

Twitter Bots and the Swedish Election



Johan Fernquist, Lisa Kaati, Ralph Schroeder, Nazar Akrami,
and Katie Cohen

Abstract In this chapter, we present a study of how political Twitter bots were used before the Swedish general election in 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 behavior. This includes a human that manually copies or retweets content repeatedly in a robot-like way to influence the interaction between a user and content or with other users.

Our results show that bots were more likely to express support towards the immigration-critical party the Sweden Democrats, compared to genuine accounts.

Keywords Bots · Automated behavior · Twitter · Social media analysis · Election

1 Introduction

The Swedish general election was held on September 9th 2018 and was surrounded by great concern regarding the role of disinformation disseminated via digital media. The background to these concerns is broader debates about the role of digital media in politics that have intensified since the Brexit referendum and the election of Donald Trump. The Internet and social media are powerful tools for influencing political campaigns and discussions. Since 2016, the Internet has been used more than TV and newspapers as a source of political information in Sweden [17]. Ever

J. Fernquist · L. Kaati (✉) · K. Cohen
Swedish Defence Research Agency, Kista, Sweden
e-mail: johan.fernquist@foi.se; lisa.kaati@foi.se; katie.cohen@foi.se

R. Schroeder
Oxford University, Oxford, England
e-mail: nazar.akrami@psyk.uu.se

N. Akrami
Uppsala University, Uppsala, Sweden
e-mail: nazar.akrami@psyk.uu.se

since the Brexit Referendum and Donald Trump's presidential election campaign in 2016, there has been a discussion about the role that digital media play in disinformation and influence operations in political campaigns. One precondition for effectively being able to counteract disinformation is a better understanding of how disinformation and attempts to influence politics work. What are the messages? How are they spread, and with what intentions?

In Sweden, the discussion about disinformation has focused on several areas: the use of alternative or partisan websites, the possibility of the spread of messages by foreigners and specifically Russia, and the use of bots to spread messages. However, bots are not only used for spreading disinformation. 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 many different definitions of 'bots'. In [13] bots have been defined as "executable software that automates the interaction between a user and content or other users". In the work by Gorwa and Guilbeault [11] a typology of bots is presented. The typology 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 operate 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 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 [11], called "trolls". According to the typology presented in [11], the bots we are studying in this work are both automated social bots and sock puppets. The aim with the definition of bots that we use is to capture an automated behavior. The automated behavior can be induced by humans working at a state-owned agency to spread propaganda (sock puppets), by fully automated software bots (spambots, social bots), and by humans that are behaving in an automated fashion (spambots). The effect of an automated behavior is the same independently of if it is a software or a human that is spreading the messages.

Political bots are used in various ways when they aim at influencing public opinion. For example, bots can be used to spread disinformation to mislead about the state-of-affairs. They can also be used to spread false news with the aim of creating uncertainty about established sources of information. Another aim of using bots is to lead users to think that specific content is more shared, more generally accepted or more mainstream than is the case. We will use *bot* and *automated account* interchangeably. As mentioned before, we do not define an automated account based on whether there is a human or a piece of software that produces the content but rather base this on the account's behavior.

Even though bots can exist on many different platforms, 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 [12]. In the rest of this chapter, we will present an analysis of how political bots were used on selected hashtags on Twitter before and after the Swedish general election in 2018. The data was collected between March 5th and September 30th.

The analysis of automated accounts presented in this chapter consist of three different parts:

- The number of automated accounts (bots)
- An analysis of the domains bot links to and what kind of messages they distribute
- How the bots communicate with other accounts on Twitter

1.1 Swedish Politics

In Sweden, a party must receive at least 4% of the votes in an election to be assigned a seat in the Swedish parliament. In 2018 eight parties received more than 4% of the votes. The largest party is the Social Democratic Party (S). S is a labor party at its core with policies based on freedom, equality, and solidarity. The party prioritizes the creation of more jobs and to provide a better education for all.

The Moderate Party (M) is the second largest party. M is a conservative party with liberal ideas. The individual's freedom to choose is central to its policies, and the party generally supports reduced taxes and economic liberalism.

The Sweden Democrats (SD) is a social conservative party based on nationalistic values. The party is associated with issues of migration, and the party's policies are based on protecting the *national identity* as a way of sustaining the Swedish welfare state.

The Centre Party (C) is a liberal and agriculture political party. The party believes that society should be built on people's responsibility for each other and nature and focus on the national economy, the environment, and integration.

The Left Party (V) defines itself as a socialist and feminist political party with an ecological basis. Focus areas are jobs, welfare services, and gender equality. The party was against Sweden joining the EU in 1995 and still advocates an exit. The Christian Democratic Party (KD) believes that stable families should form the

basis of society. The four main issues that the Christian Democrats focuses on are: improving elderly care, giving families with children the freedom to select desired childcare, simplifying regulations for companies and lowering taxes as a means to promote growth and combat unemployment.

The Liberals (L) are a liberal and social—liberal political party that holds a middle position in the Swedish political landscape. The Green Party (MP) has a clear focus on environmental issues. The party focuses on stopping climate change and protecting the environment, fighting nuclear power and promoting European integration.

Apart from the eight parties that are in the Swedish parliament, there are also several smaller parties. Two newly formed parties that appear in our analysis are Alternative for Sweden (AFS) and Citizens' Coalition (MED). AFS was founded in 2017. The party's policies are based on immigration issues, democracy and politicians, and law and order. MED considers itself liberal-conservative and green conservative. Both parties are anti-immigration.

2 Method

In our analysis of the Swedish election, we use machine learning to detect accounts with automatic behavior. In the rest of this section, we describe how the classification model is built, how it performs compared to other bot detection models, and what data we have used to train the model. Our work is also put into relation to previous work in the domain.

2.1 Classification of Bots

We have trained a classification model to identify accounts exhibiting automatic behavior. The classification problem is to determine if an account is *genuine* or a bot. Here, a genuine account is an account that is operated by a “normal” human being. To build a model that is able to recognize automatic behavior, labeled training data is needed. The training data consists of accounts that are already known as bots or genuine accounts. When training our model, we use a number of different features. Our model is language independent but we have used it to classifying tweets in Swedish. The classification is described in more detail in [9].

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 [14, 18, 19]. In [10] 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 [10] it is also shown that genuine accounts get more likes on their tweets than bots.

In [2] 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. Instead, they introduce text entropy as a feature to measure the similarity of the texts posted by an account with the hypothesis that bots 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 [19] a total of 1150 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 results show 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.

Training and Testing Data

We have 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 [19]. The dataset contains labeled information about 647 bot accounts and 1367 genuine accounts. Each account 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.

The second dataset consists of 591 bots and 1680 genuine accounts [4]. 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 dataset was manually annotated by four undergraduate students [10]. 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. We only use the two sets with users with 90 thousand to 110 thousand followers and users with 900 to 1100 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, 10, 19] 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 the datasets for comparing performance. The datasets were only used for training our model.

The dataset used in [6] (referred to as test set one 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 3000 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 purchased 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.

Features

In our classification model, we have used a total of 140 features. The features can be 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 [2, 19] we use text entropy assuming that bots might have a less complex and varied way of expressing themselves. Similar to [19], we have included time features such as statistics of time between consecutive tweets, retweets, and mentions as well as statistics for the time between posted tweets containing URLs. All features are listed in Table 1.

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 (so far) are when using random forest which is the motivation for us to also use random forest in our classification.

Model Evaluation

We have used three datasets from [4, 10, 19] together with our 140 different features to train a model using random forest. The model was tested on the only dataset we have access to in its original form. In [6], the same dataset was used to compare the performance of other bot classification models. We included a comparison from [6] and used the same dataset to test our model. The performance (Accuracy, Precision,

Table 1 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 ^a
# friends	Length of tweet ^a
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 given 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 ^a
	Time between mentions ^a
	Time between retweets ^a
	# words ^a
	Hours of day tweeting
	Weekdays tweeting
	Normalized distribution hours tweeting
	Normalized distribution weekdays tweeting
	Normalized distribution of tweet endings
	String entropy ^a
	Total entropy of all tweets' strings concatenated

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

Table 2 Performance measures for different models, the model with the best performance is marked with bold

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. [21]	Supervised	0.506	0.563	0.170	0.261
Miller et al. [16]	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

Recall, and F1-score) is shown in Table 2. The comparison includes both supervised and unsupervised models. As mentioned earlier, the 911 genuine accounts in [6] were selected randomly from a set of more than 3000 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. In [16] feature vectors with the majority of features as text features are clustered 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 Fast greedy community detection algorithm. In [5] a bio-inspired technique for modelling behavior of users online with so-called *digital DNA* sequences is presented. The sequences are string encodings of the behavior of a user, and the sequences are then compared between the different users by measuring the longest common substring to find clusters of users.

The results in Table 2 shows that the model from [5] performs best on all metrics except for recall where our model performs better. Our model performs best of the supervised models included in the comparison. 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 datasets.

Feature Importance

To get an understanding of what features that played an important role in the classification we have calculated the feature importance with *forests of trees*. The ten most important features are shown in Table 3.

The most important feature when determining whether an account is a bot or not is the number of given likes divided by the number of friends. The second most important feature is the ratio between the number of followers and friends. As mentioned earlier, this feature could however sometimes be misleading since bots can be able to unfollow accounts that not are following back. Several of the most important features are related to the time between retweets.

Table 3 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

Table 4 The different categories used to categorize the content of tweets

Criticism om media	Criticism of journalists, journalism or media
Criticism of elites	Criticism of the government, or persons or organizations that are regarded as influential political elite
Party support	Support of one or more parties
Criticism of parties	Criticism of one or more parties
Criticism of immigration	Criticism of immigration, asylum seekers and refugees
Election fraud	Discussions or observations of election fraud
Deleted	Tweets that have been deleted and can no longer be analyzed
Other/Uncategorized	Content that does not fall into any of the categories above

2.2 Content Classification

To understand the difference in communicated messages between bots and genuine Twitter accounts, a random sample of tweets was selected and then manually coded. This resulted in a number of different categories. The categories were identified by five persons that independently analyzed a set of tweets. The categories were based on the content of the tweet. The result of the five persons' classifications was combined into eight generic categories that are described in Table 4.

When the categories were defined, ten research assistants from Uppsala University classified a total of 1063 tweets, 547 tweets were published by bots and 516 published by genuine account. Each tweet was classified by at least two research assistants and each tweets category was determined by a majority decision. The research assistants did not know whether the tweet they were classifying was published by a bot or a genuine account. If the category for a tweet was selected as party support or criticism of party, the research assistant was able to enter which one or which parties that the support or criticism was expressed for. If a tweet expressed both support and criticism, the research assistants were advised to select party support.

2.3 Data

The data that is analysed was collected during the period March 5th–September 30th 2018. All tweets in Swedish that including at least one of the hashtags #valet2018, #val2018 or #valet, or one of the keywords *valet2018*, *val2018* or *valet* (all words are related to the election). The word *valet* is *the election* in Swedish. A total of 1,005,276 tweets published by 70,973 accounts were collected from the Twitter streaming API.

The classification of the accounts in the dataset was done after all the tweets had been downloaded. A number of tweets in the dataset belong to accounts that were suspended by Twitter or deleted by the users themselves. The most common reason for an account to be suspended is that the account is spamming or using a fake identity.¹ This indicates that a large part of the suspended accounts are automated, however, accounts can also be suspended by exhibiting offensive behavior.²

3 The Amount of Bots

During the period from March 5th to September 30th 2018, a total of 1,005,276 tweets linked to the Swedish general election was published by 70,973 accounts. These accounts were classified into four different classes: genuine, automatic, suspended, or deleted. The distribution between the different classes of accounts are shown in Table 5. Most of the accounts were classified as genuine accounts.

Approximately 6% of the accounts that tweet about Swedish election were classified as automated accounts, in total they were responsible for 8% of the content. If we make the somewhat extreme assumption that all suspended accounts are automated as well, then 12% of accounts are automated, and 9% of the content is automated. The real proportion of automatic accounts are most likely somewhere in that interval.

In Fig. 1, the number of tweets per day is shown. From the beginning of August until the election day, there is an increase of tweets about the election. On the election day September 9th, more than 45 thousand tweets were published. The figure shows that the interest in discussing the election decreased right after the election.

In Fig. 2, the number of active accounts per month from March to September are shown. More accounts were participating in discussing the election closer to the election. The genuine accounts make up the biggest share of accounts for the whole period. From July to August, the number of genuine accounts increased with 118%. For the same period, July to August, the number of bots increased by 167%. From

Table 5 Distribution between the different categories of accounts

Account category	Number of accounts	Number of tweets
Genuine	60,384	876,792
Bots	4084	73,723
Suspended	4052	14,902
Deleted	2453	39,859

¹See Twitter's help page about suspended accounts <https://help.twitter.com/en/managing-your-account/suspended-twitter-accounts>.

²See Twitter's end-user licence agreement <https://help.twitter.com/sv/rules-and-policies/twitter-rules>.

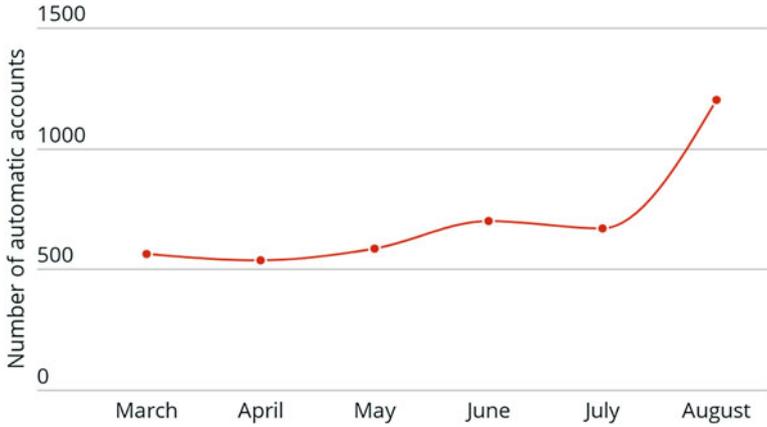


Fig. 1 The number of tweets per day for each category of accounts

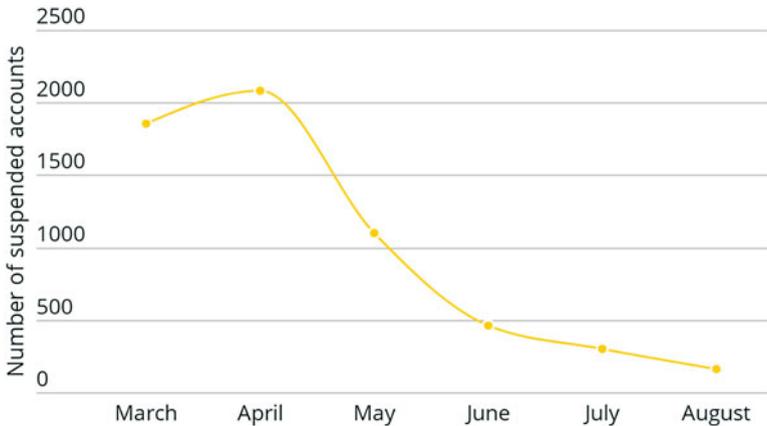


Fig. 2 The number of accounts per month for each category of accounts

August to September, the number of genuine accounts increased with 16%, and the number of bot accounts with 28%.

4 Content

To get an understanding of what kind of content that were published we conducted two analyses. We explored the domains to which different accounts link to and what messages that are spread by the different classes of accounts.

4.1 Out Links

A tweet can consist of a maximum of 280 characters, which sometimes restricts the user from actually convey the message in the text. It is quite common to include a link to an external website in the tweet. In the dataset being analyzed here, it is common to link to online versions of traditional media such as The Swedish public service television company (SVT) or the newspaper Expressen, it is also common to link to digital platforms such as Youtube and Facebook. A third category of media that is quite common to link to is immigration critical alternative media³ such as Samhällsnytt and Nyheter idag.

To get an understanding regarding the differences between the different classes of accounts, we have studied the most common domains that is linked to, for each of the different classes of accounts. The ten most linked domains for each of the categories of accounts are shown in Table 6. For each of the categories of accounts, *expressen.se* and *svt.se* are the most common domains to link to. *Youtube.com* and *aftonbladet.se* are found in the list of most linked domains for each of the class of accounts. This is also true for the immigration critical alternative media *samnytt.se*, *nyheteridag.se*, and *friatider.se*. For the suspended accounts, half of the ten most linked domains are immigration critical. Here we can find the Nordic resistance movement's news portal *nordfront.se*. One of the reasons for finding *nordfront.se* in the list for the suspended accounts is that Twitter are suspending accounts which are linking to *nordfront.se*.⁴

Table 7 shows the difference between the bots and genuine accounts most shared domains 1 week before and after the election. *Svt.se*, *expressen.se* and

Table 6 The ten most linked domains for the different categories of accounts

Genuine	Bots	Suspended	Deleted
<i>expressen.se</i>	<i>expressen.se</i>	<i>expressen.se</i>	<i>expressen.se</i>
<i>svt.se</i>	<i>svt.se</i>	<i>svt.se</i>	<i>svt.se</i>
<i>aftonbladet.se</i>	<i>aftonbladet.se</i>	<i>nordfront.se</i>	<i>samnytt.se</i>
<i>dn.se</i>	<i>omni.se</i>	<i>youtube.com</i>	<i>aftonbladet.se</i>
<i>youtube.com</i>	<i>sverigesradio.se</i>	<i>friatider.se</i>	<i>friatider.se</i>
<i>samnytt.se</i>	<i>youtube.com</i>	<i>samnytt.se</i>	<i>youtube.com</i>
<i>nyheteridag.se</i>	<i>samnytt.se</i>	<i>nyheteridag.se</i>	<i>nyheteridag.se</i>
<i>omni.se</i>	<i>friatider.se</i>	<i>aftonbladet.se</i>	<i>facebook.com</i>
<i>gp.se</i>	<i>dn.se</i>	<i>katerinamagasinet.se</i>	<i>dn.se</i>
<i>friatider.se</i>	<i>nyheteridag.se</i>	<i>sverigesradio.se</i>	<i>gp.se</i>

³In this chapter, we call these websites *immigration critical alternative media* since that is an established concept in Sweden when talking about media that are critical of what is considered to be an overly generous immigration policy.

⁴This article describes how Twitter removed accounts spreading links to *nordfront.se* (in Swedish) <https://www.dn.se/nyheter/politik/flera-konton-kopplade-till-nazistiska-nmr-raderade-av-twitter/>.

Table 7 The ten most linked domains for bots and genuine accounts 1 week before and after the election

Genuine accounts		Bots	
Before election	After election	Before election	After election
svt.se	expressen.se	expressen.se	svt.se
expressen.se	svt.se	svt.se	expressen.se
aftonbladet.se	aftonbladet.se	aftonbladet.se	aftonbladet.se
nyheteridag.se	data.val.se	samnytt.se	youtube.com
samnytt.se	metro.se	sverigesradio.se	omni.se
youtube.com	dn.se	youtube.com	sverigesradio.se
dn.se	nyadagbladet.se	omni.se	dn.se
omni.se	omni.se	nyheteridag.se	samnytt.se
sverigesradio.se	youtube.com	friatider.se	metro.se
facebook.com	samnytt.se	dn.se	data.val.se

aftonbladet.se, all traditional media, are dominating as most shared domains both before and after the election. Links to the immigration critical domains have decreased after the election for both type of accounts. Data.val.se is a website with statistics about the election, it is present for both account classes after the election.

5 Messages

To understand the difference of content spread by bots and genuine accounts, we have analyzed a sample of tweets published 1 week before and 1 week after the election (see Sect. 2.2). The distribution between the different messages for bots and genuine accounts is shown in Fig. 3.

The most common message for both categories is party support followed by criticism of parties. More than 20% of the tweets published by bots have been deleted, compared to 10% for the genuine accounts. Overall, there are only minor differences between what bots and genuine accounts communicate.

In Figs. 4 and 5 the distribution of party support and party criticism for bots and genuine accounts are shown. The figures show the proportion of support and criticism for the different parties.

Regarding party support, it is clear that the Sweden democrats (SD) and Alternative for Sweden (AFS) are the most common parties to express support for in Twitter. It is also more common that a bot is showing support for SD compared to a genuine account. Also support for the Left Party (V) is more common from a bot than a genuine account.

Regarding the party criticism that is shown in Fig. 5, the criticism for the parties is relatively even between genuine accounts and bots. The Social Democratic Party (S) is the party receiving most of the criticism. The Center Party (C) and the Sweden democrats (SD) are the parties where the difference between criticism between

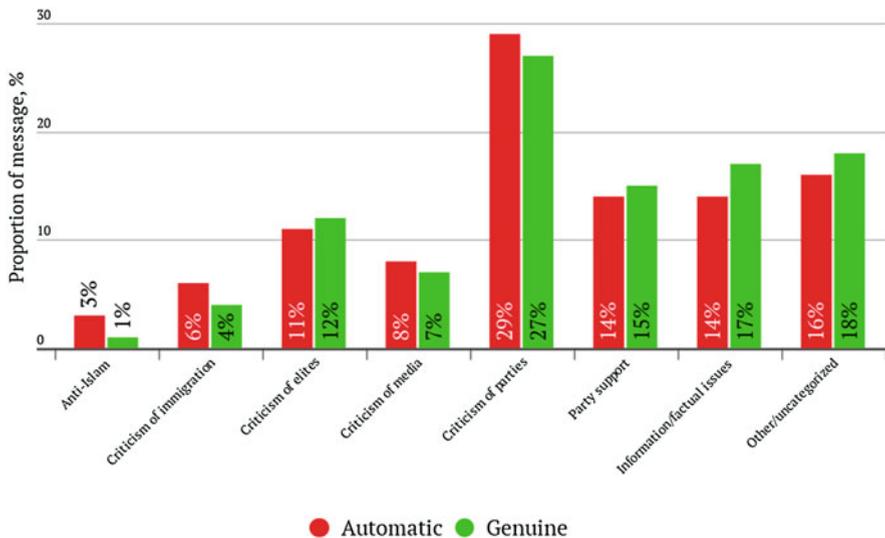


Fig. 3 The share of different messages spread by different categories of accounts

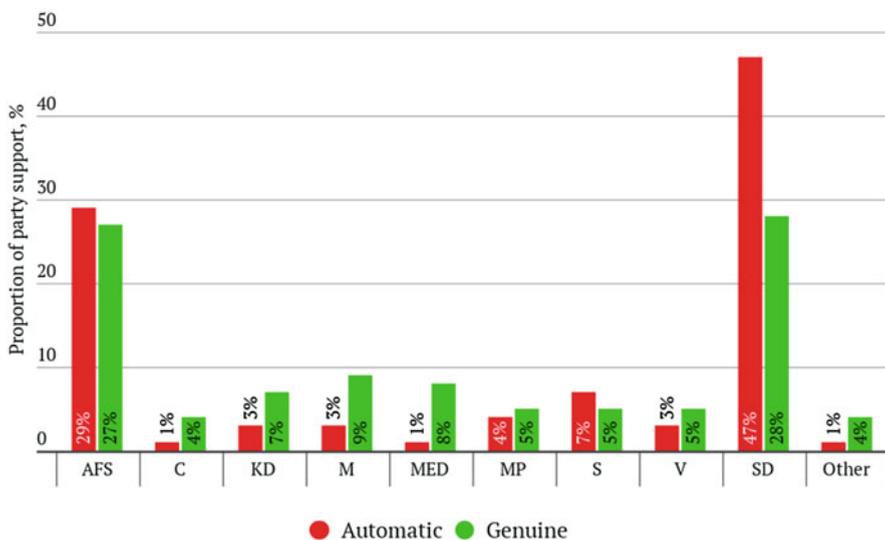


Fig. 4 The share of parties for which different categories of accounts express support

genuine and bot accounts are the largest. It is more common that a genuine account expresses criticism for SD, while it is more common for a bot to express criticism towards C.

There have been other studies on how bots were used in discussions about the Swedish election. In [3], the authors found at least 55 accounts with bot-like

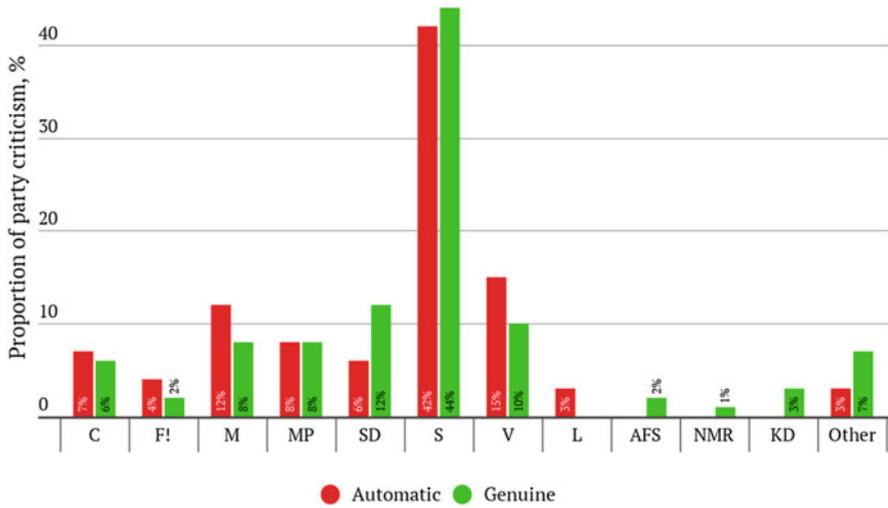


Fig. 5 Distribution of the parties for which different accounts express criticism

behavior that expressed support for AFS. Those accounts were active especially the weeks prior the election. The study showed that bot-like behavior was particularly clear for accounts criticizing S and supporting SD, which is in line with the results of our study. In [3] it is concluded that the behavior of the identified bot accounts was more genuine-like than bots which had been identified in previous elections in Germany and France.

The Swedish newspaper Dagens Nyheter published an article [15] about a Twitter analysis regarding the support of AFS. All tweets with the hashtag #afs18 was downloaded. The analysis showed that half of the 14,000 tweets with the hashtag was published by 15 accounts.

6 Spread

How messages spread and users communicate with each other can be understood by studying the network of accounts retweeting each other. Retweeting is a powerful tool for influencing others on Twitter. When retweeting another account, you are publishing someone else’s post in your own feed. This is a way to distribute someone else’s message and commonly this also means that you agree and express support for another account and the published post.

A principle in marketing and propaganda is that increased visibility leads to increased opportunity to influence. An article which because of several retweets appears in a user’s feed several times will most likely influence the user more than an article that appears only once. If the article also conveys a message which

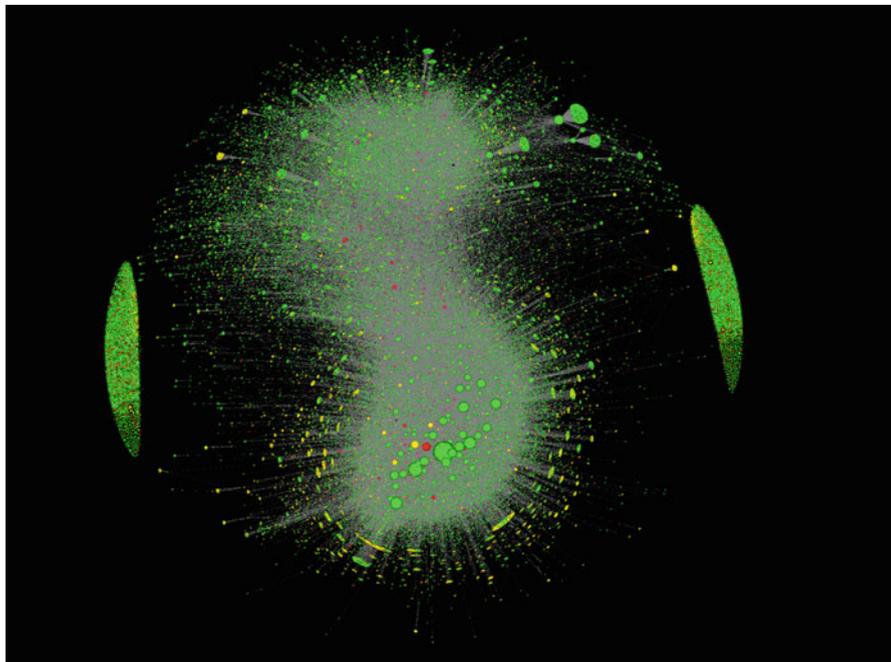


Fig. 6 The network of accounts that have discussed the election

the user recognizes from previous discussions, the article is even more likely to assimilate [8].

In Fig. 6, the 70,973 accounts tweeting about the election from March 5th until September 30th are shown. In the figure, every account is indicated by a node (circle), and every line between two accounts shows that one account has retweeted the material of another account in its own feed at least once during the time period. The size of an account's circle indicates how much an account has been retweeted by other users. The larger the circle, the more popular account to retweet. Green nodes represent genuine, red nodes represent bots, yellow nodes represent suspended, and the white nodes represent deleted accounts.

In Fig. 6, two big clusters of accounts which do not retweet anyone else in the networks are present. These two clusters appear in the top and bottom of the figure. Since these accounts have not retweeted anyone else in the network, they are not a part of the bigger cluster in the middle but are instead found outside the big cluster. This does not exclude that these users might follow, like and comment other users in the network. These accounts may also have ended up in our dataset since they have used the selected hashtags and keywords, but not in the context of discussions about the Swedish election.

Notably, there are a number of smaller clusters of accounts on the edges of the big cluster. These clusters consist of accounts that show the same behaviors as others in the groups, and they retweet only larger and more popular accounts. These

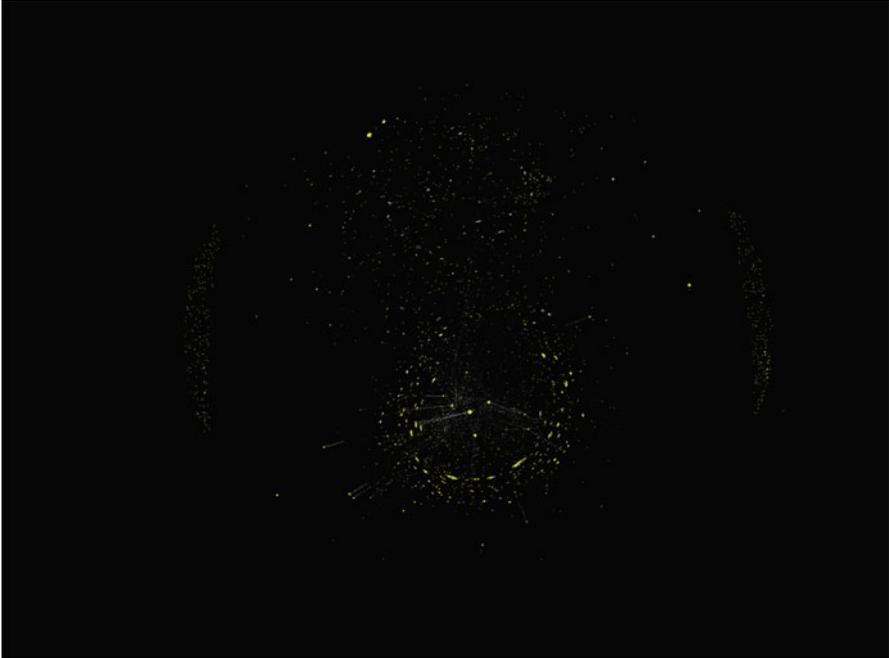


Fig. 7 The network of accounts that have discussed the election with the two clusters marked

small clusters indicate that some accounts have a group of followers which by only retweeting one account make sure that the account gets increased distribution of communicated messages. By investigating what category of accounts that appears in these smaller clusters, we can detect whether some accounts have used for example bots to get increased distribution of tweets.

Manual content analysis revealed that the main part of the network can be divided into two clusters. One of the clusters are inside the solid circle marked in Fig. 7. In this cluster, one can find the majority of the official accounts of the political parties in the Swedish parliament. The discussions are politically general and many different political issues are discussed. The second cluster can be found in the dashed circle. This cluster consists of accounts that discusses immigration policies and the negative consequences with immigration. The most popular accounts (in terms of retweets) in the network appear in the middle of the dashed cluster. These accounts have many followers and are often retweeted by several other accounts. In the middle between the two clusters (solid circled and dashed circled), we find accounts that positions themselves as political independent. These accounts are retweeted by both the clusters. Accounts that appears between the clusters can for example be news sites.

6.1 *The Impact of Bots and Suspended Accounts*

To get an understanding of whether the bots was effective in spreading messages, we investigated how the different account categories appears in the network. Since the most common reason that an account is being suspended by Twitter is that it shows what we call a bot-like behavior, we have also included the suspended accounts in this analysis.

The network analysis can be used to give answers to the following questions:

- Are the bots and the suspended accounts popular to retweet?
- Are there clusters consisting of bots and suspended accounts?

By detecting clusters of bots and suspended accounts, we can find networks connected to individual accounts that might have used bots to distribute their messages. Seven of the ten most retweeted accounts in the cluster are part of the immigration critical cluster. Nine of the accounts are genuine, and the tenth has been suspended by Twitter.

In Fig. 8, we can see where in the network the bots appear. The majority of the bots can be found outside the big cluster and appears in the two clusters of accounts

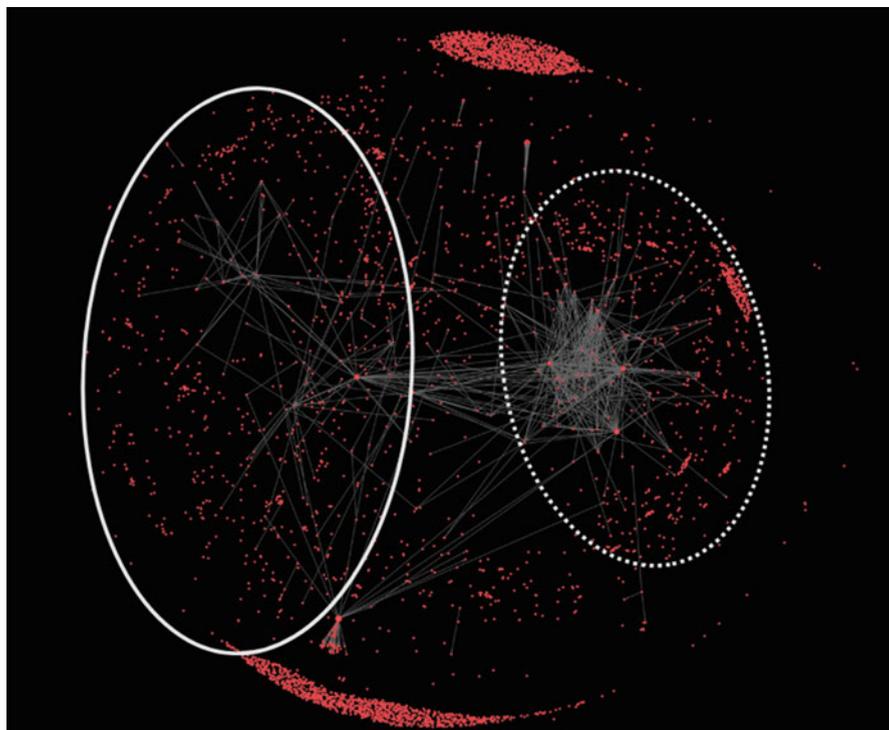


Fig. 8 The network of accounts only showing the accounts that have been classified as bots

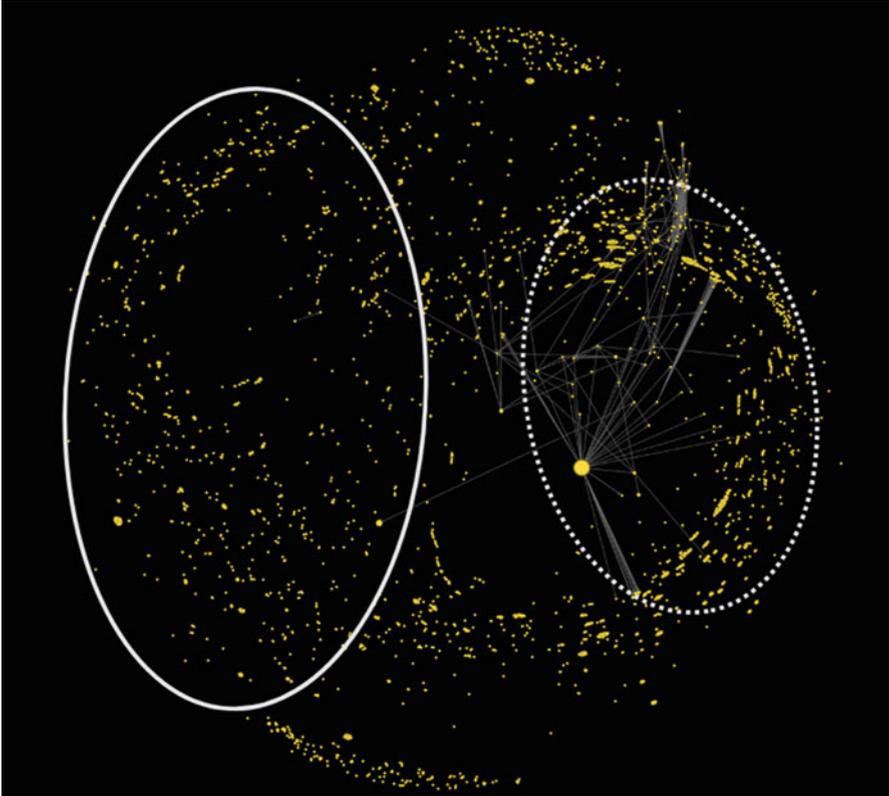


Fig. 9 The network of accounts only showing the accounts which have been suspended by Twitter

which has not retweeted nor been retweeted. The density of bots are higher in the immigration critical cluster but no large groups consisting of a large amount of bots retweeting only one account are identified.

In Fig. 9, we can see where in the network the accounts that have been suspended by Twitter appear. It is clear that the majority of the suspended accounts appear in the immigration critical cluster, and several smaller clusters of suspended accounts are found. Each of these clusters consist of accounts that have only retweeted one other account in the network and therefore acts as distributors of another account's messages. The occurrence of retweet accounts can serve as an indication that bots have been used.

Our analysis shows that it is more common for accounts discussing immigration to have other accounts distributing their messages. These accounts are more likely to be suspended by Twitter, one possible reason for the suspension is that they show a bot-like behavior.

7 Party Support on Twitter and the Election Results

In Sect. 5, we presented how bots and genuine accounts were expressing support for the different parties. In Fig. 10, the expressed party support on Twitter and the actual election results are shown. SD and AFS received the most support on Twitter while S and M got the largest share of votes in the election.

When the result of the election was reported, discussions about election fraud started on Twitter. The discussions included claims that the election was rigged and that the results were settled earlier and invalid. The term “valfusk” (election fraud in Swedish) occurred almost 2500 times the day after the election, as can be seen in Fig. 11.

The hashtag #valfusk was used frequently during the election night. Note that this hashtag was not a part of our searched hashtags or keywords and therefore all

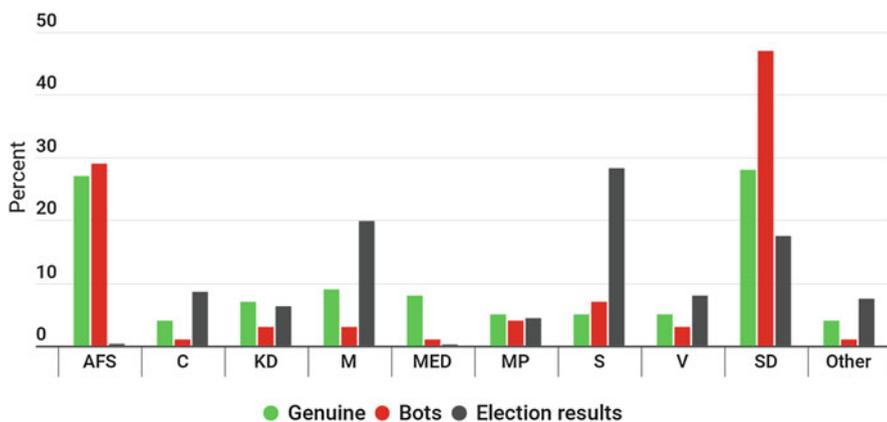


Fig. 10 Distribution of the parties for which different accounts express support, and the actual election results

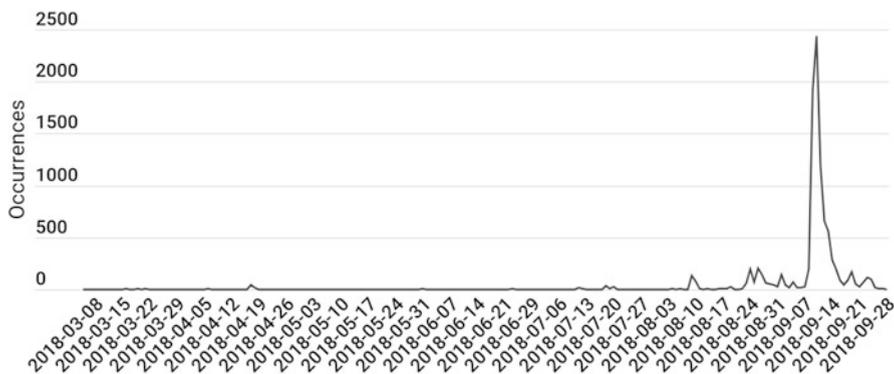


Fig. 11 Occurrences of the term valfusk (election fraud)

tweets with the hashtag #valfusk are not included in this study. The figure shows the occurrences of the term valfusk in the data with the hashtags and keywords mentioned in Sect. 2.

There are several reasons for the increase of discussion about election fraud. One reason might be many Twitter users were disappointed of the election results. Another reason can be that the experienced support of certain parties in Twitter did not agree with the results from the election.

8 Discussion and Conclusions

Among the accounts that we analyzed, the proportion of bots was between 6 and 12%, depending on whether we include accounts that were suspended among our bots or not. We can not say for sure how many of the suspended accounts that are bots, but since Twitter states that the most common reason for suspending an account is that the account shows bot-like behavior, we can assume that a high proportion of suspended accounts are bots. The number of bots that tweeted about the election more than doubled from July to August.

Our analyses shows that it has been more common for a bot account to express support for the Sweden democrats (SD), compared to a genuine account. The majority of the bots have not been a part of the cluster expressing immigration criticism, but rather a part of the clusters with accounts not retweeting any other accounts.

Our content analysis shows that criticism towards immigration and the negative consequences with immigration is frequently discussed on Twitter in conjunction with discussions about the election. However, it is important to point out that the hashtags we have used to gather our data might have been more popular to use among specific groups. This might have lead to skew in the result making the result in the study not a fully valid representation of the discussions about the election occurring on Twitter. The use of bots increased as the election date approached. The party AFS, which was founded in March 2018, did succeed on Twitter where they got around 25% of the party support from both bots and genuine accounts. This result is consistent with other studies that found that very active accounts with bot-like behavior has expressed support for the party.

It is important to add that it is not illegal to recruit bots to get an increased spread of a message. A thousand followers can be bought for around ten dollars⁵ and retweets can be bought for twice as much.⁶

The analysis of supported parties showed that the immigration critical parties SD and AFS had more support on Twitter than in the actual election. After the election, several discussions about election fraud occurred. The difference between

⁵BuyTwitterFollowersReview—<https://buytwitterfollowersreview.org/top-10/>.

⁶BuySocialMediaMarketing—<https://buisocialmediamarketing.com/twitter/retweets>.

the election results and the expressed support on Twitter have likely contributed to the discussions about election fraud. Only a few days after the election, the discussions about election fraud subsided. One obvious conclusion of our study is that in digital media certain parts of the Swedish population are more engaged. In the case of Twitter, the most engaged are those who support the SD or who are strong critics of S.

One of the questions raised by this study concerns how bots influence the democratic process and the political discussion before the election. This question cannot be answered without an understanding of how the automatic accounts analyzed here fit into a larger picture of media use in Sweden. It would also be necessary to analyze who is behind these bots and what the aims are of spreading messages from bots. While answering these question is difficult, it is clear that the use of bots for spreading various types of messages increased the closer we come to the election. This could be a sign of an attempt to influence public opinion or at least a certain part of the political discussions.

An important question concerning all types of influence is to what extent the individual who is the target of influence is aware of whether someone is attempting to exercise influence. The results of previous research indicate that attempts to influence are less effective if individuals are aware of them [20]. In other words, an awareness that someone is trying to influence us can, at least to some extent, make us less susceptible to influence. Hopefully, studies like ours contributes to better awareness of attempts to exercise influence using bots.

References

1. F. Ahmed, M. Abulaish, A generic statistical approach for spam detection in online social networks. *Comput. Commun.* **36**(10-11), 1120–1129 (2013)
2. Z. Chu, S. Gianvecchio, H. Wang, S. Jajodia, Who is tweeting on twitter: human, bot, or cyborg? in *Proceedings of the 26th annual computer security applications conference* (ACM, New York, 2010), pp. 21–30
3. C. Colliver, P. Pmerantsev, A. Applebaum, J. Birdwell, Smearing Sweden, international influence campaigns in the 2018 Swedish election. Technical report (Institute of Strategic Dialogue (ISD), London School of Economics, London, 2018)
4. S. Cresci, R.D. Pietro, M. Petrocchi, A. Spognardi, 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, M. Tesconi, Dna-inspired online behavioral modeling and its application to spambot detection. *IEEE Intell. Syst.* **31**(5), 58–64 (2016)
6. S. Cresci, R.D. Pietro, M. Petrocchi, A. Spognardi, M. Tesconi, The paradigm-shift of social spambots: evidence, theories, and tools for the arms race (2017). *CoRR*, abs/1701.03017
7. C.A. Davis, O. Varol, E. Ferrara, A. Flammini, F. Menczer, Botornot: a system to evaluate social bots (2016). *CoRR*, abs/1602.00975
8. P.M. DeMarzo, D. Vayanos, J. Zwiebel, Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* **118**(3), 909–968 (2003)
9. R.S.J. Fernquist, L. Kaati, Political bots and the Swedish general election, in *International Conference of Security Informatics (ISI)* (2018)

10. Z. Gilani, R. Farahbakhsh, G. Tyson, L. Wang, J. 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*, ASONAM '17 (ACM, New York, 2017), pp. 349–354
11. R. Gorwa, D. Guilbeault, Understanding bots for policy and research: challenges, methods, and solutions (2018). CoRR, abs/1801.06863
12. M. Grusell, Sociala medier i svenska medierörelser, in *När makten står på spel—journalistik i valrörelser. Sthlm, Inst. f. mediestudier* (2017)
13. P.N. Howard, S. Woolley, R. Calo, Algorithms, bots, and political communication in the us 2016 election: the challenge of automated political communication for election law and administration. *J. Inform. Tech. Polit.* **15**(2), 81–93 (2018)
14. K. Lee, B. David Eoff, J. Caverlee, Seven months with the devils: A long-term study of content polluters on twitter, in *Fifth International AAAI Conference on Weblogs and Social Media*, vol. 01 (2011)
15. E. Mannheimer, H. Ewald, Så sprider falska konton högerextrema budskap inför eu-valet, in *Dagens nyheter* (2018)
16. Z. Miller, B. Dickinson, W. Deitrick, W. Hu, A.H. Wang, Twitter spammer detection using data stream clustering. *Inf. Sci.* **260**, 64–73 (2014)
17. N. Newman, R. Fletche, A. Kalogeropoulos, D.A.L. Levy, R.K. Nielsen, *Reuters Institute Digital News Report 2018* (Reuters Institute for the Study of Journalism/University of Oxford, Oxford, 2018)
18. M. Singh, D. Bansal, S. Sofat, Who is who on twitter—spammer, fake or compromised account? a tool to reveal true identity in real-time. *Cybern. Syst.* **49**(1), 1–25 (2018)
19. O. Varol, E. Ferrara, C.A. Davis, F. Menczer, A. Flammini, Online human-bot interactions: Detection, estimation, and characterization (2017). CoRR, abs/1703.03107
20. W. Wood, J.M. Quinn, Forewarned and forearmed? two meta-analysis syntheses of forewarnings of influence appeals. *Psychol. Bull.* **1**(129), 119–138 (2003)
21. C. Yang, R.C. Harkreader, G. Gu, Empirical evaluation and new design for fighting evolving twitter spammers. *IEEE Trans. Inf. Forensics Secur.* **8**(8), 1280–1293 (2013)