

Misinformation and biases infect social media

By Giovanni Luca Ciampaglia and Filippo Menczer, *The Conversation*

Social media are among the primary sources of news in the U.S. and across the world. Yet users are exposed to content of questionable accuracy, including conspiracy theories, clickbait, hyperpartisan content, pseudo science and even fabricated “fake news” reports.

It’s not surprising that there’s so much disinformation published: Spam and online fraud are lucrative for criminals, and government and political propaganda yield both partisan and financial benefits. But the fact that low-credibility content spreads so quickly and easily suggests that people and the algorithms behind social media platforms are vulnerable to manipulation.

Our research has identified three types of bias that make the social media ecosystem vulnerable to both intentional and accidental misinformation. That is why our Observatory on Social Media at Indiana University is building tools to help people become aware of these biases and protect themselves from outside influences designed to exploit them.

Bias in the brain

Cognitive biases originate in the way the brain processes the information that every person encounters every day. The brain can deal with only a finite amount of information, and too many incoming stimuli can cause information overload. That in itself has serious implications for the quality of information on social media. We have found that steep competition for users’ limited attention means that some ideas go viral despite their low quality – even when people prefer to share

high-quality content.

To avoid getting overwhelmed, the brain uses a number of tricks. These methods are usually effective, but may also become biases when applied in the wrong contexts.

One cognitive shortcut happens when a person is deciding whether to share a story that appears on their social media feed. People are very affected by the emotional connotations of a headline, even though that's not a good indicator of an article's accuracy. Much more important is who wrote the piece.

To counter this bias, and help people pay more attention to the source of a claim before sharing it, we developed Fakey, a mobile news literacy game (free on Android and iOS) simulating a typical social media news feed, with a mix of news articles from mainstream and low-credibility sources. Players get more points for sharing news from reliable sources and flagging suspicious content for fact-checking. In the process, they learn to recognize signals of source credibility, such as hyperpartisan claims and emotionally charged headlines.

Bias in society

Another source of bias comes from society. When people connect directly with their peers, the social biases that guide their selection of friends come to influence the information they see.

In fact, in our research we have found that it is possible to determine the political leanings of a Twitter user by simply looking at the partisan preferences of their friends. Our analysis of the structure of these partisan communication networks found social networks are particularly efficient at disseminating information – accurate or not – when they are closely tied together and disconnected from other parts of society.

The tendency to evaluate information more favorably if it comes from within their own social circles creates “echo chambers” that are ripe for manipulation, either consciously or unintentionally. This helps explain why so many online conversations devolve into “us versus them” confrontations.

To study how the structure of online social networks makes users vulnerable to disinformation, we built Hoaxy, a system that tracks and visualizes the spread of content from low-credibility sources, and how it competes with fact-checking content. Our analysis of the data collected by Hoaxy during the 2016 U.S. presidential elections shows that Twitter accounts that shared misinformation were almost completely cut off from the corrections made by the fact-checkers.

When we drilled down on the misinformation-spreading accounts, we found a very dense core group of accounts retweeting each other almost exclusively – including several bots. The only times that fact-checking organizations were ever quoted or mentioned by the users in the misinformed group were when questioning their legitimacy or claiming the opposite of what they wrote.

Bias in the machine

The third group of biases arises directly from the algorithms used to determine what people see online. Both social media platforms and search engines employ them. These personalization technologies are designed to select only the most engaging and relevant content for each individual user. But in doing so, it may end up reinforcing the cognitive and social biases of users, thus making them even more vulnerable to manipulation.

For instance, the detailed advertising tools built into many social media platforms let disinformation campaigners exploit confirmation bias by tailoring messages to people who are already inclined to believe them.

Also, if a user often clicks on Facebook links from a particular news source, Facebook will tend to show that person more of that site's content. This so-called "filter bubble" effect may isolate people from diverse perspectives, strengthening confirmation bias.

Our own research shows that social media platforms expose users to a less diverse set of sources than do non-social media sites like Wikipedia. Because this is at the level of a whole platform, not of a single user, we call this the homogeneity bias.

Another important ingredient of social media is information that is trending on the platform, according to what is getting the most clicks. We call this popularity bias, because we have found that an algorithm designed to promote popular content may negatively affect the overall quality of information on the platform. This also feeds into existing cognitive bias, reinforcing what appears to be popular irrespective of its quality.

All these algorithmic biases can be manipulated by social bots, computer programs that interact with humans through social media accounts. Most social bots, like Twitter's Big Ben, are harmless. However, some conceal their real nature and are used for malicious intents, such as boosting disinformation or falsely creating the appearance of a grassroots movement, also called "astroturfing." We found evidence of this type of manipulation in the run-up to the 2010 U.S. midterm election.

To study these manipulation strategies, we developed a tool to detect social bots called Botometer. Botometer uses machine learning to detect bot accounts, by inspecting thousands of different features of Twitter accounts, like the times of its posts, how often it tweets, and the accounts it follows and retweets. It is not perfect, but it has revealed that as many as 15 percent of Twitter accounts show signs of being bots.

Using Botometer in conjunction with Hoaxy, we analyzed the core of the misinformation network during the 2016 U.S. presidential campaign. We found many bots exploiting both the cognitive, confirmation and popularity biases of their victims and Twitter's algorithmic biases.

These bots are able to construct filter bubbles around vulnerable users, feeding them false claims and misinformation. First, they can attract the attention of human users who support a particular candidate by tweeting that candidate's hashtags or by mentioning and retweeting the person. Then the bots can amplify false claims smearing opponents by retweeting articles from low-credibility sources that match certain keywords. This activity also makes the algorithm highlight for other users false stories that are being shared widely.

Understanding complex vulnerabilities

Even as our research, and others', shows how individuals, institutions and even entire societies can be manipulated on social media, there are many questions left to answer. It's especially important to discover how these different biases interact with each other, potentially creating more complex vulnerabilities.

Tools like ours offer internet users more information about disinformation, and therefore some degree of protection from its harms. The solutions will not likely be only technological, though there will probably be some technical aspects to them. But they must take into account the cognitive and social aspects of the problem.

Giovanni Luca Ciampaglia is an assistant research scientist at Indiana University Network Science Institute, Indiana University; Filippo Menczer is a professor of computer science and informatics, director of the Center for Complex Networks and Systems Research, Indiana University.