

CS 6474/CS 4803

Social Computing: Misinformation and Disinformation

Munmun De Choudhury

munmund@gatech.edu

Week 12 | March 25, 2024


The birth of the Obama 'birther' conspiracy

Anthony Zurcher
North America reporter
[@awzurcher](#)

🕒 16 September 2016



US election 2016



Defining “fake news”



Professor in Political Science and Computer and Information Science

DAVID LAZER

[HOME](#) [BIO](#) [RESEARCH AREAS](#) [MY NETWORK](#)



[Northeastern Home](#) [CCIS Home](#) [PollSci Home](#)
Northeastern University

BIOGRAPHICAL NOTES JRSE

Welcome! I am Professor at the Department of Political Science and College of Computer and Information Science at Northeastern University. Click here for **biographical information** and an **overview of my publications, of teaching and academic activities, and some media appearances.**

Yours,

David Lazer



RESEARCH FOLLOWUP

The objective of this website is to provide entrée into **my body of research**. Most of my work is based on the idea that how people and organizations are connected together is critical to understanding the **functioning, success and failure of actors and systems**. My teaching, research, and institution building have all centered on that theme. I've taken that essential idea and, with a variety of collaborators, examined a **wide array of domains**.



LABORATORY

My research covers everything from **very micro** (social influence processes within groups), to the **very macro** (the development of global-wide regulatory regimes).



Sources of misinformation/disinformation

- Rumors and fiction
- Governments and politicians
- Vested interests
- The media

The societal costs of misinformation

Quantifying the Impact of Misinformation and Vaccine- Skeptical Content on Facebook

Main idea

- Known association between misinformation on social media and negative societal events, such as the spread of COVID-19 misinformation.
- Questioning the causal relationship between misinformation and vaccination hesitancy.

Study Design

- Framework combining lab experiments, crowdsourcing, and machine learning.
- Estimation of the causal effect of vaccine-related misinformation on vaccination intentions.
- Analysis of 13,206 vaccine-related URLs shared on Facebook.

Impact of Vaccine Misinformation

- Misinformation flagged by fact-checkers was 50X less impactful than vaccine-skeptical content.
- Rare deaths following vaccination highlighted in unflagged stories had significant exposure.
- Unflagged content suggesting vaccines were harmful had a stronger negative impact on vaccination intentions/were more viewed on FB than flagged misinformation.

Examining the Alternative Media Ecosystem through the Production of Alternative Narratives of Mass Shooting Events on Twitter

Kate Starbird

University of Washington, HCDE
kstarbi@uw.edu

Abstract

This research explores the alternative media ecosystem through a Twitter lens. Over a ten-month period, we collected tweets related to alternative narratives—e.g. conspiracy theories—of mass shooting events. We utilized tweeted URLs to generate a domain network, connecting domains shared by the same user, then conducted qualitative analysis to understand the nature of different domains and how they connect to each other. Our findings demonstrate how alternative news sites propagate and shape alternative narratives, while mainstream media deny them. We explain how political leanings of alternative news sites do not align well with a U.S. left-right spectrum, but instead feature an anti-globalist (vs. globalist) orientation where U.S. Alt-Right sites look similar to U.S. Alt-Left sites. Our findings describe a subsection of the emerging alternative media ecosystem and provide insight in how websites that promote conspiracy theories and pseudo-science may function to conduct underlying political agendas.

Introduction

In the aftermath of major political disruptions in 2016—in Britain with the Brexit vote and in the United States with the election of Donald Trump to the presidency—there has been widespread attention to and theorizing about the problem of “fake news”. But this term is both amorphous and contested. One perspective locates the problem within the emerging ecosystem of alternative media, where the term has been applied to refer to “clickbait” content that uses tabloid-style headlines to attract viewers for financial reasons (Silverman & Alexander 2016) and to describe political propaganda intentionally planted and propagated through online spaces (Timberg 2016). Challenging these definitions, alternative media outlets have appropriated the term to attack “mainstream” media for its perceived economic and political biases and for hosting inaccurate or under-sourced content (e.g. Rappoport 2016). Beneath this

rhetoric, we are seeing traditional new providers and emergent alternative media battle not only for economic viability, but over accepted methods of how information is shared and consumed, and, more profoundly, for how narratives around that information are shaped and by whom.


This research seeks to provide a systematic lens for exploring the production of a certain type of “fake news”—*alternative narratives* of man-made crisis events. For three years, our research group has examined online rumoring during crises. Over that time, we noted the presence of very similar rumors across many man-made crisis events—including the 2013 Boston Marathon Bombings, the downing of Malaysia Airlines flight MH17, and several mass shooting events including those at Umpqua Community College in Oregon (October, 2015). For each event, rumors claimed the event had been perpetrated by someone other than the official suspects—that it was instead either a staged event performed by “crisis actors” or a “false flag” orchestrated by someone else. Both explanations claimed that a powerful individual or group was pulling the strings for political reasons. Interestingly, though the arguments and evidence used to support these alternative narratives were somewhat consistent across events, the motives cited were often very different—e.g. from the U.S. government trying to support gun control to coordinated global actors staging violence to motivate military intervention.

For this paper, we utilize this type of conspiracy theory or *alternative narrative* rumor as an entry point for understanding the ecosystem of alternative media. We examine the production of these narratives through Twitter and across the external websites that Twitter users reference as they engage in these narratives. We propose and demonstrate that this lens—Twitter data from mass shooting events and our method for utilizing that data to reveal and explore connections across web domains—provides a systematic approach for shedding light on the emerging phenomena of alternative media and “fake news”.

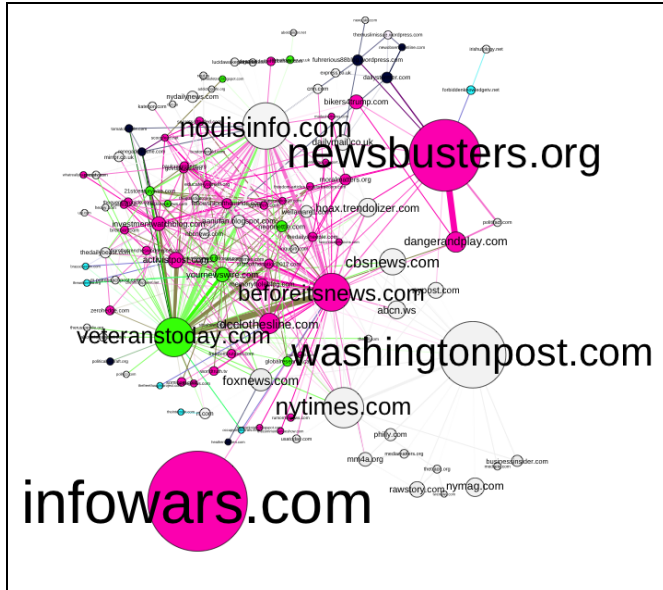
Our contributions include an increased understanding of the underlying nature of this subsection of alternative me-

Main Findings

Purple = mainstream media; Aqua = alternative media; Red = government controlled media



Ink = U.S. Alt-Right; Aqua = U.S. Alt-Left; Green = Intl. Anti-Globalist; Black = White Nationalist/Anti-Semitic; White = other



Leaning	Description
U.S. Alt Right	U.S. focused, anti-mainstream media, pro-Christian, anti-LGBT, anti-feminist, anti-globalist, climate change denying
U.S. Alt Left	U.S. focused, anti-mainstream media, anti-corporatist, critical of police, pro-prison reform, pro-BlackLivesMatter
International Anti-Globalist	Internationally focused, anti-globalist or anti-New World Order/Cabal, anti-corporatist, conspiracy-focused
White Nationalist and/or Anti-Semitic	primarily white-nationalist or anti-Semitic positions
Muslim Defense	primarily challenges mainstream narratives of terrorist attacks by Muslims
Russian Propaganda	primarily supports Russian interests, anti-globalist

Starbird found that alternative media sites may give the false perception of being exposed to a diverse information diet. Most of them, seemingly different on the surface, harp on the same political (e.g., anti-globalist) agenda.

Can social media platforms ensure information diversity while curbing the negative impact of fake news?

Renewed interest

POLITICAL SCIENCE

Fake news on Twitter during the 2016 U.S. presidential election

Nir Grinberg^{1,2*}, Kenneth Joseph^{3*}, Lisa Friedland^{1*},
Briony Swire-Thompson^{1,2}, David Lazer^{1,2†}

The spread of fake news on social media became a public concern in the United States after the 2016 presidential election. We examined exposure to and sharing of fake news by registered voters on Twitter and found that engagement with fake news sources was extremely concentrated. Only 1% of individuals accounted for 80% of fake news source exposures, and 0.1% accounted for nearly 80% of fake news sources shared. Individuals most likely to engage with fake news sources were conservative leaning, older, and highly engaged with political news. A cluster of fake news sources shared overlapping audiences on the extreme right, but for people across the political spectrum, most political news exposure still came from mainstream media outlets.

In 1925, *Harper's Magazine* published an article titled "Fake news and the public," decrying the ways in which emerging technologies had made it increasingly difficult to separate rumor from fact (1). Nearly a century later, fake news has again found its way

social media have described its spread within platforms (5, 6) and highlighted the disproportionate role played by automated accounts (7), but they have been unable to make inferences about the experiences of ordinary citizens.

Outside of social media, fake news has been

We distinguished among three classes of fake news sources to allow comparisons of different operational definitions of fake news. The three classes correspond to differences in methods of generating lists of sources as well as perceived differences in the sites' likelihoods of publishing misinformation. We labeled as "black" a set of websites taken from preexisting lists of fake news sources constructed by fact-checkers, journalists, and academics (8, 9) who identified sites that published almost exclusively fabricated stories [see supplementary materials (SM) section S.5 for details]. To measure fake news more comprehensively, we labeled additional websites as "red" or "orange" via a manual annotation process of sites identified by Snopes.com as sources of questionable claims. Sites with a red label (e.g., Infowars.com) spread falsehoods that clearly reflected a flawed editorial process, and sites with an orange label represented cases where annotators were less certain that the falsehoods stemmed from a systematically flawed process. There were 171 black, 64 red, and 65 orange fake news sources appearing at least once in our data.

Voters on Twitter

To focus on the experiences of real people on Twitter, we linked a sample of U.S. voter reg-



Measuring the news and its impact on democracy

Duncan J. Watts^{a,b,c,1} , David M. Rothschild^d, and Markus Mobius^e

^aDepartment of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104; ^bThe Annenberg School of Communication, University of Pennsylvania, Philadelphia, PA 19104; ^cOperations, Information, and Decisions Department, University of Pennsylvania, Philadelphia, PA 19104; ^dMicrosoft Research, New York, NY 10012; and ^eMicrosoft Research, Cambridge, MA 02142

Edited by Dietram A. Scheufele, University of Wisconsin–Madison, Madison, WI, and accepted by Editorial Board Member Susan T. Fiske February 21, 2021 (received for review November 8, 2019)

Since the 2016 US presidential election, the deliberate spread of misinformation online, and on social media in particular, has generated extraordinary concern, in large part because of its potential effects on public opinion, political polarization, and ultimately democratic decision making. Recently, however, a handful of papers have argued that both the prevalence and consumption of “fake news” per se is extremely low compared with other types of news and news-relevant content. Although neither prevalence nor consumption is a direct measure of influence, this work suggests that proper understanding of misinformation and its effects requires a much broader view of the problem, encompassing biased and misleading—but not necessarily factually incorrect—information that is routinely produced or amplified by mainstream news organizations. In this paper, we propose an ambitious collective research agenda to measure the origins, nature, and prevalence of misinformation, broadly construed, as well as its impact on democracy. We also sketch out some illustrative examples of completed, ongoing, or planned research projects that contribute to this agenda.

misinformation | media | democracy

It is hard to overstate the breadth and intensity of interest directed over the past 2 y at the issue of false or misleading information (also known as “fake news”) circulating on the web in general and on social media platforms such as Facebook and Twitter in particular (1–13). According to Google Scholar, since

pro-Clinton articles.” In turn, they estimated that “if one fake news article were about as persuasive as one TV campaign ad, the fake news in our database would have changed vote shares by an amount on the order of hundredths of a percentage point,” roughly two orders of magnitude less than needed to influence the election outcome. Subsequent studies have found similarly low prevalence levels for fake news relative to mainstream news on Twitter (10) and Facebook (11). Finally, our own survey of the media consumption landscape, based on a nationally representative sample of TV, desktop, and mobile media consumption (18), found three main results that undercut the conventional wisdom regarding fake news and also the dominance of online sources of news in general:

- 1) News consumption is a relatively small fraction of overall media consumption. Of the more than 7.5 h per day that Americans spend, on average, watching television or consuming content on their desktop computers or mobile devices, only about 14% is dedicated to news (“news” was defined as appearing on one of more than 400 news-relevant programs [e.g., *CBS Evening News*] and more than 800 websites [e.g., <http://www.nytimes.com/>], while “consumption” was measured in terms of minutes per person per day watching television or browsing online; see ref. 18 for details).
- 2) Online news consumption is a small fraction of overall news consumption, which is dominated by TV by a factor of five to one. Even 18 to 24 y olds consume almost twice as much TV

The spread of true and false news online

Bots and misinformation



Social bots distort the 2016 U.S. Presidential election online discussion

by Alessandro Bessi and Emilio Ferrara

Abstract

Social media have been extensively praised for increasing democratic discussion on social issues related to policy and politics. However, what happens when this powerful communication tools are exploited to manipulate online discussion, to change the public perception of political entities, or even to try affecting the outcome of political elections? In this study we investigated how the presence of social media bots, algorithmically driven entities that on the surface appear as legitimate users, affect political discussion around the 2016 U.S. Presidential election. By leveraging state-of-the-art social bot detection algorithms, we uncovered a large fraction of user population that may not be human, accounting for a significant portion of generated content (about one-fifth of the entire conversation). We inferred political partisanship from hashtag adoption, for both humans and bots, and studied spatio-temporal communication, political support dynamics, and influence mechanisms by discovering the level of network embeddedness of the bots. Our findings suggest that the presence of social media bots can indeed negatively affect democratic political discussion rather than improving it, which in turn can potentially alter public opinion and endanger the integrity of the Presidential election.

Contents

[Introduction](#)
[Methodology](#)
[Data analysis](#)
[Conclusions](#)

Introduction

Various computational social science studies demonstrated that social media have been extensively used to foster democratic conversation about social and political issues: From the Arab Spring (González-Bailón, *et al.*, 2011; Howard, *et al.*, 2011), to Occupy Wall Street (Conover, *et al.*, 2013a; Conover, *et al.*, 2013b) and many other civil protests (Varol, *et al.*, 2014; González-Bailón, *et al.*, 2013) (Bastos, *et al.*, 2014), Twitter and other social media seemed to play an instrumental role to involve the public in policy and political conversations, by collectively framing the narratives related to particular social issues, and coordinating online and off-line activities. The use of digital media to discuss politics during election times has also been the subject of various studies, covering the last four U.S. Presidential elections (Adamic and Glance, 2005; Diakopoulos and Shamma, 2010; Bekafigo and McBride, 2013; Carlisle and Patton, 2013; DiGrazia, *et al.*, 2013; Wang, *et al.*, 2016), and other countries like Australia (Gibson and McAllister, 2006; Bruns and Burgess, 2011; Burgess and Bruns, 2012), and Norway (Enli and Skogerbø, 2013). Findings that focused on the positive effects of social media such as incrementing voting turnout (Bond, *et al.*, 2012) or exposure to diverse political views (Bakshy, *et al.*, 2015) contributed to the general praise of these platforms as a tool to foster democracy and civil political engagement (Shirky, 2011; Loader and Mercea, 2011; Effing, *et al.*, 2011; Tufekci and Wilson, 2012; Tufekci, 2014; Yang, *et al.*, 2016).

However, as early as 2006, Philip Howard raised concerns regarding the possibility of manipulating public opinion and spreading political misinformation through social media (Howard, 2006). These issues have been later proved true by several studies (Ratkiewicz, *et al.*, 2011a; Ratkiewicz, *et al.*, 2011b) (Metaxas and Mustafaraj, 2012) (El-Khalili, 2013; Ferrara, 2015; Woolley and Howard, 2016; Shorey and Howard, 2016). Of particular concern is the fact social media have been demonstrated effective in influencing individuals (Aral and Walker, 2010). One way to perform such type of manipulation is by using social bots, algorithmically controlled accounts that emulate the activity of human users but operate at much higher pace (*e.g.*, automatically producing content or engaging in social interactions), while successfully keeping their artificial identity undisclosed (Hwang, *et al.*, 2012; Messias, *et al.*, 2013; Ferrara, *et al.*, 2016).

Evidence of the adoption of social media bots to attempt manipulating political communication dates back half a decade: during the 2010 U.S. midterm elections, social bots were employed to support some candidates and smear others, by injecting thousands of tweets pointing to Web sites with fake news (Ratkiewicz, *et al.*, 2011a). The research community reported another similar case around the time of the 2010 Massachusetts special election (Metaxas and Mustafaraj, 2012). Campaigns of this type are sometimes referred to as *astroturf* or *Twitter bombs*. Unfortunately, most of the times, it has proven impossible to determine who's behind these types of operations (Kollanyi, *et al.*, 2016; Ferrara, *et al.*, 2016). Governments, organizations, and other entities with sufficient resources, can obtain the technological capabilities to deploy thousands of social bots and use them to their advantage, either to support or to attack particular political figures or candidates. Indeed, it has become increasingly simpler to deploy social bots, so that, in some cases, no coding skills are required to setup accounts that perform simple automated activities: tech blogs often post tutorials and ready-to-go tools for this purposes [1], [2], [3]. Various source codes for sophisticated social media bots can be found online as well, ready to be customized and optimized by the more technical savvy users (Kollanyi, 2016). We inspected several of these readily available bots and this is a (non-comprehensive) list of the capabilities that they provide: Search Twitter for phrases/hashtags/keywords and automatically retweet them; Automatically reply to tweets that meet a certain criteria; automatically follow any users that tweet something with a specific phrase/hashtag/keyword; Automatically follow back any users that have followed the bot; Automatically follow any users that follow a specified user; Automatically add users tweeting about something to public lists; Search Google (and other engines) for articles/news according to specific criteria and post them, or link them in automatic replies to other users; Automatically aggregating public sentiment on certain topics of discussion; Buffer and post tweets automatically. Most of these bots can run in cloud services or infrastructures like Amazon Web Services (AWS) or Heroku, making it more difficult to block them. Finally, a very recent trend is that of providing Bot-As-A-Service (Baas): companies like RoboLike (<https://robo-like.com/>) provide "Easy-to-use Instagram/Twitter auto bots" performing certain automatic activities for a monthly price. Advanced conversational bots powered by more sophisticated Artificial Intelligences are provided by companies like ChatBots.io that allow anyone to "Add a bot to services like Twitter, Hubot, Facebook, Skype, Twilio, and more" (<https://developer.pandorabots.com/>).



RESEARCH ARTICLE



Bots are less central than verified accounts during contentious political events

 Sandra González-Bailón and  Manlio De Domenico

[+ See all authors and affiliations](#)

PNAS March 16, 2021 118 (11) e2013443118; <https://doi.org/10.1073/pnas.2013443118>

Edited by Nicole Ellison, University of Michigan, Ann Arbor, MI, and accepted by Editorial Board Member Margaret Levi January 30, 2021 (received for review June 27, 2020)

Article

Figures & SI

Info & Metrics

 PDF

Significance

Online networks carry benefits and risks with high-stakes consequences during contentious political events: They can be tools for organization and awareness, or tools for disinformation and conflict. We combine social media and web-tracking data to measure

Should social media platforms censor the “free speech” of harmful bots?

TECH

Zuckerberg tells Congress Facebook is not a media company: 'I consider us to be a technology company'

PUBLISHED WED, APR 11 2018·10:27 AM EDT | UPDATED WED, APR 11 2018·10:51 AM EDT



Michelle Castillo
@MISHCASTILLO

SHARE    

MEDIA 09/27/2017 08:53 pm ET

830



Mark Zuckerberg: 'I Regret' Rejecting Idea That Facebook Fake News Altered Election

He admitted this after Donald Trump claimed that Facebook was "always anti-Trump."



By Carla Herreria



Facebook CEO [Mark Zuckerberg](#) admitted on Wednesday that he was wrong to dismiss the idea that fake news shared on the giant social network affected last year's presidential election.

Zuckerberg's statement came in response to a tweeted attack from President [Donald Trump](#) hours earlier. Trump claimed that Facebook was "[always anti-Trump](#)" and accused it of colluding with news outlets that the president has deemed to be "fake news."

Class Exercise

Assume you work at Meta. Design a study to measure if disinformation campaigns from social bots or popular public figures may have impacted the outcomes of the 2016 Presidential elections.



Facebook targets 'false news' amid growing pressure from advertisers

By Marianna Spring

Specialist disinformation and social media reporter

🕒 30 June 2020



What's missing?

Get the whole story
not just a headline.

Images can be faked.

Check what other people say.

The graphic features a central smartphone screen displaying a social media post with a red 'scam alert' banner and a 'covid-19' section with three icons. The background is blue with white text and icons, including a Facebook logo and several blue virus-like icons. A red exclamation mark icon is positioned above the smartphone screen.

FACEBOOK

Facebook's new media literacy campaign will ask users questions about what they see online

Working to Stop Misinformation and False News

We know people want to see accurate information on Facebook – and so do we.

False news is harmful to our community, it makes the world less informed, and it erodes trust. It's not a new phenomenon, and all of us — tech companies, media companies, newsrooms, teachers — have a responsibility to do our part in addressing it. At Facebook, we're working to fight the spread of false news in three key areas:

- disrupting economic incentives because most false news is financially motivated;
- building new products to curb the spread of false news; and
- helping people make more informed decisions when they encounter false news.

Asia

Oct 24th 2020 edition >

Anti-social network

In Myanmar, Facebook struggles with a deluge of disinformation

Weeks before an election, Burmese social media are awash with fake news and vitriol



Social media sites are starting to label false news or take down posts. Is this a good idea or is it enough? What else can be done to stop the spread of false news?

Article Menu

[Close](#) ^[Download PDF](#) [Free EPUB View](#)

Accessing resources off campus can be a challenge. Lean Library can solve it

[Full Article](#)

Content List

- Abstract
- The Societal Cost of Misinformation
- Origins of Misinformation
- From Individual Cognition to

Misinformation and Its Correction: Continued Influence and Successful Debiasing

Stephan Lewandowsky, Ullrich K. H. Ecker, Colleen M. Seifert, more...

[Show all authors](#) ▾

First Published September 17, 2012 | Research Article | [Find in PubMed](#) |
<https://doi.org/10.1177/1529100612451018>

[Article information](#) ▾

Abstract

The widespread prevalence and persistence of misinformation in contemporary societies, such as the false belief that there is a link between childhood vaccinations and autism, is a matter of public concern. For example, the myths surrounding vaccinations, which prompted some parents to withhold immunization from their children, have led to a marked increase in vaccine-preventable disease, as well as unnecessary public expenditure on research and public-information campaigns aimed at rectifying the situation.

We first examine the mechanisms by which such misinformation is disseminated in society, both inadvertently and purposely. Misinformation can originate from rumors but also from works of fiction, governments and politicians, and vested interests. Moreover, changes in the media landscape, including the arrival of the Internet, have fundamentally influenced the ways in which information is communicated and misinformation is spread.

Social media sites are starting to label false news or take down posts. Is this a good idea or is it enough? What else can be done to stop the spread of false news?