Bots and Automation over Twitter during the First U.S. Presidential Debate

Bence Kollanyi
Corvinus University
kollanyi@gmail.com
@bencekollanyi

Philip N. Howard
Oxford University
philip.howard@oii.ox.ac.uk
@pnhoward

Samuel C. Woolley
University of Washington
samwooll@uw.edu
@samuelwoolley

ABSTRACT

Bots are social media accounts that automate interaction with other users, and political bots have been particularly active on public policy issues, political crises, and elections. We collected data on bot activity using the major hashtags related to the U.S. Presidential debate. In this brief analysis we find that (1) Twitter traffic on pro-Trump hashtags was roughly double that of the pro-Clinton hashtags, (2) about one third of the pro-Trump twitter traffic was driven by bots and highly automated accounts, compared to one fifth of the pro-Clinton twitter traffic, (3) the significant rise of Twitter traffic around debate time is mostly from real users who generate original tweets using the more neutral hashtags. In short, Twitter is much more actively pro-Trump than pro-Clinton and more of the pro-Trump twitter traffic is driven by bots, but a significant number of (human) users still use Twitter for relatively neutral political expression in critical moments.

WHAT ARE POLITICAL BOTS?

A growing number of political actors and governments worldwide are employing both people and bots to shape political conversation.[1], [2] Bots can perform legitimate tasks like delivering news and information, or undertake malicious activities like spamming, harassment and hate speech. Whatever their uses, bots are able to rapidly deploy messages, replicate themselves, and pass as human users.

Networks of such bots are called "botnets," a term combining "robot" with "networks" and describing a collection of connected computers with programs that communicate across multiple devices to perform some task. There are legitimate botnets, like the Carna botnet, which gave us our first real census of device networks, and there are malicious botnets, like those that are created to launch spam and distributed denial-of-service (DDoS) attacks and to engineer theft of confidential information, click fraud, cyber-sabotage, and cyberwarfare.[3], [4]

Social bots are particularly prevalent on Twitter, but they are found on many different platforms that increasingly form part of the system of political communication in many countries.[5] They are computer-generated programs that post, tweet, or message of their own accord. Often bot profiles lack basic account information such as screen names or profile pictures. Such accounts have become known as "Twitter eggs" because the default profile picture on the social media site is of an egg. While social media users get access from front-end websites, bots get access to such websites directly through a code-to-code connection, mainly through the site's wide-open application programming interface (API) that enables real-time posting and parsing of information.

Bots are versatile, cheap to produce, and ever evolving. Unscrupulous Internet users now deploy bots beyond mundane commercial tasks like spamming or scraping sites like eBay for bargains. Bots are the primary applications used in carrying out DDOS and virus attacks, email harvesting, and content theft. A subset of social bots are given overtly political tasks and the use of political bots varies from country to country. Political actors and governments worldwide have begun using bots to manipulate public opinion, choke off debate, and muddy political issues. Political bots tend to be developed and deployed in sensitive political moments when public opinion is polarized. How were bots used during the first Presidential debate in the United States?

SAMPLING AND METHOD

This data set contains more than 9.0 million tweets collected September 26-29 2016, using a combination of hashtags associated with the Presidential candidates or the @realDonaldTrump and @HillaryClinton account names. Since our purpose is to discern how bots are being used to amplify political communications, so we did specific analysis of hashtag use in this dataset.

Twitter provides free access to a sample of the public tweets posted on the platform. Twitter's precise sampling method is not known, but according to Twitter, the data available through the Streaming API is at most one percent of the overall global public communication on the platform at any given time.[6] In order to get the most complete and relevant data set, the tweets were collected by following particular hashtags identified by the team as being actively used during the debate. A few additional tags were added, during the debate, as they rose to prominence. The programming of the data collection and most of the analysis were done by using the statistics package R.

Selecting tweets on the basis of hashtags has the advantage of capturing the content most likely to be about this important political event. The streaming API yields (1) tweets which contain the keyword or the hashtag; (2) tweets with a link to a web source,

such as a news article, where the URL or the title of the web source includes the keyword or hashtag; (3) retweets where the text contains the original text, and the keyword or hashtag is used either in the retweet part or in the original tweet; and (4) quote tweet where the original text is not included but Twitter uses a URL to refer to the original tweet.

Our method counted tweets with selected hashtags in a simple manner. Each tweet was coded and counted if it contained one of 52 specific hashtags that were being followed. If the same hashtag was used multiple times in a tweet, this method still counted that tweet only once. If a tweet contained more than one selected hashtag, it was credited to all the relevant hashtag categories.

Unfortunately, not enough users geotag their profiles to allow analysis of the distribution of this support around the world or within the United States. In addition, analyzing sentiment on social media such as Twitter is difficult.[7], [8] Contributions using none of these hashtags are not captured in this data set, and it is possible that users who used one of these hashtags and were not discussing the debate had their tweet captured. Moreover, if people tweeted about the debate, but did not use one of these hashtags or identify a candidate account, their contributions are not analyzed here.

FINDINGS AND ANALYSIS

With this sample we can draw some conclusions about the character and process of political conversation over Twitter during the first debate. Specifically, we can parse out the amount of social media content related to the two major candidates, and we can investigate how much of this content is driven by highly automated accounts. We can parse out the volume of tweets by perspective, assess the level of automation behind the different perspectives, and evaluate the particular contribution of bots to traffic on this issue.

Comparing the Candidates on Twitter. Table 1 reveals that 4.5 million tweets used some combinations of these hashtags. This table reveals that the overall volume of pro-Twitter Trump traffic (39.1 percent), and the overall volume of neutral debate-related traffic (35.8 percent), was much greater than the volume of pro-Clinton traffic (13.6 percent). Much smaller proportions of the tweets were categorized for different mixes of hashtags. The fact that so much of the Twitter content about the debates used exclusive clusters of hashtags from each camp (52.7 percent) is evidence of how polarized and bounded the different communities of social media users are.

Figure 1 displays the rhythm of this traffic over the sample period. Interestingly, Figure 1 also reveals that the significant peak of Twitter content about the debate comes from users who do not tweet exclusively with pro-Clinton and pro-Trump hashtags. Figure 1 includes a total of 9.0m tweets

Table 1: Twitter Activity during the First U.S. Presidential Debate

	All Tweets in Sample	
	N	%
Pro-Trump	1,762,012	39.1
Pro-Clinton	612,732	13.6
Neutral	1,616,300	35.8
Trump-Neutral	197,420	4.4
Clinton-Neutral	99,172	2.2
Trump-Clinton	161,027	3.6
Trump-Clinton-Neutral	63,062	1.4
Total	4,511,725	100.0

Source: Authors' calculations from data sampled 26-29/09/16. Note: Pro-Trump hashtags include #AmericaFirst, #ImWithYou, #MAGA, #TrumpTrain, #MakeAmericaGreatAgain, #TrumpPence16, #TrumpPence2016, #Trump, #AltRight, #NeverHillary, #deplorable, #TeamTrump, #VoteTrump, #CrookedHillary, #LatinosForTrump, #ClintonFoundation, #realDonaldTrump, #LawAndOrder, #pepe, #DebateSideEffects, #WakeUpAmerica, #RNC, #tcot. Pro-Clinton hashtags include #ImWithHer, #LoveTrumpsHate, #NeverTrump, #Clinton, #ClintonKaine16, #ClintonKaine2016, #DNC, #OHHillYes, #StrongerTogether, #VoteDems, #dems, #DirtyDonald, #HillaryClinton, #Factcheck, #TrumpedUpTrickleDown, #ClintonKaine, #WhyIWantHillary, #HillarysArmy, #CountryBeforeParty, #TNTweeters, #UniteBlue, #p2, #ctl, #p2b. Neutral hashtags include #Debates2016, #Debates, #Debate, #Election2016, #POTUS.

Figure 1: Hourly Twitter Traffic, by Candidate Camp
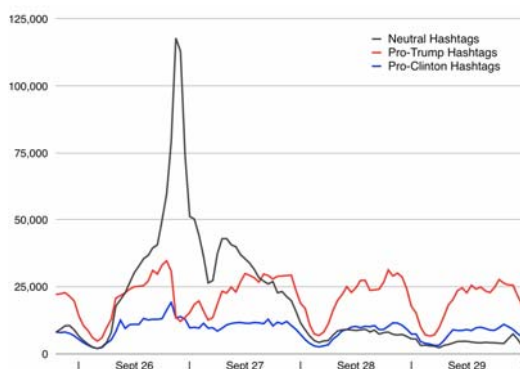
Source: Authors' calculations from data sampled 26-29/09/16. Note: This figure is made with candidate @ mentions of because of the difficulty of understanding the valence of its use.

from 2.0m users who tweeted with either our sampled hashtags or one of the candidate's user names. During the debate itself, the amount of neutrally-tagged traffic outstripped the volume of traffic using candidate-specific hashtags

Automated Political Traffic. A fairly consistent proportion of the traffic on these hashtags was generated by highly automated accounts. These accounts are often bots that see occasional human curation, or they are actively maintained by people who employ scheduling algorithms and other applications for automating social media communication. We define a high level of automation as accounts that post at least 50 times a day, meaning 200 or more tweets on at least one of these hashtags during the data collection period.

Extremely active human users might achieve this pace of social activity, especially if they are simply retweeting the content they find in their social media feed. And some bots may be relatively dormant, waiting to be activated and tweeting only occasionally. But this metric captures accounts generating large traffic with high levels of automation. Finally, self-disclosed bots were identified by searching for the term "bot" in either the tag or account description. While this is a small proportion of the overall accounts, we expect the actual number of bots to be higher—many bots, after all, would not disclose their activities. Future research will involve a more detailed analysis of the disclosed and hidden bots and searching for a wider range of terms referring to bots in the account name and description data.

Table 2 reveals the different levels of automation behind the traffic associated with clusters of hashtags. To track the activity of political bots during the Presidential debates we clustered the hashtags by their candidate associations. To evaluate the role of automation in this debate, we organize clusters of opinion based on hashtag use. Then we created a subcategory of accounts that use heavy levels of automation. Table 2 indicates the level of traffic, by political camp and associated hashtags. This table distinguishes between the messages that exclusively used a hashtag known to be associated with a perspective and then the combinations of mixed tagging that are possible.

Table 2 also reveals that automation is used at several different levels by accounts taking different perspectives in the debate. The accounts using exclusively neutral hashtags are rarely automated (only 11.4 percent use heavy automation) while one-third of all the tweets using a mixture of all hashtags are generated by accounts that use heavy automation. The exclusively neutral hashtags seem to be relatively free of highly automated traffic. Figure 2 reveals the relative flow of traffic overall, and from accounts with high levels of automation. As in many political conversations over Twitter, the most active accounts are either obvious bots or users with such high levels of automation that they are essentially bot-driven accounts.

Additional Observations on Automation. To understand the distribution of content production across these users, we then look at segments of the total population of contributors to these hashtags. There is a noticeable difference between the usage patterns of typical human users and accounts that are bots or highly automated. For example, the top 20 accounts, which were mostly bots and highly automated accounts, averaged over 1,000 tweets a day and they generated 86,000 tweets. The top 100 accounts, which still used high levels of automation, generated around 200,000 tweets at an average rate of 500 tweets per day. In contrast, the average account in the whole sample generated one tweet per day. While

	Low	High	All	
	%	%	N	%
Exclusive Hashtag Clusters				
Pro-Trump	67.3	32.7	1,762,012	100
Pro-Clinton	77.7	22.3	612,732	100
Neutral	88.6	11.4	1,616,300	100
Mixed Hashtag Clusters				
Trump-Neutral	67.0	33.0	197,420	100
Clinton-Neutral	76.2	23.8	99,172	100
Trump-Clinton	70.9	29.1	161,027	100
Trump-Clinton-Neutral	70.3	29.7	63,062	100
Sum	76.7	23.3	4,511,725	100

Table 2: Twitter Content, By Hashtag and Level of Automation

Source: Authors' calculations from data sampled 26-29/09/16.
Note: Low volume users are average human users, high volume accounts post more than 50 times per day on average.

Figure 2: Total Hourly Twitter Traffic on the First Presidential Debate, by Level of Automation
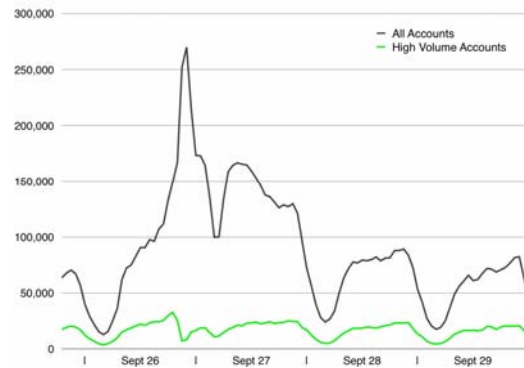
Source: Authors' calculations from data sampled 26-29/09/16.
Note: This data includes @ mentions of particular candidates. We define heavily automated accounts as tweeting 50 times or more per day.

heavily automated accounts are usually the most active, there is a long tail of human users with only occasional Twitter activity.

The accounts using a high level of automation—the accounts that tweeted 200 or more times with a related hashtag and user mention during the data collection period—generated close to 20 percent of all Twitter traffic about the Presidential debate. That volume is significant, considering that this number of posts was generated by slightly more than 4,500 users in a sample of more than 2 million users. In other words, less than half a percent of the accounts generate almost a fifth of all the content. It is difficult for human users to maintain this rapid pace of social media activity without some level of account automation, though certainly not all of these are bot accounts

CONCLUSIONS

Bots operate on many sensitive political topics during close electoral contests in many advanced democracies.[9] However, it is difficult to estimate their impact on popular opinion and to determine if they are political speech worthy of public

3

oversight.[10] Political algorithms have become a powerful means of political communication for "astroturfing" movements—defined as managing the perception of grassroots support.[11] Bots have become a means of managing citizens, by going beyond simply padding follower lists to retweeting volumes of commentary. In previous analysis, we find that bots generate a noticeable portion of all the traffic about the UK referendum, very little of it original.[2] Repeating this sample collection and study method over a longer time period and right around voting day would likely reveal additional features of bot activity on this topic.

It is likely that users tweeting from the Trump camp (a) have generated a larger block of content and (b) are better at tagging their contributions so as to link messages to a larger argument and wider community of support.

We find that political bots have a modest but strategic role in the U.S. Presidential debates. In this brief analysis we find that (1) Twitter traffic on pro-Trump hashtags was more than twice that of the pro-Clinton hashtags, (2) about one third of the pro-Trump twitter traffic was driven by bots and highly automated accounts, compared to one fifth of the pro-Clinton twitter traffic, (3) the significant rise of Twitter traffic around debate time is mostly from real users who generate original tweets using the more neutral hashtags.

ABOUT THE PROJECT
The Project on Computational Propaganda (www.politicalbots.org) involves international, interdisciplinary researchers investigating the impact of automated scripts—computational propaganda—on public life. *Data Memos* are designed to present quick snapshots of analysis on current events in a short format. They reflect methodological experience and considered analysis, but have not been peer-reviewed. *Working Papers* present deeper analysis and extended arguments that have been collegially reviewed and that engage with public issues. The Project's articles, book chapters and books are significant manuscripts that have been through peer review and formally published.

REFERENCES
[1] M. C. Forelle, P. N. Howard, A. Monroy-Hernandez, and S. Savage, "Political Bots and the Manipulation of Public Opinion in Venezuela," Project on Computational Propaganda, Oxford, UK, Working Paper 2015.1, Jul. 2015.

[2] P. N. Howard and B. Kollanyi, "Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum," *arXiv:1606.06356 [physics]*, Jun. 2016.

[3] "Carna botnet," *Wikipedia*. 24-Nov-2015.

[4] "Denial-of-service attack," *Wikipedia*. 15-Oct-2016.

[5] A. Samuel, "How Bots Took Over Twitter," *Harvard Business Review*, 19-Jun-2015. [Online]. Available: https://hbr.org/2015/06/how-bots-took-over-twitter. [Accessed: 23-Jun-2016].

[6] F. Morstatter, J. Pfeffer, H. Liu, and K. M. Carley, "Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose," *arXiv:1306.5204 [physics]*, Jun. 2013.

[7] Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia, "Who is tweeting on Twitter: human, bot, or cyborg?," in *Proceedings of the 26th annual computer security applications conference*, 2010, pp. 21–30.

[8] Cook, David, Waugh, Benjamin, Abdinpanah, Maldini, Hashimi, Omid, and Rahman, Shaquille Abdul, "Twitter Deception and Influence: Issues of Identity, Slacktivism, and Puppetry," *Journal of Information Warfare*, vol. 13, no. 1.

[9] P. N. Howard, *Pax Technica: How the Internet of Things May Set Us Free*. New Haven, CT: Yale University Press, 2015.

[10] D. W. Butrymowicz, "Loophole.com: How the Fec's Failure to Fully Regulate the Internet Undermines Campaign Finance Law," *Columbia Law Review*, pp. 1708–1751, 2009.

[11] P. N. Howard, *New Media Campaigns and the Managed Citizen*. New York, NY: Cambridge University Press, 2006.