



Bots and the Swedish election

A study of automated accounts on Twitter

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

Cover: Shutterstock/Leo Blanchette

1 Introduction

In this report, we present a study of how political bots were used on selected hashtags on Twitter before the Swedish general election in 2018. The data was collected between March 5th and August 20th and the analyses consist of three parts. The first deals with the scope of political bots. The second deals with what domain bots link to, and what kind of messages they spread. The last part concerns the spread of political bots.

1.1 Influence via Social Media

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.¹ Ever 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?

The main aim of this study is to provide better insights into how, and to what extent, bots spread their messages in the period leading up to the election. The study also includes analyses of how the content of bot messages differ from ordinary discussions on Swedish politics and the election. We focus on Twitter because it is a commonly used forum for Swedish political discussions, especially among journalists.²

1.2 Results in Brief

- The number of political bots that discuss Swedish politics and the Swedish election on Twitter has grown significantly in recent weeks
- The percentage of tweets linking to alternative or partisan websites (*Samhällsnytt* and *Fria Tider*) is higher among bots than among genuine accounts
- The Sweden Democrats received most support among the total number of accounts expressing support for political parties. This support was expressed

¹ Newman, Nic, Richard Fletcher, Antonis Kalogeropoulos, David A. L. Levy & Rasmus Kleis Nielsen (2018) 'Reuters Institute Digital News Report 2018'. Oxford: Reuters Institute for the Study of Journalism, University of Oxford, <http://www.digitalnewsreport.org/>

² M. Grusell. (2017) Sociala medier i svenska medierörelser. Dahlberg, Stefan et al. (ed). När makten står på spel journalistik i valrörelser. Stockholm, Institutet för mediestudier.

more among bots. Specifically, Sweden Democrats received support from 47% of the bots and 28% of the genuine accounts.

- Expressions of traditionalist, authoritarian or nationalist views were more common among suspended or deleted accounts.

2 Bots and Automated Accounts

The concept of bots³ typically refers to computer programs that are developed to carry out tasks that would be highly demanding and/or time-consuming to carry out otherwise. The use of bots, therefore, aims to *automate* a process. There are many different types of bots that are used with different aims. Spam bots, for example, are used on social media to spread advertising and chat bots are used in the context of customer service to quickly answer simple questions.⁴ Bots that are on social media and that behave like humans on social media are called social bots. Bots that are used to spread messages on social media without identifying the sender are called trolls or sockpuppets. Sockpuppets received considerable attention in 2017 when a Russian freelance journalist was employed at a troll factory in St. Petersburg to spread propaganda which criticized the West on social media.

2.1 Political Bots

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 a specific content is more shared, more generally accepted or more mainstream than is actually the case.

We will use *bot* and *automated account* interchangeably. 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 accounts' behavior.

An account with automated behavior can be an account that is deployed automatically by software (a *social bot*) or generated by a person who is employed to spread propaganda (a *sock puppet*) or generated by an individual who uses the account to continually and highly frequently manually copy or retweet the content. The effects of these behaviors are the same regardless of whether there is software or a human responsible for the account.

³ The term bots is derived from 'robots'.

⁴ R. Gorwa & D. Guilbeault. (2018) Understanding bots for policy and research: Challenges, methods, and solutions. Policy & Internet.

3 Automated Accounts

In this study, we focus on accounts and tweets that discuss Swedish politics and the Swedish election. This entails that we analyze all tweets that are marked with hashtags related to Swedish politics and the Swedish election, for example, #svpol and #val2018. In total, we have examined almost 600 000 tweets that were produced by more than 45 000 accounts within the period from March 5th to August 20th 2018.

We have used machine learning to classify accounts as either genuine or automated. A genuine account is an account that is handled by a human without exhibiting automated behavior. A number of tweets that we analyzed belong to accounts which, during the period of study, were suspended or deleted as a result of violating the end-user agreements. The most common reasons for suspending a Twitter account is that the account spams, uses a false identity, or exhibits offensive behavior.⁵ Whether an account has been suspended by Twitter on the basis that it is automated is something we cannot determine using our classification. However, Twitter has efficient processes for terminating accounts with malicious automation and spamming.⁶



In our analysis we have categorized accounts into ‘genuine’, ‘automated’ and ‘suspended/deleted’. The largest group of accounts that tweets about the Swedish comes from genuine accounts.

In terms of the share of genuine and automated accounts, automated accounts are responsible for 11% of the content concerning Swedish politics and the election. Approximately 6% of the accounts that tweet about Swedish politics and the election are automated accounts. If we make the somewhat extreme assumption that all suspended/deleted accounts are automated, then 17% of accounts are automated, and 14% of the content is automated.

Increased Number of Automated Accounts The number of automated accounts tweeting about Swedish politics and the election has increased significantly in the period covered by our investigation. The number of automated accounts nearly doubled

⁵ See Twitter’s end-user licence agreement <https://help.twitter.com/sv/rules-and-policies/twitter-rules>

⁶ Read more on Twitters blog: https://blog.twitter.com/official/en_us/topics/company/2018/how-twitter-is-fighting-spam-and-malicious-automation.html

–from July to August – increasing from 668 to 1201. Figure 1 shows how the number of automated accounts has increased since March.

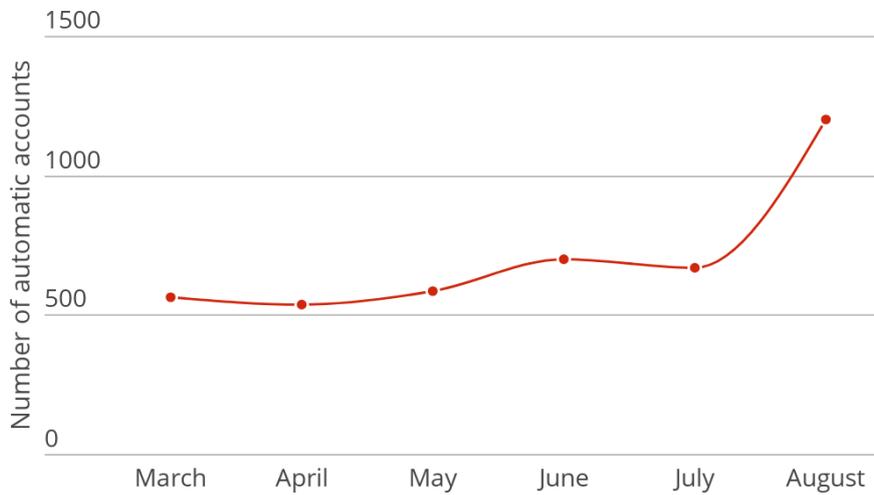


Figure 1. The number of automated accounts per month since March

We also looked at suspended accounts and found that the number of these accounts is decreasing the closer the election we get, see figure 2. The gradual decrease in the number of these accounts might be because it takes some time for Twitter before suspending an account.⁷

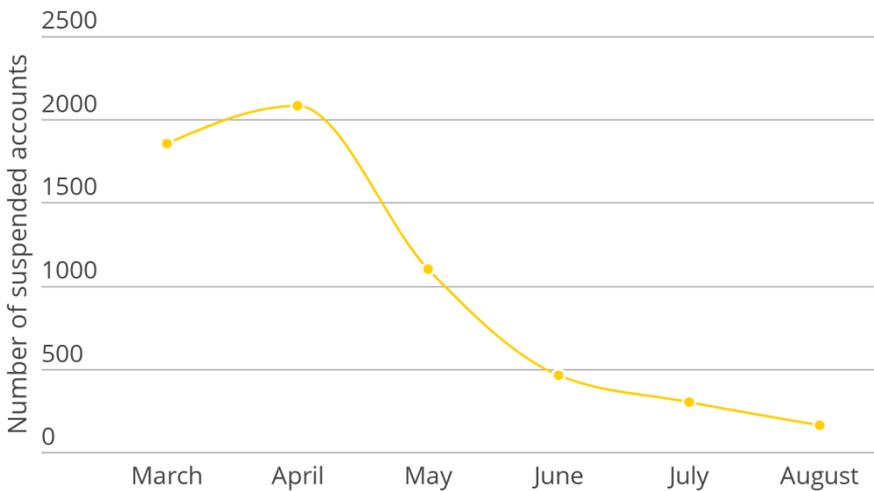


Figure 2. The number of accounts that have participated in discussions about the election and Swedish politics and that have been suspended or deleted.

⁷ More information about the process for suspending suspicious accounts can be found on Twitter's blog: https://blog.twitter.com/official/en_us/topics/company/2018/how-twitter-is-fighting-spam-and-malicious-automation.html

Assuming that all of the suspended or deleted accounts were automated (i.e. by putting together the automated and the suspended accounts), we can see that the number of automated (included the suspended) accounts increased by 40%, from 970 to 1363.

4 Content

We conducted two analyses of the content of the tweets in order to explore the domains to which different accounts link to and types/content of the messages spread by different types of accounts.

4.1 Links

In the period March 5th to August 20th 2018, all three categories of accounts link most frequently to *Expressen* (a tabloid newspaper) and *SVT* (Swedish public television). For genuine accounts, links to *Aftonbladet* (tabloid) and *Dagens Nyheter* (quality newspaper) are in third and fourth place, but automated accounts in these positions link instead to *Samhällsnytt* and *Fria Tider* (alternative or partisan news sites). Among the accounts that have been suspended or deleted the top ten sites linked to include *katerinamagasinet.se* (the website of Katerina Janouch) and *nordfront.se* (the website of the Nordic Resistance Movement). Table 1 shows the domains that are most linked to among the different types of accounts.

Table 1. The ten most linked to domains from different types of accounts

Genuine	Automated	Suspended/deleted
expressen.se	expressen.se	expressen.se
svt.se	svt.se	svt.se
aftonbladet.se	samnytt.se	samnytt.se
dn.se	friatider.se	youtube.com
samnytt.se	aftonbladet.se	friatider.se
gp.se	sverigesradio.se	nyheteridag.se
friatider.se	nyheteridag.se	aftonbladet.se
youtube.com	omni.nu	katerinamagasinet.se
nyheteridag.se	youtube.com	nordfront.se
omni.se	samtiden.nu	dn.se

4.2 Messages

To get a sense of the kind of content spread by automated accounts, we analyzed a sample of the total number of tweets spread during the period of August 1st to 20th 2018. More information about how the analysis was done can be found in section 6.3.

The same analysis was done for genuine accounts and automated accounts to allow a comparison between the two different account types. The type of content that is spread via automated versus genuine accounts is shown in figure 3. Overall, there are only minor differences between what genuine and automated accounts communicate. However, automated accounts communicate slightly more messages that are critical towards immigration, Islam, media and political parties.

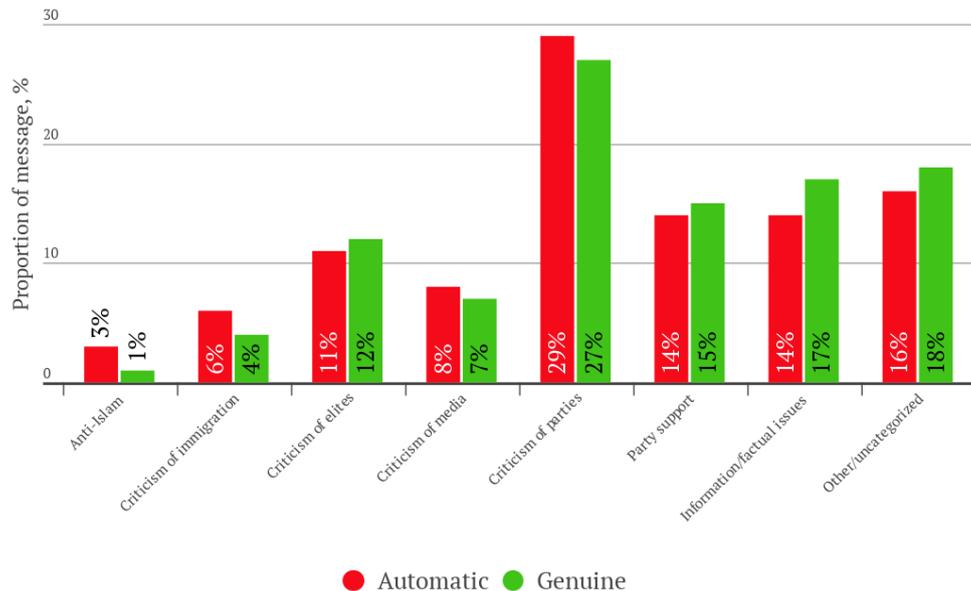


Figure 3. The share of different messages spread by different types of accounts

Support for Political Parties We examined what parties the automated and the genuine accounts expresses support for. The result is shown in figure 4. The results showed that both genuine and automatic accounts express a significant amount of support for AFS (Alternative for Sweden) – a newly formed party that is not represented in the Swedish parliament. Moreover, the Sweden Democrats is the other party for which both types of account express most support. Specifically, 47% of the bots and 28% of the genuine accounts expressed support for the Sweden Democrats. That is, automated accounts express more support for the Sweden Democrats, compared to genuine accounts.

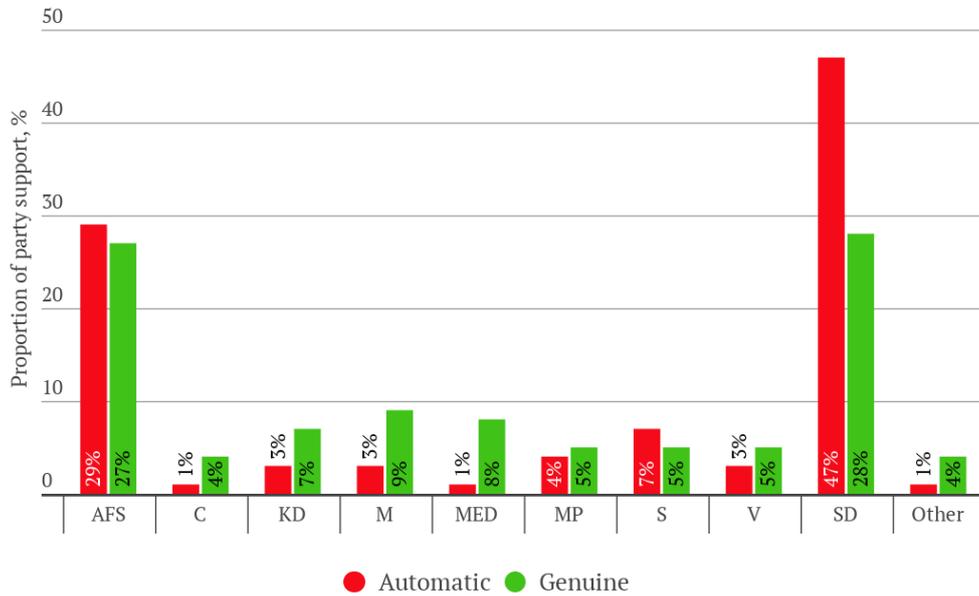


Figure 4. The share of parties for which different types of accounts express support

Criticism of Political Parties We also examined what parties that are criticized by automated and genuine accounts. The result, presented in in figure 5, show that the Social Democrats receive most criticism on both types of accounts. The automated accounts express more criticism towards the Left Party (Vänsterpartiet) and the Conservatives (Moderaterna) while the genuine accounts expressed double as much the criticism towards the Sweden Democrats, compared to the automated accounts.

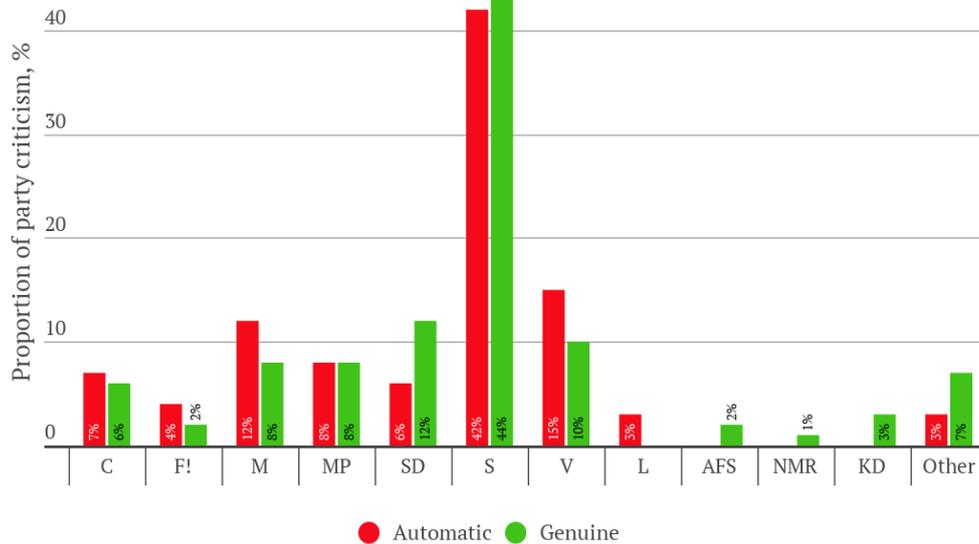


Figure 5 Distribution of the parties for which different accounts express criticism

5 Spread

One way to explore the networks of users that discuss the election and Swedish politics on Twitter is to examine how different accounts retweet each other. Therefore, we analyzed spreading during the period March through to July 2018 in terms of an associative network, see figure 6. In the figure, every account is indicated by a circle (or dot), 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 accounts' circle indicates how often an account has been retweeted by other users. The larger the circle, the more retweeted or popular the account. Green circles represent genuine, red circles represent automated, and yellow circles represent suspended or deleted accounts.

From March to July, there were 37000 accounts that discussed the election and Swedish politics on Twitter. The network of these accounts is shown in Figure 6.

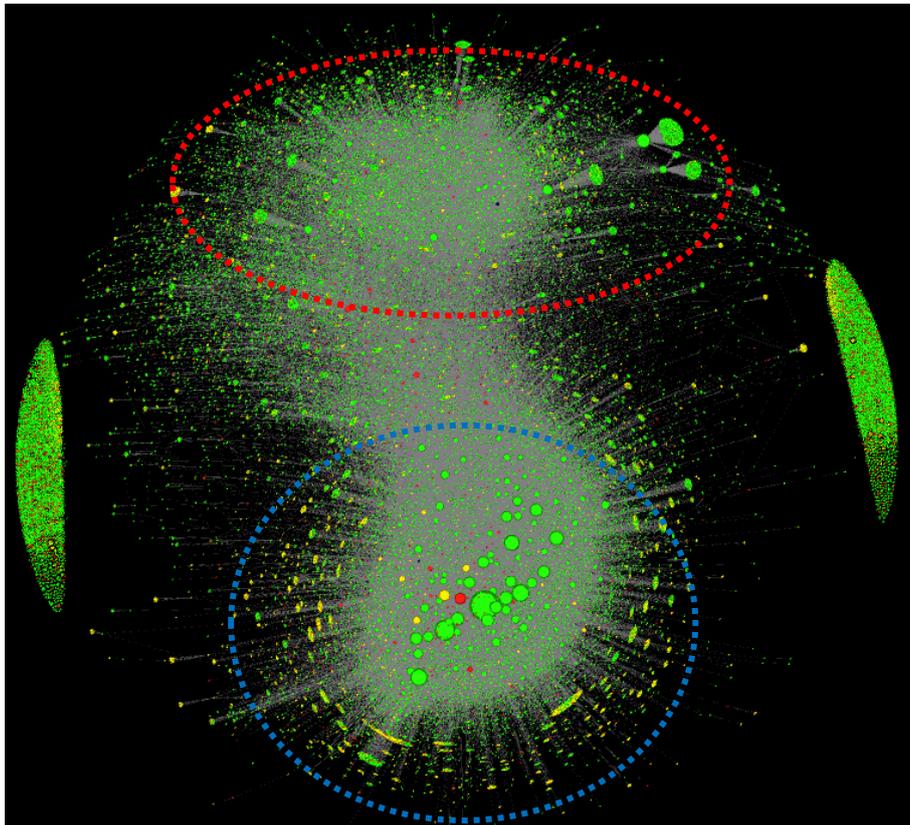


Figure 6. The network of accounts that have discussed the election and Swedish politics on Twitter where the two clusters (indicated in blue and red) are shown.

As can be seen in figure 6, there are two clusters on each side of the network (the right and the left). These two clusters consist of accounts that have tweeted about the election but that have not been retweeted and are therefore not connected to the main body of the network. The majority of the automated accounts can be found in these two clusters.

Manual content analysis revealed that the main body of the network is divided into two clusters. One of these can be found in the upper part of the network (red circle). This cluster seems to represent content with more egalitarian, progressive or liberal views. The second cluster can be found in the lower part of the network (blue circle). This cluster seems to represent content with more traditional, authoritarian or nationalistic views. In the area between the two larger clusters, there is a set of accounts that are retweeted by accounts within the red as well as the blue circle. These accounts consist mostly of news media that position themselves as politically neutral.

Notably, there are a number of smaller clusters of accounts on the edges that retweet isolated accounts. These clusters consist of accounts that show the same behaviors as others in the groups, and they retweet only larger and more popular accounts. In figure 7 we left out the genuine as against the automated accounts to illustrate where in the network the suspended and deleted accounts are located. As can be seen in figure 7, the majority of suspended/deleted accounts can be found in the blue cluster.

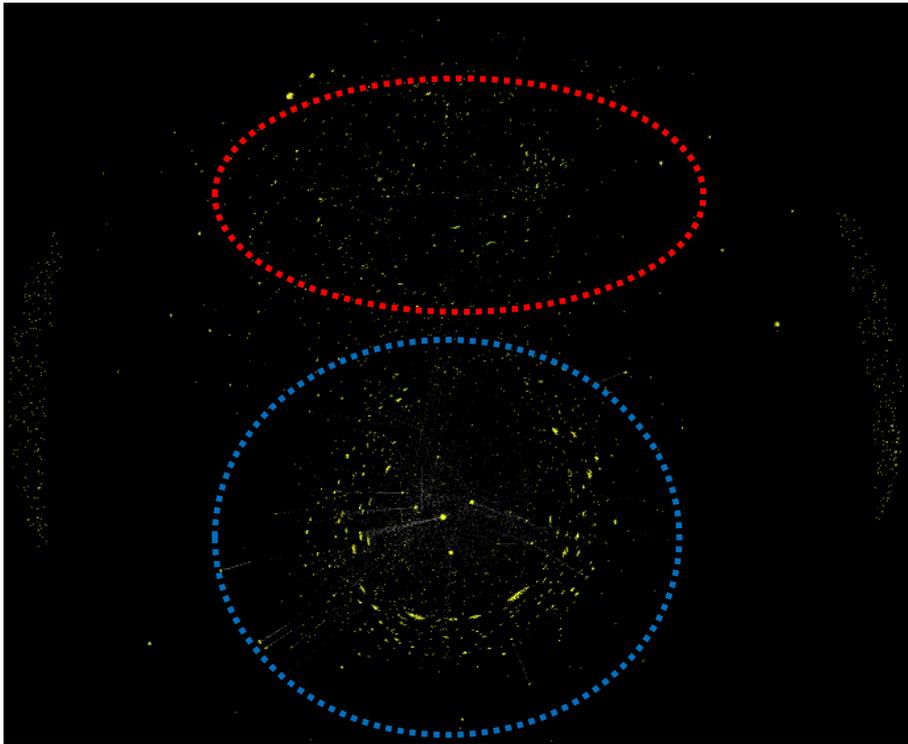


Figure 7. Suspended/deleted accounts that have discussed the election and Swedish politics on Twitter where two clusters (red and blue) are shown.

6 Method

This study is a result of a collaboration by Johan Fernquist, researcher at FOI (the Swedish Defence Research Agency), Lisa Kaati, senior researcher at FOI, Nazar Akrami, Associate professor, Department of psychology at Uppsala University, Ralph

Schroeder, professor in Social Science of the Internet at the Oxford Internet Institute, and Katie Cohen, psychologist and researcher at FOI.

The data collection was carried out by Björn Pelzer, researcher at FOI. Also, Emil Larsson, Magnus Sahlgren, and Fredrik Olsson at FOI contributed to the classification of data.

6.1 Data

In the period from March 5th to August 20th 2018, we collected all tweets with hashtags containing keywords that concern the election and Swedish politics, for example, #valet2018 or #svpol. Tweets that concern Swedish politics or the Swedish election but that do not contain our selected hashtags are outside the scope of our study. A total of 571719 tweets from 46396 accounts was collected from the Twitter streaming API.

The research in this study has been approved by the Central Ethical Review Board.

6.2 Classification of bots

To detect bots, we trained a model using machine learning that can recognize patterns that point to automated accounts. This was done by using data from accounts that we know to be automated. The model can then distinguish between automated and genuine accounts. We used a total of 140 different features of accounts to distinguish between automated and genuine ones. The classification was done using a random forest algorithm.⁸ We have in this study chosen to manually classify a set of accounts as genuine even if the account is automatic. Example of such accounts is news editorials and authorities that automatically publish information.

6.3 Content classification

We started by generating a list of categories after reviewing the content of a number of tweets. A thorough review of the list arrived at eight final categories (see below). Next, we randomly selected 935 tweets (472 from automated and 463 from genuine accounts) to be classified into the eight categories. The classification was carried out by six coders who did their coding independently of each other. The intercoder agreement (Fleiss Kappa⁹) indicated substantial agreement between the coders (0.61).

⁸ More information about the classification of algorithms can be found in J.Fernquist, L. Kaati, R. Schroeder (2018) Political Bots and the Swedish General Election. Submitted for publication.

⁹ Fleiss is a statistical method for assessing the reliability of the agreement between a certain number of assessors in the attribution of classifications to certain categories. For more information, see <http://www.real-statistics.com/reliability/fleiss-kappa/>

Table 2. Categories and their definition used to classify tweets

Category	Definition
<i>Media criticism</i>	Criticism of journalists, journalism or media.
<i>Criticism of elites</i>	Criticism of the government, or persons or organizations that are regarded as influential political elite.
<i>Party support</i>	Support for one or more parties. Subcategory: any political party.
<i>Criticism of parties</i>	Criticism of one or more parties. Subcategory: any political party.
<i>Anti-Islam</i>	Criticism of Islam, Muslims, or the Islamification of Sweden.
<i>Information, factual issues</i>	Discussion or information concerning general issues, for example, information about certain topics or announcements.
<i>Criticism of immigration</i>	Criticism of immigration, asylum seekers and refugees.
<i>Other/Uncategorized</i>	Content that does not fall into any of the categories above. This category also includes messages that have been deleted or that cannot be categorized.

All tweets were coded into the defined categories and in some cases into subcategories. If a tweet was deleted or no longer available via Twitter's web interface, the tweet was coded as deleted. Tweets that expressed both support and criticism of a party were coded as expressing support of a party.

7 Discussion and Conclusions

Among the accounts that we analyzed, the proportion of bots was between 6 and 17%, depending on whether we include accounts that were suspended or deleted among our bots. The number of bots that tweet about Swedish politics or the election has almost doubled between July and the period August 1st to 20th. It seems likely that the number of bots will continue to increase before the election, and that they will spread the same types of messages as they have done up to now.

In our analysis of which types of domains accounts most commonly link to, we can observe that bots link more to *Samhällsnytt* and *Fria Tider* than genuine accounts. Also the domains *katerinamagasin.se* and *nordfront.se* can be found among the top 10 sites

among the suspended or deleted accounts. Notably, the Twitter has suspended the account of the Nordic Resistance Movement (the owner of nordfront.se).

As regard to the sympathies of bot and genuine accounts, the Sweden Democrats receive the largest share of support from both types of accounts. A significant difference, however, is that 47% of the bots expressed support for Sweden Democrats compared to a 28% of the genuine accounts. Moreover, both genuine and automatic accounts express a significant amount of support for AFS (Alternative for Sweden) – a newly formed party that is not represented in the Swedish parliament.

Regarding criticism of political parties, the differences between bots and genuine accounts are negligible. However, criticism of the Left Party and the Conservatives is somewhat larger among bots, while criticism of the Sweden Democrats is more common among genuine accounts. More importantly, there was more than 40% of both the bots and genuine accounts expressed criticism toward the Social Democrat party. Analyses of the spread of messages among the 37 000 accounts shows that the majority of accounts can be divided into two clusters – a one with more authoritarian, traditional, and nationalist views and one with more egalitarian, progressive or liberal views. These analyses also show that a majority of the suspended or deleted accounts can be located in the cluster that expresses more authoritarian, traditional, and nationalist views.

From the above, we can draw the tentative conclusion that the cluster that expresses more authoritarian and nationalist views contains more bots than in the cluster that expresses more egalitarian, progressive or liberal views. This result points in the same direction as the other results; namely, that bots tend to share immigration critical material from partisan websites, and to a large extent express support for the Sweden Democrats. In other words, the results point towards how bots that tweet about Swedish politics and the election of 2018 tend to spread a nationalist or immigration-critical message. A question that this analysis cannot answer is who is behind these bots, or is their agenda. We also acknowledge that our conclusions are limited to the hashtags we included in our study.

7.1 Concluding remarks

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 automated accounts analyzed here fit into a larger picture of media use in Sweden. It would also be necessary to analyze who is behind these automated accounts and what the aims are of spreading messages from automated accounts. While answering these question is difficult, it is clear that the use of automated accounts for spreading various types of messages has increased markedly the closer we come to the election. Thus, there seem to be attempts to influence. 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.¹⁰ 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, this study contributes to better awareness of attempts to exercise influence using automated accounts. Ultimately this awareness would allow citizens to make decisions unaffected by the influence of Twitter-bots.

¹⁰ W. Wood, & J. M. Quinn (2003). Forewarned and forearmed? Two meta-analysis syntheses of forewarnings of influence appeals. *Psychological Bulletin*, 129, 119-138.

FOI is one of Europe's leading defence and security research agencies and acts within the entire span of defence and security issues. We perform security analyses and assessments of different kinds of threats. We are leaders in submarine research and in research relating to explosive materials. We develop decision-support systems and systems where people, technology, and organisation must interact. Our areas of research also include aero systems, ICT security, radar, laser, and other sensor systems, as well as protection against hazardous substances.

FOI
Defence Research Agency
SE-164 90 Stockholm

Tel +46 8 5550 3000
www.foi.se

