

How Twitter bots affected the US presidential campaign

November 9 2016, by Emilio Ferrara



Credit: AI-generated image ([disclaimer](#))

Key to democracy is public engagement – when people discuss the issues of the day with each other openly, honestly and without outside influence. But what happens when large numbers of participants in that conversation are biased robots created by unseen groups with unknown agendas? As my research has found, that's what has happened this

election season.

Since 2012, I have been studying how people discuss [social](#), [political](#), [ideological](#) and [policy](#) issues online. In particular, I have looked at how social media are [abused for manipulative purposes](#).

It turns out that much of the political content Americans see on social media every day is not produced by human users. Rather, about one in every five election-related [tweets](#) from Sept. 16 to Oct. 21 was generated by computer software programs called "[social bots](#)."

These artificial intelligence systems can be [rather simple](#) or [very sophisticated](#), but they share a common trait: They are set to automatically produce content following a specific political agenda determined by their controllers, who are nearly impossible to identify. These [bots](#) have [affected the online discussion](#) around the presidential election, including leading topics and how online activity was perceived by the media and the public.

How active are they?

The operators of these systems could be political parties, foreign governments, third-party organizations, or even individuals with vested interests in a particular election outcome. Their work amounts to at least four million election-related tweets during the period we studied, posted by more than 400,000 social bots.

That's at least 15 percent of all the users discussing election-related issues. It's more than twice the overall concentration of bots on Twitter – which the [company estimates](#) at 5 to 8.5 percent of all accounts.

To determine which accounts are bots and which are humans, we use [Bot Or Not](#), a publicly available bot-detection service that I developed in

collaboration with colleagues at Indiana University. Bot Or Not uses [advanced machine learning algorithms](#) to analyze multiple cues, including Twitter profile metadata, the content and topics posted by the account under inspection, the structure of its social network, the timeline of activity and much more. After considering more than 1,000 factors, Bot Or Not generates a likelihood score that the account under scrutiny is a bot. Our tool is 95 percent accurate at this determination.

There are many examples of [bot-generated tweets](#), supporting their candidates, or attacking the opponents. Here is just one:

@u_edilberto: RT @WeNeedHillary: Polls Are All Over the Place. Keep Calm & Hillary On! t.co/XwBFfLjz7x #p2 #ctl #ImWithHer #TNTweeters [https://t ...](https://t.co/XwBFfLjz7x)

How effective are they?

The effectiveness of social bots depends on the reactions of actual people. We learned, distressingly, that people were not able to ignore, or develop a sort of immunity toward, the bots' presence and activity. Instead, we found that most human users can't tell whether a tweet is posted by another real user or by a bot. We know this because bots are being retweeted at the same rate as humans. Retweeting bots' content without first verifying its accuracy can have real consequences, including spreading [rumors](#), [conspiracy theories](#) or [misinformation](#).

Some of these bots are very simple, and just retweet content produced by human supporters. Other bots, however, produce new tweets, jumping in the conversation by using existing popular hashtags (for instance, [#NeverHillary](#) or [#NeverTrump](#)). Real users who follow these Twitter hashtags will be exposed to bot-generated content seamlessly blended with the tweets produced by other actual people.

Bots produce content automatically, and therefore at a very fast and continuous rate. That means they form consistent and pervasive parts of the online discussion throughout the campaign. As a result, they were able to build significant influence, collecting [large numbers](#) of followers and having their tweets retweeted by thousands of humans.

A deeper understanding of bots

Our investigation into these politically active social bots also uncovered information that can lead us to more nuanced understanding of them. One such lesson was that bots are biased, by design. For example, Trump-supporting bots systematically produced overwhelmingly positive tweets in support of their candidate. [Previous studies](#) showed that this systematic bias alters public perception. Specifically, it creates the false impression that there is [grassroots, positive, sustained support](#) for a certain candidate.

Location provided another lesson. Twitter provides metadata about the physical location of the device used to post a certain tweet. By aggregating and analyzing their digital footprints, we discovered that bots are not uniformly distributed across the United States: They are significantly overrepresented in some states, in particular southern states like Georgia and Mississippi. This suggests that some bot operations may be based in those states.

Also, we discovered that bots can operate in multiple ways: For example, when they are not engaged in producing content supporting their respective candidates, bots can target their opponents. We discovered that bots pollute certain hashtags, like #NeverHillary or #NeverTrump, where they smear the opposing candidate.

These strategies leverage known human biases: for example, the fact that negative content travels faster on [social media](#), as one of [our recent](#)

[studies](#) demonstrated. We found that, in general, negative tweets are retweeted at a pace 2.5 times higher than positive ones. This, in conjunction with the fact that people are naturally more inclined to retweet content that [aligns with their preexisting political views](#), results in the spreading of content that is often [defamatory or based on unsupported, or even false, claims](#).

It is hard to quantify the effects of bots on the actual election outcome, but it's plausible to think that they could affect voter turnout in some places. For example, some people may think there is so much local support for their candidate (or the opponent) that they don't need to vote – even if what they're seeing is actually artificial support provided by bots.

Our study hit the limits of what can be done today by using computational methods to fight the issue of bots: Our ability to identify the bot masters is bound by technical constraints on recognizing patterns in their behavior. Social media is acquiring increasing importance in shaping [political beliefs](#) and [influencing](#) people's [online](#) and [offline behavior](#). The research community will need to continue to explore, to make these platforms as safe from abuse as possible.

This article was originally published on [The Conversation](#). Read the [original article](#).

Provided by The Conversation

Citation: How Twitter bots affected the US presidential campaign (2016, November 9) retrieved 10 April 2024 from <https://phys.org/news/2016-11-twitter-bots-affected-presidential-campaign.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private

study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.