

Characterizing the Social Media News Sphere through User Co-Sharing Practices

Mattia Samory*

Computational Social Science
GESIS, Cologne, Germany
mattia.samory@gesis.org

Vartan Kesiz Abnoui

Department of Computer Science
Virginia Tech, VA, USA
vkesizab@vt.edu

Tanushree Mitra

Department of Computer Science
Virginia Tech, VA, USA
tmitra@vt.edu

Abstract

We describe the landscape of news sources which share social media audience. We focus on 639 news sources, both credible and questionable, and characterize them according to the audience that shares their articles on Twitter. Based on user co-sharing practices, what communities of news sources emerge? We find four groups: one is home to mainstream, high-circulation sources from all sides of the political spectrum; one to satirical, left-leaning sources; one to bipartisan conspiratorial, pseudo-scientific sources; and one to right-leaning, deliberate misinformation sources. Next, we measure which assessments of credibility, impartiality, and journalistic integrity correspond to social media readers' choices of news sources, and uncover the multifaceted structure of the social news sphere. We show how news articles shared on Twitter differ across the four groups along linguistic and psycholinguistics measures. Further, we find that with a high degree of accuracy (~80%), we can classify in what news community an article belongs to. Our data-driven categorization of news sources will help to navigate the complex landscape of online news and has implications for social media platform maintainers to reliably triage questionable outlets.

Introduction

Two-thirds of American adults read news on social media, even though a majority of them expect the reported information to be mostly inaccurate (Matsa and Shearer 2018). It is a pressing concern to inform users about the quality of the news outlets in their social media feeds. In response, social media platforms, advocacy groups, and the social computing research community developed multiple assessments of the quality of news sources—for example, political slant (Elejalde, Ferres, and Herder 2017) or credibility (Soni et al. 2014). These assessments of news outlets allow researchers to study the social news sphere and to characterize its social media audience. However, on the one hand, there is no single reliable assessment of the quality of a news outlet (Zhang et al. 2018). On the other, we know little about

which characteristics of the news outlets drive the actual audience engagement on social media. What do users share and reply to most? This poses challenges for identifying which of the many assessments to prioritize when evaluating news. In particular, the assumption that users are consuming an exclusive diet of partisan-aligned news appears increasingly less relevant in the social news sphere, which calls for a more nuanced characterization (Guess et al. 2018; Starbird 2017).

This paper provides just such a characterization. We focus on 639 news sources, both credible and questionable, through a quantitative analysis of over 31 million Tweets and 1 million news articles. We adopt multiple measures of news source quality, including external assessments of sources' factuality, impartiality, and journalistic integrity. Then, we compare and contrast the external assessments of the news sources, the characteristics of their articles, and the audience that shares them on Twitter. This brings novel insight into which external assessments of news sources correspond to common audience, as well as into what types of news articles do different audience pools engage with. The paper is structured around three research questions:

RQ1: What communities of news sources emerge, when considering the sharing practices of their social media audience?

First, we connect sources to reflect how many users share links to both—a costly signal that sources cannot easily falsify (Donath 2011). We find four communities of connected sources: highly circulated news sources spanning the entire political spectrum; engaging misinformation, such as click-bait and satirical news sources; factoid misinformation, such as conspiratorial and junk science sources; and misinformation with an ulterior motive, comprised of far-right propaganda and fake news sources.

RQ2: How well external assessments of news sources distinguish the audience-based news communities?

We find that the social news sphere does not simply follow partisan polarization. We show that it is necessary to combine multiple assessments of political slant and journalistic norms to explain similar audiences on Twitter.

*A portion of this work was conducted while the author was at Virginia Tech
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

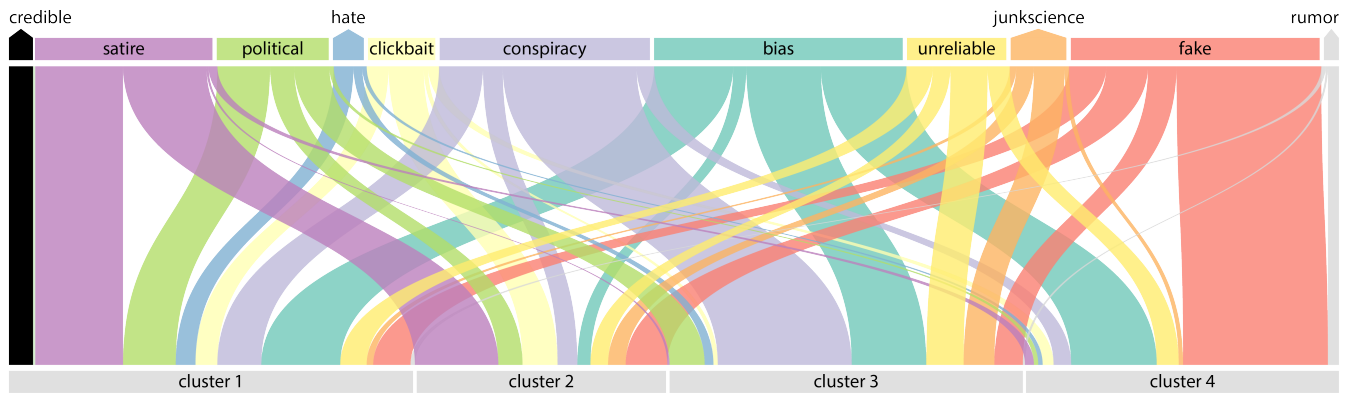


Figure 1: Cluster composition in terms of expert assessments of credibility. We map the primary category of opensources.co for each source to the corresponding cluster. Cluster 1 is the destination for all of the credible sources, as well as for political misinformation sources; cluster 2 is home to satirical and clickbait sources; cluster 3 hosts most conspiratorial, unreliable, and junk science sources; cluster 4 is largely comprised of fake news and biased sources.

Finally we look for differences in the news articles that the communities of news sources shared on Twitter.

RQ3: What kind of language and engagement attributes characterize communities of news sources?

We analyze how the articles differ in content, style, sentiment, psycholinguistic categories, as well as in how much engagement they received on Twitter.

To summarize, our data-driven study highlights a disconnect between how experts categorize news sources, and how their audience selects them on Twitter. We shed light on which expert assessments do help in differentiating between audience-based clusters of news sources—although in previously unidentified combinations. Our analysis of the content of the articles pertaining to each cluster can guide media scholars in assessing what kind of news sources bring together or divide social news media readerships. Our results also yield important practical implications. This work can help social media platforms to reliably triage new information outlets, thus alleviating the labor-intensive labeling of an ever growing information space. While creating information outlets is cheap, controlling its audience similarity space is costly; hence a reliable, hard-to-fake signal. For example, a new conspiratorial outlet will have to wield significant effort to position itself close to credible information sources in the shared social media audience space.

Related Work

Here we outline the scholarly work that informed our study. First we present a line of research studying the social news sphere from the point of view of an ecosystem of news producers and their audience. Then, we discuss the large body of work aimed at characterizing misinformation and developing indicators for identifying it.

Ecosystem of news and audiences on social media

Most related to this work is research that identifies groups of news sources by the social media audience they have in common (Mukerjee, González-Bailón, and Majó-Vázquez 2018;

Webster and Ksiazek 2012). These studies contribute fundamental methods to assess the levels of fragmentation of the social news sphere. However, results from this line of work are ambiguous, thus calling for further study.

One study hypothesized that audience behavior is the inclination of users to engage with like-minded media, resulting in a highly fragmented news media sphere (Webster and Ksiazek 2012). Indeed, most social media users are often not careful about shaping a balanced information diet for themselves, and consume content on a limited number of topics (Kulshrestha et al. 2015). Furthermore, certain topics, such as conspiratorial (Bessi et al. 2015) and political (Barberá et al. 2015) content, seem especially polarizing, and result in a fragmentation of the social media audience. In particular, several works depict social media audience as divided into conservative and liberal stances (Conover et al. 2011; Garimella et al. 2018; Bakshy, Messing, and Adamic 2015).

In contrast, another line of research finds little evidence of ideological segmentation in media use (Mukerjee, González-Bailón, and Majó-Vázquez 2018; Webster and Ksiazek 2012). For example, Mukerjee, González-Bailón, and Majó-Vázquez find no evidence of selective exposure or self-selection of the audience, and surface a tightly interconnected core of news sources—fundamentally formed by legacy brands. Opposing the view of a politically polarized audience, recent studies suggest that the dualism between conservative and liberal stances is increasingly less relevant in the social news sphere (Starbird 2017; Guess et al. 2018; Druckman, Levendusky, and McLain 2018). In particular, these studies show that groups of news producers employ mechanisms that are yet to be studied deeply to influence their readers (Horne and Adali 2018; Starbird et al. 2018).

Collectively, this body of literature highlight the importance of studying the social news sphere as an interconnected ecosystem of news sources, and not as a collection of independent actors. We build upon their intuition, and take the analyses a step further, by looking at how audience fragmentation corresponds to characteristics of news sources as well as of their content.

Differently from related work, we connect how Twitter users share news articles, how the emerging news communities from the user co-shares compare across the assessments of the articles' sources, and what are the distinguishable characteristics in the content of those articles. Similarly to (Horne and Adali 2018), we use multiple assessments of news sources to study their relationship. However, we use common Twitter audience as a measure of source similarity, which sources cannot directly control. This work also differs from (Mukerjee, González-Bailón, and Majó-Vázquez 2018), which introduces a network backbone extraction method to remove spurious audience overlap—but which, potentially, also penalizes sources with small or niche audience overall. To account for this, we employ instead an information-theoretic measure for computing audience overlap (Martin 2017).

Misinformation in online social media

The prevalence of online misinformation recently brought the social news sphere at the epicenter of the social computing research community as much as of the public discourse. On one end of the misinformation spectrum there are sites explicitly designed to deceive people, to publish provably false claims, and to propagate them through social media platforms in an attempt to increase readership and profit (Marwick and Lewis 2017). On the other end are satirical news sites which produce misleading content for entertainment. Regardless of the sources' intent and of how we label them (satire, clickbait, or the more contested term "fake news"), misleading information has a negative effect on citizen's news consumption and on their ability to make fully informed decisions (Thorson 2016). Misinformation, once assimilated, is hard to eradicate from the convictions of its audience (Lewandowsky et al. 2012). Hence, a large body of scholarly work has focused on characterizing online misinformation to prevent its spread. Notable among them are Horne and Adali's characterization of differences in the news titles of fabricated and real stories (Horne and Adali 2017), Golbeck et al.'s attempt to automatically determine whether a newswire article is satirical or not (Golbeck et al. 2018), and Ferrara et al.'s study of social bots spreading rumors on Twitter (Ferrara et al. 2014).

Despite the scholarly efforts, identifying misinformation is still challenging. In fact, even humans perform poorly in identifying non-obvious misinformation (Kumar, West, and Leskovec 2016). Defining what constitutes high-quality information requires careful considerations about content attributes and adherence to journalistic norms (Hayes, Singer, and Ceppos 2007; Diakopoulos 2015). Thus, recent research focuses on characterizing different types of misinformation also, on summarizing reliable indicators, rather than automatically identifying information (Zhang et al. 2018).

Motivated by the need for reliable indicators of misinformation, in this work we characterize the landscape of news producers on Twitter. Instead of focusing on the prevalence of one type of misinformation among news sources or their social media audience, we study the relationship between news sources and their audience on Twitter through multiple dimensions of news source quality.

Data

We first compiled a list of news sources of interest. Then, we collected tweets containing links to articles published by the sources. We complemented this with the full-text of the articles linked from Twitter, and with multifaceted assessments of the news sources.

List of news sources: We started by curating a list of news sources and combined a wide selection of both credible and questionable sources. We relied on OpenSources¹, a professionally curated list of online sources available for free for public use. News sources in this resource range from credible to misleading and outright fake websites. OpenSources heavily focuses on misinformation, with little emphasis on credible content. Therefore, we complemented this list by including additional mainstream news websites. We first referred to various online resources, including surveys by Pew Research² and NPR³. Then, a journalism and communications expert refined the list by referring to the most circulated⁴ and the most trusted news sources (Matsa and Shearer 2018). In total, the list contains 639 distinct web domains.

Tweets containing links to sources: Next, we collected all tweets linking to news articles by the sources in our list. Specifically, we accessed data through the Twitter streaming API, continuously from September 2016 to February 2017, and filtering those containing links to the news sources of interest. We collected a total of 31,567,501 tweets. We further collected all tweets that replied to them.⁵ We use this measure as a signal for audience attention to the original tweet.⁶

News articles linked from Twitter: Furthermore, we enhanced our dataset by collecting the news articles that the tweets link to. We used the library `newspaper` to scrape the text and publication date of the articles from the web. We

¹ opensources.co

² <https://www.journalism.org/2014/10/21/political-polarization-media-habits/>

³ <https://www.npr.org/sections/alltechconsidered/2016/10/28/499495517>

⁴ <https://www.cision.com/us/2017/09/top-10-u-s-daily-newspapers-2/>

⁵ The dataset shows skewed distributions typical of social media: although it includes over 2 million unique users, 60% of the tweets are authored by the top 1% most prolific users.

⁶ We relied on the list of over 300 news bots from (Lokot and Diakopoulos 2016) to check for the presence of inorganic Twitter activity. We found that bots in the list author less than 0.1% of our dataset. Out of the bots individually contributing over 1k tweets, all but two exclusively share credible news by a single source (these are bots by major outlets like New York Times and Washington Post, or sports bots). The remaining two bots are 365Arizona, sharing overwhelmingly biased tweets from counterjihad.com and centerforsecuritypolicy.org, and dubvNOW, a newsletter born out of the University of West Virginia sharing half credible (usatoday), half conspiracy (infowars) news. Given the minor impact of bots, we do not exclude them for the sake of completeness. Even if the list above might have missed some non-organic accounts, we stress that even the most prolific bots would skew the number of co-occurring users between two sources at most by 1, as we disregard contribution volume and focus only on user co-sharing practices.

widespread	satirical/clickbait	conspiratorial	right-wing/fake
washingtonexaminer.com	usuncut.com	activistpost.com	angrypatriotmovement.com
thegatewaypundit.com	countercurrentnews.com	21stcenturywire.com	usasupreme.com
wnd.com	occupydemocrats.com	thedailysheeple.com	truthandaction.org
Newsmax.com	attn.com	russia-insider.com	prntly.com
heatst.com	rawstory.com	blacklistednews.com	subjectpolitics.com
dailycaller.com	bipartisanreport.com	washingtonsblog.com	usanewsflash.com
dailywire.com	dailykos.com	veteranstoday.com	christiantimesnewspaper.com
americanthinker.com	addictinginfo.org	investmentwatchblog.com	ilovemyfreedom.org
conservativereview.com	politicususa.com	ronpaulinstitute.org	departed.co
wsj.com	altnet.org	wearechange.org	supremepatriot.com

Table 1: Top 10 central news sources in each cluster—i.e., sources that are the most cosine-similar to each KMeans centroid.

discarded articles that are non-English or too short⁷. Since multiple links within each web domain may refer to the same article, we identify duplicates by stemming and matching their plain text. We keep only the oldest instance of duplicate articles by publication date. The final dataset contains 1,212,304 articles.

External assessments of the news sources: We completed our dataset by gathering assessments of the news sources along multiple dimensions of factuality, journalistic integrity, and impartiality. Our choice of dimensions was dictated by decades of research from journalism and communications scholars (Sundar 1999).

Factuality: OpenSources, a professionally-curated list of news sources, provides annotations for each news source based on a taxonomy of twelve assessments, such as fake and junk science. We use the primary annotation of each source as a judgment of its factuality.

Journalistic Integrity: Journalists at NewsGuard⁸ label news sources according to multiple criteria, including whether sources regularly correct errors or disclose ownership and financing. NewsGuard attributes a weight to each assessment and combines them into an overall score from 0 to 100. An increasingly important assessment of the journalistic stance of online sources is whether they identify as alternative or mainstream outlets. To this end, we label outlets as “alternative” or “mainstream” according to the corresponding lists of sources from Wikipedia.

Impartiality: Finally, we also include assessments of political impartiality by Media Bias/Fact Check⁹ and AllSides¹⁰. Both services divide new sources into five categories: biased left, center-left, center, center-right and right. After obtaining the assessments from their website, we performed semi-automatic matching of the news domain names to the domains from OpenSources. We used a fuzzy string match to find the best candidate in OpenSources, and manually validated the match.

⁷less than 140 characters since we find that those mostly consist of broken links or 404 pages for URLs that no longer exists.

⁸newsguardtech.com

⁹mediabiasfactcheck.com

¹⁰allsides.com

Out of 639 new sources, we include annotations for 595 sources from OpenSources, 124 from Newsguard, 101 from MediaBias, 42 from Allsides, and 26 from Wikipedia.

RQ1: Communities of news sources based on common audience

By aggregating the news sharing practices of Twitter users, what communities of news sources emerge?

Method

We start by building a news source similarity space, based on shared Twitter audience. Then, we characterize groups of sources within this space.

Measuring audience similarity through co-shares: We base our audience similarity measure on user co-shares: two sources are more similar when more users share links to both of them. Intuitively, user-based similarity hypothesizes that users choose news sources according to their interests, such as the topics covered in the articles or their partisan affiliation; if many users choose two sources, those sources likely cater to a common interest. One issue with using the raw number of co-sharing users as a similarity measure is that it is bound by the overall number of users sharing a source. High-circulation sources are more likely to be shared independently of the users’ choices, and therefore may co-occur with other sources by chance. The converse holds for smaller less popular news outlets. Therefore, we measure instead whether a *surprisingly high* number of users share two sources, given the number of users who share each source independently from the other, using positive pointwise mutual information (PPMI) (Martin 2017).

Finding groups of news sources sharing Twitter audience: We adopt an unsupervised machine learning approach to cluster news sources according to user co-sharing practices. We use the popular K-Means method with kmeans++ initialization. One crucial step in K-means is to tune the number of clusters and evaluate their validity. We find the best number of clusters by looking for an elbow point in explained variance, which offers a natural trade-off between accuracy and number of clusters. Thus, we assume four as the optimal number of news source clusters.

We check that the final clusters remain stable with respect to different random seeds.

Characterizing groups of news sources through external assessments of journalistic norms: Next, we look at the composition of each cluster in terms of external assessments of credibility, impartiality, and journalistic practices. We aggregate the assessment for the sources within each cluster. Since assessments lack complete coverage—e.g., some sources have assessments for credibility but not for transparency—we look at each separately. We do not run statistical tests on the differences between clusters, since those would neglect the partial overlap of sources with different assessments.

RQ1: Results

We find four clusters of news sources based on Twitter co-sharing practices. Table 1 displays the top 10 news sources closest to the cluster centroids. Next, we study the clusters by looking at the external assessments of factuality, impartiality, and journalistic integrity of their members. We conclude our results by showing which of the external assessment best conforms to the Twitter audience-based clusters.

Credibility: We study the credibility of the sources in each cluster. Figure 1 summarizes the findings. We find that all clusters contain a variety of OpenSources assessments—e.g., all clusters include conspiratorial sources to some extent. However, their composition yields a clearly defined picture. Cluster 1 is the largest of the four. It contains all of the credible sources in our dataset, as well many misinformation sources of widespread appeal, such as political (Breitbart) and satirical (liberalbias.com) sources. High circulation sources like nytimes.com and wsj.com belong to the cluster. Cluster 2 is home to a significant fraction of satirical and clickbait sources such as theonion.com, as well as left-wing activist groups like dailykos.com. Cluster 3 hosts most of the conspiratorial, unreliable, and junk science sources, e.g., corbettreport.com and prisonplanet.com. The fourth and last cluster contains misinformation with an ulterior motive: it is largely comprised of fake news and biased sources such as abcnews.com.co¹¹, abcnewsgo.co, nbc.com.co¹². Although cluster 1 contains relatively widespread news sources, nine hateful news sources also belong to this cluster. Upon further inspection, we find that users piggyback on the popularity of mainstream articles to promote hate by sharing links to both in the same tweet; this explains the co-presence of hate and widespread news in the co-sharing practices.

Impartiality: Figure 2 shows the composition of the clusters in terms of partisan bias. We present results based on assessments from MediaBias; results based on Allsides are qualitatively consistent, therefore we omit them. Cluster 1 includes sources from all political sides. This is in line with

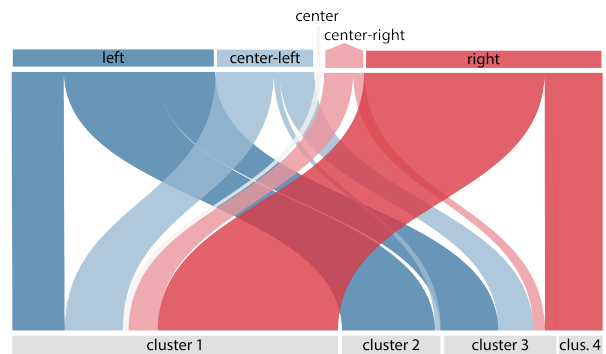


Figure 2: Cluster composition in terms of external assessments of impartiality, according to MediaBias. We omit sources for which MediaBias does not provide assessments. All political sides are represented in cluster 1—which is in line with their widespread appeal according to credibility assessments; all sources in cluster 2, containing satire and clickbait, are left-wing; cluster 3, home to conspiratorial sources, shows a mix of left and right-wing sources; all sources with an assessment in cluster 4, hosting predominantly fake news, are right-wing.

the finding that emerged from the previous credibility assessments, that cluster 1 is home to high circulation credible and misinformation news sources. In contrast, cluster 2, containing satire and clickbait, includes exclusively left-wing sources. Research from cognitive psychology suggests that left-wing political inclinations are more eager to respond to pleasing content (Dodd et al. 2012), which may explain the correlation between satirical content and left-wing political composition of the cluster. Cluster 3, home to most conspiratorial sources, shows a mix of left and right-wing sources, such as the left-leaning activistpost.com and the right-wing returnofkings.com. Conspiracy theorizing is in fact a bipartisan issue (Oliver and Wood 2014). Finally, cluster 4, hosting predominantly fake news, includes exclusively right-wing sources. Grinberg et al. also found a cluster of fake news sources sharing overlapping audiences on the extreme right of the political spectrum (Grinberg et al. 2019).

Journalistic integrity: Finally, we examine the journalistic integrity of the sources as assessed by Newsguard. We first look at the overall source quality score (0–100, a higher score means higher quality). A website needs a score of at least 60 for it to receive an overall positive assessment. The results corroborate our previous findings. Cluster 1, which contains widespread news sources, receives the best scores: it has the highest mean (40.12) and median (32.5) score, and many of its sources receive perfect scores. Cluster 2, follows a pattern similar to cluster 1. On the other hand, clusters 3 (conspiratorial) and 4 (fake) contain mostly low-quality journalism. Barely any source in cluster 3 exceeds a score of 20 (with the notable exception of counterpunch.org, a non-profit ultra-liberal magazine). Almost all sources in cluster 4 are below the cutoff score of 60 for reputable sources.

¹¹<https://en.wikipedia.org/wiki/ABCnews.com.co>

¹²<https://mediabiasfactcheck.com/nbc-com-co/>

Next, we look at the break-down of the journalistic integrity score into its different facets. Sources in clusters 1 and 2 score systematically the highest on all criteria. In particular, clusters 1 and 2 are notably higher than clusters 3 and 4 in regularly correcting errors, presenting information responsibly, and avoiding publishing false content.

To summarize We discover four clusters of news sources, according to common audience on Twitter. They show distinct compositions in terms of credibility, impartiality, and journalistic integrity. Therefore, we will interchangeably refer to the clusters with the following shorthand:

cluster 1: widespread news sources

cluster 2: satirical/clickbait news sources

cluster 3: conspiratorial news sources

cluster 4: right-wing/fake news sources

We stress that the characterization in this section applies to the clusters collectively, and individual sources may show different characteristics.

RQ2: News source assessments distinguishing the social news audience sphere on Twitter

We identified communities of news sources that differ along multiple assessments of credibility, integrity, and impartiality. However, which assessments correspond to different choices of news sources by the audience? We answer this question by ranking the assessments according to how well they distinguish the communities of shared audience. To corroborate this, we show that the top-ranking assessments also correlate highly with overall audience variance.

Which assessments best explain communities of shared audience? We measure how well grouping sources by each assessment conforms to the communities of shared audience. We label sources with the cluster they belong to, as well as with each assessment—e.g., the left- and right-wing bias labels given by MediaBias. Then, we compute how much the assessment and cluster labels overlap, according to measures of cluster quality. We repeat this process for all assessments, and we rank them accordingly.

We find that the simple assessment of alternative vs. mainstream media sources shows the highest homogeneity, completeness, and v-measure, whereas the assessments by OpenSources show the highest Rand Index and Mutual Information. This discrepancy may be because few news sources have labels for the alternative vs. mainstream assessment, and therefore are penalized in the latter chance-adjusted measures. Furthermore, we ask if all the labels in OpenSources' assessments—12 of them—are actually informative. We repeat the experiment assessing each label in a one-versus-all fashion, and we find that only few of them correspond to the audience-based clusters: bias, clickbait, conspiracy, fake, and satire rank higher than all other assessments in the study. This suggests that the opensources assessments may be redundant or, more concerningly, that

only few of them correspond to audience co-sharing practices. Surprisingly, in all cases, the often-adopted assessments of partisan bias by both MediaBias and Allsides perform poorly.

Which assessments correlate with audience variance?

To corroborate these findings, we investigate the correlation between source assessments and audience similarity. We use principal component analysis on the PPMI-weighted audience similarity matrix, and identify the direction along which audience similarity changes the most. We focus on the first principal component, which explains the most variance (18%, more than double the variance explained by the second component). The component correlates highly and positively with assessments of factuality, such as conspiracy and junk science, and only mildly and negatively with political bias. This confirms that audience similarity does not simply follow a polarized, partisan structure.

Finally, to interpret along which direction does audience vary, we inspect the sources that load highly on the principal component. On the positive end we find conspiratorial sources that propose a revolution in public institutions. For example, the anarcho-capitalist conspiratorial website corbette-report.com includes stories about deep state and globalist control, and loads among the most highly positive. On the negative end, instead, we find sources adopting the political status quo. For example, the website madpatriots.com, which leverages established conservative narratives to craft sensationalist headlines, such as “violent immigrants” and “red scare,” loads the most negative. This echoes Starbird's observation that political leanings of alternative news sites feature an anti-/pro-globalist orientation, rather than conservative/liberal.

RQ3: Identifying the source community of news articles on Twitter

The previous sections identified clusters of news sources shared by distinct audiences, and characterized the journalistic qualities of the sources within the clusters. Yet, if we see an article shared on Twitter, can we automatically identify the cluster it belongs to? We provide a fine-grained content analysis of the clusters articles. Then, we show that a classifier identifies the correct cluster with over 80% accuracy.

Method

In particular, we look at the content, style, sentiment, and psycholinguistic categories used in the articles, as well as the overall engagement that the articles produce on Twitter.

Content: First, we look at what kind of content distinguishes the clusters. We rely on newspaper python module to extract the plaintext of the articles and to discard ones whose primary language is not English. Then, we preprocess the plaintext using standard procedures of converting to lowercase, removing punctuation and accents, striping whitespaces, and removing stopwords and links. Next, we use Sparse Additive Generative models—*SAGE* (Eisenstein, Ahmed, and Xing 2011)—to find words that are specific to each cluster. *SAGE* uses a regularized log-odds ratio

measure to contrast word distributions between one corpus of interest against a baseline corpus. We contrast articles within each cluster against all the remaining. For example, we compare the distribution of words in widespread news sources vs. the distribution of words in all other clusters. The output will give the distinguishing words of the widespread news sources, *relative* to other clusters.

Style: Content analysis tells us what sources in each cluster are writing about. However, how are they writing about it? We next turn to writing style. We use two feature sets: cue verbs, which reflect the sources' choices in terms of journalistic *reporting style*, and *stylo-metric features* reflecting the sophistication of the writing style.

Cue verbs correspond to factuality judgments of journalistic reports (Sauri 2008). Sources may choose certain types of cues to nudge their audience into believing a reported claim. For example, “WSJ learns that . . .” asserts more certainty than “WSJ suspects that . . .” Following (Soni et al. 2014), we use five groups of cue verbs common in Twitter: “Report” (e.g., say, report), “Knowledge” (e.g. learn, admit), “Belief” (e.g. think, predict), “Doubt” (e.g. doubt, wonder), and “Perception” (e.g., sense, hear).

We also look at *stylo-metric* features that capture the articles' style beyond choices for reporting. On the one hand, sources may embrace more formal linguistic registers to appear more credible. On the other hand, they may make content more accessible by simplifying their language. We assess the readability of the text using the SMOG index. SMOG estimates the grade level needed for understanding the article, and correlates well with human judgments of text clarity—high SMOG implies more complex language.

Additionally, we look into markers of complexity: the overall number of words, number of informative words, type-token ratio, and long and complex words, all give us an indication of the writer's effort. Finally, we include function word types, such as conjunctions and prepositions. Function words do not carry meaning *per se*. However, they reflect how people are communicating (Tausczik and Pennebaker 2010).

Sentiment and psycholinguistic categories: *Sentiment* and *psycholinguistic* categories depict the emotional and psychological states evoked in the articles. We use two scores from VADER (Hutto and Gilbert 2014): *compound* score measures sentiment polarity, while *neutrality* captures the proportion of words that are not emotionally charged. Intuitively, sensationalist journalism uses more emotionally charged and polarized language (Zhang et al. 2018).

Additionally, LIWC (Tausczik and Pennebaker 2010) offers validated categories of words for assessing such psycholinguistic dimensions. We use the following categories. At the most superficial level, content word categories explicitly reveal the articles' focus—we include the topical category *personal concerns*. The categories *pronouns* and *verb tense* also expose the attentional focus of the writers, although more subtly. For example, pronouns like “we” and “our” often indicate references to group identity. Similarly, writers may use verb tense to signal endorsement: for example, past tense increases psychological distance with the re-

ported facts, compared to present tense. We include the categories *social* and *emotional* because they may signal less informative, more sensationalist content. We also include *motion*, *cognitive processes*, and *sense* words, since they are known signals of truthfulness. Finally, the *spoken* category exposes departure from journalistic style, capturing, for instance, the occurrence of swearwords and nonfluencies. To account for the high variability in article length, for every article we compute the fraction of words in each category.

Twitter engagement: Finally, we gauge the feedback that articles receive on Twitter. We compute the number of followers and friends (followees) of the Twitter users to assess their authority or hub status. For individual tweets, we report the number of favorites, replies, and retweets, which signal the reach of and engagement with the tweet.

Measuring differences across clusters We use a machine learning pipeline to assess differences in articles' language and engagement between clusters. We train a multinomial logistic regression model (henceforth MNLogit), penalized for dealing with sparsity and multicollinearity. We compare each cluster to the baseline “widespread news sources.” We choose “widespread news sources” as a reference class because, according to our previous findings, we expect the content to be more varied and mainstream, and because it contains the largest number of articles. Since MNLogit is a one-vs-all model, we validate the results with a post-hoc statistical significance test of the observed differences between clusters. We perform a series of non-parametric Kolmogorov–Smirnov (KS) tests for all six possible pairwise combinations of the clusters. We chose the KS test over the parametric t-student to account for the fact that none of our features are normally distributed. For all features we assessed normality using Anderson–Darling tests.

RQ3 Results: characterizing of news communities' language

We find that clusters significantly differ in the articles' language and in the engagement they produce.

Content: Table 2 reports the words that best distinguish articles from each cluster. The “widespread news” cluster, home to a wide range of sources, intuitively does not show a predominant topic: we find references to current events, like the fire which burnt the Grenfell tower and the marches against the Malaysian Prime Minister Najib. The other three clusters, instead, manifest more focused topics. Cluster “satirical/clickbait” features several polarizing terms (like *lgbtq*, *@realdonaldtrump*, *impeach*, *protectors*), as well as words often used to poke fun of stereotypical conspiracy theorists (like *pyramids*, *extraterrestrial*), and food items often featured in prodigious diets meant to generate click revenue (like *coconut*, *turmeric*).¹³ Cluster “conspiratorial” aptly uses the conspiratorial lingo, such as references to the New World Order, Zionists, Illuminati, and the

¹³the top terms like “mashshare” and “flipboard” are artifacts of the automated text extraction, and correspond to the text found in social buttons typical of clickbait sites. For integrity, we did not editorialize them from the table.

Rothschild—all being frequent actors in conspiratorial narratives of world domination. Cluster “right-wing/fake” uses terms of the hyper-conservative propaganda, such as references to the alleged lack of virility of the left (like effeminization, soy boys, and cucks) and political opponents (like antifa, leftwing, Ocasio-Cortez). In summary, we find that the content of the articles conforms with the external assessment of the sources within the news communities.

Style: Beyond content, news sources may make subtle stylistic choices to convey their message. Recall that all results are relative to the reference class, “widespread news.” All three remaining clusters use significantly fewer *report* cue words. Report cue words, such as *sources report that...* or *Whistleblower tells Congress...*, are a landmark of formal journalistic style. While “satirical/clickbait news” express more doubt in its reports than “widespread news,” “conspiratorial news” and “right-wing/fake” doubt less. In fact, “conspiratorial news” use more cue verbs expressing knowledge. Knowledge cues, such as *Scientists find that change driven largely by increased carbon dioxide...*, are part of the class of factive predicates, which imply the truth of the claim that follows (Sauri 2008).

We then turn to stylometric features to probe the sophistication of writing style. “Satirical/clickbait news” use more difficult to read language (possibly mimicking a pompous writing style), whereas “right-wing/fake news” use simpler language, according to SMOG. This is reflected in the other markers of complexity: “right-wing/fake news” use fewer long words, and a less varied vocabulary according to the type-token ratio. Yet, “right-wing/fake news” also use more polysyllabic words (complex words in the table) and more word types overall. This may signal a somewhat diluted language: more verbose, but not as informative. We find that this is indeed the case. All three clusters use a larger proportion of function words than “widespread news,” which means that the informative words are relatively fewer. In a nutshell, “widespread news” adhere to the expectations of standard journalistic style, whereas “satirical/clickbait news,” “conspiratorial news,” and “right-wing/fake news” use less formal, less informative, more accessible language. In particular, “conspiratorial news” and “right-wing/fake news” express more certainty in their reporting, possibly to appear more credible.

Sentiment and psycholinguistic categories: Abstracting from the level of content, we next analyze the sentiment and psycholinguistic categories expressed in the articles. Table 3 reports the differences between clusters. First, we summarize the characteristics that cluster “satirical/clickbait,” “conspiratorial,” and “right-wing/fake” have in common, and compare them to cluster “widespread news sources”. As expected, clusters “satirical/clickbait,” “conspiratorial,” and “right-wing/fake” use more emotionally charged and polarized language (i.e. a higher score for neutral score). In particular, we find that the clusters disproportionately express anger death-related words. In addition, they use more informal (spoken word categories like swear) language, with more references to collective identity (we pronoun) than “widespread news sources”. These psycholinguistic cate-

gories signal attempts of engaging rather than objective content (Hartung et al. 2016). Moreover, “satirical/clickbait,” “conspiratorial,” and “right-wing/fake” express more certainty (cognitive mechanism categories like certain and insight) through visual language references (see); visual language makes concepts more concrete, thus more memorable and accessible. Although journalistic news reporting style implies a temporal focus on the present or the near past, we find that “satirical/clickbait,” “conspiratorial,” and “right-wing/fake” focus more on the future. This is in line with the stylistic analyses in the previous section, suggesting that “widespread news sources” is the one that is most conforming to a formal reporting practices. Yet, psycholinguistic categories allow us to characterize clusters with higher precision. For example, “satirical/clickbait” uses more social words and references biological processes like sexuality and health. Cluster “conspiratorial” shows the most worry with the self (I pronoun) and the least focus on the present. Cluster “right-wing/fake” references conservative values like family, home, and work, and expresses the most negative sentiment (anger and anxiety).

Twitter engagement The previous sections show that the clusters departing from journalistic practices may seek the audience’s attention through multiple content, stylistic, and framing devices. Are they successful? Twitter users sharing links from the “widespread” cluster follow less, and are followed more than Twitter handles that share links from “satirical/clickbait,” “conspiratorial,” and “right-wing/fake.” In other words, users tweeting about mainstream content tend to be information authorities, rather than hubs. Links from sources in “widespread” also engage more users, both in terms of replies and retweets, than links in the other clusters. Intuitively, the only exception is that the audience especially favors links from the “satirical/clickbait.” It appears that fringe content does not, after all, elicit as many reactions as mainstream articles. All reported effects are statistically significant. We skip showing result table due to space limits.

Automatically identifying an article’s cluster We conclude by demonstrating that the differences between clusters are not just qualitative, but that we can automatically infer the cluster that produced an article with high precision. We train an XGBoost model using the style, sentiment, and psycholinguistic features presented above; we also include the top 50k terms in the articles weighted by TF-IDF as topical features. Since classes are heavily imbalanced, we use SMOTE to oversample the minority classes in the training set—test sets contain only actual and not synthetic data. We assess model performance in a 10-fold stratified shuffle-split scheme, and tune hyperparameters using a randomized, cross-validated search approach.

We find that we can classify all clusters with a weighted accuracy of 82% (weighted F1 score). We can classify “widespread news sources” almost perfectly (90%). Although recall is relatively lower for the other clusters (84%). We reach high precision for all other clusters. On the one hand we see that several articles are false positives for “widespread”, which demonstrates that alternative content may be subtly similar to mainstream. On the other hand,

widespread news	satirical/clickbait news	conspiratorial news	right-wing/fake news
triforium, gassama, otw, luzhniki, agung, moorland, sidebars, shortcode, windrush, southgate, grenfell, yas, nbsp, fitr, istockphoto, liege, playback, sidebar, abbey, prix, najib, redistributed, getty, rebounds, derby, greets, afp, heathrow, caption, rewritten	mashshare, flipboard, stumbleupon, digg, atn, anonhq, truea, screengrab, loading, republish, spoilers, impeach, featured, protectors, dapl, screenshot, queer, realdonaldtrump, pinterest, eichenwald, lgbt, extraterrestrial, lgbtq, attribution, android, tumblr, turmeric, pyramids, coconut, delicious	altnews, in5d, naturalnews, dmca, coward, analyses, glp, adsense, quoting, eyeo, abusive, nwo, ammol, button, sheeple, pravda, vibration, rothschilds, anonymous, neocon, neocons, violation, aipac, illuminati, disclaimer, click, zionist, cabal, user, oligarchy	commments, eagler, oversign, duely, effeminization, photopin, adblock, ocasio, digestible, corruptly, unmasking, gleaned, minted, quaking, antifa, hawkins, disable, turley, shoebat, czars, im, glazov, overturns, cuck, chapters, leftwing, cortez, slinging, bois, pitchforks

Table 2: Most distinguishing words for each cluster, extracted using SAGE. Intuitively, no predominant topic distinguishes “widespread,” since it is the most diverse in terms of sources. The other clusters show distinctive topics, in line with the characterization of the sources composing them. “satirical/clickbait,” include satirical and clickbait sources, sports polarizing (e.g. impeach, protectors) and stereotypical conspiratorial words (e.g. pyramids, extraterrestrial). Cluster “conspiratorial news sources,” include conspiracy and junks science sources, adopts a conspiratorial lingo (e.g. illuminati, sheeple). Cluster “right-wing/fake,” include fake news and right-wing sources, uses hyper-conservative propaganda terms (e.g. pitchforks, [soy] bois).

achieving high precision has practical implications in automating the triaging of fringe content, because it allows to distinguish between satirical, conspiracy, and fake content.

Discussion

Implications for characterizing the news sphere based on social media audience co-shares

Recent research calls for a nuanced picture of the social news sphere that goes beyond established partisan polarization (Starbird 2017; Guess et al. 2018; Druckman, Levendusky, and McLain 2018). Our results add to this line of work by detailing the social news sphere as a multifaceted environment, and quantify the need for multiple assessments to understand audience choices, instead of relying on standalone external assessments.

We find that several dimensions of credibility and journalistic stance intertwine with political partisanship to shape the social media readership. For instance, a cluster containing entertaining misinformation like satire is also home to far-left activism. A second cluster which is deeply embedded in established political narratives publishes extreme right-wing propaganda and fake news. The correlation between politicized audiences and misinformation outlets mirrors existing research, that associates individuals on the political left-wing to engagement with positive content, whereas the right-wing to threats (Dodd et al. 2012). Yet, misinformation on Twitter does not only target politicized audiences, or even audiences that are distinct from those of credible sources. For example, a cluster focusing on conspiracism and “alternative truths” comprises of sources spanning the political spectrum. Similarly, a cluster of widespread news contains all of the credible sources under study, but also many political misinformation sources, and even hateful content. In other words, quality and questionable information live in close quarters.

The present work sheds light on the relationship between users and news sources. Specifically, we find that certain *types* of misinformation are better than others at ex-

plaining this relationship. Distinct Twitter audiences seem to follow factoid misinformation (conspiracy and junk science), misinformation with an ulterior motive (fake, bias), and misinformation with a frame of engagement (clickbait, satire). In particular, the drivers for extreme left- and right-wing activism appears different. The left appears associated with satirical, socially versed and engaging misinformation. Right-wing activism instead appears associated with emotionally negative, fake stories of threats by the political adversary. Our data-driven categorization of the four clusters of news communities is corroborated by the composition of the clusters, the external assessments of the sources within them, and the analyses of the news articles they publish. We find that other assessments of news sources, e.g., labels of hate, (un)reliable, rumor, state and political, do not help analyze the social news sphere—be it because they are not as frequent in the data, not reliable as indicator, or not as informative of sharing practices. These findings advance our understanding of what assessments of news sources are relevant to navigate the way social media audience shares them.

Our results end on a positive note. We find that fringe content does not receive as many reactions as mainstream content. This suggests that, although Twitter users may be exposed to both quality and questionable information, they do not grant as much attention to the latter. However, a more conclusive analysis of the effect of fringe content should also take into account the overall reach, strength, polarity, and quality of those reactions.

It is far from our intentions to put a moral or severity judgment on different types of misinformation. Censoring right-wing fake news or conspiracy theories as more morally wrong or dangerous than clickbait frauds and hyper-liberal hacktivism might do more damage than good. For example, attempts to directly confront extreme political or conspiratorial views may prove counterproductive (Bail et al. 2018; Peter and Koch 2016; Lewandowsky et al. 2012). Instead, it is crucial to understand the types of misinformation that the audience engages with on social media. The present work is a step in this direction. This will, in turn, inform media lit-

sentiment		clickbait	conspir.	fake
intercept		-2.1 *	-1.36 *	-2.92 *
compound		-0.04 *	0.00	-0.31 *
neu		-0.33 *	-0.38 *	-0.39 *
psycholinguistic		clickbait	conspir.	fake
intercept		-2.17 *	-1.58 *	-3.01 *
we		0.01 *	0.02 *	0.02 *
they		0.01 *	-0.01 *	0.01 *
you		0.00	-0.04 *	-0.04 *
i		-0.01 *	0.02 *	-0.08 *
shehe		0.03 *	-0.18 *	-0.02 *
ipron		0.03 *	-0.00	0.08 *
friends		0.01 *	-0.10 *	-0.06 *
family		0.05 *	-0.10 *	0.04 *
anger		0.10 *	0.17 *	0.24 *
sad		0.02 *	-0.01 *	-0.01 *
anx		0.04 *	0.00	0.05 *
posemo		0.06 *	0.31 *	-0.09 *
discrep		-0.04 *	-0.06 *	-0.03 *
cause		0.13 *	0.25 *	-0.02 *
tentat		0.09 *	0.02 *	0.10 *
excl		0.01 +	0.02 *	-0.01 *
certain		0.23 *	0.23 *	0.21 *
insight		0.19 *	0.32 *	0.11 *
hear		-0.14 *	-0.41 *	-0.19 *
see		0.12 *	0.07 *	0.09 *
feel		-0.01 *	-0.10 *	-0.11 *
sexual		0.03 *	-0.14 *	-0.03 *
body		0.11 *	0.07 *	0.04 *
health		0.05 *	0.10 *	-0.03 *
ingest		0.11 *	0.05 *	-0.01 *
future		0.04 *	0.03 *	0.03 *
present		0.03 *	-0.11 *	0.18 *
past		-0.06 *	-0.30 *	0.03 *
motion		-0.01 *	-0.03 *	0.04 *
space		-0.01 *	0.03 *	-0.11 *
time		0.02 *	-0.15 *	-0.17 *
achiev		-0.03 *	-0.24 *	-0.10 *
work		-0.06 *	-0.09 *	0.03 *
leisure		-0.14 *	-0.15 *	-0.24 *
death		0.01 *	0.09 *	0.01 *
home		0.02 *	-0.05 *	0.02 *
relig		-0.07 *	0.02 *	0.16 *
money		-0.15 *	0.00 +	-0.31 *
filler		0.08 *	0.02 *	0.07 *
nonflu		-0.03 *	-0.00 +	0.00
swear		0.16 *	0.14 *	0.08 *
assent		-0.00	-0.01 *	0.05 *

Table 3: Sentiment and psycholinguistic differences between clusters. Cluster 1 being the baseline comparison cluster is not shown. All numbers should be interpreted in reference to cluster 1. Clusters 2, 3, and 4 express more emotional (less neu) and informal (spoken words) language than cluster 1. Cluster 2, satirical and clickbait, uses more social words and references sexuality and health. Cluster 3, conspiratorial, shows the most self-concerns (I pronoun) and the least focus on the present. Cluster 4, right-wing fake news, references values like family, home, and work, and expresses the most negative sentiment (anger and anxiety). * $\Rightarrow p < .001$, * $\Rightarrow p < .01$, + $\Rightarrow p < .05$

eracy campaigns to educate the audience about journalistic norms and practices.

Social media audience as a costly signal of news source similarity

One merit of using audience as a similarity measure between news sources is that the sources cannot easily manipulate their position in this space. For example, a junk science source to position itself close to credible information sources would need to control a large fraction of shares of its own content, and to also consistently share markers of credible content. These two procedures would be costly in terms of effort for the junk science source, and counterproductive in terms of audience targeting. In other words, audience similarity is a signal that is hard to falsify (Donath 2011). A word of caution: hard does not mean impossible. However costly a signal, individuals with enough resources and motivation can, for example, acquire fake audience or fake engagement.

Implications for gatekeeping the social news sphere

Online platforms are increasingly embracing source-level assessments for nudging their users towards high-quality news while preserving access to a pluralistic social news sphere. For example, Facebook offers additional information about sources appearing in the users' feed¹⁴. YouTube labels videos that come from state-funded media outlets¹⁵. Similarly, Microsoft integrates NewsGuard in their mobile browser¹⁶. Yet, it is easy to quickly create, rebrand, and shut down online news sources: all it takes is to edit a web page. Whereas the position of a news source in the audience similarity space is a costly signal, appearing as a news source in the first place is a ludicrously cheap one. Efforts like OpenSources continuously examine news sources that soon disappear. Indeed, several in our dataset were shut down within months from our data collection. Assessing news sources is time-consuming and requires the labor of experts, whose efforts must be directed towards critical cases. Our data-driven approach can help social media platforms triage new information outlets in two possible ways. First, our clustering approach discussed in RQ1 can identify sources that share the same audience with a news outlet for which external assessments is already available. Then, it can propagate those assessments to the news outlet, thus alleviating the labor-intensive annotation of the ever growing space of information sources. Second, the classifier discussed in RQ3 can classify the articles by the news outlet with high precision as either widespread content, or one of the different types of misinformation: satirical/clickbait, conspiratorial/junk science, and right-wing/fake news. Experts may use these indications to triage questionable content.

Implications for developing novel assessments for the social news sphere

Whereas automation can help scale the process of assessing news sources, the matter of communicating those as-

¹⁴<https://newsroom.fb.com/news/2018/04/inside-feed-article-context>

¹⁵<https://money.cnn.com/2018/02/02/media/youtube-state-funded-media-label/>

¹⁶<https://twitter.com/MWautier/status/1081346843487854593>

assessments to the users requires careful consideration. In our comparison of expert- versus audience-based clustering, we highlighted the necessity of using multiple existing assessments to characterize news sources accurately. However, such redundant assessments complicate the clear communication of a news source's quality to their users. We believe this challenge is best addressed through research. Communication and social computing scholars have access to exclusive domain knowledge that is essential to synthesize novel assessments, so as to better describe current user practices in their choice of news sources. To this end, our approach allows researchers to combine multiple assessments of news sources, and to interpret them in the light of how users interact with them.

Limitations

It is important for us to highlight some limitations in this work, which the readers should take into consideration when interpreting the results. One such limitation is arguably our choice of data. We focus on an expansive list of 639 English-speaking news sources. However, this list is likely not representative of the landscape of news outlets on Twitter. To address this shortcoming, one would arguably require the complete data from the platform for an extended period of time—in fact, simply accessing a subsample of the Twitter stream would result in neglecting smaller sources. Yet, even then one would be omitting large players of the wider information ecosystem that includes television and talk radio. For practical reasons, we focused on sources for which reliable external assessments were available.

A second major limitation is that we rely on expert assessments by third-party initiatives. As a byproduct, not all sources have the same assessment coverage. For instance, one source might be assessed from MediaBias but not from Allsides, and vice-a-versa. In our analyses, we address this issue by looking at different providers of assessments separately. A different approach might involve in-house human annotators to harmonize assessments for all sources. However, training annotators for journalistic norms is still a subject of research (Zhang et al. 2018). Furthermore, we rely on source-level assessments of news, although different articles by one same source might have different qualities. One example is RT, which shares a combination of high-quality news reports and state propaganda (Starbird et al. 2018). This choice is in line with our goal of studying the audience of news sources, as we focus on the aggregate characteristics of the news media sphere. Yet, practitioners should use caution when applying source annotations to individual articles.

Finally, a crucial limitation is that our results are largely correlational in nature. In particular, we find strong correlations between external assessments of news sources, and their social media audience. Yet, qualitative research in online journalism and media literacy would be essential to understand whether those characteristics of the news sources are driving the users' choice of sharing them.

Conclusion

In this research, we characterized the landscape of news sources based on their audience on Twitter. We showed that news sources aggregate in communities of shared audience, that uphold distinct factuality standards, political partisanship, and journalistic norms.

In particular, we uncovered four data-driven communities of news sources: highly circulated news spanning the entire political spectrum; engaging misinformation, such as clickbait and satire; factoid misinformation, such as conspiratorial and junk science sources; and misinformation with an ulterior motive, such as far-right propaganda and fabricated news sources. We found distinguishing stylistic and topical markers that match with the characteristics of the news sources composing the clusters. For example, whereas sources in the widespread news cluster use more formal language, the conspiratorial cluster adopts an overconfident reporting style. The difference between the clusters is measurable: classifiers can automatically distinguish between news articles coming from different clusters with high precision. Thus, the Twitter audience delineates different segments of the news media landscape, differing in both the characteristics of the news sources and of the news content that Twitter users engage with. Yet, our findings challenge the common understanding of the news media landscape, exhibiting complex interrelation between popularity, partisan lines, and journalistic quality—with deep implications for gatekeeping and triaging misinformation in social media.

Acknowledgments

We would like to thank the Social Computing lab at Virginia Tech for comments on the early versions of the paper and the anonymous reviewers for their feedback. This work is partially supported by NSF grant IIS-1755547.

References

- Bail, C. A.; Argyle, L. P.; Brown, T. W.; Bumpus, J. P.; Chen, H.; Hunzaker, M. B. F.; Lee, J.; Mann, M.; Merhout, F.; and Volfovsky, A. 2018. Exposure to opposing views on social media can increase political polarization. *PNAS* 115(37):9216–9221.
- Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348(6239).
- Barberá, P.; Jost, J. T.; Nagler, J.; Tucker, J. A.; and Bonneau, R. 2015. Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science* 26(10).
- Bessi, A.; Coletto, M.; Davidescu, G. A.; Scala, A.; Caldarelli, G.; and Quattrociocchi, W. 2015. Science vs Conspiracy: Collective Narratives in the Age of Misinformation. *PLOS ONE* 10(2).
- Conover, M. D.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2011. Political polarization on twitter. In *ICWSM*.

- Diakopoulos, N. 2015. Picking the NYT Picks: Editorial Criteria and Automation in the Curation of Online News Comments. *#ISOJ* 5(1).
- Dodd, M. D.; Balzer, A.; Jacobs, C. M.; Gruszczynski, M. W.; Smith, K. B.; and Hibbing, J. R. 2012. The political left rolls with the good and the political right confronts the bad: Connecting physiology and cognition to preferences. *Philos. Trans. Royal Soc. B* 367(1589).
- Donath, J. 2011. Signals, cues and meaning.
- Druckman, J. N.; Levendusky, M. S.; and McLain, A. 2018. No Need to Watch: How the Effects of Partisan Media Can Spread via Interpersonal Discussions. *Am. J. Pol Sci* 62(1).
- Eisenstein, J.; Ahmed, A.; and Xing, E. P. 2011. Sparse additive generative models of text. In *ICML*.
- Elejalde, E.; Ferrer, L.; and Herder, E. 2017. The Nature of Real and Perceived Bias in Chilean Media. *HT*.
- Ferrara, E.; Varol, O.; Davis, C.; Menczer, F.; and Flammini, A. 2014. The Rise of Social Bots. *arXiv preprint arXiv:1407.5225* (grant 220020274):1–11.
- Garimella, K.; Morales, G. D. F.; Gionis, A.; and Mathioudakis, M. 2018. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. *WWW*.
- Golbeck, J.; Everett, J. B.; Falak, W.; Gieringer, C.; Graney, J.; Hoffman, K. M.; Huth, L.; Ma, Z.; Jha, M.; Khan, M.; Kori, V.; Mauriello, M.; Lewis, E.; Mirano, G.; Mohn IV, W. T.; Mussenden, S.; Nelson, T. M.; Mcwillie, S.; Pant, A.; Shetye, P.; Shrestha, R.; Steinheimer, A.; Auxier, B.; Subramanian, A.; Visnansky, G.; Bhanushali, K. H.; Bonk, C.; Bouzaghane, M. A.; Buntain, C.; Chanduka, R.; and Chekalos, P. 2018. Fake News vs Satire. In *Proceedings of the 10th ACM Conference on Web Science - WebSci '18*, 17–21. New York, New York, USA: ACM Press.
- Grinberg, N.; Joseph, K.; Friedland, L.; Swire-Thompson, B.; and Lazer, D. 2019. Fake news on twitter during the 2016 U.S. Presidential election. *Science* 378.
- Guess, A.; Nyhan, B.; Lyons, B.; and Reifler, J. 2018. Avoiding the Echo Chamber about Echo Chambers.
- Hartung, F.; Burke, M.; Hagoort, P.; and Willems, R. M. 2016. Taking perspective: Personal pronouns affect experiential aspects of literary reading. *PLoS ONE* 11(5).
- Hayes, A. S.; Singer, J. B.; and Ceppos, J. 2007. Shifting Roles, Enduring Values: The Credible Journalist in a Digital Age. *Journal of Mass Media Ethics* 22(4):262–279.
- Horne, B. D., and Adali, S. 2017. This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News.
- Horne, B. D., and Adali, S. 2018. An Exploration of Verbatim Content Republishing by News Producers. *NECO*.
- Hutto, C. J., and Gilbert, E. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *ICWSM*.
- Kulshrestha, J.; Zafar, M. B.; Noboa, L. E.; Gummadi, K. P.; and Ghosh, S. 2015. Characterizing Information Diets of Social Media Users. In *ICWSM*.
- Kumar, S.; West, R.; and Leskovec, J. 2016. Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes. *World Wide Web Conference Committee* 591–602.
- Lewandowsky, S.; Ecker, U. K. H.; Seifert, C. M.; Schwarz, N.; and Cook, J. 2012. Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest* 13(3):106–131.
- Lokot, T., and Diakopoulos, N. 2016. News Bots: Automating news and information dissemination on Twitter. *Digital Journalism* 4(6):682–699.
- Martin, T. 2017. community2vec: Vector representations of online communities encode semantic relationships. In *NLP CSS Workshop*.
- Marwick, A., and Lewis, R. 2017. Media Manipulation and Disinformation Online. *Data & Society Research Institute* 1–104.
- Matsa, K. E., and Shearer, E. 2018. News use across social media platforms 2018. In *Pew Research Center*.
- Mukerjee, S.; González-Bailón, S.; and Majó-Vázquez, S. 2018. Networks of Audience Overlap in the Consumption of Digital News. *Journal of Communication* 68(1):26–50.
- Oliver, J. E., and Wood, T. J. 2014. Conspiracy theories and the paranoid style(s) of mass opinion. *Am J Pol Sci* 58(4).
- Peter, C., and Koch, T. 2016. When Debunking Scientific Myths Fails (and When It Does Not). *Science Communication* 38(1):3–25.
- Sauri, R. 2008. A factuality profiler for eventualities in text.
- Soni, S.; Mitra, T.; Gilbert, E.; and Eisenstein, J. 2014. Modeling Factuality Judgments in Social Media Text. In *ACL*.
- Starbird, K.; Arif, A.; Wilson, T.; Koevering, K. V.; Yefimova, K.; and Scarnecchia, D. 2018. Ecosystem or echo-system? exploring content sharing across alternative media domains. In *ICWSM*.
- Starbird, K. 2017. Examining the Alternative Media Ecosystem through the Production of Alternative Narratives of Mass Shooting Events on Twitter. In *ICWSM*.
- Sundar, S. S. 1999. Exploring receivers' criteria for perception of print and online news. *JMCQ* 76(2).
- Tausczik, Y. R., and Pennebaker, J. W. 2010. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *J. Lang. Soc. Psychol.* 29(1).
- Thorson, E. 2016. Belief Echoes: The Persistent Effects of Corrected Misinformation. *Political Communication* 33(3):460–480.
- Webster, J. G., and Ksiazek, T. B. 2012. The Dynamics of Audience Fragmentation: Public Attention in an Age of Digital Media. *Journal of Communication* 62(1):39–56.
- Zhang, A. X.; Robbins, M.; Bice, E.; Hawke, S.; Karger, D.; Mina, A. X.; Ranganathan, A.; Metz, S. E.; Appling, S.; Sehat, C. M.; Gilmore, N.; Adams, N. B.; Vincent, E.; and Lee, J. 2018. A Structured Response to Misinformation: Defining and Annotating Credibility Indicators in News Articles. In *WWW*.