NDN

www.ndn.org

A Primer on Social Media Bots And Their Malicious Use In U.S. Politics

By **Tim Chambers**

September 2017

1200 18th Street, NW, 7th Floor • Washington, DC 20036 • Tel: 202.842.7217 • Fax: 202.547.2929

PREFACE

As the Internet careens into its mid-twenties it remains full of extraordinary promise for the people of the world. But we can no longer overlook very real challenges which have emerged. This paper, written by our long time collaborator Tim Chambers, offers a bit of a primer on a particularly pernicious challenge – the weaponization of bots in domestic US politics. Like everything Tim writes it is thoughtful, easy to digest and provocative. I know you will find it useful.

This is likely to be the first in a series of papers from NDN on contemporary challenges facing the Internet. As someone who has been a champion of the "net" since its earliest days, I will admit to being very concerned about what I am seeing. While the Internet has grown in unimaginable ways, and spawned remarkable companies and new types of rewarding human interaction since the launch of the World Wide Web in 1994, today's Internet doesn't feel healthy and secure. Unaddressed, problems like bots could poison and weaken the Internet, leaving the promise of one of humankind's greatest inventions greatly diminished. It should be the goal of policy makers particularly here in the US to prevent that from happening. We need a new spirited commitment to ensure the Internet's emerging problems are addressed, and a sprawling, creative, innovative and above all else *healthy* Internet is left for generations to come to enjoy as we have for these past 20 plus years.

And it is in that spirit that we release this new primer on bots and their malicious use in US politics. Enjoy.

- Simon Rosenberg, NDN President

ABOUT THE AUTHOR: **TIM CHAMBERS**

Tim is the US Digital Practice Lead at the firm the Dewey Square Group, one of the nation's leading public affairs firms.

As head of digital strategy, he provides senior counsel to political campaigns, top non-profits, government agencies, and Fortune 500 corporations — to best harness the power of digital and social media, data analytics and emerging technologies.

Tim has a decade long history of working at the intersection of political, entertainment and technology industries. Prior to joining the firm, he was Senior Vice President of Advanced Media Platforms and at Sony Corp of America. Tim was awarded the 2000 Smithsonian Laureate Award.

Tim can be reached at tchambers@deweydigital.com, and be sure to follow him on Twitter at @tchambers.

A Primer on Social Media Bots and Their Malicious Use in U.S. Politics

"I would like to understand much, much better how big the impact of bot networks is. How much of the fake news distribution that we see is real people consuming fake news, and how much is it bot networks gaming algorithms and social media rankings? I think that we have only started scratching the surface there."

- Katharina Borchert, Mozilla, Chief Innovation Officer

SENATOR KING: "Was the Russian activity in the 2016 election a one-off proposition, or is this part of a long-term strategy? Will they be back?"
FBI DIRECTOR COMEY:
"Oh, it's a long-term practice of theirs. It stepped up a notch in a significant way in '16. They'll be back."

If you were Oliver Mitov's Facebook friend in 2016, you'd see him as a middle-aged man, pro-Bernie Sanders, with about 19 other Facebook friends.

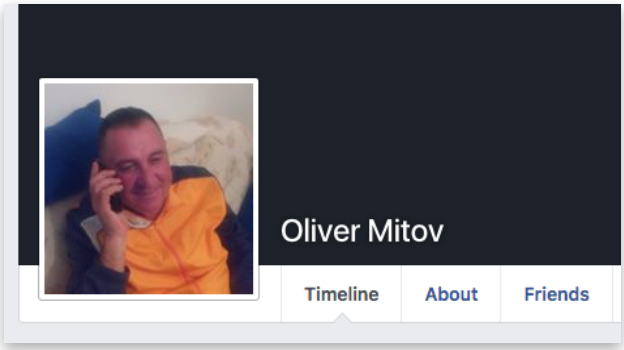

But Oliver Mitov does not seem to have existed – or more accurately, at least four Facebook users with that name existed. Collectively, these accounts were members of more than 36 pro-Bernie Sanders Facebook groups, including Oregon for Bernie Sanders 2016, Pennsylvania Progressives for Bernie Sanders and Latinos for Bernie Sanders. One of the accounts followed all the others.

The Mitov accounts reached hundreds of thousands of Facebook users via hundreds of posts to all 36 of these groups during the heat of the 2016 presidential campaign, after Sanders had dropped out and the Democrats were struggling to unite in the wake of a contentious primary. All of the automated mass posts included identical language and punctuation. All of them disparaged Hillary Clinton at a critical time. And they all linked to a junk news website, later discovered to be hosted in Macedonia.

As for Oliver Mitov, "He may be a bot. He may be one person or four. He may be living in Macedonia, laughing," wrote investigative journalist John Mattes, who first discovered Mitov. The actions of the Mitov accounts – and others like them – alarmed Sanders legitimate supporters and campaign staff as they saw forces from outside the campaign at work. Indeed,

Mitov was just one of many thousands of false Facebook users engaged in similar activity at the time, including bots, cyborgs and trolls. The threat of malicious social media bots by no means ended with the 2016 election, nor has it been confined to Facebook.

If you recently followed Angee Dixson on Twitter, you'd see her as a young woman living in the U.S. Her bio reads, "Christian first … Usually outspoken, Conservative depending on the issue. Don't mess with me, girl!" But like Mitov, Dixson does not exist.

Her Twitter account was created in August 2017 with an image stolen from Instagram user Lorena Rae, a German model. In 17 days, this account posted 891 tweets – an average of 89 a day – supporting Trump and his positions after the Charlottesville protests. The account used IFTTT.com, a service that can automate mass postings to social accounts. One post used language and imagery very similar to Russian social posts from RT and Sputnik.

After these activities were flagged by bot researchers, Twitter suspended the account.

Malicious Social Bots Pose Serious Threats to Democracy

Tens of millions of fake, automated accounts infest all of our social media platforms. Millions more real humans believe the things these bots post and share them with their friends.

There are false right-wing bots and false left-wing bots. They fake petition signatures. They skew poll results and recommendation engines. They sow discord and division. They spread rumors, falsehoods and unconfirmed "junk news." They harass dissenting voices and try to intimidate journalists.

Specifically, deceptive bots create the impression that there is grassroots, positive, sustained, human support for a certain candidate, cause, policy or idea. In doing so, they pose a real danger to the political and social fabric of our country.

"Oliver Mitov" and "Angee Dixson" are out to erode, misdirect and coopt your legitimate political power in our nation's democracy. Along with legions of other malicious bots, they're the 21st century incarnation of propaganda and misinformation—or "computational propaganda."

The threat of social bots is growing as they constantly evolve to become more capable and convincing. We have no time to lose in gaining a deep understanding of this new danger to our information platforms and news systems, and it is critical that we regard the potential harms very seriously

Behind the Bots: What They Are and How They Work

What is a bot?
A bot (derived from "robot") is a software application that performs scripted, often repetitive tasks automatically, or in some cases according to schedules or when it receives specific input.

Not every bot is malicious. Have you talked to Siri or Alexa? They're forms of bots. Or maybe you've used an automatic service to tweet out a notice to parents of the monthly school meeting. That's a bot, too. Over one million humans have used ResistBot to contact their Congressman. Bots do a lot of good, from helping online shoppers find the right sized boots to tutoring people in new skills to optimizing your web searches.

On social platforms such as Twitter, Facebook and YouTube, malicious bots are automated accounts that are programmed to tweet or post. And while there is no legitimate reason for a bot to hide the fact that it is a bot or pretend to be a human, malicious social bots do so routinely, often with remarkably creative and believable bios, fake screen names and stolen or false profile images.

"You may have an account that looks like a housewife in Nebraska, supporting a particular candidate, and then, six months later, all of their tweets are deleted, and then this person becomes a journalist in Ukraine, tweeting about a different topic. ... These attacks are sophisticated, created on a large scale, and it's very hard to see who's behind them."

– Professor Filippo Menczer, Director, Center for Complex Networks and Systems Research, Indiana University

It's important to note that bots don't just act individually; a botnet coordinates many thousands of bots as one group. Behaving like a hive, such networks tweet or post large amounts of content to a widespread audience. Recent social botnets have been discovered with 63,000+ accounts, 350,000 accounts or more.

Malicious bots have been active since the birth of the Internet. They've injected spam into discussion board comments, falsely inflated traffic counts to extract money from advertisers, distorted reviews and rankings of products and services, and shut down websites by overloading them. The most recent threats posed by malicious social bots are just the latest entries in a lengthy rap sheet.

What is a cyborg?
Behind every bot—even the most automated ones—is a human. Many bots employ a combination of human and automated activity. Such hybrid accounts are called cyborgs.

On social media, mixing in some human-created posts often helps hide the automated and scripted nature of the account and makes it appear more authentic. Cyborgs can serve almost identical functions as bots and cause the same massive spread of disinformation.

What is a troll?
Trolls are humans who use social media accounts to deliberately disrupt forums and discussions, spark disagreements or post material designed to be inflammatory, extraneous or part of a larger messaging or propaganda effort.

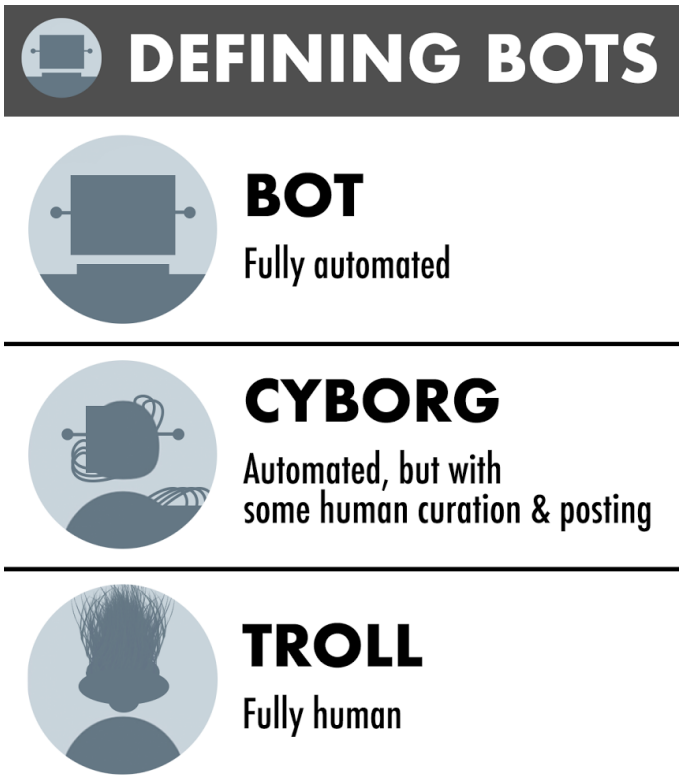

Quantifying the Threats

New studies have shown for the first time quantitatively that bots play a large and critical role in the viral spread of junk news. If fake news spreads like a contagion, bots accelerate the proliferation, just as airline travel can speed the spread of a communicable disease around the world.

Humans unwittingly add credibility to bot content by reposting it. Researchers found that when humans are exposed to bot-generated content on Twitter, "bots were retweeted by approximately 2 [actual] people on average." As University of Washington Professor Kate Starbird put it, when you see a message from multiple friends, "your brain tells you 'Hey, I got this from three different sources,' ... But you don't realize it all traces back to the same place, and might have even reached you via bots posing as real people."

Twitter bots
Twitter has estimated that more than 8.5% of its user base is "highly automated." That's about 23 million bots. Researchers outside the company have placed the estimate much higher, at about 15%, or up to nearly 48 million bots. And those researchers added that their estimates are "conservative."

On Facebook, the level of highly automated accounts is harder to research, because Facebook doesn't share much data with researchers. In the company's regulatory report, it claimed that 8.7% of users are "fake, invalid accounts," which would include bots. That equals 83.09 million false Facebook accounts.

Facebook bots running pages and groups
Recently, Facebook released a study on fake news. It found that while fully automated bots were not as prevalent on Facebook as on other platforms, the investigation "does not

contradict" the assessment of the U.S. Director of National Intelligence that "information warfare" has occurred on Facebook. The report cited "false amplification," such as creation "often en masse" of fake accounts, "coordinated rapid posts" across multiple pages and groups, and "coordinated or repeated comments, some of which may be harassing in nature," in addition to "inflammatory and sometimes racist memes."

Facebook's study called this "wide-scale coordinated human interaction," adding that the company "observed many actions by fake account operators that could only be performed by people with language skills and a basic knowledge of the political situation in the target countries, suggesting a higher level of coordination and forethought."

Of additional concern, after months of suggesting this did not occur, Facebook recently announced that over 470 "inauthentic accounts or pages" that were all "affiliated with one another and likely operated out of Russia" were serving targeted social media ads during June 2015 to May 2017. The company identified at least 3,000 and possibly more than 5,200 separate ads served, costing $100,000 to $150,000 in total. Facebook officials said some were "linked to a troll farm in St. Petersburg, referred to as the Internet Research Agency."

Facebook said the content was focused on "amplifying divisive social and political messages" and was consistent with coordinated activities they found earlier relating to 2016 election disinformation efforts.

Malicious social bots are also common on YouTube, Instagram and other platforms.

Who Are the Puppeteers Behind These Automated Armies?

Researchers from multiple fields and across the ideological spectrum are trying to track down the most important sources of politicized social bot activity. They can be elusive, cloaked behind layers of false identities, but complex data detection work has revealed three major types: state-sponsored, commercial and individual. Of course, these categories can overlap, and at times actors within them may cooperate and interrelate.

State-sponsored governmental and quasi-governmental botnets

The use of bots and botnets for propaganda by governments and government-sponsored organizations is rampant, and Russia is a major player. "Some 45 percent of Twitter activity in Russia is managed by highly automated accounts," said the University of Oxford's Samuel C. Woolley and Philip

"Marrying a hundred years of expertise in influence operations to the new world of social media, Russia may finally have gained the ability it long sought but never fully achieved in the Cold War: to alter the course of events in the U.S. by manipulating public opinion," said Massimo Calabresi in *Time* magazine.

N. Howard in Computational Propaganda Worldwide.

The Kremlin's Internet Research Agency is one of the clearest cases of state-sponsored bot and cyber-activity. According to *The New York Times'* Adrian Chen, "Russia's information war might be thought of as the biggest trolling operation in history, and its target is nothing less than the utility of the Internet as a democratic space." Similarly, *Wired Magazine* reported that Russian bots and trolls represent just the tip of a much larger coordinated messaging effort: "The bots and trolls are important, but they're maybe 10 percent of the tools that are used. They're used for amplification and they're used for bullying."

"A Russian site called buyaccs.com ("Buy Bulk Accounts at Best Prices") offers for sale a huge array of pre-existing social media accounts, including on Facebook and Twitter; like wine, the older accounts cost more, because their history makes chicanery harder to spot."

– "The Fake Americans Russia Created to Influence the Election," Scott Shane, *The New York Times*, September 7, 2017

The targets of all this activity are global. NATO strategic communications issued a study on Russian-language bot activity, which found that between March and August 2017 "about two in three users who write in Russian about the NATO presence in Eastern Europe are robotic, or 'bot' accounts." The study noted that "the English speaking space is also heavily affected: 1 in 4 active accounts were likely automated and were responsible for 46% of all English-language content."

But Russia's not the only culprit, as the Computational Propaganda Project reported:

> "We find that cybertroops are a pervasive and global phenomenon. Many different countries employ significant numbers of people and resources to manage and manipulate public opinion online, sometimes targeting domestic audiences and sometimes targeting foreign publics. … There are details on such organizations in 28 countries. … In addition to official government accounts, many cybertroop teams run fake accounts to mask their identity. … In many cases, these fake accounts are 'bots.'"

And as The Washington Post recently reported, "In some cases, these efforts involved full-blown government bureaucracies, with a steady number of employees and fixed payrolls. Other times, bands of online activists or ad hoc groups of paid workers worked together for a single campaign before being disbanded. Some efforts also get outsourced to private vendors that specialize in influencing opinion through social media."

Marketers, spammers and commercial botnets for hire
Entire shadow industries focus on spam and mass manipulation of social media via bots to achieve purely commercial goals, including click fraud, artificially driving views, and driving sales and marketing of online products and services. Twitter bots promise new followers, likes and shares – for a price – on this "dark net."

These commercial botnets are organized by various shades of cybercriminals. And it is worth noting that in 2011 Russia was one of the larger centers of cybercrime, accounting for 36% of all cybercrime profits that year.

Some of these commercial botnets make use of inflammatory political content on thousands of social accounts – not to sway public opinion, but to attract eyeballs for profit. Once the content lures viewers, the bots point users to their own websites, where each click pays off in ad revenue, malware downloads or other scams. Of course, it's not hard to imagine that these negative user experiences might dampen future political participation via social networks.

Botnets may also be commercialized by renting themselves to political campaigns. For instance, the University of Budapest Corvinus found during the run-up to the Brexit referendum that some botnets appeared to have been repurposed from previous campaigns. Bots that had been tweeting for years about Middle Eastern issues suddenly switched to pro-Brexit topics; after the Brexit vote, they returned to their previous focus.

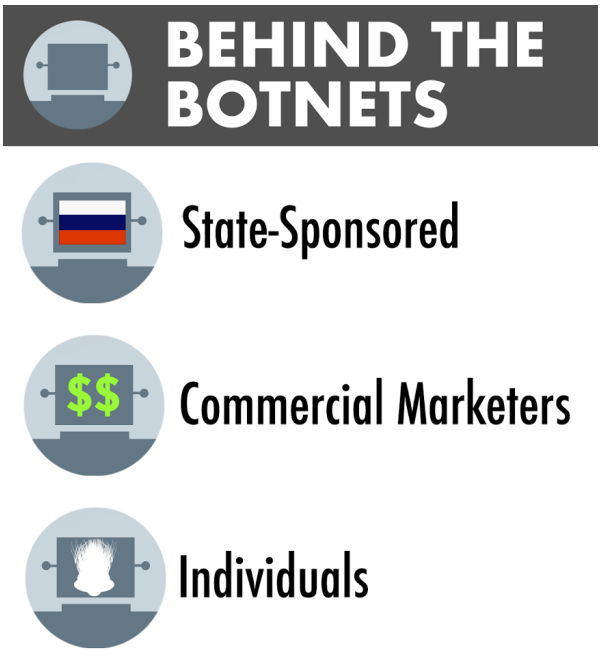

Individual citizens
Some people make their own bots and botnets to amplify their views. An average user can create a basic social bot, using no specialized programming and a few simple tools, in about 20 minutes.

An intermediate programmer can create a small army of bots in that time. Oxford's Woolley estimated that a single motivated individual could easily create and maintain a network of about 400 active Twitter bots or 10 to 20 active, highly automated mass-posting accounts on Facebook.

Emerging artificial intelligence tools and natural language processing engines will make bot creation even easier and faster.

Bot Impacts on Recent Political History

Bots have played a vigorous role in recent political events, not only in the U.S. but worldwide. The following examples are by no means comprehensive, merely illustrative.

2016: Brexit

Throughout the debate leading up to the United Kingdom's vote to withdraw from the European Union, researchers saw strong evidence of bot propaganda and concluded it had a "meaningful" effect on the conversation. Computational Propaganda Project researchers studied 314,000 accounts that tweeted on both sides of the issue. In the critical week of June 5-12, 15% of them were heavily or entirely automated. Of the most active social accounts on the Brexit debate, researchers concluded, "It is almost certain that 7 of the 10 accounts are bots."

2016: U.S. election

"We find that political bot activity reached an all-time high for the 2016 campaign," the Computational Propaganda Project wrote, reporting that of 18.9 million political hashtag tweets studied, 17.9 percent came from "highly automated" accounts that post 50 or more tweets per day. By the final debate, those tweets peaked at more than 27 percent. Moreover:

> "Not only did the pace of highly automated pro-Trump activity increase over time, but the gap between highly automated pro-Trump and pro-Clinton activity widened from 4:1 during the first debate to 5:1 by election day. The use of automated accounts was deliberate and strategic throughout the election, most clearly with pro-Trump campaigners and programmers who carefully adjusted the timing of content production during the debates, strategically colonized pro-Clinton hashtags, and then disabled activities after Election Day."

Similarly, University of Southern California expert Emilio Ferrara and his research team found that more than 400,000 social bots posted roughly 3.8 million tweets on the U.S. election between September 16th and October 21st. During this time, about one-fifth of the entire presidential election conversation on Twitter was automated.

The 2016 presidential race was so close that many factors played a role in the final result, but the tone, pace and quantity of misinformation on social networks was clearly one of them, no doubt swaying some voters, encouraging others to sit out the election and helping to split Democrats between Sanders and Clinton. One researcher looked at vote results compared to bot activity and noted that while other factors certainly applied, "It is remarkable that the states most actively targeted by misinformation-spreading bots tended to have more surprising election results."

2017: French election

Many of the Trump-supporting bots were called into duty again as the French election hit its crucial final weeks. Negative stories and documents – which would later prove false – about

candidate Emmanuel Macron appeared and were spread online. They were known as #MacronLeaks, evoking the hashtag #Hillaryleaks from the 2016 U.S. presidential campaign.

Ferrara's USC team studied 17 million tweets from this time, finding that "out of 99,378 users involved in MacronLeaks, our model classified 18,324 of them as social bots ... about 18% of the total users involved in the campaign, which is extremely consistent with results from previous studies." Much of the action came from Twitter bots that had previously posted content in support of Trump.

"Our work is the first to identify the presence of bots that existed during the 2016 U.S. Presidential election campaign period to support alt-right narratives, went dark after November 8, 2016, and came back into use in the run up days to the 2017 French presidential election," Ferrara wrote.

The Daily Beast did its own analysis around the French election and found similar results: "We found that one-third of the users tweeting about MacronLeaks last week were doing the same [posting about the DNC during the U.S. election] ... in March, with a few tweaking their profiles in between – for example, adding 'Le Pen' to their account description."

Today: A bot-driven presidency

Several forms of analysis, including TwitterAudit, count nearly half of all Trump's followers as fake users, including bots, cyborgs and trolls. Data scientist and researcher C.E. Carey, of Data for Democracy, found a similar situation among Trump's Facebook following:

> "Of the over 25,000 [Facebook-based] bot accounts active before the election, 44% continued posting after November 9, 2016…. Post-election, April 7, 2017 – the date of the Syria strike – was the most active day for bots, with just under 2,500 bot accounts posting over 6,500 messages. On that day alone, bots made up 12% of active users and nearly one in four comments."

The president himself has retweeted content from bots and cyborgs to his millions of followers. And bot-driven activity has swelled to push stories around key events of the Trump presidency, such as investigations into connections between Trump and Russia.

Further, malicious bots and fake accounts continue in other divisive arenas. Pro-Russian false Facebook accounts created two of the most popular Facebook pages that promoted "Texit" – one, called "The Heart of Texas, with over 225,000 followers as well as another, "SecuredBorders," with 133,000 followers. Both have been since shuttered by Facebook.

After the deadly Charlottesville protests, independent news organization ProPublica reported:

> "The Russian influence networks we track are definitely amplifying the broader alt-right chatter about Charlottesville. ... The major themes they have been pushing are the 'both sides are violent' argument and conspiracy theories that George Soros was behind the counter-protests, although the latter has been trending more sporadically."

After issuing this report, ProPublica itself was a victim of various cyberattacks and bot attacks.

Most recently, Russian state media messaging, individual alt-right social accounts and robotic social amplification played a key role in shaping the discussion of the Berkeley protests and Antifa.

Potential Political Impacts of Bots in the Near Future

NATO's report warned that "the democratizing possibilities of social media ... appear to have been greatly undermined." In the coming years, the dangers bots pose to democracy will likely grow further, pushed by these three trends:

1) The rise of artificial intelligence tools
Many of today's malicious social bots are unsophisticated, and since bot detection and counteraction efforts have so far been largely ineffectual, creators have been lulled into making bots that are easy to spot, often rife with errors that give them away (like Angee-bot's use of a photo of a somewhat public figure).

But researchers predict that advances in artificial intelligence will make conversation bots and automated postings "effectively indistinguishable" from human activities. Each successive generation of malicious bots will be created to better evade detection. Their makers will capitalize on improvements in bots designed for beneficial uses, adapting them to produce ever more convincing forgeries of human social media accounts.

2) Evolution across new platforms
Bots will increasingly make themselves at home in huge numbers across Twitter, Facebook and YouTube, and as the social media landscape continues to grow, bots will expand into new online territories. There are already signs of bot, cyborg and troll activity on Facebook Messenger, LinkedIn, WhatsApp and elsewhere.

3) As helpful bots become more common, malicious ones will hide among them
Increasingly, helpful bots will be developed and put to use on social media platforms. Gradually – and legitimately – they will be integrated into our social accounts. Sorting out malicious actors in the mix will become all the more challenging.

Solutions: Awareness, Action, Advocacy

Everyone – political practitioners of any stripe, media professionals, software engineers, citizens – must make ourselves more aware and more sophisticated about these issues. We need to develop more and better tools to combat malicious bots. The Hamilton 68 Dashboard, from the Alliance for Securing Democracy, is a good example of the many new utilities and tools that are needed. These new tools and technologies need to be created and to evolve as quickly as the threats they combat do.

We must demand that our social networks build for the good of the countries they act in, not

just for their own profits. And we need to take other personal, business, political and legal actions as well. These include:

What individuals should do

First, we should all be more aware and more skeptical on social media. As April Glaser wrote in *Slate* when looking at the effect of bots: "…retweets, likes, trending hashtags, and followers shouldn't be taken as a strong indication of public opinion—and moreover…virality is hardly a demonstration of genuineness."

|

Second, all social media users should learn how to identify likely bots -- and starve them out:

- Don't accept requests from strangers on Facebook. Researchers have found that "over 20% of legitimate users accept friendship requests indiscriminately and over 60% accept requests from accounts with at least one contact in common."

- Don't like or share content that you don't verify or for which you don't know the source.

- Don't retweet or repost content from accounts you don't trust, simply because you agree with their point.

- When you suspect a malicious bot, block it from your account. Report malicious Twitter bots to Twitter this way, Report malicious Facebook bots to Facebook here. Suspicious groups run by likely malicious bots can be reported here.

What social media platforms should do

As the customers of our social networks, we can all call on them to take actions on our behalf, such as these:

- In their terms of service, Twitter, Facebook and other platforms should require all bot accounts to publicly state that they are automated in their profiles and when asked.

- Twitter, Facebook and other platforms should provide users with a one-click button for reporting suspected malicious bots and blocking them from our feeds.

- Twitter should open up a version of its "verified user" status to the general public. Such status is now available only to celebrities and thought leaders. Twitter should also allow users to choose to ignore all or some comments in their feed from non-verified users and identified bots.

- Twitter should strengthen and clarify its terms of service to firmly prohibit multiple accounts, and Facebook should better police its multiple account prohibition.

- Posts, tweets, shares or comments that are posted automatically via scripts should be required to include a label to the effect that they were "posted by a robot."

What national legislation and policy should do

So far, federal legislation on bot activity has been slight, but in 2016 President Obama signed into law "The BOTS Act," prohibiting the use of bots in online ticket scalping scams. And pressure from EU potential legislation led to Google and Facebook cracking down on bots and junk news propagation on their platforms.

These are good first steps, but more must be done. For instance, legislators should:

- Require social media networks' terms of service to make all bots self-identify, both on their social profiles and when asked whether they are a bot.
- Require all social media platforms to add a tag to all automated posts labeling them as such.
- Require all political office holders (local, state and Federal) using social media accounts to continually block and unfollow any likely bots or cyborgs from their followers.
- Require social media and ad platforms to fully disclose political ad information, including data on the accounts, spend and account information.

There's No Time to Waste

As we approach the 2018 and 2020 election seasons, negative trends are growing fast. Law professor and author Tim Wu warned:

> "The problem is almost certain to get worse, spreading to even more areas of life as bots are trained to become better at mimicking humans. ... In coming years, campaign finance limits will be (and maybe already are) evaded by robot armies posing as 'small' donors. And actual voting is another obvious target — perhaps the ultimate target."

Professor Kate Starbird of the University of Washington said, "There really is an information war for your mind. And we're losing it." She fears we're moving toward "the menace of unreality—which is that nobody believes anything anymore."

But this outcome is not inevitable. "Social media has made democracy weak. It has a compromised immune system," as professor Philip Howard observed. However, we still have a chance to turn things around. "The next big step for us," he said, "is whether we can shape the internet to come…"

Our challenges now – the challenges to all democracies at the onset of every new media technology – are to counteract the threats of disinformation and propaganda in our new communication platforms, to protect the promise these new media bring, to realize their potential to open up the democratic process and to hold new media platforms accountable.

We need to be wise and act quickly to treat computational propaganda as the serious threat that it is – a "21st-century social media information war" and a "worldwide, internet-based assault".

SOURCES

Ken Stone, “Russia Duped Bernie Fans via Facebook, San Diego Dems told,” *Times of San Diego*, March 23, 2017. https://timesofsandiego.com/politics/2017/03/23/russia-duped-bernie-fans-via-facebook-san-diego-dems-told/

Sarah K. Burris, “Major Bernie Sanders backers suspect Russians flooded their social site with anti-Hillary memes, Raw Story, March 13, 2017. http://www.rawstory.com/2017/03/sanders-staffers-confirm-that-russian-trolls-were-pushing-anti-hillary-memes-on-social-networks-and-comment-threads/

The Atlantic Council’s Digital Forensic Research Lab, “Kremlin and Alt Right Share Nazi Narrative.” August 18, 2017. https://medium.com/dfrlab/kremlin-and-alt-right-share-nazi-narrative-2df4af60c749

Klint Finley, “Pro-government Twitter Bots Try to Hush Mexican Activists,” *Wired*, August 23, 2017. https://www.wired.com/2015/08/pro-government-twitter-bots-try-hush-mexican-activists/

BBC World Service, Trending, September 10, 2017. http://www.bbc.co.uk/programmes/w3csvfj6

Jacob Ratkiewicz, Michael Conover, Mark Meiss, Bruno Goncalves, Snehal Patil, Alessandro Flammini, and Filippo Menczer, “Truthy: mapping the spread of astroturf in microblog streams.” WWW ’11 Proceedings of the 20th international conference companion on World Wide Web, pg. 249-251, March 28, 2011. http://dl.acm.org/citation.cfm?doid=1963192.1963301

Ben Nimmo, The Atlantic Council’s Digital Forensic Research Lab. https://mobile.twitter.com/benimmo/status/901197575377870848?refsrc=email&s=11

The Atlantic Council’s Digital Forensic Research Lab, “The Many Facets of a Botnet.” May 25, 2017. https://medium.com/dfrlab/the-many-faces-of-a-botnet-c1a66658684

Conspirador Norteño. https://mobile.twitter.com/conspirator0/status/900159044555505665

Juan Echeverria and Shi Zhou, “Discovery, Retrieval, and Analysis of the ‘Star Wars’ Botnet in Twitter,” University College London Department of Computer Science, June 13, 2017. https://arxiv.org/pdf/1701.02405.pdf

Alessandro Flammini, Chengcheng Shao, Filippo Menczer, Giovanni Luca Ciampaglia, Onur Varol, “The spread of fake news by social bots.” Indiana University, Bloomington, July 24, 2017. https://arxiv.org/abs/1707.07592

Douglas R. Guilbeault and Samuel C. Woolley, “Computational Propaganda in the United States of America: Manufacturing Consensus Online.” University of Oxford. http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/06/Comprop-USA.pdf

Danny Westneat, “UW professor: The information war is real, and we’re losing it.” The Seattle Times, March 29, 2017. http://www.seattletimes.com/seattle-news/politics/uw-professor-the-information-war-is-real-and-were-losing-it/

Zoey Chong, “Up to 48 million Twitter accounts are bots, study says.” CNET, March 14, 2017. https://www.cnet.com/news/new-study-says-almost-15-percent-of-twitter-accounts-are-bots/

Onur Varol, Emilio Ferrara, Clayton A. Davis, Filippo Menczer, Alessandro Flammini, “Online Human-Bot Interactions: Detection, Estimation, and Characterization.” Indiana University, Bloomington. https://arxiv.org/pdf/1703.03107.pdf

Heather Kelly, “83 million Facebook accounts are fakes and dupes.” CNN, August 3, 2012. http://www.cnn.com/2012/08/02/tech/social-media/facebook-fake-accounts/index.html

Jen Weedon, William Nuland and Alex Stamos, “Information Operations and Facebook.” Facebook Security, April 27, 2017. https://fbnewsroomus.files.wordpress.com/2017/04/facebook-and-information-operations-v1.pdf

Carol D. Leonnig, Tom Hamburger and Rosalind S. Helderman, “Russian firm tied to pro-Kremlin propaganda advertised on Facebook during election.” The Washington Post, September 6, 2017. https://www.washingtonpost.com/politics/facebook-says-it-sold-political-ads-to-russian-company-during-2016-election/2017/09/06/32f01fd2-931e-11e7-89fa-bb822a46da5b_story.html?nid&utm_term=.37048c6f4caf

Alex Stamos, “An Update on Information Operations on Facebook.” Facebook Newsroom, September 6, 2017. https://newsroom.fb.com/news/2017/09/information-operations-update/

Craig Timberg, “Spreading fake news becomes standard practice for governments across the world.” The Washington Post, July 17, 2017. https://www.washingtonpost.com/news/the-switch/wp/2017/07/17/spreading-fake-news-becomes-standard-practice-for-governments-across-the-world/?utm_term=.2acda13704c9

Clint Watts. “Clint Watts’ Testimony: Russia’s Info War on the U.S. Started in 2014.” The Daily Beast, March 30, 2017. http://www.thedailybeast.com/clint-watts-testimony-russias-info-war-on-the-us-started-in-2014

Adrian Chen, “The Agency.” The New York Times Magazine, June 2, 2015. https://www.nytimes.com/2015/06/07/magazine/the-agency.html?mcubz=0

Massimo Calabresi, “Inside Russia’s Social Media War on America.” Time, May 18, 2017. http://time.com/4783932/inside-russia-social-media-war-america/

Garrett M. Graff, “A Guide to Russia’s High Tech Tool Box for Subverting US Democracy.” Wired, August 13, 2017. https://www.wired.com/story/a-guide-to-russias-high-tech-tool-box-for-subverting-us-democracy

Dr. Rolf Fredhelm, “Robotrolling.” NATO Strategic Communications Centre of Excellence. http://www.stratcomcoe.org/robotrolling-20171

Samuel C. Woolley and Philip N. Howard, “Computational Propaganda Worldwide: Executive Summary.” Samuel Woolley and Philip N. Howard, Eds. Working Paper 2017.11. Oxford, UK: Project on Computational Propaganda. comprop.oii.ox.ac.uk. http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/06/Casestudies-ExecutiveSummary.pdf

Donara Barojan, “Influence for Sale: Bot Shopping on the Darknet. How vendors are selling social media influence and reach for profit on the Darknet.” Medium, June 19, 2017. https://medium.com/dfrlab/influence-for-sale-bot-shopping-on-the-darknet-1c9ddfb3d8e6

Loek Essers, “Russian cybercriminals earned $4.5 billion in 2011.” Computerworld/IDG News Service, April 24, 2012. https://www.computerworld.com/article/2503653/cybercrime-hacking/russian-cybercriminals-earned--4-5-billion-in-2011.html

Berit Anderson and Brett Horvath, “The Rise of the Weaponized AI Propaganda Machine.” SCOUT, February 9, 2017. https://scout.ai/story/the-rise-of-the-weaponized-ai-propaganda-machine

Emilio Ferrara, “How Twitter bots affected the US presidential campaign.” The Conversation, November 8, 2016. http://theconversation.com/how-twitter-bots-affected-the-us-presidential-campaign-68406

Emilio Ferrara, “Disinformation and Social Bot Operations in the Run Up to the 2017 French Presidential Election,” *First Monday* 22(8), August 7, 2017. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2995809

Kevin Poulsen, “The Twitter Bots Who Tried to Steal France.” The Daily Beast, May 14, 2017. http://www.thedailybeast.com/the-twitter-bots-who-tried-to-steal-france

TwitterAudit, @realDonaldTrump TwitterAudit Report. https://www.twitteraudit.com/realdonaldtrump

C.E. Carey, “Continued Bot Infiltration of Trump’s Facebook Pages.” Medium, May 1, 2017. https://medium.com/data-for-democracy/continued-bot-infiltration-of-trumps-facebook-pages-2df82ca86b5b

Natasha Bertrand, “Trump retweeted a Twitter bot – then it got suspended.” Business Insider, August 7, 2017. http://www.businessinsider.com/trump-twitter-bot-nicole-protrump45-2017-8

Casey Michel, “How Russia Created the Most Popular Texas Secession Page on Facebook,” Extra Newsfeed, September 11, 2017. https://extranewsfeed.com/how-russia-created-the-most-popular-texas-secession-page-on-facebook-fd4dfd05ee5c

Ben Collins, Kevin Poulsen, Spencer Ackerman, “Exclusive: Russia Used Facebook Events to Organize Anti-Immigrant Rallies on U.S. Soil,” The Daily Beast, September 11, 2017. http://

www.thedailybeast.com/exclusive-russia-used-facebook-events-to-organize-anti-immigrant-rallies-on-us-soil

@mattmfm. “Why the attack on @ProPublica? This: Pro-Russian Bots Take Up the Right-Wing Cause After Charlottesville.” Twitter, August 24, 2017. https://twitter.com/mattmfm/status/900887102346416128

Caroline O, “How Russian & Alt-Right Twitter Accounts Worked Together to Skew the Narrative About Berkeley.” Arc Digital, September 1, 2017. https://arcdigital.media/how-russian-alt-right-twitter-accounts-worked-together-to-skew-the-narrative-about-berkeley-f03a3d04ac5d

Rob Price, “Researchers taught AI to write totally believable fake reviews, and the implications are terrifying.” Business Insider, August 29, 2017. http://www.businessinsider.com/researchers-teach-ai-neural-network-write-fake-reviews-fake-news-2017-8

Natasha Bertrand, “It looks like Russia hired internet trolls to pose as pro-Trump Americans.” Business Insider, July 27, 2016. http://www.businessinsider.com/russia-internet-trolls-and-donald-trump-2016-7

Nic Dias, “The Era of WhatsApp Propaganda Is Upon Us.” Foreign Policy, August 17, 2017. http://foreignpolicy.com/2017/08/17/the-era-of-whatsapp-propaganda-is-upon-us/

Jacob Shamsian, “There’s a bot on Tinder trying to influence votes in the British election.” Business Insider, June 8, 2017. http://www.businessinsider.com/united-kingdom-election-jeremy-corbyn-tinder-bot-labour-2017-6

Tim Wu, “Please Prove You’re Not a Robot.” New York Times, July 15, 2017. https://www.nytimes.com/2017/07/15/opinion/sunday/please-prove-youre-not-a-robot.html?mcubz=0

April Glaser, “Russian Bots Are Trying to Sow Discord on Twitter After Charlottesville.” Slate, August 24, 2017. http://www.slate.com/blogs/future_tense/2017/08/24/russian_bots_are_sharing_extreme_right_wing_information_on_twitter_after.html

Atlantic Council Digital Forensic Lab, “#BotSpot: Twelve Ways to Spot a Bot.” DFRLab, August 28, 2017. https://medium.com/dfrlab/botspot-twelve-ways-to-spot-a-bot-aedc7d9c110c

Yazan Boshmaf, Ildar Muslukhov, Konstantin Beznosov and Matei Ripeanu, “Design and analysis of a social botnet.” *Computer Networks 57*, 2 (2013), 556–578.

Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, Alessandro Flammini, “The Rise of Social Bots,” *Communications of the ACM*, Vol. 59, No. 7, pages 96-104.

Billy Steele, “President Obama signs nationwide ticket-bot ban into law.” Engadget, December 15, 2016. https://www.engadget.com/2016/12/15/president-obama-signs-ticket-bot-bill-into-law/

Reuters, “Facebook, Twitter and Google Crack Down on Hate Speech Following EU Pressure.” Newsweek, June 1, 2017. http://www.newsweek.com/facebook-twitter-and-google-crack-down-hate-speech-following-eu-pressure-618694

David Ingram, “Facebook to Keep Wraps on Political Ads Data Despite Researchers’ Demands,” Reuters, June 22, 2017. https://www.reuters.com/article/us-usa-politics-facebook/facebook-to-keep-wraps-on-political-ads-data-despite-researchers-demands-idUSKBN19D1CN

Siva Vaidhyanathan, “Facebook Wins, Democracy Loses,” *The New York Times*, September 8, 2017. https://www.nytimes.com/2017/09/08/opinion/facebook-wins-democracy-loses.html?smid=tw-nytopinion&smtyp=cur