size and diameter of the largest weakly connected component, but a smaller number of weakly connected components w.r.t to the mainstream network. Additional examples relative to other subsets are available in the Supplementary Information.

As far as the well known deceptive efforts to manipulate the regular spread of information are concerned (see for instance the documented activity of Russian trolls^{36,37} and the influence of automated accounts in general⁵), we argue that such hidden forces might indeed play to accentuate the discrepancies between the diffusion patterns of misleading and mainstream news thus enhancing the effectiveness of our methodology.

As far as the political bias of sources is concerned, a few contributions^{38–40} report differences in how conservatives and liberals socially react to relevant political events. For instance, Conovet *et al.*³⁸ report discrepancies in a few network indicators (e.g. average degree and clustering coefficient) among three specific pairs of diffusion networks (one for right-leaning users and one for left-leaning users), namely those described by follower/follower relationships between users, re-tweets and mentions, which however differ from the URL-based diffusion networks which we analyzed in this work. In addition they carried out their analysis in a different experimental setting w.r.t to ours: they used Twitter Gardenhose API (which collects a random 10% sample of tweets) and filtered tweets based on political hashtags. Although they found that right-leaning users shared hyperlinks (not necessarily news) 43% of the times compared to 36.5% of left-leaning users (percentages that grow to 51% and 62.5% in case of tweets classified as political), their findings are not commensurable to our topology comparison. Similarly, other works^{4,39,40} used different experimental settings w.r.t to our data collection, they did not specifically focus on URL-wise diffusion networks, and they carried out analyses with approaches not comparable to ours. Nevertheless, despite any possible influence of the political leaning on some features of the diffusion patterns, our methodology proves to be insensitive to the presence of political biases in news sources, as we are capable of accurately distinguishing misleading and mainstream news in several different experimental scenarios. The presence of specific outlets of misleading articles which outweigh the others in terms of data samples (respectively "breitbart.com" for right-biased networks and "politicususa.com" for left-biased networks) does not affect the classification, as results are similar even when excluding articles from these sources. Also, when we considered only left-biased (or right-biased) mainstream and misleading articles in the training data while including all sources in the test set, we observed results which are in accordance with our general aforementioned findings, for what concerns both classification performances and features distributions. Overall, our results show that mainstream news, regardless of their political bias, can be accurately distinguished from misleading news.

Conclusions

Following the latest insights on the characterization of misleading news spreading on social media compared to more traditional news, we investigated the topological structure of Twitter diffusion networks pertaining to distinct domains. Leveraging different network comparison approaches, from manually selected global properties to more elaborated network distances, we corroborate what previous research has suggested so far: misleading content spreads out differently from mainstream and reliable news, and dissimilarities can be remarkably exploited to classify the two classes of information using purely topological tools, i.e. basic global network indicators and standard machine learning.

We can qualitatively sum up these results as follows: misleading articles spread broadly and deeper than mainstream news, with a smaller global audience than mainstream, and communities sharing misleading news are more connected and clustered. We believe that future research directions might successfully exploit these results to develop real world applications that could resolve and mitigate malicious information spreading on social media.

Data availability

The datasets analysed in this work are available from the corresponding author on request.

Received: 30 June 2019; Accepted: 8 January 2020;

Published online: 28 January 2020

References

1. Lazer, D. M. J. *et al.* The science of fake news. *Sci.* **359**, 1094–1096, <https://doi.org/10.1126/science.aao2998>, <http://science.sciencemag.org/content/359/6380/1094.full.pdf> (2018).
2. Vosoughi, S., Roy, D. & Aral, S. The spread of true and false news online. *Sci.* **359**, 1146–1151, <https://doi.org/10.1126/science.aap9559>, <http://science.sciencemag.org/content/359/6380/1146.full.pdf> (2018).
3. Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B. & Lazer, D. Fake news on twitter during the 2016 u.s. presidential election. *Sci.* **363**, 374–378, <https://doi.org/10.1126/science.aau2706>, <http://science.sciencemag.org/content/363/6425/374.full.pdf> (2019).
4. Bovet, A. & Makse, H. A. Influence of fake news in Twitter during the 2016 US presidential election. *Nat. Commun.* **10**, 7, <https://doi.org/10.1038/s41467-018-07761-2> (2019).
5. Shao, C. *et al.* The spread of low-credibility content by social bots. *Nat. Commun.* **9**, 4787 (2018).
6. Allcott, H. & Gentzkow, M. Social media and fake news in the 2016 election. *J. Econ. Perspectives* **31**, 211–36 (2017).
7. Nickerson, R. S. Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* **2**, 175 (1998).
8. Fernandez, M. & Alani, H. Online misinformation: Challenges and future directions. In *Companion of the The Web Conference 2018 on The Web Conference 2018*, 595–602 (International World Wide Web Conferences Steering Committee, 2018).
9. Reed, E. S., Turiel, E. & Brown, T. Naïve realism in everyday life: Implications for social conflict and misunderstanding. In *Values and knowledge*, 113–146 (Psychology Press, 2013).
10. Del Vicario, M. *et al.* The spreading of misinformation online. *Proc. Natl. Acad. Sci.* **113**, 554–559 (2016).
11. Sunstein, C. R. *Echo chambers: Bush v. Gore, impeachment, and beyond* (Princeton University Press, 2001).
12. Sunstein, C. *On Rumors: How Falsehoods Spread, Why We Believe Them, What Can Be Done*. (Yale University Press. Stowe, New Haven, 2007).
13. Pariser, E. *The filter bubble: What the Internet is hiding from you* (Penguin UK, 2011).
14. Pierr, F. & Ceri, S. False news on social media: a data-driven perspective. *ACM Sigmod Rec.* **48**(2) (2019).
15. Shao, C. *et al.* Anatomy of an online misinformation network. *PLOS ONE* **13**, 1–23, <https://doi.org/10.1371/journal.pone.0196087> (2018).
16. Wardle, C. & Derakhshan, H. Information disorder: Toward an interdisciplinary framework for research and policy making. *Counc. Eur. Rep.* **27** (2017).
17. Wikipedia, <https://en.wikipedia.org/wiki/Disinformation>.

18. Shao, C., Ciampaglia, G. L., Flammini, A. & Menczer, F. Hoaxy: A platform for tracking online misinformation. In *Proceedings of the 25th International Conference Companion on World Wide Web, WWW '16 Companion*, 745–750, <https://doi.org/10.1145/2872518.2890098> (International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 2016).
19. Goel, S., Anderson, A., Hofman, J. & Watts, D. J. The structural virality of online diffusion. *Manag. Sci.* **62**, 180–196 (2015).
20. Mitchell, A., Gottfried, J., Kiley, J. & Matsa, K. E. Political polarization & media habits. *Pew Res. Cent.* **21** (2014).
21. Barabási, A.-L. *Network science* (Cambridge University press, 2016).
22. Newman, M. *Networks* (Oxford University press, 2018).
23. Sarajlić, A., Malod-Dognin, N., Yaveroglu, Ö. N. & Pržulj, N. Graphlet-based characterization of directed networks. *Sci. Reports* **6**, 35098 (2016).
24. Itzkovitz, S., Milo, R., Kashtan, N., Ziv, G. & Alon, U. Subgraphs in random networks. *Phys. Review E* **68**, 026127 (2003).
25. Bagrow, J. P. & Boltt, E. M. An information-theoretic, all-scales approach to comparing networks. *arXiv preprint arXiv:1804.03665* (2018).
26. Bagrow, J. P., Boltt, E. M., Skufca, J. D. & benAvraham, D. Portraits of complex networks. *EPL (Europhysics Lett.)* **81**, 68004, <https://doi.org/10.1209/0295-5075/81/68004> (2008).
27. Yaveroglu, Ö. N., Milenković, T. & Pržulj, N. Proper evaluation of alignment-free network comparison methods. *Bioinforma.* **31**, 2697–2704 (2015).
28. Tantardini, M., Ieva, F., Tajoli, L. & Piccardi, C. Comparing methods for comparing networks. *Sci. Reports* **9**, 17557, <https://doi.org/10.1038/s41598-019-53708-y> (2019).
29. Aksoy, S. & Haralick, R. M. Feature normalization and likelihood-based similarity measures for image retrieval. *Pattern Recognit. Lett.* **22**, 563–582 (2001).
30. Lemaitre, G., Nogueira, F. & Aridas, C. K. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *J. Mach. Learn. Res.* **18**, 1–5 (2017).
31. Ratkiewicz, J. *et al.* Detecting and Tracking Political Abuse in Social Media. *ICWSM 2011* 249, <https://doi.org/10.1145/1963192.1963301> 1011.3768 (2011).
32. Monti, F., Frasca, F., Eynard, D., Mannion, D. & Bronstein, M. M. Fake news detection on social media using geometric deep learning. *arXiv preprint arXiv:1902.06673* (2019).
33. Zhao, Z. *et al.* Fake news propagate differently from real news even at early stages of spreading. *arXiv preprint arXiv:1803.03443* (2018).
34. Li, J. *et al.* Feature selection: A data perspective. *ACM Comput. Surv.* **50**, 94:1–94:45, <https://doi.org/10.1145/3136625> (2017).
35. Jang, S. M. *et al.* A computational approach for examining the roots and spreading patterns of fake news: Evolution tree analysis. *Comput. Hum. Behav.* **84**, 103–113 (2018).
36. Badawy, A., Ferrara, E. & Lerman, K. Analyzing the digital traces of political manipulation: the 2016 Russian interference Twitter campaign. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 258–265 (IEEE, 2018).
37. Stewart, L. G., Arif, A. & Starbird, K. Examining trolls and polarization with a retweet network. In *Proceedings ACM WSDM, Workshop on Misinformation and Misbehavior Mining on the Web* (2018).
38. Conover, M. D., Gonçalves, B., Flammini, A. & Menczer, F. Partisan asymmetries in online political activity. *EPJ Data Sci.* **1**, 6 (2012).
39. Bovet, A., Morone, F. & Makse, H. A. Validation of Twitter opinion trends with national polling aggregates: Hillary clinton vs donald trump. *Sci. Rep.* **8**, 8673 (2018).
40. Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A. & Bonneau, R. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychol. science* **26**, 1531–1542 (2015).

Acknowledgements

F.P. and S.C. are supported by the PRIN grant HOPE (FP6, Italian Ministry of Education). S.C. is partially supported by ERC Advanced Grant 693174. The authors are grateful to Hoaxy developers team for their kind assistance to use Hoaxy API and to Simone Vantini for general discussions on the statistical methods.

Author contributions

F.P. designed and performed the experiments. C.P. and S.C. supervised the experiments. All authors wrote and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-58166-5>.

Correspondence and requests for materials should be addressed to S.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020