



BEHAVIOR & SOCIETY

How Fake News Goes Viral—Here's the Math

Models similar to those used to track disease show what happens when too much information hits social media networks

.....
By Madhusree Mukerjee on July 14, 2017



Credit: Peter Dazeley Getty Images

ADVERTISEMENT | REPORT AD

NASA runs a child-slave colony on Mars!

Photos taken by a Chinese orbiter reveal an alien settlement on the moon!

Shape-shifting reptilian extraterrestrials that can control human minds are running the U.S. government!

What drives the astonishing popularity of such stories? Are we a particularly gullible species? Perhaps not—maybe we're just overwhelmed. A bare-bones model of how news spreads on social media, published in June in *Nature Human Behavior*, indicates that just

about anything can go viral. Even in a perfect world, where everyone wants to share real news and is capable of evaluating the veracity of every claim, some fake news would still reach thousands (or even millions) of people, simply because of information overload. It is often impossible to see everything that comes into one's news feed, let alone confirm it. "If you live in a world where you are bombarded with junk—even if you're good at discriminating—you're only seeing a portion of what's out there, so you still may share misinformation," explains computer scientist Filippo Menczer of Indiana University Bloomington (I.U.), one of the model's co-authors. "The competition is so harsh that the good stuff cannot bubble to the top."

Chances are that in the virtual world, the beauty of a photograph or the persuasiveness of an article do help to spread a "meme"—the term Menczer and his colleagues use for a link, video, phrase or other unit of online information. The researchers demonstrate, however, that just three inexorable factors can explain a network's inability to distinguish truth from falsehood in memes, even if individuals can. They are: the enormous amount of information out there; the limited amount of time and attention people can devote to scrolling through their news feeds and choosing what to share; and the structure of the underlying social networks. All three conspire to spread some of the worst memes at the expense of the best ones.

Mathematical models for exploring how memes spread on social media networks are known as agent-based models because they require the active participation of "agents," a techie term for individuals. These models originate from an older class of simulations that study how diseases spread through a community. Think of a diagram in which each agent is represented by a dot, or node, and is linked via lines to other nodes, representing friends or followers. If, say, Alice is "infected" by an influenza virus or a piece of fake news, she may transmit the contagion along these links to her friends Bob and Clive by shaking hands or sharing the meme with them, respectively. Bob and Clive could in turn pass the contagion to their contacts, and so on. By fleshing out this skeletal framework, scientists try to simulate how far a meme can spread under different conditions.

"Information is not a virus," however, cautions information scientist Kristina Lerman of the University of Southern California, who was not involved in creating the new model. Whereas we are usually dealing with one flu strain at a time, or at worst a few, the number of memes competing to infect us is staggering. The modelers incorporate this abundance by imagining that each person has a screen on which he or she views incoming memes.

The model assigns a value to the probability that Alice will create and share a new meme—say, a video she has made of her dancing cockatoo—and it also does so for all the possible new memes originating from all other users. Because new memes increase the total quantity of information in the system, these values measure the information load experienced by those viewing their screens.

Another parameter tracks the number of items Alice views on her long news feed before choosing to simply pass an existing meme along to her connections, instead of creating a new one. This parameter serves as a proxy for the attention span—the information that Alice focused on. Once Alice sends along a message, it appears on the screens of Bob, Clive and others, who in turn choose whether to create memes of their own or to transmit one of them from their feeds.

Using an earlier version of this model, Menczer and others at I.U. showed in 2012 that a few memes will go viral even if all memes are equally “contagious”—that is, equally likely to be shared each time they are viewed. The memes in both models roughly follow what is called a “power law,” meaning that the chance of a meme being tweeted or otherwise shared a certain number of times decreases as an inverse power of that number. For example, a meme is four times less likely to be tweeted twice than once. “If you look at the distribution of pictures on Flickr or articles on Facebook or hashtags on Twitter—all of these have power laws,” Menczer says. Still, memes reaching thousands of recipients are surprisingly commonplace.

In 2014 mathematician James Gleeson of the University of Limerick in Ireland and others demonstrated a mathematical similarity between models of the kind concocted by Menczer, among others, and “sandpiles”—canonical systems for what physicists call “self-organized criticality.” If one gently dribbles sand onto a flat surface, it will pile up until its slopes reach a critical angle. A few additional grains of sand may cause nothing much to happen, but all of a sudden yet another grain will trigger an avalanche: the equivalent of a meme going viral. Gleeson's analysis suggests the intrinsic properties of the system, as opposed to the particularities of a meme, are driving virality.

In the latest paper Menczer, Xiaoyan Qiu and others at I.U. examine what happens if some memes are more contagious than others. They find that if the information load is low and the attention span is high, the more attractive memes prevail. Actual tracking of attention

and information overload, obtained from Twitter and Tumblr data, however, indicate that in real life the sheer quantity of information usually overwhelms us. “You don’t have to assume that the reason why junk spreads is because people like it or because they can’t tell the difference,” Menczer explains. “You could assume that people do know the difference, and still the fake stuff would go viral, simply because of information overload.”

One key factor influencing the spread of memes is the pattern of connections in the underlying social media network. “Some network structures will promote the spread being fast and others will inhibit the spread,” says mathematician Mason Porter of the University of California, Los Angeles. If the simulated network in the competition-driven model is assumed to be random, for instance—meaning the connections are randomly distributed among nodes on the network—no memes go viral. Real social media networks display a roughly power-law distribution of links, however—a feature Menczer and his colleagues incorporate into their simulation. So whereas most of us—each a node on Twitter, for example—have a handful of followers, a few outliers may have tens of thousands. If any of these “superconnected” individuals, or hubs, becomes infected with a fake meme, they can presumably transmit it far and wide.

But U.S.C.’s Lerman begs to differ. In disease models, highly connected people are called “superspreaders” because they help drive epidemics. By examining the behavior of actual Twitter users, however, she demonstrated in 2016 that superconnected agents pass on very few of the memes they receive. This is because they cannot possibly see, let alone read, everything in their staggeringly lengthy feeds. “People who are highly connected are unlikely to see anything that is even five minutes old because it is so far down their feed,” she notes. Thus information overload ensures they are less likely to get infected in the first place. In her view, hubs suppress the vast majority of memes but may help to spread the few they let through.

Also playing a role in virality: friends tend to form clusters. So, for instance, because Alice knows Bob and Clive, the latter likely know each other as well, and likely share similar views on many issues. These clusters help establish what social media aficionados think of as an “echo chamber.” Most of us tend to see some memes several times, increasing the likelihood that we too will share them. Making matters worse, the contagiousness of a meme—unlike that of a flu virus—depends on how often it has been shared. In a Web-based experiment involving more than 14,000 volunteers, sociologist Matthew Salganik, then at Columbia University, and others showed in 2006 that recruits were much more

likely to download a particular song if they were aware that their peers liked it.

Such “social reinforcement” can ensure that contagiousness increases sharply once a certain threshold of exposure is crossed. “You see one person post, ‘NASA’s got slave colonies on Mars,’ and you think, ‘That’s ridiculous,’” Porter explains. “You see a second person post, ‘NASA’s got slave colonies on Mars.’ You see this many times, and it somehow becomes more plausible the more times you see it.” And so you share it, too. Several research groups are exploring the intricate cognitive processes that lead to one meme being chosen over another.

Debate persists, though, on the accuracy of the models used in this research. “In general, I tend to be skeptical of agent-based models because there are so many knobs you can tweak,” Lerman says. Menczer concedes that any model used in attempt to reproduce all the subtleties of human cognitive behavior would have many unknown parameters—or “knobs”—which would make their results hard to interpret. But that is less of a problem with such minimalistic models (often called “toy models”), which seek only to explore broad-brush features. “As long as they are very simple, they are useful,” Menczer says—because they reveal surprisingly powerful truths.

[Rights & Permissions](#)



ABOUT THE AUTHOR(S)

Madhusree Mukerjee

Recent Articles

5 Years Later, the Fukushima Nuclear Disaster Site Continues to Spill Waste

Crippled Fukushima Reactors Are Still a Danger, 5 Years after the Accident

The Impending Dam Disaster in the Himalayas

READ THIS NEXT

BEHAVIOR & SOCIETY

The Ultimate Cure for the Fake News Epidemic Will Be More Skeptical Readers

February 1, 2017 — David Pogue



ELECTRONICS

What Facebook Is Doing to Combat Fake News

February 1, 2017 — David Pogue



ARTS & CULTURE

3 Historical Examples of "Fake News"

December 1, 2016 — Krystal [REDACTED]

