MIT INITIATIVE ON THE DIGITAL ECONOMY RESEARCH BRIEF

THE SPREAD OF TRUE AND FALSE NEWS ONLINE

By Soroush Vosoughi, Deb Roy, and Sinan Aral



FALSE NEWS IS BIG NEWS.

Barely a day goes by without a new development about the veracity of social media, foreign meddling in U.S. elections, or questionable science.

Adding to the confusion is speculation about what's behind such developments—is the motivation deliberate and political, or is it a case of uninformed misinformation? And who is spreading the word online—rogue AI bots or agitated humans?

These were among the questions we sought to address in the largest-ever longitudinal study of the spread of false news online. Until now, few large-scale empirical investigations existed on the diffusion of misinformation or its social origins. Studies about the spread of misinformation were limited to analyses of small, ad hoc samples. But these ad hoc studies ignore two of the most important scientific questions: How do truth and falsity diffuse differently, and what factors related to human judgment explain these differences?

Understanding how false news spreads is the first step toward containing it. With this research in hand, we can consider the implications of false news on hotly debated issues –– from the regulation of social media sites such as Facebook and Twitter, to social media's role in elections.

REDEFINING NEWS

The basic concepts of truth and accuracy are central to theories of decision-making [1, 2, 3], cooperation [4], communication [5], and markets [6]. Today's online media adds new dimensions and complexity to this field of study.

There has been a lot of attention given to the impact of social media on our democracy and our politics. In addition to politics, false rumors have affected stock prices and the motivation for large scale investments. Indeed, our responses to everything from natural disasters [7, 8] to terrorist attacks [9] have been disrupted by the spread of false news online.



RESEARCH HIGHLIGHTS

We investigated the differential diffusion of all the verified, true and false news stories distributed on Twitter from 2006 to 2017. The data comprise approximately 126,000 cascades of news stories spreading on Twitter, tweeted by about 3 million people over 4.5 million times.

We classified news as true or false using information from six independent fact-checking organizations that exhibited 95% -98% agreement on the classifications.

Falsehood diffused significantly farther, faster, deeper, and more broadly than the truth in all categories. The effects were most pronounced for false political news than for news about terrorism, natural disasters, science, urban legends, or financial information.

Controlling for many factors, false news was 70% more likely to be retweeted than the truth.

Novelty is an important factor. False news was perceived as more novel than true news, which suggests that people are more likely to share novel information.

Contrary to conventional wisdom, robots accelerated the spread of true and false news at the same rate, implying that humans, not robots, are more likely responsible for the dramatic spread of false news.

New social technologies, notably Twitter, Facebook, and photo-sharing apps, facilitate rapid information-sharing and large-scale information "cascades" that can also spread misinformation, or information that is inaccurate or misleading.

THE SPREAD OF TRUE AND FALSE NEWS ONLINE

By Soroush Vosoughi, Deb Roy, and Sinan Aral

But, while more of our access to information and news is guided by these new technologies [10] we know little about their exact contribution to the spread of falsity online. Anecdotal analyses of false news by the media [11] are getting lots of attention, but there are few large-scale empirical investigations of the diffusion of misinformation or its social origins.

Current research has analyzed the spread of single rumors, like the discovery of the Higgs boson [12], or the Haitian earthquake of 2010 [13]. Others have studied multiple rumors from a single disaster event, like the Boston Marathon bombing of 2013. Theoretical models of rumor diffusion [14], or methods for rumor detection [15], credibility evaluation [16, 17], or interventions to curtail the spread of rumors, can also be found.

Yet, almost no studies comprehensively evaluate differences in the spread of truth and falsity across topics nor do they examine why false news may spread differently than the truth. That was our goal.

To understand the spread of false news, our research examines the diffusion of true and false news on Twitter.

FACT-CHECKING THE RUMORS

A rumor cascade begins on Twitter when a user makes a statement about a topic in a tweet, which could include written text, photos, or links to articles online. Other users propagate the rumor by retweeting it. A rumor's diffusion process can be characterized as having one or more "cascades," which we define as "a rumorspreading pattern that exhibit an unbroken retweet chain with a common, singular origin."

For example, an individual could start a rumor cascade by tweeting a story or claim with an assertion in it, and another individual independently starts a second cascade of the same rumor that is completely independent of the first, except that it pertains to the same story or claim.





Our investigation looked at a highly comprehensive dataset of all of the fact-checked rumor cascades that spread on Twitter from its inception in 2006 until 2017. The data include approximately 126,000 rumor cascades spread by about 3 million people over 4.5 million times.

The next problem we addressed was how to fact-check the tweets. All rumor cascades were investigated by six independent fact-checking organizations: snopes.com, politifact.com, factcheck.org, truthorfiction.com, hoax-slayer. com, and urbanlegends.about.com. Then, we parsed the title, body, and verdict (true, false or mixed) of each rumor investigation reported on their websites, and automatically collected the cascades corresponding to those rumors on Twitter. The result was a sample of rumor cascades whose veracity had been agreed upon by these organizations 95% to 98% of the time.

We quantified the cascades into four categories:

1. Depth: The number of retweet hops from the origin tweet over time;

2. Size: The number of users involved in the cascade over time;

3. Maximum breadth: The full number of users involved in the cascade at any depth;

4. Structural virality: A measure that interpolates between content spread through a single, large broadcast and content spread through multiple generations, with any one individual directly responsible for only a fraction of the total spread. [19]

Our results were dramatic: Analysis found that it took the truth approximately six times as long as falsehood to reach 1,500 people and 20 times as long as falsehood to reach a cascade depth of ten.

As the truth never diffused beyond a depth of ten, we saw that falsehood reached a depth of 19 nearly ten times faster than the truth reached a depth of ten. Falsehood also diffused significantly more broadly and was retweeted by more unique users than the truth at every cascade depth. MIT INITIATIVE ON THE DIGITAL ECONOMY RESEARCH BRIEF

THE SPREAD OF TRUE AND FALSE NEWS ONLINE

By Soroush Vosoughi, Deb Roy, and Sinan Aral

THE VIRALITY AND NOVELTY OF FALSE NEWS

In particular, we determined that false political news traveled deeper and more broadly, reached more people, and was more viral than any other category of false information. False political news also diffused deeper more quickly, and reached more than 20,000 people nearly three times faster than all other types of false news reached 10,000 people.

Furthermore, analysis of all news categories showed that news about politics, urban legends, and science spread to the most people, while news about politics and urban legends spread the fastest and were the most viral. When we estimated a model of the likelihood of retweeting we found that falsehoods were fully 70% more likely to be retweeted than the truth.

What could explain such surprising results? One explanation emerges from information theory and Bayesian decision theory: People thrive on novelty. As others have noted, novelty attracts human attention [20], contributes to productive decision making [21], and encourages information-sharing [22]. In essence, it can update our understanding of the world. When information is novel, it is not only surprising, but also more valuable-both from an information theory perspective (it provides the greatest aid to decision-making), and from a social perspective (it conveys social status that one is 'in the know,' or has access to unique 'inside' information).

To check the results, we tested whether falsity was more novel than the truth, and whether Twitter users were more likely to retweet information that was more novel. The tests confirmed our findings. Numerous diagnostic statistics and checks validated our results and confirmed their robustness. Moreover, in case there was concern that our conclusions about human judgment were biased by the presence of bots in our analysis, we employed a sophisticated bot-detection algorithm [23] to identify and remove all bots before running the analysis. When we added bot traffic back into the analysis, we found that none of our main conclusions changed—false news still spread







farther, faster, deeper, and more broadly than the truth in all categories of information.

Although the inclusion of bots accelerated the spread of both true and false news, it affected their spread roughly equally. This suggests that contrary to what many believe, false news spreads farther, faster, deeper, and more broadly than the truth because humans, not robots, are more likely to spread it.

SIGNIFICANCE AND RAMIFICATIONS

There are enormous potential ramifications to these results. False news can drive the misallocation of resources during terror attacks and natural disasters, the misalignment of business investments, and can misinform elections. And while the amount of false news online is clearly increasing, our scientific understanding of how and why false news spreads is still largely based on ad hoc rather than large-scale, systematic analyses. Our analysis sheds new light on these trends and affirms that false news spreads more pervasively online than the truth. It also upends conventional wisdom about how false news spreads.

Though one might expect network structure and the characteristics of users to favor and promote false news, the opposite is true. What drives the spread of false news, despite network and individual factors that favor the truth, is the greater likelihood of people to retweet falsity.

Furthermore, while recent testimony before congressional committees on misinformation in the U.S. has focused on the role of bots in spreading false news [24], we conclude that human behavior contributes more to the differential spread of falsity and truth than automated robots do. This implies that misinformation containment policies should emphasize behavioral interventions, like labeling and incentives, rather than focusing exclusively on curtailing bots.

We hope our work inspires more large-scale research into the causes and consequences of the spread of false news as well as its potential cures.

THE SPREAD OF TRUE AND FALSE NEWS ONLINE

By Soroush Vosoughi, Deb Roy, and Sinan Aral

REFERENCES

1. Savage, L.J. (1951) The theory of statistical decision. Journal of the American Statistical Association 46.253 (1951): 55-67.

Simon, H.A. (1960) The new science of management decision. Harper & Brothers, New York, NY, US.
Wedgwood, R. (2002) The aim of belief. Noûs, 36(s16): 267–297. http://onlinelibrary.wiley.com/doi/10.1111/1468-0068.36.s16.10/full

4. Fehr, E., & Fischbacher, U. (2003) The nature of human altruism. Nature, 425(6960), 785–791. https://www.nature.com/articles/nature02043 5. Shannon, C.E. (1948) A mathematical theory of communication. The Bell System Technical Journal, 27: 379–423, 623–656. http://ieeexplore.ieee.org/ document/6773024/

6. Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of political Economy, 100(5): 992-1026. http://www.journals.uchicago.edu/ doi/abs/10.1086/261849

7. Mendoza, M., Poblete, B., & Castillo, C. (2010) Twitter Under Crisis: Can we trust what we RT?. In Proceedings of the first workshop on social media analytics (pp. 71– 79). ACM. https://dl.acm.org/citation.cfm?id=1964869 8. Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013) Faking Sandy: characterizing and identifying fake images on twitter during hurricane sandy. In Proceedings of the 22nd international conference on World Wide Web (pp. 729–736). ACM. https://dl.acm. org/citation.cfm?id=2488033

9. Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, R. M. (2014) Rumors, false flags, and digital vigilantes: Misinformation on twitter after the 2013 Boston marathon bombing. iConference 2014 Proceedings (pp. 654–662) http://hdl.handle. net/2142/47257

10. Gottfried, J., & Shearer, E. (2016) News use across social media platforms. Pew Research Center, 26. http:// www.journalism.org/2016/05/26/news-use-acrosssocial-media-platforms- 2016/

11. Silverman, C. (2016) This Analysis Shows How Viral Fake Election News Stories Outperformed Real News on Facebook. Buzzfeed News

https://www.buzzfeed.com/craigsilverman/viral-fakeelection-news-outperformed-real-news-onfacebook/ 12. De Domenico, M., Lima, A., Mougel, P., & Musolesi, M. (2013). The anatomy of a scientific rumor. Scientific reports, 3, 2980. https://www.nature.com/articles/ srep02980

13. Oh, O., Kwon, K. H., & Rao, H. R. (2010). An Exploration of Social Media in Extreme Events: Rumor





Theory and Twitter during the Haiti Earthquake 2010. In Proceedings of the International Conference on Information Systems, St. Louis, MO: 231. https://asu.pure.elsevier.com/en/publications/anexploration-of-social-media-in-extreme-eventsrumor-

theory-and 14. Tambuscio, M., Ruffo, G., Flammini, A., & Menczer, F. (2015) Fact-checking effect on viral hoaxes: A model of misinformation spread in social networks. In Proceedings of the 24th International Conference on World Wide Web (pp. 977-982). ACM.

https://dl.acm.org/citation.cfm?id=2742572 15. Zhao, Z., Resnick, P., & Mei, Q. (2015) Enquiring minds: Early detection of rumors in social media from enquiry posts. In Proceedings of the 24th International Conference on World Wide Web (pp. 1395–1405). ACM https://dl.acm.org/citation.cfm?id=2741637 16. Gupta, M., Zhao, P., & Han, J. (2012) Evaluating event credibility on twitter. In Proceedings of the 2012 Society for Industrial and Applied Mathematics International

Conference on Data Mining (pp. 153–164). http://epubs. siam.org/doi/abs/10.1137/1.9781611972825.14 17. Ciampaglia, G. L., Shiralkar, P., Rocha, L. M., Bollen, J., Menczer, F., & Flammini, A. (2015) Computational fact checking from knowledge networks. PloS one, 10(6), e0128193. http://journals.plos.org/plosone/

article?id=10.1371/journal.pone.0128193 18. Friggeri, A., Adamic, L. A., Eckles, D., & Cheng, J. (2014) Rumor Cascades. In Proceedings of the International Conference on Weblogs and Social Media. https://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/ paper/viewFile/8122/8110

 Goel, S., Anderson, A., Hofman, J., & Watts, D.
J. (2015) The structural virality of online diffusion. Management Science, 62(1), 180–196 https://pubsonline. informs.org/doi/abs/10.1287/mnsc.2015.2158
Itti, L., & Baldi, P. (2009) Bayesian surprise attracts human attention. Vision research, 49(10): 1295–1306. http://www.sciencedirect.com/science/article/pii/ S0042698908004380

21. Aral, S., & Van Alstyne, M. (2011) The Diversity-Bandwidth Trade-off. American Journal of Sociology, 117(1): 90-171. http://www.journals.uchicago.edu/doi/ abs/10.1086/661238

22. Berger, J., & Milkman, K. L. (2012) What makes online content viral? Journal of marketing research, 49(2), 192–205. http://journals.ama.org/doi/abs/10.1509/ jmr.10.0353

23. Davis, C. A., Varol, O., Ferrara, E., Flammini, A., & Menczer, F. (2016) BotOrNot: A system to evaluate social bots. In Proceedings of the 25th International Conference on the World Wide Web (pp. 273–274). https://dl.acm.org/citation.cfm?id=2889302 24. For example, https://www.intelligence.senate.gov/ sites/default/files/documents/os-cwatts-033017.pdf MIT INITIATIVE ON THE DIGITAL ECONOMY RESEARCH BRIEF

THE SPREAD OF TRUE AND FALSE NEWS ONLINE

By Soroush Vosoughi, Deb Roy, and Sinan Aral

RESEARCH

Read the full research in Science here.

VIDEO

Watch Sinan Aral, MIT, discuss the research here.

ABOUT THE AUTHORS

Soroush Vosoughi is a postdoctoral associate at the MIT Media Lab and a fellow at the Berkman Klein Center at Harvard University. He received his PhD, MSc, and BSc from MIT.

Deb Roy is an Associate Professor at MIT where he directs the Laboratory for Social Machines (LSM) based at the Media Lab. His lab explores new methods in media analytics (natural language processing, social network analysis, speech, image, and video analysis) and media design (information visualization, games, communication apps) with applications in children's learning and social listening.

Sinan Aral is the David Austin Professor of Management at MIT, where he is a Professor of IT & Marketing, and Professor in the Institute for Data, Systems and Society where he co-leads the MIT Initiative on the Digital Economy.

MIT INITIATIVE ON THE DIGITAL ECONOMY

The MIT IDE is solely focused on the digital economy. We conduct groundbreaking research, convene the brightest minds, promote dialogue, expand knowledge and awareness, and implement solutions that provide critical, actionable insight for people, businesses, and government. We are solving the most pressing issues of the second machine age, such as defining the future of work in this time of unprecendented disruptive digital transformation.

SUPPORT THE MIT IDE

The generous support of individuals, foundations, and corporations are critical to the success of the IDE. Their contributions fuel cutting-edge research by MIT faculty and graduate students, and enables new faculty hiring, curriculum development, events, and fellowships. Contact Christie Ko (cko@mit.edu) to learn how you or your organization can support the IDE.

TO LEARN MORE ABOUT THE IDE, INCLUDING UPCOMING EVENTS, VISIT IDE.MIT.EDU



