



Manipulation During the French Presidential Campaign: Coordinated Inauthentic Behaviors and Astroturfing Analysis on Text and Images

Victor Chomel^{1,2}(✉), Mazyar Panahi¹, and David Chavalarias^{1,2}

¹ Centre National de la Recherche Scientifique (CNRS), Institut des Systèmes Complexes Paris Ile-de-France (ISC-PIF), Paris, France

victor.chomel@iscpif.fr

² Ecole des Hautes Etudes en Sciences Sociales (EHESS), Paris, France

Abstract. In April 2022, the French presidential election took place, and social media played a prominent role in it. By analyzing more than 150 million interactions on French Twitter, this study aims to provide evidence of coordinated behaviors from political parties. We find that extreme parties left and right, appear with a particular internal structure compared to moderate parties. Moreover, by examining similar patterns in community structures but also in duplicated tweets, we unveil online astroturfing strategies of the main parties online, and in particular the extreme right.

Keywords: Computational social science · Misinformation · Disinformation · Astroturfing

1 Introduction

It is believed that social media exerts an increasing influence on the outcome of elections [3]. Indeed, the recent 2022 French presidential election anticipated social media to play a crucial role in shaping election results. During the political campaign season, political parties across the spectrum sought to impose their platform in the public discourse and mobilize their constituencies online in a coordinated, strategic manner. Our study proposes and validates a method to measure the scale of these online information campaigns by probing the coordination patterns and structure of online communities. For instance, in order to reach a larger audience, political campaigns may use coordination techniques that are considered *disinformation* [18]. One such case is posting numerous identical media instead of retweeting. Contrary to a retweet that explicitly references another user, this method masks the real coordination behind these tweets and can artificially enlarge the audience. These *astroturfing* methods can be detected and measured for each community. Astroturfing can be seen as an attempt to give the false impression that a message or organization enjoys widespread grassroots

support of the community when little such support exists. Due to astroturfing, some communities consistently push their narratives into trending topics for the entire country. With only a few hundred accounts, political parties manage daily to reach out to hundreds of thousands of accounts outside of their original community.

Related Work Misinformation and disinformation have gained momentum in public discourse during the last few years. A broad definition of misinformation is incorrect or misleading information. It is differentiated from disinformation, which is deliberately deceptive. Information spread regarding COVID-19 has accelerated both the production of disinformation and misinformation as well as their detection [9]. There is a growing literature that seeks to understand the roots of misinformation and disinformation, including evaluating their volume [11], determining fake news ecosystems [1], detecting misleading content [7], characterizing astroturfing [17], viewing activism as advertising [5], or disinformation as distraction [16]. The worldwide prevalence of astroturfing and coordinated inauthentic behaviors [14, 24], especially in a political context [8], demonstrates a need to analyze the last French presidential election in this framework.

Notably, methods have been developed to uncover coordinated disinformation networks based on Twitter data [22]. These astroturfing campaigns [19] have sometimes been seen as related to foreign countries trying to undermine the trust in the political system [2] or promoting their authoritarian views [21]. Most of them are now cross-platform, duplicating content and creating narratives before coming to a large audience [13, 20]. These coordination patterns can come from several potential coordination mechanisms such as several accounts run by the same person, a shared channel with instructions on how/when to share, or semi-organic behavior by users that rapidly amplifies target accounts.

Many papers analyze bot activity to answer the astroturfing problem [6, 25, 26]. However, analyzing disinformation using bot detection is being increasingly called into question [23]. Though bots participate in astroturfing campaigns, they do not represent the entirety of astroturfing activity [15]. These information campaigns are largely planned and relayed by true accounts and even accounts with an already substantial audience.

Most of these studies focus only on one factor such as time coordination [4], structure, or identical content. In this study, we propose to gather these solutions to show that, across different modalities (text, image, video), they all point in the same direction— that is, to the presence of coordinated information campaigns by online political communities.

1.1 Data

The study period starts in October 1st, 2021 and ends on the day of the second round of the French presidential election, April 24th, 2022, covering the French presidential electoral campaign. The data has been collected with the Twitter Public Stream API using the same methodology than Gaumont et al. [10]. Most

of top French political accounts (Governments, Members of parliament, Members of European parliament, Senators, Mayors of big cities, etc.) from all political parties are followed with the API. There are more than 2000 selected accounts. Moreover, thousands of elections keywords such as names of candidates and parties are tracked. If a selected account tweets, retweets or quotes, or if it is retweeted, mentioned, quoted, the interaction is collected. If any message (tweet, retweet or quote) has a keyword in it, it is collected. About 160 million interactions (tweets, retweets, and quotes) have been collected over this period.

1.2 French Political Landscape

Political communities are defined based on the retweet graph. A retweet graph is a graph with users as nodes and retweets without quotes as edges. We compute the retweet graph over the whole period (from October 2021 to April 2022), then perform a clustering, using the Louvain algorithm, to which we assign political labels. In addition to the previous political communities, we define a “media community” in our clustering, consisting of users who primarily engage with major media outlets. Communities size and average degrees are reported in Table 1. In previous work, this method has been tested and had an averaged precision above 90%, so with few false positives [10]. During these seven months, the set of political communities remained stable and thus their internal structures can be studied.

Table 1. Main political communities with number of users and the average number of retweets per user

Political community with their leaders	Number of nodes	Average number of retweets per user over the whole period
Alt. Right—Dupont-Aignan	39,330	27.4
Ext. Left—Mélenchon	39,098	22.4
Center—Macron	33,306	19.2
Ext. RightZemmour	30,783	45.4
Right—Péresse	11,341	18.7
Ext. Right—Le Pen	9060	21.2
Left—Hidalgo	5770	8.3
Green—Jadot	5582	9.4
Ext. Left—Roussel	3634	10.3
Media	22,848	4.1
Other	45,808	3.4

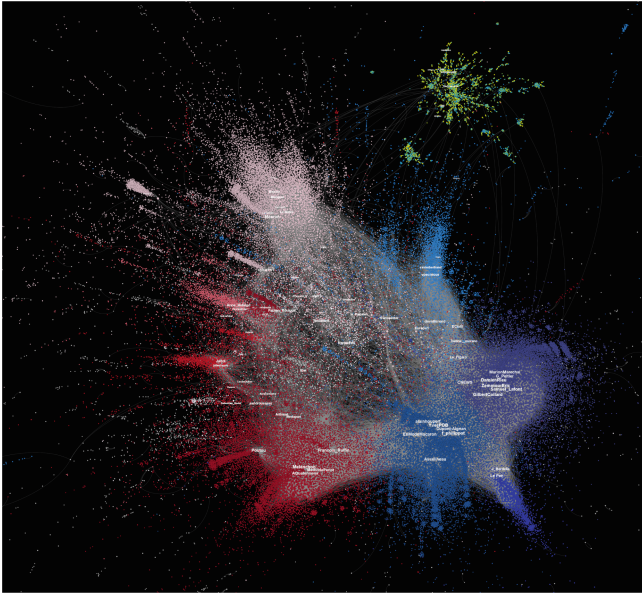


Fig. 1. Retweet graph of the French political debate—March 2022. Nodes are located with ForceAtlas2 in Gephi [12]. Ukraine community, on the top right in yellow/green, gathers users from Ukraine mentioning and retweeting Emmanuel Macron to support the war effort.

The French political landscape (Fig. 1) is composed of 9 online political communities, with their candidates in italics, divided into three main blocks:

- Left: Parti socialiste—*Hidalgo* (Left party, pink on the left of Fig. 1), EELV—*Jadot* (Green party, red on the left), Parti communiste français—*Roussel* (Extreme Left, red on the bottom left) and France Insoumise—*Mélenchon* (Extreme Left, red on the bottom)
- Center: République En Marche—*Macron* (Government, Light pink on the top)
- Right: Les Républicains—*Pécresse* (Right, light blue), Rassemblement National—*Le Pen* (Extreme Right, blue on the bottom right), Reconquête—*Zemmour* (Extreme Right, dark blue on the right) and an alt-right community including several parties (dark blue on the bottom).

2 Astroturfing Criterion and Definition of Communities

Intra-community Dynamics We first compare the extent to which each community engages its users. To do so, we measure the online social mobility of each community using a method based on Markov chains [27]. Each community’s




Fig. 2. Successive layers on a toy example. In blue are the remaining nodes for the next step. In red are the nodes in the layer.

users are categorized into layers of increasing importance. This is done by iteratively removing nodes with the lowest degree (Fig. 2). As this layer distribution is recomputed each two weeks, the layer of each node can be computed for all time steps. Then, the probability of transitioning between layers at a given time step may be computed for a given community. If a node leaves a community, it goes to layer 0. For each community, the transition matrices are then averaged over all time steps, from October 2021 to April 2022.

The probability matrix of moving from layer i to layer j may be formulated as the transition matrix of a Markov chain, with layers as the state. This allows us to capture the internal dynamics of communities, the extent to which users remain in a given community, and the extent to which users participate during the political campaign. This is summarized by the stationary distribution of the Markov chain.

Communities which experience high artificial tweet volume and coordinated campaigns get a particular type of signature in this method. To measure these coordinations, several other indicators are used such as co-retweet graphs, identical tweets, or identical images.

Co-retweet Graphs A co-retweet graph \mathcal{G}_T is a graph where nodes are user accounts. Two nodes are linked if the two accounts retweet the same tweet within T seconds from each other, that is to say, they *co-retweet* each other. To capture changes in the network on an informative timescale T , we use small values of T because we are looking for coordination. The lower the T , the more there is coordination. The weight of an edge is defined as the number of co-retweets between two users that occur within a time difference of T . A co-retweet graph is an aggregation over a longer period, typically a week. As we seek to find clusters of nodes that are strongly connected—knowing that they all participate in the same coordinated information campaign—we excise all spurious or weak connections. The probability of two users retweeting the same content within a small amount of time is really low, it rarely happens twice, spurious connections are thus defined by edge weight ≤ 2 .

Identical Tweets Campaign During the presidential campaign, a strategy used to give the impression of volume was to post numerous exactly identical messages without retweeting the original poster, which we will refer to as an *information campaign*. Users coordinated on private discussion but also on other

platforms such as 4chan to publish the same content at the same time [29]. Another strategy is to have a trending hashtag although the tweets may originate from a very small group of coordinated users. In what follows, we will consider all coordinated campaigns with least five different users.

Coordinated Identical Images and Video Campaign Posted images and videos need not be exactly identical to be considered part of the same information campaign. For example, images may be cropped or otherwise slightly edited. Therefore, we expand our criterion for image “sameness” to include these altered images. We collect images and videos attached to tweets with at least one retweet in order to gather “relevant” content, or content that has been shared at least once. For the sake of simplicity, only the first frame of videos is kept and used as image. Due to the volume of the data, only the last two weeks before the presidential elections (28/03/22–10/04/22) and the two weeks between the two rounds of voting (11/04/22–24/04/22) are collected. About 100,000 images and videos are collected for each period.

The media are embedded as 512 dimensional vectors using a pre-trained ResNet18 from img2vec. For each image, cosine similarities with all other images are computed. Only almost identical images (similarity above 0.95) are kept. To understand information campaign coordination, we order images by their number of tweets (and not retweets) to see how many times they have been posted. For each image, we collect the posters and the users who retweet the image. The 100 most tweeted and retweeted images and videos were checked by hand to validate the method. On this small sample, the method has a precision of 1, meaning that all videos and images labeled as identical or near identical are correctly labeled (no false positives). Accuracy or recall requires to count false negatives which is extremely hard due to the size of the dataset.

3 Results

3.1 Intra-community Dynamics

Stationary distributions of user engagement for seven political blocs and the media can be found in Fig. 3. Here, the user engagement is defined as the ability of a party to mobilize its users and push them towards a higher online activity. The higher the layer, the more engaged users are. By taking a look at the values at $x = 0$, or “user attrition”, one can see that the extreme parties have low attrition, thus more stable communities. This is confirmed by the percentage of supporters (value at $x = 1$). The only two anomalous cases (the right and




Fig. 3. Stationary probability of being in each layer by community.

Table 2. Level of political engagement quantified by the mean layer of the stationary distribution by community, ordered descending from left to right.

Community	ExtRight Zemmour	AltRight	ExtLeft	ExtRight Le Pen	Center	Right	Left	Media
Mean layer of stationary distribution	3.44	3.29	3.15	2.93	2.67	2.35	1.67	0.93

the extreme-right - Le Pen) reflect the fact that a substantial part of their supporters left to join Zemmour, reflecting higher attrition. As expected, the media community is mainly composed of users that retweet occasionally, leading to a high attrition and a high percentage of supporters.

The mean of the distribution can also be a good proxy to characterize the engagement (Table 2). The mean engagement of extreme parties is higher, while on the other hand, that of centrist parties and the media is lower.




Fig. 4. Stationary probability of being in each layer by community. Zoom on high layers.

3.2 Co-retweet Graphs

We compute co-retweet graphs from the last part of the political campaign using several time differences T between retweets.

We quantify the level of coordination in online communities by computing a *coordination ratio*, defined as the number of nodes in the co-retweet graph divided by the number of nodes in the retweet graph by community. This value is plotted as a function of time between retweets (Fig. 5). This ratio captures coordination as it represents the share of the community engaged in coordination (measured by co-retweets), correcting the community size effect.




Fig. 5. Coordination ratio against the time between retweets

We find that, mirroring our results on levels of political engagement, extreme communities tend to have higher coordination ratios, with the main exception being the Alt-Right community. For lower timescales such as 30 or 60s, the Alt-Right community appears less coordinated than the Center. This may be explained by the fact that it is not a well organised community but mainly a gathering of people coming for different reasons (vaccine criticism, conspiracy theories, sanitary pass opponents, etc.).

However, this ratio does not take into account the volume of the activity of the community. A large part of the network could be involved in a low-intensity coordinated campaign for example. The average degree in the co-retweet graph per user against the time between retweets can be found in Fig. 6.




Fig. 6. Average degree in co-retweet graph by community against time between retweets

Based on these first indications and on previous work [24], it appears that the right value to detect the coordination (for example in the case of the Alt-Right) is at $T = 60$ s. The value is large enough to let users retweet but low enough to exhibit coordination. We fix $T = 60$ s and plot the coordination ratio and average degree, computed over snapshots of two weeks, from October 2021 to April 2022.

The coordination ratio and the average degree over time track political campaign events in each community (Figs. 7 and 8). For example, we see the main political parties appear to get increasingly coordinated with the incoming deadline of the first round ($t = 23$), then witness a dip in coordination ratio for several communities explained by their loss in the first round. Compare to the Center, headed by Macron, and the Extreme Right with Le Pen, whose coordination continue to increase to week 25. Moreover, the only primary election that occurred during the period was the Right party primary located in the beginning at the sixth week.




Fig. 7. Moving average (window = 3 time steps) of the coordination ratio as a function of timescale T in seconds (from Oct. 2021 to April 2022).

3.3 Identical Tweets

Most of the information campaigns are coming from abroad (Tigray war, Russo-Ukrainian war, etc.) but their influence on French public discourse is limited, as they do not come from users who are well-rooted in political communities. To measure the impact of campaigns,




Fig. 8. Moving average (window = 3 time steps) of the average degree in the co-retweet graph against time steps (form Oct. 2021 to April 2022).

we thus rank users according to the number of campaigns they participate in. The political colors of each of them are computed based on the results of the clustering from the retweet graph. A limit of this analysis is that a user who, on purpose, only tweets and never retweets (for example a bot), would not be politically labeled but it can clearly support a candidate.

We remark that the extreme right community of Zemmour is over-represented in top users: all the top 100 users belong to this community. Among the top 500 users, 257 of them come from this community and 213 are not politically affiliated. The extreme right community behind Eric Zemmour highly push content and hashtags on a daily basis. Indeed, their intention was to break into trending topics as much as possible to reach users outside of their community.

Users posting same tweets can be seen as a graph with users as nodes and edges between users having shared at least a campaign. Two of such graphs have been computed with two thresholds on edge weights. Two users are linked if they share at least a defined number of campaigns. Many values have been tested that can be gathered in two cases. A low value (typically 2 or 3) leads to many users with weak ties. A higher value (typically 5 or higher) is useful to detect the core of users that always act together. The presented values are 2 and 5.

As one may expect, the resulting graph is a union of many small connected components. Most of them are related to an above mentioned political community or to a special issue such as the war in Ukraine.

Interestingly, we have found a large connected component (more than 1000 users) that doesn't belong to a unique political community, in the lower threshold settings. It gathers accounts that campaigns in favor of ecological issues and try to push users towards petitions or crowdfunding. By changing the threshold of the degrees from 2 to 5, the community disappears meaning that it gathers a lot of users participating to low intensity campaigns.

To verify if accounts in this community are labeled as bots, the Botometer [28] has been used. This tool gives a score between 0 and 1 to each account. The bot score distribution on this large component has a bimodal pattern as reported in the original paper but with slightly more bots. About one third of the accounts scores above 0.7 (probably bots) and half of them below 0.3 (not detected as bots). As expected, many bots belong to this community but surprisingly, about half of the accounts seems to, at least partially, be run by humans.

3.4 Identical Images

Images and videos are key vectors during a political campaign to communicate campaign goals, or push narratives. Media that are part of the campaign: meetings, posters or images of the candidates, are widely posted and shared intra-community. However, other images which, e.g. criticise the government, are shared extra-community. We compare who creates, repeatedly posts, and retweets such content and examine the communities of origin and dissemination. Images can be found in Appendix 4.

When examining information campaigns, e.g. the relationship between Putin and Le Pen, it appears that most posting users come from the center community

Table 3. Several main image-narratives developed during the campaign, ordered by the number of tweets.

Period	Information campaign	Tweets	Retweets
Two weeks before the first round	Campaign Zemmour	967	30,158
	Videos Zemmour	640	7223
	Macron McKinsey	582	5787
	State Violence	361	7621
	Putin with Lepen	314	5667
	Ecology from Macron	476	4410
Between the two rounds	Putin with Le Pen	433	14,645
	State violence	360	2566
	Putin with Macron	254	20,906
	Anti-elite images	71	11,963

but with retweets from non-affiliated nodes. Indeed, among the retweets, more than 60% of the retweets come from another community than center. This shows that the political center manages to take root in public opinion.

After the first round, the two remaining parties were the center of Macron and the extreme right of Le Pen. To be elected, both candidates attempted to garner the vote of the extreme left (around 20% of votes in the first round). To do so, Macron defended his performance on environmental issues during his last presidential term. Even if they posted hundreds of messages (Table 3), the candidate did not receive the expected response. Among the more than 4000 retweets, more than 95% of them circulated within the original community.

On the other hand, the extreme right launched a narrative around images of Macron with Putin as an answer to information campaign concerning Putin and Le Pen. Even though this narrative has been pushed by the extreme right parties, more than half of the retweets are outside of extreme right parties.

4 Discussion

Our examination of internal community dynamics over time reveals a complex picture of online political discourse. We detect and quantify coordination strategies used by communities (especially the extreme right), such as posting numerous identical tweets to lend the community artificial volume, strategies which have contributed to a never-before-seen level of *astroturfing* in French politics. To compare, when the same method for identical messages has been applied to the 2017 presidential election, only one tweet was identically tweeted more than 1000 times (compared to 26 in 2022), 2975 tweets were tweeted more than 5 times (26489 in 2022) and 1101 users participated to campaigns with at least 5 tweets (6507 in 2022).

Coordination, as it has been defined in this paper, has some limits. For example, several accounts tweeting the same image because it's coming from another

platform can happen. Co-retweeting popular content isn't rare. To avoid these pitfalls, the accounts and the contents appearing the more "coordinated" according to the metrics have been manually checked. Almost all of them denoted an intention to disinform in some ways (fake accounts, repeatedly posting and deleting contents, extremely high activity, etc.). Another limit of this approach is that a counterfactual network without coordination can only be theoretical. The question of the expected level of coordination of an uncoordinated political community is open.

If election outcomes do not seem to reflect a candidate's online influence, especially for the extreme-right candidate Zemmour, it has pushed the boundaries of the extreme right sphere due to visibility given by inflated media coverage.

Acknowledgement. We would like to express our very great appreciation to Emily S. Cheng and Jeanne Bruneau-Bongard for their valuable and constructive suggestions during the writing of this research work. We would also like to thank reviewers for taking the time and effort necessary to review the manuscript and appreciate their valuable comments and suggestions for the discussion.

Appendix



Fig. 9. Example of repeatedly posted images within two narratives : the proximity of Marine Le Pen and Putin and the violence of the French police, especially during the Yellow Jackets movement



Fig. 10. Most posted images during the two weeks before the first round (left), between the two rounds (right). The first one is a manipulation of the Macron political campaign poster, with a caption about the McKinsey scandal. The second one is part of Macron's campaign against Le Pen.

References

1. Allen, J., Howland, B., Mobius, M., Rothschild, D., Watts, D.J.: Evaluating the fake news problem at the scale of the information ecosystem. *Sci. Adv.* **6**(14), eaay3539 (2020). <https://www.science.org/doi/abs/10.1126/sciadv.aay3539>
2. Bail, C.A., Guay, B., Maloney, E., Combs, A., Hillygus, D.S., Merhout, F., Freelon, D., Volfovsky, A.: Assessing the Russian internet research agency impact on the political attitudes and behaviors of American twitter users in late 2017. *Proc. Nat. Acad. Sci.* **117**(1), 243–250 (2020). <https://www.pnas.org/doi/abs/10.1073/pnas.1906420116>
3. Bright, J., Hale, S., Ganesh, B., Bulovsky, A., Margetts, H., Howard, P.: Does campaigning on social media make a difference? evidence from candidate use of twitter during the 2015 and 2017 U.K. elections. *Commun. Res.* **47**(7), 988–1009 (2020). <https://doi.org/10.1177/0093650219872394>
4. Cinelli, M., Cresci, S., Quattrociocchi, W., Tesconi, M., Zola, P.: Coordinated inauthentic behavior and information spreading on twitter. In: *Decision Support Systems*, p. 113819 (2022). <https://doi.org/10.1016/j.dss.2022.113819>
5. Enos, R.D., Hersh, E.D.: Party activists as campaign advertisers: the ground campaign as a principal-agent problem. *Am. Polit. Sci. Rev.* **109**(2), 252–278 (2015). <https://doi.org/10.1017/s0003055415000064>
6. Ferrara, E., Varol, O., Davis, C., Menczer, F., Flammini, A.: The rise of social bots. *Commun. ACM* **59**(7), 96–104 (2016). <https://doi.org/10.1145/2818717>
7. Ferreira, R.R.: Liquid disinformation tactics: overcoming social media countermeasures through misleading content. *J. Pract.* 1–22 (2021). <https://doi.org/10.1080/17512786.2021.1914707>

8. Freelon, D., Wells, C.: Disinformation as political communication. *Polit. Commun.* **37**(2), 145–156 (2020). <https://doi.org/10.1080/10584609.2020.1723755>
9. Gallotti, R., Valle, F., Castaldo, N., Sacco, P., Domenico, M.D.: Assessing the risks of ‘infodemics’ in response to COVID-19 epidemics. *Nature Human Behav.* **4**(12), 1285–1293 (2020). <https://doi.org/10.1038/s41562-020-00994-6>
10. Gaumont, N., Panahi, M., Chavalarias, D.: Reconstruction of the socio-semantic dynamics of political activist twitter networks-method and application to the 2017 French presidential election. *PLOS ONE* **13**(9), 1–38 (2018). <https://doi.org/10.1371/journal.pone.0201879>
11. Guess, A.M., Nyhan, B., Reifler, J.: Exposure to untrustworthy websites in the 2016 US election. *Nature Human Beh.* **4**(5), 472–480 (2020). <https://doi.org/10.1038/s41562-020-0833-x>
12. Jacomy, M., Venturini, T., Heymann, S., Bastian, M.: Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software. *PLOS ONE* (2014). <https://doi.org/10.1371/journal.pone.0098679>
13. K., S., Wilson, T.: Cross-platform disinformation campaigns: lessons learned and next steps. *Harvard Kennedy School (HKS) Misinf. Rev.* (2020). <https://doi.org/10.37016/mr-2020-002>
14. Keller, F., Schoch, D., Stier, S., Yang, J.: How to manipulate social media: analyzing political astroturfing using ground truth data from South Korea. *Proc. Int. AAAI Conf. Web Social Media* **11**(1), 564–567 (2017). <https://ojs.aaai.org/index.php/ICWSM/article/view/14941>
15. Keller, F.B., Schoch, D., Stier, S., Yang, J.: Political astroturfing on twitter: how to coordinate a disinformation campaign. *Polit. Commun.* **37**(2), 256–280 (2020). <https://doi.org/10.1080/10584609.2019.1661888>
16. King, G., Pan, J., Roberts, M.E.: How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *Am. Polit. Sci. Rev.* **111**(3), 484–501 (2017). <https://doi.org/10.1017/s0003055417000144>
17. Kovic, M., Rauchfleisch, A., Sele, M., Caspar, C.: Digital astroturfing in politics: definition, typology, and countermeasures. *Stud. Commun. Sci.* **18**(1) (2018). <https://doi.org/10.24434/j.scoms.2018.01.005>
18. Lazer, D.M.J., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Slosman, S.A., Sunstein, C.R., Thorson, E.A., Watts, D.J., Zittrain, J.L.: The science of fake news. *Science* **359**(6380), 1094–1096 (2018). <https://www.science.org/doi/abs/10.1126/science.aao2998>
19. Linvill, D.L., Warren, P.L.: Troll factories: manufacturing specialized disinformation on twitter. *Polit. Commun.* **37**(4), 447–467 (2020). <https://doi.org/10.1080/10584609.2020.1718257>
20. Lukito, J.: Coordinating a multi-platform disinformation campaign: internet research agency activity on three U.S. social media platforms, 2015–2017. *Polit. Commun.* **37**(2), 238–255 (2020). <https://doi.org/10.1080/10584609.2019.1661889>
21. Miller, B.A.P.: Automatic detection of comment propaganda in Chinese media. *SSRN Electron. J.* (2015). <https://doi.org/10.2139/ssrn.2738325>
22. Pacheco, D., Hui, P., Torres-Lugo, C., Truong, B.T., Flammini, A., Menczer, F.: Uncovering coordinated networks on social media. *CoRR abs/2001.05658* (2020). <https://arxiv.org/abs/2001.05658>
23. Rauchfleisch, A., Kaiser, J.: The false positive problem of automatic bot detection in social science research. *PLOS ONE* **15**(10), 1–20 (2020). <https://doi.org/10.1371/journal.pone.0241045>

24. Schoch, D., Keller, F.B., Stier, S., Yang, J.: Coordination patterns reveal online political astroturfing across the world. *Sci. Rep.* **12**(1) (2022). <https://doi.org/10.1038/s41598-022-08404-9>
25. Stella, M., Ferrara, E., Domenico, M.D.: Bots increase exposure to negative and inflammatory content in online social systems. *Proc. Nat. Acad. Sci.* **115**(49), 12435–12440 (2018). <https://www.pnas.org/doi/abs/10.1073/pnas.1803470115>
26. Stukal, D., Sanovich, S., Bonneau, R., Tucker, J.A.: Detecting bots on Russian political twitter. *Big Data* **5**(4), 310–324 (2017). <https://doi.org/10.1089/big.2017.0038>
27. Chomel, V., Panahi, M., Chavalarias, D.: A Markov chain based method to model the dynamical properties of misinformation networks—application to climate change on twitter. In: *Proceedings of Disinformation in Open Online Media: Second Multidisciplinary International Symposium* (2022)
28. Yang, K.C., Ferrara, E., Menczer, F.: Botometer 101: social bot practicum for computational social scientists (2022). <https://arxiv.org/abs/2201.01608>
29. Zannettou, S., Caulfield, T., De Cristofaro, E., Sirivianos, M., Stringhini, G., Blackburn, J.: Disinformation warfare: understanding state-sponsored trolls on twitter and their influence on the web. In: *Companion Proceedings of The 2019 World Wide Web Conference*, pp. 218–226. WWW '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3308560.3316495>