Political Bots and the Swedish General Election

Johan Fernquist and Lisa Kaati
Decision Support Systems
Swedish Defence Research Agency
Kista, Sweden
firstnamne.lastname@foi.se

Ralph Schroeder
Oxford Internet Institute
University of Oxford
Oxford, England
ralph.schroeder@oii.ox.ac.uk

***Abstract*—In this paper we present a study of political bots on Twitter and their influence on the Swedish general election in September 2018. We have not restricted our study to bots that are a software program, instead we are interested in any type of bot-like automated behaviour. This includes a human being manually copying or retweeting content repeatedly in a robot-like way with the aim of influencing the interaction between a user and content or with other users.**

We examined how bots influence discussions about the Swedish general election on Twitter during a three month period starting half a year before the election. For this purpose we trained a machine learning model to recognize automatic behaviour and studied accounts that tweet about the Swedish election. Our results show, among other things, that bot activity is at a similar level to that found in other studies, and that distinctive clusters, for example among right-wing Twitter accounts, can be identified.

***Index Terms*—bots, bot network, bot detection, social media, Twitter, election, political bots**

I. Introduction

Bots have been used for a variety of purposes. While they were initially designed to automate otherwise unwieldy online processes which could not be done manually, they have come to be most commonly used for commercial purposes such as directing internet users to advertisements and the like. Bots are also often used to further illegal activity such as collecting data from users for criminal gain.

There are also other uses of bots for practical and benign purposes, such as checking content on Wikipedia pages for spelling mistakes or malicious activities. Political bots can similarly be used for a variety of benign and malicious purposes, though the latter have understandably received the bulk of attention. A recent paper [11] defines political bots as automated scripts designed to influence public opinion. The authors point out that political bots have been used for various purposes, including to spread misinformation, inflate the visibility of political actors or issues, or spread junk news simply to create an environment of informational uncertainty. Their known history goes back to around 2010, and they have been used by many actors. Roberts points out in her analysis of Chinese disinformation efforts [16] that the practice of spreading rumours or other forms of disinformation on a large scale to distract or divert attention what she calls flooding is not limited to the Chinese regime, but it has also become a more prominent tool for political actors in democracies. It is difficult to counteract this process, as it is not clear, for example in the US, how they fall under existing legislation regarding campaign finance. Another difficulty here is identifying their source, for example, if political bots originate in a foreign country: do they fall into the category of foreign electoral interference? These issues, though they had been discussed most widely for the US, affect political bots more broadly, and also outside of elections.

A. Bots and bot detection

There are many different definitions of 'bots'. In [11] bots have been defined as executable software that automates the interaction between a user and content or other users. In [9] a typology of bots is presented. The topology suggests six different types of bots:

- web robots (crawlers and scrapers)
- chatbots (human-computer dialog system which operates through natural language via text or speech)
- spambots (bots that post on online comment sections and spread advertisements or malware on social media platforms)
- social bots (various forms of automation that operates on social media platforms)
- sock puppets and trolls (fake identities used to interact with ordinary users on social networks)
- cyborgs and hybrid accounts (a combination of automation and of human curation)

Web robots do not interact with users on a social platform and are therefore considered to be different from automated social media accounts. Social bots are bots that generally act in ways that are similar to how a real human may act in an online space. Social bots that are used for political purposes are called political bots. The term sock puppet refers to fake identities used to interact with ordinary users on social networks. Politically motivated sock puppets, especially when coordinated by governments or interrelated actors, are according to [9], called 'trolls'.

For the purposes of this paper, we widen the definition of bots used by [11] since we are interested in understanding any automated behaviour; in other words, not necessarily a software program. That is because, in this research, we have sought out any bot- like automated behaviour of this type, which could include a human being manually copying or retweeting content repeatedly, in a robot-like way, to achieve

978-1-5386-7848-0/18/$31.00

the same aims as a software bot (i.e., to influence the interaction between a user and content or other users). The authors of [9] suggests that fully automated social media accounts should be referred to as social bots and accounts with manual curation and control should be referred to as sock puppets. Using the typology of Gorwa and Guilbeaul, the bots we are interested in detecting and studying are both automated social bots and sock puppets.

Our widened definition of bots is necessary because our research relies on machine learning techniques to identify bots as being similar to existing bots, something that could be achieved by a human being. In the rest of this paper, we use the term *bot* when referring to accounts that have an automated behaviour, including automated social bots and sock puppets. Some estimates presented in [18] have put the proportion of bots among active Twitter accounts at between 9% and 15%.

II. Bots and the Swedish Election

There are many reasons for using bots during an election. Two possible reasons are to control the narrative during political debates and to skew the discourse. However, it is still not clear to what extent bots drive social media activity and what kind of influence they have on public opinion and election outcomes. We are interested in detecting bots automatically to get an understanding of what kind of content is distributed by bots.

When using the definition of bots that we have decided on, it is clear that bots can exist in all kinds on social media. Bots can, for example, be active in various forms of discussion forums or commentary fields. However, we only focus on studying bots on Twitter. One of the reasons for studying Twitter is because it is a widely used public forum for political discussion in Sweden, especially among journalists [10].

Specifically, the aim of our study is to:

- Develop a model that recognizes Twitter accounts with an automated behavior (bots).
- Study themes that are spread by bots.
- Study how content generated by bots is spread on Twitter.
- Reason about possible effects of Twitter accounts with an automated behavior.

III. Classification of Bots

To be able to detect accounts with automated behavior automatically, we have trained a classification model to identify accounts exhibiting automatic behavior. The classification problem is to determine if an account has an automatic behavior or not. To build a model that is able to recognize automatic behavior, labeled training data is needed. The training data consists of accounts that are either bots or *genuine* accounts. Here, a genuine account is an account that is operated by a ”normal” human being. When training our model, we use a number of different features. Our model is language independent but in this paper we have used it to classifying tweets in Swedish.

In the rest of this section, we describe how the classification model is built and what data we have used to train the model. Our work is also put into relation to previous work in the domain.

A. Related work

There have been several efforts dedicated to bot detection on Twitter. Random forest is the classification algorithm that has been proven to give the best performance for bot detection for the supervised problem when several different classifiers have been tested [12], [17], [18]. In [8] features that included user meta-features and tweet features were used when training a classification algorithm. Their results indicated that bots have more URLs in their tweets and that they have a higher *follower-friend ratio*. The terminology used is from the Twitter API, where friend indicates the number of users that the user is following, as opposed that the number of followers he or she has. In [8] it is also shown that genuine accounts get more likes on their tweets than bots.

In [3] bots and cyborgs are studied. The author states that follower-friend ratio might be a bad feature since the bots might be able to unfollow accounts which not are following them back automatically. They instead introduce text entropy as a feature to measure the similarity of the texts posted by an account with the hypothesis that a bot might have more uniform content in their tweets. Another feature that is considered is what kind of devices the different accounts are using when tweeting. Most of the genuine accounts are using the web or the mobile application while bots are using other applications such as the API. It was also noted that genuine accounts have a more complex timing behavior compared to bots and cyborgs.

In [18] a total of 1,150 different features are used to train a model that recognizes bots. One set of features that is used is time features, including the statistics of times between consecutive tweets, retweets, and mentions. The research shows that the two most informative feature types are user meta-data and content features. The content features include frequency and proportion of part-of-speech-tags (POS), number of words in a tweet and entropy of words in a tweet.

B. Data

We used a number of different datasets to train our classification model. The first dataset was originally crawled during October and November 2015 and is described in [18]. The dataset contains labeled information about 647 bot accounts and 1,367 genuine accounts. Each of the accounts has produced at least 200 tweets, of which at least 90 occurred during the crawling period. The accounts were manually annotated as bots or genuine. The annotation was based on characteristics such as profile appearance, produced content and the interaction with other profiles.

Second, we used a dataset [4] consisting of 591 bots and 1,680 genuine accounts. The genuine accounts are Italian users that through a survey accepted to be a part of the study or accounts that were regularly active for a long period. The bot accounts were bought from a bot-service provider.

The third data set is used in [8]. Four undergraduate students manually annotated the dataset. The users in the dataset were divided into four subsets depending on the number of followers. The subsets were divided into users with more than 10 million followers, users with between 900 thousand and 1,1 million followers, users with 90 thousand to 110 thousand followers and users with 900 to 1100 followers. In our study, we only use the two sets with users with 90 thousand to 110 thousand followers and users with 900 to 1,100 followers since we believe that it is unlikely that a Swedish bot account has more than 1 million followers. In total the two sets consist of 519 human accounts and 355 bot accounts.

The datasets used in [4], [8], [18] are not available in the original form. Either data is missing, or only the annotated labels of the accounts are given. Since it is not possible to obtain these datasets in their original form, we cannot use these datasets for comparing performance with previous work. These datasets were only used for training our model.

The dataset used in [6] (referred to as test set 1 by the authors of the paper) is the only dataset available in its original form. The dataset consists of 991 social spam bots and 991 genuine accounts. The genuine accounts are randomly selected from a set of more than 3,000 accounts to get a 50/50 distribution of bots and genuine accounts. This means that we do not have the same set as the authors. The bots are collected in conjunction with a mayoral election in Rome 2014 where one candidate bought 1000 automatic accounts. The bought accounts all had (stolen) profile pictures, (fake) profile description and a (fake) location. Genuine accounts were identified by sending out a question to randomly selected Twitter users. The ones that replied were considered genuine.

C. Features

In our classification model, we used a total of 140 different features. The features are divided into two different types. The first type is *User Meta Data features* where information about the characteristics of the profile such as the number of followers and friends and a total number of tweets is gathered. The second feature type is the *tweet features* that holds information about the actual content and when and how the content is posted. Similar to what is done in [3], [18] we use text entropy assuming that bots might have a simpler way of expressing themselves. We also include time features as in [18]. This includes statistics of time between consecutive tweets, retweets, and mentions but we also include statistics for the time between posted tweets containing URLs. All features are listed in Table I.

D. Classification algorithm

There have been several approaches to build classification models for bot detection. Different algorithms such as AdaBoost, logistic regression, support vector machines and naive Bayes have been tested. The best results are when using random forest and, therefore, we have decided to use random forest in our classification.

TABLE I
LIST OF THE 140 FEATURES EXTRACTED FROM EACH TWITTER USER.

Meta features	Content features
Age of account	# unique hashtags per tweet
# tweets	# unique mentions per tweet
# tweets per day	# unique Urls per tweet
Friends-account age ratio	Normalized distribution of sources
# followers	Time between tweets*
# friends	Length of tweet *
Follower-friends ratio	# unique sources
Has location	retweet-tweet ratio
Has default profile description	# hashtags per tweets
Has default profile image	# urls per tweet
# likes given	# mentions per tweet
# likes geven per # followers	# media per tweet
# likes given per # friends	# symbols per tweet
# likes per day	# retweets achieved per # tweet
Length of user name	Time between urls *
	Time between mentions *
	Time between retweets *
	# words *
	Hours of day tweeting
	Weekdays tweeting
	Normalized distribution hours tweeting
	Normalized distribution weekdays tweeting
	Normalized distribution of tweet endings
	String entropy *
	Total entropy of all tweets strings concatenated

* Statistics of an array of values (mean, median, population standard deviation, standard deviation, maximum value and minimum value).

TABLE II
PERFORMANCE MEASURE OF TRAINED MODEL

Model	Type	A	P	R	F1
Our model	supervised	0.957	0.941	**0.976**	0.958
Davis *et al.* [7]	supervised	0.734	0.471	0.208	0.288
Yang *et al.* [19]	supervised	0.506	0.563	0.170	0.261
Miller *et al.* [13]	unsupervised	0.526	0.555	0.358	0.435
Ahmed *et al.* [1]	unsupervised	0.943	0.945	0.944	0.944
Cresci *et al.* [5]	unsupervised	**0.976**	**0.982**	0.972	**0.977**

E. Model evaluation

We have used three datasets from [4], [8], [18] together with our 140 different features to train a model using random forest. The model was tested on the only dataset we managed to gather in its original form. In [6], the same dataset was used to compare the performance of other bot classification models. We included the comparison from [6] and used the same dataset to test our model. The performance (Accuracy, Precision, Recall and F1-score) is shown in Table II. The comparison includes both supervised and unsupervised models. As mentioned earlier, the 911 genuine accounts in [6] was selected randomly from a set of more than 3,000 genuine accounts. Since our genuine accounts were selected randomly, we (most likely) ended up with a slightly different testing set.

The supervised methods use cluster algorithms to identify clusters of bots. [13] clusters feature vectors with the majority of features as text features using DenStream and StremKM++ as clustering algorithms.

In [1] graph clustering on statistical features related to hashtags, URLs, mentions, and retweets are used. The feature vectors were compared to each other using Euclidean distance and then clustered using the *fastgreedy* community detection algorithm.

[5] use a bio-inspired technique for modelling of behaviour of users online with so-called *digital DNA* sequences. The sequences are string encodings of the behaviour of a user, and the sequences are then compared between the different users by measuring the longest common substring to find clusters of users.

In Table II we notice that the model from [5] performs best of all the metrics except for recall where our model performs better. Of the supervised models included in the comparison, our model performs best on all metrics.

We want to point out that we are aware that one of our training sets has the same author as our test set which might be a reason for the high accuracy that we obtain. However, we have verified that none of the users are found in both data sets.

F. Feature importance

When studying the most important features in the model that determines if an account is genuine or a bot, we can get an idea of the features that separates bots from genuine accounts. We calculated feature importances with *forests of trees* and the ten most important features are shown in Table III.

TABLE III
MOST IMPORTANT FEATURES FOR OUR BOT CLASSIFICATION

Top 10	Feature
1	# given likes per # friends
2	Followers-friends ratio
3	Maximum time between retweets
4	# retweets achieved per tweet
5	Standard deviation of time between retweets
6	Median time between retweets
7	Population standard deviation of time between retweets
8	Mean time between retweets
9	# given likes
10	# given likes per # followers

We can notice that the most important feature for determining whether an account is a bot or not is the number of likes the account has given divided by the number of friends the account has. The second most important feature is the ratio between the number of followers and friends. It was discussed earlier that this feature can sometimes be misleading since bots can be able to unfollow accounts which not are following back - which not is the case here. Several of the most important features are related to the time between retweets which is highly certain to be because the bots are busy retweeting rather than posting tweets produced by themselves.

IV. A STUDY OF THE SWEDISH ELECTION

To get an understanding of the extent to which the Swedish election is influenced by political bots, we are interested in finding out 1) the number of bots that tweet about the Swedish election and 2) what kind of messages the bots are distributing.

TABLE IV
DISTRIBUTION OF LINKS PER ACCOUNT TYPE TO DIFFERENT TYPES OF MEDIA.

Media type	Bots	Suspended	Genuine
News	39.1% (739)	34.0% (405)	40.8% (6,696)
Local news	7.1% (133)	5.1% (61)	6.91% (1,135)
Political blogs	0,4%(8)	1.1% (13)	1.0% (163)
Left-wing	0.1%(2)	0.3% (3)	1.3% (219)
Right-wing	2.7 (52)	1.9% (22)	2.6% (431)
Immigration critical	19.0% (354)	26.8% (317)	18.0% (2,952)
Social media	6.3% (115)	7.7% (91)	5.7% (932)
Other	25.3% (471)	22.9% (271)	23.7% (3,897)

We collected all tweets using hashtags such as #valet2018 and #svpol in their communication through Twitter's public streaming API. The aim is to capture accounts that are discussing Swedish politics. The dataset collected consists of 194,792 tweets from 24,930 accounts collected between March 1st and May 31, 2018. The Swedish general election took place in September 2018. Since we are not interested in accounts connected to news sites, community information etc. that are automatically generated but not trying to influence discussions, we have created a white-list of known accounts that we have manually classified as genuine. The list includes accounts of major Swedish news sites that have Twitter accounts that automatically post tweets about articles as soon as they are published on the website. To get an understanding of the number of bots present in the dataset we used our bot classification model. The results show that around 6% (1,429) of the accounts were identified as bots according to our model (not including the accounts from the news sites). The bots produced around 5% (8,954 tweets) of the content related to the Swedish election.

Something worth noting is that around 16% (3,985) of the accounts in our study are suspended. The suspended accounts produced around 6% (11,468) of the content related to the Swedish election. There are many reasons why an account may have been suspended: it could for example be the case that Twitter has suspended the account for some reason, or that the user has terminated the account by him or herself.

A. Retweets and links

If we consider the number of retweets we notice that 68% (119,074) of the tweets from genuine accounts are retweets. 63% (5,673) of the tweets from bots are retweets and 80% (9,220) of the tweets from the suspended accounts are retweets. The suspended accounts retweet significantly more than the bots and the genuine accounts.

For the bots, 20% (1,827) of the tweets contain links to external URLs. For suspended accounts, the proportion of these is 10% (1,175) and for genuine accounts 9% (16,290). The bots include URLs more than twice as often than suspended and genuine accounts.

We conducted a categorization of sites into a number of manually identified categories to get an understanding of what external sites the different accounts link to. The proportion of links to the different types of media is presented in Table IV.

TABLE V
THE DIFFERENT THEMES AND SOME SAMPLE WORDS

Theme	Sample words (eng)	Sample words (swe)
Islam	muslim, islam, jihad, sharia	muslim, islam, jihad, sharia
Migration	immigration,newly arrived	invandring, nyanlända
Russia	Russia, Putin, Moskow	Ryssland, Putin, ryska, Moskva

TABLE VI
OCCURRENCES OF THE DIFFERENT THEMES. FIGURES ARE PRESENTED IN PERCENT AND THE ABSOLUTE VALUE.

	Islam	Migration	Russia
Bots	1.84% (165)	6.37% (570)	1.38% (124)
Suspended	1.72% (197)	7.68% (881)	1.13% (130)
Genuine	1.36% (2,373)	6.17% (10,764)	1.17% (2,036)

B. *Thematic analysis*

To get an understanding of the kind of messages the bots are distributing, we conducted a thematic analysis of all tweets sent by bots. The thematic analysis was done using different dictionaries containing words that represent the different themes, and then counting the relative frequencies of these words in the text material. The relative frequencies of the dictionary words are then averaged, producing an aggregated score for each theme that represents its relative frequency of occurrence. The relative frequency indicates the themes that are more prevalent in the data, and how they develop over time (if the data has a temporal dimension). There are many drawbacks of using a dictionary-based approach. One obvious drawback is that the themes are decided beforehand, and new themes that arise in the data will not be analyzed. Another drawback is that the meaning of words can be context-dependent, which means that words may have several different meanings depending on the context.

To develop dictionaries that capture the different themes, we included experts with domain knowledge of the environment that we study. The dictionaries were augmented with words using a distributional semantic model that was pre-trained on relevant data. Experts manually verified each of the words suggested by the distributional semantic models before inclusion in the dictionary. The themes we have studied with some sample words from the dictionaries can be seen in Table V. The choice of themes was motivated by the fact that Islam, migration and Russia have been flashpoints in relation to misinformation campaigns in recent elections.

Table VI shows the percentage of tweets containing at least one word from each theme. The actual number of tweets are also listed in parentheses.

C. *Spread*

We visualized the network of users that are tweeting about the Swedish election. The visualization can be seen in Figure 1 where each node corresponds to a user. An edge between two users indicates that one of the users (or both) has retweeted tweets from the other. The size of the node corresponds to the number of outgoing edges, i.e., the bigger the node, the more users are retweeting that user. The color indicates whether the user is classified as genuine (green), bots (red) or suspended (yellow).

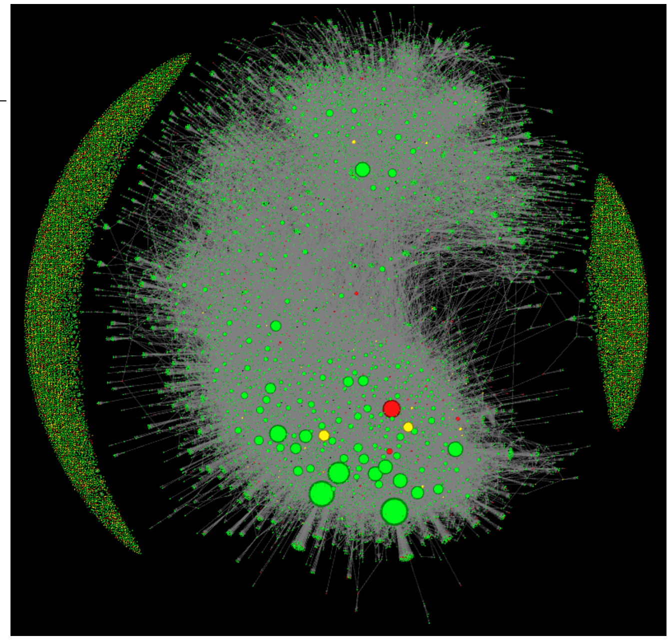

Fig. 1. Network of Twitter users discussions about the Swedish election.

In the figure, two clusters of users without edges on both sides can be noted. These clusters include users who are tweeting but do not get any retweets. These users might have followers who read their tweets, but they have not been retweeted when talking about the election.

In the middle, one big cluster with all the users retweeting each other can be seen. On the border around the cluster, small communities with few users retweeting different accounts can be noticed. In the middle of the lower part of the cluster, several large nodes corresponding to accounts which have been retweeted by several different users can be noticed. By a manual inspection, we notice that the lower part of the big cluster is users who mainly discuss and sympathize with right-wing politics. In the upper part of the cluster, users discussing left-wing politics can be found. The cluster can be divided into a right-wing part and a left-wing part. In between, we have a lot of independent political users such as news channels and political commentators.

From the visualization, we can draw the conclusion that the right-wing cluster is larger and has more frequent discussions than the left-wing cluster. We also notice by manual inspection that the right-wing cluster has many more key-users with many followers. A similar result was found in [14] where it was also noted that the bots in the German election were engaging more frequently with the alt-right English-speaking users.

V. IMPACT AND RELATION TO OTHER ELECTIONS

During the political discussions around the 2016 U.S. Presidential election, researchers [2] found that nearly 15% of the accounts were likely to be bots and that the bots were responsible for nearly 19% of the total conversations.

The German election was studied in [15]. A set of 984,713 tweets related to the German election was collected and accounts with bot-driven automation (accounts that post at least 50 times a day, meaning 500 or more tweets on at least one of these hashtags used for collection during a certain period) was identified. In total, around 7.4% of the total traffic around the German election was from bot-driven automation accounts. This result is in line with our findings from the Swedish election even though we acknowledge that the method used in [15] for detecting bots is entirely different from ours.

The same study [15] also identified the number of external URLs. In total, 11 646 URLs were analyzed and categorized. The results show that 40% of the links were pointing to news media and 28% were links pointing at other sources of political news and information, including junk news. The study showed that during the general election, German Twitter users shared many links to political news and information. Links to professional news outnumbered links to news by a ratio of four to one. Junk news are defined as websites that deliberately publish misleading, deceptive or incorrect information purporting to be real news about politics, economics or culture. Professional news content is political news and information by outlets that display the qualities of professional journalism, with fact checking and credible standards of production.

Another study that used Twitter data related to the German general election 2017 was presented in [14]. Data was collected from users who posted tweets related to the German election during a two-month period. To identify social bots in the dataset, the authors checked the status of accounts. If accounts had been suspended from Twitter, they were considered to be bots, otherwise they were considered to be legitimate (genuine) users. In the data collected from the German election, around 11% of the accounts were considered to be bots, and the amount of content they produced was 9%.

VI. Conclusions and Future Work

In this work, we have presented a classification model that recognizes automated behaviour among Twitter users. We have used our classification model on a set of Twitter accounts that show interest in the Swedish general election. Our analysis is in line with [15] in terms of the level of bot activity during the German election. The authors of [15] note that this is a far lower level of bot activity than during the US election of 2016. The closeness between the German and Swedish results makes intuitive sense, given that the media systems and campaign styles are more similar than to the US. Nevertheless, this is an emerging area, and the study presented here and [15] used different methods, so further validation will be needed in the future. Based on the data in our study we can conclude that bots retweet slightly less than genuine accounts. At the same time, bots include more external URLs in their tweets compared to genuine accounts. Much work remains to be done: while we have shown that a robust analysis of bots is possible using advanced methods, using these methods can be seen as a trial run for a period that leads up to the Swedish election, but does not yet cover the period of the most intense campaigning, when bot activity, if it is aimed at influencing the election, can also be expected to intensify sharply. Whether it does so, and to what extent our results can be further validated by comparing them to other, similar efforts, will be the subject of future research. Such future research, which includes the period through the election, should also take the larger media environment into account in which bots operate, and this includes other social media, traditional media, alternative news websites, and their audiences.

References

[1] F. Ahmed and M. Abulaish. A generic statistical approach for spam detection in online social networks. *Computer Communications*, 36(10-11):1120–1129, 2013.

[2] A. Bessi and E. Ferrara. Social bots distort the 2016 u.s. presidential election online discussion. *First Monday*, 21(11), 2016.

[3] Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia. Who is tweeting on twitter: human, bot, or cyborg? In *Proceedings of the 26th annual computer security applications conference*, pages 21–30. ACM, 2010.

[4] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi. Fame for sale: efficient detection of fake twitter followers. *CoRR*, abs/1509.04098, 2015.

[5] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi. Dna-inspired online behavioral modeling and its application to spambot detection. *IEEE Intelligent Systems*, 31(5):58–64, 2016.

[6] S. Cresci, R. D. Pietro, M. Petrocchi, A. Spognardi, and M. Tesconi. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. *CoRR*, abs/1701.03017, 2017.

[7] C. A. Davis, O. Varol, E. Ferrara, A. Flammini, and F. Menczer. Botornot: A system to evaluate social bots. *CoRR*, abs/1602.00975, 2016.

[8] Z. Gilani. Of bots and humans (on twitter), 08 2017.

[9] R. Gorwa and D. Guilbeault. Understanding bots for policy and research: Challenges, methods, and solutions. *CoRR*, abs/1801.06863, 2018.

[10] M. Grusell. Sociala medier i svenska medierörelser. In *När makten står på spel journalistik i valrörelser. Sthlm, Inst. f. mediestudier*, 2017.

[11] P. N. Howard, S. Woolley, and R. Calo. Algorithms, bots, and political communication in the us 2016 election: The challenge of automated political communication for election law and administration. *Journal of Information Technology & Politics*, 15(2):81–93, 2018.

[12] K. Lee, B. David Eoff, and J. Caverlee. Seven months with the devils: A long-term study of content polluters on twitter. 01 2011.

[13] Z. Miller, B. Dickinson, W. Deitrick, W. Hu, and A. H. Wang. Twitter spammer detection using data stream clustering. *Inf. Sci.*, 260:64–73, 2014.

[14] F. Morstatter, Y. Shao, A. Galstyan, and S. Karunasekera. From alt-right to alt-rechts: Twitter analysis of the 2017 german federal election. In *Companion Proceedings of the The Web Conference 2018*, WWW '18, pages 621–628, 2018.

[15] L.-M. Neudert, B. Kollanyi, and P. N. Howard. Junk news and bots during the german parliamentary election: What are german voters sharing over twitter? In *Data Memo, UK: Project on Computational Propaganda. comprop.oii.ox.ac.uk.* Oxford, UK, 2017.

[16] M. Roberts. Censored: Distraction and diversion insider chinas great firewall. In *Princeton: Princeton University Press*, 2018.

[17] M. Singh, D. Bansal, and S. Sofat. Who is who on twitter–spammer, fake or compromised account? a tool to reveal true identity in real-time. *Cybernetics and Systems*, pages 1–25, 2018.

[18] O. Varol, E. Ferrara, C. A. Davis, F. Menczer, and A. Flammini. Online human-bot interactions: Detection, estimation, and characterization. *CoRR*, abs/1703.03107, 2017.

[19] C. Yang, R. C. Harkreader, and G. Gu. Empirical evaluation and new design for fighting evolving twitter spammers. *IEEE Trans. Information Forensics and Security*, 8(8):1280–1293, 2013.