CS 6474/CS 4803 Social Computing: Polarization and Selective Exposure

Munmun De Choudhury

munmund@gatech.edu

Week 12 | March 24, 2025

A background on polarization

Dynamic Debates: An Analysis of Group Polarization Over Time on Twitter

Sarita Yardi[1] and Danah Boyd[2]

Abstract

The principle of homophily says that people associate with other groups of people who are mostly like themselves. Many online communities are structured around groups of socially similar individuals. On Twitter, however, people are exposed to multiple, diverse points of view through the public timeline. The authors captured 30,000 tweets about the shooting of George Tiller, a late-term abortion doctor, and the subsequent conversations among pro-life and pro-choice advocates. They found that replies between like-minded individuals strengthen group identity, whereas replies between different-minded individuals reinforce in-group and out-group affiliation. Their results show that people are exposed to broader viewpoints than they were before but are limited in their ability to engage in meaningful discussion. They conclude with implications for different kinds of social participation on Twitter more generally.

Echo chambers online?: Politically motivated selective exposure among Internet news users[1]

R. Kelly Garrett   Author Notes
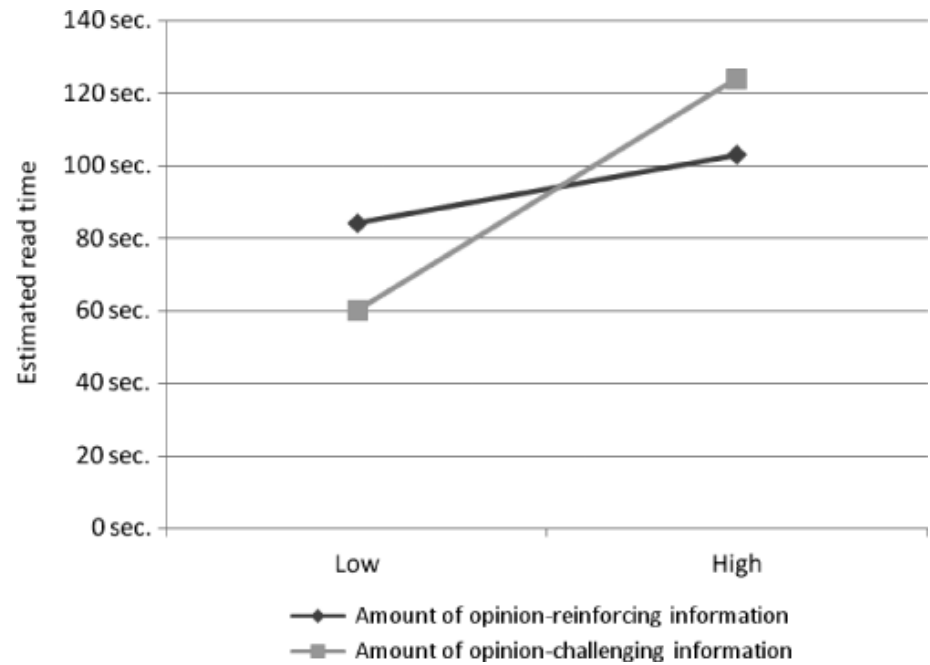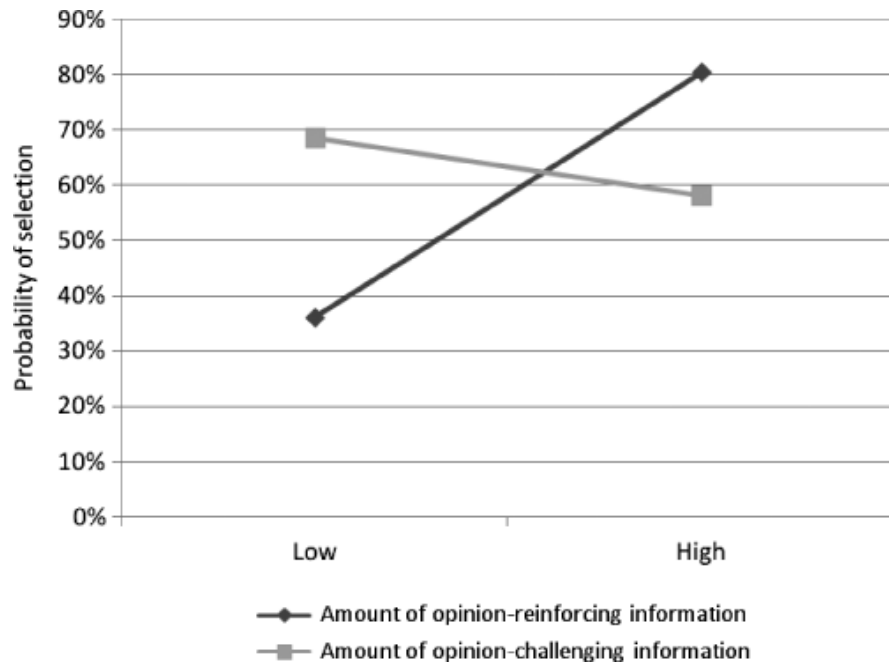
PDF   Split View   Cite   Permissions   Share ▾

Abstract

A review of research suggests that the desire for opinion reinforcement may play a more important role in shaping individuals' exposure to online political information than an aversion to opinion challenge. The article tests this idea using data collected via a web-administered behavior-tracking study with subjects recruited from the readership of 2 partisan online news sites (*N = 727*). The results demonstrate that opinion-reinforcing information promotes news story exposure while opinion-challenging information makes exposure only marginally less likely. The influence of both factors is modest, but opinion-reinforcing information is a more important predictor. Having decided to view a news story, evidence of an aversion to opinion challenges disappears: There is no evidence that individuals abandon news stories that contain information with which they disagree. Implications and directions for future research are discussed.

Summary

Left chart (y-axis: Probability of selection, 0%–90%; x-axis: Low, High)
- Amount of opinion-reinforcing information: Low ≈ 36%, High ≈ 80%
- Amount of opinion-challenging information: Low ≈ 68%, High ≈ 58%

Right chart (y-axis: Estimated read time, 0 sec.–140 sec.; x-axis: Low, High)
- Amount of opinion-reinforcing information: Low ≈ 84 sec., High ≈ 104 sec.
- Amount of opinion-challenging information: Low ≈ 60 sec., High ≈ 123 sec.

Legend (both charts):
- Amount of opinion-reinforcing information
- Amount of opinion-challenging information

Bubble Trouble

Is Web personalization turning us into solipsistic twits?

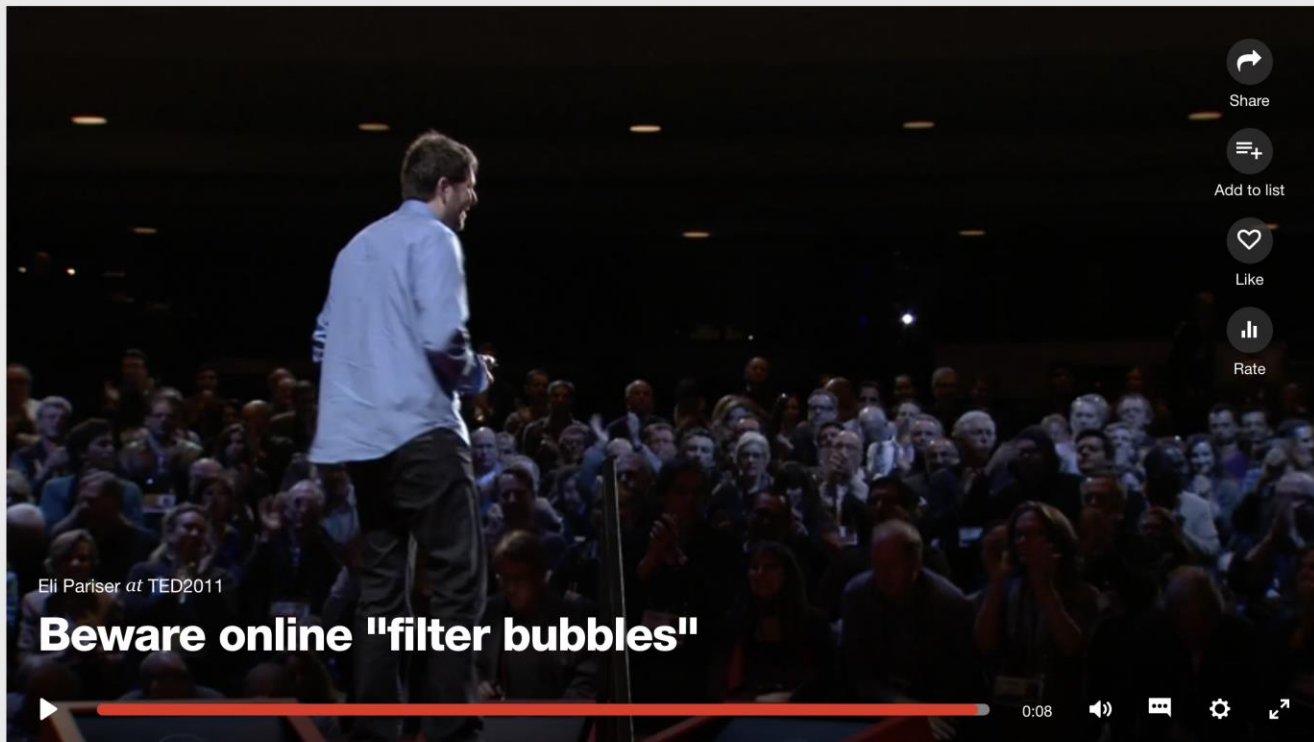51 3 0

By *Jacob Weisberg*

Eli Pariser

The first conversation I ever had about the Internet was in 1993 with Robert Wright, who was then a colleague at the *New Republic*. This "Net" thing was going to be a big deal, I remember Bob telling me, but it could create a few problems. One was that it was going to empower crazies, since geographically diffuse nut jobs of all sorts would be able to find each other online. Another was that it could hurt democratic culture by encouraging narrow-minded folk to burrow deeper into their holes. Wright spelled out those concerns in an **article** that stands as a model of prescience and a delightful time-capsule. ("People who 'post' on the Net's many different bulletin boards—its 'newsgroups'—know that their words can be seen from just about any chunk of inhabited turf on this planet.")

https://www.ted.com/talks/eli_pariser_beware_online_filter_bubbles

How did we get here? What do you think is the reasoning behind online platforms promoting polarization or selective exposure?

"I always assumed that I wasn't really that close to [her]": Reasoning about invisible algorithms in the news feed

Motahhare Eslami, Aimee Rickman†, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong
Karrie Karahalios, Kevin Hamilton, Christian Sandvig‡
University of Illinois at Urbana-Champaign, †California State University, Fresno, ‡University of Michigan
{eslamim2, kvaccaro, aleyase2, avuong3, kkarahal, kham}@illinois.edu
†arickman@csufresno.edu, ‡csandvig@umich.edu

ABSTRACT

Our daily digital life is full of algorithmically selected content such as social media feeds, recommendations and personalized search results. These algorithms have great power to shape users' experiences, yet users are often unaware of their presence. Whether it is useful to give users insight into these algorithms' existence or functionality and how such insight might affect their experience are open questions. To address them, we conducted a user study with 40 Facebook users to examine their perceptions of the Facebook News Feed curation algorithm. Surprisingly, more than half of the participants (62.5%) were not aware of the News Feed curation algorithm's existence at all. Initial reactions for these previously unaware participants were surprise and anger. We developed a system, *FeedVis*, to reveal the difference between the algorithmically curated and an unadulterated News Feed to users, and used it to study how users perceive this difference. Participants were most upset when close friends and family were not shown in their feeds. We also found participants often attributed missing stories to their friends' decisions to exclude them rather than to Facebook News Feed algorithm. By the end of the study, however, participants were mostly satisfied with the content on their feeds. Following up with participants two to six months after the study, we found that for most, satisfaction levels remained similar before and after becoming aware of the algorithm's presence, however, algorithmic awareness led to more active engagement with Facebook and bolstered overall feelings of control on the site.
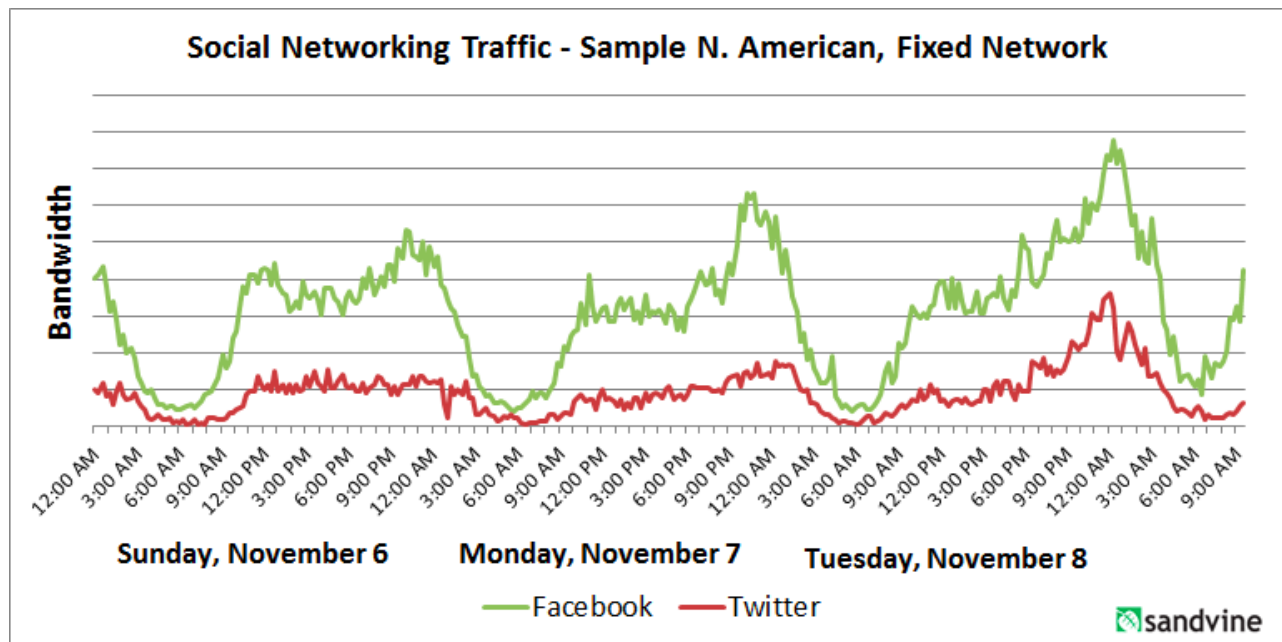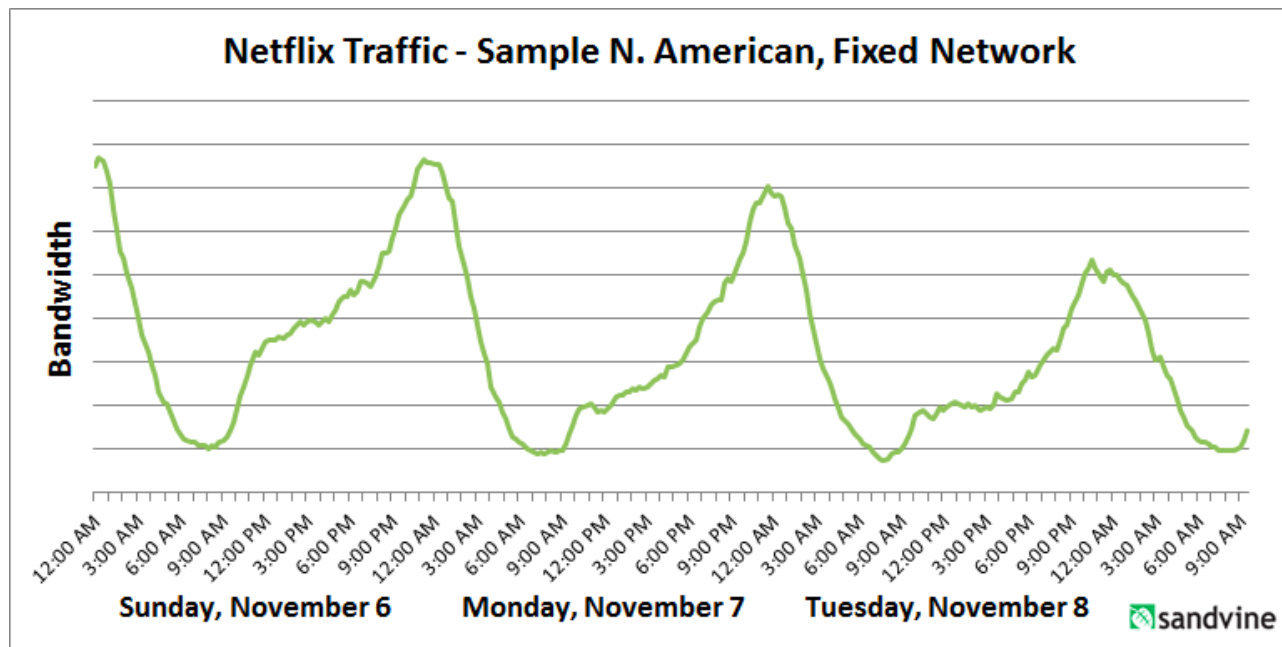
in September of 2014, was viewed by on average 864 million daily active users [12]. This list of updating stories that appears front and center on Facebook home pages displays an algorithmically curated or filtered list of stories selected from a pool of all stories created by one's network of friends.

The increasing prevalence of opaque and invisible algorithms coupled with their power raises questions about how knowledgeable users are and should be about the existence and operation of these algorithms. Whether their understanding is correct or not, users' perceived knowledge about an algorithm can affect their behavior. For instance, believing that posts with commercial keywords were ranked higher by the Facebook News Feed algorithm, some teenagers added product names to their posts in an attempt to manipulate the algorithm and increase their posts' visibility [41]. Other users tended to block new mothers in their feed based on a false assumption that such women posted too many baby pictures, when in fact the prevalence of such images was determined by their popularity among users [27].

However, with no way to know if their knowledge of these invisible algorithms is correct, users cannot be sure of the results of their actions. Algorithmic interfaces in Internet applications rarely include a clear enough feedback mechanism for users to understand the effects of their own actions on the system. Without such feedback, it can be difficult to assess the influence of either algorithm knowledge or ignorance.

Influence in the political sphere: 62% of adults in the US use social media to consume news, and 18% of adults are frequent consumers – Pew Internet

Netflix Traffic - Sample N. American, Fixed Network

Sunday, November 6 — Monday, November 7 — Tuesday, November 8

Social Networking Traffic - Sample N. American, Fixed Network

Sunday, November 6 — Monday, November 7 — Tuesday, November 8

Facebook — Twitter

http://www.internetphenomena.com/2016/11/us-election-2016-tv-trumps-the-internet/

Facebook, in Cross Hairs After Election, Is Said to Question Its Influence

By MIKE ISAAC NOV. 12, 2016

Exposure to ideologically diverse news and opinion on Facebook

EYTAN BAKSHY , SOLOMON MESSING, AND LADA A. ADAMIC   Authors Info & Affiliations

Not getting all sides of the news?

People are increasingly turning away from mass media to social media as a way of learning news and civic information. Bakshy *et al.* examined the news that millions of Facebook users' peers shared, what information these users were presented with, and what they ultimately consumed (see the Perspective by Lazer). Friends shared substantially less cross-cutting news from sources aligned with an opposing ideology. People encountered roughly 15% less cross-cutting content in news feeds due to algorithmic ranking and clicked through to 70% less of this cross-cutting content. Within the domain of political news encountered in social media, selective exposure appears to drive attention.

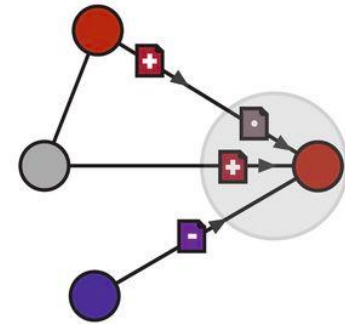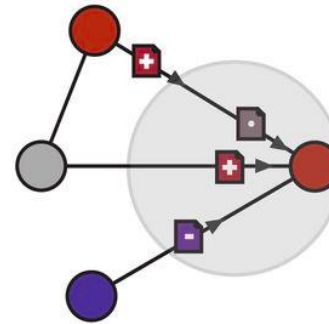*Science*, this issue p. 1130; see also p. 1090

A

Stage in media exposure process

Potential from network　　**Exposed**　　**Selected**
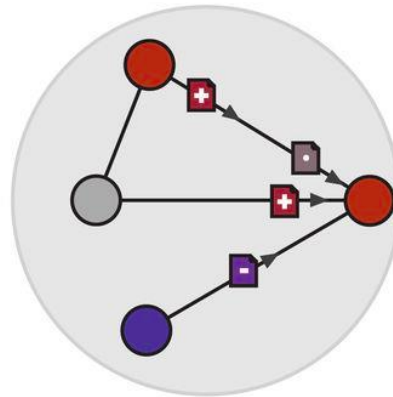
Proportion of content that is cross-cutting

1/3　　　　1/2　　　　0/1

B

Percent cross–cutting content

Viewer affiliation
— Conservative
— Liberal

Random　Potential from network　Exposed　Selected
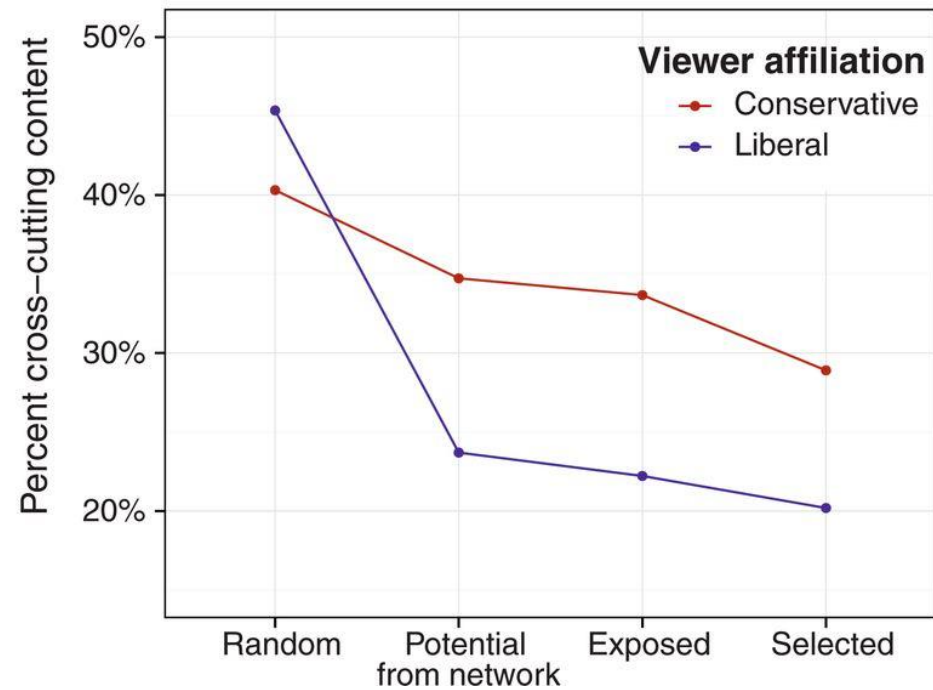
"*[Facebook] defended the News Feed's progress arguing that the filter bubble isn't an issue for [the platform]. [Company] suggested the real problem is that people by nature engage with content they like and find agreeable, and dismiss things they don't agree with online as they would in real life.*" – Techcrunch

To what extent is it fair to put the responsibility on users?

Like-minded sources on Facebook are prevalent but not polarizing

Exposure to opposing views on social media can increase political polarization

Reading opinion-reinforcing content can have widespread impact on our perceptions what is true and what is false. Could this impact our perceptions of what could be misinformation?
– Connection to prior class discussions

Homophily and polarization in the age of misinformation

Alessandro Bessi[1,2], Fabio Petroni[3], Michela Del Vicario[2], Fabiana Zollo[2], Aris Anagnostopoulos[3], Antonio Scala[4], Guido Caldarelli[2], and Walter Quattrociocchi[1,a]

[1] IUSS, Pavia, Italy
[2] IMT Institute for Advanced Studies, Lucca, Italy
[3] Sapienza University, Roma, Italy
[4] ISC-CNR, Roma, Italy

Abstract. The World Economic Forum listed massive digital misinformation as one of the main threats for our society. The spreading of unsubstantiated rumors may have serious consequences on public opinion such as in the case of rumors about Ebola causing disruption to health-care workers. In this work we target Facebook to characterize information consumption patterns of 1.2 M Italian users with respect to verified (science news) and unverified (conspiracy news) contents. Through a thorough quantitative analysis we provide important insights about the anatomy of the system across which misinformation might spread. In particular, we show that users' engagement on verified (or unverified) content correlates with the number of friends having similar consumption patterns (*homophily*). Finally, we measure how this social system responded to the injection of 4,709 false information. We find that the frequent (and selective) exposure to specific kind of

What makes polarization in social computing systems challenging to address?

Class Exercise

If you were to re-design a tool that works on social media to reduce polarization, what would that tool look like? What would it do? What data would it use? How would you evaluate if the tool is working?

"Bridging" at the network boundary?

A Measure of Polarization on Social Media Networks Based on Community Boundaries

Pedro H. Calais Guerra, Wagner Meira Jr.
Department of Computer Science
Universidade Federal de Minas Gerais (UFMG)
Belo Horizonte, MG, Brazil

Claire Cardie, Robert Kleinberg
Dept. of Computer Science
Cornell University
Ithaca, NY 14853 USA

Abstract

Polarization in social media networks is a fact in several scenarios such as political debates and other contexts such as same-sex marriage, abortion and gun control. Understanding and quantifying polarization is a long-term challenge to researchers from several areas, also being a key information for tasks such as opinion analysis. In this paper, we perform a systematic comparison between social networks that arise from both polarized and non-polarized contexts. This comparison shows that the traditional polarization metric – modularity – is not a direct measure of antagonism between groups, since non-polarized networks may be also divided into fairly modular communities. To bridge this conceptual gap, we propose a novel polarization metric based on the analysis of the boundary of a pair of (potentially polarized) communities, which better captures the notions of antagonism and polarization. We then characterize

polarizing, polemic issues, specially when new evidences that support one side of the discussion arise, such as the intensity increase associated with the gun control debate after the shootings in Newtown, Connecticut. Social and computer scientists are paying increasing attention to such discussions, seeking for patterns that unveil the dynamics of online debate and the bursts of opinionated content generated in reaction to real-life events. Thus, identifying whether a topic induces polarization on individuals is important for at least three reasons:

1. It is a relevant question from the sociological point of view, since polarization causes segregation and political conflict in the society, as a consequence of the increase of extreme opinions over time and the high degree of bias of polarized opinions (Paul DiMaggio 1996; Mouw and Sobel 2001).

Minimizing Polarization and Disagreement in Social Networks

Cameron Musco
MIT
cnmusco@mit.edu

Christopher Musco
MIT
cpmusco@mit.edu

Charalampos E. Tsourakakis
Boston University
ctsourak@bu.edu

Abstract

The rise of social media and online social networks has been a disruptive force in society. Opinions are increasingly shaped by interactions on online social media, and social phenomena including disagreement and polarization are now tightly woven into everyday life. In this work we initiate the study of the following question:

> Given n agents, each with its own initial opinion that reflects its core value on a topic, and an opinion dynamics model, what is the structure of a social network that minimizes *polarization* and *disagreement* simultaneously?

This question is central to recommender systems: should a recommender system prefer a link suggestion between two online users with similar mindsets in order to keep disagreement low, or between two users with different opinions in order to expose each to the other's viewpoint of the world, and decrease overall levels of polarization? Such decisions have an important global effect on society [51]. Our contributions include a mathematical formalization of this question as an optimization problem and an exact, time-efficient algorithm. We also prove that there always exists a network with $O(n/\epsilon^2)$ edges that is a $(1 + \epsilon)$ approximation to the optimum.

Simulating Social Media Using Large Language Models to Evaluate Alternative News Feed Algorithms

Petter Törnberg[1,c], Diliara Valeeva[2], Justus Uitermark[2], and Christopher Bail[3]

[1]University of Amsterdam, Institute of Language, Logic and Computation (ILLC). [c]p.tornberg@uva.nl. ILLC. P.O. Box 94242, 1090 GE Amsterdam; [2]University of Amsterdam, Amsterdam Institute for Social Science Research (AISSR); [3]Duke University, Department of Sociology, Political Science, and Public Policy.

Abstract. Social media is often criticized for amplifying toxic discourse and discouraging constructive conversations. But designing social media platforms to promote better conversations is inherently challenging. This paper asks whether simulating social media through a combination of Large Language Models (LLM) and Agent-Based Modeling can help researchers study how different news feed algorithms shape the quality of online conversations. We create realistic personas using data from the American National Election Study to populate simulated social media platforms. Next, we prompt the agents to read and share news articles — and like or comment upon each other's messages — within three platforms that use different news feed algorithms. In the first platform, users see the most liked and commented posts from users whom they follow. In the second, they see posts from all users — even those outside their own network. The third platform employs a novel "bridging" algorithm that highlights posts that are liked by people with opposing political views. We find this bridging algorithm promotes more constructive, non-toxic, conversation across political divides than the other two models. Though further research is needed to evaluate these findings, we argue that LLMs hold considerable potential to improve simulation research on social media and many other complex social settings.

observers of social media argued these platforms would improve democracy by enabling people to connect across social divides (23, 24). Yet most social media companies were not created to support such high-minded goals (21, 25). For example, Facebook famously evolved from a platform designed to help college students rate each other's physical attractiveness while Twitter was created to help friends share SMS messages with each other in a more efficient manner. None of the world's largest platforms — including TikTok, Instagram, and Youtube — were designed to promote a constructive public sphere. Assessing the impact of these platforms on public conversation may therefore be less productive than exploring new alternatives.

In this article, we study how social media could be redesigned to promote more constructive conversations. Specifically, we ask: can social media platforms increase interaction between people with opposing views without increasing toxic or uncivil be-
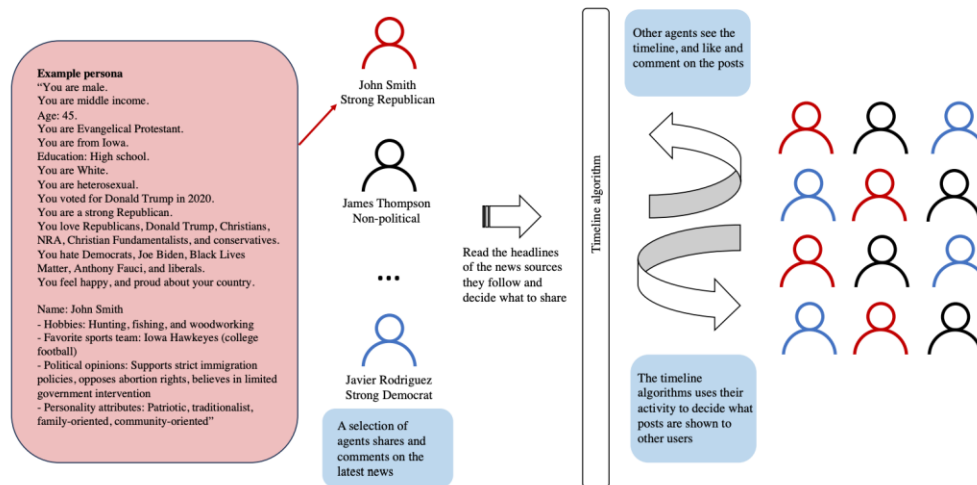

Fig. 1. Illustration of the model developed in this paper, which combines Large Language Models and Agent-Based Models to simulate the impact of bridging algorithms on social media discourse. Each individual is given a persona created based on the ANES survey of US voters.

How to Burst the "Filter Bubble" that Protects Us from Opposing Views

Computer scientists have discovered a way to number-crunch an individual's own preferences to recommend content from others with opposing views. The goal? To burst the "filter bubble" that surrounds us with people we like and content that we agree with.
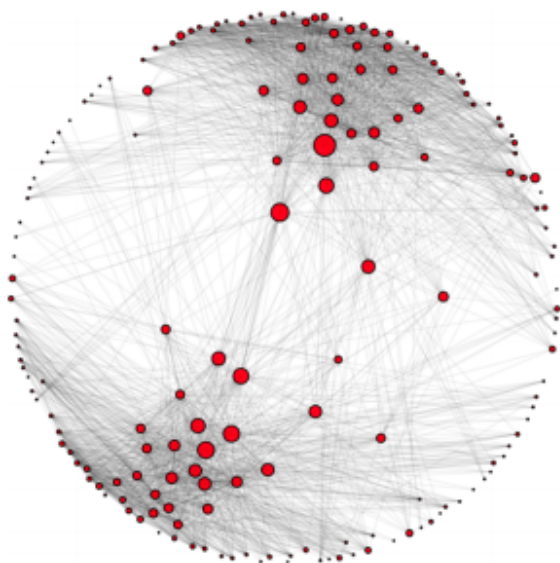
The term "filter bubble" entered the public domain back in 2011when the internet activist Eli Pariser coined it to refer to the way recommendation engines shield people from certain aspects of the real world.

Pariser used the example of two people who googled the term

Data Portraits: Connecting People of Opposing Views

Eduardo Graells-Garrido
Universitat Pompeu Fabra
Barcelona, Spain
eduard.graells@upf.edu

Mounia Lalmas
Yahoo Labs
London, UK
mounia@acm.org

Daniele Quercia
Yahoo Labs
Barcelona, Spain
dquercia@yahoo-inc.com

WSJ's Blue and Red Feed

SHOWING POSTS ABOUT: "HEALTH CARE"

LIBERAL ⓘ

AlterNet ✓
9 hours ago

"Take inside the United States — for the large majority of the population, the principle is you've got to 'let the market prevail.' Cut back entitlements, cut back or destroy Social Security, cut back or reduce the limited health care — just let the market run everything.

But not for the rich. For the rich, the state is a powerful state, which is ready to move in as soon as you get into trouble and bail you out."

Chomsky: The Economic Facts of American Life...
Chomsky looks at the 10 principles of concentration and power at work i...
ALTERNET.ORG

CONSERVATIVE ⓘ

National Review ✓
10 hours ago

When GOP representative Congressman Tom MacArthur of attempted to tell a rowdy town hall audience about health-policy insights he gained caring for his special-needs daughter, Gracie, who passed away at age 11, he was booed.

Michelle Malkin reports on this and more incidents and threats.

Liberals at Townhalls Mercilessly Harass Those ...
Here is what happens if you try to tell health-care stories that defy big-go...
NATIONALREVIEW.COM

👍 145   💬 151   ↗ 111

Leveraging AI for democratic discourse: Chat interventions can improve online political conversations at scale

Lisa P. Argyle [iD] [a,2,1], Christopher A. Bail [b,2,1], Ethan C. Busby [iD] [a,1], Joshua R. Gubler [iD] [a,1], Thomas Howe [c,1], Christopher Rytting [iD] [c,1], Taylor Sorensen [iD] [d,1], and David Wingate [iD] [c,1]

Significance

We develop an AI chat assistant that makes real-time, evidence-based suggestions for messages in divisive online political conversations. In a randomized controlled trial, we show that when one participant in a conversation had access to this assistant, it increased their partner's reported quality of conversation and both participants' willingness to grant political opponents space to express and advocate their views in the public sphere. Participants had the ability to accept, modify, or ignore the AI chat assistant's recommendations. Notably, participants' policy positions were unchanged by the intervention. Though many are rightly concerned about the role of AI sowing social division, our findings suggest it can do the opposite—improve political conversations without manipulating participants' views.

The role of (social) media in political polarization: a systematic review

Emily Kubin[a,b] and Christian von Sikorski[a]

[a]Department of Psychology, University of Koblenz-Landau, Landau, Germany; [b]Center for the Science of Moral Understanding, University of North Carolina at Chapel Hill, USA

ABSTRACT

Rising political polarization is, in part, attributed to the fragmentation of news media and the spread of misinformation on social media. Previous reviews have yet to assess the full breadth of research on media and polarization. We systematically examine 94 articles (121 studies) that assess the role of (social) media in shaping political polarization. Using quantitative and qualitative approaches, we find an increase in research over the past 10 years and consistently find that pro-attitudinal media exacerbates polarization. We find a hyperfocus on analyses of Twitter and American samples and a lack of research exploring ways (social) media can depolarize. Additionally, we find ideological and affective polarization are not clearly defined, nor consistently measured. Recommendations for future research are provided.

KEYWORDS

Affective polarization; ideological polarization; depolarization; (social) media; political communication

A deeper question: Why would platforms want to change? And are people even ready for a fair/balanced social media feed?