

Beyond Fact-Checking: Network Analysis Tools for Monitoring Disinformation in Social Media

Stefano Guarino^{1,2}, Noemi Trino², Alessandro Chessa^{2,3}, and Gianni Riotta²

¹ Institute for Applied Computing, National Research Council, Rome, Italy
s.guarino@iac.cnr.it

² Data Lab, Luiss “Guido Carli” University, Rome, Italy

³ Linkalab, Cagliari, Italy

Abstract. Operated by the H2020 SOMA Project, the recently established *Social Observatory for Disinformation and Social Media Analysis* supports researchers, journalists and fact-checkers in their quest for quality information. At the core of the Observatory lies the DisInfoNet Toolbox, designed to help a wide spectrum of users understand the dynamics of (fake) news dissemination in social networks. DisInfoNet combines text mining and classification with graph analysis and visualization to offer a comprehensive and user-friendly suite. To demonstrate the potential of our Toolbox, we consider a Twitter dataset of more than 1.3M tweets focused on the Italian 2016 constitutional referendum and use DisInfoNet to: (i) track relevant news stories and reconstruct their prevalence over time and space; (ii) detect central debating communities and capture their distinctive polarization/narrative; (iii) identify influencers both globally and in specific “disinformation networks”.

Keywords: social network analysis, disinformation, classification

1 Introduction

“SOMA – Social Observatory for Disinformation and Social Media Analysis” is a H2020 Project aimed at supporting, coordinating and guiding the efforts of researchers, fact-checkers and journalists contrasting online and social disinformation, to shield a fair political debate and a responsible, shared, set of information for our citizens. At the core of the Observatory is a web-based collaborative platform for the verification of digital (user-generated) content and the analysis of its prevalence in the social debate, based on a special instance of (SOMA partner) ATC’s Truly Media⁴. In this paper, we present the first prototype of the DisInfoNet Toolbox, designed to support the users of the SOMA verification platform in understanding the dynamics of (fake) news dissemination in social media and tracking down the origin and the broadcasters of false information. We overview current features, preview future extensions, and report on the insights provided by our tools in the analysis of a Twitter dataset.

⁴ <https://www.truly.media/>

Data collected on social media is paramount for understanding disinformation disorders [7] as it is instrumental to: (i) quantitative analyses of the diffusion of unreliable news stories [1]; (ii) comprehending the relevance of disinformation in the social debate, possibly incorporating thematic, polarity or sentiment classification [34]; (iii) unveiling the structure of social ties and their impact on (dis)information flows [3]. DisInfoNet was designed to allow all of the above and more, as it allows tracking specific news pieces in the data and visualizing their prevalence over time/space, classifying content in a semi-automatic fashion (relying on clustering a keyword/hashtag co-occurrence graph), and extracting, analyzing and visualizing social interaction graphs, embedding community-detection and user classification. Additional features will soon enrich the Toolbox, such as a user-friendly interface for Structural Topic Model [29], supporting sentiment analysis both globally and at topic level [16].

To demonstrate the potential of DisInfoNet, we also present an analysis of a dataset of over 1.3M Italian tweets dating back to November 2016 and focused on the constitutional referendum held on December 4, 2016. The significant diffusion of fake news in the phase of political campaign before the vote, together with the dichotomic structure of referendums fostering user polarization, make this dataset especially fit for purpose. Additionally, the distance in time of such a crucial political event makes it easier treating sensitive issues like disinformation while preventing the risk of recentism in analyzing social phenomena. We found evidence of a few relevant false stories in our dataset and, by relating polarization and network analysis, we were able to gain a better understanding of their patterns of production/propagation and contrast, and of the role of renowned authoritative accounts as well as outsiders and bots in driving the production and sharing of news stories. From a purely quantitative point of view, it is worth noting that our findings diverge significantly from what observed by (SOMA partner) Pagella Politica at the time [26], underlining once more that Twitter and Facebook provide very different perspectives on society and that further support of social media platforms is paramount for the research community.

2 Related Work

As reported by a recent Science Policy Forum article [21], stemming the viral diffusion of fake news and characterizing disinformation networks largely remain open problems. Besides the technical setbacks, the existence of the so-called “continued influence effect of misinformation” is widely acknowledged among socio-political scholars [31], thus questioning the intrinsic potential of debunking in contrasting the proliferation of fake news. Yet, the body of research work on fake news detection and (semi-)automatic debunking is vast and heterogeneous, relying on linguistics [22], deep syntax analysis [14], knowledge networks [11], or data mining [30]. Attempts at designing an end-to-end fact-checking system exist [19], but are mostly limited to detecting and evaluating strictly factual claims. Even supporting professional fact-checkers by automating *stance detection* is problematic, due to relatedness being far easier to capture than agree-

ment/disagreement [18]. Approaches specifically conceived for measuring the credibility of social media rumours appear to benefit from the combined effectiveness of analyzing textual features, classifying users’ posting and re-posting behaviors, examining external citations patterns, and comparing same-topic messages [10,35,5]. Unfortunately, this is well beyond what social media analytics and editorial fact-checking tools on the market permit.

In this context, DisInfoNet was designed to help researchers, journalists and fact-checkers characterizing the prevalence and dynamics of disinformation on social media. Recent work confirmed the general perception that, on average, fake news get diffused farther, faster, deeper and more broadly than true news [34,1]. The prevalence of false information is often deemed to be caused by the presence of “fake” and automated profiles, usually called *bots* [6]. The role of bots in disinformation campaigns is however far from being sorted out: albeit bots seem to be the main responsible for fake news production and are used to boost the perceived authority of successful (human) sources of disinformation [3], they have been found to accelerate the spread of true and false news at the same rate [34]. Models for explaining the success of false information without a direct reference to bots have also been recently proposed, either based on information overload vs. limited attention [28], or on information theory and (adversarial) noise decoding [8]. Finally, investigating the relation between polarization and information spreading has been shown to be instrumental for both uncovering the role of disinformation in a country’s political life [7] and predicting potential targets for hoaxes and fake news [33].

3 The Toolbox

DisInfoNet is a Python library built on top of well-known packages (*e.g.*, *igraph*, *scikit-learn*, *NumPy*, *Gensim*), soon to be available under the GPL on GitLab⁵. It provides modules for managing archives, elaborating and classifying text, building and analyzing graphs, and more. It is memory-efficient to support large datasets and, albeit a few functions are optimized for Twitter data, generally flexible. At the same time, DisInfoNet implements a pipeline designed to enable journalists and fact-checkers with no coding expertise assessing the prevalence of disinformation in social media data. This pipeline, depicted in Figure 1, consists of three main tools which may be controlled by a single configuration file – soon to be replaced by a user-friendly dashboard embedded in the SOMA platform. One of DisInfoNet’s main features is the ability to extract and examine both keyword co-occurrence graphs and user interaction graphs induced by a specific set of themes of interest, thus providing valuable insights into the contents and the actors of the social debate around disinformation stories.

The first tool of DisInfoNet’s pipeline is the **Subject Finder**. It filters a dataset and returns information about the prevalence of themes or news pieces of interest. It uses keyword-based queries (migration to document similarity

⁵ Please, contact the authors if you wish to be notified when the code is released.

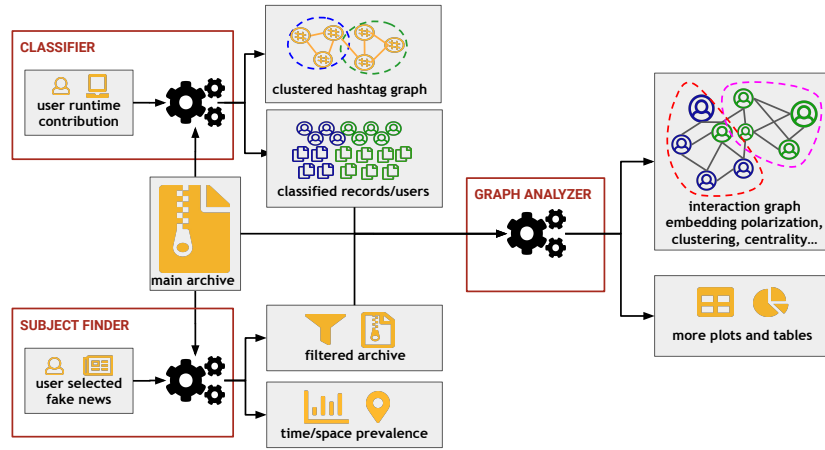


Fig. 1: DisInfoNet’s main pipeline.

is in progress) to extract (parsed) records into a CSV file. For instance, for Twitter data it returns tweets with covariates such as author, timestamp, geo-localization, retweet count, hashtags, mentions. It also plots the temporal and spatial distribution of all and query-matching records.

The **Classifier** partitions records into classes based on a semi-automatic “self-training” process. By building and clustering a keyword co-occurrence graph (that the user may prune of central yet generic and/or out-of-context keywords, detrimental to clustering), it presents the user with an excerpt of the keywords associated with the obtained classes. Significantly, this means using far more keywords than any fully manual approach would permit, without sacrificing accuracy, but rather possibly discovering previously unknown and highly informative keywords. The user can select and label the classes of interest, which are used to automatically extract a training set. The Classifier then selects the best performing model among a few alternative (currently, Logistic Regression and Gradient Boosting Classifier, with 10-fold cross-validation) and predicts a label for all records. When only two classes are used (*e.g.*, republican vs. democratic, right- vs. left-wing, pro vs. against; discussing theme A vs. theme B), the obtained classification may also be extended to users (*e.g.*, authors) by averaging over the classification of all records associated to a specific user.

Finally, the **Graph Analyzer** incorporates functions for graph mining and visualization. It first extracts a directed user interaction graph, wherein two users (*e.g.*, authors) are connected based on how often they interact (*e.g.*, cite each other). It then computes a set of global and local metrics, including: distances, eccentricity, radius and diameter; clustering coefficient; degree and assortativity; PageRank, closeness and betweenness centrality [24]. It also partitions the graph into communities, relying on the well-known Louvain [4] or Leading Eigenvector [25] algorithms, and applies the Guimerà-Amalal cartography [17], based on

discerning inter- and intra-community connections. This results into a number of tables and plots.

4 Politics and Information in 2016 Italy

The 2013 election imposed an unprecedented tri-polar equilibrium in the Italian political scene, with the 5 Stars Movement (5SM) breaking the traditional left-right framework, and the rise of the populist right party Northern League (NL). In 2016, the Italian government guided by the center-left Democratic Party (PD) promoted a constitutional reform which led to a referendum, held on December 4, 2016. Both the 5SM and the NL opposed the referendum, making the NO faction a composite front supported by a wide spectrum of formations with alternative yet sometimes overlapping political justifications.

In this framework, populist movements showed an extraordinary ability in setting the agenda, by imposing carefully selected instrumental news-frames and narratives that found the perfect breeding ground in Italy – the country of political disaffection par excellence [12]. New media, in particular, offered an unprecedented opportunity: to maintain a critical – even conspiratorial – attitude towards the establishment-dominated media, while enhancing the role of alternative/social media as strategic resources for community-building and alternative agenda setting [2]. In these contexts, Twitter plays a strategic role for newly born political parties, that through the activation of the two-way street mediatization may incorporate their proposals into conventional media [9]. The dichotomous structuring of referendum was however instrumental to both sides for aligning the various issues along a pro-anti/status quo spectrum.

The final victory of the NO caused Renzi’s resignation from Head of Government and paved the way for the definite affirmation of the 5SM and the NL, who in 2018 joined forces in forming a so-called “government of change”.

4.1 Disinformation Stories

In order to identify relevant themes of disinformation of the political campaigning we relied on the activity of fact-checking and news agencies, who reported lists of fake news that went viral during the referendum campaign. Mostly based on the work by fact-checking web portal *Bufale.net* [23], online newspaper *Il Post* [27], and SOMA partner and political fact-checking agency *Pagella Politica* [26], we were able to identify the twelve main pieces of disinformation related to the referendum. To widen the scope of the analysis, we considered stories and speculations that reflect information disorders in a broader sense, from rumors, hearsays, clickbait items and unintentionally propagated misinformation, to conspiracy theories and organized propaganda, often used by the two sides to accuse one another. We then classified these disinformation stories into four categories: (i) the **QUOTE** category includes entirely fabricated quotes of public figures endorsing one or the other faction or defaming voters of the other side; (ii) the **CONSQ** group of news contains manipulated interpretations of

genuine information about the (potential) consequences of the reform; (iii) the **PROPG** category includes news inserted in a typical populist frame, opposing people vs the élite; (iv) finally, the **FRAUD** category involves the integrity of the electoral process, gaining unauthorized access to voting machines and altering voting results. Due to page restrictions, in this paper we only study disinformation at this category level, deferring a detailed analysis at news-story level to future work. Significantly, this type of category-based approach is fully supported by DisInfoNet and easily available through the configuration file.

5 Findings

In this section, we demonstrate the potential of DisInfoNet by analyzing a dataset of more than 1.3M tweets to shed light on the dynamics of social disinformation as Italy approached the referendum.

5.1 Disinformation Prevalence

With each of the selected news stories represented by a suitable keyword-based query, we ran the Subject Finder to identify our set of *disinformation tweets*, have them labelled with categories, and obtain the plots in Figure 2 showing their temporal and geographical distribution.

In Figure 2a we see the one-day rolling mean of the four classes across November 2016, compared with the overall trend. The presence of disinformation in the dataset is limited, yet non-negligible: except for QUOTE tweets, each of the other three classes accounts for $\approx 5\%$ of the records. The volume of discussion about fake/distorted news stories does not seem to simply increase at the approach of the referendum as for the general discussion, but different stories have different spikes, possibly related with events (*e.g.*, a politician giving an interview) or with the activity of some influencer. Regarding the geography of the debate, we found that only 29716 tweets – that is, 2.21% of the whole dataset – were geo-tagged, and this percentage is even lower ($\approx 1\%$) among disinformation tweets (see Table 1 for details), possibly due to users involved in this type of discussions being more concerned about privacy than the average. The map, reported in Figure 2b, shows some activity in Great Britain and the Benelux area, but disinformation topics appear to be substantially absent outside Italy.

5.2 Polarization and Disinformation

The Classifier can now be used to gain a better understanding of the relation between polarization and disinformation in our dataset.

During the semi-automatic self-training process, we pruned a few central but out-of-context hashtags (*e.g.*, “#photo” and “#trendingtopic”) and let the Classifier run Louvain’s algorithm and plot the hashtag graph. This graph, reported in Figure 3, shows that: (i) hashtags used by the NO and YES supporters are strongly clustered; (ii) “neutral” hashtags (such as those used by international

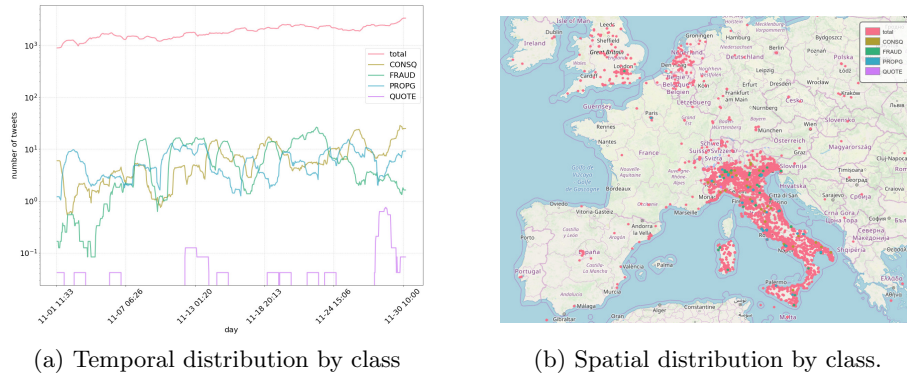


Fig. 2: The temporal and spatial distribution of disinformation tweets.

reporters) also cluster together; (iii) a few hashtags are surprisingly high-ranked, such as “#ottoemezzo”, a popular and supposedly impartial political talk-show being central in the NO cluster – thus confirming regular patterns of behavior in the “second-screen” use of social network sites to comment television programs [32]. In particular, it is easy to identify two large clusters of hashtags clearly characterizing the two sides: the *YES cluster* is dominated by the hashtags “#bastaunsi” (“a yes is enough”) and “#iovotosi” (“I vote yes”), whereas the *NO cluster* by “#iovotono” (“I vote no”), “#iodicono” (“I say no”) and “#renziacasa” (“Renzi go home”). In this perspective, both communities show clear segregation and high levels of clustering by political alignments, thus confirming the hypothesis of social-media platforms as *echo chambers*, with political exchanges exhibiting “a highly partisan community structure with two homogeneous clusters of users who tend to share the same political identity” [12].

By interacting with the Classifier, we selected the aforementioned YES and NO clusters as the sets of hashtags to be used for building a training set. Labelling works as follows: -1 (NO) if the tweet only contains hashtags from the NO cluster; $+1$ (YES) if the tweet only contains hashtags from the YES cluster; 0 (UNK) if the tweet contains a mix of hashtags from the two clusters. Significantly, we also obtained a continuous score in $[-1, 1]$ for each user, as the average score of the user’s tweets. When ran after the Subject Finder, the Classifier also plots a histogram that helps relating classification and disinformation, reported in Figure 3b. We immediately see that UNK tweets are substantially negligible, while NO tweets are almost $1.5\times$ more frequent than YES tweets, supporting the diffused belief that the NO front was significantly more active than its counterpart in the social debate. Disinformation news stories mostly follow the general trend, but: (i) topics of the QUOTE and PROPG classes, which gather attack vectors frequently used by the populist parties, are especially popular among NO supporters (hence, debunking efforts are invisible); (ii) on the other hand, YES supporters are more active than the average in the CONSQ topics, probably due

to the concurrent attempts at promoting the referendum and at tackling the fears of potential NO voters.

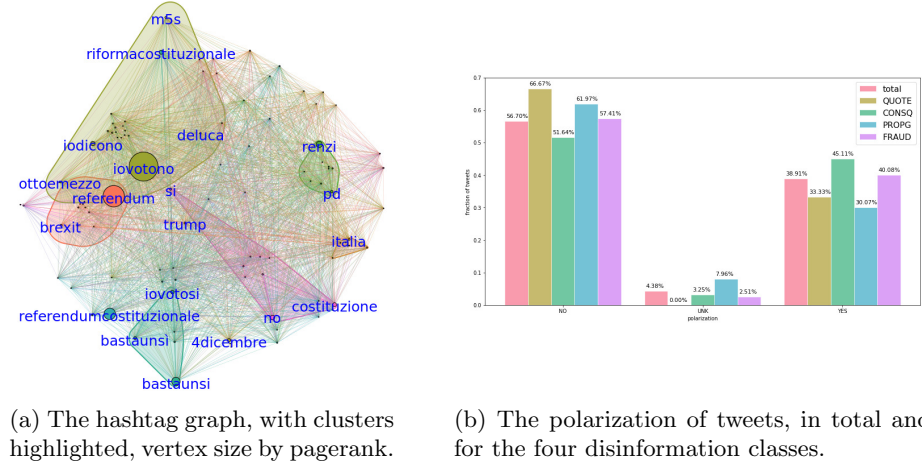


Fig. 3: The hashtag graph and the classification results.

5.3 Interaction Graphs and Disinformation

Finally, we used the Graph Analyzer to better understand the dynamics of disinformation networks in our dataset. Due to page restrictions, in the following we only focus on retweets and on the CONSQ and PROPG disinformation classes, leaving a more detailed analysis to future work. Among the three supported types of interaction, in fact, retweeting is the simplest endorsement tool [20], commonly used for promoting ideas and campaigns and for community building, possibly relying on semi-automatic accounts. On the other hand, the CONSQ and PROPG classes appeared to be the most informative, for both their different polarity distribution and their almost non-intersecting sets of influencers. First of all, we obtained a number of macroscopic descriptors that yield insights into the structural similarities and differences of the two graphs, reported in Table 1. The CONSQ and PROPG are similar in size (2755 vertices and 3786 edges vs. 2126 and 2886) and have similarly sized in- and out-hubs (628 and 16 vs. 653 and 18), but the diameter of the CONSQ graph is significantly smaller (12 vs. 30) despite it having a larger average distance (2.73 vs. 1.64). These numbers suggest that PROPG disinformation stories travelled less on average, but were sporadically able to reach very peripheral users. Additionally, we see that the clustering coefficient of the two graphs is almost identical and rather small (≈ 0.004), more than one order of magnitude smaller than the clustering coefficient of the whole graph. This suggests that these disinformation networks

may not be “self-organizing” and their structure might be governed by artificial diffusion patterns.

Table 1: Dataset overview.

	Tweets	Geotags (%)	Retweet graph						
			vertices	edges	deg_{in}^{\max}	deg_{out}^{\max}	clustering	diam.	avg. dist.
Dataset	1344216	29716 (2.21%)	72574	451423	4813	1541	0.0483	149	4.81044
CONSQ	7909	71 (0.90%)	2755	3786	628	16	0.0039	12	2.72581
PROPG	4345	47 (1.08%)	2126	2886	653	18	0.00385	30	1.63941
FRAUD	5362	69 (1.29%)	2195	3452	692	13	0.00321	8	2.45673
QUOTE	57	1 (1.75%)	9	8	8	1	0.0	1	1.0

For a more close-up analysis, Figure 4 shows, for both classes, the network composed of the top 500 users by pagerank. In these plots, users are colored by their polarity and edges take the average color of the connected vertices. The size of a vertex is proportional to its pagerank, whereas the width of an edge to its weight, *i.e.*, number of interactions between the two users. These plots highlight a number of interesting aspects. First of all, the NO front appears to be generally dominant, with relevant YES actors only emerging in the debate on the alleged consequences of the referendum. Also, there seems to be limited interaction between YES and NO supporters, as can be noted by the fact that edges almost always link vertices of similar or even identical color. Among the leaders of the NO front, we find well-known public figures (*e.g.*, politicians Renato Brunetta and Fabio Massimo Castaldo in the PROPG graph) along with accounts not associated with any publicly known individual. In most cases, these are militants of the NO front, sometimes having multiple aliases, and whose activity is characterized by a high number of retweets and mentions of well-known actors belonging to the same community (*e.g.*, Antonio Bordin, Claudio Degl’Innocenti, Angelo Sisca, Liberati Linda). Additional insights can be gained by using Truthnest⁶, a tool developed by SOMA partner ATC, which reports analytics on the usage patterns of a specified account summarized into a bot-likelihood score. One of the most influential nodes of the PROPG graph, @INarratore, came out having a suspiciously high 60% bot-score, other than only 1% of original tweets and a considerable number of “suspicious followers”. In the same graph, @dukana2 has a 50% bot-score, while the account @advalita has been suspended from Twitter. In the CONSQ graph, the most central user is @ClaudioDeglinn2, characterized by a relatively low 10% bot-score, but apparently in control of at least other 7 aliases and strongly connected with other amplification accounts. Two of these “amplifiers” are especially noteworthy: @IPredicatore, having a 40% bot-score, and @Patriotall, having a 30% bot-score, mentioning @ClaudioDeglinn2 in more than 20% of his tweets, and producing only 3% original tweets. Altogether, we seem to have found indicators of coordinated efforts to avoid bot detection tools while reaching peripheral users and expanding the network.

⁶ <https://app.truthnest.com/>

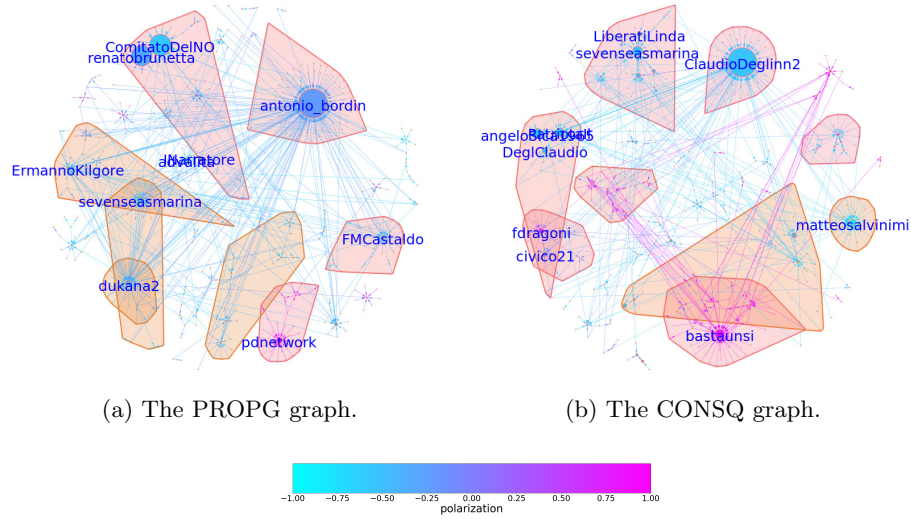


Fig. 4: 500 top users by pagerank. Color is by polarity, size by pagerank.

6 Conclusion

In this paper, we publicly presented – to both the scientific and fact-checking community – an integrated toolbox for monitoring social disinformation, conceived as part of the H2020 Social Observatory for Disinformation and Social Media Analysis. Our DisInfoNet Toolbox builds on well-established techniques for text and graph mining to provide a wide spectrum of users instruments for quantifying the prevalence of disinformation and understanding its dynamics of diffusion on social media. We presented a case study analysis focused on the 2016 Italian constitutional referendum, wherein the natural bipolar political structure of the debate helps in reducing one of the most frequent problem in opinion detection on social media, related to the identification of all possible political orientations (associated to communities). Following the literature [12,15], we resorted to retweets in order to analyze accounts and their interactions according to their possible political orientation. The combined analysis of political communities and network clustering and centrality shows how the referendum caused a clear segregation by political alignment [13], configuring the existence of different *echo-chambers*. From a thematic point of view, news stories related to conspiracy theories and distrust with political élite were especially popular and traveled deeper than any other category of disinformation. We found evidence of a correlation between users’ polarization and participation to disinformation campaigns, and by highlighting the primary actors of disinformation production and propagation we could manually tell apart public figures, activists and potential bots. Our DisInfoNet Toolbox will soon be available online and extended

in the next future. We believe that the state-of-the-art techniques for classification and network analysis embedded in the Toolbox will pave the way for future research in the area, crucial to the preservation of our public conversation and the future of our democracies.

References

1. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31(2), 211–36 (2017)
2. Alonso-Muñoz, L., Casero-Ripollés, A.: Communication of european populist leaders on twitter: Agenda setting and the ‘more is less’ effect. *El profesional de la información* 27(6), 1193–1202 (2018)
3. Bessi, A., Ferrara, E.: Social bots distort the 2016 us presidential election online discussion (2016)
4. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008(10), P10008 (2008)
5. Boididou, C., Middleton, S.E., Jin, Z., Papadopoulos, S., Dang-Nguyen, D.T., Boato, G., Kompatsiaris, Y.: Verifying information with multimedia content on twitter. *Multimedia Tools and Applications* 77(12), 15545–15571 (2018)
6. Boshmaf, Y., Musluhkhov, I., Beznosov, K., Ripeanu, M.: Design and analysis of a social botnet. *Computer Networks* 57(2), 556–578 (2013)
7. Bovet, A., Makse, H.A.: Influence of fake news in twitter during the 2016 us presidential election. *Nature communications* 10(1), 7 (2019)
8. Brody, D.C., Meier, D.M.: How to model fake news. *arXiv preprint arXiv:1809.00964* (2018)
9. Casero-Ripollés, A., Feenstra, R.A., Tormey, S.: Old and new media logics in an electoral campaign: The case of podemos and the two-way street mediatization of politics. *The International Journal of Press/Politics* 21(3), 378–397 (2016)
10. Castillo, C., Mendoza, M., Poblete, B.: Information credibility on twitter. In: *Proceedings of the 20th international conference on World wide web*. pp. 675–684. ACM (2011)
11. Ciampaglia, G.L., Shiralkar, P., Rocha, L.M., Bollen, J., Menczer, F., Flammini, A.: Computational fact checking from knowledge networks. *PloS one* 10(6), e0128193 (2015)
12. Conover, M., Ratkiewicz, J., Francisco, M.R., Gonçalves, B., Menczer, F., Flammini, A.: Political polarization on twitter. *Icwsm* 133, 89–96 (2011)
13. Conover, M.D., Gonçalves, B., Flammini, A., Menczer, F.: Partisan asymmetries in online political activity. *EPJ Data Science* 1(1), 6 (Jun 2012)
14. Feng, V.W., Hirst, G.: Detecting deceptive opinions with profile compatibility. In: *Proceedings of the Sixth International Joint Conference on Natural Language Processing*. pp. 338–346 (2013)
15. Garimella, K., Weber, I.: A long-term analysis of polarization on twitter. *CoRR abs/1703.02769* (2017)
16. Guarino, S., Santoro, M.: Multi-word structural topic modelling of tor drug marketplaces. In: *2018 IEEE 12th International Conference on Semantic Computing (ICSC)*. pp. 269–273. IEEE (2018)
17. Guimerà, R., Nunes Amaral, L.: Functional cartography of complex metabolic networks. *Nature* 433(895) (2005)

18. Hanselowski, A., PVS, A., Schiller, B., Caspelherr, F., Chaudhuri, D., Meyer, C.M., Gurevych, I.: A retrospective analysis of the fake news challenge stance detection task. arXiv:1806.05180 (2018)
19. Hassan, N., Zhang, G., Arslan, F., Caraballo, J., Jimenez, D., Gawsane, S., Hasan, S., Joseph, M., Kulkarni, A., Nayak, A.K., et al.: Claimbuster: the first-ever end-to-end fact-checking system. *Proceedings of the VLDB Endowment* 10(12), 1945–1948 (2017)
20. Kantrowitz, A.: The Man Who Built The Retweet: “We Handed A Loaded Weapon To 4-Year-Olds”. www.buzzfeednews.com/article/alexkantrowitz/how-the-retweet-ruined-the-internet (2019), [accessed 05-Aug-2019]
21. Lazer, D.M., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., et al.: The science of fake news. *Science* 359(6380), 1094–1096 (2018)
22. Markowitz, D.M., Hancock, J.T.: Linguistic traces of a scientific fraud: The case of diederik stapel. *PloS one* 9(8), e105937 (2014)
23. Mastinu, L.: TOP 10 Bufale e disinformazione sul Referendum. www.bufale.net/top-10-bufale-e-disinformazione-sul-referendum/ (2016), [accessed 05-Jul-2019]
24. Newman, M., Barabasi, A.L., Watts, D.J. (eds.): *The Structure and Dynamics of Networks*. Princeton University Press (2006)
25. Newman, M.E.: Finding community structure in networks using the eigenvectors of matrices. *Physical review E* 74(3), 036104 (2006)
26. Politica, R.P.: La notizia più condivisa sul referendum? È una bufala. <https://pagellapolitica.it/blog/show/148/la-notizia-pi%C3%B9-condivisa-sul-referendum-%C3%A8-una-bufala> (2016), [accessed 05-Jul-2019]
27. Post, R.I.: Nove bufale sul referendum. www.ilpost.it/2016/12/02/bufale-referendum/ (2016), [accessed 05-Jul-2019]
28. Qiu, X., Oliveira, D.F., Shirazi, A.S., Flammini, A., Menczer, F.: Limited individual attention and online virality of low-quality information. *Nature Human Behaviour* 1(7), 0132 (2017)
29. Roberts, M.E., Stewart, B.M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S.K., Albertson, B., Rand, D.G.: Structural topic models for open-ended survey responses. *American Journal of Political Science* 58(4), 1064–1082 (2014)
30. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter* 19(1), 22–36 (2017)
31. Skurnik, I., Yoon, C., Park, D.C., Schwarz, N.: How warnings about false claims become recommendations. *Journal of Consumer Research* 31(4), 713–724 (2005)
32. Trilling, D.: Two different debates? investigating the relationship between a political debate on tv and simultaneous comments on twitter. *Social Science Computer Review* 33(3), 259–276 (2015)
33. Vicario, M.D., Quattrociocchi, W., Scala, A., Zollo, F.: Polarization and fake news: Early warning of potential misinformation targets. *ACM Transactions on the Web (TWEB)* 13(2), 10 (2019)
34. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. *Science* 359(6380), 1146–1151 (2018)
35. Zubiaga, A., Aker, A., Bontcheva, K., Liakata, M., Procter, R.: Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys (CSUR)* 51(2), 32 (2018)