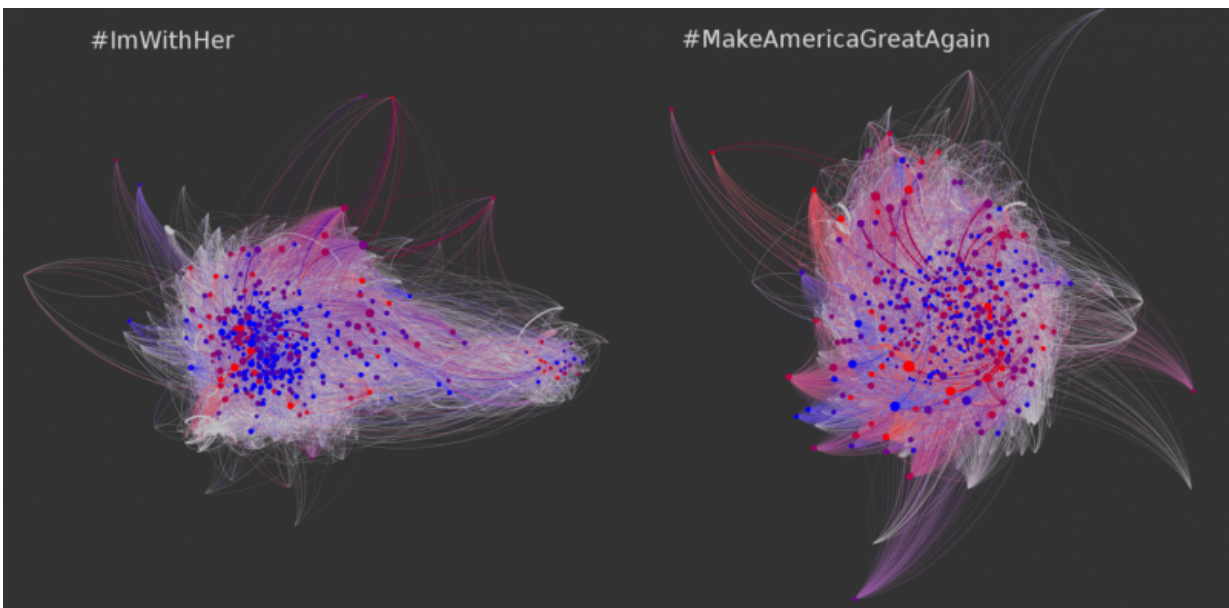


Misinformation on social media—can technology save us?

November 28 2016, by Filippo Menczer



Sharing election hashtags: Dots are Twitter accounts; lines show retweeting; larger dots are retweeted more. Red dots are likely bots; blue ones are likely humans. Credit: Clayton Davis, CC BY-ND

If you get your news from social media, as [most Americans do](#), you are exposed to a daily dose of hoaxes, rumors, conspiracy theories and misleading news. When it's all mixed in with reliable information from honest sources, the truth can be very hard to discern.

In fact, my research team's analysis of data from Columbia University's [Emergent](#) rumor tracker suggests that this [misinformation is just as likely to go viral](#) as reliable information.

Many are asking whether this onslaught of digital misinformation [affected the outcome](#) of the 2016 U.S. election. The truth is we do not know, although there are reasons to believe it is entirely possible, based on [past analysis](#) and [accounts from other countries](#). Each piece of misinformation contributes to the shaping of our opinions. Overall, the harm can be very real: If people can be conned into jeopardizing our children's lives, [as they do when they opt out of immunizations](#), why not our democracy?

As a researcher on the spread of misinformation through social media, I know that limiting news fakers' ability to sell ads, as recently announced by [Google](#) and [Facebook](#), is a step in the right direction. But it will not curb abuses driven by political motives.

Exploiting social media

About 10 years ago, my colleagues and I ran an [experiment](#) in which we learned 72 percent of college students trusted links that appeared to originate from friends – even to the point of entering personal login information on phishing sites. This widespread vulnerability suggested another form of malicious manipulation: People might also believe misinformation they receive when clicking on a link from a social contact.

To explore that idea, I created a [fake web page](#) with random, computer-generated gossip news – things like "Celebrity X caught in bed with Celebrity Y!" Visitors to the site who searched for a name would trigger the script to automatically fabricate a story about the person. I included on the site a disclaimer, saying the site contained meaningless text and

made-up "facts." I also placed ads on the page. At the end of the month, I got a check in the mail with earnings from the ads. That was my proof: Fake news could make money by polluting the internet with falsehoods.

Sadly, I was not the only one with this idea. Ten years later, we have an [industry of fake news](#) and [digital misinformation](#). Clickbait sites manufacture hoaxes to make money from ads, while so-called hyperpartisan sites publish and spread rumors and [conspiracy theories](#) to influence public opinion.

This industry is bolstered by how easy it is to create [social bots](#), fake accounts controlled by software that look like real people and therefore can have real influence. Research in my [lab](#) uncovered many examples of fake grassroots campaigns, also called [political astroturfing](#).

In response, we developed the [BotOrNot](#) tool to detect social bots. It's not perfect, but [accurate enough](#) to uncover persuasion campaigns in the Brexit and antivax movements. Using BotOrNot, our colleagues found that a large portion of online chatter about the 2016 elections was generated by bots.

Creating information bubbles


We humans are vulnerable to manipulation by digital misinformation thanks to a complex set of social, cognitive, economic and algorithmic biases. Some of these have evolved for good reasons: Trusting signals from our social circles and rejecting information that contradicts our experience served us well when our species adapted to evade predators. But in today's shrinking online networks, a social network connection with a conspiracy theorist on the other side of the planet does not help inform my opinions.

Copying our friends and unfollowing those with different opinions give

us echo chambers so [polarized](#) that researchers can tell with high accuracy whether you are [liberal or conservative](#) by just looking at your friends. The network structure is so [dense](#) that any misinformation spreads almost instantaneously within one group, and so segregated that it does not reach the other.

Inside our bubble, we are selectively exposed to information aligned with our beliefs. That is an ideal scenario to maximize engagement, but a detrimental one for developing healthy skepticism. [Confirmation bias](#) leads us to share a headline [without even reading](#) the article.

Our lab got a personal lesson in this when our own research project became the subject of a [vicious misinformation campaign](#) in the run-up to the 2014 U.S. midterm elections. When we investigated what was happening, we found [fake news](#) stories about our research being predominantly shared by Twitter users within one partisan echo chamber, a large and homogeneous community of politically active users. These people were quick to retweet and impervious to debunking information.



In this graph of echo chambers in the Twittersphere, purple dots represent people spreading false claims about the Truthy research project; the two accounts that sought to debunk the false information are in orange on the far left. Credit: Giovanni Luca Ciampaglia, CC BY-ND

Viral inevitability

Our research shows that given the structure of our social networks and our limited attention, it is [inevitable](#) that some memes will go viral, irrespective of their quality. Even if individuals tend to share information of higher quality, the network as a whole is not effective at discriminating between reliable and fabricated information. This helps explain all the viral hoaxes we observe in the wild.

The [attention economy](#) takes care of the rest: If we pay attention to a certain topic, more information on that topic will be produced. It's cheaper to fabricate information and pass it off as fact than it is to report actual truth. And fabrication can be tailored to each group: Conservatives read that the pope endorsed Trump, liberals read that he endorsed Clinton. [He did neither](#).

Beholden to algorithms

Since we cannot pay attention to all the posts in our feeds, algorithms determine what we see and what we don't. The algorithms used by [social media platforms](#) today are designed to prioritize engaging posts – ones we're likely to click on, react to and share. But a recent analysis found intentionally misleading pages got at least [as much online sharing](#) and reaction as real news.

This algorithmic bias toward engagement over truth reinforces our social and cognitive biases. As a result, when we follow links shared on social media, we tend to visit a smaller, [more homogeneous](#) set of sources than when we conduct a search and visit the top results.

Existing research shows that being in an echo chamber can make people [more gullible](#) about accepting unverified rumors. But we need to know a lot more about how different people respond to a single hoax: Some

share it right away, others fact-check it first.

We are [simulating a social network](#) to study this competition between sharing and fact-checking. We are hoping to help untangle conflicting evidence about [when fact-checking helps stop hoaxes](#) from spreading and when it doesn't. Our preliminary results suggest that the more segregated the community of hoax believers, the longer the hoax survives. Again, it's not just about the hoax itself but also about the network.

Many people are trying to figure out [what to do about all this](#). According to Mark Zuckerberg's latest [announcement](#), Facebook teams are testing potential options. And a group of college students has proposed a way to simply [label shared links](#) as "verified" or not.

Some solutions remain out of reach, at least for the moment. For example, we can't yet teach artificial intelligence systems how to [discern between truth and falsehood](#). But we can tell ranking algorithms to give higher priority to more reliable sources.

Studying the spread of fake news

We can make our fight against fake news more efficient if we better understand how bad information spreads. If, for example, bots are responsible for many of the falsehoods, we can focus attention on detecting them. If, alternatively, the problem is with echo chambers, perhaps we could design recommendation systems that don't exclude differing views.

To that end, our lab is building a platform called [Hoaxy](#) to track and visualize the spread of unverified claims and corresponding fact-checking on social media. That will give us real-world data, with which we can inform our simulated social networks. Then we can test possible

approaches to fighting fake news.

Hoaxy may also be able to show people how easy it is for their opinions to be manipulated by online information – and even how likely some of us are to share falsehoods online. Hoaxy will join a suite of tools in our [Observatory on Social Media](#), which allows anyone to see how memes spread on Twitter. Linking tools like these to human fact-checkers and [social media](#) platforms could make it easier to minimize duplication of efforts and [support](#) each other.

It is imperative that we invest resources in the study of this phenomenon. We need all hands on deck: Computer scientists, social scientists, economists, journalists and industry partners must [work together](#) to stand firm against the spread of misinformation.

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

Provided by The Conversation

Citation: Misinformation on social media—can technology save us? (2016, November 28) retrieved 26 June 2024 from <https://phys.org/news/2016-11-misinformation-social-media-can-technology.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.