

Political Communities on Twitter: Case Study of the 2022 French Presidential Election

Hadi Abdine^{1*}, Yanzhu Guo^{1*}, Virgile Rennard^{1,2*}, Michalis Vazirgiannis^{1,3}

¹Ecole Polytechnique, ²Linagora, ³AUEB

{hadi.abdine, yanzhu.guo, virgile.rennard}@polytechnique.edu
mvazirg@lix.polytechnique.fr

Abstract

With the significant increase in users on social media platforms, a new means of political campaigning has appeared. Twitter and Facebook are now notable campaigning tools during elections. Indeed, the candidates and their parties now take to the internet to interact and spread their ideas. In this paper, we aim to identify political communities formed on Twitter during the 2022 French presidential election and analyze each respective community. We create a large-scale Twitter dataset containing 1.2 million users and 62.6 million tweets that mention keywords relevant to the election. We perform community detection on a retweet graph of users and propose an in-depth analysis of the stance of each community. Finally, we attempt to detect offensive tweets and automatic bots, comparing across communities in order to gain insight into each candidate’s supporter demographics and online campaign strategy.

Keywords: French Presidential Election 2022, Natural Language Processing, Political Community Detection, Social Media

1. Introduction

Social media has created a forum for everyone to express themselves, bringing disputes to a wide audience and playing an increasingly crucial part in today’s information economy. With the 2008 U.S. presidential election, a relatively new paradigm was observed, where a large part of the political campaign was held on either Facebook or Twitter. The extensive outreach of these platforms has been shown to bring multiple benefits to politicians, such as increases in donations (Petrova et al., 2021), or an amplified impact on the politically inattentive youth (Utz, 2009). The tremendous amount of data provided by social media platforms gives us insight into the inner workings of the online political horizon.

This paper aims to present an in-depth study on the Twitter landscape of the 2022 French elections. We start by creating a Twitter dataset containing more than 60 million tweets from more than a million users. The tweets are extracted based on keywords related to the election. We use this dataset to build a retweet graph among the users and run a graph-based algorithm for community detection. By analyzing the top hashtags and word clouds of tweets posted by users of each community, we are able to interpret which candidate they each support. We go on to visualize the geographical distribution of each candidate’s online supporters across different regions, making comparisons between communities. Eventually, we perform offensiveness detection and bot detection in all the communities. The detection of offensive tweets reveals that supporters of certain candidates are more likely to post offensive contents. The results of bot detection also indicates that there are higher levels of bot activities in certain on-

line communities than in others. However, we would like to emphasize that the results of both offensiveness detection and bot detection are produced by automatic classification models reflecting patterns of the datasets they were trained on. Such models are subject to various limitations and by no means reflect our personal opinions.

The rest of this paper is organized as follows. Section 2 provides an overview of the related work. Section 3 describes our the dataset we use and how we collected it. Section 4 supplies a detailed description of the graph-based communities. Section 5 detects offensive tweets in each community while section 6 studies the use of automated bot accounts. Finally, section 7 summarizes our research and presents potential future work.

2. Related Work

Early work on Twitter analysis of political elections dates back to when Twitter was founded. An example is a study on the 2008 U.S. presidential election (Diakopoulos and Shamma, 2010) where the debate performance of presidential candidates is characterized by aggregated Twitter sentiment. This initiated a branch of research centered around monitoring online public reactions during election periods. Relevant studies have been carried out on a wide range of elections in different countries, including the 2012 South Korea presidential election (Bae et al., 2013), the 2013 German parliamentary election (Rill et al., 2014), and the 2017 UK general election (Yaqub et al., 2020).

Another popular branch of research aims at forecasting election results based on Twitter data. For example, a study on the 2009 German federal election (Tumasjan et al., 2010) claims that the respective shares of Twitter volume can accurately reflect the distribution of electoral votes for the six main parties. However, another

*These authors contributed equally to this work

study on the 2011 Singapore presidential election using Twitter sentiment succeeded in picking out the top two candidates but failed to predict the final ranking (Choy et al., 2011). It is generally agreed upon that Twitter analysis for election outcome prediction cannot substitute traditional polling approaches (Birmingham and Smeaton, 2011), and that explainable models should accompany the predictive results (Gayo-Avello et al., 2011).

More recently, attention has been drawn to the diffusion of misinformation and toxicity during online campaigns. Relevant topics include the detection of fake news (Cinelli et al., 2020), social bots (Pastor-Galindo et al., 2020), political trolls (Badawy et al., 2018) as well as hate speech (Siegel et al., 2021) and offensive language (Grimminger and Klinger, 2021). Our work is closely related to this field of study while also drawing upon graph-based analysis of Twitter network structures (Radicioni et al., 2021).

3. Dataset

We create a novel Twitter dataset with the 2022 French presidential election as the central topic. The tweets used for building this dataset date from February 14, 2022 to April 5, 2022. This dataset corresponds to a large and coherent corpus consisting of small pieces of text related to the election candidates and major events during their campaigns. In our case, we focus on tweets that include tokens such as “présidentielle” and “élection”, as well as the names of candidates and their parties. We extract tweets in the French language containing the keywords mentioned above through Twitter’s public streaming API. The public streaming API is able to extract a subset of the real-time Twitter stream. The resulting dataset consists of 62.6 million tweets and 1.2 million users. Our dataset of Tweet IDs is available upon request. We adhere to Twitter’s Developer Agreement and Policy^{*} and therefore can only distribute up to a total of 1,500,000 Tweet IDs to a single entity within a 30 day period.

4. Graph-based Community Detection

Given that the election has been at the center of attention in France since the beginning of 2022, daily campaigns, scandals, and debates are widespread through social media. We believe that there would be detectable online communities related to the different candidates and that each community will contain important information about each party’s campaign. Therefore, we construct a retweet graph of Twitter users based on our dataset.

4.1. Graph Creation

With our set of extracted tweets, their authors, and the users who retweet these tweets, we create a directed weighted graph $G = (V, E)$ where $n = |V|$ denotes

^{*}<https://developer.twitter.com/en/developer-terms/agreement-and-policy>

the number of nodes. Specifically, each node represents a user on Twitter, and a weighted edge connects two nodes if one user retweets the other. For instance, the weight $A_{u,v}$ of the edge (u, v) from vertex v to vertex u notes the number of times that the user u retweeted the user v . We believe that a normal retweet indicates approval of the tweet’s content, unlike the quote retweets and the replies. Therefore, the graph does not model the textual similarity of the tweets but only the relation between users. Note that the obtained graph G might contain self-loops that correspond to self-retweets. The obtained graph contains 1.2M nodes and 12.4M edges. To avoid small, non-dense clusters formed from few numbers of users with few retweets, we decided to work on a dense subset of the graph that we obtained using the k-core decomposition algorithm instead of working on a full graph. The k-core decomposition algorithm (Seidman, 1983) aims to find subsets of a graph G . The subsets are called k-cores of G and are obtained by a recursive pruning strategy. Each node inside a k-core is connected to at least k other nodes inside this subset. The hyperparameter k is chosen so that we do not get a cluster with less than ten users. The final used k-core graph has 47,578 nodes and 8.2M edges. Seven communities are found in our dataset through applying this approach.

4.2. Community Detection

In graph mining, community detection helps to reveal the hidden relations among the nodes in a graph. Hence, to discern opinion groups inside the k-core graph, we apply the Louvain community detection method (Blondel et al., 2008) on an undirected version of our graph without self-loops. The Louvain method is chosen as it has reasonable computation costs while maximizing modularity. Moreover, it does not require fine-tuning of hyperparameters. It is therefore the only applicable method for finding communities in large graphs.

Seven communities were found after applying this approach to the user-based k-core graph. To get a general overview of the communities, we compute the frequencies of hashtags used by the users inside each community. We define the frequency of a hashtag as the number of users using this hashtag inside a community. We notice that first six out of these seven communities have remarkable hashtags with high frequencies that relate each community to a candidate in the election. For example, in one community with 16,190 users, the top three hashtags are #melenchonvagagner, #melenchonsecondtour and #jevotemelenchon with respectively 16,073, 14,378 and 10,127 users. In contrast, no other hashtags supporting other candidates appear in this community’s 50 most frequent hashtags. Thus, we label this community with “Mélenchon”. Finally, the seventh and last community only holds keywords against the current President of France in the top 50 hashtags




Figure 2: Choropleth map of users in each community.

| Community | # Accounts | Frequent hashtags |
|-------------|------------|--|
| Mélenchon | 15,001 | #melenchonvagagner, #melenchonsecondtour, #jevotemelenchon |
| Anti-Macron | 12,428 | #macrondehors, #toutsaufmacron, #macrondegage |
| Zemmour | 12,101 | #zemmourpresident, #jevotezemmour, #zemmourpresident2022 |
| Macron | 5,816 | #avecvous, #macron2022, #5ansdeplus |
| Pécresse | 1,035 | #valeriepresidente, #pecresse2022, #nouvellefrance |
| Le Pen | 1,001 | #marinepresidente, #dimanchejevotemarine, #jevotemarine |
| Jadot | 196 | #jadot2022, #jevotepadot, #totalsoutienàjadot |

Table 1: Statistics of different communities.

as we can see in the second row in table 1. We thus label it as the "Anti-Macron" community. In table 1 we detail the labels, the number of users, as well as the most biased and frequent hashtags for each community.

4.3. Word Cloud Analysis

To further analyze the political stances of the seven communities, we group the tweets posted by users of each community and generate a word cloud for each group of tweets. The resulting word clouds assemble the set of the most frequent unigrams and bigrams (after removing stop words) in each group of tweets, with their sizes proportional to their frequencies. Our analysis shows that except for the community with the "Anti-Macron" theme, all the other communities are distinctively composed of users who support one of the candidates among Emmanuel Macron, Valérie Pécresse, Éric Zemmour, Jean-Luc Mélenchon, Marine Le Pen and Yannick Jadot. The word clouds for each of these communities are shown in Figure 1. In addition to identifying which candidate each community of users supports, we can also use the word clouds to obtain a general idea of each candidate's campaign program. For example, the word cloud for Jadot's community contains words such as "justice social (social justice)" and "écologiste (ecologist)" with significant weights, and these concepts precisely lie at the center of his campaign program.

4.4. Geolocation Analysis

For each community of users, we plot the distribution of their geolocations within different regions of Metropolitan France. We select users with geolocation information and feed the declared locations into the geolocator from `geopy`* to obtain the region each area belongs to. We only consider users who are located within Metropolitan France. Although this limits our analysis to a subset of users from the dataset, the distribution can still reflect the overall situation.

Given that users are remarkably concentrated in the Île de France region for all the communities, we plot the choropleth map with the number of users scaled by logarithm to visualize the variations for the other areas more clearly. The choropleth map is shown in Figure 2. Through observing these maps, we can gain insights into each candidate's respective heartland. For instance, Le Pen has a much higher proportion of supporters in Occitanie, Provence-Alpes-Côte d'Azur and Hauts-de-France than all the other candidates. We also find that the supporters of Macron, Pécresse and Mélenchon are more evenly distributed among different regions outside of Île de France than the other three candidates. Zemmour and Le Pen both lack Twitter supporters in the middle and northwestern areas, while Jadot lacks supporters in the whole northern half of France except the region Hauts-de-France.

*<https://github.com/geopy/geopy>

Our study was carried out right before the first turn of the election. We later did a follow-up of the actual results and found that the choropleth map of votes* bear a striking resemblance to our choropleth maps of communities.

5. Detection of Offensive Tweets

Along with the growing popularity of social media and online platforms, the use of offensive online language has become a significant problem. Under such circumstances, automatic detection of offensive language has received much research attention (Risch et al., 2020). With presidential elections being such a controversial topic, relevant tweets are bound to contain offensive language. It is expected that online supporters of a given candidate would make offensive comments towards other candidates and their supporters. However, supporters of the more extreme candidates might also be more inclined to use offensive language. The targets of offensive tweets might include other public groups in addition to opposing candidates and their supporters. We build an automatic classification model to detect offensive tweets in each political community and eventually compare the results across communities.

5.1. Detection Model

We initialize our model using BERTweetFR (Guo et al., 2021) and fine-tune it on the MLMA Hate Speech Dataset (Ousidhoum et al., 2019). BERTweetFR is a French RoBERTa model (Liu et al., 2019), initialized using the general-domain French language model CamemBERT (Martin et al., 2020) and further fine-tuned on 16GB of French tweets. It achieves the state-of-the-art performance on French Twitter tasks. The MLMA Hate Speech Dataset is a multilingual multi-aspect Twitter dataset for hate speech analysis. We take a subset of this dataset selecting only French tweets labeled as either "normal" or "offensive". The resulting subset contains 821 normal tweets and 1690 offensive tweets. After fine-tuning for 3 epochs, our classification model achieves a f1 score of 83.96% on a 80/20 train-test split.

5.2. Detection Results

We run our classification model for tweets posted by users in each political community. We only consider unique tweets, discarding retweets by deduplicating them based on the text content. This choice is made because we aim to detect the origination of offensive language rather than to analyze its propagation pattern. The detection results are shown in Table 2. A key observation is that users from the Anti-Macron and Zemmour communities are the most likely to post offensive tweets, reaching respective proportions of 0.307 and 0.305. This is in line with our expectations: the Anti-Macron community is naturally supposed to be more offensive as the main goal is to oppose and defy; as for

*<https://tinyurl.com/ycxya5dx>

| Community | # Unique Tweets | # Offensive Tweets | Proportion of Offensive Tweets |
|-------------|-----------------|--------------------|--------------------------------|
| Mélenchon | 756,318 | 208,178 | 0.275 |
| Anti-Macron | 549,138 | 168,685 | 0.307 |
| Zemmour | 1,034,538 | 316,214 | 0.305 |
| Macron | 468,138 | 126,122 | 0.269 |
| Pécresse | 80,365 | 19,487 | 0.242 |
| Le Pen | 86,272 | 25,368 | 0.294 |
| Jadot | 12,340 | 1,632 | 0.132 |

Table 2: Offensive Tweets in Each Community

Zemmour, he is a far-right candidate who has been personally fined €10,000 for hate speech by a Paris court*. We also observe that the communities of right-wing candidates tend to have more offensive content in general, with the only exception being Pécresse’s who is a more moderate candidate. A possible explanation is the employment of automatic bots in her community. We will further elaborate on the topic of bots in the following section.

6. Detection of Automatic Bots

It has come to light in recent years that a significant amount of Twitter accounts are controlled, at least partly, by software. Some research estimate that between 9% to 15% of all twitter accounts are somewhat automated (Varol et al., 2017). Bots are an important tool for opinion manipulation (Subrahmanian et al., 2016), and are being used to influence important subjects such as political elections ((Ferrara, 2017),(Deb et al., 2019)). Bots also help spread misinformation ((Shao et al., 2017)), and have impacted the online debate on vaccination (Broniatowski et al., 2018), with an estimated 45% of COVID-19 related Twitter accounts exhibiting bot-like behavior ((Memon and Carley, 2020)). This section proposes to estimate the role bots are playing in the 2022 French election by comparing their relative use within each community.

6.1. Detection model

There is a multitude of available Twitter bot detection models ((Lee et al., 2021), (Kudugunta and Ferrara, 2018), (Miller et al., 2014), (Ali Alhosseini et al., 2019)), and APIs ((Davis et al., 2016)) that use semantic, statistical or neighborhood properties to evaluate the likelihood of an account being a bot. However, due to the large amount of data we have to process, we need to use a scalable and generalizable models. This study is therefore going to rely on statistical features available in the user metadata object given by the Twitter API. Our employed model is similar to the one presented in (Yang et al., 2019b), which is scalable and yields adequate generalization results on the task of bot detection (Feng et al., 2021) in different scenarios.

*<https://www.bbc.com/news/world-europe-60022996>

6.1.1. Feature Selection

The list of available user metadata features relevant to bot detection is listed as below:

- STATUSES COUNT
- FOLLOWERS COUNT
- FRIENDS COUNT
- FAVOURITES COUNT
- LISTED COUNT
- DEFAULT PROFILE
- VERIFIED
- GEOGRAPHICAL LOCATION ENABLED

These available features give other interesting statistical information that we compute as additional derived features for the model. Such features include the frequency of tweets (statuses count/user age), the respective growth rate of followers, friends, favorites and listed accounts (respective counts/user age). We also take into account information from the username, such as its length and the number of digits it contains. The length of user description is also proven to be a relevant feature (Yang et al., 2019b).

We choose random forest as our classifier, as it yields near-perfect results on any individually labeled dataset.

6.1.2. Training Data

The choice of training data for such a task is crucial. There are many different types of bots for different domains, and there is generally poor classification generalization across datasets (Echeverría et al., 2018). Considering the task at hand which is the classification of politically oriented Twitter users into bots or human labels, we decided to train our model on a concatenation of multiple available datasets: **Political-bots-2019** (Yang et al., 2019a) (a compendium of political bots), **midterm-2018** (Yang et al., 2019b) (a hand-labeled dataset of users and bots during the 2018 American midterm elections), **botwiki** (Yang et al., 2019b) (a collection of self identified Twitter bots), **verified-2019** (Yang et al., 2019b) (a collection of verified Twitter users), **Cresci 2019-2018** ((Mazza et al., 2019), (Cresci et al., 2018)) (datasets of manually annotated bots), and finally **Twibot-20** (Feng et al., 2021) (a comprehensive hand labeled dataset of Twitter bots). The statistics of each dataset is shown in Table 4.

| Community | # Accounts | # Bots | Proportion of bots in community | # Tweets | # Automated Tweets | proportion Automated tweets |
|-------------|------------|--------|---------------------------------|------------------|--------------------|-----------------------------|
| Mélenchon | 15,001 | 2,507 | 0.167 | 5,755,664 | 1,273,656 | 0.284 |
| Anti-Macron | 12,428 | 2,181 | 0.175 | 5,435,820 | 1,268,240 | 0.304 |
| Zemmour | 12,101 | 2,217 | 0.183 | 6,160,153 | 1,501,207 | 0.322 |
| Macron | 5,816 | 1,001 | 0.172 | 2,219,491 | 514,820 | 0.302 |
| Pécresse | 1,035 | 208 | 0.200 | 408,319 | 134,373 | 0.490 |
| Le Pen | 1001 | 184 | 0.184 | 463,290 | 127,633 | 0.380 |
| Jadot | 196 | 30 | 0.153 | 66,541 | 19,876 | 0.426 |

Table 3: Bots statistics of different communities

| Datasets | # human | # bots |
|---------------------|---------|--------|
| POLITICAL-BOTS-2019 | 0 | 62 |
| MIDTERM-2018 | 8,092 | 42,446 |
| BOTWIKI | 0 | 698 |
| VERIFIED-2019 | 1,987 | 0 |
| TWIBOT-20 | 5,237 | 6,589 |
| CRESCI-18/19 | 6,514 | 7,455 |
| TOTAL | 21,830 | 57,250 |

Table 4: Statistics of training datasets.

6.1.3. Training results

Correlation It is crucial to consider the most discriminative features when attempting a task like bot detection, the model should be interpretable. For example, as we see in Figure 3, there is a strong correlation between the automation of an account and the age of the account, whether or not the account is verified or geolocalisation enabled. Another strong indicator of automation is the presence of a default-profile, which means an account with a lack of personalization (i.e. custom banner or profile picture).

Feature importance By studying the importance of each feature, we find that the number of statuses, the information of followers (such as its raw count and growth rate) and the age of the user are the most critical features for the classifier.

Results Our random forest classifier achieves a 95.0% f1 score with 10-fold cross validation.

6.2. Bot Detection Results

From our experiments, we observe that there is a significant amount of automated accounts in our dataset – with an estimation of at least 15%, with a conservative labeling threshold, of accounts partaking in the debate coming from bots. We have tuned the classification threshold, which is usually at 50% certainty to 75%, considering the importance of precision in the case of bot classification.

As shown in Table 3, while the number of users in each community is significantly different, we find a similar

proportion of bots for each cluster. Therefore, we conjecture that there has not been any large-scale operation to influence the election with bots from any side.

Another insight we can get from Table 3 is the campaigning approach of each community. We do not deduplicate tweets to remove retweets, considering the importance of retweeting in automated accounts. For example, the cluster supporting Zemmour, while being smaller than some of the others, has significantly more tweets, including retweets per person, showing the particular engagement Zemmour supporters seem to offer online. Similarly, we can see that it is the only large cluster with the most considerable bot activity, albeit by a small margin. On the other side, we can see that the Pécresse cluster, smaller in scale, has heavy activities from bots. This difference in bot activity for the three smallest communities may come from factors such as some very dynamic automated news pages.

6.2.1. Limitations of Automatic Bot Detection

It is important to keep in mind that bot detection, while being effective, is a limited approach, especially in the case of political elections. While a lot of bots can be found, a nuance is to be made, as bots in a cluster are not necessarily promoting the candidate. Some of the more basic bots that promote cryptocurrency or fishing sites usually simply post the same messages repeatedly, along with all the popular hashtag at a time t . Such a behavior artificially inflates the number of bots we find in the community of candidates that are naturally more active on twitter. The same goes for "automatized behavior", as we see in figure 3 and in our feature importance section. While algorithms are accurate, their most discriminative features are the number of statuses, followers, and the age of the user. The situation can be more complicated in the case of politics, as there are actual people who are willing to tweet with the hashtag #MélenchonPrésident one hundred times a day simply because they are extremely passionate about the campaign.

7. Conclusion and Future Work

In this paper, we have leveraged graph-based community detection methods to gather insights into each of the most significant candidates' online campaigns for the 2022 French presidential election. We have been able to build a portrait of the average voter for each can-




Figure 3: Heatmap of Pearson correlation between features and labels.)

didate, the interest they carry in different political subjects, their geolocation, and their language habits. We have also presented results on the usage of automated accounts, or the lack thereof, in each community.

Many future tasks are possible to be performed on this dataset. Based on the community detection of political communities, a relevant study could be to analyze of the impact of major political events or debates. Considering that we have collected tweets from February to April on a daily basis, we could quantify the shift in the communities after debates between two candidates or how the start of the Ukraine war influenced electors. In the same way, we could also investigate the shift between the two turn of votes. French elections are based on a two-turn system, with the first turn aiming at narrowing down the the list of candidates and only keeping the two largest ones. The continued gathering of data and community detection could show us which communities turn to which candidate during the period between the two turns and how their language habits evolve.

Acknowledgment

This research was supported by the ANR chair AML/HELAS (ANR-CHIA-0020-01).

8. Bibliographical References

- Ali Alhosseini, S., Bin Tareaf, R., Najafi, P., and Meinel, C. (2019). Detect me if you can: Spam bot detection using inductive representation learning. In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, page 148–153, New York, NY, USA. Association for Computing Machinery.
- Badawy, A., Ferrara, E., and Lerman, K. (2018). 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)*, pages 258–265. IEEE.
- Bae, J.-h., Son, J.-e., and Song, M. (2013). Analysis of twitter for 2012 south korea presidential election by text mining techniques. *Journal of Intelligence and Information Systems*, 19(3):141–156.

- Bermingham, A. and Smeaton, A. (2011). On using twitter to monitor political sentiment and predict election results. In *Proceedings of the Workshop on Sentiment Analysis where AI meets Psychology (SAAIP 2011)*, pages 2–10.
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, oct.
- Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., Quinn, S. C., and Dredze, M. (2018). Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate. *American Journal of Public Health*, 108(10):1378–1384. PMID: 30138075.
- Choy, M., Cheong, M. L., Laik, M. N., and Shung, K. P. (2011). A sentiment analysis of singapore presidential election 2011 using twitter data with census correction. *arXiv preprint arXiv:1108.5520*.
- Cinelli, M., Cresci, S., Galeazzi, A., Quattrociocchi, W., and Tesconi, M. (2020). The limited reach of fake news on twitter during 2019 european elections. *PLoS one*, 15(6):e0234689.
- Cresci, S., Lillo, F., Regoli, D., Tardelli, S., and Tesconi, M. (2018). Cashtag piggybacking: uncovering spam and bot activity in stock microblogs on twitter. *CoRR*, abs/1804.04406.
- Davis, C. A., Varol, O., Ferrara, E., Flammini, A., and Menczer, F. (2016). Botornot: A system to evaluate social bots. *CoRR*, abs/1602.00975.
- Deb, A., Luceri, L., Badawy, A., and Ferrara, E. (2019). Perils and challenges of social media and election manipulation analysis: The 2018 US midterms. *CoRR*, abs/1902.00043.
- Diakopoulos, N. A. and Shamma, D. A. (2010). Characterizing debate performance via aggregated twitter sentiment. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1195–1198.
- Echeverría, J., Cristofaro, E. D., Kourtellis, N., Leontiadis, I., Stringhini, G., and Zhou, S. (2018). LOBO - evaluation of generalization deficiencies in twitter bot classifiers. *CoRR*, abs/1809.09684.
- Feng, S., Wan, H., Wang, N., Li, J., and Luo, M. (2021). Twibot-20: A comprehensive twitter bot detection benchmark. *CoRR*, abs/2106.13088.
- Ferrara, E. (2017). Disinformation and social bot operations in the run up to the 2017 french presidential election. *CoRR*, abs/1707.00086.
- Gayo-Avello, D., Metaxas, P., and Mustafaraj, E. (2011). Limits of electoral predictions using twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5, pages 490–493.
- Grimminger, L. and Klinger, R. (2021). Hate towards the political opponent: A twitter corpus study of the 2020 us elections on the basis of offensive speech and stance detection. In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 171–180.
- Guo, Y., Rennard, V., Xypolopoulos, C., and Vazirgianis, M. (2021). BERTweetFR : Domain adaptation of pre-trained language models for French tweets. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 445–450, Online, November. Association for Computational Linguistics.
- Kudugunta, S. and Ferrara, E. (2018). Deep neural networks for bot detection. *CoRR*, abs/1802.04289.
- Lee, K., Eoff, B., and Caverlee, J. (2021). Seven months with the devils: A long-term study of content polluters on twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1):185–192, Aug.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Martin, L., Muller, B., Ortiz Suárez, P. J., Dupont, Y., Romary, L., de la Clergerie, É., Seddah, D., and Sagot, B. (2020). CamemBERT: a tasty French language model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online, July. Association for Computational Linguistics.
- Mazza, M., Cresci, S., Avvenuti, M., Quattrociocchi, W., and Tesconi, M. (2019). Rtbust: Exploiting temporal patterns for botnet detection on twitter. *CoRR*, abs/1902.04506.
- Memon, S. A. and Carley, K. M. (2020). Characterizing COVID-19 misinformation communities using a novel twitter dataset. *CoRR*, abs/2008.00791.
- Miller, Z., Dickinson, B., Deitrick, W., Hu, W., and Wang, A. (2014). Twitter spammer detection using data stream clustering. *Information Sciences*, 260:64–73, 03.
- Ousidhoum, N., Lin, Z., Zhang, H., Song, Y., and Yeung, D.-Y. (2019). Multilingual and multi-aspect hate speech analysis. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4675–4684, Hong Kong, China, November. Association for Computational Linguistics.
- Pastor-Galindo, J., Zago, M., Nespoli, P., Bernal, S. L., Celdrán, A. H., Pérez, M. G., Ruipérez-Valiente, J. A., Pérez, G. M., and Mármol, F. G. (2020). Spotting political social bots in twitter: A use case of the 2019 spanish general election. *IEEE Transactions on Network and Service Management*, 17(4):2156–2170.
- Petrova, M., Sen, A., and Yildirim, P. (2021). Social media and political contributions: The impact

- of new technology on political competition. *Manag. Sci.*, 67:2997–3021.
- Radicioni, T., Saracco, F., Pavan, E., and Squartini, T. (2021). Analysing twitter semantic networks: the case of 2018 italian elections. *Scientific Reports*, 11(1):1–22.
- Rill, S., Reinel, D., Scheidt, J., and Zicari, R. V. (2014). Politwi: Early detection of emerging political topics on twitter and the impact on concept-level sentiment analysis. *Knowledge-Based Systems*, 69:24–33.
- Risch, J., Ruff, R., and Krestel, R. (2020). Offensive language detection explained. In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 137–143, Marseille, France, May. European Language Resources Association (ELRA).
- Seidman, S. (1983). Network structure and minimum degree. *Soc. Netw.* 5:269–287. *Social Networks*, 5:269–287, 09.
- Shao, C., Ciampaglia, G. L., Varol, O., Flammini, A., and Menczer, F. (2017). The spread of fake news by social bots. *CoRR*, abs/1707.07592.
- Siegel, A. A., Nikitin, E., Barberá, P., Sterling, J., Pullen, B., Bonneau, R., Nagler, J., Tucker, J. A., et al. (2021). Trumping hate on twitter? online hate speech in the 2016 us election campaign and its aftermath. *Quarterly Journal of Political Science*, 16(1):71–104.
- Subrahmanian, V. S., Azaria, A., Durst, S., Kagan, V., Galstyan, A., Lerman, K., Zhu, L., Ferrara, E., Flammini, A., Menczer, F., Waltzman, R., Stevens, A., Dekhtyar, A., Gao, S., Hogg, T., Kooti, F., Liu, Y., Varol, O., Shiralkar, P., Vydiswaran, V. G. V., Mei, Q., and Huang, T. (2016). The DARPA twitter bot challenge. *CoRR*, abs/1601.05140.
- Tumasjan, A., Sprenger, T., Sandner, P., and Welp, I. (2010). Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 4, pages 178–185.
- Utz, S. (2009). The (Potential) Benefits of Campaigning via Social Network Sites. *Journal of Computer-Mediated Communication*, 14(2):221–243, 01.
- Varol, O., Ferrara, E., Davis, C. A., Menczer, F., and Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. *CoRR*, abs/1703.03107.
- Yang, K., Varol, O., Davis, C. A., Ferrara, E., Flammini, A., and Menczer, F. (2019a). Arming the public with AI to counter social bots. *CoRR*, abs/1901.00912.
- Yang, K., Varol, O., Hui, P., and Menczer, F. (2019b). Scalable and generalizable social bot detection through data selection. *CoRR*, abs/1911.09179.
- Yaqub, U., Sharma, N., Pabreja, R., Chun, S. A., Atluri, V., and Vaidya, J. (2020). Location-based sentiment analyses and visualization of twitter election data. *Digit. Gov.: Res. Pract.*, 1(2), apr.