

Analyzing the Behaviour of Twitter Bots in Post Brexit Politics

Bello Shehu Bello, Reiko Heckel

Department of Informatics, University of Leicester, UK

bsbello.cs@buk.edu.ng, reiko@mcsle.ac.uk

Abstract—Twitter bots are automated user accounts widely used in political campaigns to promote opinions and attack opponents. In this paper, we study the behaviour of bots in the Brexit debates after the 2016 referendum, including on issues such as a second Brexit referendum and Scottish independence. Our findings, which are useful for understanding the role of bots in political debates more generally, encourage further research on the effects of bots on public opinion.

We collected our data over a set of Brexit related hashtags and searched for bot accounts by selecting strongly opinionated accounts, then analysed their strategies, intended influence and policy position on Brexit issues. There are more than 1,962 bot accounts currently engaged in the Brexit debate. Using a novel approach for reverse engineering twitter bots [5] we show how to uncover the bots' implemented strategies. We found that @StillYesScot, @IsThisABot and @FAO_Scotbot, all bots that promote Scottish independence, use similar strategies. Such an analysis can contribute to identifying their political background and affiliation.

Our differential sentiment analysis shows how the bots attempt to influence, constructively or destructively, the opinions of users connected to their network. Our geolocation analysis shows that the bots propagate their information to the correct target audience, with the majority of their followers living in the UK.

Index Terms—Twitter bots, Brexit, Scottish independence

I. INTRODUCTION

Social media platforms such as Twitter are now widely used by political figures and government officials to make announcements and reach out their supporters [21]. This greatly increases the wider acceptance of the medium. Twitter and other social media platforms play an important role in steering public participation in social policy and public activities [3], [8], [16]. The use of social media in social and political campaigns was examined by many studies, including on the online mobilisation of protest [25], in the Occupy Wall Street movement [11] and other political campaigns [32], [33]. The extent to which social media is used in this space and the acceptance of the public to actively participate in discussions on social media mean that this channel is used for spreading propaganda with the aim of manipulating public opinion [26]. This use of social media was reported in [13], [15], [30] and found to be effective in influencing public opinion [2], [4], [12]. However, the rise of bots changes how politicians use social media. Political actors employ social bots to engage in political conversations [19], [26]. This offers them an elevated opportunity for promoting their agenda. Social bots are automated user accounts programmed to emulate the activity of human users, automatically interacting with other users

and promoting a specific agenda [18], [23], [31]. The role of social bots in manipulating public opinion became apparent in 2010 during the US midterm elections, where they were used to spread thousands of tweets intended to damage political opponents [30]. Similarly, bots were used in the 2016 US presidential elections [3], [6], the UK Brexit referendum [26], and campaigns in France [17], Australia [22], Norway [28] and Venezuela [19]. Political actors, organizations and other entities with adequate resources can deploy thousands of bot accounts to support or attack certain opinion [3]. Several approaches have been developed for the detection of bots [9], [14], [23], [31], the majority using the ability to distinguish features of a typical human user based on metadata, e.g., the length or auto-generated nature of screen names, the temporal distribution or volume of tweets, changes in sentiment, network and lexical features [23], [27], [31]. This helped the detection and blocking of thousands of bots accounts by Twitter. However, we argued [5] that detection of automated accounts without considering their detailed behaviour can endanger thousands of legitimated automated accounts, where bots are used for tasks such as delivering news updates, weather alerts, advertisements of products and services [18], [26].

Reverse engineering the behaviour and strategies of these accounts has proven a challenging research area [1], [14], [20], [31], but can eventually lead to techniques to identify their masters [3]. Recently, [5] proposed an approach for reverse engineering the behaviour of Twitter bots. In this paper, we utilize the reverse engineering approach to study the strategies of bots promoting positions on Scottish independence and Brexit.

A. Research Questions

We aim to answer the following questions regarding the use of bots after the UK vote to leave the EU. They are prototypical of many political debates and can serve as a template for the use of our approach in their analyses.

- Are there bots promoting post Brexit referendum issues? If yes, what are their strategies? In order to understand how bots carry out their campaigns and to infer their potential effects, we will use our reverse engineering approach to study strategies of bots to propagate information.
- What are their policy positions on specific issues, in addition to what they are generally known for? Here

we investigate whether bots engage in discussing issues beyond their core purpose and discover their respective positions. This is important to understand their scope and detect attempted cross-fertilization between topics.

- Who are bots trying to influence, and how? While it is difficult to evidence actual impact of bots on their target audience (e.g., due to external factors such as news and political events) we will analyze intended influence and assess possible impact.
- Do bots focus their messages on users from certain regions? We will use the geolocation of bots' target audiences to gain insight into regional factors.

B. Organisation

We collected Twitter data over a set of Brexit related hashtags from 4th of March to 9th of May 2019 and analysed bot accounts. We obtained 2,520,663 tweets from 143,332 distinct users. Details about the data and the collection are given in Section II-A. We proposed an effective method to identify strongly opinionated accounts in Section II-C. Despite announcements such as "Twitter is sweeping out fake accounts like never before"¹ in June 2018, there are still many bot accounts. We found a total of 1,962 bots in a sample of 7,000 accounts of our 143,332 user accounts. In Section III-B we discuss how @StillYesScot, a bot which promotes Scottish independence, has strongly engaged in promoting tweets which oppose a Brexit deal. This is important to understanding its role during the UK MPs' deliberations on Brexit.

In Section III-C we analyse the potential impact of bots on their networks (friends and followers). We find that the intensity at which @StillYesScot argues against a Brexit deal and promotes a second referendum is higher than with the average users on its network. This provides information about its role in the network and signals how it might want to influence them.

Finally, In Section III-A we use the reverse engineering approach [5] to uncover the bots' implemented strategies. We find that @StillYesScot, @IsThisABot and @FAO_Scotbot, all bots promoting Scottish independence, use very similar strategies. In Section IV we conclude that further monitoring of these accounts will lead to the identification of their masters and more importantly their impact on other issues currently under deliberation following the 2016 UK referendum on EU membership.

C. Summary of Contributions

- We analyse tweets on post Brexit politics and show an effective way to identify strongly opinionated accounts.
- Using bot detection mechanism, we find bots that promote Scottish independence.
- Using a novel approach for reverse engineering Twitter bots [5] we uncover their implemented strategies. This

reveals their sources of information to promote their agenda, helping to understand their impact.

- Using the reverse engineering technique, we identify bots that share the same strategies. This can contribute to identifying their masters if continuously monitored.
- Using sentiment analysis on their content, we analyse the bots' policy position on other issues and find that, apart from promoting Scottish independence, bots are actively involved in the discussion of a Brexit deal and a second referendum.
- We study the impact of bots on their network and identify bots in a position of high potential influence. This encourages further research on these accounts.
- Our analysis of regional factors shows that most followers are from the UK (many from Scotland). This shows us that the bots are adequately targeted.

II. METHODOLOGY

Bots are often created to promote a specific agenda. They are deployed with goals, targets and strategies to achieve the goals. We intend to study the bots' campaign strategies to understand their roles and potentially understand their impact. Following the steps described below, we use a tool² based on our approach for reverse engineering bots [5] to study the bots' strategies. Given a bot account, the tool analyses user features, content, network and sentiment features of tweets promoted and/or reacted to by a bot. The tool constructs a decision tree-based machine learning model from which it extracts rules describing the strategies used by the bot to tweet, retweet, reply, favourite and follow other users. Depending on the bot and the type of action, the rules provided by the tool reveal vital information about the nature of the content used by the bots (e.g., hashtags, topics), the sources of its contents, its masters, target users and their locations.

A. Data collection

We manually set a list of hashtags related to Brexit and Scottish independence. We include twitter handles of the leaders of both the Conservative and the Labour parties. To form a comprehensive list, first we crawled data for 3 days, analyse the most frequent hashtags used by online users and then updated our list of search terms. The search terms used for the data collection are: *brexit*, *brexit Second referendum*, *#NoDealBrexit*, *Brexit deal*, *Cancel Brexit*, *No Deal Brexit*, *#scottishindependence*, *Indyref2*, *scotref*, *scottish independence*, *#euref*, *@theSNP*, *#Brexit*, *#RoadtoBrexit*, *@NicolaSturgeon*, *@theresa_may*, *@Conservatives*, *@jeremycorbyn*, *@UKLabour*. We obtained a total of 2,520,663 tweets from 143,332 distinct users by querying the Twitter search API from 4th of March to 9th of May 2019. We chose to use the Twitter search API to ensure that we obtain all tweets related to the search terms rather than a sample of unfiltered tweets provided in real time by the streaming API. This is to avoid issues reported in [29] with collecting data using the Twitter Stream API.

¹https://www.washingtonpost.com/technology/2018/07/06/twitter-is-sweeping-out-fake-accounts-like-never-before-putting-user-growth-risk/?utm_term=.96abcb0330f0

²<https://github.com/bellobichi2/botscope>

B. Bot detection

We used the Python API of Botometer³ to check for accounts that are likely to be bots. Botometer [14] is a service available to check the likelihood of an account to be a bot. The Botometer API uses the Twitter API to extract the top 300 tweets from a given account, analysing its content, temporal and network features to produce a bot score. Since the Botometer API incurs limitations imposed by the Twitter API⁴, it is difficult to test all user accounts. Instead, we obtained a bot score for the top 7,000 most active Twitter accounts ranked by volume of tweets in our dataset. A score of 50% has been shown to be effective to label an account as a bot [3], [6]. In this way we detect a total of 1,962 potential bot accounts out of the 7,000 accounts.

C. Identification of strongly opinionated accounts

We expect that strongly opinionated accounts would produce a high volume of tweets with an intense opinion. First, we use topic analysis to find the top five topics in the dataset, then analyse the sentiment intensity of the accounts over those topics. For the topic analysis, we used a topic model based on latent Dirichlet allocation (LDA) to compute the top five topics [7], [34]. Before applying the topic model, we convert the tweets into a corpus of aggregated bi-grams. The use of bi-grams instead of uni-grams allows us to keep the correlation between the words and obtain good topics. For the sentiment analysis, we used VADER (Valence Aware Dictionary and sEntiment Reasoner) to compute the sentiment intensity of the tweets. VADER is a sentiment analysis model designed specifically for social media text [24]. To identify the most opinionated accounts, we compute the absolute sentiment of each account for each topic. Fig. 1 is a visualization of the accounts based on their absolute sentiments and the volume of tweets produced. Our idea of using absolute sentiment (the sum of negative and positive sentiments) pushes all the less opinionated accounts on to the zero line while those with high polarity stand out. Since the majority of the accounts tend to be less opinionated, the further analysis concentrates on the top 100 accounts.

Then, we investigate the relationship between accounts that promote opinions in favour or against a Brexit deal with those that are calling for a second referendum. Fig. 2 is a visualization of these accounts. We notice that @StillYesScot is clearly against a Brexit deal while promoting Scottish Independence. @EdinburghWatch supports a Brexit deal and @AnalyticaGlobal tweets against it. @_Max_Baring_ is featured as an account that is against a Brexit deal and supports a second referendum. We can also notice that, while @botanic_my promotes a Brexit deal and a second referendum it is against Scottish Independence.

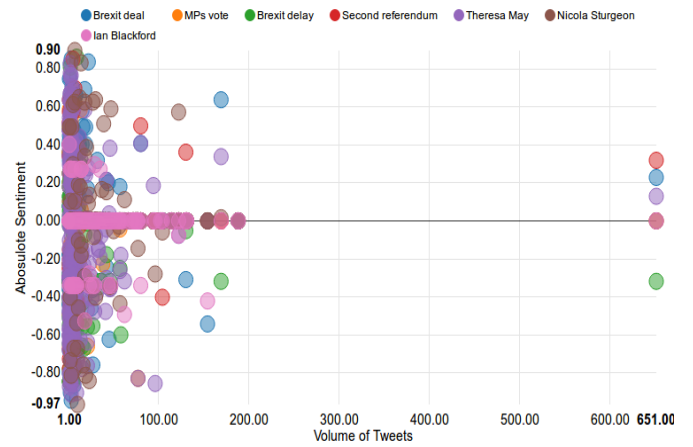


Fig. 1: Opinion polarization of the accounts

III. RESULTS AND DISCUSSION


A. Implemented Strategies

Since the recent success in the detection of bots [10], [14], [31] the most challenging task is understanding their strategies, targets and influence [18], [26]. Recently, [5] proposed a rule-based approach for reverse engineering Twitter bots to understand their behaviour. The behaviour (actions) are represented in the form of rules encoded with sentiments. The rule describes the behaviour of an account in a simple way. The account @StillYesScot is known for promoting Scottish independence from its account description, but how does it do this? We utilise the approach proposed in [5] to understand its implemented strategies. Given a bot account, the approach analyses traces of the bot's actions, builds a machine learning model based on the decision tree algorithm, extracts rules from the tree and visualizes the behaviour in the form of rules encoded with sentiments. Fig. 3a shows the result of learning the rule on @StillYesScot. Apart from understanding the underlying construct of its retweet action, we see the main sources of tweets propagated by the account, and their hashtags: The account is retweeting positive tweets with hashtags #indref2 and #ScotRef, both promoting Scottish independence. The tweets contain photos and URLs of the Guardian and Scotsman newspapers. Analysing URLs can help us assess the credibility of the news an account is propagating based on their sources.

Using the reverse engineering technique we found other bots which have similar behaviour as @StillYesScot. In Fig 3 we show the strategies used by @StillYesScot, @IsThisABot and @FAO_Scotbot. They all use similar hashtags (#indref2, #DisolveTheUnion) and nearly the same rule construct. This suggests that they may have the same political masters. Further monitoring of these accounts through periodic analysis of rules describing their strategies could lead to their identification, for example, if we found a rule indicating that these bots are all retweeting from a particular Twitter account which is known to belong to the leader of the campaign they are supporting.

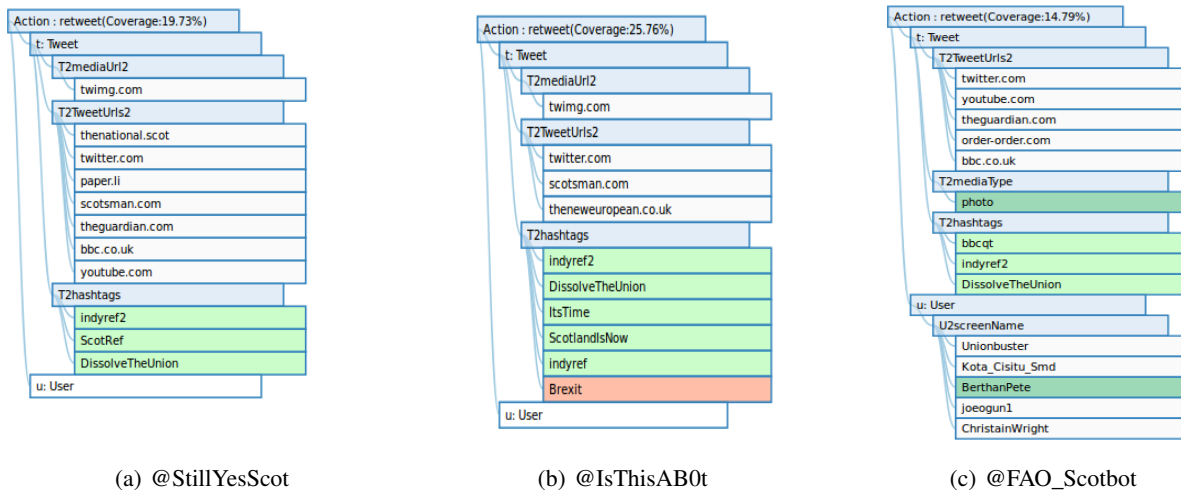
³<https://botometer.iuni.iu.edu!/api>

⁴<https://developer.twitter.com/en/docs/basics/rate-limits>



(a) Opinion polarization of top 100 accounts on the Brexit deal, Second referendum and Scottish independence (b) Opinion polarization of top 100 accounts on Scottish independence only

Fig. 2: Opinion polarization of top 100 accounts on the Brexit deal, Second referendum and Scottish independence



(a) @StillYesScot

(b) @IsThisAB0t

(c) @FAO_Scotbot

Fig. 3: Retweets strategy

B. Policy positions on specific issues

We investigated the policy position of the accounts on specific issues beyond what they are popularly known for by analyzing their sentiments on other related topics. This is to understand their scope and attempted influence in related areas. For example, we asked about the position of @StillYesScot on the Brexit deal and the relevant votes by MPs. Fig. 4 shows absolute sentiments of @StillYesScot on these topics. The result shows that it is against the Brexit deal more strongly than it promotes Scottish independence. This could be because the Brexit deal was the major issue on social media during that period. We also looked at its sentiments on those topics over time. Fig. 5 shows the result where between 20th February and 7th March, a period when the UK MPs began their deliberations on the Brexit deal, the account campaigns strongly against the deal while from the 2nd to the 7th of March the account continues with its usual campaign on Scottish independence.

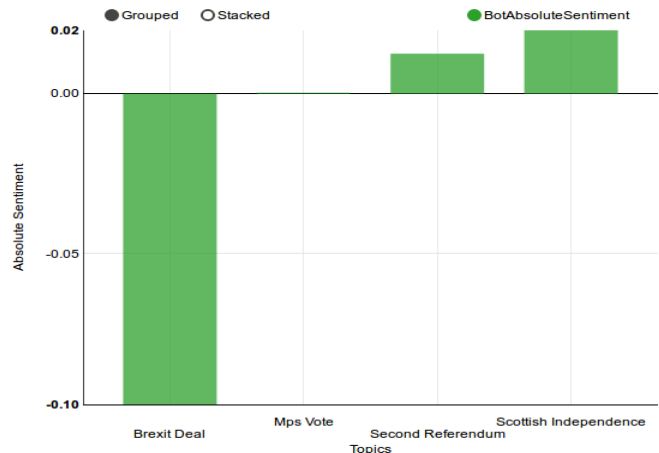


Fig. 4: @StillYesScot policy on Brexit deal, MPs vote, Second referendum and Scottish Independence




Fig. 5: @StillYesScot policy overtime




Fig. 7: Opinion of @StillYesScot Vs Its followers

C. Bias and Attempt of Influence

Social connections can exist between parties of the same or different opinions. While the former is more common in a physical social network, in an online social network the latter can be considered as an attempt by one party to influence the other. We analyse the opinion of users (friends) followed by the bots to investigate an attempt of influence. Fig. 6 shows the sentiments of @StillYesScot versus its friends. While we can not establish actual influence due to other external factors, we notice that among the users followed by the bot there are some who support the Brexit deal while the bot is known to be against it. This could be an attempt to share its negative opinion with these users while promoting a second referendum and Scottish independence.

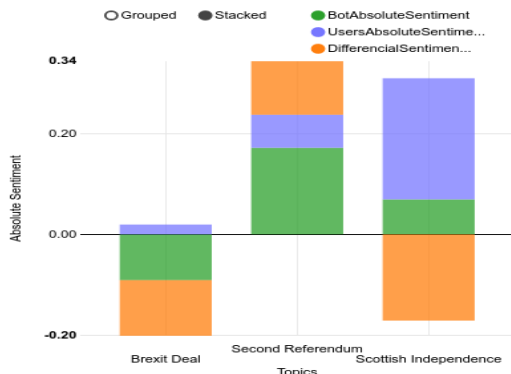


Fig. 6: Opinion of @StillYesScot Vs Its Friends

We also analysed the sentiments of tweets produced by the bots vs that of their followers to investigate the role of the bots in their networks. Fig. 7 shows that the absolute sentiment of @StillYesScot on the Brexit deal and a second referendum is higher than that of all the users in the network. This is a sign that tries to exert impact on other users in this direction.

D. Regional targets

We analysed the geographical locations of users connected to the accounts to discover if the accounts are trying to

influence or favour users from certain regions. To identify the users' location, we utilised the account location provided by the users in their Twitter profile. Since this is mostly noisy and some users provide an imprecise location, we took an extra pre-processing step, removed special characters, tokenized the text and used the Google location API with geocoder⁵, a Python geocoding library, to identify a valid country and/or city from the given text. Through this process we identify the country names of more than 90% of the users. While we did not find a strong geo-location policy attached to the accounts, it is worth mentioning that the majority of their followers and friends are from the UK, many from the cities of Scotland (Glasgow, Edinburgh and Aberdeen). This is an indication that they are disseminating their information to the correct target community.

IV. CONCLUSION

Social media platforms such as Twitter play an important role in the participation of the public on social issues and other democratic activities. Twitter has now become a prominent channel used by TV stations to show audiences' reactions in near real-time during political debates [28]. However, the rate at which social media is used for political campaigns and the use of social bots to promote political agendas raises concerns about the possibility of manipulating public opinion using bots.

In this paper, we studied the behaviour of bots in post referendum Brexit campaigning, addressing issues such as a second referendum and Scottish independence. We collected tweets posted between 4th of March to 9th of May 2019 based on a manually compiled list of post Brexit related keywords and hashtags. Using a bot detection method, we found that there are more than 1,962 bots currently engaged with our issues of interest. Using a novel approach for reverse engineering Twitter bots, we have shown how to uncover the bots' implemented strategies. We found that @StillYesScot, @IsThisAB0t, @FAO_Scotbot employed similar strategies.

⁵<https://pypi.org/project/geocoder/>

This could lead to the identification of their masters if continuously monitored. Our influence and geolocation analysis shows us that the bots are in a good position of influencing, constructively or destructively, the opinions of their audience.

Social media has been proven to be an effective tool to influence public opinions [2], [4], [12]. Social bots can have severe negative consequences on public opinion. They can enhance the spread of false information, cause confusion and polarize political conversation [3]. We plan to explore more details on targets and strategies used by social bots to carry out their campaigns and alert the general public. Concluding, our analysis encourages further research to monitor the use of bots for manipulating public opinion on other issues currently under deliberation. This could provide a powerful tool to help political researchers, journalists and government create transparency in the use of social bots in public debate.

REFERENCES

- [1] Abdulrahman Alarifi, Mansour Alsaleh, and AbdulMalik S. Al-Salman. Twitter turing test: Identifying social machines. *Inf. Sci.*, 372:332–346, 2016.
- [2] Sinan Aral, Lev Muchnik, and Arun Sundararajan. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549, 2009.
- [3] Adam Badawy, Emilio Ferrara, and Kristina Lerman. Analyzing the digital traces of political manipulation: The 2016 russian interference twitter campaign. In *IEEE/ACM 2018 International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2018.*, pages 258–265, 2018.
- [4] Eytan Bakshy, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. Everyone’s an influencer: quantifying influence on twitter. In *Proceedings of the Forth International Conference on Web Search and Web Data Mining, WSDM 2011.*, pages 65–74, 2011.
- [5] Bello Shehu Bello, Reiko Heckel, and Leandro Minku. Reverse engineering the behaviour of twitter bots. In *2018 Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pages 27–34. IEEE, 2018.
- [6] Alessandro Bessi and Emilio Ferrara. Social bots distort the 2016 us presidential election online discussion. *first monday*, 21(11), 2016.
- [7] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- [8] Robert M Bond, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295, 2012.
- [9] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. On-demand bot detection and archival system. In *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017*, pages 183–187, 2017.
- [10] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. Temporal patterns in bot activities. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 1601–1606. International World Wide Web Conferences Steering Committee, 2017.
- [11] Michael D. Conover, Emilio Ferrara, Filippo Menczer, and Alessandro Flammini. The digital evolution of occupy wall street. *CoRR*, abs/1306.5474, 2013.
- [12] Michael D Conover, Bruno Gonçalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. Predicting the political alignment of twitter users. In *2011 IEEE third international conference on social computing*, pages 192–199. IEEE, 2011.
- [13] Michael D Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*, 2011.
- [14] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. Botnot: A system to evaluate social bots. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11-15, 2016, Companion Volume*, pages 273–274, 2016.
- [15] Sara El-Khalili. Social media as a government propaganda tool in post-revolutionary egypt. *First Monday*, 18(3), 2013.
- [16] Gunn Sara Enli and Eli Skogerbø. Personalized campaigns in party-centred politics: Twitter and facebook as arenas for political communication. *Information, communication & society*, 16(5):757–774, 2013.
- [17] Emilio Ferrara. Disinformation and social bot operations in the run up to the 2017 french presidential election. *First Monday*, 22(8), 2017.
- [18] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. The rise of social bots. *Communications of the ACM*, 59(7):96–104, 2016.
- [19] Michelle Forelle, Philip N. Howard, Andrés Monroy-Hernández, and Saiph Savage. Political bots and the manipulation of public opinion in venezuela. *CoRR*, abs/1507.07109, 2015.
- [20] Carlos A. Freitas, Fabrício Benevenuto, Saptarshi Ghosh, and Adriano Veloso. Reverse engineering socialbot infiltration strategies in twitter. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2015.*, pages 25–32, 2015.
- [21] Jennifer Fromm, Stefanie Melzer, Björn Ross, and Stefan Stieglitz. Trump versus clinton: Twitter communication during the US primaries. In *4th European Network Intelligence Conference, ENIC 2017, Duisburg, Germany.*, pages 201–217, 2017.
- [22] Rachel K Gibson and Ian McAllister. Does cyber-campaigning win votes? online communication in the 2004 australian election. *Journal of Elections, Public Opinion and Parties*, 16(3):243–263, 2006.
- [23] Zafar Gilani, Reza Farahbakhsh, Gareth Tyson, Liang Wang, and Jon Crowcroft. Of bots and humans (on twitter). In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pages 349–354. ACM, 2017.
- [24] CJ Hutto Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International Conference on Weblogs and Social Media (ICWSM-14)*, 2014.
- [25] Sandra González-Bailón, Javier Borge-Holthoefer, and Yamir Moreno. Broadcasters and hidden influentials in online protest diffusion. *CoRR*, abs/1203.1868, 2012.
- [26] Philip N Howard and Bence Kollanyi. Bots,# strongerin, and# brexit: computational propaganda during the uk-eu referendum. 2016.
- [27] Isa Inuwa-Dutsea, Bello Shehu Bello, Ioannis Korkontzelos, and Heckel Reiko. The effect of engagement intensity and lexical richness in identifying bot accounts on twitter. *IADIS International Journal on WWW/Internet*, 16(2), 2018.
- [28] Bente Kalsnes, Arne H. Krumsvik, and Tanja Storsul. Social media as a political backchannel: Twitter use during televised election debates in norway. *Aslib J. Inf. Manag.*, 66(3):313–328, 2014.
- [29] Fred Morstatter, Jürgen Pfeffer, Huan Liu, and Kathleen M. Carley. Is the sample good enough? comparing data from twitter’s streaming API with twitter’s firehose. In *Proceedings of the Seventh International Conference on Weblogs and Social Media, ICWSM 2013.*, 2013.
- [30] Jacob Ratkiewicz, Michael D. Conover, Mark R. Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. Detecting and tracking political abuse in social media. In *Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Catalonia, Spain, July 17-21, 2011*, 2011.
- [31] VS Subrahmanian, Amos Azaria, Skylar Durst, Vadim Kagan, Aram Galstyan, Kristina Lerman, Linhong Zhu, Emilio Ferrara, Alessandro Flammini, Filippo Menczer, et al. The darpa twitter bot challenge. *arXiv preprint arXiv:1601.05140*, 2016.
- [32] Zeynep Tufekci and Christopher Wilson. Social media and the decision to participate in political protest: Observations from tahrir square. *Journal of communication*, 62(2):363–379, 2012.
- [33] Onur Varol, Emilio Ferrara, Christine L Ogan, Filippo Menczer, and Alessandro Flammini. Evolution of online user behavior during a social upheaval. In *Proceedings of the 2014 ACM conference on Web science*, pages 81–90. ACM, 2014.
- [34] Xiaohui Yan, Jiafeng Guo, Yanyan Lan, and Xueqi Cheng. A bitern topic model for short texts. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1445–1456. ACM, 2013.