Bots By Topic: Exploring Differences in Bot Activity by Conversation Topic

Kurt Wirth
American University
Washington, D.C.
pubs@kurtwirth.com

Ericka Menchen-Trevino
American University
Washington, D.C.
pubs@ericka.cc

Ryan T. Moore
American University
Washington, D.C.
rtm@american.edu

ABSTRACT

This study introduces a new tool to compare bot levels in real-time across conversation topics or hashtags. With the data collected, we measured higher levels of bot activity in some topics of conversation as compared to others and propose a novel application of bot detection analysis to advance research in this fast-changing field.

CCS CONCEPTS

• **General and reference → Measurement**;

KEYWORDS

Twitter, social bots, Twitter bots, bots, botscan, BotOMeter, computational propaganda, misinformation, disinformation

ACM Reference format:
Kurt Wirth, Ericka Menchen-Trevino, and Ryan T. Moore. 2019. Bots By Topic: Exploring Differences in Bot Activity by Conversation Topic. In *Proceedings of 10th International Conference on Social Media Society, Toronto, Canada, July 2019 (SMS 2019),* 6 pages.
https://doi.org/10.475/123_4

1 INTRODUCTION

Social media play an important function in democratic dialogue. Donald Trump relied heavily on direct access to voters in his 2016 campaign, which led to his successful nomination in 2017, and created a media ecosystem that fed off a variety of outlets and platforms [27]. The process Trump employed is formally known as "intermedia agenda setting" where information in one medium, such as social media, influences the topics covered in another medium, such as cable news or print news [18]. As information and misinformation enter this hyper-connected media environment, the features of social media, including the ability to create, share, and react to content in real time, impact media writ large [1].

The continued and rapid evolution of social media results in a competition that drives the marketplace of ideas and is bolstered by the flow of information [14]. Those who control information — including how it is created, propagated and consumed — control significant power to influence societal discourse and could potentially use their control to disrupt the idea economy. Restricting and gathering data from information flows can be weaponized to exploit users, damage an individual's reputation, and manipulate perspectives of individuals and/or groups, as well as other damaging purposes [3]. Twitter, specifically, has proven potent in its ability to shift national discussion. Hashtags like #BlackLivesMatter, #MAGA, and #MeToo have become rallying cries for social commentary, deliberative action, and counter-public creation. Online conversations, particularly on Twitter, can be manipulated by automated accounts known as bots. Though the creators of social media platforms may have viewed the creation of those platforms as an attempt to move closer toward an inclusive marketplace of ideas, existing powerful entities have continuously attempted to consolidate their influence using social media. Bots have been deployed by state actors to help citizens during crises as well as to create and disseminate propaganda to build support for incumbent regimes [28]. Misinformation spread by bots is threaded into existing conversations and creates confusion about topics ranging from politics to health [21]. Research has shown that bots manipulate the flows of conversations to the extent that they affect individuals and, as a result, the flow of information online [20].

For bot creators seeking profit, cheap platforms available to reach large audiences help them profit from selling conceptions of reality to those that most benefit from perceiving reality in the way offered by the seller [2]. Other bot networks deploy campaigns to sabotage election processes or otherwise introduce misinformation in an effort to reduce trust in democratic processes and influence international politics [28]. Inauthentic information, often called "fake news", can then be disseminated through various media systems, sometimes resulting in serious real-world effects [13]. When information warfare is directed at democratic institutions such as elections, the health of the nation is at stake [5]. Informational ecosystems corrupted with ongoing creation and sharing of fake news weaken the structure of democratic government [4].

2 IDENTIFYING BOTS AND THEIR EFFECTS

Much of the societal disruption caused by bots is done through the creation and dissemination of misleading or untrue stories disguised as news. The term "fake news" encompasses a broad swath of meanings depending on the context of its use. In *Disinformation, 'Fake News', And Influence Campaigns On Twitter*, Matthew Hindman and Vlad Barash only analyzed sites that "overwhelmingly publish false articles" [19]. The authors produced a thorough examination of bot behavior by retroactively following the behaviors of accounts confirmed by Twitter to be bots, using network analysis to show the densities of their relationships and behaviors across the nodes of the network. Their analysis includes sites like InfoWars and Sputnik News, but excludes sites that were merely strongly biased or ideologically extreme. Using this narrow definition, the researchers

SMS 2019, July 2019, Toronto, Canada

ACM ISBN 123-4567-24-567/08/06.
https://doi.org/10.475/123_4

found that bots are not only propagating fake news and other forms of misinformation, the bots are being used in a concerted effort by those that control them. Hindman and Barash establish that dense bot networks, or "bot armies", are designed to create and amplify misinformation in an attempt to manipulate the online discourse. This method is tremendously useful but leaves room for improvement. For example, it is possible that the sample provided by Twitter simply was not representative of the larger Twitter bot ecosystem. The study is limited in scope, focusing on a specific non-generalizable sample to detail a quantitative case study of how specific types of bots behave in specific circumstances. Larger questions like where bots are centralizing right now, are there differences in their localization within conversations, and what sorts of bots are in what sorts of conversations are difficult to answer with this type of study because of its lack of generalizability.

Samuel Woolley has similar findings in *Automating Power: Social Bot Interference In Global Politics*. Woolley's study skirts the troublesome task of identifying bots by using media coverage of bot activity to gather data. Seeking to uncover specific tactics and strategies used by bots, his content analysis finds that "computational propaganda, proliferated by political actors using political bots, enables control globally" [28]. According to his analysis, inflated popularity of issues and user audience size and spam flooding are among tactics used by bot armies. Additionally, Woolley finds that bots are often deployed by large organizations and regimes when under pressure. In *Algorithms, Bots, And Political Communication In The U.S. 2016 Election: The Challenge Of Automated Political Communication For Election Law And Administration*, Philip Howard, Woolley, and Ryan Calo focus in on political bots. They build a narrative of bots and their tactics by connecting studies to present a logical chain of reasoning behind bot creators' motivations. Both studies are crucial to countering bot activity, however they are limited in scope.

Identifying bots on social media is difficult. A number of factors confound the process. Applying algorithms to identify automated accounts is particularly troubling for researchers because no genuine sample of the population of bots can be identified. Without true and absolute knowledge of whether any given account is fully automated, this area of research lacks what is often called a "ground truth". This, in turn, limits the generalizability of studies in this field [10]. Furthermore, even advanced and highly accurate algorithms may more quickly identify accounts as bots with specific traits (i.e., posting every hour, on the hour) while tending to ignore others (i.e., stock photography used for profile picture). Without exhaustive testing, it is difficult for researchers to detect these biases. Additionally, a program that performs accurately when it is first created, as the algorithms were tested on then-recent data, may degrade as bot makers improve their own code to avoid detection and account suspension. The ability to detect constantly changing bots is an uphill battle for researchers. Lastly, methodologies driven by bot-detection software are limited by the packages' inability to access deleted or suspended accounts, which limits retrospective analysis [22]. Despite these concerns, studies using tools that deploy computational bot detection methodologies have consistently shown results pointing to the efforts of bot creators to manipulate discourse on social media [12]. Periodic rigorous validation tests of bot detection software are needed, but research in this area must continue because automated accounts are shaping conversation in our societies around the world today.

BotOMeter, formerly named BotOrNot, was first published by five researchers led by Indiana University's Clayton Davis in 2016 [9]. BotOMeter was the first-ever publicly-accessible bot detection tool, originally developed to compete in a challenge issued by the United States Defense Advanced Research Projects Agency (DARPA) [23]. While previous attempts to create bot detection tools relied on by-hand verification [8] or earlier data sets [6], BotOMeter analyzes over 1,000 features of every account and was trained using over 30,000 confirmed bot and non-bot accounts. Currently, BotOMeter is researchers' premier tool in efforts to identify bots on Twitter. Attempts to better understand bot behavior focus on identifying individual bot accounts [6], the effect of interacting with bots [26], or the content produced by the bot accounts [11]. Programs like DeBot [7] and BotOMeter [9] are designed to identify accounts individually.

In summary, communication technologies are used to create added uncertainty in political deliberation. People struggle to distinguish authentic dialogue from inauthentic dialogue driven by disinformation campaigns. Creating a tool to track bots at the conversation-level and in real-time is the focus of this project. Technology has been used to create the problem of bot influence in social media discourse and our botscan tool aims to assist researchers, journalists, and the general public to establish if a deliberation in online discourse is driven by authentic or inauthentic communication. Our tool, botscan, allows researchers to easily measure bot activity at the conversation level in real-time for the first time and is supported by the BotOMeter algorithms.

3 RESEARCH QUESTIONS

(1) How does bot activity volume differ on Twitter between different types of conversations?
(2) How do bot account "types" differ between topic categories as well as a random sample?

4 BOTSCAN

To address these questions we created the botscan tool [1] for R [24], to implement topic-level monitoring for bot activity on Twitter. Since botscan does not rely on its own algorithm, it uses BotOMeter to identify bots because it is the best-tested engine currently available for real-time Twitter data. Originally released in 2014 and later published in 2016, BotOMeter analyzes over 1,000 metadata features of each account tested, including network, user, friends, temporal, content, and sentiment features. Davis and his colleagues validated the tool using ten-fold cross-validation, obtaining an area under the ROC curve of 0.95. BotOMeter does not sort or attempt to categorize types of bots and instead focuses solely on the identification of automated influence in a given account.

The botscan tool reviews every account in any Twitter-compatible search string, including a hashtag (#MeToo) or a term without a hashtag (Alyssa Milano). We refer to the tweets that are returned when one searches for strings as *conversations* or *topics*. Like other Twitter-based software, botscan requires the user to acquire their

[1] Available at https://github.com/kurtawirth/botscan

own unique Twitter keys, as well as a BoM Mashape key, [2] which requires a Mashape account. An upgraded and free version of the BotOMeter API key allows a user to increase their rate limit substantially. Given the rate limits of Twitter and BotOMeter, a search processes approximately 10-12 accounts per minute. [3]

Our tool uses Twitter's streaming application programming interface (API), which can only gather tweets going forward in time. Though Twitter's search API allows for retrieval of past data up to about one week old, it provides incomplete data that has been vetted using quality assurance software, which removes a significant amount of bot content. However, botscan allows for the use of the search API if the user chooses. Our pre-testing using search API data found much lower levels of bot activity compared to the streaming API. While the lack of ability to collect past data is a limitation, it allows for real-time bot detection in a fast-paced information environment.

The botscan tool accepts three primary parameters: 1) the search term, 2) the number of tweets that contain the search term to gather, and 3) the threshold for bot detection by BotOMeter.

5 METHODS

We selected five topic areas to explore: conservative politics; liberal[4] politics; both conservative and liberal politics; top news; and top trends unrelated to news. Using several search terms in each topic area, we examine trends in the fraction of a conversation produced by bots (see Table 2). We selected current topics that had a sufficient volume of data to collect tweets within a 24 hour time period. To select a conversation within each area, we first chose an appropriate conversation within the worldwide trending top 10. If no conversations within worldwide trends were appropriate for one or more of our topics, we then selected trends from the American nationwide trending top 10. If again no conversations were deemed appropriate for one or more of our topics, we selected a known high-volume conversation within that topic (i.e., #MAGA). Researcher-selected conversations were tested to ensure their current volumes were high enough for collection of sufficient sample sizes. The selection method of each search term is identified in Table 2 (global trend, U.S. trend, or researcher-selected).

To determine the number of tweets per search term to collect, we began by testing botscan's reliability in providing a stable percentage of bots in a conversation given different sample sizes. After collecting 10 rounds of two thousand random tweets via Twitter's streaming API, successive tests revealed the bot prevalence of samples within each round with sizes of 100, 200, 500, 1,000, and 2,000 respectively. The results are displayed in Table 1.

In order to test whether some samples from a given conversation are significantly more or less generated by bots than others, we collect 10 rounds of data, calculate the fraction of the conversation generated by bots, and examine the distributions for outliers. We further examine samples of different sizes, sensitive to the fact that researchers may prefer smaller samples due to the infrequency of some search terms, or due to computational constraints. We find no evidence of outlying samples, either across rounds or in smaller samples. We find that the means roughly follow a normal distribution, across rounds and sample sizes, as we would expect from the sampling distribution of the mean. Below, we examine the samples of size 2,000, due to the decreased variability of the test data with this sample size.

We used a bot identification threshold of 0.43 based on previous research on the Complete Automation Probability scale, or CAP [16]. BotOMeter does not sort or attempt to categorize types of bots. Twitter, in fact, openly allows many types of automated accounts. Accounts that "broadcast helpful information", engage automatically with followers, and generally "help people" fall into this category [25]. These types of accounts are measured against the same yardstick as accounts that tweet spam, hijack political hashtags, and manufacture virality through retweet networks. Because we are interested only in what Gorwa calls "social bots", or automated accounts meant to engage on social media [15], we conducted a qualitative review of 50 accounts identified as bots in each category to ascertain the various types of bots prevalent in each conversation and how they differ. (See Figure 1)

6 RESULTS

Bot levels by round for random samples are found in Table 1. Each round consists of 2,000 tweets, where the first 100 are analyzed, followed by the first 200 (inclusive), etc. Bot levels by conversation are shown in Table 2.

Table 1: botscan Validation. Cells contain percentage of conversation produced by bots.

Round	$n = 100$	$n = 200$	$n = 500$	$n = 1000$	$n = 2000$
1	8.0%	8.0	7.0	7.6	6.5
2	5.0	5.0	3.8	3.8	4.6
3	7.0	10.5	8.6	7.6	6.1
4	9.0	9.0	10.2	8.5	6.9
5	7.0	6.5	6.2	6.6	6.3
6	4.0	3.5	4.6	5.6	5.5
7	6.0	6.0	5.6	7.1	6.5
8	7.0	4.5	5.4	4.9	5.8
9	6.0	6.0	7.2	6.9	6.0
10	4.0	6.5	6.6	5.6	6.3
Mean	6.30%	6.55	6.52	6.42	6.04
St. Dev.	1.64%	2.11	1.88	1.43	0.64
Pct. From Mean	25.97%	32.28	28.86	22.25	10.61

We tested conversations for bots in three rounds. Each round of data acquisition was completed over the course of 24 hours and consisted of 2,000 tweets per conversation per round. The results and conversations are found in Table 2. A summary of results by type of conversation are found in Table 3.

We performed a χ^2 test of independence to examine the relation between conversation type and whether a given conversation's

[2]Available with instructions, respectively, at https://apps.twitter.com/ and https://market.mashape.com/OSoMe/botometer.

[3]Expected updates and additional functionality can be viewed on the tool's GitHub "Issues" page.

[4]The term liberal is used here as it is used in the U.S. political system, similar to center-left in European politics.

accounts are likely to be labeled as a bot by botscan. The relation between these variables was statistically significant $\chi^2(392, N = 25{,}391) = 3{,}139.8$, $p<.01$. Bot activity appears significantly denser within conservative political conversations. We hypothesize that the increased percentage of bots in trending but non-political conversations is attributed to a larger overall Twitter volume of non-political bot types (e.g., spam and niche). Conservative conversations contain on average 254% as many bots as liberal conversations and 306% as many bots as cross-spectrum conversations.

Following the data gathering stage, the top 50 accounts within each conversation type were inductively categorized according to their apparent purpose. The apparent bot purposes we categorize include

- Political: Shares mostly politically-motivated posts, and potentially includes political messaging in bio.
- Spam: Seemingly unrelated content and/or repeatedly asking for an action such as clicks or views.
- Niche: Promotes specific topic regardless of usefulness and asks for little or no action (e.g., clicks or views).
- Porn: Predominantly posts material pornographic in nature.
- Organization: Seems to belong to an organizational entity.
- Informative: Shares information meant to be useful while rarely or never asking for an action.
- Suspended: Account was suspended at the time of data collection or deleted after tweets were pulled and before late January, 2018.
- Undetermined: Accounts that have too little activity and/or a language barrier that renders accurate classification impossible.

We performed a χ^2 test of independence to examine the relation between conversation type and the category a bot was assigned. The relation between these variables is statistically significant $\chi^2(28, N = 250) = 181.1$, $p<.01$. Political bots are significantly more likely to appear in political conversations (here represented by the cross-spectrum, conservative, and liberal categories) than non-political ones.

Table 2: Bots By Conversation

Round	Topic	Conversation	Percent Bots
1	Conservative	"#Bolsonaro2022"[a]	11.15%
1	Liberal	"#WheresMitch"[b]	1.20
1	Cross-Spectrum	"Cohen"[b]	2.15
1	Trending/News	"Mary Oliver"[a]	2.15
1	Trending/Not News	"#AllStars4"[a]	5.15
2	Conservative	"#GloboLixo"[a]	2.60
2	Liberal	"#Kamala2020"[c]	1.95
2	Cross-Spectrum	"Erykah Badu"[b]	0.90
2	Trending/News	"#Itzy"[a]	5.25
2	Trending/Not News	"#TheVoiceKids"[a]	4.45
3	Conservative	"#MAGA"[c]	5.35
3	Liberal	"#Resist"[c]	2.70
3	Cross-Spectrum	"Venezuela"[a]	1.65
3	Trending/News	"Dourado"[a]	0.35
3	Trending/Not News	"#NationalComplimentDay"[a]	3.65

[a]Trending in world. [b]Trending in U.S. [c]Researcher-selected.

Table 3: Bots By Conversation, Summary

Topic	Average Percent Bots
Conservative	6.37%
Liberal	1.80
Cross-Spectrum	1.57
Trending/News	2.58
Trending/Not News	4.42

Figure 1: Bots By Type

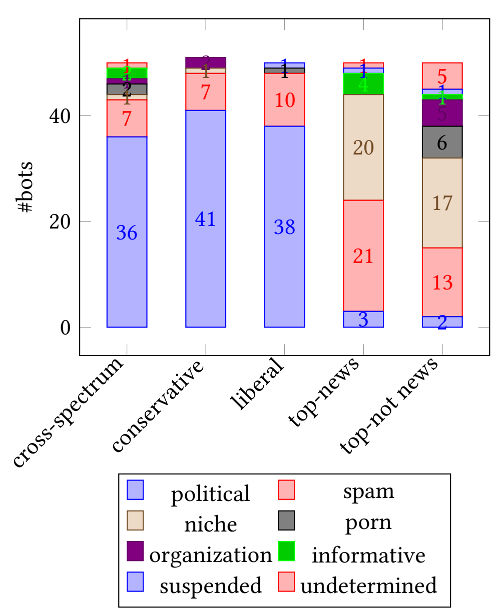

7 DISCUSSION

Bot armies can weaken society's ability to verify what is true and what is not, and the worst types of stress on a democracy are unpredictability and uncertainty [4]. The botscan tool helps conceptualize bot activity as part of a holistic interconnected and interactive ecosystem. New approaches to bot monitoring methodology introduced by botscan have the potential to help researchers, journalists and citizens understand bot account activity in a particular conversation in real-time.

Using botscan, this study shows that there are more bots active in the conservative deliberative conversations than in a random sample, and that the types of bot accounts differ significantly across various conversations, which is consistent with previous research [17]. Further research into a much larger variety of political conversations over a longer time period and across platforms is needed to make any general claim about bot presence by political ideology, which is a contested and difficult to define concept particularly across countries. In our small study, however, the results are clear. Though botscan's results from the #Bolsonaro2022 search are extreme, the lowest bot presence detected in any conservative conversation (#GloboLixo at 2.60%) was the only conservative result lower than the highest level of bot presence in liberal conversations (#Resist at 2.70%). Removing #Bolsonaro2022 produces an average bot presence of 3.975%, which is 121% higher than the 1.80% average for liberal conversations. This demonstrates the potential benefits in understanding the ecology of bot behavior by comparing topics.

Researchers now have an open-source tool to gather bot activity data on the conversational level. This enables data analysis on a wide variety of questions including: bot behavior differences between conversation types, bot propagation over time, bot behavior changes over time (in response to real-world events, for example), and bot "types" composition differences per conversation types.

Likewise, journalists could use this type of tool to prevent their coverage of social media trends from being influenced by bots, and they could cover trends in bot behavior as it shifts from one topic of conversation to another. This kind of information could enable a new media literacy of real-time bot behavior that is necessary in the social media age.

8 CONCLUSION AND FUTURE RESEARCH

The introduction of botscan enables a new and important perspective and further exploration of how discourse online is vulnerable to manipulation through interconnectivity and interactivity. This paper contributes a method to study conversational - rather than individual - bot behavior, and implements a new tool, botscan, to explore another level of bot analysis. Our tool produces the first estimates at the conversational level about automated social media accounts. However, botscan's reliance on the BotOMeter algorithm and limitation to Twitter limit its explanatory potential. When monitoring bot activity, there is simply no "ground truth", making verification and measurement of accuracy difficult.

Other confounds exist. Specifically, the volume of conversational traffic for a topic may bias botscan's results. Indeed, our findings suggest this may be so. Our study found the average for trending conversations universally fell below the 6.04% average of 10 tests of 2,000 random samples each. Likewise, if the level of bot presence in ongoing conversations - like #MAGA - is relatively consistent, fluxes in human-led conversational traffic skew botscan results of that conversation. Future research should investigate how conversational volume affects botscan results and perhaps produce a more reliable metric with which conversations could be compared.

Despite these considerations, botscan - and the methods it enables for measuring and tracking bot activity - have the potential to expand knowledge in this field. Early findings, including this study, have suggested that conservative Twitter conversations contain more politically-motivated bots than those on other topics. Using and expanding upon methods like that used in our study, monitoring the nature of these accounts and their activities could reveal a great deal about the intentions of bot creators and how to counter such efforts in order to protect public discourse from bad actors.

REFERENCES

[1] Luca Maria Aiello, Martina Deplano, Rossano Schifanella, and Giancarlo Ruffo. 2012. People Are Strange When You're a Stranger. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* (2012-06-04).
[2] Hunt Allcott and Matthew Gentzkow. 2017. Social Media and Fake News in the 2016 Election. 31, 2 (2017), 211–236. https://doi.org/10.1257/jep.31.2.211
[3] Mark Andrejevic. 2007. *ISpy: Surveillance and Power in the Interactive Era.* University Press of Kansas. Google-Books-ID: S2y5AAAAIAAJ.
[4] Yochai Benkler. 2006. *The Wealth Of Networks: How Social Production Transforms Markets And Freedom.* Yale University Press. OCLC: ocm61881089.
[5] Alessandro Bessi and Emilio Ferrara. 2016. Social Bots Distort The 2016 U.S. Presidential Election Online Discussion. 21, 11 (2016). http://uncommonculture.org/ojs/index.php/fm/article/view/7090
[6] Qiang Cao, Michael Sirivianos, Xiaowei Yang, and Tiago Pregueiro. 2012. Aiding The Detection Of Fake Accounts In Large Scale Social Online Services. In *NSDI'12 Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation.*
[7] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2016. Identifying Correlated Bots in Twitter. In *Social Informatics*, Emma Spiro and Yong-Yeol Ahn (Eds.). Vol. 10047. Springer International Publishing, 14–21. https://doi.org/10.1007/978-3-319-47874-6_2
[8] Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia. 2012. Detecting Automation of Twitter Accounts: Are You a Human, Bot, or Cyborg? 9, 6 (2012), 811–824. https://doi.org/10.1109/TDSC.2012.75
[9] Clayton Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A System To Evaluate Social Bots. https://arxiv.org/pdf/1602.00975.pdf
[10] Juan Echeverría, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Shi Zhou. 2018. LOBO – Evaluation of Generalization Deficiencies in Twitter Bot Classifiers. (2018). arXiv:1809.09684 http://arxiv.org/abs/1809.09684
[11] Chad Edwards, Autumn Edwards, Patric R. Spence, and Ashleigh K. Shelton. 2014. Is That A Bot Running The Social Media Feed? Testing The Differences In Perceptions Of Communication Quality For A Human Agent And A Bot Agent On Twitter. 33 (2014), 372–376. https://doi.org/10.1016/j.chb.2013.08.013
[12] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The Rise Of Social Bots. 59, 7 (2016), 96–104. https://doi.org/10.1145/2818717
[13] Marc Fisher, John Woodrow Cox, and Peter Hermann. 2016. Pizzagate: From Rumor, To Hashtag, To Gunfire In D.C. (2016). https://www.washingtonpost.com/local/pizzagate-from-rumor-to-hashtag-to-gunfire-in-dc/2016/12/06/4c7def50-bbd4-11e6-94ac-3d324840106c_story.html
[14] Christian Fuchs. 2007. *Internet And Society: Social Theory In The Information Age.* Routledge. https://doi.org/10.4324/9780203937778
[15] Robert Gorwa and Douglas Guilbeault. 2018. Unpacking The Social Media Bot: A Typology To Guide Research And Policy. (2018). https://doi.org/10.1002/poi3.184
[16] John Gramlich. 2018. How We Identified Bots On Twitter. http://www.pewresearch.org/fact-tank/2018/04/19/qa-how-pew-research-center-identified-bots-on-twitter/
[17] Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. Fake News On Twitter During The 2016 U.S. Presidential Election. 363, 6425 (2019), 374–378. https://doi.org/10.1126/science.aau2706
[18] Raymond A. Harder, Julie Sevenans, and Peter Van Aelst. 2017. Intermedia Agenda Setting In The Social Media Age: How Traditional Players Dominate The News Agenda In Election Times. 22, 3 (2017), 275–293. https://doi.org/10.1177/1940161217704969

[19] Matthew Hindman and Vlad Barash. 2018. Disinformation, 'Fake News', And Influence Campaigns On Twitter. (2018). https://kf-site-production.s3.amazonaws.com/media_elements/files/000/000/238/original/KF-DisinformationReport-final2.pdf

[20] Kathleen Hall Jamieson. 2018. *Cyberwar: How Russian Hackers And Trolls Helped Elect A President What We Don't, Can't, And Do Know*. Oxford University Press. Google-Books-ID: bh1rDwAAQBAJ.

[21] Fang Jin, Wei Wang, Liang Zhao, Edward Dougherty, Yang Cao, Chang-Tien Lu, and Naren Ramakrishnan. 2014. Misinformation Propagation in the Age of Twitter. 47, 12 (2014), 90–94. https://doi.org/10.1109/MC.2014.361

[22] Denis Stukal, Sergey Sanovich, Richard Bonneau, and Joshua A. Tucker. 2017. Detecting Bots on Russian Political Twitter. 5, 4 (2017), 310–324. https://doi.org/10.1089/big.2017.0038

[23] V.S. Subrahmanian, Amos Azaria, Skylar Durst, Vadim Kagan, Aram Galstyan, Kristina Lerman, Linhong Zhu, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. The DARPA Twitter Bot Challenge. 49, 6 (2016), 38–46. https://doi.org/10.1109/MC.2016.183

[24] R Core Team. 2014. R: A Language and Environment for Statistical Computing. http://www.R-project.org/

[25] Twitter. 2018. Automation Rules. https://help.twitter.com/en/rules-and-policies/twitter-automation

[26] R. Wald, T. M. Khoshgoftaar, A. Napolitano, and C. Sumner. 2013. Predicting Susceptibility To Social Bots On Twitter. In *2013 IEEE 14th International Conference on Information Reuse & Integration (IRI)*. 6–13. https://doi.org/10.1109/IRI.2013.6642447

[27] Chris Wells, Dhavan V. Shah, Jon C. Pevehouse, JungHwan Yang, Ayellet Pelled, Frederick Boehm, Josephine Lukito, Shreenita Ghosh, and Jessica L. Schmidt. 2016. How Trump Drove Coverage to the Nomination: Hybrid Media Campaigning. 33, 4 (2016), 669–676. https://doi.org/10.1080/10584609.2016.1224416

[28] Samuel C. Woolley. 2016. Automating Power: Social Bot Interference In Global Politics. 21, 4 (2016). http://journals.uic.edu/ojs/index.php/fm/article/view/6161