

Deceiving the Masses on Social Media

The social media platforms like their freedom, but information gerrymandering may require legislation to fix.

CONSIDERABLE ATTENTION HAS been paid to the impact of social media on the electoral process, given that 55% of U.S. adults now get their news from social media either “often” or “sometimes”—an 8% increase from the previous year, according to a Pew Research study published in October 2019, which was conducted in July 2019. This is concerning because, according to the Pew data, 88% of Americans understand and realize that social media companies now have at least some control over the mix of the news consumed each day, and 62% believe social media companies have far too much control over the content mix of news that is seen each day.

Much of the concern about social media companies controlling the news is visceral. However, a study published in the journal *Nature* in September 2019 identified and explained mathematically how social media companies may unwittingly become a disruptive force to the democratic process, via a concept called information gerrymandering.

In electoral gerrymandering, political district boundaries are drawn by the party or group in power to create an unfair political advantage for them. In one scenario, voting district lines are drawn so the voting power of the opposing party’s supporters are spread out across many districts, thereby leaving the party in power with a majority of voters in a single district. Another tactic used in gerrymandering is to redraw district boundary lines so the voting power of an opposing group is concentrated in one district, thereby reducing their voting power in other districts. A third tactic is when a majority party seeks to manipulate district boundaries so that in each district, the majority power will always retain a population advantage.



Researchers from the University of Pennsylvania, the University of Houston, and the Massachusetts Institute of Technology have hypothesized that a similar phenomenon may be occurring in online social networks, driven by a mismatch or inequality of viewpoints reaching users, thereby creating information networks that may skew user perceptions, unbeknownst to the members of that social media group. A social media group could be defined as the followers of a particular social media personality, people who follow a particular hashtag, or the people who read or retweet a specific news article or periodical online. In this scenario, the information gerrymandering occurs by grouping users not just by their linkages to each other (such as by whom they follow on Twitter, which hashtags they follow, or which news sources they follow), but by the information viewpoint to which they are exposed.

The researchers began studying the power of information to change opinions by placing research participants in simulated elections, then conducting experiments to see how a group of people, evenly divided on an issue, might change their opinions if exposed to additional information that was weighted to one side of an issue, or if exposed to “zealots,” humans or bots that would argue for only one side of an issue.

After conducting repeated simulated elections around a specific issue, the researchers found that people may change their views based on additional and repeated exposure to information, even if it is contrary to their long-held viewpoint. Further, people may change their views to “go along with the group,” particularly if it appears that the majority of the group is voting a certain way. They then applied their research to an analysis of Twitter users.

“So what we found in this paper is that even when two parties have the exact same number of members, and when everyone is equally influential on a social network, the network can bias the outcome of a vote in favor of one party,” says Joshua Plotkin, Walter H. and Leonore C. Annenberg Professor of the Natural Sciences at the University of Pennsylvania, a co-author of the study.

Plotkin and his team noted that the structure and opacity of how people are connected on Twitter and other online social networks may lead to users being over-exposed to news and viewpoints that don’t line up with the expected viewpoints of the hashtags, influencers, or news sources they follow; for example, if one user follows another user who constantly retweets news stories, opinion pieces, or other content that comes from sources that the first user is not familiar with or cannot verify. Repeated exposure to this information or viewpoint may be able to sway susceptible users, and they may not even be aware it’s happening, according to Plotkin.

“The basic idea is that one party can be at a disadvantage even if it has the same number of members as the other party,” Plotkin says, based on how the network is structured, the type or bias of information delivered to the group, and the frequency of that information delivery.

This is not a problem, per se; in a democracy, different viewpoints, particularly those that challenge people to think rather than simply go along with their preconceived notions, can be a good thing. The problem, according to Alexander Stewart, an assistant professor in the Department of Biology and Biochemistry at the University of Houston and a co-author of the study, is that social networks do not make clear how their algorithms connect or separate users.

This is particularly the case on Twitter, Stewart says, as that service will suggest that a user follow other users or news sources without explicitly indicating the rationale, or what the linkage may be between them. Twitter—or any other online platform, for that matter—may not be able to police or stop a user

from posting information that comes from dubious sources, or stating one’s opinion as a fact.

Stewart doesn’t believe Twitter or any of the other social networks have set up their networks to favor one side or the other of certain arguments deliberately. “I don’t think it’s plausible to suggest that there is anybody deliberately wiring networks in a way to give one side an advantage over the other,” Stewart says. “Rather, what is being reflected are naturally emerging asymmetries which reflect dominance of one side of the discussion over another which arise both due to human behavior, but also it is due to choices made by the platforms that we use. Whether or not I see your tweets depends on whether I follow you, but also on whether your tweets appear in my news feed in a more prominent or less prominent way.”

Still, he says his work around information gerrymandering is useful for citizens to consider, particularly if they are getting the majority of their news and information from tweets or their Facebook news feeds, and if they are actively participating by retweet-

Milestones

National Academy of Engineering Names Three from ACM among Newest Members

Among the 87 new members and 18 international members recently elected to the National Academy of Engineering (NAE) are three who also belong to ACM, including past president Vicki L. Hanson.

NAE membership honors those who have made outstanding contributions to “engineering research, practice, or education, including, where appropriate, significant contributions to the engineering literature” and to “the pioneering of new and developing fields of technology, making major advancements in traditional fields of engineering, or developing/implementing innovative approaches to engineering education.”

Hanson was honored with NAE membership “for contributions to the design of accessible systems, and for leadership in the computer science and engineering community.” She served as ACM president from 2016 to 2018,

was a Distinguished Professor of the Rochester Institute of Technology in the HCI and Accessibility research groups, and also professor and chair of inclusive technologies at Scotland’s University of Dundee.

Hanson is an ACM Fellow, as well as a Fellow Chartered Information Technology Professional of the British Computer Society, and a Fellow of the Royal Society of Edinburgh. She has received the Royal Society Wolfson Research Merit Award, the ACM SIGCHI Social Impact Award, the Women of Vision ABIE Award for Social Impact, and the ACM SIGACCESS Award for Outstanding Contributions to Computing and Accessibility. She was elected to the ACM SIGCHI Academy in 2017.

Also newly elected to NAE membership were James F. Kurose and Fei Fei Li.

Kurose, a distinguished professor in the College of

Information and Computer Science at the University of Massachusetts, Amherst, was honored with NAE membership “for contributions to the design and analysis of network protocols for multimedia communication.”

Kurose has been on leave from the University of Massachusetts since 2015, serving as the assistant director of the National Science Foundation for Computer and Information Science and Engineering. Kurose also co-chairs the Networking and Information Technology Research and Development Subcommittee of the National Science and Technology Council Committee on Technology.

In the course of his career, Kurose has been awarded the IEEE’s Taylor Booth Award, the IEEE’s INFOCOM Achievement Award, and the ACM SIGCOMM Special Interest Group Lifetime achievement award.

Li, a professor at Stanford University and co-director of Stanford’s Human-Centered AI Institute and the Stanford Vision and Learning Lab, was honored with NAE membership “for contributions in building large knowledge bases for machine learning and visual understanding.”

Li served as director of the Stanford Artificial Intelligence Laboratory from 2013 to 2018. In 2017, she co-founded AI4ALL, a nonprofit aimed at increasing diversity in the field of artificial intelligence.

Among the awards Li has received are the IEEE PAMI Mark Everingham Prize, the J.K. Aggarwal Prize of the International Association for Pattern Recognition, the WITI@UC Athena Award for Academic Leadership of the University of California, and the Technical Leadership Abie Award of AnitaB.org. She was named a Fellow by ACM, and one of America’s Top 50 Women in Tech by *Forbes*.

ing, forwarding, and commenting on news stories. Indeed, an August 2019 Pew Research Center study found Facebook, along with Twitter and Reddit, had the highest proportion of respective users seeking news, with 73% of Facebook users seeking news on the platform, compared with 71% of Twitter users, and 65% of Reddit users.

This greater level of engagement on social media platforms increases the potential for misinformation, compared with 20 or 30 years ago, when most people passively watched television or read newspapers and did not actively share content across a wide circle of influence. Implications within a democracy may be profound, for those who may be influencing political thoughts and viewpoints, as well as the citizens who consume, repeat, and amplify them. Indeed, information gerrymandering may create social networks or groups in which repeated exposure to a particular viewpoint is “wasted,” as the group consists primarily of like-minded individuals. It also may inadvertently create groups of people that are particularly susceptible to changing their minds on an issue, if constantly and repeatedly exposed to another point of view.

“Are these platforms skewing our conversations, not necessarily deliberately, but in ways that makes it harder for us to reach informed decisions, or engage in compromise and collective decision making?” Stewart asks. “If people are being exposed to asymmetric information, and they can be made aware of that, then you can suggest to them how they might be able to combat that” by encouraging them to consider the sources of the material they consume.

Plotkin says additional regulation of how social media sites connect people, serve up news stories, and suggest tweets, should focus on making these algorithms more transparent so users can see how they are connected and why they are being fed specific new stories or content. That said, social media companies are unlikely to share the ‘secret sauce’ behind how they match users and content, given that is how they generate revenue.

“Technology data analysis and social media companies often claim that their algorithms are intellectual property and confidential trade secrets,” says Linda Priebe, a U.S.-EU data pri-

Information gerrymandering could inadvertently create groups of people susceptible to changing their minds on an issue, if repeatedly exposed to another point of view.

vacy/security and federal relations attorney with Dallas, TX-based law firm Culhane Meadows PLLC. As a result, the social media companies “don’t want to lose what they believe to be their competitive advantage by being transparent about what the algorithms do, how they function, what information they rely on, [and even] to what extent they’re accurate.”

Priebe also notes that many companies feel, particularly around the distribution of news or political information, it’s not the social media companies’ role to become an arbiter of thought or discussion. “When it comes to maybe what people consider to be fake news and political advertising, they feel that they have a First Amendment right, or their customers have a First Amendment right, to express their political opinion and it’s not their role to censor that,” Priebe says. “So, there’s a concern there and I think all of that needs to be balanced in some way; fair to companies, but also fair to consumers.”

It should be noted that for its part, Twitter has prohibited political advertising, though the platform still allows people to share political news and views, so long as it does not advocate for or support a specific position or candidate. As such, information gerrymandering can (and likely still does) occur by users linking to or sharing content among their connections.

Other social media platforms have taken differing approaches. Facebook’s Mark Zuckerberg has taken the

position that political content is fair game, and relies on its users to determine the veracity and accuracy of political advertisements. Google, which includes YouTube, bans misinformation (via a team of fact-checkers who compare ad content to known and verified information sources) in some ads, such as around voting procedures, but does not have a policy prohibiting politicians from running false or misleading ads. Reddit, meanwhile, continues to permit political issue ads, and allows ads from political candidates at the federal level, but bans advertising around state or local elections.

Balancing transparency against competitive needs may make enacting legislation addressing information gerrymandering via regulation of algorithms extremely challenging.

“I think transparency is really key,” Priebe says. “What the Europeans have done with GDPR [the General Data Protection Regulation] is, they’ve striven to balance the business interests and concerns with individual and consumer privacy and protection concerns.”

Indeed, GDPR mandates specific consumer information disclosures, while also protecting the trade secrets of companies. For many U.S.-based companies, however, “my sense is that they’ve been very comfortable with the fact that their internal operations are pretty much behind the scenes, and not subject to any oversight,” Priebe says. As a result, she adds, “They can generate a lot of revenue and feel that European approach would restrict their revenue streams.”

Further Reading

Stewart, A.J., Mosleh, M., Diakonova, M. et al. Information gerrymandering and undemocratic decisions. *Nature* 573, 117–121 (2019) doi:10.1038/s41586-019-1507-6 <https://www.nature.com/articles/s41586-019-1507-6>

How Social Media Sites Handle Political Ads: <https://reut.rs/2TTBz9t>

What is Information Gerrymandering: <https://www.youtube.com/watch?v=SOIUHQkdaa0>

What is GDPR: <https://gdpr.eu/>

Keith Kirkpatrick is principal at 4K Research & Consulting, LLC, based in New York City, NY, USA.

© 2020 ACM 0001-0782/20/5 \$15.00